

NSX-T Data Center Administration Guide

Modified on 24 JAN 2024
VMware NSX-T Data Center 3.0

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

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

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

Copyright © 2017-2023 Broadcom. All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries. For more information, go to <https://www.broadcom.com>. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies. [Copyright and trademark information](#).

Contents

About Administering VMware NSX-T Data Center	13
1 NSX Manager	14
View Monitoring Dashboards	17
2 Tier-0 Gateways	20
Add a Tier-0 Gateway	21
Create an IP Prefix List	25
Create a Community List	26
Configure a Static Route	27
Create a Route Map	28
Using Regular Expressions to Match Community Lists When Adding Route Maps	30
Configure BGP	31
Configure BFD	35
Configure Multicast	35
Configure IPv6 Layer 3 Forwarding	36
Create SLAAC and DAD Profiles for IPv6 Address Assignment	36
Changing the HA Mode of a Tier-0 Gateway	38
Add a VRF Gateway	38
Configuring EVPN	40
3 Tier-1 Gateway	42
Add a Tier-1 Gateway	42
4 Segments	45
Segment Profiles	46
Understanding QoS Segment Profile	47
Understanding IP Discovery Segment Profile	49
Understanding SpoofGuard Segment Profile	51
Understanding Segment Security Segment Profile	52
Understanding MAC Discovery Segment Profile	54
Add a Segment	56
Types of DHCP on a Segment	58
Configure DHCP on a Segment	59
Configure DHCP Static Bindings on a Segment	66
Layer 2 Bridging	69
Create an Edge Bridge Profile	70
Configure Edge-Based Bridging	71

- Create a Layer 2 Bridge-Backed Segment 73
- Add a Metadata Proxy Server 73

5 Host Switches 75

- Managing NSX-T on a vSphere Distributed Switch 75
 - Configuring a vSphere Distributed Switch 76
 - Managing NSX Distributed Virtual Port Groups 78
 - NSX-T Cluster Prepared with VDS 79
 - APIs to Configure vSphere Distributed Switch on NSX-T Data Center 80
 - Feature Support in a vSphere Distributed Switch Enabled to Support NSX-T Data Center 83
- Enhanced Networking Stack 86
 - Automatically Assign ENS Logical Cores 86
 - Configure Guest Inter-VLAN Routing 87
- Migrate Host Switch to vSphere Distributed Switch 89
- NSX Virtual Distributed Switch 95

6 Virtual Private Network (VPN) 96

- Understanding IPsec VPN 97
 - Using Policy-Based IPsec VPN 98
 - Using Route-Based IPsec VPN 99
- Understanding Layer 2 VPN 100
 - Enable and Disable L2 VPN Path MTU Discovery 102
- Adding VPN Services 102
 - Add an IPsec VPN Service 104
 - Add an L2 VPN Service 105
- Adding IPsec VPN Sessions 108
 - Add a Policy-Based IPsec Session 108
 - Add a Route-Based IPsec Session 111
 - About Supported Compliance Suites 115
 - Understanding TCP MSS Clamping 116
- Adding L2 VPN Sessions 117
 - Add an L2 VPN Server Session 117
 - Add an L2 VPN Client Session 119
 - Download the Remote Side L2 VPN Configuration File 121
- Add Local Endpoints 122
- Adding Profiles 124
 - Add IKE Profiles 125
 - Add IPsec Profiles 127
 - Add DPD Profiles 129
- Add an Autonomous Edge as an L2 VPN Client 130
- Check the Realized State of an IPsec VPN Session 133

Monitor and Troubleshoot VPN Sessions 136

7 Network Address Translation (NAT) 137

Configure NAT on a Gateway 138

8 Load Balancing 142

Key Load Balancer Concepts 143

Scaling Load Balancer Resources 143

Supported Load Balancer Features 144

Load Balancer Topologies 145

Setting Up Load Balancer Components 147

Add Load Balancers 147

Add an Active Monitor 149

Add a Passive Monitor 152

Add a Server Pool 154

Setting Up Virtual Server Components 157

Groups Created for Server Pools and Virtual Servers 189

9 Distributed Load Balancer 190

Understanding Traffic Flow with a Distributed Load Balancer 192

Create and Attach a Distributed Load Balancer Instance 193

Create a Server Pool for Distributed Load Balancer 194

Create a Virtual Server with a Fast TCP or UDP Profile 195

Verifying Distributed Load Balancer Configuration on ESXi Hosts 196

Monitoring Distributed Load Balancer Statistics 198

10 Forwarding Policies 200

Add or Edit Forwarding Policies 201

11 IP Address Management (IPAM) 203

Add a DNS Zone 203

Add a DNS Forwarder Service 204

Add a DHCP Profile 205

Add a DHCP Server Profile 205

Add a DHCP Relay Profile 208

Attach a DHCP Profile to a Tier-0 or Tier-1 Gateway 209

Scenarios: Selection of Edge Cluster for DHCP Service 210

Scenarios: Impact of Changing Segment Connectivity on DHCP 215

Add an IP Address Pool 218

Add an IP Address Block 219

12 Networking Settings 220

- Configuring Multicast 220
 - Create an IGMP Profile 222
 - Create a PIM Profile 222
- Add a VNI Pool 223
- Configure Gateway Settings 223
- Add a Gateway QoS Profile 224
- Add a BFD Profile 225

13 Security 226

- Security Configuration Overview 226
- Security Overview 227
- Security Terminology 228
- Identity Firewall 229
 - Identity Firewall Workflow 230
- Layer 7 Context Profile 232
 - Layer 7 Firewall Rule Workflow 233
 - Attributes 234
- Distributed Firewall 238
 - Firewall Drafts 238
 - Add a Distributed Firewall 241
 - Distributed Firewall Packet Logs 244
 - Manage a Firewall Exclusion List 246
 - Filtering Specific Domains (FQDN/URLs) 247
 - Extending Security Policies to Physical Workloads 249
 - Shared Address Sets 255
- Distributed IDS 256
 - Distributed IDS Settings and Signatures 256
 - Distributed IDS Profiles 259
 - Distributed IDS Rules 262
 - Distributed IDS Events 263
 - Verify Distributed IDS Status on Host 265
- East-West Network Security - Chaining Third-party Services 267
 - Key Concepts of Network Protection East-West 267
 - NSX-T Data Center Requirements for East-West Traffic 268
 - High-Level Tasks for East-West Network Security 268
 - Deploy a Service for East-West Traffic Introspection 269
 - Add Redirection Rules for East-West Traffic 270
 - Uninstall an East-West Traffic Introspection Service 272
- Gateway Firewall 273
 - Add a Gateway Firewall Policy and Rule 273

URL Analysis Workflow	276
Gateway Firewall Packet Logs	278
North-South Network Security - Inserting Third-party Service	281
High-Level Tasks for North-South Network Security	281
Deploy a Service for North-South Traffic Introspection	281
Add Redirection Rules for North-South Traffic	283
Uninstall a North-South Traffic Introspection Service	284
Endpoint Protection	285
Understand Endpoint Protection	285
Configure Endpoint Protection	289
Manage Endpoint Protection	306
Security Profiles	319
Create a Session Timer	319
Flood Protection	321
Configure DNS Security	323
Manage Group to Profile Precedence	324
Time-Based Firewall Policy	325
Network Introspection Settings	326
Add a Service Segment	326
Add a Service Profile	327
Add a Service Chain	328
Troubleshooting Firewall	329
Monitor and Troubleshoot Firewall on NSX Manager	329
Troubleshooting Distributed Firewall on ESX Hosts	329
Troubleshooting Distributed Firewall on KVM Hosts	339
Troubleshooting Gateway Firewall	342
Check Rule Realization Status	346
Distributed Firewall Packet Logs	348
Bare Metal Server Security	351

14 Inventory 353

Add a Service	353
Add a Group	354
Add a Context Profile	357
Containers	358
Public Cloud Services	360
Physical Servers	360
Tags	360
Add Tags to an Object	364
Add a Tag to Multiple Objects	364
Unassign Tags from an Object	366

Unassign a Tag from Multiple Objects 366

15 Multisite and Federation 368

NSX-T Data Center Multisite 368

Working with VMware Site Recovery Manager 379

NSX Federation 380

Overview of NSX Federation 380

Networking in NSX Federation 389

Security in NSX Federation 404

Backup and Restore in NSX Federation 419

16 System Monitoring 421

Monitor NSX Edge Nodes 421

Working with Events and Alarms 423

About Events and Alarms 423

View Alarm Information 455

View Alarm Definitions 456

Configuring Alarm Definition Settings 458

Managing Alarm States 459

Using vRealize Log Insight for System Monitoring 460

Using vRealize Operations Manager for System Monitoring 461

Using vRealize Network Insight Cloud for System Monitoring 465

17 Network Monitoring 476

Add an IPFIX Collector 476

Add a Firewall IPFIX Profile 477

Add a Switch IPFIX Profile 477

IPFIX Monitoring on a vSphere Distributed Switch 479

Add a Port Mirroring Profile 479

Port Mirroring on a vSphere Distributed Switch 480

Perform a Traceflow 481

Simple Network Management Protocol (SNMP) 484

Monitor Fabric Nodes 485

Network Latency Statistics 485

Measure Network Latency Statistics 489

Export Network Latency Statistics 490

Monitoring Tools in Manager Mode 492

View Port Connection Information in Manager Mode 492

Traceflow 493

Monitor Port Mirroring Sessions in Manager Mode 496

Configure Filters for a Port Mirroring Session 499

- Configure IPFIX in Manager Mode 501
- Monitor a Logical Switch Port Activity in Manager Mode 671

18 Authentication and Authorization 672

- Local User Accounts 673
 - Manage a User's Password or Name 673
 - Resetting the Passwords of an Appliance 674
 - Authentication Policy Settings 675
- Integration with VMware Identity Manager/Workspace ONE Access 676
 - Time Synchronization between NSX Manager, vIDM, and Related Components 676
 - Obtain the Certificate Thumbprint from a vIDM Host 678
 - Configure VMware Identity Manager/Workspace ONE Access Integration 679
 - Validate VMware Identity Manager Functionality 682
- Integration with LDAP 683
 - LDAP Identity Source 684
- Add a Role Assignment or Principal Identity 686
- Configuring Both vIDM and LDAP or Transitioning from vIDM to LDAP 688
- Role-Based Access Control 688
- Logging User Account Changes 700

19 Certificates 702

- Types of Certificates 702
- Certificates for NSX Federation 704
- Create a Certificate Signing Request File 706
- Creating Self-signed Certificates 707
 - Create a Self-Signed Certificate 707
 - Import a Certificate for a CSR 708
- Importing and Replacing Certificates 709
 - Import a Self-signed or CA-signed Certificate 709
 - Import a CA Certificate 710
 - Set Checks for Certificate Imports 711
 - Replace Certificates 711
- Importing and Retrieving CRLs 713
 - Import a Certificate Revocation List 713
 - Configuring NSX Manager to Retrieve a Certificate Revocation List 714
- Storage of Public Certificates and Private Keys for Load Balancer or VPN service 715
- Alarm Notification for Certificate Expiration 715

20 Configuring NSX-T Data Center in Manager Mode 716

- Logical Switches in Manager Mode 716
 - Understanding BUM Frame Replication Modes 717

Create a Logical Switch in Manager Mode	719
Connecting a VM to a Logical Switch in Manager Mode	720
Create a Logical Switch Port In Manager Mode	729
Test Layer 2 Connectivity in Manager Mode	730
Create a VLAN Logical Switch for the NSX Edge Uplink in Manager Mode	733
Switching Profiles for Logical Switches and Logical Ports	735
Layer 2 Bridging in Manager Mode	752
Logical Routers in Manager Mode	758
Tier-1 Logical Router	759
Tier-0 Logical Router	769
NAT in Manager Mode	802
Network Address Translation	802
Grouping Objects in Manager Mode	815
Create an IP Set in Manager Mode	815
Create an IP Pool in Manager Mode	816
Create a MAC Set in Manager Mode	816
Create an NSGroup in Manager Mode	817
Configuring Services and Service Groups	819
Manage Tags for a VM in Manager Mode	820
DHCP in Manager Mode	821
DHCP	821
Metadata Proxies	826
IP Address Management in Manager Mode	828
Manage IP Blocks in Manager Mode	828
Manage Subnets for IP Blocks in Manager Mode	829
Load Balancing in Manager Mode	830
Key Load Balancer Concepts	831
Configuring Load Balancer Components	831
Firewall in Manager Mode	862
Add or Delete a Firewall Rule to a Logical Router in Manager Mode	862
Configure Firewall for a Logical Switch Bridge Port in Manager Mode	863
Firewall Sections and Firewall Rules	864
About Firewall Rules	868

21 Backing Up and Restoring NSX Manager 876

Configure Backups	877
Remove Old Backups	879
Restore a Backup	880
Listing Available Backups	883
Certificate Management after Restore	884

22 Operations and Management 886

- View the Usage and Capacity of Categories of Objects 887
- Configure User Interface Settings 889
- Configure a Node Profile 889
- Checking the Realized State of a Configuration Change 891
- View Network Topology 895
- Search for Objects 895
- Filter by Object Attributes 896
- Add a Compute Manager 897
- Add an Active Directory 900
- Add an LDAP Server 901
- Synchronize Active Directory 902
- Remove NSX-T Data Center Extension from vCenter Server 903
- Managing the NSX Manager Cluster 904
 - View the Configuration and Status of the NSX Manager Cluster 904
 - Update API Service Configuration of the NSX Manager Cluster 907
 - Shut Down and Power On the NSX Manager Cluster 908
 - Reboot an NSX Manager 908
 - Change the IP Address of an NSX Manager 908
 - Resize an NSX Manager Node 910
- Replacing an NSX Edge Transport Node in an NSX Edge Cluster 911
 - Replace an NSX Edge Transport Node Using the NSX Manager UI 911
 - Replace an NSX Edge Transport Node Using the API 912
- Managing Resource Reservations for an Edge VM Appliance 914
 - Tune Resource Reservations for an NSX Edge Appliance 915
- Adding and Removing an ESXi Host Transport Node to and from vCenter Servers 916
- Changing the Distributed Router Interfaces' MAC Address 917
- Configuring Appliances 919
- Add a License Key and Generate a License Usage Report 919
- Compliance-Based Configuration 922
 - View Compliance Status Report 923
 - Compliance Status Report Codes 923
 - Configure Global FIPS Compliance Mode for Load Balancer 926
- Collect Support Bundles 929
- Log Messages and Error Codes 930
 - Configure Remote Logging 934
 - Log Message IDs 941
 - Troubleshooting Syslog Issues 942
 - Configure Serial Logging on an Appliance VM 943
 - Firewall Audit Log Messages 943
- Customer Experience Improvement Program 960

Edit the Customer Experience Improvement Program Configuration	960
Find the SSH Fingerprint of a Remote Server	960
Configuring an External Load Balancer	961
Configure Proxy Settings	963
View Container-Related Information	964

23 Using NSX Cloud 965

Cloud Service Manager: UI Walkthrough	965
Clouds	965
System	971
Threat Detection using the NSX Cloud Quarantine Policy	977
Quarantine Policy in the NSX Enforced Mode	977
Quarantine Policy in the Native Cloud Enforced Mode	982
User Managed List for VMs	982
NSX Enforced Mode	983
Currently Supported Operating Systems for Workload VMs	984
Onboarding VMs in the NSX Enforced Mode	984
Managing VMs in the NSX Enforced Mode	993
Native Cloud Enforced Mode	995
Managing VMs in the Native Cloud Enforced Mode	995
NSX-T Data Center Features Supported with NSX Cloud	998
Group VMs using NSX-T Data Center and Public Cloud Tags	999
Use Native-Cloud Services	1003
Service Insertion for your Workload VMs in the NSX Enforced Mode	1004
Enable NAT on NSX-managed VMs	1013
Enable Syslog Forwarding	1014
Set up VPN in the Native Cloud Enforced Mode	1014
Set up VPN in the NSX Enforced Mode	1023
NSX Cloud FAQs and Troubleshooting	1029

About Administering VMware NSX-T Data Center

The *NSX-T Data Center Administration Guide* provides information about configuring and managing networking for VMware NSX-T™ Data Center, including how to create logical switches and ports and how to set up networking for tiered logical routers, configure NAT, firewalls, SpoofGuard, grouping and DHCP. It also describes how to configure NSX Cloud.

Intended Audience

This information is intended for anyone who wants to configure NSX-T Data Center. The information is written for experienced Windows or Linux system administrators who are familiar with virtual machine technology, networking, and security operations.

VMware Technical Publications Glossary

VMware Technical Publications provides a glossary of terms that might be unfamiliar to you. For definitions of terms as they are used in VMware technical documentation, go to <https://www.vmware.com/topics/glossary>.

Related Documentation

You can find the VMware NSX® Intelligence™ documentation at <https://docs.vmware.com/en/VMware-NSX-Intelligence/index.html>. The NSX Intelligence 1.0 content was initially included and released with the NSX-T Data Center 2.5 documentation set.

NSX Manager

1


The NSX Manager provides a web-based user interface where you can manage your NSX-T environment. It also hosts the API server that processes API calls.

The NSX Manager interface provides two modes for configuring resources:

- Policy mode
- Manager mode

Accessing Policy Mode and Manager Mode

If present, you can use the **Policy** and **Manager** buttons to switch between the Policy and Manager modes. Switching modes controls which menus items are available to you.



- By default, if your environment contains only objects created through Policy mode, your user interface is in Policy mode and you do not see the **Policy** and **Manager** buttons.
- By default, if your environment contains any objects created through Manager mode, you see the **Policy** and **Manager** buttons in the top-right corner.

These defaults can be changed by modifying the user interface settings. See [Configure User Interface Settings](#) for more information.

The same **System** tab is used in the Policy and Manager interfaces. If you modify Edge nodes, Edge clusters, or transport zones, it can take up to 5 minutes for those changes to be visible in Policy mode. You can synchronize immediately using `POST /policy/api/v1/infra/sites/default/enforcement-points/default?action=reload`.

When to Use Policy Mode or Manager Mode

Be consistent about which mode you use. There are a few reasons to use one mode over the other.

- If you are deploying a new NSX-T Data Center environment, using **Policy** mode to create and manage your environment is the best choice in most situations.
 - Some features are not available in Policy mode. If you need these features, use **Manager** mode for all configurations.
- If you plan to use NSX Federation, use **Policy** mode to create all objects. Global Manager supports only Policy mode.
- If you are upgrading from an earlier version of NSX-T Data Center and your configurations were created using the Advanced Networking & Security tab, use **Manager** mode.

The menu items and configurations that were found under the Advanced Networking & Security tab are available in NSX-T Data Center 3.0 in **Manager** mode.

Important If you decide to use Policy mode, use it to create all objects. Do not use Manager mode to create objects.

Similarly, if you need to use Manager mode, use it to create all objects. Do not use Policy mode to create objects.

Table 1-1. When to Use Policy Mode or Manager Mode

Policy Mode	Manager Mode
Most new deployments should use Policy mode. Federation supports only Policy mode. If you want to use Federation, or might use it in future, use Policy mode.	Deployments which were created using the advanced interface, for example, upgrades from versions before Policy mode was available.
NSX Cloud deployments	Deployments which integrate with other plugins. For example, NSX Container Plug-in, Openstack, and other cloud management platforms.
Networking features available in Policy mode only: ■ DNS Services and DNS Zones ■ VPN ■ Forwarding policies for NSX Cloud	Networking features available in Manager mode only: ■ Forwarding up timer
Security features available in Policy mode only: ■ Endpoint Protection ■ Network Introspection (East-West Service Insertion) ■ Context Profiles ■ L7 applications ■ FQDN ■ New Distributed Firewall and Gateway Firewall Layout ■ Categories ■ Auto service rules ■ Drafts	Security features available in Manager mode only: ■ Bridge Firewall

Names for Objects Created in Policy Mode and Manager Mode

The objects you create have different names depending on which interface was used to create them.

Table 1-2. Object Names

Objects Created Using Policy Mode	Objects Created Using Manager Mode
Segment	Logical switch
Tier-1 gateway	Tier-1 logical router
Tier-0 gateway	Tier-0 logical router
Group	NSGroup, IP Sets, MAC Sets
Security Policy	Firewall section
Gateway firewall	Edge firewall

Policy and Manager APIs

The NSX Manager provides two APIs: Policy and Manager.

- The Policy API contains URLs that begin with `/policy/api`.
- The Manager API contains URLs that begin with `/api`.

For more information about using the Policy API, see the [NSX-T Policy API Getting Started Guide](#).

Security

NSX Manager has the following security features:

- NSX Manager has a built-in user account called `admin`, which has access rights to all resources, but does not have rights to the operating system to install software. NSX-T upgrade files are the only files allowed for installation. You cannot edit the rights of or delete the `admin` user. Note that you can change the username `admin`.
- NSX Manager supports session timeout and automatic user logout. NSX Manager does not support session lock. Initiating a session lock can be a function of the workstation operating system being used to access NSX Manager. Upon session termination or user logout, users are redirected to the login page.
- Authentication mechanisms implemented on NSX-T follow security best practices and are resistant to replay attacks. The secure practices are deployed systematically. For example, sessions IDs and tokens on NSX Manager for each session are unique and expire after the user logs out or after a period of inactivity. Also, every session has a time record and the session communications are encrypted to prevent session hijacking.

You can view and change the session timeout value with the following CLI commands:

- The command `get service http` displays a list of values including session timeout.
- To change the session timeout value, run the following commands:

```
set service http session-timeout <timeout-value-in-seconds>
restart service ui-service
```

Read the following topics next:

- [View Monitoring Dashboards](#)

View Monitoring Dashboards

The NSX Manager interface provides numerous monitoring dashboards showing details regarding system status, networking and security, and compliance reporting. This information is displayed or accessible throughout the NSX Manager interface, but can be accessed together in the **Home > Monitoring Dashboards** page.

You can access the monitoring dashboards from the Home page of the NSX Manager interface. From the dashboards, you can click through and access the source pages from which the dashboard data is drawn.

Procedure

- 1 Log in as administrator to the NSX Manager interface.
- 2 Click **Home** if you are not already on the Home page.
- 3 Click Monitoring Dashboards and select the desired category of dashboards from the drop-down menu.

The page displays the dashboards in the selected categories. The dashboard graphics are color-coded, with color code key displayed directly above the dashboards.

- 4 To access a deeper level of detail, click the title of the dashboard, or one of the elements of the dashboard, if activated.

The following tables describe the default dashboards and their sources.

Table 1-3. System Dashboards

Dashboard	Sources	Description
System	System > Appliances > Overview	Shows the status of the NSX Manager cluster and resource (CPU, memory, disk) consumption.
Fabric	System > Fabric > Nodes System > Fabric > Transport Zones System > Fabric > Compute Managers	Shows the status of the NSX-T fabric, including host and edge transport nodes, transport zones, and compute managers.

Table 1-3. System Dashboards (continued)

Dashboard	Sources	Description
Backups	System > Backup & Restore	Shows the status of NSX-T backups, if configured. It is strongly recommended that you configure scheduled backups that are stored remotely to an SFTP site.
Endpoint Protection	System > Service Deployments	Shows the status of endpoint protection deployment.

Table 1-4. Networking & Security Dashboards in Policy Mode

Dashboard	Sources	Description
Security	Inventory > Groups Security > Distributed Firewall	Shows the status of groups and security policies. A group is a collection of workloads, segments, segment ports, and IP addresses, where security policies, including East-West firewall rules, may be applied.
Gateways	Networking > Tier-0 Gateways Networking > Tier-1 Gateways	Shows the status of Tier-0 and Tier-1 gateways.
Segments	Networking > Segments	Shows the status of network segments.
Load Balancers	Networking > Load Balancing	Shows the status of the load balancer VMs.
VPNs	Networking > VPN	Shows the status of virtual private networks.

Table 1-5. Networking & Security Dashboards in Manager Mode

Dashboard	Sources	Description
Load Balancers	Networking > Load Balancing	Shows the status of the load balancer services, load balancer virtual servers, and load balancer server pools. A load balancer can host one or more virtual servers. A virtual server is bound to a server pool that includes members hosting applications.
Firewall	Security > Distributed Firewall Security > Bridge Firewall Networking > Tier-0 Logical Routers and Networking > Tier-1 Logical Routers	Indicates if the firewall is enabled, and shows the number of policies, rules, and exclusions list members. Note Each detailed item displayed in this dashboard is sourced from a specific sub-tab in the source page cited.
VPN	Not applicable.	Shows the status of virtual private networks and the number of IPSec and L2 VPN sessions open.
Switching	Networking > Logical Switches	Shows the status of logical switches and logical ports, including both VM and container ports.

Table 1-6. Compliance Report Dashboard

Column	Description
Non-Compliance Code	Displays the specific non-compliance code.
Description	Specific cause of non-compliance status.
Resource Name	The NSX-T resource (node, switch, and profile) in non-compliance.
Resource Type	Resource type of cause.
Affected Resources	Number of resources affected. Click the number value to view a list.

See the [Compliance Status Report Codes](#) for more information about each compliance report code.

Tier-0 Gateways

2

A tier-0 gateway performs the functions of a tier-0 logical router. It processes traffic between the logical and physical networks.

NSX Cloud Note If using NSX Cloud, see [NSX-T Data Center Features Supported with NSX Cloud](#) for a list of auto-generated logical entities, supported features, and configurations required for NSX Cloud.

An Edge node can support only one tier-0 gateway or logical router. When you create a tier-0 gateway or logical router, make sure you do not create more tier-0 gateways or logical routers than the number of Edge nodes in the NSX Edge cluster.

Note When connecting tier-0 uplinks to multi-chassis port-channel topologies such as vPC (virtual PortChannel) or VSS (Virtual Switching System) from Cisco, or MLAG (Multi-Chassis Link Aggregation) from Arista, be sure to consult with the network provider to understand the limitations of the topology when it is being used for transit routing.

Read the following topics next:

- [Add a Tier-0 Gateway](#)
- [Create an IP Prefix List](#)
- [Create a Community List](#)
- [Configure a Static Route](#)
- [Create a Route Map](#)
- [Using Regular Expressions to Match Community Lists When Adding Route Maps](#)
- [Configure BGP](#)
- [Configure BFD](#)
- [Configure Multicast](#)
- [Configure IPv6 Layer 3 Forwarding](#)
- [Create SLAAC and DAD Profiles for IPv6 Address Assignment](#)
- [Changing the HA Mode of a Tier-0 Gateway](#)
- [Add a VRF Gateway](#)

■ [Configuring EVPN](#)

Add a Tier-0 Gateway

A tier-0 gateway has downlink connections to tier-1 gateways and uplink connections to physical networks.

If you are adding a tier-0 gateway from Global Manager in NSX Federation, see [Add a Tier-0 Gateway from Global Manager](#).

You can configure the HA (high availability) mode of a tier-0 gateway to be active-active or active-standby. The following services are only supported in active-standby mode:

- NAT
- Load balancing
- Stateful firewall
- VPN

Note Active-standby tier-0 gateways are supported starting in NSX-T Data Center 3.0.1.

Tier-0 and tier-1 gateways support the following addressing configurations for all interfaces (uplinks, service ports and downlinks) in both single tier and multi-tiered topologies:

- IPv4 only
- IPv6 only
- Dual Stack - both IPv4 and IPv6

To use IPv6 or dual stack addressing, enable **IPv4 and IPv6** as the L3 Forwarding Mode in **Networking > Networking Settings > Global Networking Config**.

You can configure the tier-0 gateway to support EVPN (Ethernet VPN) type-5 routes. For more information about configuring EVPN, see [Configuring EVPN](#).

If you configure route redistribution for the tier-0 gateway, you can select from two groups of sources: tier-0 subnets and advertised tier-1 subnets. The sources in the tier-0 subnets group are:

Source Type	Description
Connected Interfaces and Segments	These include external interface subnets, service interface subnets and segment subnets connected to the tier-0 gateway.
Static Routes	Static routes that you have configured on the tier-0 gateway.
NAT IP	NAT IP addresses owned by the tier-0 gateway and discovered from NAT rules that are configured on the tier-0 gateway.
IPSec Local IP	Local IPSEC endpoint IP address for establishing VPN sessions.
DNS Forwarder IP	Listener IP for DNS queries from clients and also used as source IP used to forward DNS queries to upstream DNS server.
EVPN TEP IP	This is used to redistribute EVPN local endpoint subnets on the tier-0 gateway.

The sources in the advertised tier-1 subnets group are:

Source Type	Description
Connected Interfaces and Segments	These include segment subnets connected to the tier-1 gateway and service interface subnets configured on the tier-1 gateway.
Static Routes	Static routes that you have configured on the tier-1 gateway.
NAT IP	NAT IP addresses owned by the tier-1 gateway and discovered from NAT rules that are configured on the tier-1 gateway.
LB VIP	IP address of the load balancing virtual server.
LB SNAT IP	IP address or a range of IP addresses used for source NAT by the load balancer.
DNS Forwarder IP	Listener IP for DNS queries from clients and also used as source IP used to forward DNS queries to upstream DNS server.
IPSec Local Endpoint	IP address of the IPSec local endpoint.

On a tier-0 gateway, proxy ARP handles ARP queries for the external and service interface IPs. Starting with NSX-T Data Center 3.0.2, proxy ARP also handles ARP queries for service IPs that are in an IP prefix list that is configured with the `Permit` action.

Prerequisites

If you plan to configure multicast, see [Configuring Multicast](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Tier-0 Gateways**.
- 3 Click **Add Tier-0 Gateway**.
- 4 Enter a name for the gateway.
- 5 Select an HA (high availability) mode.

The default mode is active-active. In the active-active mode, traffic is load balanced across all members. In active-standby mode, all traffic is processed by an elected active member. If the active member fails, a new member is elected to be active.

- 6 If the HA mode is active-standby, select a failover mode.

Option	Description
Preemptive	If the preferred node fails and recovers, it will preempt its peer and become the active node. The peer will change its state to standby.
Non-preemptive	If the preferred node fails and recovers, it will check if its peer is the active node. If so, the preferred node will not preempt its peer and will be the standby node.

- 7 (Optional) Select an NSX Edge cluster.

8 (Optional) Click **Additional Settings.**

- a In the **Internal Transit Subnet** field, enter a subnet.

This is the subnet used for communication between components within this gateway. The default is 169.254.0.0/24.

- b In the **TO-T1 Transit Subnets** field, enter one or more subnets.

These subnets are used for communication between this gateway and all tier-1 gateways that are linked to it. After you create this gateway and link a tier-1 gateway to it, you will see the actual IP address assigned to the link on the tier-0 gateway side and on the tier-1 gateway side. The address is displayed in **Additional Settings > Router Links** on the tier-0 gateway page and the tier-1 gateway page. The default is 100.64.0.0/16.

9 Click **Route Distinguisher for VRF Gateways to configure a route distinguisher admin address.**

This is only needed for EVPN and for the automatic route distinguisher use case.

10 (Optional) Add one or more tags.**11 Click **Save**.****12 For IPv6, under **Additional Settings**, you can select or create an **ND Profile** and a **DAD Profile**.**

These profiles are used to configure Stateless Address Autoconfiguration (SLAAC) and Duplicate Address Detection (DAD) for IPv6 addresses.

13 (Optional) Click **EVPN Settings to configure EVPN.**

- a Select a VNI pool.

You can click the menu icon (3 dots) to create a VNI pool if you have not previously created one.

- b In the **EVPN Tunnel Endpoint** field click **Set** to add EVPN local tunnel endpoints.

For the tunnel endpoint, select an Edge node and specify an IP address.

Optionally, you can specify the MTU.

Note Ensure that the uplink interface has been configured on the NSX Edge node that you select for the EVPN tunnel endpoint.

14 To configure route redistribution, click **Route Redistribution and **Set**.**

Select one or more of the sources:

- Tier-0 subnets: **Static Routes, NAT IP, IPsec Local IP, DNS Forwarder IP, EVPN TEP IP, Connected Interfaces & Segments.**

Under **Connected Interfaces & Segments**, you can select one or more of the following: **Service Interface Subnet, External Interface Subnet, Loopback Interface Subnet, Connected Segment.**

- Advertised tier-1 subnets: **DNS Forwarder IP, Static Routes, LB VIP, NAT IP, LB SNAT IP, IPsec Local Endpoint, Connected Interfaces & Segments.**

Under **Connected Interfaces & Segments**, you can select **Service Interface Subnet** and/or **Connected Segment**.

15 To configure interfaces, click **Interfaces** and **Set**.

- a Click **Add Interface**.
- b Enter a name.
- c Select a type.

If the HA mode is active-standby, the choices are **External**, **Service**, and **Loopback**. If the HA mode is active-active, the choices are **External** and **Loopback**.

- d Enter an IP address in CIDR format.
- e Select a segment.
- f If the interface type is not **Service**, select an NSX Edge node.
- g (Optional) If the interface type is not **Loopback**, enter an MTU value.
- h (Optional) If the interface type is **External**, you can enable multicast by setting **PIM** (Protocol Independent Multicast) to **Enabled**.

PIM can be enabled only on a single uplink interface.

Note: If you later disable **PIM** on this interface, then multicast will be disabled on all interfaces including the downlinks on this gateway.

- i (Optional) Add tags and select an ND profile.
- j (Optional) If the interface type is **External**, for **URPF Mode**, you can select **Strict** or **None**.

URPF (Unicast Reverse Path Forwarding) is a security feature.

- k After you create an interface, you can download the ARP table by clicking the menu icon (three dots) for the interface and selecting **Download ARP table**.

16 (Optional) If the HA mode is active-standby, click **Set** next to **HA VIP Configuration** to configure HA VIP.

With HA VIP configured, the tier-0 gateway is operational even if one uplink is down. The physical router interacts with the HA VIP only. HA VIP is intended to work with static routing and not with BGP.

- a Click **Add HA VIP Configuration**.
- b Enter an IP address and subnet mask.

The HA VIP subnet must be the same as the subnet of the interface that it is bound to.

- c Select two interfaces from two different Edge nodes.

17 Click **Routing** to add IP prefix lists, community lists, static routes, and route maps.

- 18 Click **Multicast** to configure multicast routing.
- 19 Click **BGP** to configure BGP.
- 20 (Optional) To download the routing table or forwarding table, click the menu icon (three dots) and select a download option. Enter values for **Transport Node**, **Network** and **Source** as required, and save the .CSV file.

What to do next

After the tier-0 gateway is added, you can optionally enable dynamic IP management on the gateway by selecting either a DHCP server profile or a DHCP relay profile. For more information, see [Attach a DHCP Profile to a Tier-0 or Tier-1 Gateway](#).

Create an IP Prefix List

An IP prefix list contains single or multiple IP addresses that are assigned access permissions for route advertisement. The IP addresses in this list are processed sequentially. IP prefix lists are referenced through BGP neighbor filters or route maps with in or out direction.

For example, you can add the IP address 192.168.100.3/27 to the IP prefix list and deny the route from being redistributed to the northbound router. You can also append an IP address with less-than-or-equal-to (le) and greater-than-or-equal-to (ge) modifiers to grant or limit route redistribution. For example, 192.168.100.3/27 ge 24 le 30 modifiers match subnet masks greater than or equal to 24-bits and less than or equal to 30-bits in length.

Note The default action for a route is **Deny**. When you create a prefix list to deny or permit specific routes, be sure to create an IP prefix with no specific network address (select **Any** from the dropdown list) and the **Permit** action if you want to permit all other routes.

Prerequisites

Verify that you have a tier-0 gateway configured. See [Create a Tier-0 Logical Router in Manager Mode](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Tier-0 Gateways**.
- 3 To edit a tier-0 gateway, click the menu icon (three dots) and select **Edit**.
- 4 Click **Routing**.
- 5 Click **Set** next to **IP Prefix List**.
- 6 Click **Add IP Prefix List**.
- 7 Enter a name for the IP prefix list.
- 8 Click **Set** to add IP prefixes.

- 9 Click **Add Prefix**.
 - a Enter an IP address in CIDR format.
For example, 192.168.100.3/27.
 - b (Optional) Set a range of IP address numbers in the **le** or **ge** modifiers.
For example, set **le** to 30 and **ge** to 24.
 - c Select **Deny** or **Permit** from the drop-down menu.
 - d Click **Add**.
- 10 Repeat the previous step to specify additional prefixes.
- 11 Click **Save**.

Create a Community List

You can create BGP community lists so that you can configure route maps based on community lists.

Community lists are user-defined lists of community attribute values. These lists can be used for matching or manipulating the communities attribute in BGP update messages.

Both the BGP Communities attribute (RFC 1997) and the BGP Large Communities attribute (RFC 8092) are supported. The BGP Communities attribute is a 32-bit value split into two 16-bit values. The BGP Large Communities attribute has 3 components, each 4 octets in length.

In route maps we can match on or set the BGP Communities or Large Communities attribute. Using this feature, network operators can implement network policy based on the BGP communities attribute.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Tier-0 Gateways**.
- 3 To edit a tier-0 gateway, click the menu icon (three dots) and select **Edit**.
- 4 Click **Routing**.
- 5 Click **Set** next to **Community List**.
- 6 Click **Add Community List**.
- 7 Enter a name for the community list.

- 8 Specify a list of communities. For a regular community, use the aa:nn format, for example, 300:500. For a large community, use the format aa:bb:cc, for example, 11:22:33. Note that the list cannot have both regular communities and large communities. It must contain only regular communities, or only large communities.

In addition, you can select one or more of the following regular communities. Note that they cannot be added if the list contains large communities.

- NO_EXPORT_SUBCONFED - Do not advertise to EBGp peers.
- NO_ADVERTISE - Do not advertise to any peer.
- NO_EXPORT - Do not advertise outside BGP confederation

- 9 Click **Save**.

Configure a Static Route

You can configure a static route on the tier-0 gateway to external networks. After you configure a static route, there is no need to advertise the route from tier-0 to tier-1, because tier-1 gateways automatically have a static default route towards their connected tier-0 gateway.

Recursive static routes are supported.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Tier-0 Gateways**.
- 3 To edit a tier-0 gateway, click the menu icon (three dots) and select **Edit**.
- 4 Click **Routing**.
- 5 Click **Set** next to **Static Routes**.
- 6 Click **Add Static Route**.
- 7 Enter a name and network address in CIDR format. Static routes based on IPv6 are supported. IPv6 prefixes can only have an IPv6 next hop.
- 8 Click **Set Next Hops** to add next-hop information.
- 9 Click **Add Next Hop**.
- 10 Enter an IP address or select **NULL**.
If **NULL** is selected, the route is called a device route.
- 11 Specify the administrative distance.
- 12 Select a scope from the drop-down list. A scope can be an interface, a gateway, an IPsec session, or a segment.
- 13 Click **Add**.

What to do next

Check that the static route is configured properly. See [Verify the Static Route on a Tier-0 Router](#).

Create a Route Map

A route map consists of a sequence of IP prefix lists, BGP path attributes, and an associated action. The router scans the sequence for an IP address match. If there is a match, the router performs the action and scans no further.

Route maps can be referenced at the BGP neighbor level and for route redistribution.

Prerequisites

- Verify that an IP prefix list or a community list is configured. See [Create an IP Prefix List in Manager Mode](#) or [Create a Community List](#).
- For details about using regular expressions to define route-map match criteria for community lists, see [Using Regular Expressions to Match Community Lists When Adding Route Maps](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Tier-0 Gateways**.
- 3 To edit a tier-0 gateway, click the menu icon (three dots) and select **Edit**.
- 4 Click **Routing**.
- 5 Click **Set** next to **Route Maps**.
- 6 Click **Add Route Map**.
- 7 Enter a name and click **Set** to add match criteria.
- 8 Click **Add Match Criteria** to add one or more match criteria.

9 For each criterion, select **IP Prefix** or **Community List** and click **Set** to specify one or more match expressions.

a If you selected **Community List**, specify match expressions that define how to match members of community lists. For each community list, the following match options are available:

- **MATCH ANY** - perform the set action in the route map if any of the communities in the community list is matched.
- **MATCH ALL** - perform the set action in the route map if all the communities in the community list are matched regardless of the order.
- **MATCH EXACT** - perform the set action in the route map if all the communities in the community list are matched in the exact same order.
- **MATCH COMMUNITY REGEXP** - perform the set action in the route map if all the regular communities associated with the NRLI match the regular expression.
- **MATCH LARGE COMMUNITY REGEXP** - perform the set action in the route map if all the large communities associated with the NRLI match the regular expression.

You should use the match criterion `MATCH_COMMUNITY_REGEX` to match routes against standard communities, and use the match criterion `MATCH_LARGE_COMMUNITY_REGEX` to match routes against large communities. If you want to permit routes containing either the standard community or large community value, you must create two match criteria. If the match expressions are given in the same match criterion, only the routes containing both the standard and large communities will be permitted.

For any match criterion, the match expressions are applied in an AND operation, which means that all match expressions must be satisfied for a match to occur. If there are multiple match criteria, they are applied in an OR operation, which means that a match will occur if any one match criterion is satisfied.

10 Set BGP attributes.

BGP Attribute	Description
AS-path Prepend	Prepend a path with one or more AS (autonomous system) numbers to make the path longer and therefore less preferred.
MED	Multi-Exit Discriminator indicates to an external peer a preferred path to an AS.
Weight	Set a weight to influence path selection. The range is 0 - 65535.

BGP Attribute	Description
Community	Specify a list of communities. For a regular community use the aa:nn format, for example, 300:500. For a large community use the aa:bb:cc format, for example, 11:22:33. Or use the drop-down menu to select one of the following: ■ NO_EXPORT_SUBCONFED - Do not advertise to EBGp peers. ■ NO_ADVERTISE - Do not advertise to any peer. ■ NO_EXPORT - Do not advertise outside BGP confederation
Local Preference	Use this value to choose the outbound external BGP path. The path with the highest value is preferred.

- 11 In the Action column, select **Permit** or **Deny**.

You can permit or deny IP addresses matched by the IP prefix lists or community lists from being advertised.

- 12 Click **Save**.

Using Regular Expressions to Match Community Lists When Adding Route Maps

You can use regular expressions to define the route-map match criteria for community lists. BGP regular expressions are based on POSIX 1003.2 regular expressions.

The following expressions are a subset of the POSIX regular expressions.

Expression	Description
.	Matches any single character.
*	Matches 0 or more occurrences of pattern.
+	Matches 1 or more occurrences of pattern.
?	Matches 0 or 1 occurrence of pattern.
^	Matches the beginning of the line.
\$	Matches the end of the line.
_	This character has special meanings in BGP regular expressions. It matches to a space, comma, AS set delimiters { and } and AS confederation delimiters (and). It also matches to the beginning of the line and the end of the line. Therefore this character can be used for an AS value boundaries match. This character technically evaluates to (^ [,{}()])\$).

Here are some examples for using regular expressions in route maps:

Expression	Description
^101	Matches routes having community attribute that starts with 101.
^[0-9]+	Matches routes having community attribute that starts with a number between 0-9 and has one or more instances of such a number.
.*	Matches routes having any or no community attribute.

Expression	Description
.+	Matches routes having any community value.
^\$	Matches routes having no/null community value.

Configure BGP

To enable access between your VMs and the outside world, you can configure an external or internal BGP (eBGP or iBGP) connection between a tier-0 gateway and a router in your physical infrastructure.

When configuring BGP, you must configure a local Autonomous System (AS) number for the tier-0 gateway. You must also configure the remote AS number. EBGP neighbors must be directly connected and in the same subnet as the tier-0 uplink. If they are not in the same subnet, BGP multi-hop should be used.

BGPv6 is supported for single hop and multihop. A BGPv6 neighbor only supports IPv6 addresses. Redistribution, prefix list, and route maps are supported with IPv6 prefixes.

A tier-0 gateway in active-active mode supports inter-SR (service router) iBGP. If gateway #1 is unable to communicate with a northbound physical router, traffic is re-routed to gateway #2 in the active-active cluster. If gateway #2 is able to communicate with the physical router, traffic between gateway #1 and the physical router will not be affected.

The implementation of ECMP on NSX Edge is based on the 5-tuple of the protocol number, source and destination address, and source and destination port.

The iBGP feature has the following capabilities and restrictions:

- Redistribution, prefix lists, and routes maps are supported.
- Route reflectors are not supported.
- BGP confederation is not supported.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Tier-0 Gateways**.
- 3 To edit a tier-0 gateway, click the menu icon (three dots) and select **Edit**.

4 Click BGP.

- a Enter the local AS number.

In active-active mode, the default ASN value, 65000, is already filled in. In active-standby mode, there is no default ASN value.

- b Click the **BGP** toggle to enable or disable BGP.

In active-active mode, **BGP** is enabled by default. In active-standby mode, **BGP** is disabled by default.

- c If this gateway is in active-active mode, click the **Inter SR iBGP** toggle to enable or disable inter-SR iBGP. It is enabled by default.

If the gateway is in active-standby mode, this feature is not available.

- d Click the **ECMP** toggle button to enable or disable ECMP.

- e Click the **Multipath Relax** toggle button to enable or disable load-sharing across multiple paths that differ only in AS-path attribute values but have the same AS-path length.

Note **ECMP** must be enabled for **Multipath Relax** to work.

- f In the **Graceful Restart** field, select **Disable**, **Helper Only**, or **Graceful Restart and Helper**.

You can optionally change the **Graceful Restart Timer** and **Graceful Restart Stale Timer**.

By default, the Graceful Restart mode is set to **Helper Only**. Helper mode is useful for eliminating and/or reducing the disruption of traffic associated with routes learned from a neighbor capable of Graceful Restart. The neighbor must be able to preserve its forwarding table while it undergoes a restart.

For EVPN, only the **Helper Only** mode is supported.

The Graceful Restart capability is not recommended to be enabled on the tier-0 gateways because BGP peerings from all the gateways are always active. On a failover, the Graceful Restart capability will increase the time a remote neighbor takes to select an alternate tier-0 gateway. This will delay BFD-based convergence.

Note: Unless overridden by neighbor-specific configuration, the tier-0 configuration applies to all BGP neighbors.

5 Configure Route Aggregation by adding IP address prefixes.

- a Click **Add Prefix**.
- b Enter a IP address prefix in CIDR format.
- c For the option **Summary Only**, select **Yes** or **No**.

6 Click Save.

You must save the global BGP configuration before you can configure BGP neighbors.

7 Configure **BGP Neighbors**.

- a Enter the IP address of the neighbor.
- b Enable or disable **BFD**.
- c Enter a value for **Remote AS number**.

For iBGP, enter the same AS number as the one in step 4a. For eBGP, enter the AS number of the physical router.

- d Under **Route Filter**, click **Set** to add one or more route filters.

For **IP Address Family**, you can select **IPv4**, **IPv6**, or **L2VPN EVPN**. You can have at most two route filters, with one address family being **IPv4** and the other being **L2VPN EVPN**. No other combinations (**IPv4** and **IPv6**, **IPv6** and **L2VPN EVPN**) are allowed.

For **Maximum Routes**, you can specify a value between 1 and 1,000,000. This is the maximum number of BGP routes that the gateway will accept from the BGP neighbor.

Note: If you configure a BGP neighbor with one address family, for example, **L2VPN EVPN**, and then later add a second address family, the established BGP connection will be reset.

- e Enable or disable the **Allowas-in** feature.

This is disabled by default. With this feature enabled, BGP neighbors can receive routes with the same AS, for example, when you have two locations interconnected using the same service provider. This feature applies to all the address families and cannot be applied to specific address families.

- f In the **Source Addresses** field, you can select a source address to establish a peering session with a neighbor using this specific source address. If you do not select any, the gateway will automatically choose one.
- g Enter a value for **Max Hop Limit**.

- h In the **Graceful Restart** field, you can optionally select **Disable**, **Helper Only**, or **Graceful Restart and Helper**.

Option	Description
None selected	The Graceful Restart for this neighbor will follow the Tier-0 gateway BGP configuration.
Disable	■ If the tier-0 gateway BGP is configured with Disable, Graceful Restart will be disabled for this neighbor. ■ If the tier-0 gateway BGP is configured with Helper Only, Graceful Restart will be disabled for this neighbor. ■ If the tier-0 gateway BGP is configured with Graceful Restart and Helper, Graceful Restart will be disabled for this neighbor.
Helper Only	■ If the tier-0 gateway BGP is configured with Disable, Graceful Restart will be configured as Helper Only for this neighbor. ■ If the tier-0 gateway BGP is configured with Helper Only, Graceful Restart will be configured as Helper Only for this neighbor. ■ If the tier-0 gateway BGP is configured with Graceful Restart and Helper, Graceful Restart will be configured as Helper Only for this neighbor.
Graceful Restart and Helper	■ If the tier-0 gateway BGP is configured with Disable, Graceful Restart will be configured as Graceful Restart and Helper for this neighbor. ■ If the tier-0 gateway BGP is configured with Helper Only, Graceful Restart will be configured as Graceful Restart and Helper for this neighbor. ■ If the tier-0 gateway BGP is configured with Graceful Restart and Helper, Graceful Restart will be configured as Graceful Restart and Helper for this neighbor.

Note: For EVPN, only the **Helper Only** mode is supported.

- i Click **Timers & Password**.

- j Enter a value for **BFD Interval**.

The unit is milliseconds. For an Edge node running in a VM, the minimum value is 500. For a bare-metal Edge node, the minimum value is 50.

- k Enter a value for **BFD Multiplier**.

- l Enter a value, in seconds, for **Hold Down Time** and **Keep Alive Time**.

The **Keep Alive Time** specifies how frequently KEEPALIVE messages will be sent. The value can be between 0 and 65535. Zero means no KEEPALIVE messages will be sent.

The **Hold Down Time** specifies how long the gateway will wait for a KEEPALIVE message from a neighbor before considering the neighbor dead. The value can be 0 or between 3 and 65535. Zero means no KEEPALIVE messages are sent between the BGP neighbors and the neighbor will never be considered unreachable.

Hold Down Time must be at least three times the value of the **Keep Alive Time**.

- m Enter a password.

This is required if you configure MD5 authentication between BGP peers.

- 8 Click **Save**.

Configure BFD

BFD (Bidirectional Forwarding Detection) is a protocol that can detect forwarding path failures.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Tier-0 Gateways**.
- 3 To edit a tier-0 gateway, click the menu icon (three dots) and select **Edit**.
- 4 Click **Routing** and **Set** for **Static Route BFD Peer**.
- 5 Click **Add Static Route BFD Peer**.
- 6 Select a **BFD profile**. See [Add a BFD Profile](#).
- 7 Enter the peer IP address and optionally the source addresses.
- 8 Click **Save**.

Configure Multicast

IP multicast routing enables a host (source) to send a single copy of data to a single multicast address. Data is then distributed to a group of recipients using a special form of IP address called the IP multicast group address. You can configure multicast on a tier-0 gateway for an IPv4 network to enable multicast routing.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Tier-0 Gateways**.
- 3 To edit a tier-0 gateway, click the menu icon (three dots) and select **Edit**.
- 4 Click the **Multicast** toggle to enable multicast.
- 5 In the **Replication Multicast Range** field, enter an address range in CIDR format.

Replication Multicast Range is a range of multicast group addresses (GENEVE outer destination IP) that is used in the underlay to replicate workload/tenant multicast group addresses. It is recommended that there is no overlap between the Replication Multicast Range and workload/tenant multicast group addresses.
- 6 In the **IGMP Profile** drop-down list, select an IGMP profile.
- 7 In the **PIM Profile** drop-down list, select a PIM profile.

Configure IPv6 Layer 3 Forwarding

IPv4 layer 3 forwarding is enabled by default. You can also configure IPv6 layer 3 forwarding.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Networking Settings**.
- 3 Click the **Global Networking Config** tab.
- 4 Edit the Global Gateway Configuration and select **IPv4 and IPv6** for the **L3 Forwarding Mode**.
IPv6 only is not supported.
- 5 Click **Save**.
- 6 Select **Networking > Tier-0 Gateways**.
- 7 Edit a tier-0 gateway by clicking the menu icon (three dots) and select **Edit**.
- 8 Go to **Additional Settings**.
 - a There are no configurable IPv6 addresses for **Internal Transit Subnet**. The system automatically uses IPv6 link local addresses.
 - b Enter an IPv6 subnet for **TO-T1 Transit Subnets**.
- 9 Go to **Interfaces** and add an interface for IPv6.

Create SLAAC and DAD Profiles for IPv6 Address Assignment

When using IPv6 on a logical router interface, you can set up Stateless Address Autoconfiguration (SLAAC) for the assignment of IP addresses. SLAAC enables the addressing of a host, based on a network prefix advertised from a local network router, through router advertisements. Duplicate Address Detection (DAD) ensures the uniqueness of IP addresses.

Prerequisites

Navigate to **Networking > Networking Settings**, click the **Global Gateway Config** tab and select **IPv4 and IPv6** as the **L3 Forwarding Mode**

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Tier-0 Gateways**.
- 3 To edit a tier-0 gateway, click the menu icon (three dots) and select **Edit**.
- 4 Click **Additional Settings**.

- 5 To create an **ND Profile** (SLAAC profile), click the menu icon (three dots) and select **Create New**.
 - a Enter a name for the profile.
 - b Select a mode:
 - **Disabled** - Router advertisement messages are disabled.
 - **SLAAC with DNS Through RA** - The address and DNS information is generated with the router advertisement message.
 - **SLAAC with DNS Through DHCP** - The address is generated with the router advertisement message and the DNS information is generated by the DHCP server.
 - **DHCP with Address and DNS through DHCP** - The address and DNS information is generated by the DHCP server.
 - **SLAAC with Address and DNS through DHCP** - The address and DNS information is generated by the DHCP server. This option is only supported by NSX Edge and not by KVM hosts or ESXi hosts.
 - c Enter the reachable time and the retransmission interval for the router advertisement message.
 - d Enter the domain name and specify a lifetime for the domain name. Enter these values only for the **SLAAC with DNS Through RA** mode.
 - e Enter a DNS server and specify a lifetime for the DNS server. Enter these values only for the **SLAAC with DNS Through RA** mode.
 - f Enter the values for router advertisement:
 - **RA Interval** - The interval of time between the transmission of consecutive router advertisement messages.
 - **Hop Limit** - The lifetime of the advertised routes.
 - **Router Lifetime** - The lifetime of the router.
 - **Prefix Lifetime** - The lifetime of the prefix in seconds.
 - **Prefix Preferred Time** - The time that a valid address is preferred.
- 6 To create a **DAD Profile**, click the menu icon (three dots) and select **Create New**.
 - a Enter a name for the profile.
 - b Select a mode:
 - **Loose** - A duplicate address notification is received but no action is taken when a duplicate address is detected.
 - **Strict** - A duplicate address notification is received and the duplicate address is no longer used.

- c Enter the **Wait Time (seconds)** that specifies the interval of time between the NS packets.
- d Enter the **NS Retries Count** that specifies the number of NS packets to detect duplicate addresses at intervals defined in **Wait Time (seconds)**

Changing the HA Mode of a Tier-0 Gateway

You can change the high availability (HA) mode of a tier-0 gateway in certain circumstances.

Changing the HA mode is allowed only if there is no more than one service router running on the gateway. This means that you must not have uplinks on more than one Edge transport node. However, you can have more than one uplink on the same Edge transport node.

After you set the HA mode from active-active to active-standby, you can set the failover mode. The default is non-preemptive.

HA mode change is not allowed if the following services or features are configured.

- DNS Forwarder
- IPsec VPN
- L2 VPN
- HA VIP
- Stateful Firewall
- SNAT, DNAT, NO_SNAT, or NO_DNAT
- Reflexive NAT applied on an interface
- Service Insertion
- VRF
- Centralized Service Port

Add a VRF Gateway

A virtual routing and forwarding (VRF) gateway makes it possible for multiple instances of a routing table to exist within the same gateway at the same time. VRFs are the layer 3 equivalent of a VLAN. A VRF gateway must be linked to a tier-0 gateway. From the tier-0 gateway, the VRF gateway inherits the failover mode, Edge cluster, internal transit subnet, T0-T1 transit subnets, and BGP routing configuration.

Prerequisites

For VRF gateways on EVPN, ensure that you configure the EVPN settings for the tier-0 gateway that you want to link to. These settings are only needed to support EVPN:

- Specify a VNI pool on the tier-0 gateway.
- Set the EVPN local tunnel endpoints on the tier-0 gateway.

For more information, see [Configuring EVPN](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Tier-0 Gateway**.
- 3 Click **Add Gateway > VRF**.
- 4 Enter a name for the gateway.
- 5 Select a tier-0 gateway.
- 6 Click **VRF Settings**.

These settings are only needed to support EVPN.

- a Specify a **Route Distinguisher**.

If the connected tier-0 gateway has **RD Admin Address** configured, the **Route Distinguisher** is automatically populated. Enter a new value if you want to override the assigned Route Distinguisher.

- b Specify an **EVPN Transit VNI**.

The VNI must be unique and belong to the VNI pool configured on the linked tier-0 gateway.

- c In the **Route Targets** field, click **Set** to add route targets.

For each route target, select a mode, which can be **Auto** or **Manual**. Specify one or more **Import Route Targets**. Specify one or more **Export Route Targets**.

- 7 Click **Save** and then **Yes** to continue configuring the VRF gateway.
- 8 For VRF-lite, configure one or more external interfaces on the VRF gateway with an **Access VLAN ID** and connect to a VLAN Segment. For EVPN, configure one or more service interfaces on the VRF gateway with an Access VLAN ID and connect to an Overlay Segment. See [Add a Segment](#). VRF interfaces require existing external interfaces on the linked tier-0 gateway to be mapped to each edge node. The Segment connected to the Access interface needs to have VLAN IDs configured in range or list format.
- 9 Click **BGP** to set **BGP**, **ECMP**, **Route Aggregation**, and **BGP Neighbours**. You can add a route filter with IPv4/IPv6 address families. See [Add a Tier-0 Gateway](#).
- 10 Click **Routing** and complete routing configuration. For supporting route leaking between the VRF gateway and linked tier-0 gateway/peer VRF gateway, you can add a static route and select Next Hop scope as the linked tier-0 gateway, or as one of the existing peer VRF gateways. See [Add a Tier-0 Gateway](#).

Configuring EVPN

EVPN (Ethernet VPN) is a standards-based BGP control plane that provides the ability to extend Layer 2 and Layer 3 connectivity between different data centers.

The EVPN feature has the following capabilities and limitations:

- Multi-Protocol BGP (MP-BGP) EVPN between NSX Edge and physical routers.
- VXLAN used as the overlay for MP-BGP EVPN.
- Multi-tenancy in MP-BGP EVPN by using VRF instances.
- Support for EVPN type-5 routes only.
- NSX-T generates unique router MAC for every NSX edge VTEP in the EVPN domain. However, there may be other nodes in the network that are not managed by NSX-T, for example, physical routers. You must make sure that the router MACs are unique across all the VTEPs in the EVPN domain.
- The EVPN feature supports NSX Edge to be either the ingress or the egress of the EVPN virtual tunnel endpoint. If an NSX Edge node receives EVPN type-5 prefixes from its eBGP peer that need to be redistributed to another eBGP peer, the routes will be re-advertised without any change to the nexthop.
- In a multi-path network topology, it is recommended that you do not enable ECMP on the gateway that has EVPN configured.

Configuration Prerequisites

- Virtual Router (vRouter) deployed on VMware ESXi hypervisor.
- Peer physical router supporting EVPN type-5 routes.

Configuration Steps

- Create a VNI pool. See [Add a VNI Pool](#).
- Configure a VLAN Segment. See [Add a Segment](#).
- Configure an overlay Segment and specify one or more VLAN ranges. See [Add a Segment](#).
- Configure a tier-0 gateway to support EVPN. See [Add a Tier-0 Gateway](#).
- Under EVPN Settings, select a VNI pool and create EVPN Tunnel Endpoints.
- Under Route Distinguisher for VRF Gateways, configure RD Admin Address for the automatic route distinguisher use case.
- Configure one or more external interfaces on the tier-0 gateway and connect to the VLAN Segment.
- Configure BGP neighbors with the peer physical router. Add route filter with IPv4 and L2VPN EVPN Address Families.

- Configure Route Re-Distribution. Select EVPN TEP IP under Tier-0 Subnets along with other sources.
- Configure VRF to support EVPN. See [Add a VRF Gateway](#).
- Under VRF Settings, specify an EVPN Transit VNI.
- Specify Route Distinguisher for a manual route distinguisher.
- Specify Import/Export Route Targets for manual route targets.
- Add service interface on VRF for each edge node and connect to the Overlay Segment. Specify an Access VLAN ID for each service interface.
- Configure per VRF BGP neighbors with the peer vRouter. The routes learned over the VRF BGP sessions are redistributed by the NSX Edge to the peer physical router over the MP-BGP EVPN session.

Tier-1 Gateway

3

A tier-1 gateway has downlink connections to segments and uplink connections to tier-0 gateways.

You can configure route advertisements and static routes on a tier-1 gateway. Recursive static routes are supported.

Read the following topics next:

- [Add a Tier-1 Gateway](#)

Add a Tier-1 Gateway

A tier-1 gateway is typically connected to a tier-0 gateway in the northbound direction and to segments in the southbound direction.

If you are adding a tier-1 gateway from Global Manager in NSX Federation, see [Add a Tier-1 Gateway from Global Manager](#).

Tier-0 and tier-1 gateways support the following addressing configurations for all interfaces (uplinks, service ports and downlinks) in both single tier and multi-tiered topologies:

- IPv4 only
- IPv6 only
- Dual Stack - both IPv4 and IPv6

To use IPv6 or dual stack addressing, enable **IPv4 and IPv6** as the L3 Forwarding Mode in **Networking > Networking Settings > Global Networking Config**.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Tier-1 Gateways**.
- 3 Click **Add Tier-1 Gateway**.
- 4 Enter a name for the gateway.
- 5 (Optional) Select a tier-0 gateway to connect to this tier-1 gateway to create a multi-tier topology.

- 6 (Optional) Select an NSX Edge cluster if you want this tier-1 gateway to host stateful services such as NAT, load balancer, or firewall.

If an NSX Edge cluster is selected, a service router will always be created (even if you do not configure stateful services), affecting the north/south traffic pattern.

- 7 (Optional) In the Edges field, click **Set** to select an NSX Edge node.
- 8 If you selected an NSX Edge cluster, select a failover mode or accept the default.

Option	Description
Preemptive	If the preferred NSX Edge node fails and recovers, it will preempt its peer and become the active node. The peer will change its state to standby.
Non-preemptive	If the preferred NSX Edge node fails and recovers, it will check if its peer is the active node. If so, the preferred node will not preempt its peer and will be the standby node. This is the default option.

- 9 If you plan to configure a load balancer on this gateway, select an **Edges Pool Allocation Size** setting according to the size of the load balancer.

The options are **Routing**, **LB Small**, **LB Medium**, **LB Large**, and **LB XLarge**. The default is **Routing** and is suitable if no load balancer will be configured on this gateway. This parameter allows the NSX Manager to place the tier-1 gateway on the Edge nodes in a more intelligent way. With this setting the number of load balancing and routing functions on each node is taken into consideration. Note that after you create the gateway, you can change this setting if you have not configured a load balancer.

- 10 (Optional) Click the **Enable StandBy Relocation** toggle to enable or disable standby relocation.

Standby relocation means that if the Edge node where the active or standby logical router is running fails, a new standby logical router is created on another Edge node to maintain high availability. If the Edge node that fails is running the active logical router, the original standby logical router becomes the active logical router and a new standby logical router is created. If the Edge node that fails is running the standby logical router, the new standby logical router replaces it.

- 11 (Optional) Click **Route Advertisement**.

Select one or more of the following:

- **All Static Routes**
- **All NAT IP's**
- **All DNS Forwarder Routes**
- **All LB VIP Routes**
- **All Connected Segments and Service Ports**
- **All LB SNAT IP Routes**

■ All IPsec Local Endpoints

12 Click **Save**.

13 (Optional) Click **Route Advertisement**.

- a In the **Set Route Advertisement Rules** field, click **Set** to add route advertisement rules.

14 (Optional) Click **Additional Settings**.

- a For IPv6, you can select or create an **ND Profile** and a **DAD Profile**.

These profiles are used to configure Stateless Address Autoconfiguration (SLAAC) and Duplicate Address Detection (DAD) for IPv6 addresses.

- b Select an **Ingress QoS Profile** and an **Egress QoS Profile** for traffic limitations.

These profiles are used to set information rate and burst size for permitted traffic. See [Add a Gateway QoS Profile](#) for more information on creating QoS profiles.

If this gateway is linked to a tier-0 gateway, the **Router Links** field shows the link addresses.

15 (Optional) Click **Service Interfaces** and **Set** to configure connections to segments. Required in some topologies such as VLAN-backed segments or one-arm load balancing.

- a Click **Add Interface**.
- b Enter a name and IP address in CIDR format.
- c Select a segment.
- d In the **MTU** field, enter a value between 64 and 9000.
- e For **URPF Mode**, you can select **Strict** or **None**.
URPF (Unicast Reverse Path Forwarding) is a security feature.
- f Add one or more tags.
- g In the **ND Profile** field, select or create a profile.
- h Click **Save**.

16 (Optional) Click **Static Routes** and **Set** to configure static routes.

- a Click **Add Static Route**.
- b Enter a name and a network address in the CIDR or IPv6 CIDR format.
- c Click **Set Next Hops** to add next hop information.
- d Click **Save**.

What to do next

After the tier-1 gateway is added, you can optionally enable dynamic IP management on the gateway by selecting either a DHCP server profile or a DHCP relay profile. For more information, see [Attach a DHCP Profile to a Tier-0 or Tier-1 Gateway](#).

Segments

4

In NSX-T Data Center, segments are virtual layer 2 domains. A segment was earlier called a logical switch.

There are two types of segments in NSX-T Data Center:

- VLAN-backed segments
- Overlay-backed segments

A VLAN-backed segment is a layer 2 broadcast domain that is implemented as a traditional VLAN in the physical infrastructure. This means that traffic between two VMs on two different hosts but attached to the same VLAN-backed segment is carried over a VLAN between the two hosts. The resulting constraint is that you must provision an appropriate VLAN in the physical infrastructure for those two VMs to communicate at layer 2 over a VLAN-backed segment.

In an overlay-backed segment, traffic between two VMs on different hosts but attached to the same overlay segment have their layer 2 traffic carried by a tunnel between the hosts. NSX-T Data Center instantiates and maintains this IP tunnel without the need for any segment-specific configuration in the physical infrastructure. As a result, the virtual network infrastructure is decoupled from the physical network infrastructure. That is, you can create segments dynamically without any configuration of the physical network infrastructure.

The default number of MAC addresses learned on an overlay-backed segment is 2048. The default MAC limit per segment can be changed through the API field `remote_overlay_mac_limit` in `MacLearningSpec`. For more information see the `MacSwitchingProfile` in the *NSX-T Data Center API Guide*.

Read the following topics next:

- [Segment Profiles](#)
- [Add a Segment](#)
- [Types of DHCP on a Segment](#)
- [Configure DHCP on a Segment](#)
- [Configure DHCP Static Bindings on a Segment](#)
- [Layer 2 Bridging](#)
- [Add a Metadata Proxy Server](#)

Segment Profiles

Segment profiles include Layer 2 networking configuration details for segments and segment ports. NSX Manager supports several types of segment profiles.

The following types of segment profiles are available:

- QoS (Quality of Service)
- IP Discovery
- SpoofGuard
- Segment Security
- MAC Management

Note You cannot edit or delete the default segment profiles. If you require alternate settings from what is in the default segment profile you can create a custom segment profile. By default all custom segment profiles except the segment security profile will inherit the settings of the appropriate default segment profile. For example, a custom IP discovery segment profile by default will have the same settings as the default IP discovery segment profile.

Each default or custom segment profile has a unique identifier. You use this identifier to associate the segment profile to a segment or a segment port.

A segment or segment port can be associated with only one segment profile of each type. You cannot have, for example, two QoS segment profiles associated with a segment or segment port.

If you do not associate a segment profile when you create a segment, then the NSX Manager associates a corresponding default system-defined segment profile. The children segment ports inherit the default system-defined segment profile from the parent segment.

When you create or update a segment or segment port you can choose to associate either a default or a custom segment profile. When the segment profile is associated or disassociated from a segment the segment profile for the children segment ports is applied based on the following criteria.

- If the parent segment has a profile associated with it, the child segment port inherits the segment profile from the parent.
- If the parent segment does not have a segment profile associated with it, a default segment profile is assigned to the segment and the segment port inherits that default segment profile.
- If you explicitly associate a custom profile with a segment port, then this custom profile overrides the existing segment profile.

Note If you have associated a custom segment profile with a segment, but want to retain the default segment profile for one of the child segment port, then you must make a copy of the default segment profile and associate it with the specific segment port.

You cannot delete a custom segment profile if it is associated to a segment or a segment port. You can find out whether any segments and segment ports are associated with the custom segment profile by going to the Assigned To section of the Summary view and clicking on the listed segments and segment ports.

Understanding QoS Segment Profile

QoS provides high-quality and dedicated network performance for preferred traffic that requires high bandwidth. The QoS mechanism does this by prioritizing sufficient bandwidth, controlling latency and jitter, and reducing data loss for preferred packets even when there is a network congestion. This level of network service is provided by using the existing network resources efficiently.

For this release, shaping and traffic marking namely, CoS and DSCP is supported. The Layer 2 Class of Service (CoS) allows you to specify priority for data packets when traffic is buffered in the segment due to congestion. The Layer 3 Differentiated Services Code Point (DSCP) detects packets based on their DSCP values. CoS is always applied to the data packet irrespective of the trusted mode.

NSX-T Data Center trusts the DSCP setting applied by a virtual machine or modifying and setting the DSCP value at the segment level. In each case, the DSCP value is propagated to the outer IP header of encapsulated frames. This enables the external physical network to prioritize the traffic based on the DSCP setting on the external header. When DSCP is in the trusted mode, the DSCP value is copied from the inner header. When in the untrusted mode, the DSCP value is not preserved for the inner header.

Note DSCP settings work only on tunneled traffic. These settings do not apply to traffic inside the same hypervisor.

You can use the QoS switching profile to configure the average ingress and egress bandwidth values to set the transmit limit rate. The peak bandwidth rate is used to support burst traffic a segment is allowed to prevent congestion on the northbound network links. These settings do not guarantee the bandwidth but help limit the use of network bandwidth. The actual bandwidth you will observe is determined by the link speed of the port or the values in the switching profile, whichever is lower.

The QoS switching profile settings are applied to the segment and inherited by the child segment port.

Create a QoS Segment Profile

You can define the DSCP value and configure the ingress and egress settings to create a custom QoS switching profile.

Prerequisites

- Familiarize yourself with the QoS switching profile concept. See [Understanding QoS Switching Profile](#).

- Identify the network traffic you want to prioritize.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Segments > Segment Profiles**.
- 3 Click **Add Segment Profile** and select **QoS**.
- 4 Complete the QoS switching profile details.

Option	Description
Name	Name of the profile.
Mode	Select either a Trusted or Untrusted option from the Mode drop-down menu. When you select the Trusted mode the inner header DSCP value is applied to the outer IP header for IP/IPv6 traffic. For non IP/IPv6 traffic, the outer IP header takes the default value. Trusted mode is supported on an overlay-based logical port. The default value is 0. Untrusted mode is supported on overlay-based and VLAN-based logical port. For the overlay-based logical port, the DSCP value of the outbound IP header is set to the configured value irrespective to the inner packet type for the logical port. For the VLAN-based logical port, the DSCP value of IP/IPv6 packet will be set to the configured value. The DSCP values range for untrusted mode is between 0 to 63. Note DSCP settings work only on tunneled traffic. These settings do not apply to traffic inside the same hypervisor.
Priority	Set the CoS priority value. The priority values range from 0 to 63, where 0 has the highest priority.
Class of Service	Set the CoS value. CoS is supported on VLAN-based logical port. CoS groups similar types of traffic in the network and each type of traffic is treated as a class with its own level of service priority. The lower priority traffic is slowed down or in some cases dropped to provide better throughput for higher priority traffic. CoS can also be configured for the VLAN ID with zero packet. The CoS values range from 0 to 7, where 0 is the best effort service.
Ingress	Set custom values for the outbound network traffic from the VM to the logical network. You can use the average bandwidth to reduce network congestion. The peak bandwidth rate is used to support burst traffic and the burst size is based on the duration with peak bandwidth. You set burst duration in the burst size setting. You cannot guarantee the bandwidth. However, you can use the Average, Peak, and Burst Size settings to limit network bandwidth. For example, if the average bandwidth is 30 Mbps, peak bandwidth is 60 Mbps, and the allowed duration is 0.1 second, then the burst size is $60 * 1000000 * 0.10 / 8 = 750000$ Bytes. The default value 0 disables rate limiting on the ingress traffic.

Option	Description
Ingress Broadcast	Set custom values for the outbound network traffic from the VM to the logical network based on broadcast. For example, when you set the average bandwidth for a logical switch to 3000 Kbps, peak bandwidth is 6000 Kbps, and the allowed duration is 0.1 second, then the burst size is $6000 * 1000 * 0.10/8 = 75000$ Bytes. The default value 0 disables rate limiting on the ingress broadcast traffic.
Egress	Set custom values for the inbound network traffic from the logical network to the VM. The default value 0 disables rate limiting on the egress traffic.

If the ingress, ingress broadcast, and egress options are not configured, the default values are used.

- 5 Click **Save**.

Understanding IP Discovery Segment Profile

IP Discovery uses DHCP and DHCPv6 snooping, ARP (Address Resolution Protocol) snooping, ND (Neighbor Discovery) snooping, and VM Tools to learn MAC and IP addresses.

Note IP discovery methods for IPv6 are disabled in the default IP discovery segment profile. To enable IP discovery for IPv6 for segments, you must create an IP discovery profile with the IPv6 options enabled and attach the profile to the segments. In addition, make sure that distributed firewall allows IPv6 Neighbor Discovery packets between all workloads (allowed by default).

The discovered MAC and IP addresses are used to achieve ARP/ND suppression, which minimizes traffic between VMs connected to the same segment. The number of IPs in the ARP/ND suppression cache for any given port is determined by the settings in the port's IP Discovery profile. The relevant settings are ARP Binding Limit, ND Snooping Limit, Duplicate IP Detection, ARP ND Binding Limit Timeout, and Trust on First Use (TOFU).

The discovered MAC and IP addresses are also used by the SpoofGuard and distributed firewall (DFW) components. DFW uses the address bindings to determine the IP address of objects in firewall rules.

DHCP/DHCPv6 snooping inspects the DHCP/DHCPv6 packets exchanged between the DHCP/DHCPv6 client and server to learn the IP and MAC addresses.

ARP snooping inspects the outgoing ARP and GARP (gratuitous ARP) packets of a VM to learn the IP and MAC addresses.

VM Tools is software that runs on an ESXi-hosted VM and can provide the VM's configuration information including MAC and IP or IPv6 addresses. This IP discovery method is available for VMs running on ESXi hosts only.

ND snooping is the IPv6 equivalent of ARP snooping. It inspects neighbor solicitation (NS) and neighbor advertisement (NA) messages to learn the IP and MAC addresses.

Duplicate address detection checks whether a newly discovered IP address is already present on the realized binding list for a different port. This check is performed for ports on the same segment. If a duplicate address is detected, the newly discovered address is added to the discovered list, but is not added to the realized binding list. All duplicate IPs have an associated discovery timestamp. If the IP that is on the realized binding list is removed, either by adding it to the ignore binding list or by disabling snooping, the duplicate IP with the oldest timestamp is moved to the realized binding list. The duplicate address information is available through an API call.

By default, the discovery methods ARP snooping and ND snooping operate in a mode called trust on first use (TOFU). In TOFU mode, when an address is discovered and added to the realized bindings list, that binding remains in the realized list forever. TOFU applies to the first 'n' unique <IP, MAC, VLAN> bindings discovered using ARP/ND snooping, where 'n' is the binding limit that you can configure. You can disable TOFU for ARP/ND snooping. The methods will then operate in trust on every use (TOEU) mode. In TOEU mode, when an address is discovered, it is added to the realized bindings list and when it is deleted or expired, it is removed from the realized bindings list. DHCP snooping and VM Tools always operate in TOEU mode.

Note TOFU is not the same as SpoofGuard, and it does not block traffic in the same way as SpoofGuard. For more information, see [Understanding SpoofGuard Segment Profile](#).

For Linux VMs, the ARP flux problem might cause ARP snooping to obtain incorrect information. The problem can be prevented with an ARP filter. For more information, see <http://linux-ip.net/html/ether-arp.html#ether-arp-flux>.

For each port, NSX Manager maintains an ignore bindings list, which contains IP addresses that cannot be bound to the port. If you navigate to **Networking > Logical Switches > Ports** in **Manager** mode and select a port, you can add discovered bindings to the ignore bindings list. You can also delete an existing discovered or realized binding by copying it to **Ignore Bindings**.

Create an IP Discovery Segment Profile

NSX-T Data Center has several default IP Discovery segment profiles. You can also create additional ones.

Prerequisites

Familiarize yourself with the IP Discovery segment profile concepts. See [Understanding IP Discovery Segment Profile](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Segments > Segment Profiles**.
- 3 Click **Add Segment Profile** and select **IP Discovery**.

4 Specify the IP Discovery segment profile details.

Option	Description
Name	Enter a name.
ARP Snooping	For an IPv4 environment. Applicable if VMs have static IP addresses.
ARP Binding Limit	The maximum number of IPv4 IP addresses that can be bound to a port. The minimum value allowed is 1 and the maximum is 256. The default is 1.
ARP ND Binding Limit Timeout	The timeout value, in minutes, for IP addresses in the ARP/ND binding table if TOFU is disabled. If an address times out, a newly discovered address replaces it.
DHCP Snooping	For an IPv4 environment. Applicable if VMs have IPv4 addresses.
DHCP Snooping - IPv6	For an IPv6 environment. Applicable if VMs have IPv6 addresses.
VM Tools	Available for ESXi-hosted VMs only.
VM Tools - IPv6	Available for ESXi-hosted VMs only.
ND Snooping	For an IPv6 environment. Applicable if VMs have static IP addresses.
ND Snooping Limit	The maximum number of IPv6 addresses that can be bound to a port.
Trust on First Use	Applicable to ARP and ND snooping.
Duplicate IP Detection	For all snooping methods and both IPv4 and IPv6 environments.

5 Click **Save**.

Understanding SpoofGuard Segment Profile

SpoofGuard helps prevent a form of malicious attack called "web spoofing" or "phishing." A SpoofGuard policy blocks traffic determined to be spoofed.

SpoofGuard is a tool that is designed to prevent virtual machines in your environment from sending traffic with an IP address it is not authorized to send traffic from. In the instance that a virtual machine's IP address does not match the IP address on the corresponding logical port and segment address binding in SpoofGuard, the virtual machine's vNIC is prevented from accessing the network entirely. SpoofGuard can be configured at the port or segment level. There are several reasons SpoofGuard might be used in your environment:

- Preventing a rogue virtual machine from assuming the IP address of an existing VM.
- Ensuring the IP addresses of virtual machines cannot be altered without intervention – in some environments, it's preferable that virtual machines cannot alter their IP addresses without proper change control review. SpoofGuard facilitates this by ensuring that the virtual machine owner cannot simply alter the IP address and continue working unimpeded.
- Guaranteeing that distributed firewall (DFW) rules will not be inadvertently (or deliberately) bypassed – for DFW rules created utilizing IP sets as sources or destinations, the possibility always exists that a virtual machine could have its IP address forged in the packet header, thereby bypassing the rules in question.

NSX-T Data Center SpoofGuard configuration covers the following:

- MAC SpoofGuard - authenticates MAC address of packet
- IP SpoofGuard - authenticates MAC and IP addresses of packet
- Dynamic Address Resolution Protocol (ARP) inspection, that is, ARP and Gratuitous Address Resolution Protocol (GARP) SpoofGuard and Neighbor Discovery (ND) SpoofGuard validation are all against the MAC source, IP Source and IP-MAC source mapping in the ARP/GARP/ND payload.

At the port level, the allowed MAC/VLAN/IP allow-list is provided through the Address Bindings property of the port. When the virtual machine sends traffic, it is dropped if its IP/MAC/VLAN does not match the IP/MAC/VLAN properties of the port. The port level SpoofGuard deals with traffic authentication, i.e. is the traffic consistent with VIF configuration.

At the segment level, the allowed MAC/VLAN/IP allow-list is provided through the Address Bindings property of the segment. This is typically an allowed IP range/subnet for the segment and the segment level SpoofGuard deals with traffic authorization.

Traffic must be permitted by port level AND segment level SpoofGuard before it will be allowed into segment. Enabling or disabling port and segment level SpoofGuard, can be controlled using the SpoofGuard segment profile.

Create a SpoofGuard Segment Profile

When SpoofGuard is configured, if the IP address of a virtual machine changes, traffic from the virtual machine may be blocked until the corresponding configured port/segment address bindings are updated with the new IP address.

Enable SpoofGuard for the port group(s) containing the guests. When enabled for each network adapter, SpoofGuard inspects packets for the prescribed MAC and its corresponding IP address.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Segments > Segment Profiles**.
- 3 Click **Add Segment Profile** and select **Spoof Guard**.
- 4 Enter a name.
- 5 To enable port level SpoofGuard, set **Port Bindings** to **Enabled**.
- 6 Click **Save**.

Understanding Segment Security Segment Profile

Segment security provides stateless Layer2 and Layer 3 security by checking the ingress traffic to the segment and dropping unauthorized packets sent from VMs by matching the IP address, MAC address, and protocols to a set of allowed addresses and protocols. You can use segment

security to protect the segment integrity by filtering out malicious attacks from the VMs in the network.

Note that the default segment security profile has the DHCP settings **Server Block** and **Server Block - IPv6** enabled. This means that a segment that uses the default segment security profile will block traffic from a DHCP server to a DHCP client. If you want a segment that allows DHCP server traffic, you must create a custom segment security profile for the segment.

Create a Segment Security Segment Profile

You can create a custom segment security segment profile with MAC destination addresses from the allowed BPDU list and configure rate limiting.

Prerequisites

Familiarize yourself with the segment security segment profile concept. See [Understanding Switch Security Switching Profile](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Segments > Segment Profiles**.
- 3 Click **Add Segment Profile** and select **Segment Security**.
- 4 Complete the segment security profile details.

Option	Description
Name	Name of the profile.
BPDU Filter	Toggle the BPDU Filter button to enable BPDU filtering. Disabled by default. When the BPDU filter is enabled, all of the traffic to BPDU destination MAC address is blocked. The BPDU filter when enabled also disables STP on the logical switch ports because these ports are not expected to take part in STP.
BPDU Filter Allow List	Click the destination MAC address from the BPDU destination MAC addresses list to allow traffic to the permitted destination. You must enable BPDU Filter to be able to select from this list.
DHCP Filter	Toggle the Server Block button and Client Block button to enable DHCP filtering. Both are disabled by default. DHCP Server Block blocks traffic from a DHCP server to a DHCP client. Note that it does not block traffic from a DHCP server to a DHCP relay agent. DHCP Client Block prevents a VM from acquiring a DHCP IP address by blocking DHCP requests.

Option	Description
DHCPv6 Filter	Toggle the Server Block - IPv6 button and Client Block - IPv6 button to enable DHCP filtering. Both are disabled by default. DHCPv6 Server Block blocks traffic from a DHCPv6 server to a DHCPv6 client. Note that it does not block traffic from a DHCP server to a DHCP relay agent. Packets whose UDP source port number is 547 are filtered. DHCPv6 Client Block prevents a VM from acquiring a DHCP IP address by blocking DHCP requests. Packets whose UDP source port number is 546 are filtered.
Block Non-IP Traffic	Toggle the Block Non-IP Traffic button to allow only IPv4, IPv6, ARP, and BPDU traffic. The rest of the non-IP traffic is blocked. The permitted IPv4, IPv6, ARP, GARP and BPDU traffic is based on other policies set in address binding and SpoofGuard configuration. By default, this option is disabled to allow non-IP traffic to be handled as regular traffic.
RA Guard	Toggle the RA Guard button to filter out ingress IPv6 router advertisements. ICMPv6 type 134 packets are filtered out. This option is enabled by default.
Rate Limits	Set a rate limit for broadcast and multicast traffic. This option is enabled by default. Rate limits can be used to protect the logical switch or VMs from events such as broadcast storms. To avoid any connectivity problems, the minimum rate limit value must be ≥ 10 pps.

5 Click **Save**.

Understanding MAC Discovery Segment Profile

The MAC discovery segment profile supports two functionalities: MAC learning and MAC address change.

The MAC address change feature allows a VM to change its MAC address. A VM connected to a port can run an administrative command to change the MAC address of its vNIC and still send and receive traffic on that vNIC. This feature is supported on ESXi only and not on KVM. In the default MAC discovery segment profile, this property is enabled by default.

MAC learning provides network connectivity to deployments where multiple MAC addresses are configured behind one vNIC, for example, in a nested hypervisor deployment where an ESXi VM runs on an ESXi host and multiple VMs run inside the ESXi VM. Without MAC learning, when the ESXi VM's vNIC connects to a segment port, its MAC address is static. VMs running inside the ESXi VM do not have network connectivity because their packets have different source MAC addresses. With MAC learning, the vSwitch inspects the source MAC address of every packet coming from the vNIC, learns the MAC address and allows the packet to go through. If a MAC address that is learned is not used for a certain period of time, it is removed. This time period is not configurable. The field **MAC Learning Aging Time** displays the pre-defined value, which is 600.

MAC learning also supports unknown unicast flooding. Normally, when a packet that is received by a port has an unknown destination MAC address, the packet is dropped. With unknown unicast flooding enabled, the port floods unknown unicast traffic to every port on the switch that has MAC learning and unknown unicast flooding enabled. This property is enabled by default, but only if MAC learning is enabled.

The number of MAC addresses that can be learned is configurable. The maximum value is 4096, which is the default. You can also set the policy for when the limit is reached. The options are:

- **Drop** - Packets from an unknown source MAC address are dropped. Packets inbound to this MAC address will be treated as unknown unicast. The port will receive the packets only if it has unknown unicast flooding enabled.
- **Allow** - Packets from an unknown source MAC address are forwarded although the address will not be learned. Packets inbound to this MAC address will be treated as unknown unicast. The port will receive the packets only if it has unknown unicast flooding enabled.

If you enable MAC learning or MAC address change, to improve security, configure SpoofGuard as well.

Create a MAC Discovery Segment Profile

You can create a MAC discovery segment profile to manage MAC addresses.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Segments > Segment Profiles**.
- 3 Click **Add Segment Profile** and select **MAC Discovery**.
- 4 Complete the MAC discovery profile details.

Option	Description
Name	Name of the profile.
MAC Change	Enable or disable the MAC address change feature. The default is disabled.
MAC Learning	Enable or disable the MAC learning feature. The default is disabled.
MAC Limit Policy	Select Allow or Drop . The default is Allow . This option is available if you enable MAC learning
Unknown Unicast Flooding	Enable or disable the unknown unicast flooding feature. The default is enabled. This option is available if you enable MAC learning
MAC Limit	Set the maximum number of MAC addresses. The default is 4096. This option is available if you enable MAC learning
MAC Learning Aging Time	For information only. This option is not configurable. The pre-defined value is 600.

- 5 Click **Save**.

Add a Segment

You can add two kinds of segments: overlay-backed segments and VLAN-backed segments.

Segments are created as part of a transport zone. There are two types of transport zones: VLAN transport zones and overlay transport zones. A segment created in a VLAN transport zone is a VLAN-backed segment, and a segment created in an overlay transport zone is an overlay-backed segment.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Segments**.
- 3 Click **Add Segment**.
- 4 Enter a name for the segment.
- 5 Select the type of connectivity for the segment.

Connectivity	Description
None	Select this option when you do not want to connect the segment to any upstream gateway (tier-0 or tier-1). Typically, you want to add a standalone segment in the following scenarios: ■ When you want to create a local testing environment for users that are running workloads on the same subnet. ■ When east-west connectivity with users on the other subnets is not necessary. ■ When north-south connectivity to users outside the data center is not necessary. ■ When you want to configure layer 2 bridging or guest VLAN tagging.
Tier-1	Select this option when you want to connect the segment to a tier-1 gateway.
Tier-0	Select this option when you want to connect the segment to a tier-0 gateway.

Note You can change the connectivity of a gateway-connected segment from one gateway to another gateway (same or different gateway type). In addition, you can change the connectivity of segment from "None" to a tier-0 or tier-1 gateway. The segment connectivity changes are permitted only when the gateways and the connected segments are in the same transport zone. However, if the segment has DHCP configured on it, some restrictions and caveats apply on changing the segment connectivity. For more information, see [Scenarios: Impact of Changing Segment Connectivity on DHCP](#).

- 6 Enter the Gateway IP address of the subnet in a CIDR format. A segment can contain an IPv4 subnet, or an IPv6 subnet, or both.
 - If a segment is not connected to a gateway, subnet is optional.

- If a segment is connected either to a tier-1 or tier-0 gateway, subnet is required.

Subnets of one segment must not overlap with the subnets of other segments in your network. A segment is always associated with a single virtual network identifier (VNI) regardless of whether it is configured with one subnet, two subnets, or no subnet.

- 7 Select a transport zone, which can be an overlay or a VLAN.

To create a VLAN-backed segment, add the segment in a VLAN transport zone. Similarly, to create an overlay-backed segment, add the segment in an overlay transport zone.

- 8 (Optional) To configure DHCP on the segment, click **Set DHCP Config**.

For detailed steps on configuring **DHCP Settings** and **DHCP Options**, see [Configure DHCP on a Segment](#).

- 9 If the transport zone is of type VLAN, specify a list of VLAN IDs. If the transport zone is of type Overlay, and you want to support layer 2 bridging or guest VLAN tagging, specify a list of VLAN IDs or VLAN ranges

- 10 (Optional) Select an uplink teaming policy for the segment.

This drop-down menu displays the named teaming policies, if you have added them in the VLAN transport zone. If no uplink teaming policy is selected, the default teaming policy is used.

- Named teaming policies are not applicable to overlay segments. Overlay segments always follow the default teaming policy.
- For VLAN-backed segments, you have the flexibility to override the default teaming policy with a selected named teaming policy. This capability is provided so that you can steer the infrastructure traffic from the host to specific VLAN segments in the VLAN transport zone. Before adding the VLAN segment, ensure that the named teaming policy names are added in the VLAN transport zone.

- 11 (Optional) Enter the fully qualified domain name.

DHCPv4 server and DHCPv4 static bindings on the segment automatically inherit the domain name from the segment configuration as the Domain Name option.

- 12 If you want to use Layer 2 VPN to extend the segment, click the **L2 VPN** text box and select an L2 VPN server or client session.

You can select more than one.

- 13 In **VPN Tunnel ID**, enter a unique value that is used to identify the segment.

- 14 (Optional) In the **Metadata Proxy** field, click **Set** to attach or detach a metadata proxy to this segment.

To attach a metadata proxy, select an existing metadata proxy. To detach a metadata proxy, deselect the metadata proxy that is selected.

- 15 Click **Save**.

16 To add segment ports, click **Yes** when prompted if you want to continue configuring the segment.

- a Click **Ports** and **Set**.
- b Click **Add Segment Port**.
- c Enter a port name.
- d For **ID**, enter the VIF UUID of the VM or server that connects to this port.
- e Select a type: **Child**, or **Static**.

Leave this text box blank except for use cases such as containers or VMware HCX. If this port is for a container in a VM, select **Child**. If this port is for a bare metal container or server, select **Static**.

- f Enter a context ID.

Enter the parent VIF ID if **Type** is **Child**, or transport node ID if **Type** is **Static**.

- g Enter a traffic tag.

Enter the VLAN ID in container and other use cases.

- h Select an address allocation method: **IP Pool**, **MAC Pool**, **Both**, or **None**.

- i Specify tags.

- j Apply address binding by specifying the IP (IPv4 address, IPv6 address, or IPv6 subnet) and MAC address of the logical port to which you want to apply address binding. For example, for IPv6, 2001::/64 is an IPv6 subnet, 2001::1 is a host IP, whereas 2001::1/64 is an invalid input. You can also specify a VLAN ID.

Manual address bindings, if specified, override the auto discovered address bindings.

- k Select segment profiles for this port.

17 To select segment profiles, click **Segment Profiles**.

18 (Optional) To bind a static IP address to the MAC address of a VM on the segment, expand **DHCP Static Bindings**, and then click **Set**.

Both DHCP for IPv4 and DHCP for IPv6 static bindings are supported. For detailed steps on configuring static binding settings, see [Configure DHCP Static Bindings on a Segment](#).

19 Click **Save**.

Types of DHCP on a Segment

NSX-T Data Center supports three types of DHCP on a segment: DHCP local server, Gateway DHCP, and DHCP relay.

DHCP Local Server

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. The IP address of a local DHCP server must be in the subnet that is configured on the segment.

Gateway DHCP

It 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. The IP address of a Gateway DHCP server can be different from the subnets that are configured in the segments.

DHCP Relay

It is a DHCP relay service that is local to the segment and not available to the other segments in the network. The DHCP relay service relays the DHCP requests of the VMs that are attached to the segment to the remote DHCP servers. The remote DHCP servers can be in any subnet, outside the SDDC, or in the physical network.

You can configure DHCP on each segment regardless of whether the segment is connected to a gateway. Both DHCP for IPv4 (DHCPv4) and DHCP for IPv6 (DHCPv6) servers are supported.

For a gateway-connected segment, all the three DHCP types are supported. However, Gateway DHCP is supported only in the IPv4 subnet of a segment.

For a standalone segment that is not connected to a gateway, only local DHCP server is supported.

The following restrictions apply to DHCPv6 server configuration on an IPv6 subnet:

- Segments configured with an IPv6 subnet can have either a local DHCPv6 server or a DHCPv6 relay. Gateway DHCPv6 is not supported.
- DHCPv6 Options (classless static routes and generic options) are not supported.

Configure DHCP on a Segment

You can configure DHCP on each segment regardless of whether the segment is connected to a gateway. Both DHCP for IPv4 (DHCPv4) and DHCP for IPv6 (DHCPv6) servers are supported.

For a gateway-connected segment, all the following DHCP types are supported:

- DHCP local server
- DHCP relay
- Gateway DHCP (supported only for IPv4 subnets in a segment)

For a standalone segment that is not connected to a gateway, only local DHCP server is supported.

The following restrictions apply to DHCPv6 server configuration on an IPv6 subnet:

- Segments configured with an IPv6 subnet can have either a local DHCPv6 server or a DHCPv6 relay. Gateway DHCPv6 is not supported.
- DHCPv6 Options (classless static routes and generic options) are not supported.

Prerequisites

- DHCP profile is added in the network.
- If you are configuring Gateway DHCP on a segment, a DHCP profile must be attached to the directly connected tier-1 or tier-0 gateway.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Segments**.
- 3 Either add or edit a segment.
 - To configure a new segment, click **Add Segment**.
 - To modify the properties of an existing segment, click the vertical ellipses next to the name of an existing segment, and then click **Edit**.
- 4 If you are adding a segment, ensure that the following segment properties are specified:
 - Segment name
 - Connectivity
 - Transport zone
 - Subnets (required for a gateway-connected segment, optional for a standalone segment)If you are editing an existing segment, go directly to the next step.
- 5 Click **Set DHCP Config**.

6 In the **DHCP Type** drop-down menu, select any one of the following types.

On a segment, IPv6 and IPv4 subnets always use the same DHCP type. Mixed configuration is not supported.

DHCP Type	Description
DHCP Local 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 standalone 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. When you use a DHCP relay on a segment, you cannot configure DHCP Settings and DHCP Options. The UI does not prevent you from configuring DHCP static bindings. However, in NSX-T Data Center 3.0, static binding with a DHCP relay is an unsupported configuration.
Gateway DHCP	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. If the IPv4 subnet uses Gateway DHCP, you cannot configure DHCPv6 in the IPv6 subnet of the same segment because Gateway DHCPv6 is not supported. In this case, the IPv6 subnet cannot support any DHCPv6 server configuration, including the IPv6 static bindings.

Note In NSX-T Data Center 3.0 and 3.0.1, after the segment is created and the DHCP service is in use, you cannot change the DHCP type of a gateway-connected segment. Starting in version 3.0.2, you can change the DHCP type of a gateway-connected segment.

7 In the **DHCP Profile** drop-down menu, select the name of the DHCP server profile or DHCP relay profile.

- If the segment is connected to a gateway, Gateway DHCP server is selected by default. The DHCP profile that is attached to the gateway is autoselected. The name and server IP address are fetched automatically from that DHCP profile and displayed in a read-only mode.

When a segment is using a Gateway DHCP server, ensure that an edge cluster is selected either in the gateway, or DHCP server profile, or both. If an edge cluster is unavailable in either the profile or the gateway, an error message is displayed when you save the segment.

- If you are configuring a local DHCP server or a DHCP relay on the segment, you must select a DHCP profile from the drop-down menu. If no profiles are available in the drop-down menu, click the vertical ellipses icon and create a DHCP profile. After the profile is created, it is automatically attached to the segment.

When a segment is using a local DHCP server, ensure that the DHCP server profile contains an edge cluster. If an edge cluster is unavailable in the profile, an error message is displayed when you save the segment.

Note In NSX-T Data Center 3.0 and 3.0.1, after the segment is created and the DHCP service is in use, you cannot change the DHCP profile of the segment. Starting in version 3.0.2, you can change the DHCP profile of the segment that uses a DHCP local server or a DHCP relay.

8 Click the **IPv4 Server** or **IPv6 Server** tab.

If the segment contains an IPv4 subnet and an IPv6 subnet, you can configure both DHCPv4 and DHCPv6 servers on the segment.

9 Configure the DHCP settings.

- 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.

- c (Optional) In the **DHCP Ranges** text box, enter one or more IP address ranges.

Both IP ranges and IP addresses are allowed. IPv4 addresses must be in a CIDR /32 format, and IPv6 addresses must be in a CIDR /128 format. You can also enter an IP address as a range by entering the same IP address in the start and the end of the range. For example, 172.16.10.10-172.16.10.10.

Ensure that DHCP ranges meet the following requirements:

- IP addresses in the DHCP ranges must belong to the subnet that is configured on the segment. That is, 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.

Note The following types of IPv6 addresses are not permitted in DHCP for IPv6 ranges:

- Link Local Unicast addresses (FE80::/64)
 - Multicast addresses (FF00::/8)
 - Unspecified address (0:0:0:0:0:0:0:0)
 - Address with all F (F:F:F:F:F:F:F)
-

Caution After a DHCP server is created, you can update existing ranges, append new IP ranges, or delete existing ranges. However, it is a good practice to avoid deleting, shrinking, or expanding the existing IP ranges. For example, do not try to combine multiple smaller IP ranges to create a single large IP range. You might accidentally miss including IP addresses, which are already leased to the DHCP clients from the larger IP range. Therefore, when you modify existing ranges after the DHCP service is running, it might cause the DHCP clients to lose network connection and result in a temporary traffic disruption.

- d (Optional) (Only for DHCPv6): In the **Excluded Ranges** text box, enter IPv6 addresses or a range of IPv6 addresses that you want to exclude for dynamic IP assignment to DHCPv6 clients.

In IPv6 networks, the DHCP ranges can be large. Sometimes, you might want to reserve certain IPv6 addresses, or multiple small ranges of IPv6 addresses from the large DHCP range for static binding. In such situations, you can specify excluded ranges.

- e (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.

- f (Optional) (Only for DHCPv6): Enter the preferred time in seconds.

Preferred time is the length of time that a valid IP address is preferred. When the preferred time expires, the IP address becomes deprecated. If no value is entered, preferred time is autocalculated as (lease time * 0.8).

Lease time must be > preferred time.

Valid range of values is 60–4294967295. Default is 69120.

- g (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.

- h (Optional) (Only for DHCPv6): Enter one or more domain names.

DHCPv4 configuration automatically fetches the domain name that you specified in the segment configuration.

- i (Optional) (Only for DHCPv6): Enter the IP address of the Simple Network Time Protocol (SNTP) server. A maximum of two SNTP servers are permitted.

In NSX-T Data Center 3.0, DHCPv6 server does not support NTP.

DHCPv4 server supports only NTP. To add an NTP server, click **Options**, and add the Generic Option (Code 42 - NTP Servers).

- 10 (Optional) Click **Options**, and specify the Classless Static Routes (Option 121) and Generic Options.

In NSX-T Data Center 3.0, DHCP Options for IPv6 are not supported.

- 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. For information about classless static routes in DHCPv4, see RFC 3442 specifications. 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.

- 11 Click **Apply** to save the DHCP configuration, and then click **Save** to save the segment configuration.

What to do next

- After a segment has DHCP configured on it, some restrictions and caveats apply on changing the segment connectivity. For more information, see [Scenarios: Impact of Changing Segment Connectivity on DHCP](#).

- When a DHCP server profile is attached to a segment that uses a DHCP local server, the DHCP service is created in the edge cluster that you specified in the DHCP profile. However, if the segment uses a Gateway DHCP server, the edge cluster in which the DHCP service is created depends on a combination of several factors. For a detailed information about how an edge cluster is selected for DHCP service, see [Scenarios: Selection of Edge Cluster for DHCP Service](#).

Configure DHCP Static Bindings on a Segment

You can configure static bindings on both DHCP for IPv4 and DHCP for IPv6 servers.

In a typical network environment, you have VMs that run services, such as FTP, email servers, application servers, and so on. You might not want the IP address of these VMs to change in your network. In this case, you can bind a static IP address to the MAC address of each VM (DHCP client). The static IP address must not overlap with the DHCP IP ranges and the DHCP Server IP addresses.

DHCP static bindings are allowed when you are configuring either a local DHCP server or a Gateway DHCP server on the segment. The UI does not prevent you from configuring DHCP static bindings when the segment is using a DHCP relay. However, in NSX-T Data Center 3.0, static binding with a DHCP relay is an unsupported configuration.

Prerequisites

The segment on which you want to configure DHCP static bindings must be saved in the network.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Segments**.
- 3 Next to the segment that you want to edit, click the vertical ellipses, and click **Edit**.
- 4 Expand the **DHCP Static Bindings** section, and next to **DHCP Static Bindings**, click **Set**.

By default, the **IPv4 Static Binding** page is displayed. To bind IPv6 addresses, make sure that you first click the **IPv6 Static Binding** tab, and then proceed to the next step.

5 Click **Add Static Binding**.

a Specify the DHCP static binding options.

The following table describes the static binding options that are common to DHCP for IPv4 and DHCP for IPv6 servers.

Option	Description
Name	Enter a unique display name to identify each static binding. The name must be limited to 255 characters.
IP Address	■ Required for IPv4 static binding. Enter a single IPv4 address to bind to the MAC address of the client. ■ Optional for IPv6 static binding. Enter a single Global Unicast IPv6 address to bind to the MAC address of the client. When no IPv6 address is specified for static binding, Stateless Address Autoconfiguration (SLAAC) is used to auto-assign an IPv6 address to the DHCPv6 clients. Also, you can use Stateless DHCP to assign other DHCP configuration options, such as DNS, domain names, and so on, to the DHCPv6 clients. For more information about Stateless DHCP for IPv6, read the RFC 3376 specifications. The following types of IPv6 addresses are not permitted in IPv6 static binding: ■ Link Local Unicast addresses (FE80::/64) ■ Multicast IPv6 addresses (FF00::/8) ■ Unspecified address (0:0:0:0:0:0:0:0) ■ Address with all F (F:F:F:F:F:F:F) The static IP address must belong to the subnet (if any) that is configured on the segment.
MAC Address	Required. Enter the MAC address of the DHCP client to which you want to bind a static IP address. The following validations apply to MAC address in static bindings: ■ MAC address must be unique in all the static bindings on a segment that uses a local DHCP server. ■ MAC address must be unique in all the static bindings across all the segments that are connected to the gateway and which use the Gateway DHCP server. For example, consider that you have 10 segments connected to a tier-1 gateway. You use a Gateway DHCP server for four segments (Segment1 to Segment4), and a local DHCP server for the remaining six segments (Segment5 to Segment10). Assume that you have a total of 20 static bindings across all the four segments (Segment1 to Segment4), which use the Gateway DHCP server. In addition, you have five static bindings in each of the other six segments (Segment5 to Segment10), which use a local DHCP server. In this example: ■ The MAC address in each of the 20 static bindings must be unique across all the segments (Segment1 to Segment4) that use the Gateway DHCP server. ■ The MAC address in the five static bindings must be unique on each segment (Segment5 to Segment10) that use a local DHCP server.

Option	Description
Lease Time	Optional. Enter the amount of time in seconds for which the IP address is bound to the DHCP client. When the lease time expires, the IP address becomes invalid and the DHCP server can assign the address to other DHCP clients on the segment. Valid range of values is 60–4294967295. Default is 86400.
Description	Optional. Enter a description for the static binding.
Tags	Optional. Add tags to label static bindings so that you can quickly search or filter bindings, troubleshoot and trace binding-related issues, or do other tasks. For more information about adding tags and use cases for tagging objects, see Tags .

The following table describes the static binding options that are available only in a DHCP for IPv4 server.

DHCP For IPv4 Option	Description
Gateway Address	Enter the default gateway IP address that the DHCP for IPv4 server must provide to the DHCP client.
Host Name	Enter the host name of the DHCP for IPv4 client so that the DHCPv4 server can always bind the client (host) with the same IPv4 address each time. The host name must be limited to 63 characters. The following validations apply to host name in static bindings: ■ Host name must be unique in all the static bindings on a segment that uses a local DHCP server. ■ Host name must be unique in all the static bindings across all the segments that are connected to the gateway and which use the Gateway DHCP server. For example, consider that you have 10 segments connected to a tier-1 gateway. You use a Gateway DHCP server for four segments (Segment1 to Segment4), and a local DHCP server for the remaining six segments (Segment5 to Segment10). Assume that you have a total of 20 static bindings across all the four segments (Segment1 to Segment4), which use the Gateway DHCP server. In addition, you have five static bindings in each of the other six segments (Segment5 to Segment10), which use a local DHCP server. In this example: ■ The host name in each of the 20 static bindings must be unique across all the segments (Segment1 to Segment4) that use the Gateway DHCP server. ■ The host name in the five static bindings must be unique on each segment (Segment5 to Segment10) that use a local DHCP server.
DHCP Options	Optional. Click Set to configure DHCP for IPv4 Classless Static Routes and other Generic Options.

Some additional notes for DHCPv4 static binding:

- IPv4 static bindings automatically inherit the domain name that you configured on the segment.

- To specify DNS servers in the static binding configuration, add the Generic Option (Code 6 - DNS Servers).
- To synchronize the system time on DHCPv4 clients with DHCPv4 servers, use NTP. DHCP for IPv4 server does not support SNTP.
- If DHCP Options are not specified in the static bindings, the DHCP Options from the DHCPv4 server on the segment are automatically inherited in the static bindings. However, if you have explicitly added one or more DHCP Options in the static bindings, these DHCP Options are not autoinherited from the DHCPv4 server on the segment.

The following table describes the static binding options that are available only in a DHCP for IPv6 server.

DHCP For IPv6 Option	Description
DNS Servers	Optional. Enter a maximum of two domain name servers to use for the name resolution. When not specified, no DNS is assigned to the DHCP client.
SNTP Servers	Optional. Enter a maximum of two Simple Network Time Protocol (SNTP) servers. The clients use these SNTP servers to synchronize their system time to that of the standard time servers.
Preferred Time	Optional. Enter the length of time that a valid IP address is preferred. When the preferred time expires, the IP address becomes deprecated. If no value is entered, preferred time is auto-calculated as (lease time * 0.8). Lease time must be > preferred time. Valid range of values is 60–4294967295. Default is 69120.
Domain Names	Optional. Enter the domain name to provide to the DHCPv6 clients. Multiple domain names are supported in an IPv6 static binding. When not specified, no domain name is assigned to the DHCP clients.

Some additional notes for DHCPv6 static binding:

- Gateway IP address configuration is unavailable in IPv6 static bindings. IPv6 client learns about its first-hop router from the ICMPv6 router advertisement (RA) message.
 - NTP is not supported in DHCPv6 static bindings.
- b Click **Save** after configuring each static binding.

Layer 2 Bridging

With layer 2 bridging, you can have a connection to a VLAN-backed port group or a device, such as a gateway, that resides outside of your NSX-T Data Center deployment. A layer 2 bridge is also useful in a migration scenario, in which you need to split a subnet across physical and virtual workloads.

A layer 2 bridge requires an Edge cluster and an Edge Bridge profile. An Edge Bridge profile specifies which Edge cluster to use for bridging and which Edge transport node acts as the primary and backup bridge. When you configure a segment, you can specify an Edge bridge profile to enable layer 2 bridging.

Create an Edge Bridge Profile

An Edge bridge profile makes an NSX Edge cluster capable of providing layer 2 bridging to a segment.

When you create an edge bridge profile, if you set the failover mode to be preemptive and a failover occurs, the standby node becomes the active node. After the failed node recovers, it becomes the active node again. If you set the failover mode to be non-preemptive and a failover occurs, the standby node becomes the active node. After the failed node recovers, it becomes the standby node. You can manually set the standby edge node to be the active node by running the CLI command `set l2bridge-port <uuid> state active` on the standby edge node. The command can only be applied in non-preemptive mode. Otherwise, there will be an error. In non-preemptive mode, the command will trigger an HA failover when applied on a standby node, and it will be ignored when applied on an active node. For more information, see the *NSX-T Data Center Command-Line Interface Reference*.

Prerequisites

- Verify that you have an NSX Edge cluster with two NSX Edge transport nodes.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Segments > Edge Bridge Profiles**.
- 3 Click **Add Edge Bridge Profile**.
- 4 Enter a name for the Edge bridge profile and optionally a description.
- 5 Select an NSX Edge cluster.
- 6 Select a primary node.
- 7 Select a backup node.
- 8 Select a failover mode.

The options are **Preemptive** and **Non-Preemptive**.

- 9 Click **Save**.

What to do next

Configure Edge-based bridging. See [Configure Edge-Based Bridging](#).

Configure Edge-Based Bridging

When you configure Edge-based bridging, after creating an Edge bridge profile for an Edge cluster, some additional configurations are required for an Edge node running in a VM.

Note that bridging a segment twice on the same Edge node is not supported. However, you can bridge two VLANs to the same segment on two different Edge nodes.

Depending on your environment, choose one of the following options.

Note If you are bridging a segment to VLAN 0 and you are using a distributed router on this segment, VLAN 0 traffic might not be routed by the gateway when using MAC learning. In this scenario you should avoid option 3 (and even option 2a if the edge VM is attached to the portgroup of a VDS prepared for NSX for vSphere).

Option 1: Edge VM is on a VSS portgroup

This option is for when the Edge VM is connected to a VSS (vSphere Standard Switch). You must enable promiscuous mode and forged transmit.

- Set promiscuous mode on the portgroup.
- Allow forged transmit on the portgroup.
- Run the following command to enable reverse filter on the ESXi host where the Edge VM is running:

```
esxcli system settings advanced set -o /Net/ReversePathFwdCheckPromisc -i 1
```

Then disable and enable promiscuous mode on the portgroup with the following steps:

- Edit the portgroup's settings.
- Disable promiscuous mode and save the settings.
- Edit the portgroup's settings again.
- Enable promiscuous mode and save the settings.
- Do not have other port groups in promiscuous mode on the same host sharing the same set of VLANs.
- The active and standby Edge VMs should be on different hosts. If they are on the same host the throughput might be reduced because VLAN traffic needs to be forwarded to both VMs in promiscuous mode.

Option 2a: Edge VM is on a VDS 6.6.0 (or later) portgroup

This option is for when the Edge VM is connected to a VDS (vSphere Distributed Switch). You must be running ESXi 6.7 or later, and VDS 6.6.0 or later.

- Enable MAC learning with the option “allow unicast flooding” on the portgroup using the VIM API `DVSMacLearningPolicy` and setting `allowUnicastFlooding` to `true`.

Option 2b: Edge VM is on a VDS 6.5.0 (or later) portgroup

This option is for when the Edge VM is connected to a VDS (vSphere Distributed Switch). You enable promiscuous mode and forged transmit.

- Set promiscuous mode on the portgroup.
- Allow forged transmit on the portgroup.
- Run the following command to enable reverse filter on the ESXi host where the Edge VM is running:

```
esxcli system settings advanced set -o /Net/ReversePathFwdCheckPromisc -i 1
```

Then disable and enable promiscuous mode on the portgroup with the following steps:

- Edit the portgroup's settings.
- Disable promiscuous mode and save the settings.
- Edit the portgroup's settings again.
- Enable promiscuous mode and save the settings.
- Do not have other port groups in promiscuous mode on the same host sharing the same set of VLANs.
- The active and standby Edge VMs should be on different hosts. If they are on the same host the throughput might be reduced because VLAN traffic needs to be forwarded to both VMs in promiscuous mode.

Option 3: Edge VM is connected to an NSX-T segment

If the Edge is deployed on a host with NSX-T installed, it can connect to a VLAN segment and use MAC Learning which is the preferred configuration option.

- Create a new MAC Discovery segment profile by navigating to **Networking > Segments > Segment Profiles**.
 - Click **Add Segment Profile > MAC Discovery > .**
 - Enable **MAC Learning**. This will also enable **Unknown Unicast Flooding**. Keep the flooding option enabled for bridging to work in all scenarios.
- Edit the segment used by the Edge by navigating to **Networking > Segments**.
 - Click the menu icon (3 dots) and select **Edit** to edit the segment.
 - In the **Segment Profiles** section, set the **MAC Discovery** profile to the one created above.

What to do next

Associate a segment with the bridge profile. See [Create a Layer 2 Bridge-Backed Segment](#).

Create a Layer 2 Bridge-Backed Segment

When you have VMs that are connected to the NSX-T Data Center overlay, you can configure a bridge-backed segment to provide layer 2 connectivity with other devices or VMs that are outside of your NSX-T Data Center deployment.

Prerequisites

- Verify that you have an Edge bridge profile.
- Configure one of the following options: promiscuous mode, MAC learning, or a sink port. See [Configure Edge-Based Bridging](#).
- At least one ESXi or KVM host to serve as a regular transport node. This node has hosted VMs that require connectivity with devices outside of a NSX-T Data Center deployment.
- A VM or another end device outside of the NSX-T Data Center deployment. This end device must be attached to a VLAN port matching the VLAN ID of the bridge-backed segment.
- One segment in an overlay transport zone to serve as the bridge-backed segment.

Procedure

- 1 From a browser, log in to an NSX Manager at `https://<nsx-mgr>`.
- 2 Select **Networking > Segments**
- 3 Click the menu icon (three dots) of the overlay segment that you want to configure layer 2 bridging on and select **Edit**.
- 4 In the **Edge Bridges** field, click **Set**.
- 5 Click **Add Edge Bridge**.
You can add one or more Edge bridge profiles.
- 6 Select an Edge bridge profile.
- 7 Select a transport zone.
- 8 Enter a VLAN ID or a VLAN trunk specification (specify VLAN ranges and not individual VLANs).
- 9 (Optional) Select a teaming policy.
- 10 Click **Add**.

Results

You can test the functionality of the bridge by sending a ping from a VM attached to the segment to a device that is external to the NSX-T deployment.

Add a Metadata Proxy Server

A metadata proxy server enables VMs to retrieve metadata from an OpenStack Nova API server.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Segments > Metadata Proxies**.
- 3 Click **Add Metadata Proxy**.
- 4 Enter a name for the metadata proxy server.
- 5 In the **Server Address** field, enter the URL and port for the Nova server.
The valid port range is 3000 - 9000.
- 6 Select an Edge cluster.
- 7 (Optional) Select Edge nodes.
If you select any Edge node, you cannot enable **Standby Relocation** in the next step.
- 8 (Optional) Enable **Standby Relocation**.
Standby relocation means that if the Edge node running the metadata proxy fails, the metadata proxy will run on a standby Edge node. You can only enable standby relocation if you do not select any Edge node.
- 9 In the **Shared Signature Secret** field, enter the secret that the metadata proxy will use to access the Nova server.
- 10 (Optional) Select a certificate for encrypted communication with the Nova server.
- 11 (Optional) Select a cryptographic protocol.
The options are TLSv1, TLSv1.1, and TLSv1.2. TLSv1.1 and TLSv1.2 are supported by default.

Host Switches

5

A host switch managed object is a virtual network switch that provides networking services to the various hosts in the network. It is instantiated on every host that participates in NSX-T networking

The following host switches are supported in NSX-T:

- **NSX-T Virtual Distributed Switch:** NSX-T introduces a host switch that normalizes connectivity among various compute domains, including multiple VMware vCenter Server instances, KVM, containers, and other off premises or cloud implementations.

NSX-T Virtual Distributed Switch can be configured based on the performance required in your environment:

- **Standard:** Configured for regular workloads, where normal traffic throughput is expected on the workloads.
- **Enhanced:** Configured for telecom workloads, where high traffic throughput is expected on the workloads.
- **vSphere Distributed Virtual Switch:** Provides centralized management and monitoring of the networking configuration of all hosts that are associated with the switch in a vCenter Server environment.

Read the following topics next:

- [Managing NSX-T on a vSphere Distributed Switch](#)
- [Enhanced Networking Stack](#)
- [Migrate Host Switch to vSphere Distributed Switch](#)
- [NSX Virtual Distributed Switch](#)

Managing NSX-T on a vSphere Distributed Switch

You can configure and run NSX-T on a vSphere Distributed Switch (VDS) switch.

In NSX-T 3.0, a host transport node can be prepared by installing NSX-T on a VDS switch. To prepare an NSX Edge VM as a transport node, you can only use an N-VDS switch. But, you can connect an NSX Edge VM to any of the supported switches (VSS, VDS, or N-VDS) depending on the topology in your network.

After you prepare a cluster of transport node hosts with VDS as the host switch, you can do the following:

- Manage NSX-T transport nodes on a VDS switch.
- Realize a segment created in NSX-T as an NSX Distributed Virtual port group in vCenter Server.
- Migrate VMs between vSphere Distributed Virtual port groups and NSX Distributed Virtual port groups.
- Send VMs traffic running on both these type of port groups.

Configuring a vSphere Distributed Switch

When a transport node is configured on a VDS host switch, some network parameters can only be configured in vCenter Server.

The following requirements must be met to install NSX-T on a VDS host switch:

- vCenter Server 7.0 or a later version
- ESXi 7.0 or a later version

The created VDS switch can be configured to centrally manage networking for NSX-T hosts.

Configuring a VDS switch for NSX-T networking requires objects to be configured on NSX-T and in vCenter Server.

- In vSphere:
 - Create a VDS switch.
 - Set MTU to at least 1600
 - Add ESXi hosts to the switch. These hosts are later prepared as NSX-T transport nodes.
 - Assign uplinks to physical NICs.
- In NSX-T:
 - When configuring a transport node, map uplinks created in NSX-T uplink profile with uplinks in VDS.

For more details on preparing a host transport node on a VDS switch, see the *NSX-T Data Center Installation Guide*.

The following parameters can only be configured in a vCenter Server on a VDS backed host switch:


Configuration	VDS	NSX-T	Description
MTU	In vCenter Server, set an MTU value on the switch. Note A VDS switch must have an MTU of 1600 or higher. In vCenter Server, select VDS, click Actions → Settings → Edit Settings.	Any MTU value set in an NSX-T uplink profile is overridden.	As a host transport node that is prepared using VDS as the host switch, the MTU value needs to be set on the VDS switch in vCenter Server.
Uplinks/LAGs	In vCenter Server, configure Uplinks/LAGs on a VDS switch. In vCenter Server, select VDS, click Actions → Settings → Edit Settings.	When a transport node is prepared, the teaming policy on NSX-T is mapped to uplinks/LAGs configured on a VDS switch. Note A transport node prepared on an N-VDS switch, the teaming policy is mapped to physical NICs.	As a host transport node that is prepared using VDS as the host switch, the uplink or LAG are configured on the VDS switch. During configuration, NSX-T requires teaming policy be configured for the transport node. This teaming policy is mapped to the uplinks/LAGs configured on the VDS switch.
NIOC	Configure in vCenter Server. In vCenter Server, select VDS, click Actions → Settings → Edit Settings.	NIOC configuration is not available when a host transport node is prepared using a VDS switch.	As a host transport node that is prepared using VDS as the host switch, the NIOC profile can only be configured in vCenter Server.
Link Layer Discovery Protocol (LLDP)	Configure in vCenter Server. In vCenter Server, select VDS, click Actions → Settings → Edit Settings.	LLDP configuration is not available when a host transport node is prepared using a VDS switch.	As a host transport node that is prepared using VDS as the host switch, the LLDP profile can only be configured in vCenter Server.
Add or Manage Hosts	Manage in vCenter Server. In vCenter Server, go to Networking → VDS Switch → Add and Manage Host..	Prepared as transport nodes in NSX-T.	Before preparing a transport node using a VDS switch, that node must be added to the VDS switch in vCenter Server.

Note NIOC profiles, Link Layer Discovery Protocol (LLDP) profile, and Link Aggregation Group (LAG) for these virtual machines are managed by VDS switches and not by NSX-T. As a vSphere administrator, configure these parameters from vCenter Server UI or by calling VDS API commands.

After preparing a host transport node with VDS as a host switch, the host switch type displays VDS as the host switch. It displays the configured uplink profile in NSX-T and the associated transport zones.



In vCenter Server, the VDS switch used to prepare NSX-T hosts is created as an NSX



Switch.


Managing NSX Distributed Virtual Port Groups

A transport node prepared with VDS as a host switch ensures that segments created in NSX-T is realized as an NSX Distributed Virtual port group on a VDS switch and Segment in NSX-T.

In earlier versions of NSX-T Data Center, a segment created in NSX-T are represented as an opaque network in vCenter Server. When running NSX-T on a VDS switch, a segment is represented as an NSX Distributed Virtual Port Groups.

Any changes to the segments on the NSX-T network are synchronized in vCenter Server.

In vCenter Server, an NSX Distributed Virtual Port Group is represented as .



Any NSX-T segment created in NSX-T is realized in vCenter Server as an NSX-T object. A vCenter Server displays the following details related to NSX-T segments:

- NSX Manager
- Virtual network identifier of the segment
- Transport zone
- Attached virtual machines

The port binding for the segment is by default set to **Ephemeral**. Switching parameters for the switch that are set in NSX-T cannot be edited in vCenter Server and conversely.

Important In a vCenter Server, an NSX Distributed Virtual port group realized does not require a unique name to differentiate it with other port groups on a VDS switch. So, multiple NSX Distributed Virtual port groups can have the same name. Any vSphere automations that use port group names might result in errors.

In vCenter Server, you can perform these actions on an NSX Distributed Virtual Port Group:

- Add VMkernel Adapters.
- Migrate VMs to Another Network.

However, NSX-T objects related to an NSX Distributed Virtual port group can only be edited in NSX Manager. You can edit these segment properties:


- Replication Mode for the segment
- VLAN trunk ID used by the segment
- Switching Profiles (for example, Port Mirroring)
- Ports created on the segment

Logical Switch	ID	Admin Status	Logical Ports	Traffic Type	Config State	Transport Zone
☒ 3-switch-1533	d370...e16c	Up	4	Overlay : 73729	Success	3-transportzone-164
☐ 4-switch-351	b1b3...a09a	Up	4	VLAN : 21	Success	4-transportzone-679

For details on how to configure a vSphere Distributed Virtual port group, refer to the *vSphere Networking Guide*.

NSX-T Cluster Prepared with VDS

An example of an NSX-T cluster prepared using VDS as the host switch.



In the sample topology diagram, two VDS switches are configured to manage NSX-T traffic and vSphere traffic.

VDS-1 and VDS-2 are configured to manage networking for ESXi hosts from Cluster-1, Cluster-2, and Cluster-3. Cluster-1 is prepared to run only vSphere traffic, whereas, Cluster-2 and Cluster-3 are prepared as host transport nodes with these VDS switches.

In vCenter Server, uplink port groups on VDS switches are assigned physical NICs. In the topology, uplinks on VDS-1 and VDS-2 are assigned to physical NICs. Depending on the hardware configuration of the ESXi host, you might want to plan out how many physical NICs to be assigned to the switch. In addition to assigning uplinks to the VDS switch, MTU, NIOC, LLDP, LAG profiles are configured on the VDS switches.

After VDS switches are configured, in NSX-T, add an uplink profile.

When preparing a cluster by applying a transport node profile (on a VDS switch), the uplinks from the transport node profile is mapped to VDS uplinks. In contrast, when preparing a cluster on an N-VDS switch, the uplinks from the transport node profile is directly mapped to physical NICs.

After preparing the clusters, ESXi hosts on cluster-2 and Cluster-3 manage NSX-T traffic, while cluster-1 manage vSphere traffic.

APIs to Configure vSphere Distributed Switch on NSX-T Data Center

NSX-T Data Center API commands to support vSphere Distributed Switch on NSX-T Data Center.

API Changes for vSphere Distributed Switch

For detailed information related to API calls, see the *NSX-T Data Center API Guide*.

Note Configuration done using API commands is also possible from the vCenter Server user interface. For more information on creating a NSX-T Data Center transport node using Sphere Distributed Switch as host switch, refer to the *Configure a Managed Host Transport Node* topic in the *NSX-T Data Center Installation Guide*.

API	NSX-T on vSphere Distributed Switch (VDS)
Create a Transport node for a Discovered node.	<pre data-bbox="740 239 1596 1656">/api/v1/fabric/discovered-nodes/<external-id/discovered-node-id>?action=create { "node_id": "d7ef478b-752c-400a-b5f0-207c04567e5d", "host_switch_spec": { "host_switches": [{ "host_switch_name": "vds-1", "host_switch_id": "50 2b 92 54 e0 80 d8 d1-ee ab 8d a6 7b fd f9 4b", "host_switch_type": "VDS", "host_switch_mode": "STANDARD", "host_switch_profile_ids": [{ "key": "UplinkHostSwitchProfile", "value": "159353ae-c572-4aca-9469-9582480a7467" }], "pnics": [], "uplinks": [{ "vds_uplink_name": "Uplink 2", "uplink_name": "nsxuplink1" }], "is_migrate_pnics": false, "ip_assignment_spec": { "resource_type": "AssignedByDhcp" }, "cpu_config": [], "transport_zone_endpoints": [{ "transport_zone_id": "06ba5326-67ac-4f2c-9953-a8c5d326b51e", "transport_zone_profile_ids": [{ "resource_type": "BfdHealthMonitoringProfile", "profile_id": "52035bb3-ab02-4a08-9884-18631312e50a" }] }], "vmk_install_migration": [], "pnics_uninstall_migration": [], "vmk_uninstall_migration": [], "not_ready": false }], "resource_type": "StandardHostSwitchSpec" }, "transport_zone_endpoints": [], "maintenance_mode": "DISABLED", "is_overridden": false, "resource_type": "TransportNode", "display_name": "TestTN", }</pre>
VM configuration	vim.vm.device.VirtualEthernetCard.DistributedVirtualPortBackingInfo

API	NSX-T on vSphere Distributed Switch (VDS)
VMkernel NIC	<code>vim.dvs.DistributedVirtualPort</code>
Physical NIC to Uplink Mapping	API: <code>vim.host.NetworkSystem:networkSystem.updateNetworkConfig</code> Property: <code>vim.host.NetworkConfig.proxySwitch</code>
MTU	API: <code>vim.dvs.VmwareDistributedVirtualSwitch.reconfigure</code> Property: <code>VmwareDistributedVirtualSwitch.ConfigSpec.maxMtu</code>
LAG	API: <code>vim.dvs.VmwareDistributedVirtualSwitch.updateLacpGroupConfig</code> Property: <code>vim.dvs.VmwareDistributedVirtualSwitch.LacpGroupSpec</code>
NIOC	API: <code>vim.dvs.VmwareDistributedVirtualSwitch.reconfigure</code> Property: <code>vim.dvs.VmwareDistributedVirtualSwitch.ConfigSpec.infrastructureT</code>
LLDP	API: <code>vim.dvs.VmwareDistributedVirtualSwitch.reconfigure</code> Property: <code>vim.dvs.VmwareDistributedVirtualSwitch.ConfigSpec.linkDiscoveryPr</code>

Feature Support in a vSphere Distributed Switch Enabled to Support NSX-T Data Center

Comparison of features supported by a VDS switch version earlier to 7.0 and VDS version 7.0 or later (NSX-T Data Center enabled).

IPFIX and Port Mirroring

An NSX-T Data Center transport node prepared with a VDS switch supports IPFIX, port mirroring.

See [Port Mirroring on a vSphere Distributed Switch](#).

See [IPFIX Monitoring on a vSphere Distributed Switch](#).

SR-IOV support

SR-IOV is supported on a vSphere Distributed Switch but not on a NSX Virtual Distributed Switch.

Feature	NSX Virtual Distributed Switch	vSphere Distributed Switch
SR-IOV	No	Yes (vSphere 7.0 and later)

Stateless Cluster Host Profile Support

Feature	NSX Virtual Distributed Switch	vSphere Distributed Switch
Host Profile Stateless	Yes	Yes (vSphere 7.0 and later) No (when VMkernel adapters are connected to NSX Port Group on vSphere Distributed Switch).

Distributed Resource Scheduler Support

Source Host	Destination Host	DRS (NIOC Configured)	vSphere
vSphere Distributed Switch-A	vSphere Distributed Switch-B	No	No
Opaque Network (N-VDS-A)	Opaque Network (N-VDS-B)	Yes	6.7
vSphere Distributed Switch	Opaque Network (N-VDS)	Yes	7.0
vSphere Distributed Switch-A	vSphere Distributed Switch-A	Yes	7.0
Opaque Network (N-VDS)	vSphere Distributed Switch	No	No

vMotion Support

vMotion between source vSphere Distributed Switch and destination vSphere Distributed Switch. Both VDS switches are enabled to support NSX-T Data Center.

Source / VDS	Destination / VDS	Compute vMotion	Storage vMotion
vSphere Distributed Switch-A (vCenter Server -A)	vSphere Distributed Switch-A (vCenter Server -A)	Yes	Yes
vSphere Distributed Switch-A (vCenter Server -A)	vSphere Distributed Switch-B (vCenter Server -A)	Yes	Yes
vSphere Distributed Switch-A (vCenter Server -A)	vSphere Distributed Switch-B (vCenter Server -B)	Yes	Yes
Segment-A (vCenter Server -A)	Segment-B (vCenter Server -A)	No	No
Segment-A (vCenter Server -A)	Segment-B (vCenter Server -B)	No	No
Transport Zone-A	Transport Zone-B	No	No
NSX-T Data Center-A	NSX-T Data Center-B	No	No

vMotion between vSphere Distributed Switch (NSX-T Data Center enabled) and NSX Virtual Distributed Switch

Source / VDS	Destination / NSX Virtual Distributed Switch	Compute vMotion	Storage vMotion
vCenter Server-A	vCenter Server-A	Yes	Yes
vCenter Server-A	vCenter Server-B	Yes	Yes
Segment-A (vCenter Server-A)	Segment-B (vCenter Server-A)	No	No
Segment-A (vCenter Server-A)	Segment-B (vCenter Server-B)	No	No
Transport Zone-A	Transport Zone-B	No	No
NSX-T Data Center-A	NSX-T Data Center-B	No	No

vMotion between vSphere Distributed Switch (NSX-T Data Center enabled) and vSphere Standard Switch or vSphere Distributed Switch

Source / VDS	Destination / NSX Virtual Distributed Switch	Compute vMotion	Storage vMotion
vCenter Server-A	vCenter Server-A	Yes	Yes
vCenter Server-A	vCenter Server-B	Yes	Yes
Segment-A (vCenter Server-A)	Segment-B (vCenter Server-A)	No	No
Segment-A (vCenter Server-A)	Segment-B (vCenter Server-B)	No	No
Transport Zone-A	Transport Zone-B	No	No
NSX-T Data Center-A	NSX-T Data Center-B	No	No

Enhanced Networking Stack

Both VDS and NSX Virtual Distributed Switches support all features of the enhanced networking stack.

Scale Supported in vSphere 7.0

Parameter	NSX Virtual Distributed Switch
Logical Switch	■ NSX Distributed Virtual port groups (in vCenter Server) support 10000 X N, where N is the number of VDS switches in vCenter Server. ■ NSX-T Data Center supports 10000 segments.

Relationship between NSX Distributed Virtual port groups and Hostd memory on the host.

NSX Distributed Virtual Port Groups	Minimum Hostd Memory	Supported VMs
5000	600 MB	191
10000	1000 MB	409
15000	1500 MB	682

Enhanced Networking Stack

Enhanced data path is a networking stack mode, which when configured provides superior network performance. It is primarily targeted for NFV workloads, which requires the performance benefits provided by this mode.

The N-VDS switch can be configured in the enhanced data path mode only on an ESXi host. ENS also supports traffic flowing through Edge VMs. In the enhanced data path mode, you can configure overlay traffic and VLAN traffic.

Automatically Assign ENS Logical Cores

Automatically assign logical cores to vNICs such that dedicated logical cores manage the incoming traffic to and outgoing traffic from vNICs.

With the N-VDS switch configured in the enhanced datapath mode, if a single logical core is associated to a vNIC, then that logical core processes bidirectional traffic coming into or going out of a vNIC. When multiple logical cores are configured, the host automatically determines which logical core must process a vNIC's traffic.

Assign logical cores to vNICs based on one of these parameters.

- **vNIC-count:** Host assumes transmission of incoming or outgoing traffic for a vNIC direction requires same amount of the CPU resource. Each logical core is assigned the same number of vNICs based on the available pool of logical cores. It is the default mode. The vNIC-count mode is reliable, but is not optimal for an asymmetric traffic.
- **CPU-usage:** Host predicts the CPU usage to transmit incoming or outgoing traffic at each vNIC direction by using internal statistics. Based on the usage of CPU to transmit traffic, host changes the logical core assignments to balance load among logical cores. The CPU usage mode is more optimal than vNIC-count, but unreliable when traffic is not steady.

In CPU usage mode, if the traffic transmitted changes frequently, then the predicted CPU resources required and vNIC assignment might also change frequently. Too frequent assignment changes might cause packet drops.

If the traffic patterns are symmetric among vNICs, the vNIC-count option provides reliable behavior, which is less vulnerable to frequent changes. However, if the traffic patterns are asymmetric, vNIC-count might result in packet drops since it does not distinguish the traffic difference among vNICs.

In vNIC-count mode, it is recommended to configure an appropriate number of logical cores so that each logical core is assigned to the same number of vNICs. If the number vNIC associated to each logical core is different, CPU assignment is unfair and performance is not deterministic.

When a vNIC is connected or disconnected or when a logical core is added or removed, hosts automatically detect the changes and rebalance.

Procedure

- ◆ To switch from one mode to another mode, run the following command.

```
set ens lcore-assignment-mode <host-switch-name> <ens-lc-mode>
```

Where, *<ens-lc-mode>* can be set to the mode **vNIC-count** or **cpu-usage**.

vNIC-count is vNIC/Direction count-based logical core assignment.

cpu-usage is CPU usage-based logical core assignment.

Configure Guest Inter-VLAN Routing

On overlay networks, NSX-T supports routing of inter-VLAN traffic on an L3 domain. During routing, virtual distributed router (VDR) uses VLAN ID to route packets between VLAN subnets.

Inter-VLAN routing overcomes the limitation of 10 vNICs that can be used per VM. NSX-T supporting inter-VLAN routing ensures that many VLAN subinterfaces can be created on the vNIC and consumed for different networking services. For example, one vNIC of a VM can be divided into several subinterfaces. Each subinterface belongs to a subnet, which can host a networking service such as SNMP or DHCP. With Inter-VLAN routing, for example, a subinterface on VLAN-10 can reach a subinterface on VLAN-10 or any other VLAN.

Each vNIC on a VM is connected to the N-VDS through the parent logical port, which manages untagged packets.

To create a subinterface, on the Enhanced N-VDS switch, create a child port using the API with an associated VIF using the API call described in the procedure. The subinterface tagged with a VLAN ID is associated to a new logical switch, for example, VLAN10 is attached to logical switch LS-VLAN-10. All subinterfaces of VLAN10 have to be attached to LS-VLAN-10. This 1-1 mapping between the VLAN ID of the subinterface and its associated logical switch is an important prerequisite. For example, adding a child port with VLAN20 to logical switch LS-VLAN-10 mapped to VLAN-10 makes routing of packets between VLANs non-functional. Such configuration errors make the inter-VLAN routing non-functional.

Prerequisites

- Before you associate a VLAN subinterface to a logical switch, ensure that the logical switch does not have any other associations with another VLAN subinterface. If there is a mismatch, inter-VLAN routing on overlay networks might not work.
- Ensure that hosts run ESXi v 6.7 U2 or later versions.

Procedure

- 1 To create subinterfaces for a vNIC, ensure that the vNIC is updated to a parent port. Make the following REST API call.

```
PUT https://<nsx-mgr-ip>/api/v1/logical-ports/<Logical-Port UUID-of-the-vNIC>
{
  "resource_type" : "LogicalPort",
  "display_name" : "parentport",
  "attachment" : {
    "attachment_type" : "VIF",
    "context" : {
      "resource_type" : "VifAttachmentContext",
      "vif_type": "PARENT"
    },
    "id" : "<Attachment UUID of the vNIC>"
  },
  "admin_state" : "UP",
  "logical_switch_id" : "UUID of Logical Switch to which the vNIC is connected",
  "_revision" : 0
}
```

- 2 To create child ports for a parent vNIC port on the N-VDS that is associated to the subinterfaces on a VM, make the API call. Before making the API call, verify that a logical switch exists to connect child ports with the subinterfaces on the VM.

```
POST https://<nsx-mgr-ip>/api/v1/logical-ports/
{
  "resource_type" : "LogicalPort",
  "display_name" : "<Name of the Child PORT>",
  "attachment" : {
    "attachment_type" : "VIF",
    "context" : {
      "resource_type" : "VifAttachmentContext",
      "parent_vif_id" : "<UUID of the PARENT port from Step 1>",
      "traffic_tag" : <VLAN ID>,
      "app_id" : "<ID of the attachment>", ==> display id(can give any string). Must be
unique.
      "vif_type" : "CHILD"
    },
    "id" : "<ID of the CHILD port>"
  },

  "logical_switch_id" : "<UUID of the Logical switch(not the PARENT PORT's logical switch)
to which Child port would be connected to>",
  "address_bindings" : [ { "mac_address" : "<vNIC MAC address>", "ip_address" : "<IP
address to the corresponding VLAN>", "vlan" : <VLAN ID> } ],
  "admin_state" : "UP"
}
```

Results

NSX-T Data Center creates subinterfaces on VMs.

Migrate Host Switch to vSphere Distributed Switch

When using N-VDS as the host switch, NSX-T is represented as an opaque network in vCenter Server. N-VDS owns one or more of the physical interfaces (pNICs) on the transport node, and port configuration is performed from NSX-T Data Center. You can migrate your host switch to vSphere Distributed Switch (VDS) 7.0 for optimal pNIC usage, and managing the networking for NSX-T hosts from vCenter Server. When running NSX-T on a VDS switch, a segment is represented as an NSX Distributed Virtual Port Groups. Any changes to the segments on the NSX-T network are synchronized in vCenter Server. Note that in NSX-T Data Center 3.0, N-VDS to VDS migration is not supported for an N-VDS-based Collapsed Cluster environment. Also, migration does not support named teaming policy.

You can scale up migration by parallelly migrating hosts through vSphere Update Manager or manually through APIs. By default, 64 hosts per cluster, with a thread pool size of 22 per manager in a cluster, can be migrated in parallel mode. For migration through vSphere Update Manager, "In Queue" status is shown for any host waiting for an available thread. For migration through API, any request over 64 active migration is rejected. You can set the thread count in the properties file through `PARALLEL_HOST_MIGRATION_COUNT` field. By default, it is set to 64.

Note N-VDS to NSX-T on VDS migration is triggered only for ESX upgrades that cross the 7.0.2 (X.Y.Z-U.P) release. Migration will not be triggered for any "U.P" (update-patch) upgrades. ESX version is specified as X.Y.Z-U.P where,

- X = Major
 - Y = Minor
 - Z = Maintenance
 - U = Update
 - P = Patch
-

Prerequisites

Contact VMware support to assess the impact of migrating to VDS 7.0.

The following requirements must be met to migrate to a VDS 7.0 host switch:

- vCenter Server 7.0 or later
- ESXi 7.0 or later
- NSX-T is no longer represented as an opaque network after migration. You may need to update your scripts to manage the migrated representation of the NSX-T hosts.

Procedure

- 1 Migrate your host switch using API calls or run the commands for migration from the CLI. You can also initiate migration from the NSX Manager user interface.

- Make the following API calls to perform the migration:
 - a To verify that the hosts are ready for migration, make the following API call and run a pre-check:

```
POST https://<nsx-mgr>/api/v1/nvds-urt/precheck
```

Starting with NSX-T 3.1.1, you can also migrate host switches where separate VDSes will be generated for each cluster. For this use case, you must call the following API instead of `POST https://<nsx-mgr>/api/v1/nvds-urt/precheck`.

```
POST https://<nsx-mgr>/api/v1/nvds-urt/precheck-by-cluster
```

Example response for both API calls:

```
{ "precheck_id": "166959af-7f4b-4d49-b294-907000eef889" }
```

- b Address any configuration inconsistencies and run the pre-check again.
- c Verify the status of the pre-check.

```
POST https://<nsx-mgr>/api/v1/nvds-urt/status-summary/<precheck-id>
```

Example response:

```
{
  "precheck_id": "166959af-7f4b-4d49-b294-907000eef889",
  "precheck_status": "PENDING_TOPOLOGY"
}
```

- d For stateless hosts, nominate one of the hosts as the source host and initiate the migration.
- e To retrieve the recommended topology, make the following API call:

```
GET https://<nsx-mgr>/api/v1/nvds-urt/topology/<precheck-id>
```

Example response:

```
{
  "topology": [
    {
      "nvds_id": "21d4fd9b-7214-46b7-ab16-c4e7138f011f",
      "nvds_name": "nsxvswitch",
      "compute_manager_topology": [
        {
          "compute_manager_id": "fa1421d9-54a7-418e-9e18-7d0ff0d2f771",
          "dvswitch": [

```

```

    {
      "data_center_id": "datacenter-3",
      "vds_name": "VDS-nsxvswitch-datacenter-3",
      "vmknics": [
        "vmk1"
      ],
      "transport_node_id": [
        "4a6161af-7eec-4780-8faf-0e0610c33c2e",
        "5a78981a-03a6-40c0-8a77-28522bbf07a9",
        "f9c6314d-9b99-48aa-bfc8-1b3a582162bb"
      ]
    }
  ]
}
]
}

```

- f Make the following API call to create a VDS with the recommended topology:

```
POST https://<nsx-mgr>/api/v1/nvds-urt/topology?action=apply
```

Note that you can only replace an N-VDS with a new VDS and cannot use an existing VDS.

Example input:

```

{
  "topology": [
    {
      "nvds_id": "c8ff4053-502a-4636-8a38-4413c2a2d52f",
      "nvds_name": "nsxvswitch",
      "compute_manager_topology": [
        {
          "compute_manager_id": "fa1421d9-54a7-418e-9e18-7d0ff0d2f771",
          "dvswitch": [
            {
              "data_center_id": "datacenter-3",
              "vds_name": "test-dvs",
              "transport_node_id": [
                "65592db5-adad-47a7-8502-1ab548c63c6d",
                "e57234ee-1d0d-425e-b6dd-7dbc5f6e6527",
                "70f55855-6f81-45a8-bd40-d8b60ae45b82"
              ]
            }
          ]
        }
      ]
    }
  ]
}

```

- g To track the status of the migration, make the following API call:

```
POST https://<nsx-mgr>/api/v1/nvds-urt/status-summary/<precheck-id>
```

When the host is ready for migration, `precheck_status` changes from `APPLYING_TOPOLOGY` to `UPGRADE_READY`.

Refer to the *NSX-T Data Center API Guide* guide for more information on API parameters.

- h Place the ESXi host in maintenance mode from vCenter.
- i To initiate the N-VDS to VDS migration, make the following API call:

```
POST https://<nsx-mgr>/api/v1/transport-nodes/<tn-id>?action=migrate_to_vds
```

The hosts are migrated asynchronously. You can upgrade multiple transport nodes in parallel by calling the API for a desired set of hosts. Services like DRS continue to run as expected during the process of migration.

- j Make the following API call to track the status of migration:

```
POST https://<nsx-mgr>/api/v1/nvds-urt/status-summary/<precheck-id>
```

The host `migration_state` changes from `UPGRADE_IN_PROGRESS` to `SUCCESS` after a successful migration.

Example response:

```
{
  "precheck_id": "c306e279-8b75-4160-919c-6c40030fb3d0",
  "precheck_status": "READY",
  "migration_state": [
    {
      "host": "65592db5-adad-47a7-8502-1ab548c63c6d",
      "overall_state": "UPGRADE_READY"
    },
    {
      "host": "e57234ee-1d0d-425e-b6dd-7dbc5f6e6527",
      "overall_state": "UPGRADE_READY"
    },
    {
      "host": "70f55855-6f81-45a8-bd40-d8b60ae45b82",
      "overall_state": "SUCCESS"
    }
  ]
}
```

In the event of failures, the `overall_state` changes to `FAILED`, indicating the reason for the migration failure. Run the `migrate_to_vds` action to run the migration task again.

- k For stateless hosts:
 - 1 Extract the host profile from the migrated host and attach it to the cluster.
 - 2 Reboot the remaining hosts in the cluster.
- Perform the migration from the NSX Manager CLI.
 - a To verify that the hosts are ready for migration, run the following command and run a pre-check:

```
vds-migrate precheck
```

Sample output:

```
Precheck Id: 0a26d126-7116-11e5-9d70-feff819cdc9f
```

- b Address any configuration inconsistencies and run the pre-check again.
- c To retrieve the recommended topology, run the following command:

```
vds-migrate show-topology
```

Sample output:

```
Precheck Id: 137d2a87-0544-4914-829d-d8b7e33b13f2
NVDS: nvds1(19cca902-9455-4316-92e2-65f4f5b4b138)
Compute Manager Topology:
[
  {
    "compute_manager_id": "fd37ed6e-0eae-4d65-b29a-d40eee1d5d47",
    "dvswitch": [
      {
        "transport_node_id": [
          "4d011ade-a010-4eea-b45a-b2569c0bb9ad"
        ],
        "data_center_id": "datacenter-3",
        "vmknics": [],
        "vds_name": "VDS-nvds1-datacenter-3"
      }
    ]
  }
]
```

- d Run the following command to create a VDS with the recommended topology:
- e Log in to vCenter Server and verify that the VDS is created.
- f To initiate the N-VDS to VDS migration, run the following command:

```
vds-migrate esxi-cluster-name <cluster-name>
```

Sample output:

```
VDS Migration Done:
  3 Transport-Nodes Migrate Successfully
  0 Transport-Nodes Migrate Failed
```

You can also use the transport node ID to initiate migration:

```
vds-migrate tn-list <file-path>
```

where <file-path> includes the Transport Node IDs.

Sample output:

```
nsx-manager-1> vds-migrate tn-list /opt/tnid
VDS Migration Done:
  3 Transport-Nodes Migrate Successfully
  0 Transport-Nodes Migrate Failed
```

- Starting NSX-T Data Center 3.1.1, you can use the NSX Manager to prepare your hosts for migration and then migrate to VDS as part of the host OS upgrade, using vSphere Update Manager.

To migrate your host switch as part of the host OS upgrade using vSphere Update Manager, you need vCenter Server 7.0.U2.

Note The ESXi hosts must be upgraded across ESXi update releases for this migration to be triggered. For example,

- Upgrade from ESXi 7.0 to ESXi 7.0 U2 - switch migration can be triggered.
 - Upgrade from ESXi 7.0 U2 to ESXi 7.0 U2a - switch migration cannot be triggered as upgrade is within the same ESXi update version.
-

- a Log in as a local admin user to an NSX Manager at <https://nsx-manager-ip-address/login.jsp?local=true>.
- b Select **System > Quick Start**.
- c Click **Get Started** to prepare your hosts for migration from N-VDS to VDS.
- d Click **Precheck** to verify if the hosts are ready for migration.
- e Address any configuration inconsistencies and run the pre-check again.
- f Review the recommended network topology.
- g Click **Create** to prepare the selected hosts for migration by creating a corresponding VDS switch in vCenter Server.
- h Log in to vCenter Server and upgrade your ESXi hosts using vSphere Update Manager. The switch migration is completed when the host OS upgrade is done.
- i Monitor the progress of migration from the Monitor tab.

- 2 Move the migrated hosts out of maintenance mode. This step is not required for migration of host switch as part of host upgrade using vSphere Update Manager.

NSX Virtual Distributed Switch

The primary component involved in the data plane of the transport nodes is the NSX Virtual Distributed Switch (N-VDS). On ESXi hypervisors, the N-VDS implementation is derived from VMware vSphere® Distributed Switch™ (VDS). With KVM hypervisors, the N-VDS implementation is derived from the Open vSwitch (OVS).

The N-VDS is required on NSX-T overlay and VLAN-backed networks.

The NVDS forwards traffic between:

- Components running on the transport node (for example, between virtual machines).
- Internal components and the physical network.

If N-VDS is used to forward traffic between internal components and physical network, the NVDS must own one or more physical interfaces (pNICs) on the transport node. As with other virtual switches, an N-VDS cannot share a physical interface with another N-VDS, it may coexist with another N-VDS (or other vSwitch) when using a separate set of pNICs. While N-VDS behavior in realizing connectivity is identical regardless of the specific implementation, data plane realization and enforcement capabilities differ based on compute manager and associated hypervisor capability.

By default, IGMP snooping (IGMPv1/v2/v3, MLDv1/v2) is enabled on N-VDS configured on an ESXi host.

To change settings of IGMP snooping on an N-VDS switch, run the following CLI commands.

```
get host-switch nvds mcast-filter
```

```
set host-switch nvds mcast-filter
  legacy mcast filter mode: {legacy|snooping}
  snooping mcast filter mode: {legacy|snooping}
```

To change settings at a per-port level, run the following CLI commands.

```
get host-switch <host-switch-name> dvport <dvport-id> mcast-filter
```

```
get host-switch <host-switch-name> dvport <dvport-id> mcast-filter <entry-
mode> <entry-group>
```

For more details, see the *NSX-T Data Center Command-Line Interface Reference*.

Virtual Private Network (VPN)

6

NSX-T Data Center supports IPsec Virtual Private Network (IPsec VPN) and Layer 2 VPN (L2 VPN) on an NSX Edge node. IPsec VPN offers site-to-site connectivity between an NSX Edge node and remote sites. With L2 VPN, you can extend your data center by enabling virtual machines to keep their network connectivity across geographical boundaries while using the same IP address.

Note IPsec VPN and L2 VPN are not supported in the NSX-T Data Center limited export release.

You must have a working NSX Edge node, with at least one configured Tier-0 or Tier-1 gateway, before you can configure a VPN service. For more information, see "NSX Edge Installation" in the *NSX-T Data Center Installation Guide*.

Beginning with NSX-T Data Center 2.4, you can also configure new VPN services using the NSX Manager user interface. In earlier releases of NSX-T Data Center, you can only configure VPN services using REST API calls.

Important When using NSX-T Data Center 2.4 or later to configure VPN services, you must use new objects, such as Tier-0 gateways, that were created using the NSX Manager UI or Policy APIs that are included with NSX-T Data Center 2.4 or later release. To use existing Tier-0 or Tier-1 logical routers that were configured before the NSX-T Data Center 2.4 release, you must continue to use API calls to configure a VPN service.

System-default configuration profiles with predefined values and settings are made available for your use during a VPN service configuration. You can also define new profiles with different settings and select them during the VPN service configuration.

The Intel QuickAssist Technology (QAT) feature on a bare metal server is supported for IPsec VPN bulk cryptography. Support for this feature began with NSX-T Data Center 3.0. For more information on support of the QAT feature on bare metal servers, see the *NSX-T Data Center Installation Guide*.

Read the following topics next:

- [Understanding IPsec VPN](#)
- [Understanding Layer 2 VPN](#)
- [Adding VPN Services](#)

- [Adding IPsec VPN Sessions](#)
- [Adding L2 VPN Sessions](#)
- [Add Local Endpoints](#)
- [Adding Profiles](#)
- [Add an Autonomous Edge as an L2 VPN Client](#)
- [Check the Realized State of an IPsec VPN Session](#)
- [Monitor and Troubleshoot VPN Sessions](#)

Understanding IPsec VPN

Internet Protocol Security (IPsec) VPN secures traffic flowing between two networks connected over a public network through IPsec gateways called endpoints. NSX Edge only supports a tunnel mode that uses IP tunneling with Encapsulating Security Payload (ESP). ESP operates directly on top of IP, using IP protocol number 50.

IPsec VPN uses the IKE protocol to negotiate security parameters. The default UDP port is set to 500. If NAT is detected in the gateway, the port is set to UDP 4500.

NSX Edge supports a policy-based or a route-based IPsec VPN.

Beginning with NSX-T Data Center 2.5, IPsec VPN services are supported on both Tier-0 and Tier-1 gateways. See [Add a Tier-0 Gateway](#) or [Add a Tier-1 Gateway](#) for more information. The Tier-0 or Tier-1 gateway must be in `Active-Standby` high-availability mode when used for an IPsec VPN service. You can use segments that are connected to either Tier-0 or Tier-1 gateways when configuring an IPsec VPN service.

An IPsec VPN service in NSX-T Data Center uses the gateway-level failover functionality to support a high-availability service at the VPN service level. Tunnels are re-established on failover and VPN configuration data is synchronized. Before NSX-T Data Center 3.0 release, the IPsec VPN state is not synchronized as tunnels are being re-established. Beginning with NSX-T Data Center 3.0 release, the IPsec VPN state is synchronized to the standby NSX Edge node when the current active NSX Edge node fails and the original standby NSX Edge node becomes the new active NSX Edge node without renegotiating the tunnels. This feature is supported for both policy-based and route-based IPsec VPN services.

Pre-shared key mode authentication and IP unicast traffic are supported between the NSX Edge node and remote VPN sites. In addition, certificate authentication is supported beginning with NSX-T Data Center 2.4. Only certificate types signed by one of the following signature hash algorithms are supported.

- SHA256withRSA
- SHA384withRSA
- SHA512withRSA

Using Policy-Based IPsec VPN

Policy-based IPsec VPN requires a VPN policy to be applied to packets to determine which traffic is to be protected by IPsec before being passed through the VPN tunnel.

This type of VPN is considered static because when a local network topology and configuration change, the VPN policy settings must also be updated to accommodate the changes.

When using a policy-based IPsec VPN with NSX-T Data Center, you use IPsec tunnels to connect one or more local subnets behind the NSX Edge node with the peer subnets on the remote VPN site.

You can deploy an NSX Edge node behind a NAT device. In this deployment, the NAT device translates the VPN address of an NSX Edge node to a publicly accessible address facing the Internet. Remote VPN sites use this public address to access the NSX Edge node.

You can place remote VPN sites behind a NAT device as well. You must provide the remote VPN site's public IP address and its ID (either FQDN or IP address) to set up the IPsec tunnel. On both ends, static one-to-one NAT is required for the VPN address.

Note DNAT is not supported on a tier-1 gateway where policy-based IPsec VPN is configured.

IPsec VPN can provide a secure communications tunnel between an on-premises network and a network in your cloud software-defined data center (SDDC). For policy-based IPsec VPN, the local and peer networks provided in the session must be configured symmetrically at both endpoints. For example, if the cloud-SDDC has the *local* networks configured as **x**, **y**, **z** subnets and the *peer* network is **a**, then the on-premises VPN configuration must have **a** as the *local* network and **x**, **y**, **z** as the *peer* network. This case is true even when **a** is set to **ANY** (**0.0.0.0/0**). For example, if the cloud-SDDC policy-based VPN session has the *local* network configured as **10.1.1.0/24** and the *peer* network as **0.0.0.0/0**, at the on-premises VPN endpoint, the VPN configuration must have **0.0.0.0/0** as the local network and **10.1.1.0/24** as the *peer* network. If misconfigured, the IPsec VPN tunnel negotiation might fail.

The size of the NSX Edge node determines the maximum number of supported tunnels, as shown in the following table.

Table 6-1. Number of IPsec Tunnels Supported

Edge Node Size	# of IPsec Tunnels Per VPN Session (Policy-Based)	# of Sessions Per VPN Service	# of IPsec Tunnels Per VPN Service (16 tunnels per session)
Small	N/A (POC/Lab Only)	N/A (POC/Lab Only)	N/A (POC/Lab Only)
Medium	128	128	2048

Table 6-1. Number of IPSec Tunnels Supported (continued)

Edge Node Size	# of IPSec Tunnels Per VPN Session (Policy-Based)	# of Sessions Per VPN Service	# of IPSec Tunnels Per VPN Service (16 tunnels per session)
Large	128 (soft limit)	256	4096
Bare Metal	128 (soft limit)	512	6000

Restriction The inherent architecture of policy-based IPSec VPN restricts you from setting up a VPN tunnel redundancy.

For information about configuring a policy-based IPSec VPN, see [Add an IPSec VPN Service](#).

Using Route-Based IPSec VPN

Route-based IPSec VPN provides tunneling on traffic based on the static routes or routes learned dynamically over a special interface called virtual tunnel interface (VTI) using, for example, BGP as the protocol. IPSec secures all the traffic flowing through the VTI.

Note

- OSPF dynamic routing is not supported for routing through IPSec VPN tunnels.
- Dynamic routing for VTI is not supported on VPN that is based on Tier-1 gateways.

Route-based IPSec VPN is similar to Generic Routing Encapsulation (GRE) over IPSec, with the exception that no additional encapsulation is added to the packet before applying IPSec processing.

In this VPN tunneling approach, VTIs are created on the NSX Edge node. Each VTI is associated with an IPSec tunnel. The encrypted traffic is routed from one site to another site through the VTI interfaces. IPSec processing happens only at the VTI.

VPN Tunnel Redundancy


You can configure VPN tunnel redundancy with a route-based IPSec VPN session that is configured on a Tier-0 gateway. With tunnel redundancy, multiple tunnels can be set up between two sites, with one tunnel being used as the primary with failover to the other tunnels when the primary tunnel becomes unavailable. This feature is most useful when a site has multiple connectivity options, such as with different ISPs for link redundancy.

Important

- In NSX-T Data Center, IPSec VPN tunnel redundancy is supported using BGP only.
- Do not use static routing for route-based IPSec VPN tunnels to achieve VPN tunnel redundancy.

The following figure shows a logical representation of IPsec VPN tunnel redundancy between two sites. In this figure, Site A and Site B represent two data centers. For this example, assume that NSX-T Data Center is not managing the Edge VPN Gateways in Site A, and that NSX-T Data Center is managing an Edge Gateway virtual appliance in Site B.

Figure 6-1. Tunnel Redundancy in Route-Based IPsec VPN



As shown in the figure, you can configure two independent IPsec VPN tunnels by using VTIs. Dynamic routing is configured using BGP protocol to achieve tunnel redundancy. If both IPsec VPN tunnels are available, they remain in service. All the traffic destined from Site A to Site B through the NSX Edge node is routed through the VTI. The data traffic undergoes IPsec processing and goes out of its associated NSX Edge node uplink interface. All the incoming IPsec traffic received from Site B VPN Gateway on the NSX Edge node uplink interface is forwarded to the VTI after decryption, and then usual routing takes place.

You must configure BGP HoldDown timer and KeepAlive timer values to detect loss of connectivity with peer within the required failover time. See [Configure BGP](#).

Understanding Layer 2 VPN

With Layer 2 VPN (L2 VPN), you can extend Layer 2 networks (VNIs or VLANs) across multiple sites on the same broadcast domain. This connection is secured with a route-based IPsec tunnel between the L2 VPN server and the L2 VPN client.

Note This L2 VPN feature is available only for NSX-T Data Center and does not have any third-party interoperability.

The extended network is a single subnet with a single broadcast domain, which means the VMs remain on the same subnet when they are moved between sites. The VMs' IP addresses do not change when they are moved. So, enterprises can seamlessly migrate VMs between network sites. The VMs can run on either VNI-based networks or VLAN-based networks. For cloud providers, L2 VPN provides a mechanism to onboard tenants without modifying existing IP addresses used by their workloads and applications.

In addition to supporting data center migration, an on-premises network extended with an L2 VPN is useful for a disaster recovery plan and dynamically engaging off-premise compute resources to meet the increased demand.

L2 VPN services are supported on both Tier-0 and Tier-1 gateways. Only one L2 VPN service (either client or server) can be configured for either Tier-0 or Tier-1 gateway.

Each L2 VPN session has one Generic Routing Encapsulation (GRE) tunnel. Tunnel redundancy is not supported. An L2 VPN session can extend up to 4094 L2 segments.

VLAN-based and VNI-based segments can be extended using L2 VPN service on an NSX Edge node that is managed in an NSX-T Data Center environment. You can extend L2 networks from VLAN to VNI, VLAN to VLAN, and VNI to VNI.

Segments can be connected to either Tier-0 or Tier-1 gateways and use L2 VPN services.

Also supported is VLAN trunking using an ESX NSX-managed virtual distributed switch (N-VDS). If there are sufficient compute and I/O resources, an NSX Edge cluster can extend multiple VLAN networks over a single interface using VLAN trunking.

Beginning with NSX-T Data Center 3.0, the L2 VPN path MTU discovery (PMTUD) feature is enabled by default. With the PMTUD enabled, the source host learns the path MTU value for the destination host through the L2 VPN tunnel and limits the length of the outgoing IP packet to the learned value. This feature helps avoid IP fragmentation and reassembly within the tunnel, as a result improving the L2 VPN performance.

The L2 VPN PMTUD feature is not applicable for non-IP, non-unicast, and unicast packets with the DF (Don't Fragment) flag cleared. The global PMTU cache timer expires every 10 minutes. To disable or enable L2 VPN PMTUD feature, see [Enable and Disable L2 VPN Path MTU Discovery](#).

The L2 VPN service support is provided in the following deployment scenarios.

- Between an NSX-T Data Center L2 VPN server and an L2 VPN client hosted on an NSX Edge that is managed in an NSX Data Center for vSphere environment. A managed L2 VPN client supports both VLANs and VNIs.
- Between an NSX-T Data Center L2 VPN server and an L2 VPN client hosted on a standalone or unmanaged NSX Edge. An unmanaged L2 VPN client supports VLANs only.
- Between an NSX-T Data Center L2 VPN server and an L2 VPN client hosted on an autonomous NSX Edge. An autonomous L2 VPN client supports VLANs only.

- Beginning with NSX-T Data Center 2.4 release, L2 VPN service support is available between an NSX-T Data Center L2 VPN server and NSX-T Data Center L2 VPN clients. In this scenario, you can extend the logical L2 segments between two on-premises software-defined data centers (SDDCs).

Enable and Disable L2 VPN Path MTU Discovery

You can enable or disable the L2 VPN path MTU (PMTU) discovery feature using CLI commands. By default L2 VPN PMTU discovery is enabled.

Prerequisites

You must have the user name and password for the admin account to log in to the NSX Edge node.

Procedure

- 1 Log in with admin privileges to the CLI of the NSX Edge node .
- 2 To check the status of the L2 VPN PMTU discovery feature, use the following command.

```
Nsxedge> get dataplane l2vpn-pmtu config
```

If the feature is enabled, you see the following output: `l2vpn_pmtu_enabled : True.`

If the feature is disabled, you see the following output: `l2vpn_pmtu_enabled : False.`

- 3 To disable the L2 VPN PMTU discovery feature, use the following command.

```
nsxedge> set dataplane l2vpn-pmtu disabled
```

Adding VPN Services

You can add either an IPSec VPN (policy-based or route-based) or an L2 VPN using the NSX Manager user interface (UI).

The following sections provide information about the workflows required to set up the VPN service that you need. The topics that follow these sections provide details on how to add either an IPSec VPN or an L2 VPN using the NSX Manager user interface.

Policy-Based IPSec VPN Configuration Workflow

Configuring a policy-based IPSec VPN service workflow requires the following high-level steps.

- 1 Create and enable an IPSec VPN service using an existing Tier-0 or Tier-1 gateway. See [Add an IPSec VPN Service](#).
- 2 Create a DPD (dead peer detection) profile, if you prefer not to use the system default. See [Add DPD Profiles](#).

- 3 To use a non-system default IKE profile, define an IKE (Internet Key Exchange) profile . See [Add IKE Profiles](#).
- 4 Configure an IPSec profile using [Add IPSec Profiles](#).
- 5 Use [Add Local Endpoints](#) to create a VPN server hosted on the NSX Edge.
- 6 Configure a policy-based IPSec VPN session, apply the profiles, and attach the local endpoint to it. See [Add a Policy-Based IPSec Session](#). Specify the local and peer subnets to be used for the tunnel. Traffic from a local subnet destined to the peer subnet is protected using the tunnel defined in the session.

Route-Based IPSec VPN Configuration Workflow

A route-based IPSec VPN configuration workflow requires the following high-level steps.

- 1 Configure and enable an IPSec VPN service using an existing Tier-0 or Tier-1 gateway. See [Add an IPSec VPN Service](#).
- 2 Define an IKE profile if you prefer not to use the default IKE profile. See [Add IKE Profiles](#).
- 3 If you decide not to use the system default IPSec profile, create one using [Add IPSec Profiles](#).
- 4 Create a DPD profile if you want to do not want to use the default DPD profile. See [Add DPD Profiles](#).
- 5 Add a local endpoint using [Add Local Endpoints](#).
- 6 Configure a route-based IPSec VPN session, apply the profiles, and attach the local endpoint to the session. Provide a VTI IP in the configuration and use the same IP to configure routing. The routes can be static or dynamic (using BGP). See [Add a Route-Based IPSec Session](#).

L2 VPN Configuration Workflow

Configuring an L2 VPN requires that you configure an L2 VPN service in Server mode and then another L2 VPN service in Client mode. You also must configure the sessions for the L2 VPN server and L2 VPN client using the peer code generated by the L2 VPN Server. Following is a high-level workflow for configuring an L2 VPN service.

- 1 Create an L2 VPN Service in Server mode.
 - a Configure a route-based IPSec VPN tunnel with a Tier-0 or Tier-1 gateway and an L2 VPN Server service using that route-based IPSec tunnel. See [Add an L2 VPN Server Service](#).
 - b Configure an L2 VPN server session, which binds the newly created route-based IPSec VPN service and the L2 VPN server service, and automatically allocates the GRE IP addresses. See [Add an L2 VPN Server Session](#).
 - c Add segments to the L2 VPN Server sessions. This step is also described in [Add an L2 VPN Server Session](#).
 - d Use [Download the Remote Side L2 VPN Configuration File](#) to obtain the peer code for the L2 VPN Server service session, which must be applied on the remote site and used to configure the L2 VPN Client session automatically.

- 2 Create an L2 VPN Service in Client mode.
 - a Configure another route-based IPSec VPN service using a different Tier-0 or Tier-1 gateway and configure an L2 VPN Client service using that Tier-0 or Tier-1 gateway that you just configured. See [Add an L2 VPN Client Service](#) for information.
 - b Define the L2 VPN Client sessions by importing the peer code generated by the L2 VPN Server service. See [Add an L2 VPN Client Session](#).
 - c Add segments to the L2 VPN Client sessions defined in the previous step. This step is described in [Add an L2 VPN Client Session](#).

Add an IPSec VPN Service

NSX-T Data Center supports a site-to-site IPSec VPN service between a Tier-0 or Tier-1 gateway and remote sites. You can create a policy-based or a route-based IPSec VPN service. You must create the IPSec VPN service first before you can configure either a policy-based or a route-based IPSec VPN session.

Note IPSec VPN is not supported in the NSX-T Data Center limited export release.

IPSec VPN is not supported when the local endpoint IP address goes through NAT in the same logical router that the IPSec VPN session is configured.

Prerequisites

- Familiarize yourself with the IPSec VPN. See [Understanding IPSec VPN](#).
- You must have at least one Tier-0 or Tier-1 gateway configured and available for use. See [Add a Tier-0 Gateway](#) or [Add a Tier-1 Gateway](#) for more information.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Navigate to **Networking > VPN > VPN Services**.
- 3 Select **Add Service > IPSec**.
- 4 Enter a name for the IPSec service.
This name is required.
- 5 From the **Tier-0/Tier-1 Gateway** drop-down menu, select the Tier-0 or Tier-1 gateway to associate with this IPSec VPN service.
- 6 Enable or disable **Admin Status**.
By default, the value is set to `Enabled`, which means the IPSec VPN service is enabled on the Tier-0 or Tier-1 gateway after the new IPSec VPN service is configured.
- 7 Set the value for **IKE Log Level**.
The default is set to the `Info` level.

- 8 Enter a value for **Tags** if you want to include this service in a tag group.
- 9 To enable or disable the stateful synchronization of VPN sessions, toggle **Session sync**.
By default, the value is set to `Enabled`.
- 10 Click **Global Bypass Rules** if you want to allow data packets to be exchanged between the specified local and remote IP addresses without any IPSec protection. In the **Local Networks** and **Remote Networks** text boxes, enter the list of local and remote subnets between which the bypass rules are applied.

If you enable these rules, data packets are exchanged between the specified local and remote IP sites even if their IP addresses are specified in the IPSec session rules. The default is to use the IPSec protection when data is exchanged between local and remote sites. These rules apply for all IPSec VPN sessions created within this IPSec VPN service.

- 11 Click **Save**.

After the new IPSec VPN service is created successfully, you are asked whether you want to continue with the rest of the IPSec VPN configuration. If you click **Yes**, you are taken back to the Add IPSec VPN Service panel. The **Sessions** link is now enabled and you can click it to add an IPSec VPN session.

What to do next

Use information in [Adding IPSec VPN Sessions](#) to guide you in adding an IPSec VPN session. You also provide information for the profiles and local endpoint that are required to finish the IPSec VPN configuration.

Add an L2 VPN Service

You configure an L2 VPN service on a Tier-0 or Tier-1 gateway. To enable the L2 VPN service, you must first create an IPSec VPN service on the Tier-0 or Tier-1 gateway, if it does not exist yet. You then configure an L2 VPN tunnel between an L2 VPN server (destination gateway) and an L2 VPN client (source gateway).

To configure an L2 VPN service, use the information in the topics that follow in this section.

Prerequisites

- Familiarize yourself with IPsec VPN and L2 VPN. See [Understanding IPSec VPN](#) and [Understanding Layer 2 VPN](#).
- You must have at least one Tier-0 or Tier-1 gateway configured and available for use. See [Add a Tier-0 Gateway](#) or [Add a Tier-1 Gateway](#).

Procedure

1 [Add an L2 VPN Server Service](#)

To configure an L2 VPN Server service, you must configure the L2 VPN service in server mode on the destination NSX Edge to which the L2 VPN client is to be connected.

2 Add an L2 VPN Client Service

After configuring the L2 VPN Server service, configure the L2 VPN service in the client mode on another NSX Edge instance.

Add an L2 VPN Server Service

To configure an L2 VPN Server service, you must configure the L2 VPN service in server mode on the destination NSX Edge to which the L2 VPN client is to be connected.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 (Optional) If an IPsec VPN service does not exist yet on either a Tier-0 or Tier-1 gateway that you want to configure as the L2 VPN server, create the service using the following steps.
 - a Navigate to the **Networking > VPN > VPN Services** tab and select **Add Service > IPsec**.
 - b Enter a name for the IPsec VPN service.
 - c From the **Tier-0/Tier-1 Gateway** drop-down menu, select the gateway to use with the L2 VPN server.
 - d If you want to use values different from the system defaults, set the rest of the properties on the Add IPsec Service pane, as needed.
 - e Click **Save** and when prompted if you want to continue configuring the IPsec VPN service, select **No**.
- 3 Navigate to the **Networking > VPN > VPN Services** tab and select **Add Service > L2 VPN Server** to create an L2 VPN server.
- 4 Enter a name for the L2 VPN server.
- 5 From the **Tier-0/Tier-1 Gateway** drop-down menu, select the same Tier-0 or Tier-1 gateway that you used with the IPsec service you created a moment ago.
- 6 Enter an optional description for this L2 VPN server.
- 7 Enter a value for **Tags** if you want to include this service in a tag group.
- 8 Enable or disable the **Hub & Spoke** property.

By default, the value is set to `Disabled`, which means the traffic received from the L2 VPN clients is only replicated to the segments connected to the L2 VPN server. If this property is set to `Enabled`, the traffic from any L2 VPN client is replicated to all other L2 VPN clients.

- 9 Click **Save**.

After the new L2 VPN server is created successfully, you are asked whether you want to continue with the rest of the L2 VPN service configuration. If you click **Yes**, you are taken back to the Add L2 VPN Server pane and the **Session** link is enabled. You can use that link to create an L2 VPN server session or use the **Networking > VPN > L2 VPN Sessions** tab.

What to do next

Configure an L2 VPN server session for the L2 VPN server that you configured using information in [Add an L2 VPN Server Session](#) as a guide.

Add an L2 VPN Client Service

After configuring the L2 VPN Server service, configure the L2 VPN service in the client mode on another NSX Edge instance.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 (Optional) If an IPsec VPN service does not exist yet on either a Tier-0 or Tier-1 gateway that you want to configure as the L2 VPN client, create the service using the following steps.
 - a Navigate to the **Networking > VPN > VPN Services** tab and select **Add Service > IPsec**.
 - b Enter a name for the IPsec VPN service.
 - c From the **Tier-0/Tier-1 Gateway** drop-down menu, select a Tier-0 or Tier-1 gateway to use with the L2 VPN client.
 - d If you want to use values different from the system defaults, set the rest of the properties on the Add IPsec Service pane, as needed.
 - e Click **Save** and when prompted if you want to continue configuring the IPsec VPN service, select **No**.
- 3 Navigate to the **Networking > VPN > VPN Services** tab and select **Add Service > L2 VPN Client**.
- 4 Enter a name for the L2 VPN Client service.
- 5 From the **Tier-0/Tier-1 Gateway** drop-down menu, select the same Tier-0 or Tier-1 gateway that you used with the route-based IPsec tunnel you created a moment ago.
- 6 Optionally set the values for **Description** and **Tags**.
- 7 Click **Save**.

After the new L2 VPN client service is created successfully, you are asked whether you want to continue with the rest of the L2 VPN client configuration. If you click **Yes**, you are taken back to the Add L2 VPN Client pane and the **Session** link is enabled. You can use that link to create an L2 VPN client session or use the **Networking > VPN > L2 VPN Sessions** tab.

What to do next

Configure an L2 VPN client session for the L2 VPN Client service that you configured. Use the information in [Add an L2 VPN Client Session](#) as a guide.

Adding IPsec VPN Sessions

After you have configured an IPsec VPN service, you must add either a policy-based IPsec VPN session or a route-based IPsec VPN session, depending on the type of IPsec VPN you want to configure. You also provide the information for the local endpoint and profiles to use to finish the IPsec VPN service configuration.

Add a Policy-Based IPsec Session

When you add a policy-based IPsec VPN, IPsec tunnels are used to connect multiple local subnets that are behind the NSX Edge node with peer subnets on the remote VPN site.

The following steps use the **IPsec Sessions** tab on the NSX Manager UI to create a policy-based IPsec session. You also add information for the tunnel, IKE, and DPD profiles, and select an existing local endpoint to use with the policy-based IPsec VPN.

Note You can also add the IPsec VPN sessions immediately after you have successfully configured the IPsec VPN service. You click **Yes** when prompted to continue with the IPsec VPN service configuration and select **Sessions > Add Sessions** on the Add IPsec Service panel. The first few steps in the following procedure assume you selected **No** to the prompt to continue with the IPsec VPN service configuration. If you selected **Yes**, proceed to step 3 in the following steps to guide you with the rest of the policy-based IPsec VPN session configuration.

Prerequisites

- You must have configured an IPsec VPN service before proceeding. See [Add an IPsec VPN Service](#).
- Obtain the information for the local endpoint, IP address for the peer site, local network subnet, and remote network subnet to use with the policy-based IPsec VPN session you are adding. To create a local endpoint, see [Add Local Endpoints](#).
- If you are using a Pre-Shared Key (PSK) for authentication, obtain the PSK value.
- If you are using a certificate for authentication, ensure that the necessary server certificates and corresponding CA-signed certificates are already imported. See [Chapter 19 Certificates](#).
- If you do not want to use the defaults for the IPsec tunnel, IKE, or dead peer detection (DPD) profiles provided by NSX-T Data Center, configure the profiles you want to use instead. See [Adding Profiles](#) for information.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Navigate to the **Networking > VPN > IPsec Sessions** tab.
- 3 Select **Add IPsec Session > Policy Based**.
- 4 Enter a name for the policy-based IPsec VPN session.

- 5 From the **VPN Service** drop-down menu, select the IPsec VPN service to which you want to add this new IPsec session.

Note If you are adding this IPsec session from the **Add IPsec Sessions** dialog box, the VPN Service name is already indicated above the **Add IPsec Session** button.

- 6 Select an existing local endpoint from the drop-down menu.
This local endpoint value is required and identifies the local NSX Edge node. If you want to create a different local endpoint, click the three-dot menu () and select **Add Local Endpoint**.
- 7 In the **Remote IP** text box, enter the required IP address of the remote site.
This value is required.
- 8 Enter an optional description for this policy-based IPsec VPN session.
The maximum length is 1024 characters.
- 9 To enable or disable the IPsec VPN session, click **Admin Status** .
By default, the value is set to `Enabled`, which means the IPsec VPN session is to be configured down to the NSX Edge node.
- 10 (Optional) From the **Compliance suite** drop-down menu, select a security compliance suite.

Note Compliance suite support is provided beginning with NSX-T Data Center 2.5. See [About Supported Compliance Suites](#) for more information.

The default value selected is `None`. If you select a compliance suite, the **Authentication Mode** is set to `Certificate` and in the **Advanced Properties** section, the values for **IKE profile** and **IPsec profile** are set to the system-defined profiles for the selected security compliance suite. You cannot edit these system-defined profiles.

- 11 If the **Compliance Suite** is set to `None`, select a mode from the **Authentication Mode** drop-down menu.
The default authentication mode used is `PSK`, which means a secret key shared between NSX Edge and the remote site is used for the IPsec VPN session. If you select `Certificate`, the site certificate that was used to configure the local endpoint is used for authentication.
- 12 In the Local Networks and Remote Networks text boxes, enter at least one IP subnet address to use for this policy-based IPsec VPN session.
These subnets must be in a CIDR format.
- 13 If **Authentication Mode** is set to `PSK`, enter the key value in the **Pre-shared Key** text box.
This secret key can be a string with a maximum length of 128 characters.

Caution Be careful when sharing and storing a PSK value because it contains some sensitive information.

14 To identify the peer site, enter a value in **Remote ID**.

For peer sites using PSK authentication, this ID value must be the public IP address or the FQDN of the peer site. For peer sites using certificate authentication, this ID value must be the common name (CN) or distinguished name (DN) used in the peer site's certificate.

Note If the peer site's certificate contains an email address in the DN string, for example,

```
C=US, ST=California, O=MyCompany, OU=MyOrg, CN=Site123/emailAddress=user1@mycompany.com
```

then enter the **Remote ID** value using the following format as an example.

```
C=US, ST=California, O=MyCompany, OU=MyOrg, CN=Site123, MAILTO=user1@mycompany.com"
```

If the local site's certificate contains an email address in the DN string and the peer site uses the strongSwan IPsec implementation, enter the local site's ID value in that peer site. The following is an example.

```
C=US, ST=California, O=MyCompany, OU=MyOrg, CN=Site123, E=user1@mycompany.com"
```

15 To change the profiles, initiation mode, TCP MSS clamping mode, and tags used by the policy-based IPsec VPN session, click **Advanced Properties**.

By default, the system generated profiles are used. Select another available profile if you do not want to use the default. If you want to use a profile that is not configured yet, click the three-dot menu (⋮) to create another profile. See [Adding Profiles](#).

- a If the **IKE Profiles** drop-down menu is enabled, select the IKE profile.
- b Select the IPsec tunnel profile, if the **IPSec Profiles** drop-down menu is not disabled.
- c Select the preferred DPD profile if the **DPD Profiles** drop-down menu is enabled.

- d Select the preferred mode from the **Connection Initiation Mode** drop-down menu.

Connection initiation mode defines the policy used by the local endpoint in the process of tunnel creation. The default value is **Initiator**. The following table describes the different connection initiation modes available.

Table 6-2. Connection Initiation Modes

Connection Initiation Mode	Description
Initiator	The default value. In this mode, the local endpoint initiates the IPSec VPN tunnel creation and responds to incoming tunnel setup requests from the peer gateway.
On Demand	In this mode, the local endpoint initiates the IPSec VPN tunnel creation after the first packet matching the policy rule is received. It also responds to the incoming initiation request.
Respond Only	The IPSec VPN never initiates a connection. The peer site always initiates the connection request and the local endpoint responds to that connection request.

- e If you want 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](#) for more information.

- f If you want to include this session as part of a specific group, enter the tag name in **Tags**.

16 Click **Save**.

Results

When the new policy-based IPSec VPN session is configured successfully, it is added to the list of available IPSec VPN sessions. It is in read-only mode.

What to do next

- Verify that the IPSec VPN tunnel status is Up. See [Monitor and Troubleshoot VPN Sessions](#) for information.
- If necessary, manage the IPSec VPN session information by clicking the three-dot menu (⋮) on the left-side of the session's row. Select one of the actions you are allowed to perform.

Add a Route-Based IPSec Session

When you add a route-based IPSec VPN, tunneling is provided on traffic that is based on routes that were learned dynamically over a virtual tunnel interface (VTI) using a preferred protocol, such as BGP. IPSec secures all the traffic flowing through the VTI.

The steps described in this topic use the **IPSec Sessions** tab to create a route-based IPSec session. You also add information for the tunnel, IKE, and DPD profiles, and select an existing local endpoint to use with the route-based IPSec VPN.

Note You can also add the IPSec VPN sessions immediately after you have successfully configured the IPSec VPN service. You click **Yes** when prompted to continue with the IPSec VPN service configuration and select **Sessions > Add Sessions** on the Add IPsec Service panel. The first few steps in the following procedure assume you selected **No** to the prompt to continue with the IPSec VPN service configuration. If you selected **Yes**, proceed to step 3 in the following steps to guide you with the rest of the route-based IPSec VPN session configuration.

Prerequisites

- You must have configured an IPSec VPN service before proceeding. See [Add an IPSec VPN Service](#).
- Obtain the information for the local endpoint, IP address for the peer site, and tunnel service IP subnet address to use with the route-based IPSec session you are adding. To create a local endpoint, see [Add Local Endpoints](#).
- If you are using a Pre-Shared Key (PSK) for authentication, obtain the PSK value.
- If you are using a certificate for authentication, ensure that the necessary server certificates and corresponding CA-signed certificates are already imported. See [Chapter 19 Certificates](#).
- If you do not want to use the default values for the IPSec tunnel, IKE, or dead peer detection (DPD) profiles provided by NSX-T Data Center, configure the profiles you want to use instead. See [Adding Profiles](#) for information.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Navigate to **Networking > VPN > IPSec Sessions**.
- 3 Select **Add IPSec Session > Route Based**.
- 4 Enter a name for the route-based IPSec session.
- 5 From the **VPN Service** drop-down menu, select the IPSec VPN service to which you want to add this new IPSec session.

Note If you are adding this IPSec session from the **Add IPSec Sessions** dialog box, the VPN Service name is already indicated above the **Add IPSec Session** button.

- 6 Select an existing local endpoint from the drop-down menu.

This local endpoint value is required and identifies the local NSX Edge node. If you want to create a different local endpoint, click the three-dot menu () and select **Add Local Endpoint**.

- 7 In the **Remote IP** text box, enter the IP address of the remote site.

This value is required.

- 8 Enter an optional description for this route-based IPSec VPN session.

The maximum length is 1024 characters.

- 9 To enable or disable the IPSec session, click **Admin Status**.

By default, the value is set to `Enabled`, which means the IPSec session is to be configured down to the NSX Edge node.

- 10 (Optional) From the **Compliance suite** drop-down menu, select a security compliance suite.

Note Compliance suite support is provided beginning with NSX-T Data Center 2.5. See [About Supported Compliance Suites](#) for more information.

The default value is set to `None`. If you select a compliance suite, the **Authentication Mode** is set to `Certificate` and in the **Advanced Properties** section, the values for **IKE profile** and **IPSec profile** are set to the system-defined profiles for the selected compliance suite. You cannot edit these system-defined profiles.

- 11 Enter an IP subnet address in **Tunnel Interface** in the CIDR notation.

This address is required.

- 12 If the **Compliance Suite** is set to `None`, select a mode from the **Authentication Mode** drop-down menu.

The default authentication mode used is `PSK`, which means a secret key shared between NSX Edge and the remote site is used for the IPSec VPN session. If you select `Certificate`, the site certificate that was used to configure the local endpoint is used for authentication.

- 13 If you selected `PSK` for the authentication mode, enter the key value in the **Pre-shared Key** text box.

This secret key can be a string with a maximum length of 128 characters.

Caution Be careful when sharing and storing a PSK value because it contains some sensitive information.

14 Enter a value in **Remote ID**.

For peer sites using PSK authentication, this ID value must be the public IP address or the FQDN of the peer site. For peer sites using certificate authentication, this ID value must be the common name (CN) or distinguished name (DN) used in the peer site's certificate.

Note If the peer site's certificate contains an email address in the DN string, for example,

```
C=US, ST=California, O=MyCompany, OU=MyOrg, CN=Site123/emailAddress=user1@mycompany.com
```

then enter the **Remote ID** value using the following format as an example.

```
C=US, ST=California, O=MyCompany, OU=MyOrg, CN=Site123, MAILTO=user1@mycompany.com"
```

If the local site's certificate contains an email address in the DN string and the peer site uses the strongSwan IPsec implementation, enter the local site's ID value in that peer site. The following is an example.

```
C=US, ST=California, O=MyCompany, OU=MyOrg, CN=Site123, E=user1@mycompany.com"
```

15 If you want to include this IPsec session as part of a specific group tag, enter the tag name in **Tags**.**16** To change the profiles, initiation mode, TCP MSS clamping mode, and tags used by the route-based IPsec VPN session, click **Advanced Properties**.

By default, the system-generated profiles are used. Select another available profile if you do not want to use the default. If you want to use a profile that is not configured yet, click the three-dot menu (⋮) to create another profile. See [Adding Profiles](#).

- a If the **IKE Profiles** drop-down menu is enabled, select the IKE profile.
- b Select the IPsec tunnel profile, if the **IPsec Profiles** drop-down menu is not disabled.

- c Select the preferred DPD profile if the **DPD Profiles** drop-down menu is enabled.
- d Select the preferred mode from the **Connection Initiation Mode** drop-down menu.

Connection initiation mode defines the policy used by the local endpoint in the process of tunnel creation. The default value is **Initiator**. The following table describes the different connection initiation modes available.

Table 6-3. Connection Initiation Modes

Connection Initiation Mode	Description
Initiator	The default value. In this mode, the local endpoint initiates the IPsec VPN tunnel creation and responds to incoming tunnel setup requests from the peer gateway.
On Demand	Do not use with the route-based VPN. This mode applies to policy-based VPN only.
Respond Only	The IPsec VPN never initiates a connection. The peer site always initiates the connection request and the local endpoint responds to that connection request.

- 17 If you want 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](#) for more information.

- 18 If you want to include this IPsec session as part of a specific group tag, enter the tag name in **Tags**.

- 19 Click **Save**.

Results

When the new route-based IPsec VPN session is configured successfully, it is added to the list of available IPsec VPN sessions. It is in read-only mode.

What to do next

- Verify that the IPsec VPN tunnel status is Up. See [Monitor and Troubleshoot VPN Sessions](#) for information.
- Configure routing using either a static route or BGP. See [Configure a Static Route](#) or [Configure BGP](#).
- If necessary, manage the IPsec VPN session information by clicking the three-dot menu (⋮) on the left-side of the session's row. Select one of the actions you can perform.

About Supported Compliance Suites

Beginning with NSX-T Data Center 2.5, you can specify a security compliance suite to use to configure the security profiles used for an IPsec VPN session.

A security compliance suite has predefined values that are used for different security parameters and that cannot be modified. When you select a compliance suite, the predefined values are automatically used for the security profile of the IPSec VPN session you are configuring.

The following table lists the compliance suites that are supported for IKE profiles in NSX-T Data Center and the values that are predefined for each.

Compliance Suite Name	IKE Version	Encryption Algorithm	Digest Algorithm	Diffie Hellman Group
CNSA	IKE V2	AES 256	SHA2 384	Group 15, Group 20
FIPS	IKE FLEX	AES 128	SHA2 256	Group 20
Foundation	IKE V1	AES 128	SHA2 256	Group 14
PRIME	IKE V2	AES GCM 128	Not Set	Group 19
Suite-B-GCM-128	IKE V2	AES 128	SHA2 256	Group 19
Suite-B-GCM-256	IKE V2	AES 256	SHA2 384	Group 20

Note The AES 128 and AES 256 algorithms use the CBC mode of operation.

The following table lists the compliance suites that are supported for IPSec profiles in NSX-T Data Center and the values that are predefined for each.

Compliance Suite Name	Encryption Algorithm	Digest Algorithm	PFS Group	Diffie-Hellman Group
CNSA	AES 256	SHA2 384	Enabled	Group 15, Group 20
FIPS	AES GCM 128	Not Set	Enabled	Group 20
Foundation	AES 128	SHA2 256	Enabled	Group 14
PRIME	AES GCM 128	Not Set	Enabled	Group 19
Suite-B-GCM-128	AES GCM 128	Not Set	Enabled	Group 19
Suite-B-GCM-256	AES GCM 256	Not Set	Enabled	Group 20

Note The AES 128 and AES 256 algorithms use the CBC mode of operation.

Understanding TCP MSS Clamping

TCP MSS clamping enables you to reduce the maximum segment size (MSS) value used by a TCP session during connection establishment through an IPSec tunnel. This feature is supported starting with NSX-T Data Center 2.5.

TCP MSS is the maximum amount of data in bytes that a host is willing to accept in a single TCP segment. Each end of a TCP connection sends its desired MSS value to its peer-end during a three-way handshake, where MSS is one of the TCP header options used in a TCP SYN packet. TCP MSS is calculated based on the maximum transmission unit (MTU) of the egress interface of the sender host.

When a TCP traffic goes through an IPsec VPN or any kind of VPN tunnel, additional headers are added to the original packet to keep it secure. For IPsec tunnel mode, additional headers used are IP, ESP, and optionally UDP (if port translation is present in the network). Because of these additional headers, the size of the encapsulated packet goes beyond the MTU of the VPN interface. The packet can get fragmented or dropped based on the DF policy.

To avoid packet fragmentation or drop, you can adjust the MSS value for the IPsec session by enabling the TCP MSS clamping feature. Navigate to **Networking > VPN > IPsec Sessions**. When you are adding an IPsec session or editing an existing one, expand the **Advance Properties** section, and enable **TCP MSS Clamping**.

You can configure the pre-calculated MSS value suitable for the IPsec session by setting both **TCP MSS Direction** and **TCP MSS Value**. The configured MSS value is used for MSS clamping. You can opt to use the dynamic MSS calculation by setting the **TCP MSS Direction** and leaving **TCP MSS Value** blank. The MSS value is auto-calculated based on the VPN interface MTU, VPN overhead, and the path MTU (PMTU) when it is already determined. The effective MSS is recalculated during each TCP handshake to handle the MTU or PMTU changes dynamically.

Adding L2 VPN Sessions

After you have configured an L2 VPN server and an L2 VPN client, you must add L2 VPN sessions for both to complete the L2 VPN service configuration.

Add an L2 VPN Server Session

After creating an L2 VPN Server service, you must add an L2 VPN session and attach it to an existing segment.

The following steps use the **L2 VPN Sessions** tab on the NSX Manager UI to create an L2 VPN Server session. You also select an existing local endpoint and segment to attach to the L2 VPN Server session.

Note You can also add an L2 VPN Server session immediately after you have successfully configured the L2 VPN Server service. You click **Yes** when prompted to continue with the L2 VPN Server configuration and select **Sessions > Add Sessions** on the Add L2 VPN Server panel. The first few steps in the following procedure assume you selected **No** to the prompt to continue with the L2 VPN Server configuration. If you selected **Yes**, proceed to step 3 in the following steps to guide you with the rest of the L2 VPN Server session configuration.

Prerequisites

- You must have configured an L2 VPN Server service before proceeding. See [Add an L2 VPN Server Service](#).
- Obtain the information for the local endpoint and remote IP to use with the L2 VPN Server session you are adding. To create a local endpoint, see [Add Local Endpoints](#).
- Obtain the values for the pre-shared key (PSK) and the tunnel interface subnet to use with the L2 VPN Server session.
- Obtain the name of the existing segment you want to attach to the L2 VPN Server session you are creating. See [Add a Segment](#) for information.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Navigate to the **Networking > VPN > L2 VPN Sessions** tab.
- 3 Select **Add L2 VPN Session > L2 VPN Server**.
- 4 Enter a name for the L2 VPN Server session.
- 5 From the **L2 VPN Service** drop-down menu, select the L2 VPN Server service for which the L2 VPN session is being created.

Note If you are adding this L2 VPN Server session from the Set L2VPN Server Sessions dialog box, the L2 VPN Server service is already indicated above the **Add L2 Session** button.

- 6 Select an existing local endpoint from the drop-down menu.

If you want to create a different local endpoint, click the three-dot menu () and select **Add Local Endpoint**.

- 7 Enter the IP address of the remote site.
- 8 To enable or disable the L2 VPN Server session, click **Admin Status**.

By default, the value is set to **Enabled**, which means the L2 VPN Server session is to be configured down to the NSX Edge node.

- 9 Enter the secret key value in **Pre-shared Key**.

Caution Be careful when sharing and storing a PSK value because it is considered sensitive information.

- 10 Enter an IP subnet address in the **Tunnel Interface** using the CIDR notation.
For example, 4.5.6.6/24. This subnet address is required.

11 Enter a value in **Remote ID**.

For peer sites using certificate authentication, this ID must be the common name in the peer site's certificate. For PSK peers, this ID can be any string. Preferably, use the public IP address of the VPN or an FQDN for the VPN services as the `Remote ID`.

12 If you want to include this session as part of a specific group, enter the tag name in **Tags**.**13** Click **Save** and click **Yes** when prompted if you want to continue with the VPN service configuration.

You are returned to the Add L2VPN Sessions panel and the **Segments** link is now enabled.

14 Attach an existing segment to the L2 VPN Server session.

- a Click **Segments > Set Segments**.
- b In the **Set Segments** dialog box, click **Set Segment** to attach an existing segment to the L2 VPN Server session.
- c From the **Segment** drop-down menu, select the VNI-based or VLAN-based segment that you want to attach to the session.
- d Enter a unique value in the **VPN Tunnel ID** that is used to identify the segment that you selected.
- e In the **Local Egress Gateway IP** text box, enter the IP address of the local gateway that your workload VMs on the segment use as their default gateway. The same IP address can be configured in the remote site on the extended segment.
- f Click **Save** and then **Close**.

In the Set L2VPN Sessions pane or dialog box, the system has incremented the **Segments** count for the L2 VPN Server session.

15 To finish the L2 VPN Server session configuration, click **Close Editing**.**Results**

In the **VPN Services** tab, the system incremented the **Sessions** count for the L2 VPN Server service that you configured.

What to do next

To complete the L2 VPN service configuration, you must also create an L2 VPN service in Client mode and an L2 VPN client session. See [Add an L2 VPN Client Service](#) and [Add an L2 VPN Client Session](#).

Add an L2 VPN Client Session

You must add an L2 VPN Client session after creating an L2 VPN Client service, and attach it to an existing segment.

The following steps use the **L2 VPN Sessions** tab on the NSX Manager UI to create an L2 VPN Client session. You also select an existing local endpoint and segment to attach to the L2 VPN Client session.

Note You can also add an L2 VPN Client session immediately after you have successfully configured the L2 VPN Client service. Click **Yes** when prompted to continue with the L2 VPN Client configuration and select **Sessions > Add Sessions** on the Add L2 VPN Client panel. The first few steps in the following procedure assume you selected **No** to the prompt to continue with the L2 VPN Client configuration. If you selected **Yes**, proceed to step 3 in the following steps to guide you with the rest of the L2 VPN Client session configuration.

Prerequisites

- You must have configured an L2 VPN Client service before proceeding. See [Add an L2 VPN Client Service](#).
- Obtain the IP addresses information for the local IP and remote IP to use with the L2 VPN Client session you are adding.
- Obtain the peer code that was generated during the L2 VPN server configuration. See [Download the Remote Side L2 VPN Configuration File](#).
- Obtain the name of the existing segment you want to attach to the L2 VPN Client session you are creating. See [Add a Segment](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select the **Networking > VPN > L2 VPN Sessions**.
- 3 Select **Add L2 VPN Session > L2 VPN Client**.
- 4 Enter a name for the L2 VPN Client session.
- 5 From the **VPN Service** drop-down menu, select the L2 VPN Client service with which the L2 VPN session is to be associated.

Note If you are adding this L2 VPN Client session from the Set L2VPN Client Sessions dialog box, the L2 VPN Client service is already indicated above the **Add L2 Session** button.

- 6 In the **Local IP address** text box, enter the IP address of the L2 VPN Client session.
- 7 Enter the remote IP address of the IPsec tunnel to be used for the L2 VPN Client session.
- 8 In the **Peer Configuration** text box, enter the peer code generated when you configured the L2 VPN Server service.
- 9 Enable or disable **Admin Status**.

By default, the value is set to **Enabled**, which means the L2 VPN Server session is to be configured down to the NSX Edge node.

- 10 Click **Save** and click **Yes** when prompted if you want to continue with the VPN service configuration.
- 11 Attach an existing segment to the L2 VPN Client session.
 - a Select **Segments > Add Segments**.
 - b In the **Set Segments** dialog box, click **Add Segment**.
 - c From the **Segment** drop-down menu, select the VNI-based or VLAN-based segment you want to attach to the L2 VPN Client session.
 - d Enter a unique value in the **VPN Tunnel ID** that is used to identify the segment that you selected.
 - e Click **Close**.
- 12 To finish the L2 VPN Client session configuration, click **Close Editing**.

Results

In the **VPN Services** tab, the sessions count is updated for the L2 VPN Client service that you configured.

Download the Remote Side L2 VPN Configuration File

To configure the L2 VPN client session, you must obtain the peer code that was generated when you configured the L2 VPN server session.

Prerequisites

- You must have configured an L2 VPN server service and a session successfully before proceeding. See [Add an L2 VPN Server Service](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Navigate to the **Networking > VPN > L2 VPN Sessions** tab.
- 3 In the table of L2 VPN sessions, expand the row for the L2 VPN server session you plan to use for the L2 VPN client session configuration.
- 4 Click **Download Config** and click **Yes** on the Warning dialog box.

A text file with the name `L2VPNSession_<name-of-L2-VPN-server-session>_config.txt` is downloaded. It contains the peer code for the remote side L2 VPN configuration.

Caution Be careful when storing and sharing the peer code because it contains a PSK value, which is considered sensitive information.

For example, `L2VPNSession_L2VPNServer_config.txt` contains the following configuration.

```
[
  {
    "transport_tunnel_path": "/infra/tier-0s/ServerT0_AS/locale-services/1-
policyconnectivity-693/ipsec-vpn-services/IpsecService1/sessions/Routebase1",
    "peer_code":
    "MCw3ZjBjYzdjLHsic2l0ZU5hbWUiOiJSb3V0ZWJhc2UxIiwic3JjVGFwSXAiOiIxNjkuMjU0LjY0LjIiLCJkc3RUYXB
BJcCI6IjE2OS4yNTQuNjQuMSIsImlrZU9wdG1
vbiI6ImlrZXZyYiwiZW5jYXBQcm90byI6ImdyZS9pcHNlYyIsImRoR3JvdXAiOiJkaDE0IiwiZW5jcnlwdEFuZERpZ2
VzdCI6ImFlcylnY20vc2hhLTI1NiIsInBzayI
6IlZNd2FyZTEyMyIsInRlbn5lbHMiOlt7ImxvY2FsSWQiOiI2MC42MC42MC4xIiwicGVlcklkIjoNTAuNTAuNTAuMS
IsImxvY2FsVnRpSXAiOiIxNjkuMi4yLjMvMzEifV19"
  }
]
```

- 5 Copy the peer code, which you use to configure the L2 VPN client service and session.

Using the preceding configuration file example, the following peer code is what you copy to use with the L2 VPN client configuration.

```
MCw3ZjBjYzdjLHsic2l0ZU5hbWUiOiJSb3V0ZWJhc2UxIiwic3JjVGFwSXAiOiIxNjkuMjU0LjY0LjIiLCJkc3RUYXB
JcCI6IjE2OS4yNTQuNjQuMSIsImlrZU9wdG1
vbiI6ImlrZXZyYiwiZW5jYXBQcm90byI6ImdyZS9pcHNlYyIsImRoR3JvdXAiOiJkaDE0IiwiZW5jcnlwdEFuZERpZ2
VzdCI6ImFlcylnY20vc2hhLTI1NiIsInBzayI
6IlZNd2FyZTEyMyIsInRlbn5lbHMiOlt7ImxvY2FsSWQiOiI2MC42MC42MC4xIiwicGVlcklkIjoNTAuNTAuNTAuMS
IsImxvY2FsVnRpSXAiOiIxNjkuMi4yLjMvMzEifV19
```

What to do next

Configure the L2 VPN Client service and session. See [Add an L2 VPN Client Service](#) and [Add an L2 VPN Client Session](#).

Add Local Endpoints

You must configure a local endpoint to use with the IPsec VPN that you are configuring.

The following steps use the **Local Endpoints** tab on the NSX Manager UI. You can also create a local endpoint while in the process of adding an IPsec VPN session by clicking the three-dot menu (⋮) and selecting **Add Local Endpoint**. If you are in the middle of configuring an IPsec VPN session, proceed to step 3 in the following steps to guide you with creating a new local endpoint.

Prerequisites

- If you are using a certificate-based authentication mode for the IPsec VPN session that is to use the local endpoint you are configuring, obtain the information about the certificate that the local endpoint must use.
- Ensure that you have configured an IPsec VPN service to which this local endpoint is to be associated.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Navigate to **Networking > VPN > Local Endpoints** and click **Add Local Endpoint**.
- 3 Enter a name for the local endpoint.
- 4 From the **VPN Service** drop-down menu, select the IPsec VPN service to which this local endpoint is to be associated.
- 5 Enter an IP address for the local endpoint.

For an IPsec VPN service running on a Tier-0 gateway, the local endpoint IP address must be different from the Tier-0 gateway's uplink interface IP address. The local endpoint IP address you provide is associated with the loopback interface for the Tier-0 gateway and is also published as a routable IP address over the uplink interface. For IPsec VPN service running on a Tier-1 gateway, in order for the local endpoint IP address to be routable, the route advertisement for IPsec local endpoints must be enabled in the Tier-1 gateway configuration. See [Add a Tier-1 Gateway](#) for more information.

- 6 If you are using a certificate-based authentication mode for the IPsec VPN session, from the **Site Certificate** drop-down menu, select the certificate that is to be used by the local endpoint.
- 7 (Optional) Optionally add a description in **Description**.
- 8 Enter the **Local ID** value that is used for identifying the local NSX Edge instance.

This local ID is the peer ID on the remote site. The local ID must be either the public IP address or FQDN of the remote site. For IPsec VPN sessions with certificate-based authentication and are associated with the local endpoint, the **Local ID** is derived from the certificate associated with the local endpoint. The ID specified in the **Local ID** text box is ignored. The local ID derived from the certificate for a VPN session depends on the extensions present in the certificate.

- If the X509v3 extension `X509v3 Subject Alternative Name` is not present in the certificate, then the Distinguished Name (DN) is used as the local ID value.

For example, if the certificate does not have any Subject Alternative Name (SAN) fields and its DN string is:

```
C=US, ST=California, O=MyCompany, OU=MyOrg, CN=Site123
```

then the DN string is used as the local ID. This local ID is the peer ID on the remote site.

- If the X509v3 extension `X509v3 Subject Alternative Name` is found in the certificate, then one of the SAN fields is taken as the local ID value.

If the certificate has multiple SAN fields, then the following order is used to select the local ID.

Order	SAN Field
1	IP Address
2	DNS
3	Email Address

For example, if the configured site certificate has the following SAN fields:

```
x509v3 Subject Alternative Name:
DNS:Site123.vmware.com, email:user@company.com, IP Address:1.1.1.1
```

then the IP address 1.1.1.1 is used as the local ID. If the IP address is not available, then the DNS string is used. And if the IP address and DNS are not available, then the email address is used.

To see the local ID that is used for an IPSec VPN session, do the following:

- a Navigate to **Networking > VPN** and then click the **IPSec Sessions** tab.
 - b Expand the IPSec VPN session.
 - c Click **Download Config** to download the configuration file which contains the local ID.
- 9 From the **Trusted CA Certificates** and **Certificate Revocation List** drop-down menus, select the appropriate certificates that are required for the local endpoint.
 - 10 (Optional) Specify a tag.
 - 11 Click **Save**.

Adding Profiles

NSX-T Data Center provides the system-generated IPSec tunnel profile and an IKE profile that are assigned by default when you configure either an IPSec VPN or L2 VPN service. A system-generated DPD profile is created for an IPSec VPN configuration.

The IKE and IPSec profiles provide information about the algorithms that are used to authenticate, encrypt, and establish a shared secret between network sites. The DPD profile provides information about the number of seconds to wait in between probes to detect if an IPSec peer site is alive or not.

If you decide not to use the default profiles provided by NSX-T Data Center, you can configure your own profile using the information in the topics that follow in this section.

Add IKE Profiles

The Internet Key Exchange (IKE) profiles provide information about the algorithms that are used to authenticate, encrypt, and establish a shared secret between network sites when you establish an IKE tunnel.

NSX-T Data Center provides system-generated IKE profiles that are assigned by default when you configure an IPSec VPN or L2 VPN service. The following table lists the default profiles provided.

Table 6-4. Default IKE Profiles Used for IPSec VPN or L2 VPN Services

Default IKE Profile Name	Description
nsx-default-l2vpn-ike-profile	■ Used for an L2 VPN service configuration. ■ Configured with IKE V2, AES CBC 128 encryption algorithm, SHA2 256 algorithm, and Diffie-Hellman group14 key exchange algorithm.
nsx-default-l3vpn-ike-profile	■ Used for an IPSec VPN service configuration. ■ Configured with IKE V2, AES CBC 128 encryption algorithm, SHA2 256 algorithm, and Diffie-Hellman group 14 key exchange algorithm.

Instead of the default IKE profiles used, you can also select one of the compliance suites supported starting with NSX-T Data Center 2.5. See [About Supported Compliance Suites](#) for more information.

If you decide not to use the default IKE profiles or compliance suites provided, you can configure your own IKE profile using the following steps.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > VPN** and then click the **Profiles** tab.
- 3 Select the **IKE Profiles** profile type, and click **Add IKE Profile**.
- 4 Enter a name for the IKE profile.
- 5 From the **IKE Version** drop-down menu, select the IKE version to use to set up a security association (SA) in the IPSec protocol suite.

Table 6-5. IKE Versions

IKE Version	Description
IKEv1	When selected, the IPSec VPN initiates and responds to an IKEv1 protocol only.
IKEv2	This version is the default. When selected, the IPSec VPN initiates and responds to an IKEv2 protocol only.
IKE-Flex	If this version is selected and if the tunnel establishment fails with the IKEv2 protocol, the source site does not fall back and initiate a connection with the IKEv1 protocol. Instead, if the remote site initiates a connection with the IKEv1 protocol, then the connection is accepted.

- 6 Select the encryption, digest, and Diffie-Hellman group algorithms from the drop-down menus. You can select multiple algorithms to apply or deselect any selected algorithms you do not want to be applied.

Table 6-6. Algorithms Used

Type of Algorithm	Valid Values	Description
Encryption	■ AES 128 (default) ■ AES 256 ■ AES GCM 128 ■ AES GCM 192 ■ AES GCM 256	The encryption algorithm used during the Internet Key Exchange (IKE) negotiation. The AES 128 and AES 256 algorithms use the CBC mode of operation. The AES-GCM algorithms are supported when used with IKEv2. They are not supported when used with IKEv1.
Digest	■ SHA2 256 (default) ■ SHA1 ■ SHA2 384 ■ SHA2 512	The secure hashing algorithm used during the IKE negotiation. If AES-GCM is the only encryption algorithm selected in the Encryption Algorithm text box, then no hash algorithms can be specified in the Digest Algorithm text box, per section 8 in RFC 5282. In addition, the Pseudo-Random Function (PRF) algorithm PRF-HMAC-SHA2-256 is implicitly selected and used in the IKE security association (SA) negotiation. The PRF-HMAC-SHA2-256 algorithm must also be configured on the peer gateway in order for the phase 1 of the IKE SA negotiation to succeed. If more algorithms are specified in the Encryption Algorithm text box, in addition to the AES-GCM algorithm, then one or more hash algorithms can be selected in the Digest Algorithm text box. In addition, the PRF algorithm used in the IKE SA negotiation is implicitly determined based on the hash algorithms configured. At least one of the matching PRF algorithms must also be configured on the peer gateway in order for the phase 1 of the IKE SA negotiation to succeed. For example, if the Encryption Algorithm text box contains AES 128 and AES GCM 128 and SHA1 is specified in the Digest Algorithm text box, then the PRF-HMAC-SHA1 algorithm is used during the IKE SA negotiation. It must also be configured in the peer gateway.
Diffie-Hellman Group	■ Group 14 (default) ■ Group 2 ■ Group 5 ■ Group 15 ■ Group 16 ■ Group 19 ■ Group 20 ■ Group 21	The cryptography schemes that the peer site and the NSX Edge use to establish a shared secret over an insecure communications channel.

Note When you attempt to establish an IPSec VPN tunnel with a GUARD VPN Client (previously QuickSec VPN Client) using two encryption algorithms or two digest algorithms, the GUARD VPN Client adds additional algorithms in the proposed negotiation list. For example, if you specified AES 128 and AES 256 as the encryption algorithms and SHA2 256 and SHA2 512 as the digest algorithms to use in the IKE profile you are using to establish the IPSec VPN tunnel, the GUARD VPN Client also proposes AES 192 (using CBC mode) and SHA2 384 in the negotiation list. In this case, NSX-T Data Center uses the first encryption algorithm you selected when establishing the IPSec VPN tunnel.

- 7 Enter a security association (SA) lifetime value, in seconds, if you want it different from the default value of 86400 seconds (24 hours).
- 8 Provide a description and add a tag, as needed.
- 9 Click **Save**.

Results

A new row is added to the table of available IKE profiles. To edit or delete a non-system created profile, click the three-dot menu (⋮) and select from the list of actions available.

Add IPSec Profiles

The Internet Protocol Security (IPSec) profiles provide information about the algorithms that are used to authenticate, encrypt, and establish a shared secret between network sites when you establish an IPSec tunnel.

NSX-T Data Center provides system-generated IPSec profiles that are assigned by default when you configure an IPSec VPN or L2 VPN service. The following table lists the default IPSec profiles provided.

Table 6-7. Default IPSec Profiles Used for IPSec VPN or L2 VPN Services

Name of Default IPSec Profile	Description
nsx-default-l2vpn-tunnel-profile	■ Used for L2 VPN. ■ Configured with AES GCM 128 encryption algorithm and Diffie-Hellman group 14 key exchange algorithm.
nsx-default-l3vpn-tunnel-profile	■ Used for IPSec VPN. ■ Configured with AES GCM 128 encryption algorithm and Diffie-Hellman group 14 key exchange algorithm.

Instead of the default IPSec profile, you can also select one of the compliance suites supported starting with NSX-T Data Center 2.5. See [About Supported Compliance Suites](#) for more information.

If you decide not to use the default IPSec profiles or compliance suites provided, you can configure your own using the following steps.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > VPN** and then click the **Profiles** tab.
- 3 Select the **IPSec Profiles** profile type, and click **Add IPSec Profile**.
- 4 Enter a name for the IPSec profile.
- 5 From the drop-down menus, select the encryption, digest, and Diffie-Hellman algorithms. You can select multiple algorithms to apply.

Deselect the ones you do not want used.

Table 6-8. Algorithms Used

Type of Algorithm	Valid Values	Description
Encryption	■ AES GCM 128 (default) ■ AES 128 ■ AES 256 ■ AES GCM 192 ■ AES GCM 256 ■ No Encryption Auth AES GMAC 128 ■ No Encryption Auth AES GMAC 192 ■ No Encryption Auth AES GMAC 256 ■ No Encryption	The encryption algorithm used during the Internet Protocol Security (IPSec) negotiation. The AES 128 and AES 256 algorithms use the CBC mode of operation.
Digest	■ SHA1 ■ SHA2 256 ■ SHA2 384 ■ SHA2 512	The secure hashing algorithm used during the IPSec negotiation.
Diffie-Hellman Group	■ Group 14 (default) ■ Group 2 ■ Group 5 ■ Group 15 ■ Group 16 ■ Group 19 ■ Group 20 ■ Group 21	The cryptography schemes that the peer site and NSX Edge use to establish a shared secret over an insecure communications channel.

- 6 Deselect **PFS Group** if you decide not to use the PFS Group protocol on your VPN service. It is selected by default.
- 7 In the **SA Lifetime** text box, modify the default number of seconds before the IPSec tunnel must be re-established.

By default, an SA lifetime of 24 hours (86400 seconds) is used.

- 8 Select the value for **DF Bit** to use with the IPsec tunnel.

The value determines how to handle the "Don't Fragment" (DF) bit included in the data packet received. The acceptable values are described in the following table.

Table 6-9. DF Bit Values

DF Bit Value	Description
COPY	The default value. When this value is selected, NSX-T Data Center copies the value of the DF bit from the received packet into the packet which is forwarded. This value implies that if the data packet received has the DF bit set, after encryption, the packet also has the DF bit set.
CLEAR	When this value is selected, NSX-T Data Center ignores the value of the DF bit in the data packet received, and the DF bit is always 0 in the encrypted packet.

- 9 Provide a description and add a tag, if necessary.
- 10 Click **Save**.

Results

A new row is added to the table of available IPsec profiles. To edit or delete a non-system created profile, click the three-dot menu (⋮) and select from the list of actions available.

Add DPD Profiles

A DPD (Dead Peer Detection) profile provides information about the number of seconds to wait in between probes to detect if an IPsec peer site is alive or not.

NSX-T Data Center provides a system-generated DPD profile, named `nsx-default-l3vpn-dpd-profile`, that is assigned by default when you configure an IPsec VPN service. This default DPD profile is a periodic DPD probe mode.

If you decide not to use the default DPD profile provided, you can configure your own using the following steps.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Navigate to **Networking > VPN > Profiles**.
- 3 Select **DPD Profiles** from the **Select Profile Type** drop-down menu, and click **Add DPD Profile**.
- 4 Enter a name for the DPD profile.
- 5 From the **DPD Probe Mode** drop-down menu, select **Periodic** or **On Demand** mode.

For a periodic DPD probe mode, a DPD probe is sent every time the specified DPD probe interval time is reached.

For an on-demand DPD probe mode, a DPD probe is sent if no IPSec packet is received from the peer site after an idle period. The value in **DPD Probe Interval** determines the idle period used.

- 6 In the **DPD Probe Interval** text box, enter the number of seconds you want the NSX Edge node to wait before sending the next DPD probe.

For a periodic DPD probe mode, the valid values are between 3 and 360 seconds. The default value is 60 seconds.

For an on-demand probe mode, the valid values are between 1 and 10 seconds. The default value is 3 seconds.

When the periodic DPD probe mode is set, the IKE daemon running on the NSX Edge sends a DPD probe periodically. If the peer site responds within half a second, the next DPD probe is sent after the configured DPD probe interval time has been reached. If the peer site does not respond, then the DPD probe is sent again after waiting for half a second. If the remote peer site continues not to respond, the IKE daemon resends the DPD probe again, until a response is received or the retry count has been reached. Before the peer site is declared to be dead, the IKE daemon resends the DPD probe up to a maximum of times specified in the **Retry Count** property. After the peer site is declared dead, the NSX Edge node then tears down the security association (SA) on the dead peer's link.

When the on-demand DPD mode is set, the DPD probe is sent only if no IPSec traffic is received from the peer site after the configured DPD probe interval time has been reached.

- 7 In the **Retry Count** text box, enter the number of retries allowed.

The valid values are between 1 and 100. The default retry count is 5.

- 8 Provide a description and add a tag, as needed.
- 9 To enable or disable the DPD profile, click the **Admin Status** toggle.

By default, the value is set to **Enabled**. When the DPD profile is enabled, the DPD profile is used for all IPSec sessions in the IPSec VPN service that uses the DPD profile.

- 10 Click **Save**.

Results

A new row is added to the table of available DPD profiles. To edit or delete a non-system created profile, click the three-dot menu (⋮) and select from the list of actions available.

Add an Autonomous Edge as an L2 VPN Client

You can use L2 VPN to extend your Layer 2 networks to a site that is not managed by NSX-T Data Center. An autonomous NSX Edge can be deployed on the site, as an L2 VPN client. The autonomous NSX Edge is simple to deploy, easily programmable, and provides high-performance VPN. The autonomous NSX Edge is deployed using an OVF file on a host that is not managed by

NSX-T Data Center. You can also enable high availability (HA) for VPN redundancy by deploying primary and secondary autonomous Edge L2 VPN clients.

Prerequisites

- Create a port group and bind it to the vSwitch on your host. Ensure that this port group accepts promiscuous mode and forged transmits from the port group's security settings.
- Create a port group for your internal L2 extension port.
- Obtain the IP addresses for the local IP and remote IP to use with the L2 VPN Client session you are adding.
- Obtain the peer code that was generated during the L2 VPN server configuration.

Procedure

- 1 Using vSphere Web Client, log in to the vCenter Server that manages the non-NSX environment.
- 2 Select **Hosts and Clusters** and expand clusters to show the available hosts.
- 3 Right-click the host where you want to install the autonomous NSX Edge and select **Deploy OVF Template**.
- 4 Enter the URL to download and install the OVF file from the Internet or click **Browse** to locate the folder on your computer that contains the autonomous NSX Edge OVF file and click **Next**.
- 5 On the **Select name and folder** page, enter a name for the autonomous NSX Edge and select the folder or data center where you want to deploy. Then click **Next**.
- 6 On the **Select a compute resource** page, select the destination of the compute resource.
- 7 On the OVF Template Details page, review the template details and click **Next**.
- 8 On the **Configuration** page, select a deployment configuration option.
- 9 On the **Select storage** page, select the location to store the files for the configuration and disk files.
- 10 On the **Select networks** page, configure the networks that the deployed template must use. Select the port group you created for the uplink interface, the port group that you created for the L2 extension port, and enter an HA interface. Click **Next**.
- 11 On the **Customize Template** page, enter the following values and click **Next**.
 - a Type and retype the CLI admin password.
 - b Type and retype the CLI enable password.
 - c Type and retype the CLI root password.
 - d Enter the IPv4 address for the Management Network.
 - e Enable the option to deploy an autonomous Edge.

- f Enter the **External Port** details for VLAN ID, exit interface, IP address, and IP prefix length such that the exit interface maps to the Network with the port group of your uplink interface.

If the exit interface is connected to a trunk port group, specify a VLAN ID. For example, **20,eth2,192.168.5.1,24**. You can also configure your port group with a VLAN ID and use VLAN 0 for the **External Port**.

- g (Optional) To configure High Availability, enter the **HA Port** details where the exit interface maps to the appropriate HA Network.
- h (Optional) When deploying an autonomous NSX Edge as a secondary node for HA, select **Deploy this autonomous-edge as a secondary node**.

Use the same OVF file as the primary node and enter the primary node's IP address, user name, password, and thumbprint.

To retrieve the thumbprint of the primary node, log in to the primary node and run the following command:

```
get certificate api thumbprint
```

Ensure that the VTEP IP addresses of the primary and secondary nodes are in the same subnet and that they connect to the same port group. When you complete the deployment and start the secondary-edge, it connects to the primary node to form an edge-cluster .

- 12 On the **Ready to complete** page, review the autonomous Edge settings and click **Finish**.

Note If there are errors during the deployment, a message of the day is displayed on the CLI. You can also use an API call to check for errors:

```
GET https://<nsx-mgr>/api/v1/node/status
```

The errors are categorized as soft errors and hard errors. Use API calls to resolve the soft errors as required. You can clear the message of day using an API call:

```
POST /api/v1/node/status?action=clear_bootup_error
```

- 13 Power on the autonomous NSX Edge appliance.
- 14 Log in to the autonomous NSX Edge client.
- 15 Select **L2VPN > Add Session** and enter the following values:
 - a Enter a session name.
 - b Enter the local IP address and the remote IP address.
 - c Enter the peer code from the L2VPN server. See [Download the Remote Side L2 VPN Configuration File](#) for details on obtaining the peer code.
- 16 Click **Save**.

- 17 Select **Port > Add Port** to create an L2 extension port.
- 18 Enter a name, a VLAN, and select an exit interface.
- 19 Click **Save**.
- 20 Select **L2VPN > Attach Port** and enter the following values:
 - a Select the L2 VPN session that you created.
 - b Select the L2 extension port that you created.
 - c Enter a tunnel ID.
- 21 Click **Attach**.

You can create additional L2 extension ports and attach them to the session if you need to extend multiple L2 networks.

- 22 Use the browser to log in to the autonomous NSX Edge or use API calls to view the status of the L2VPN session.

Note If the L2VPN server configuration changes, ensure that you download the peer code again and update the session with the new peer code.

Check the Realized State of an IPSec VPN Session

After you send a configuration update request for an IPSec VPN session, you can check to see if the requested state has been successfully processed in the NSX-T Data Center local control plane on the transport nodes.

When you create an IPSec VPN session, multiple entities are created: IKE profile, DPD profile, tunnel profile, local endpoint, IPSec VPN service, and IPSec VPN session. These entities all share the same `IPSecVPNSession` span, so you can obtain the realization state of all the entities of the IPSec VPN session by using the same `GET` API call. You can check the realization state using only the API.

Prerequisites

- Familiarize yourself with IPSec VPN. See [Understanding IPSec VPN](#).
- Verify the IPSec VPN is configured successfully. See [Add an IPSec VPN Service](#).
- You must have access to the NSX Manager API.

Procedure

- 1 Send a `POST`, `PUT`, or `DELETE` request API call.

For example:

```
PUT https://<nsx-mgr>/api/v1/vpn/ipsec/sessions/8dd1c386-9b2c-4448-85b8-51ff649fae4f
{
  "resource_type": "PolicyBasedIPSecVPNSession",
  "id": "8dd1c386-9b2c-4448-85b8-51ff649fae4f",
  "display_name": "Test RZ_UPDATED",
  "ipsec_vpn_service_id": "7adfa455-a6fc-4934-a919-f5728957364c",
  "peer_endpoint_id": "17263ca6-dce4-4c29-bd8a-e7d12bd1a82d",
  "local_endpoint_id": "91ebfa0a-820f-41ab-bd87-f0fb1f24e7c8",
  "enabled": true,
  "policy_rules": [
    {
      "id": "1026",
      "sources": [
        {
          "subnet": "1.1.1.0/24"
        }
      ],
      "logged": true,
      "destinations": [
        {
          "subnet": "2.1.4..0/24"
        }
      ],
      "action": "PROTECT",
      "enabled": true,
      "_revision": 1
    }
  ]
}
```

- 2 Locate and copy the value of `x-nsx-requestid` from the response header returned.

For example:

```
x-nsx-requestid    e550100d-f722-40cc-9de6-cf84d3da3ccb
```

- 3 Request the realization state of the IPsec VPN session using the following `GET` call.

```
GET https://<nsx-mgr>/api/v1/vpn/ipsec/sessions/<ipsec-vpn-session-id>/state?request_id=<request-id>
```

The following API call uses the `id` and `x-nsx-requestid` values in the examples used in the previous steps.

```
GET https://<nsx-mgr>/api/v1/vpn/ipsec/sessions/8dd1c386-9b2c-4448-85b8-51ff649fae4f/state?request_id=e550100d-f722-40cc-9de6-cf84d3da3ccb
```

Following is an example of a response you receive when the realization state is `in_progress`.

```
{
  "details": [
    {
      "sub_system_type": "TransportNode",
```

```

    "sub_system_id": "fe651e63-04bd-43a4-a8ec-45381a3b71b9",
    "state": "in_progress",
    "failure_message": "CCP Id:ab5958df-d98a-468e-a72b-d89dcdae5346, Message:State
realization is in progress at the node."
  },
  {
    "sub_system_type": "TransportNode",
    "sub_system_id": "ebe174ac-e4f1-4135-ba72-3dd2eb7099e3",
    "state": "in_sync"
  }
],
"state": "in_progress",
"failure_message": "The state realization is in progress at transport nodes."
}

```

Following is an example of a response you receive when the realization state is `in_sync`.

```

{
  "details": [
    {
      "sub_system_type": "TransportNode",
      "sub_system_id": "7046e8f4-a680-11e8-9bc3-020020593f59",
      "state": "in_sync"
    }
  ],
  "state": "in_sync"
}

```

The following are examples of possible responses you receive when the realization state is `unknown`.

```

{
  "state": "unknown",
  "failure_message": "Unable to get response from any CCP node. Please retry operation
after some time."
}

```

```

{
  "details": [
    {
      "sub_system_type": "TransportNode",
      "sub_system_id": "3e643776-5def-11e8-94ae-020022e7749b",
      "state": "unknown",
      "failure_message": "CCP Id:ab5958df-d98a-468e-a72b-d89dcdae5346, Message: Unable
to get response from the node. Please retry operation after some time."
    },
    {
      "sub_system_type": "TransportNode",
      "sub_system_id": "4784ca0a-5def-11e8-93be-020022f94b73",
      "state": "in_sync"
    }
  ]
}

```

```

    ],
    "state": "unknown",
    "failure_message": "The state realization is unknown at transport nodes"
  }

```

After you perform an entity `DELETE` operation, you might receive the status of `NOT_FOUND`, as shown in the following example.

```

{
  "http_status": "NOT_FOUND",
  "error_code": 600,
  "module_name": "common-services",
  "error_message": "The operation failed because object identifier LogicalRouter/61746f54-7ab8-4702-93fe-6ddeb804 is missing: Object identifiers are case sensitive.."
}

```

If the IPSec VPN service associated with the session is disabled, you receive the `BAD_REQUEST` response, as shown in the following example.

```

{
  "httpStatus": "BAD_REQUEST",
  "error_code": 110199,
  "module_name": "VPN",
  "error_message": "VPN service f9cfe508-05e3-4e1d-b253-fed096bb2b63 associated with the session 8dd1c386-9b2c-4448-85b8-51ff649fae4f is disabled. Can not get the realization status."
}

```

Monitor and Troubleshoot VPN Sessions

After you configure an IPSec or L2 VPN session, you can monitor the VPN tunnel status and troubleshoot any reported tunnel issues using the NSX Manager user interface.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Navigate to the **Networking > VPN > IPSec Sessions** or **Networking > VPN > L2 VPN Sessions** tab.
- 3 Expand the row for the VPN session that you want to monitor or troubleshoot.
- 4 To view the status of the VPN tunnel status, click the info icon.
The Status dialog box appears and displays the available statuses.
- 5 To view the VPN tunnel traffic statistics, click **View Statistics** in the Status column.
The Statistics dialog box displays the traffic statistics for the VPN tunnel.
- 6 To view the error statistics, click the **View More** link in the Statistics dialog box.
- 7 To close the **Statistics** dialog box, click **Close**.

Network Address Translation (NAT)

7

Network address translation (NAT) maps one IP address space to another. You can configure NAT on tier-0 and tier-1 gateways.

The following types of NAT are supported, in addition to NAT64:

- Source NAT (SNAT) - translates a source IP address of outbound packets so that packets appears as originating from a different network. Supported on tier-0/tier-1 gateways running in active-standby mode. For one-to-one SNAT, the SNAT translated IP address is not programmed on the loopback port, and there is no forwarding entry with an SNAT translated IP as the prefix. For n-to-one SNAT, the SNAT translated IP address is programmed on the loopback port, and users will see a forwarding entry with an SNAT translated IP address prefix.
- Destination NAT (DNAT) - translates the destination IP address of inbound packets so that packets are delivered to a target address into another network. Supported on tier-0/tier-1 gateways running in active-standby mode.
- Reflexive NAT - (sometimes called stateless NAT) translates addresses passing through a routing device. Inbound packets undergo destination address rewriting, and outbound packets undergo source address rewriting. It is not keeping a session as it is stateless. Supported on tier-0 gateways running in active-active mode. Stateful NAT is not supported in active-active mode.

You can also disable SNAT or DNAT for an IP address or a range of addresses. If an address has multiple NAT rules, the rule with the highest priority is applied.

Note DNAT is not supported on a tier-1 gateway where policy-based IPsec VPN is configured.

SNAT configured on a tier-0 gateway's external interface processes traffic from a tier-1 gateway, and from another external interface on the tier-0 gateway.

Note NAT is configured on the uplinks of the tier-0/tier-1 gateways and processes traffic going through this interface. This implies that tier-0 gateway NAT rules will not apply between two tier-1 gateways connected to the tier-0.

NAT64 is a mechanism for translating IPv6 packets to IPv4 packets, and vice versa. NAT 64 allows IPv6-only clients to contact IPv4 servers using unicast UDP, or TCP. NAT64 only allows an IPv6-only client to initiate communications to an IPv4-only server. To perform IPv6-IPv4 translation, binding and session information are saved. NAT64 is stateful.

- NAT64 is only supported for external IPv6 traffic coming in through the NSX-T edge uplink to the IPv4 server in the overlay.
- NAT64 supports TCP and UDP, all other protocol type packets are discarded. NAT64 does not support: ICMP, Fragmentation, and IPV6 packets that have extension headers.

Note When a NAT64 rule and an inline load balancer are configured on the same edge node, using the NAT64 rule to direct IPV6 packets to the IPv4 inline load balancer is not supported.

Read the following topics next:

- [Configure NAT on a Gateway](#)

Configure NAT on a Gateway

You can configure NAT and NAT 64 rules on a tier-0 or tier-1 gateway.

Note If there is a service configured in this NAT rule, the translated_port will be realized on NSX Manager as the destination_port. This means the service will be the translated port while the translated port is used to match the traffic as destination port. If there is no service configured, the port will be ignored

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > NAT**.
- 3 Select a gateway.
- 4 Next to **View**, select **NAT** or **NAT64**.
- 5 Click **Add NAT Rule** or **Add NAT 64 Rule**.
- 6 Enter a **Name**.
- 7 If you are configuring NAT, select an action. For **NAT 64**, the action is NAT64.

NAT Option	Description
Tier-1 gateway	Available actions are SNAT , DNAT , Reflexive , NO SNAT , and NO DNAT .
Tier-0 gateway in active-standby mode	Available actions are SNAT , DNAT , NO SNAT , and NO DNAT .
Tier-0 gateway in active-active mode	The available action is Reflexive .

- 8 Enter a **Source**. If this text box is left blank, the NAT rule applies to all sources outside of the local subnet.

Option	Description
NAT	Specify an IP address, or an IP address range in CIDR format. For SNAT, NO_SNAT and REFLEXIVE rules, this is a mandatory text box and represents the source network of the packets leaving the network.
NAT64	Enter an IPv6 address, or an IPv6 CIDR.

- 9 (Required) Enter a **Destination**.

Option	Description
NAT	Specify an IP address, or an IP address range in CIDR format.
NAT64	Enter an IPv6 address, or an IPv6 address range in CIDR format with the prefix /96. The prefix /96 is supported because the destination IPv4 IP is embedded as the last 4 bytes in the IPv6 address

- 10 Enter a value for **Translated IP**.

Option	Description
NAT	Specify an IPv4 address, or an IP address range in CIDR format.
NAT64	Specify an IPv4 address, a comma-separated list of IPv4 addresses, or an IPv4 address range. IPV4 CIDR is not supported.

- 11 Toggle **Enable** to enable the rule.
- 12 In the **Service** column, click **Set** to select services. See [Add a Service](#) for more information. For NAT 64, select a pre-defined service or create a user-defined service with TCP or UDP, with the source/destination port as **Any**, or a specific port.

13 For **Apply To**, click **Set** and select objects that this rule applies to.

The available objects are **Tier-0 Gateways**, **Interfaces**, **Labels**, **Service Instance Endpoints**, and **Virtual Endpoints**.

Note If you are using NSX Federation and creating a NAT rule from a Global Manager appliance, you can select site-specific IP addresses for NAT. You can apply the NAT rule to any of the following location spans:

- Do not click **Set** if you want to use the default option of applying the NAT rule to all locations.
- Click **Set**. In the **Apply To** dialog box, select the locations whose entities you want to apply the rule to and then select **Apply NAT rule to all entities**.
- Click **Set**. In the **Apply To** dialog box, select a location and then select **Interfaces** from the **Categories** drop-down menu. You can select specific interfaces to which you want to apply the NAT rule.

See [Features and Configurations Supported in Federation](#) for more details.

14 Enter a value for **Translated Port**.

15 Select a firewall setting.

Option	Description
NAT	Available settings are: ■ Match External Address - The packet is processed by firewall rules that match the combination of translated IP address, and translated port. ■ For SNAT, the external address is the translated source address after NAT is done. ■ For DNAT, the external address is the original destination address before NAT is done. ■ For REFLEXIVE, to egress traffic, the firewall is applied to the translated source address after NAT is done. For ingress traffic, the firewall is applied to the original destination address before NAT is done. ■ Match Internal Address - The packet is processed by firewall rules that match the combination of original IP address, and original port. ■ For SNAT, the internal address is the original source address before NAT is done. ■ For DNAT, the internal address is the translated destination address after NAT is done. ■ For REFLEXIVE, for egress traffic, the firewall is applied to the original source address before NAT is done. For ingress traffic, the firewall is applied to the translated destination address after NAT is done. ■ Bypass - The packet bypasses firewall rules.
NAT64	The available setting is Bypass - the packet bypasses firewall rules.

16 (Optional) Toggle the logging button to enable logging.

17 Specify a priority value.


A lower value means a higher priority. The default is 0.

18 Click **Save**.

Load Balancing

8

The NSX-T Data Center logical load balancer offers high-availability service for applications and distributes the network traffic load among multiple servers.



The load balancer distributes incoming service requests evenly among multiple servers in such a way that the load distribution is transparent to users. Load balancing helps in achieving optimal resource utilization, maximizing throughput, minimizing response time, and avoiding overload.

You can map a virtual IP address to a set of pool servers for load balancing. The load balancer accepts TCP, UDP, HTTP, or HTTPS requests on the virtual IP address and decides which pool server to use.

Depending on your environment needs, you can scale the load balancer performance by increasing the existing virtual servers and pool members to handle heavy network traffic load.

Note Logical load balancer is supported only on the tier-1 gateway. One load balancer can be attached only to a tier-1 gateway.

Read the following topics next:


- [Key Load Balancer Concepts](#)
- [Setting Up Load Balancer Components](#)
- [Groups Created for Server Pools and Virtual Servers](#)

Key Load Balancer Concepts

Load balancer includes virtual servers, server pools, and health checks monitors.

A load balancer is connected to a Tier-1 logical router. The load balancer hosts single or multiple virtual servers. A virtual server is an abstract of an application service, represented by a unique combination of IP, port, and protocol. The virtual server is associated to single to multiple server pools. A server pool consists of a group of servers. The server pools include individual server pool members.

To test whether each server is correctly running the application, you can add health check monitors that check the health status of a server.



Scaling Load Balancer Resources

When you configure a load balancer, you can specify a size (small, medium, large, or extra large). The size determines the number of virtual servers, server pools, and pool members the load balancer can support.

A load balancer runs on a tier-1 gateway, which must be in active-standby mode. The gateway runs on NSX Edge nodes. The form factor of the NSX Edge node (bare metal, small, medium, large, or extra large) determines the number of load balancers that the NSX Edge node can support. Note that in Manager mode, you create logical routers, which have similar functionality to gateways. See [Chapter 1 NSX Manager](#).

For more information about what the different load balance sizes and NSX Edge form factors can support, see <https://configmax.vmware.com>.

Note that using a small NSX Edge node to run a small load balancer is not recommended in a production environment.

You can call an API to get the load balancer usage information of an NSX Edge node. If you use Policy mode to configure load balancing, run the following command:

```
GET /policy/api/v1/infra/lb-node-usage?node_path=<node-path>
```

If you use Manager mode to configure load balancing, run the following command:

```
GET /api/v1/loadbalancer/usage-per-node/<node-id>
```


The usage information includes the number of load balancer objects (such as load balancer services, virtual servers, server pools, and pool members) that are configured on the node. For more information, see the *NSX-T Data Center API Guide*.

Supported Load Balancer Features

NSX-T Data Center load balancer supports the following features.

- Layer 4 - TCP and UDP
- Layer 7 - HTTP and HTTPS with load balancer rules support
- Server pools - static and dynamic with NSGroup
- Persistence - Source-IP and Cookie persistence mode
- Health check monitors - Active monitor which includes HTTP, HTTPS, TCP, UDP, and ICMP, and passive monitor
- SNAT - Transparent, Automap, and IP List
- HTTP upgrade - For applications using HTTP upgrade such as WebSocket, the client or server requests for HTTP Upgrade, which is supported. By default, NSX-T Data Center supports and accepts HTTPS upgrade client request using the HTTP application profile.

Note: SSL -Terminate-mode and proxy-mode is not supported in NSX-T Data Center limited export release.




Load Balancer Topologies

Load balancers are typically deployed in either inline or one-arm mode. One-arm mode requires virtual server Source NAT (SNAT) configuration, and inline mode does not.

Inline Topology


In the inline mode, the load balancer is in the traffic path between the client and the server. Clients and servers should not be connected to overlay segments on the same tier-1 logical router if SNAT on the load balancer is not desired. If clients and servers are connected to overlay segments on the same tier-1 logical router, SNAT is required.



One-Arm Topology

In one-arm mode, the load balancer is not in the traffic path between the client and the server. In this mode, the client and the server can be anywhere. The load balancer performs Source NAT (SNAT) to force return traffic from the server destined to the client to go through the load balancer. This topology requires virtual server SNAT to be enabled.

When the load balancer receives the client traffic to the virtual IP address, the load balancer selects a server pool member and forwards the client traffic to it. In the one-arm mode, the load balancer replaces the client IP address with the load balancer IP address so that the server response is always sent to the load balancer. The load balancer forwards the response to the client.



Special Use Case When no overlay is used and everything is VLAN, overlay still must be configured on the edge nodes hosting the tier-1 one-arm load balancer. This is because edge nodes must have at least one tunnel end point UP for high availability between edge nodes. When the tunnels are UP they will agree on which edge node will run the active or standby role of each tier-0 and tier-1 gateway.

Tier-1 Service Chaining

If a tier-1 gateway or logical router hosts different services, such as NAT, firewall, and load balancer, the services are applied in the following order:

- Ingress

DNAT - Firewall - Load Balancer

Note: If DNAT is configured with Firewall Bypass, firewall is skipped but not load balancer.


- Egress

Load Balancer - Firewall - SNAT

Setting Up Load Balancer Components

To use logical load balancers, you must start by configuring a load balancer and attaching it to a tier-1 gateway.

Next, you set up health check monitoring for your servers. You must then configure server pools for your load balancer. Finally, you must create a layer 4 or layer 7 virtual server for your load balancer and attach the newly created virtual server to the load balancer.




Add Load Balancers

Load balancer is created and attached to the tier-1 gateway.

You can configure the level of error messages you want the load balancer to add to the error log.

Note Avoid setting the log level to DEBUG on load balancers with a significant traffic due to the number of messages printed to the log that affect performance.



Prerequisites

Verify that a tier-1 gateway is configured. See [Chapter 3 Tier-1 Gateway](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Load Balancing > Add Load Balancer**.
- 3 Enter a name and a description for the load balancer.
- 4 Select the load balancer size based on your needs of virtual servers and pool members and available resources.
- 5 Select the already configured tier-1 gateway to attach to this load balancer from the drop-down menu.

The tier-1 gateway must be in the Active-Standby mode.

- 6 Define the severity level of the error log from the drop-down menu.

Load balancer collects information about encountered issues of different severity levels to the error log.

- 7 (Optional) Enter tags to make searching easier.

You can specify a tag to set a scope of the tag.

- 8 Click **Save**.

The load balancer creation and attaching the load balancer to the tier-1 gateway takes about three minutes and the configuration status to appear green and Up.

If the status is Down, click the information icon and resolve the error before you proceed.

- 9 (Optional) Delete the load balancer.

- a Detach the load balancer from the virtual server and tier-1 gateway.
- b Select the load balancer.

- c Click the vertical ellipses button.
- d Select **Delete**.


Add an Active Monitor

The active health monitor is used to test whether a server is available. The active health monitor uses several types of tests such as sending a basic ping to servers or advanced HTTP requests to monitor an application health.

Servers that fail to respond within a certain time period or respond with errors are excluded from future connection handling until a subsequent periodic health check finds these servers to be healthy.

Active health checks are performed on server pool members after the pool member is attached to a virtual server and that virtual server is attached to a tier-1 gateway. The tier-1 uplink IP address is used for the health check.

Note More than one active health monitor can be configured per server pool.



Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Load Balancing > Monitors > Active > Add Active Monitor**.
- 3 Select a protocol for the server from the drop-down menu.
You can also use predefined protocols; HTTP, HTTPS, ICMP, TCP, and UDP for NSX Manager.
- 4 Select the **HTTP** protocol.

5 Configure the values to monitor a service pool.

You can also accept the default active health monitor values.

Option	Description
Name and Description	Enter a name and description for the active health monitor.
Monitoring Port	Set the value of the monitoring port. There must be a port in either the pool member or the monitor. Otherwise, the monitor will not work.
Monitoring Interval	Set the time in seconds that the monitor sends another connection request to the server.
Timeout Period	Set the time the load balancer will wait for the Pool Member monitor response before considering failed.
Fail Count	Set a value when the consecutive failures reach this value, the server is considered temporarily unavailable.
Rise Count	Set the value of consecutive successful monitors to reach before changing the Pool Member Status from Down to Up.
Tags	Enter tags to make searching easier. You can specify a tag to set a scope of the tag.

For example, if the monitoring interval is set as 5 seconds and the timeout as 15 seconds, the load balancer send requests to the server every 5 seconds. In each probe if the expected response is received from the server within 15 seconds, the health check result is OK. If not, then the result is CRITICAL. If the recent three health check results are all UP, the server is considered as UP.

6 To configure the HTTP Request, click **Configure**.

7 Enter the HTTP request and response configuration details.

Option	Description
HTTP Method	Select the method to detect the server status from the drop-down menu, GET, OPTIONS, POST, HEAD, and PUT.
HTTP Request URL	Enter the request URI for the method. ASCII control characters (backspace, vertical tab, horizontal tab, line feed, etc), unsafe characters such as a space, \, <, >, {, }, and any character outside the ASCII character set are not allowed in the request URL and should be encoded. For example, replace a space with a plus (+) sign, or with %20.
HTTP Request Version	Select the supported request version from the drop-down menu. You can also accept the default version, HTTP_VERSION_1.
HTTP Request Header	Click Add and enter the HTTP request header name and corresponding value.
HTTP Request Body	Enter the request body. Valid for the POST and PUT methods.

Option	Description
HTTP Response Code	Enter the string that the monitor expects to match in the status line of HTTP response body. The response code is a comma-separated list. For example, 200,301,302,401.
HTTP Response Body	If the HTTP response body string and the HTTP health check response body match, then the server is considered as healthy.

- 8 Click **Save**.
- 9 Select the **HTTPS** protocol from the drop-down list.
- 10 Complete step 5.
- 11 Click **Configure**.
- 12 Enter the HTTP request and response and SSL configuration details.

Option	Description
Name and Description	Enter a name and description for the active health monitor.
HTTP Method	Select the method to detect the server status from the drop-down menu, GET, OPTIONS, POST, HEAD, and PUT.
HTTP Request URL	Enter the request URI for the method. ASCII control characters (backspace, vertical tab, horizontal tab, line feed, etc), unsafe characters such as a space, \, <, >, {, }, and any character outside the ASCII character set are not allowed in the request URL and should be encoded. For example, replace a space with a plus (+) sign, or with %20.
HTTP Request Version	Select the supported request version from the drop-down menu. You can also accept the default version, HTTP_VERSION_1.
HTTP Request Header	Click Add and enter the HTTP request header name and corresponding value.
HTTP Request Body	Enter the request body. Valid for the POST and PUT methods.
HTTP Response Code	Enter the string that the monitor expects to match in the status line of HTTP response body. The response code is a comma-separated list. For example, 200,301,302,401.
HTTP Response Body	If the HTTP response body string and the HTTP health check response body match, then the server is considered as healthy.
Server SSL	Toggle the button to enable the SSL server.
Client Certificate	(Optional) Select a certificate from the drop-down menu to be used if the server does not host multiple host names on the same IP address or if the client does not support an SNI extension.
Server SSL Profile	(Optional) Assign a default SSL profile from the drop-down menu that defines reusable and application-independent client-side SSL properties. Click the vertical ellipses and create a custom SSL profile.

Option	Description
Trusted CA Certificates	(Optional) You can require the client to have a CA certificate for authentication.
Mandatory Server Authentication	(Optional) Toggle the button to enable server authentication.
Certificate Chain Depth	(Optional) Set the authentication depth for the client certificate chain.
Certificate Revocation List	(Optional) Set a Certificate Revocation List (CRL) in the client-side SSL profile to reject compromised client certificates.

13 Select the **ICMP** protocol.

14 Complete step 5 and assign the data size in byte of the ICMP health check packet.

15 Select the **TCP** protocol.

16 Complete step 5 and you can leave the TCP data parameters empty.

If both the data sent and expected are not listed, then a three-way handshake TCP connection is established to validate the server health. No data is sent.

Expected data if listed has to be a string. Regular expressions are not supported.

17 Select the **UDP** protocol.

18 Complete step 5 and configure the UDP data.

Required Option	Description
UDP Data Sent	Enter the string to be sent to a server after a connection is established.
UDP Data Expected	Enter the string expected to receive from the server. Only when the received string matches this definition, is the server is considered as UP.

What to do next

Associate the active health monitor with a server pool. See [Add a Server Pool](#).

Add a Passive Monitor

Load balancers perform passive health checks to monitor failures during client connections and mark servers causing consistent failures as DOWN.

Passive health check monitors client traffic going through the load balancer for failures. For example, if a pool member sends a TCP Reset (RST) in response to a client connection, the load balancer detects that failure. If there are multiple consecutive failures, then the load balancer considers that server pool member to be temporarily unavailable and stops sending connection requests to that pool member for some time. After some time, the load balancer sends a connection request to verify that the pool member has recovered. If that connection is successful, then the pool member is considered healthy. Otherwise, the load balancer waits for some time and tries again.

Passive health check considers the following scenarios to be failures in the client traffic.

- For server pools associated with Layer 7 virtual servers, if the connection to the pool member fails. For example, if the pool member sends a TCP RST when the load balancer tries to connect or perform an SSL handshake between the load balancer and the pool member fails.
- For server pools associated with Layer 4 TCP virtual servers, if the pool member sends a TCP RST in response to client TCP SYN or does not respond at all.
- For server pools associated with Layer 4 UDP virtual servers, if a port is unreachable or a destination unreachable ICMP error message is received in response to a client UDP packet.

Server pools associated to Layer 7 virtual servers, the failed connection count is incremented when any TCP connection errors, for example, TCP RST failure to send data or SSL handshake failures occur.

Server pools associated to Layer 4 virtual servers, if no response is received to a TCP SYN sent to the server pool member or if a TCP RST is received in response to a TCP SYN, then the server pool member is considered as DOWN. The failed count is incremented.

For Layer 4 UDP virtual servers, if an ICMP error such as, port or destination unreachable message is received in response to the client traffic, then it is considered as DOWN.

Note One passive health monitor can be configured per server pool.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Load Balancing > Monitors > Passive > Add Passive Monitor**.
- 3 Enter a name and description for the passive health monitor.
- 4 Configure the values to monitor a service pool.

You can also accept the default active health monitor values.

Option	Description
Fail Count	Set a value when the consecutive failures reach this value, the server is considered temporarily unavailable.
Timeout Period	Set the number of times the server is tested before it is considered as DOWN.
Tags	Enter tags to make searching easier. You can specify a tag to set a scope of the tag.

For example, when the consecutive failures reach the configured value 5, that member is considered temporarily unavailable for 5 seconds. After this period, that member is tried again for a new connection to see if it is available. If that connection is successful, then the member is considered available and the failed count is set to zero. However, if that connection fails, then it is not used for another timeout interval of 5 seconds.

What to do next

Associate the passive health monitor with a server pool. See [Add a Server Pool](#).

Add a Server Pool

A server pool consists of one or more servers that are configured and running the same application. A single pool can be associated to both Layer 4 and Layer 7 virtual servers.





Figure 8-1. Server Pool Parameter Configuration



Prerequisites

- If you use dynamic pool members, a NSGroup must be configured. See [Create an NSGroup in Manager Mode](#).
- Verify that a passive health monitor is configured. See [Add a Passive Monitor](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Load Balancing > Server Pools > Add Server Pool**.

3 Enter a name and description for the load balancer server pool.

You can optionally describe the connections managed by the server pool.

4 Select the algorithm balancing method for the server pool.

Load balancing algorithm controls how the incoming connections are distributed among the members. The algorithm can be used on a server pool or a server directly.

All load balancing algorithms skip servers that meet any of the following conditions:

- Admin state is set to DISABLED.
- Admin state is set to GRACEFUL_DISABLED and no matching persistence entry.
- Active or passive health check state is DOWN.
- Connection limit for the maximum server pool concurrent connections is reached.

Option	Description
ROUND_ROBIN	Incoming client requests are cycled through a list of available servers capable of handling the request. Ignores the server pool member weights even if they are configured.
WEIGHTED_ROUND_ROBIN	Each server is assigned a weight value that signifies how that server performs relative to other servers in the pool. The value determines how many client requests are sent to a server compared to other servers in the pool. This load balancing algorithm focuses on fairly distributing the load among the available server resources.
LEAST_CONNECTION	Distributes client requests to multiple servers based on the number of connections already on the server. New connections are sent to the server with the fewest connections. Ignores the server pool member weights even if they are configured.
WEIGHTED_LEAST_CONNECTION	Each server is assigned a weight value that signifies how that server performs relative to other servers in the pool. The value determines how many client requests are sent to a server compared to other servers in the pool. This load balancing algorithm focuses on using the weight value to distribute the load among the available server resources. By default, the weight value is 1 if the value is not configured and slow start is enabled.
IP-HASH	Selects a server based on a hash of the source IP address and the total weight of all the running servers.

5 Click **Select Members** and elect the server pool members.

A server pool consists of single or multiple pool members.

Option	Description
Enter individual members	Enter a pool member name, IPv4 or IPv6 address, and a port. IP addresses can be either IPv4 or IPv6. Mixed addressing is not supported. Note that the pool members IP version must match the VIP IP version. For example, VIP-IPv4 with Pool-IPv4, and IPv6 with Pool-IPv6. Each server pool member can be configured with a weight for use in the load balancing algorithm. The weight indicates how much more or less load a given pool member can handle relative to other members in the same pool. You can set the server pool admin state. By default, the option is enable when a server pool member is added. If the option is disabled, active connections are processed, and the server pool member is not selected for new connections. New connections are assigned to other members of the pool. If gracefully disabled, it allows you to remove servers for maintenance. The existing connections to a member in the server pool in this state continue to be processed. Toggle the button to designate a pool member as a backup member to work with the health monitor to provide an Active-Standby state. Traffic failover occurs for backup members if active members fail a health check. Backup members are skipped during the server selection. When the server pool is inactive, the incoming connections are sent to only the backup members that are configured with a sorry page indicating an application is unavailable. Max Concurrent Connection value assigns a connection maximum so that the server pool members are not overloaded and skipped during server selection. If a value is not specified, then the connection is unlimited.
Select a group	Select a pre-configured group of server pool members. Enter a group name and an optional description. Set the compute member from existing list or create one. You can specify membership criteria, select members of the group, add IP addresses, and MAC addresses as group members, and add Active Directory groups. IP addresses can be either IPv4 or IPv6. Mixed addressing is not supported. The identity members intersect with the compute member to define membership of the group. Select a tag from the drop-down menu. You can optionally define the maximum group IP address list.

6 Click **Set Monitors** and select one or more active health check monitors for the server. Click **Apply**.

The load balancer periodically sends an ICMP ping to the servers to verify health independent of data traffic. You can configure more than one active health check monitor per server pool.

7 Select the Source NAT (SNAT) translation mode.

Depending on the topology, SNAT might be required so that the load balancer receives the traffic from the server destined to the client. SNAT can be enabled per server pool.

Mode	Description
Auto Map Mode	Load Balancer uses the interface IP address and ephemeral port to continue the communication with a client initially connected to one of the server's established listening ports. SNAT is required. Enable port overloading to allow the same SNAT IP and port to be used for multiple connections if the tuple (source IP, source port, destination IP, destination port, and IP protocol) is unique after the SNAT process is performed. You can also set the port overload factor to allow the maximum number of times a port can be used simultaneously for multiple connections.
Disable	Disable the SNAT translation mode.
IP Pool	Specify a single IPv4 or IPv6 address range, for example, 1.1.1.1-1.1.1.10 to be used for SNAT while connecting to any of the servers in the pool. IP addresses can be either IPv4 or IPv6. Mixed addressing is not supported. By default, from 4000 through 64000-port range is used for all configured SNAT IP addresses. Port ranges from 1000 through 4000 are reserved for purposes such as, health checks and connections initiated from Linux applications. If multiple IP addresses are present, then they are selected in a Round Robin manner. Enable port overloading to allow the same SNAT IP and port to be used for multiple connections if the tuple (source IP, source port, destination IP, destination port, and IP protocol) is unique after the SNAT process is performed. You can also set the port overload factor to allow the maximum number of times a port can be used simultaneously for multiple connections.

8 Click **Additional Properties**, and toggle the button to enable TCP Multiplexing.

With TCP multiplexing, you can use the same TCP connection between a load balancer and the server for sending multiple client requests from different client TCP connections.

9 Set the maximum number of TCP multiplexing connections per server that are kept alive to send future client requests.

10 Enter the minimum number of active members the server pool must always maintain.

11 Select a passive health monitor for the server pool from the drop-down menu.

12 Select a tag from the drop-down menu.

Setting Up Virtual Server Components

You can set up the Layer 4 and Layer 7 virtual servers and configure several virtual server components such as, application profiles, persistent profiles, and load balancer rules.





Figure 8-2. Virtual Server Components



Add an Application Profile

Application profiles are associated with virtual servers to enhance load balancing network traffic and simplify traffic-management tasks.

Application profiles define the behavior of a particular type of network traffic. The associated virtual server processes network traffic according to the values specified in the application profile. Fast TCP, Fast UDP, and HTTP application profiles are the supported types of profiles.

TCP application profile is used by default when no application profile is associated to a virtual server. TCP and UDP application profiles are used when an application is running on a TCP or UDP protocol and does not require any application level load balancing such as, HTTP URL load balancing. These profiles are also used when you only want Layer 4 load balancing, which has a faster performance and supports connection mirroring.

HTTP application profile is used for both HTTP and HTTPS applications when the load balancer must take actions based on Layer 7 such as, load balancing all images requests to a specific server pool member or stopping HTTPS to offload SSL from pool members. Unlike the TCP application profile, the HTTP application profile stops the client TCP connection before selecting the server pool member.

Figure 8-3. Layer 4 TCP and UDP Application Profile





Figure 8-4. Layer 7 HTTPS Application Profile



Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Load Balancing > Profiles > Application > Add Application Profiles**.

3 Select a **Fast TCP** application profile and enter the profile details.

You can also accept the default FAST TCP profile settings.

Option	Description
Name and Description	Enter a name and a description for the Fast TCP application profile.
Idle Timeout	Enter the time in seconds on how long the server can remain idle after a TCP connection is established. Set the idle time to the actual application idle time and add a few more seconds so that the load balancer does not close its connections before the application does.
HA Flow Mirroring	Toggle the button to make all the flows to the associated virtual server mirrored to the HA standby node.
Connection Close Timeout	Enter the time in seconds that the TCP connection both FINs or RST must be kept for an application before closing the connection. A short closing timeout might be required to support fast connection rates.
Tags	Enter tags to make searching easier. You can specify a tag to set a scope of the tag.

4 Select a **Fast UDP** application profile and enter the profile details.

You can also accept the default UDP profile settings.

Option	Description
Name and Description	Enter a name and a description for the Fast UDP application profile.
Idle Timeout	Enter the time in seconds on how long the server can remain idle after a UDP connection is established. UDP is a connectionless protocol. For load balancing purposes, all the UDP packets with the same flow signature such as, source and destination IP address or ports and IP protocol received within the idle timeout period are considered to belong to the same connection and sent to the same server. If no packets are received during the idle timeout period, the connection which is an association between the flow signature and the selected server is closed.
HA Flow Mirroring	Toggle the button to make all the flows to the associated virtual server mirrored to the HA standby node.
Tags	Enter tags to make searching easier. You can specify a tag to set a scope of the tag.

5 Select a **HTTP** application profile and enter the profile details.

You can also accept the default HTTP profile settings.

HTTP application profile is used for both HTTP and HTTPS applications.

Option	Description
Name and Description	Enter a name and a description for the HTTP application profile.
Idle Timeout	Enter the time in seconds on how long an HTTP application can remain idle, instead of the TCP socket setting which must be configured in the TCP application profile.
Request Header Size	Specify the maximum buffer size in bytes used to store HTTP request headers.
Response Header Size	Specify the maximum buffer size in bytes used to store HTTP response headers. The default is 4096, and the maximum is 65536.
X-Forwarded-For (XFF)	■ Insert - If the XFF HTTP header is not present in the incoming request, the load balancer inserts a new XFF header with the client IP address. If the XFF HTTP header is present in the incoming request, the load balancer appends the XFF header with the client IP address. ■ Replace - If the XFF HTTP header is present in the incoming request, the load balancer replaces the header. Web servers log each request they handle with the requesting client IP address. These logs are used for debugging and analytics purposes. If the deployment topology requires SNAT on the load balancer, then server uses the client SNAT IP address which defeats the purpose of logging. As a workaround, the load balancer can be configured to insert XFF HTTP header with the original client IP address. Servers can be configured to log the IP address in the XFF header instead of the source IP address of the connection.
Request Body Size	Enter value for the maximum size of the buffer used to store the HTTP request body. If the size is not specified, then the request body size is unlimited.

Option	Description
Redirection	■ None - If a website is temporarily down, user receives a page not found error message. ■ HTTP Redirect - If a website is temporarily down or has moved, incoming requests for that virtual server can be temporarily redirected to a URL specified here. Only a static redirection is supported. For example, if HTTP Redirect is set to <code>http://sitedown.abc.com/sorry.html</code>, then irrespective of the actual request, for example, <code>http://original_app.site.com/home.html</code> or <code>http://original_app.site.com/somepage.html</code>, incoming requests are redirected to the specified URL when the original website is down. ■ HTTP to HTTPS Redirect - Certain secure applications might want to force communication over SSL, but instead of rejecting non-SSL connections, they can redirect the client request to use SSL. With HTTP to HTTPS Redirect, you can preserve both the host and URI paths and redirect the client request to use SSL. For HTTP to HTTPS redirect, the HTTPS virtual server must have port 443 and the same virtual server IP address must be configured on the same load balancer. For example, a client request for <code>http://app.com/path/page.html</code> is redirected to <code>https://app.com/path/page.html</code>. If either the host name or the URI must be modified while redirecting, for example, redirect to <code>https://secure.app.com/path/page.html</code>, then load balancing rules must be used.
NTLM Authentication	Toggle the button for the load balancer to turn off TCP multiplexing and enable HTTP keep-alive. NTLM is an authentication protocol that can be used over HTTP. For load balancing with NTLM authentication, TCP multiplexing must be disabled for the server pools hosting NTLM-based applications. Otherwise, a server-side connection established with one client's credentials can potentially be used for serving another client's requests. If NTLM is enabled in the profile and associated to a virtual server, and TCP multiplexing is enabled at the server pool, then NTLM takes precedence. TCP multiplexing is not performed for that virtual server. However, if the same pool is associated to another non-NTLM virtual server, then TCP multiplexing is available for connections to that virtual server. If the client uses HTTP/1.0, the load balancer upgrades to HTTP/1.1 protocol and the HTTP keep-alive is set. All HTTP requests received on the same client-side TCP connection are sent to the same server over a single TCP connection to ensure that reauthorization is not required.
Tags	Enter tags to make searching easier. You can specify a tag to set a scope of the tag.

Add a Persistence Profile


To ensure stability of stateful applications, load balancers implement persistence which directs all related connections to the same server. Different types of persistence are supported to address different types of application needs.

Some applications maintain the server state such as, shopping carts. Such state might be per client and identified by the client IP address or per HTTP session. Applications might access or modify this state while processing subsequent related connections from the same client or HTTP session.

The source IP persistence profile tracks sessions based on the source IP address. When a client requests a connection to a virtual server that enables the source address persistence, the load balancer checks if that client was previously connected, if so, returns the client to the same server. If not, you can select a server pool member based on the pool load balancing algorithm. Source IP persistence profile is used by Layer 4 and Layer 7 virtual servers.

The cookie persistence profile inserts a unique cookie to identify the session the first time a client accesses the site. The client forwards the HTTP cookie in subsequent requests and the load balancer uses that information to provide the cookie persistence. Layer 7 virtual servers can only use the cookie persistence profile. Note that a blank space in a cookie name is **not** supported.

The generic persistence profile supports persistence based on the HTTP header, cookie, or URL in the HTTP request. Therefore, it supports app session persistence when the session ID is part of the URL. This profile is not associated with a virtual server directly. You can specify this profile when you configure a load balancer rule for request forwarding and response rewrite.



Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Load Balancing > Profiles > Persistence > Add Persistence Profiles**.

3 Select **Source IP** to add a source IP persistence profile and enter the profile details.

You can also accept the default Source IP profile settings.

Option	Description
Name and Description	Enter a name and a description for the Source IP persistence profile.
Share Persistence	Toggle the button to share the persistence so that all virtual servers this profile is associated with can share the persistence table. If the persistence sharing is not enabled in the Source IP persistence profile associated to a virtual server, each virtual server that the profile is associated to maintains a private persistence table.
Persistence Entry Timeout	Enter the persistence expiration time in seconds. The load balancer persistence table maintains entries to record that client requests are directed to the same server. The very first connection from new client IP is load balanced to a pool member based on the load balancing algorithm. NSX will store that persistence entry on the LB persistence-table which is viewable on the Edge Node hosting the T1-LB active via the CLI command:<code>get load-balancer <LB-UUID> persistence-tables</code>. ■ When there are connections from that client to the VIP, the persistence entry is kept. ■ When there are no more connections from that client to the VIP, the persistence entry begins the timer count down specified in the "Persistence Entry Timeout" value. If no new connection from that client to the VIP is made before the timer expires, the persistence entry for that client IP is deleted. If that client comes back after the entry is deleted, it will be load balanced again to a pool member based on the load balancing algorithm.
Purge Entries When Full	A large timeout value might lead to the persistence table quickly filling up when the traffic is heavy. When this option is enabled, the oldest entry is deleted to accept the newest entry. When this option is disabled, if the source IP persistence table is full, new client connections are rejected.
HA Persistence Mirroring	Toggle the button to synchronize persistence entries to the HA peer. When HA persistence mirroring is enabled, the client IP persistence remains in the case of load balancer failover.
Tags	Enter tags to make searching easier. You can specify a tag to set a scope of the tag.

4 Select a **Cookie** persistence profile and enter the profile details.

Option	Description
Name and Description	Enter a name and a description for the Cookie persistence profile.
Share Persistence	Toggle the button to share persistence across multiple virtual servers that are associated to the same pool members. The Cookie persistence profile inserts a cookie with the format, <code><name>.<profile-id>.<pool-id></code>. If the persistence shared is not enabled in the Cookie persistence profile associated with a virtual server, the private Cookie persistence for each virtual server is used and is qualified by the pool member. The load balancer inserts a cookie with the format, <code><name>.<virtual_server_id>.<pool_id></code>.
Cookie Mode	Select a mode from the drop-down menu. ■ INSERT - Adds a unique cookie to identify the session. ■ PREFIX - Appends to the existing HTTP cookie information. ■ REWRITE - Rewrites the existing HTTP cookie information.
Cookie Name	Enter the cookie name. A blank space in a cookie name is not supported.
Cookie Domain	Enter the domain name. HTTP cookie domain can be configured only in the INSERT mode.
Cookie Fallback	Toggle the button so that the client request is rejected if cookie points to a server that is in a DISABLED or is in a DOWN state. Selects a new server to handle a client request if the cookie points to a server that is in a DISABLED or is in a DOWN state.
Cookie Path	Enter the cookie URL path. HTTP cookie path can be set only in the INSERT mode.
Cookie Garbling	Toggle the button to disable encryption. When garbling is disabled, the cookie server IP address and port information is in a plain text. Encrypt the cookie server IP address and port information.
Cookie Type	Select a cookie type from the drop-down menu. Session Cookie - Not stored. Will be lost when the browser is closed. Persistence Cookie - Stored by the browser. Not lost when the browser is closed.
Max Idle Time	Enter the time in seconds that the cookie type can be idle before a cookie expires.
Max Cookie Age	For the session cookie type, enter the time in seconds a cookie is available.
Tags	Enter tags to make searching easier. You can specify a tag to set a scope of the tag.

5 Select **Generic** to add a generic persistence profile and enter the profile details.

Option	Description
Name and Description	Enter a name and a description for the Source IP persistence profile.
Share Persistence	Toggle the button to share the profile among virtual servers.

Option	Description
Persistence Entry Timeout	Enter the persistence expiration time in seconds. The load balancer persistence table maintains entries to record that client requests are directed to the same server. The very first connection from new client IP is load balanced to a pool member based on the load balancing algorithm. NSX will store that persistence entry on the LB persistence-table which is viewable on the Edge Node hosting the T1-LB active via the CLI command:<code>get load-balancer <LB-UUID> persistence-tables</code>. ■ When there are connections from that client to the VIP, the persistence entry is kept. ■ When there are no more connections from that client to the VIP, the persistence entry begins the timer count down specified in the "Persistence Entry Timeout" value. If no new connection from that client to the VIP is made before the timer expires, the persistence entry for that client IP is deleted. If that client comes back after the entry is deleted, it will be load balanced again to a pool member based on the load balancing algorithm.
HA Persistence Mirroring	Toggle the button to synchronize persistence entries to the HA peer.
Tags	Enter tags to make searching easier. You can specify a tag to set a scope of the tag.

Add an SSL Profile

SSL profiles configure application-independent SSL properties such as, cipher lists and reuse these lists across multiple applications. SSL properties are different when the load balancer is acting as a client and as a server, as a result separate SSL profiles for client-side and server-side are supported.

Note SSL profile is not supported in the NSX-T Data Center limited export release.

Client-side SSL profile refers to the load balancer acting as an SSL server and stopping the client SSL connection. Server-side SSL profile refers to the load balancer acting as a client and establishing a connection to the server.

You can specify a cipher list on both the client-side and server-side SSL profiles.

SSL session caching allows the SSL client and server to reuse previously negotiated security parameters avoiding the expensive public key operation during the SSL handshake. SSL session caching is disabled by default on both the client-side and server-side.

SSL session tickets are an alternate mechanism that allows the SSL client and server to reuse previously negotiated session parameters. In SSL session tickets, the client and server negotiate whether they support SSL session tickets during the handshake exchange. If supported by both, server can send an SSL ticket, which includes encrypted SSL session parameters to the client. The client can use that ticket in subsequent connections to reuse the session. SSL session tickets are enabled on the client-side and disabled on the server-side.

Figure 8-5. SSL Offloading





Figure 8-6. End-to-End SSL



Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Load Balancing > Profiles > SSL Profile**.
- 3 Select a **Client SSL Profile** and enter the profile details.

Option	Description
Name and Description	Enter a name and a description for the Client SSL profile.
SSL Suite	Select the SSL Cipher group from the drop-down menu and available SSL Ciphers and SSL protocols to be included in the Client SSL profile are populated. Balanced SSL Cipher group is the default.
Session Caching	Toggle the button to allow the SSL client and server to reuse previously negotiated security parameters avoiding the expensive public key operation during an SSL handshake.

Option	Description
Tags	Enter tags to make searching easier. You can specify a tag to set a scope of the tag.
Supported SSL Ciphers	Depending on the SSL suite, you assigned the supported SSL Ciphers are populated here. Click View More to view the entire list. If you selected Custom , you must select the SSL Ciphers from the drop-down menu.
Supported SSL Protocols	Depending on the SSL suite, you assigned the supported SSL protocols are populated here. Click View More to view the entire list. If you selected Custom , you must select the SSL Ciphers from the drop-down menu.
Session Cache Entry Timeout	Enter the cache timeout in seconds to specify how long the SSL session parameters must be kept and can be reused.
Prefer Server Cipher	Toggle the button so that the server can select the first supported cipher from the list it can support. During an SSL handshake, the client sends an ordered list of supported ciphers to the server.

4 Select a **Server SSL Profile** and enter the profile details.

Option	Description
Name and Description	Enter a name and a description for the Server SSL profile.
SSL Suite	Select the SSL Cipher group from the drop-down menu and available SSL Ciphers and SSL protocols to be included in the Server SSL profile are populated. Balanced SSL Cipher group is the default.
Session Caching	Toggle the button to allow the SSL client and server to reuse previously negotiated security parameters avoiding the expensive public key operation during an SSL handshake.
Tags	Enter tags to make searching easier. You can specify a tag to set a scope of the tag.
Supported SSL Ciphers	Depending on the SSL suite, you assigned the supported SSL Ciphers are populated here. Click View More to view the entire list. If you selected Custom , you must select the SSL Ciphers from the drop-down menu.
Supported SSL Protocols	Depending on the SSL suite, you assigned the supported SSL protocols are populated here. Click View More to view the entire list. If you selected Custom , you must select the SSL Ciphers from the drop-down menu.
Session Cache Entry Timeout	Enter the cache timeout in seconds to specify how long the SSL session parameters must be kept and can be reused.
Prefer Server Cipher	Toggle the button so that the server can select the first supported cipher from the list it can support. During an SSL handshake, the client sends an ordered list of supported ciphers to the server.

Add Layer 4 Virtual Servers

Virtual servers receive all the client connections and distribute them among the servers. A virtual server has an IP address, a port, and a protocol. For Layer 4 virtual servers, lists of ports ranges can be specified instead of a single TCP or UDP port to support complex protocols with dynamic ports.

A Layer 4 virtual server must be associated to a primary server pool, also called a default pool.

If a virtual server status is disabled, any new connection attempts to the virtual server are rejected by sending either a TCP RST for the TCP connection or ICMP error message for UDP. New connections are rejected even if there are matching persistence entries for them. Active connections continue to be processed. If a virtual server is deleted or disassociated from a load balancer, then active connections to that virtual server fail.

Prerequisites

- Verify that application profiles are available. See [Add an Application Profile](#).
- Verify that persistent profiles are available. See [Add a Persistence Profile](#).
- Verify that SSL profiles for the client and server are available. See [Add an SSL Profile](#).
- Verify that server pools are available. See [Add a Server Pool](#).
- Verify that load balancer is available. See [Add Load Balancers](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Load Balancing > Virtual Servers > Add Virtual Server**.
- 3 Select a **L4 TCP** or a **L4 UDP** protocol and enter the protocol details.

Layer 4 virtual servers support either the Fast TCP or Fast UDP protocol, but not both.

For Fast TCP or Fast UDP protocol support on the same IP address and port, for example DNS, a virtual server must be created for each protocol.

L4 TCP Option	L4 TCP Description
Name and Description	Enter a name and a description for the Layer 4 virtual server.
IP Address	Enter the virtual server IP address. Both IPv4 and IPv6 addresses are supported. Note that the pool members IP version must match the VIP IP version. For example, VIP-IPv4 with Pool-IPv4, and IPv6 with Pool-IPv6.
Ports	Enter the virtual server port number.
Load Balancer	Select an existing load balancer to attach to this Layer 4 virtual server from the drop-down menu.

L4 TCP Option	L4 TCP Description
Server Pool	Select an existing server pool from the drop-down menu. The server pool consists of one or more servers, also called pool members that are similarly configured and running the same application. You can click the vertical ellipses to create a server pool.
Application Profile	Based on the protocol type, the existing application profile is automatically populated. Click the vertical ellipses to create an application profile.
Persistence	Select an existing persistence profile from the drop-down menu. Persistence profile can be enabled on a virtual server to allow Source IP related client connections to be sent to the same server.
Access List Control	When you enable Access List Control (ALC), all traffic flowing through the load balancer is compared with the ACL statement, which either drops or allows the traffic. ACL is disabled by default. To enable, click Configure, and select Enabled. Select an Action: ■ Allow - Allows connections matching the selected group. All other connections are dropped. ■ Drop - Allows connections not matching the selected group. A dropped connection generates a log entry if access log is enabled. Select a Group. The IP addresses included in this group are either dropped or allowed by the ACL.
Max Concurrent Connection	Set the maximum concurrent connection allowed to a virtual server so that the virtual server does not deplete resources of other applications hosted on the same load balancer.
Max New Connection Rate	Set the maximum new connection to a server pool member so that a virtual server does not deplete resources.
Sorry Server Pool	Select an existing sorry server pool from the drop-down menu. The sorry server pool serves the request when a load balancer cannot select a backend server to serve the request from the default pool. You can click the vertical ellipses to create a server pool.
Default Pool Member Port	Enter a default pool member port if the pool member port for a virtual server is not defined. For example, if a virtual server is defined with a port range of 2000–2999 and the default pool member port range is set as 8000–8999, then an incoming client connection to the virtual server port 2500 is sent to a pool member with a destination port set to 8500.
Admin State	Toggle the button to disable the admin state of the Layer 4 virtual server.

L4 TCP Option	L4 TCP Description
Access Log	Toggle the button to enable logging for the Layer 4 virtual server.
Tags	Enter tags to make searching easier. You can specify a tag to set a scope of the tag.
L4 UDP Option	L4 UDP Description
Name and Description	Enter a name and a description for the Layer 4 virtual server.
IP Address	Enter the virtual server IP address. Both IPv4 and IPv6 addresses are supported. Note that the pool members IP version must match the VIP IP version. For example, VIP-IPv4 with Pool-IPv4, and IPv6 with Pool-IPv6.
Ports	Enter the virtual server port number.
Load Balancer	Select an existing load balancer to attach to this Layer 4 virtual server from the drop-down menu.
Server Pool	Select an existing server pool from the drop-down menu. The server pool consists of one or more servers, also called pool members that are similarly configured and running the same application. You can click the vertical ellipses to create a server pool.
Application Profile	Based on the protocol type, the existing application profile is automatically populated. You can click the vertical ellipses to create an application profile.
Persistence	Select an existing persistence profile from the drop-down menu. Persistence profile can be enabled on a virtual server to allow Source IP related client connections to be sent to the same server.
Max Concurrent Connection	Set the maximum concurrent connection allowed to a virtual server so that the virtual server does not deplete resources of other applications hosted on the same load balancer.
Access List Control	When you enable Access List Control (ALC) all traffic flowing through the load balancer will be compared with the ACL statement, which will either drop it or allow it. ACL is disabled by default. To enable, click Configure , and check Enabled . Select an Action: ■ Allow - Allows connections matching the selected group. All other connections are dropped ■ Drop - Allows connections not matching the selected group. A dropped connection generates a log entry if access log is enabled. Select a Group . The IP addresses included in this group are either dropped or allowed by the ACL.
Max New Connection Rate	Set the maximum new connection to a server pool member so that a virtual server does not deplete resources.
Sorry Server Pool	Select an existing sorry server pool from the drop-down menu. The sorry server pool serves the request when a load balancer cannot select a backend server to serve the request from the default pool. You can click the vertical ellipses to create a server pool.

L4 UDP Option	L4 UDP Description
Default Pool Member Port	Enter a default pool member port if the pool member port for a virtual server is not defined. For example, if a virtual server is defined with port range 2000–2999 and the default pool member port range is set as 8000–8999, then an incoming client connection to the virtual server port 2500 is sent to a pool member with a destination port set to 8500.
Admin State	Toggle the button to disable the admin state of the Layer 4 virtual server.
Access Log	Toggle the button to enable logging for the Layer 4 virtual server.
Log Significant Event Only	This field can only be configured if access logs are enabled. Connections that cannot be sent to a pool member are treated as a significant event such as "max connection limit," or "Access Control drop."
Tags	Enter tags to make searching easier. You can specify a tag to set a scope of the tag.

Add Layer 7 HTTP Virtual Servers

Virtual servers receive all the client connections and distribute them among the servers. A virtual server has an IP address, a port, and a protocol TCP.

If a virtual server status is disabled, any new connection attempts to the virtual server are rejected by sending either a TCP RST for the TCP connection or ICMP error message for UDP. New connections are rejected even if there are matching persistence entries for them. Active connections continue to be processed. If a virtual server is deleted or disassociated from a load balancer, then active connections to that virtual server fail.

Note SSL profile is not supported in the NSX-T Data Center limited export release.

If a client-side SSL profile binding is configured on a virtual server but not a server-side SSL profile binding, then the virtual server operates in an SSL-terminate mode, which has an encrypted connection to the client and plain text connection to the server. If both the client-side and server-side SSL profile bindings are configured, then the virtual server operates in SSL-proxy mode, which has an encrypted connection both to the client and the server.

Associating server-side SSL profile binding without associating a client-side SSL profile binding is currently not supported. If a client-side and a server-side SSL profile binding is not associated with a virtual server and the application is SSL-based, then the virtual server operates in an SSL-unaware mode. In this case, the virtual server must be configured for Layer 4. For example, the virtual server can be associated to a fast TCP profile.

Prerequisites

- Verify that application profiles are available. See [Add an Application Profile](#).
- Verify that persistent profiles are available. See [Add a Persistence Profile](#).
- Verify that SSL profiles for the client and server are available. See [Add an SSL Profile](#).
- Verify that server pools are available. See [Add a Server Pool](#).

- Verify that CA and client certificate are available. See [Create a Certificate Signing Request File](#).
- Verify that a certification revocation list (CRL) is available. See [Import a Certificate Revocation List](#).
- Verify that load balancer is available. See [Add Load Balancers](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Load Balancing > Virtual Servers > Add Virtual Server**.
- 3 Select **L7 HTTP** from the drop-down list and enter the protocol details.

Layer 7 virtual servers support the HTTP and HTTPS protocols.

Option	Description
Name and Description	Enter a name and a description for the Layer virtual server.
IP Address	Enter the virtual server IP address. Both IPv4 and IPv6 addresses are supported.
Ports	Enter the virtual server port number.
Load Balancer	Select an existing load balancer to attach to this Layer 4 virtual server from the drop down menu.
Server Pool	Select an existing server pool from the drop-down menu. The server pool consists of one or more servers, also called pool members that are similarly configured and running the same application. You can click the vertical ellipses to create a server pool.
Application Profile	Based on the protocol type, the existing application profile is automatically populated. You can click the vertical ellipses to create an application profile.
Persistence	Select an existing persistence profile from the drop-down menu. Persistence profile can be enabled on a virtual server to allow Source IP and Cookie related client connections to be sent to the same server.

- 4 Click **Configure** to set the Layer 7 virtual server SSL.

You can configure the Client SSL and Server SSL.

5 Configure the Client SSL.

Option	Description
Client SSL	Toggle the button to enable the profile. Client-side SSL profile binding allows multiple certificates, for different host names to be associated to the same virtual server.
Default Certificate	Select a default certificate from the drop-down menu. This certificate is used if the server does not host multiple host names on the same IP address or if the client does not support Server Name Indication (SNI) extension.
Client SSL Profile	Select the client-side SSL Profile from the drop-down menu.
SNI Certificates	Select the available SNI certificate from the drop-down menu.
Trusted CA Certificates	Select the available CA certificate.
Mandatory Client Authentication	Toggle the button to enable this menu item.
Certificate Chain Depth	Set the certificate chain depth to verify the depth in the server certificates chain.
Certificate Revocation List	Select the available CRL to disallow compromised server certificates.

6 Configure the Server SSL.

Option	Description
Server SSL	Toggle the button to enable the profile.
Client Certificate	Select a client certificate from the drop-down menu. This certificate is used if the server does not host multiple host names on the same IP address or if the client does not support Server Name Indication (SNI) extension.
Server SSL Profile	Select the Server-side SSL Profile from the drop-down menu.
Trusted CA Certificates	Select the available CA certificate.
Mandatory Server Authentication	Toggle the button to enable this menu item. Server-side SSL profile binding specifies whether the server certificate presented to the load balancer during the SSL handshake must be validated or not. When validation is enabled, the server certificate must be signed by one of the trusted CAs whose self-signed certificates are specified in the same server-side SSL profile binding.
Certificate Chain Depth	Set the certificate chain depth to verify the depth in the server certificates chain.
Certificate Revocation List	Select the available CRL to disallow compromised server certificates. OCSP and OCSP stapling are not supported on the server-side.

- 7 Click **Additional Properties** to configure additional Layer 7 virtual server properties.

Option	Description
Max Concurrent Connection	Set the maximum concurrent connection allowed to a virtual server so that the virtual server does not deplete resources of other applications hosted on the same load balancer.
Max New Connection Rate	Set the maximum new connection to a server pool member so that a virtual server does not deplete resources.
Sorry Server Pool	Select an existing sorry server pool from the drop-down menu. The sorry server pool serves the request when a load balancer cannot select a backend server to serve the request from the default pool. You can click the vertical ellipses to create a server pool.
Default Pool Member Port	Enter a default pool member port, if the pool member port for a virtual server is not defined. For example, if a virtual server is defined with port range 2000-2999 and the default pool member port range is set as 8000-8999, then an incoming client connection to the virtual server port 2500 is sent to a pool member with a destination port set to 8500.
Admin State	Toggle the button to disable the admin state of the Layer 7 virtual server.
Access Log	Toggle the button to enable logging for the Layer 7 virtual server.
Log Significant Event Only	This field can only be configured if access logs are enabled. Requests with an HTTP response status of ≥ 400 are treated as a significant event.
Tags	Select a tag from the drop-down list. You can specify a tag to set a scope of the tag.

- 8 Click **Save**.

Add Load Balancer Rules

With Layer 7 HTTP virtual servers, you can optionally configure load balancer rules and customize load balancing behavior using match or action rules.

Load balancer rules are supported for only Layer 7 virtual servers with an HTTP application profile. Different load balancer services can use load balancer rules.

Each load balancer rule consists of single or multiple match conditions and single or multiple actions. If the match conditions are not specified, then the load balancer rule always matches and is used to define default rules. If more than one match condition is specified, then the matching strategy determines if all conditions must match or any one condition must match for the load balancer rule to be considered a match.

Each load balancer rule is implemented at a specific phase of the load balancing processing; Transport, HTTP Access, Request Rewrite, Request Forwarding, and Response Rewrite. Not all the match conditions and actions are applicable to each phase.

Up to 4,000 load balancer rules can be configured with the API, if the `skip_scale_validation` flag in LbService is set. Note that the flag can be set via API. Refer to the *NSX-T Data Center API Guide* for more information. Up to 512 load balancer rules can be configured through the user interface.

Load Balancer rules support REGEX for match types. For more information, see [Regular Expressions in Load Balancer Rules](#).

Prerequisites

Verify a Layer 7 HTTP virtual server is available. See [Add Layer 7 HTTP Virtual Servers](#).

- [Configure Transport Phase Load Balancer Rules](#)

Transport phase is the first phase of a client HTTP request.

- [Configure HTTP Access Load Balancer Rules](#)

A JSON web token (JWT) is a standardized, optionally validated and/or encrypted format that is used to securely transfer information between two parties.

- [Configure Request Rewrite Load Balancer Rules](#)

An HTTP request rewrite is applied to the HTTP request coming from the client.

- [Configure Request Forwarding Load Balancer Rules](#)

Request forwarding redirects a URL or host to a specific server pool.

- [Configure Response Rewrite Load Balancer Rules](#)

An HTTP response rewrite is applied to the HTTP response going out from the servers to the client.

- [Regular Expressions in Load Balancer Rules](#)

Regular expressions (REGEX) are used in match conditions for load balancer rules.

Configure Transport Phase Load Balancer Rules

Transport phase is the first phase of a client HTTP request.

Load Balancer virtual server SSL configuration is found under **SSL Configuration**. There are two possible configurations. In both modes, the load balancer sees the traffic, and applies load balancer rules based on the client HTTP traffic.

- **SSL Offload**, configuring only the SSL client. In this mode, the client to VIP traffic is encrypted (HTTPS), and the load balancer decrypts it. The VIP to Pool member traffic is clear (HTTP).
- **SSL End-to-End**, configuring both the Client SSL and Server SSL. In this mode, the client to VIP traffic is encrypted (HTTPS), and the load balancer decrypts it and then re-encrypts it. The VIP to Pool member traffic is encrypted (HTTPS).

The Transport Phase is complete when the virtual server receives the client SSL hello message virtual server. this occurs before SSL is ended, and before HTTP traffic.

The Transport Phase allows administrators to select the SSL mode, and specific server pool based on the client SSL hello message. There are three options for the virtual server SSL mode:

- SSL Offload
- End-to-End
- SSL-Passthrough (the load balancer does not end SSL)

Load Balancer rules support REGEX for match types. PCRE style REGEX patterns are supported with a few limitations on advanced use cases. When REGEX is used in match conditions, named capturing groups are supported. See [Regular Expressions in Load Balancer Rules](#).

Prerequisites

Verify that a Layer 7 HTTP virtual server is available. See [Add Layer 7 HTTP Virtual Servers](#).

Procedure

- 1 Open the Layer 7 HTTP virtual server.
- 2 In the Load Balancer Rules section, next to Transport Phase, click **Set > Add Rule** to configure the load balancer rules for the Transport Phase.
- 3 SSL SNI is the only match condition supported. Match conditions are used to match application traffic passing through load balancers.
- 4 From the drop-down list, select a **Match Type**: starts with, ends with, equals, contains, matches regex.
- 5 Enter a **SNI Name**.
- 6 Toggle the **Case Sensitive** button to set a case-sensitive flag for HTTP header value comparison.
- 7 Toggle the **Negate** button to enable it.
- 8 From the drop-down list, select a **Match Strategy**:

Match Strategy	Description
Any	Either host or path may match for this rule to be considered a match.
All	Both host and path must match for this rule to be considered a match.

9 From the drop-down menu, select the **SSL Mode Selection**.

SSL Mode	Description
SSL Passthrough	SSL Passthrough passes HTTP traffic to a backend server without decrypting the traffic on the load balancer. The data is kept encrypted as it travels through the load balancer. If SSL Passthrough is selected, a server pool can be selected. See Add a Server Pool for Load Balancing in Manager Mode .
SSL Offloading	SSL Offloading decrypts all HTTP traffic on the load balancer. SSL offloading allows data to be inspected as it passes between the load balancer and server. If NTLM and multiplexing are not configured, the load balancer establishes a new connection to the selected backend server for each HTTP request.
SSL End-to End	After receiving the HTTP request, the load balancer connects to the selected backend server and talks with it using HTTPS. If NTLM and multiplexing are not configured, the load balancer establishes a new connection to the selected backend server for each HTTP request.

10 Click **SAVE** and **APPLY**.

Configure HTTP Access Load Balancer Rules

A JSON web token (JWT) is a standardized, optionally validated and/or encrypted format that is used to securely transfer information between two parties.

In the HTTP ACCESS phase, users can define the action to validate JWT from clients and pass, or remove JWT to backend servers.

Load Balancer rules support REGEX for match types. PCRE style REGEX patterns is supported with a few limitations on advanced use cases. When REGEX is used in match conditions, named capturing groups are supported. See [Regular Expressions in Load Balancer Rules](#).

Prerequisites

Verify that a Layer 7 HTTP virtual server is available. See [Add Layer 7 HTTP Virtual Servers](#).

Procedure

- 1 Open the Layer 7 HTTP virtual server.
- 2 In the Load Balancer Rules section, next to HTTP Access Phase, click **Set > Add Rule** to configure the load balancer rules for the HTTP Request Rewrite phase.

- From the drop-down menu, select a match condition. Match conditions are used to match application traffic passing through load balancers. Multiple match conditions can be specified in one load balancer rule. Each match condition defines a criterion for application traffic.

Supported Match Condition	Description
HTTP Request Method	Match an HTTP request method. http_request.method - value to match
HTTP Request URI	Match an HTTP request URI without query arguments. http_request.uri - value to match
HTTP Request URI Arguments	Match an HTTP request URI query argument. http_request.uri_arguments - value to match
HTTP Request Version	Match an HTTP request version. http_request.version - value to match
HTTP Request Header	Match any HTTP request header. http_request.header_name - header name to match http_request.header_value - value to match
HTTP Request Cookie	Match any HTTP request cookie. http_request.cookie_value - value to match
HTTP Request Body	Match an HTTP request body content. http_request.body_value - value to match
TCP Header Port	Match a TCP source or the destination port. tcp_header.source_port - source port to match tcp_header.destination_port - destination port to match
IP Header Source	Matches IP header text boxes in of HTTP messages. The source type must be either a single IP address, a range of IP addresses, or a group. See Add a Group . ■ If IP Header Source is selected, with an IP Address source type, the source IP address of HTTP messages should match IP addresses which are configured in groups. Both IPv4 and IPv6 addresses are supported. ■ If IP Header Source is selected with a Group source type, select the group from the drop-down menu. ip_header.source_address - source address to match ip_header.destination_address - destination address to match
Variable	Create a variable and assign a value to the variable.
Client SSL	Match client SSL profile ID. ssl_profile_id - value to match
Case Sensitive	Set a case-sensitive flag for HTTP header value comparison. If true, case is significant when comparing HTTP body value.

- From the drop-down list, select a **Match Type**: starts with, ends with, equals, contains, matches regex.
- If needed, enter the URI.

6 From the drop-down list, select a **Match Strategy**:

Match Strategy	Description
Any	Either host or path may match for this rule to be considered a match.
All	Both host and path must match for this rule to be considered a match.

7 From the drop-down menu select an **Action**:

Action	Description
JWT Authentication	JSON Web Token (JWT) is an open standard that defines a compact and self-contained way for securely transmitting information between parties as a JSON object. This information can be verified and trusted because it is digitally signed. ■ Realm - A description of the protected area. If no realm is specified, clients often display a formatted hostname. The configured realm is returned when a client request is rejected with 401 http status. The response is: "WWW-Authentication: Bearer realm=<realm>". ■ Tokens - This parameter is optional. Load balancer searches for every specified token one-by-one for the JWT message until found. If not found, or if this text box is not configured, load balancer searches the Bearer header by default in the http request "Authorization: Bearer <token>". ■ Key Type - Symmetric key or asymmetric public key (certificate-id) ■ Preserve JWT - This is a flag to preserve JWT and pass it to backend server. If disabled, the JWT key to the backend server is removed.
Connection Drop	If negate is enabled, when Connection Drop is configured, all requests not matching the specified match condition are dropped. Requests matching the specified match condition are allowed.
Variable Assignment	Enables users to assign a value to a variable in HTTP Access Phase, in such a way that the result can be used as a condition in other load balancer rule phases.

8 Click **Save** and **Apply**.

Configure Request Rewrite Load Balancer Rules

An HTTP request rewrite is applied to the HTTP request coming from the client.

Prerequisites

Verify that a Layer 7 HTTP virtual server is available. See [Add Layer 7 HTTP Virtual Servers](#).

Load Balancer rules support REGEX for match types. PCRE style REGEX patterns is supported with a few limitations on advanced use cases. When REGEX is used in match conditions, named capturing groups are supported. See [Regular Expressions in Load Balancer Rules](#).

Procedure

1 Open the Layer 7 HTTP virtual server.

- 2 In the Load Balancer Rules section, next to Request Rewrite Phase, click **Set > Add Rule** to configure the load balancer rules for the HTTP Request Rewrite phase.
- 3 From the drop-down list, select a match condition. Match conditions are used to match application traffic passing through load balancers. Multiple match conditions can be specified in one load balancer rule. Each match condition defines a criterion for application traffic.

Supported Match Condition	Description
HTTP Request Method	Match an HTTP request method. http_request.method - value to match
HTTP Request URI	Match an HTTP request URI without query arguments. http_request.uri - value to match
HTTP Request URI Arguments	Used to match URI arguments aka query string of HTTP request messages, for example, in URI http://exaple.com?foo=1&bar=2, the "foo=1&bar=2" is the query string containing URI arguments. In an URI scheme, query string is indicated by the first question mark ("?") character and terminated by a number sign("#") character or by the end of the URI. http_request.uri_arguments - value to match
HTTP Request Version	Used to match the HTTP protocol version of the HTTP request messages http_request.version - value to match
HTTP Request Header	Used to match HTTP request messages by HTTP header fields. HTTP header fields are components of the header section of HTTP request and response messages. They define the operating parameters of an HTTP transaction. http_request.header_name - header name to match http_request.header_value - value to match
HTTP Request Cookie	Used to match HTTP request messages by cookie which is a specific type of HTTP header. The match_type and case_sensitive define how to compare cookie value. http_request.cookie_value - value to match
HTTP Request Body	Match an HTTP request body content. http_request.body_value - value to match
Client SSL	Match client SSL profile ID. ssl_profile_id - value to match
TCP Header Port	Match a TCP source or the destination port. tcp_header.source_port - source port to match tcp_header.destination_port - destination port to match
IP Header Source	Matches IP header fields in of HTTP messages. The source type must be either a single IP address, or a range of IP addresses, or a group. See Add a Group . ■ If IP Header Source is selected, with an IP Address source type, the source IP address of HTTP messages should match IP addresses which are configured in groups. Both IPv4 and IPv6 addresss are supported ■ If IP Header Source is selected with a Group source type, seelct the group from the drop-down list. ip_header.source_address - source address to match ip_header.destination_address - destination address to match

Supported Match Condition	Description
Variable	Create a variable and assign a value to the variable.
Case Sensitive	Set a case-sensitive flag for HTTP header value comparison. If true, case is significant when comparing HTTP body value.

- 4 From the drop-down menu, select a **Match Type**: starts with, ends with, equals, contains, or matches regex. Match type is used to match a condition with a specified action.

Match Type	Description
Starts With	If the match condition starts with the specified value, the condition matches.
Ends With	If the match condition ends with the specified value, the condition matches.
Equals	If the match condition is the same as the specified value, the condition matches.
Contains	If the match condition contains the specified value, the condition matches.
Matches Regex	If the match condition matches the specified values, the condition matches.

- 5 Specify the URI.

- 6 From the drop-down menu, select a **Match Strategy**:

Match Strategy	Description
Any	Indicates that either host or path can match for this rule to be considered a match.
All	Indicates that both host and path must match for this rule to be considered a match.

- 7 Select an Action from the drop-down menu:

Actions	Description
HTTP Request URI Rewrite	This action is used to rewrite URIs in matched HTTP request messages. Specify the URI and URI Arguments in this condition to rewrite the matched HTTP request message's URI and URI arguments to the new values. Full URI scheme of HTTP messages have following syntax: Scheme: [//[user[:password]@]host[:port]][/path][?query][#fragment] The URI field of this action is used to rewrite the /path part in the above scheme. The URI Arguments field is used to rewrite the query part. Captured variables and built-in variables can be used in the URI and URI Arguments fields. - a Enter the URI of the HTTP request - b Enter the query string of URI, which typically contains key value pairs, for example: foo1=bar1&foo2=bar2.
HTTP Request Header Rewrite	This action is used to rewrite header fields of matched HTTP request messages to specified new values. - a Enter the name of a header text box HTTP request message. - b Enter the header value.

Actions	Description
HTTP Request Header Delete	This action is used to delete header fields of HTTP request messages at HTTP_REQUEST_REWRITE phase. One action can be used to delete all headers with same header name. To delete headers with different header names, multiple actions must be defined. ■ Enter the name of a header field of HTTP request message.
Variable Assignment	Create a variable and assign it a name and value.

- 8 Toggle the **Case Sensitive** button to set a case-sensitive flag for HTTP header value comparison.
- 9 Toggle the **Negate** button to enable it.
- 10 Click **Save** and **Apply**.

Configure Request Forwarding Load Balancer Rules

Request forwarding redirects a URL or host to a specific server pool.

Prerequisites

Verify that a Layer 7 HTTP virtual server is available. See [Add Layer 7 HTTP Virtual Servers](#).

Load Balancer rules support REGEX for match types. PCRE style REGEX patterns is supported with a few limitations on advanced use cases. When REGEX is used in match conditions, named capturing groups are supported. See [Regular Expressions in Load Balancer Rules](#).

Procedure

- 1 Open the Layer 7 HTTP virtual server.
- 2 Click **Request Forwarding > Add Rule** to configure the load balancer rules for the HTTP Request Forwarding.
- 3 From the drop-down list, select a match condition. Match conditions are used to match application traffic passing through load balancers. Multiple match conditions can be specified in one load balancer rule. Each match condition defines a criterion for application traffic.

Supported Match Condition	Description
HTTP Request Method	Match an HTTP request method. http_request.method - value to match
HTTP Request URI	Match an HTTP request URI without query arguments. http_request.uri - value to match
HTTP Request URI Arguments	Used to match URI arguments aka query string of HTTP request messages, for example, in URI http://example.com?foo=1&bar=2, the "foo=1&bar=2" is the query string containing URI arguments. In an URI scheme, query string is indicated by the first question mark ("?") character and terminated by a number sign ("#") character or by the end of the URI. http_request.uri_arguments - value to match
HTTP Request Version	Used to match the HTTP protocol version of the HTTP request messages http_request.version - value to match

Supported Match Condition	Description
HTTP Request Header	Used to match HTTP request messages by HTTP header fields. HTTP header fields are components of the header section of HTTP request and response messages. They define the operating parameters of an HTTP transaction. http_request.header_name - header name to match http_request.header_value - value to match
HTTP Request Cookie	Used to match HTTP request messages by cookie which is a specific type of HTTP header. The match_type and case_sensitive define how to compare cookie value. http_request.cookie_value - value to match
HTTP Request Body	Match an HTTP request body content. http_request.body_value - value to match
Client SSL	Match client SSL profile ID. ssl_profile_id - value to match
TCP Header Port	Match a TCP source or the destination port. tcp_header.source_port - source port to match tcp_header.destination_port - destination port to match
IP Header Source	Matches IP header fields in of HTTP messages. The source type must be either a single IP address, or a range of IP addresses, or a group. See Add a Group . ■ If IP Header Source is selected, with an IP Address source type, the source IP address of HTTP messages should match IP addresses which are configured in groups. Both IPv4 and IPv6 addresses are supported ■ If IP Header Source is selected with a Group source type, select the group from the drop-down list. ip_header.source_address - source address to match ip_header.destination_address - destination address to match
Variable	Create a variable and assign a value to the variable.
Case Sensitive	Set a case-sensitive flag for HTTP header value comparison. If true, case is significant when comparing HTTP body value.

4 Select an action:

Action	Description
HTTP Reject	Used to reject HTTP request messages. The specified reply_status value is used as the status code for the corresponding HTTP response message. The response message is sent back to client (usually a browser) indicating the reason it was rejected. http_forward.reply_status - HTTP status code used to reject http_forward.reply_message - HTTP rejection message
HTTP Redirect	Used to redirect HTTP request messages to a new URL. The HTTP status code for redirection is 3xx, for example, 301, 302, 303, 307, etc. The redirect_url is the new URL that the HTTP request message is redirected to. http_forward.redirect_status - HTTP status code for redirect http_forward.redirect_url - HTTP redirect URL

Action	Description
Select Pool	Force the request to a specific server pool. Specified pool member's configured algorithm (predictor) is used to select a server within the server pool. The matched HTTP request messages are forwarded to the specified pool. When HTTP keep-alive is enabled and forwarding rules are configured in the load balancer, the server keep-alive setting takes precedence. As a result, HTTP requests are sent to servers already connected with keep-alive. If you always want to give priority to the forwarding rules when the load balancer rule conditions are met, disable the keep-alive setting. Note that the persistence setting takes precedence over the keep-alive setting. Processing is done in the order of Persistence > Keep-Alive > Load Balancer Rules http_forward.select_pool - server pool UUID
Variable Persistence On	Select a generic persistence profile and enter a variable name. You can also enable Hash Variable. If the variable value is long, hashing the variable ensures that it is correctly stored in the persistence table. If the Hash Variable is not enabled, only the fixed prefix part of the variable value is stored in the persistence table if the variable value is long. As a result, two different requests with long variable values might be dispatched to the same backend server because their variable values have the same prefix part, when they should be dispatched to different backend servers.
Connection Drop	If negate is enabled in condition, when Connection Drop is configured, all requests not matching the condition are dropped. Requests matching the condition are allowed.
Reply Status	Shows the status of the reply.
Reply Message	Server responds with a reply message that contains confirmed addresses and configuration.

5 Click **Save** and **Apply**.

Configure Response Rewrite Load Balancer Rules

An HTTP response rewrite is applied to the HTTP response going out from the servers to the client.

Prerequisites

Verify that a Layer 7 HTTP virtual server is available. See [Add Layer 7 HTTP Virtual Servers](#).

Load Balancer rules support REGEX for match types. PCRE style REGEX patterns is supported with a few limitations on advanced use cases. When REGEX is used in match conditions, named capturing groups are supported. See [Regular Expressions in Load Balancer Rules](#).

Procedure

1 Open the Layer 7 HTTP virtual server.

- 2 Click **Response Rewrite > Add Rule** to configure the load balancer rules for the HTTP Response Rewrite.

All match values accept regular expressions.

Supported Match Condition	Description
HTTP Response Header	This condition is used to match HTTP response messages from backend servers by HTTP header fields. http_response.header_name - header name to match http_response.header_value - value to match
HTTP Response Method	Match an HTTP response method. http_response.method - value to match
HTTP Response URI	Match an HTTP response URI. http_response.uri - value to match
HTTP Response URI Arguments	Match an HTTP response URI arguments. http_response.uri_args - value to match
HTTP Response Version	Match an HTTP response version. http_response.version - value to match
HTTP Response Cookie	Match any HTTP response cookie. http_response.cookie_value - value to match
Client SSL	Match client SSL profile ID. ssl_profile_id - value to match
TCP Header Port	Match a TCP source or the destination port. tcp_header.source_port - source port to match tcp_header.destination_port - destination port to match
IP Header Source	Matches IP header fields in of HTTP messages. The source type must be either a single IP address, or a range of IP addresses, or a group. See Add a Group . The source IP address of HTTP messages should match IP addresses which are configured in groups. Both IPv4 and IPv6 addresses are supported. ip_header.source_address - source address to match ip_header.destination_address - destination address to match
Variable	Create a variable and assign a value to the variable.
Case Sensitive	Set a case-sensitive flag for HTTP header value comparison.

3 Select an action:

Action	Description
HTTP Response Header Rewrite	This action is used to rewrite header fields of HTTP response messages to specified new values. http_response.header_name - header name http_response.header_value - value to write
HTTP Response Header Delete	This action is used to delete header fields of HTTP response messages. http_request.header_delete - header name http_request.header_delete - value to write
Variable Persistence Learn	Select a generic persistence profile and enter a variable name. You can also enable Hash Variable . If the variable value is long, hashing the variable ensures that it will be correctly stored in the persistence table. If Hash Variable is not enabled, only the fixed prefix part of the variable value is stored in the persistence table if the variable value is long. As a result, two different requests with long variable values might be dispatched to the same backend server (because their variable values have the same prefix part) when they should be dispatched to different backend servers.

4 Click **Save** and **Apply**.

Regular Expressions in Load Balancer Rules

Regular expressions (REGEX) are used in match conditions for load balancer rules.

Perl Compatible Regular Expressions (PCRE) style REGEX patterns is supported with a few limitations on advanced use cases. When REGEX is used in match conditions, named capturing groups are supported.

REGEX restrictions include:

- Character unions and intersections are not supported. For example, do not use `[a-z[0-9]]` and `[a-z&&[aeiou]]` instead use `[a-z0-9]` and `[aeiou]` respectively.
- Only 9 back references are supported and `\1` through `\9` can be used to refer to them.
- Use `\Odd` format to match octal characters, not the `\ddd` format.
- Embedded flags are not supported at the top level, they are only supported within groups. For example, do not use `"Case (?i:s)ensitive"` instead use `"Case ((?i:s)ensitive)"`.
- Preprocessing operations `\l`, `\u`, `\L`, `\U` are not supported. Where `\l` - lowercase next char `\u` - uppercase next char `\L` - lower case until `\E` `\U` - upper case to `\E`.
- `(?(condition)X)`, `(?{code})`, `(?#{Code})` and `(?#comment)` are not supported.
- Predefined Unicode character class `\X` is not supported
- Using named character construct for Unicode characters is not supported. For example, do not use `\N{name}` instead use `\u2018`.

When REGEX is used in match conditions, named capturing groups are supported. For example, REGEX match pattern `/news/(?<year>\d+)-(?<month>\d+)-(?<day>\d+)/(?<article>.*)` can be used to match a URI like `/news/2018-06-15/news1234.html`.

Then variables are set as follows, \$year = "2018" \$month = "06" \$day = "15" \$article = "news1234.html". After the variables are set, these variables can be used in load balancer rule actions. For example, URI can be rewritten using the matched variables like, /news.py?year=\$year&month=\$month&day=\$day&article=\$article. Then the URI gets rewritten as /news.py?year=2018&month=06&day=15&article=news1234.html.

Rewrite actions can use a combination of named capturing groups and built-in variables. For example, URI can be written as /news.py?year=\$year&month=\$month&day=\$day&article=\$article&user_ip=\$_remote_addr. Then the example URI gets rewritten as /news.py?year=2018&month=06&day=15&article=news1234.html&user_ip=1.1.1.1.

Note For named capturing groups, the name cannot start with an _ character.

In addition to named capturing groups, the following built-in variables can be used in rewrite actions. All the built-in variable names start with _.

- \$_args - arguments from the request
- \$_arg_<name> - argument <name> in the request line
- \$_cookie_<name> - value of <name> cookie
- \$_upstream_cookie_<name> - cookie with the specified name sent by the upstream server in the "Set-Cookie" response header field
- \$_upstream_http_<name> - arbitrary response header field and <name> is the field name converted to lower case with dashes replaced by underscores
- \$_host - in the order of precedence - host name from the request line, or host name from the "Host" request header field, or the server name matching a request
- \$_http_<name> - arbitrary request header field and <name> is the field name converted to lower case with dashes replaced by underscores
- \$_https - "on" if connection operates in SSL mode, or "" otherwise
- \$_is_args - "?" if a request line has arguments, or "" otherwise
- \$_query_string - same as \$_args
- \$_remote_addr - client address
- \$_remote_port - client port
- \$_request_uri - full original request URI (with arguments)
- \$_scheme - request scheme, "http" or "https"
- \$_server_addr - address of the server which accepted a request
- \$_server_name - name of the server which accepted a request
- \$_server_port - port of the server which accepted a request
- \$_server_protocol - request protocol, usually "HTTP/1.0" or "HTTP/1.1"

- `$_ssl_client_escaped_cert` - returns the client certificate in the PEM format for an established SSL connection.
- `$_ssl_server_name` - returns the server name requested through SNI
- `$_uri` - URI path in request
- `$_ssl_ciphers`: returns the client SSL ciphers
- `$_ssl_client_i_dn`: returns the "issuer DN" string of the client certificate for an established SSL connection according to RFC 2253
- `$_ssl_client_s_dn`: returns the "subject DN" string of the client certificate for an established SSL connection according to RFC 2253
- `$_ssl_protocol`: returns the protocol of an established SSL connection
- `$_ssl_session_reused`: returns "r" if an SSL session was reused, or "." otherwise

Groups Created for Server Pools and Virtual Servers

NSX Manager automatically creates groups for load balancer server pools and VIP ports.

Load Balancer created groups are visible under **Inventory > Groups**.

Server pool groups are created with the name `NLB.PoolLB.Pool_Name LB_Name` with group member IP addresses assigned:

- Pool configured with no LB-SNAT (transparent): 0.0.0.0/0
- Pool configured with LB-SNAT Automap: T1-Uplink IP 100.64.x.y and T1-ServiceInterface IP
- Pool configured with LB-SNAT IP-Pool: LB-SNAT IP-Pool

VIP Groups are created with the name `NLB.VIP.virtual server name` and the VIP group member IP addresses are VIP IP@.

For server pool groups, you can create an allow traffic distributed firewall rule from the load balancer (`NLB.PoolLB. Pool_Name LB_Name`). For Tier-1 gateway firewall, you can create an allow traffic from clients to LB VIP `NLB.VIP.virtual server name`.

Distributed Load Balancer

9

A Distributed Load Balancer configured in NSX-T Data Center can help you effectively load balance East-West traffic and scale traffic because it runs on each ESXi host.


Important Distributed Load Balancer is supported only for Kubernetes (K8s) cluster IPs managed by vSphere with Kubernetes. Distributed Load Balancer is not supported for any other workload types. As an administrator, you cannot use NSX Manager GUI to create or modify Distributed Load Balancer objects. These objects are pushed by vCenter Server through NSX-T API when K8 cluster IPs are created in vCenter Server.

Note Do not enable Distributed Intrusion Detection Service (IDS) in an environment that is using Distributed Load Balancer. NSX-T Data Center does not support using IDS with a Distributed Load Balancer.

In traditional networks, a central load balancer deployed on an NSX Edge node is configured to distribute traffic load managed by virtual servers that are configured on the load balancer.

If you are using a central balancer, increasing the number of virtual servers in the load balancer pool might not always meet scale or performance criteria for a multi-tier distributed application. A distributed load balancer is realized on each hypervisor where load balancing workloads, such as clients and servers are deployed, ensuring traffic is load balanced on each hypervisor in a distributed way.

A distributed load balancer can be configured on the NSX-T Data Center network along with a central load balancer.



In the diagram, an instance of the Distributed Load Balancer is attached to a VM group. As the VMs are downlinks to the distributed logical router, Distributed Load Balancer only load balances east-west traffic. In contrast, the central load balancer, manages north-south traffic.


To cater load balancing requirements of each component or module of an application, a distributed load balancer can be attached to each tier of an application. For example, to serve a user request, a frontend of the application needs to reach out to the middle module to get data. However, the middle layer might not be deployed to serve the final data to the user, so it needs to reach out the backend layer to get additional data. For a complex application, many modules might need to interact with each other to get information. Along with complexity, when the number of user request increase exponentially, a distributed load balancer can efficiently meet the user needs without taking a performance hit. Configuring a Distributed Load Balancer on every host achieves issues of scale and packet transmission efficiency.

Read the following topics next:

- [Understanding Traffic Flow with a Distributed Load Balancer](#)
- [Create and Attach a Distributed Load Balancer Instance](#)
- [Create a Server Pool for Distributed Load Balancer](#)
- [Create a Virtual Server with a Fast TCP or UDP Profile](#)
- [Verifying Distributed Load Balancer Configuration on ESXi Hosts](#)
- [Monitoring Distributed Load Balancer Statistics](#)

Understanding Traffic Flow with a Distributed Load Balancer

Understand how traffic flows between VMs that are connected to an instance of a distributed load balancer (DLB).



As an administrator ensure:

- Virtual IP addresses and pool members connected to a DLB instance must have unique IP address for traffic to be routed correctly.

Traffic flow between Web VM1 and APP VM2.

- When Web VM1 sends out a packet to APP VM2 it is received by the VIP-APP.
The DLB APP is attached to the policy group consisting of Web tier VMs. Similarly, DLB-APP hosting VIP-DB must be attached to the policy group consisting of App tier VMs.
- The VIP-APP hosted on DLB APP receives the request from Web VM1.
- Before reaching the destination VM group, the packet is filtered by distributed firewall rules.
- After the packets are filtered based on the firewall rules, it is sent to the Tier-1 router.
- It is further routed to the physical router.
- The route is completed when the packet is delivered to the destination App VM2 group.

As DLB VIPs can only be accessed from VMs connected to downlinks of Tier-0 or Tier-1 logical routers, DLB provides load balancing services to east-west traffic.

A DLB instance can co-exist with an instance of DFW. With DLB and DFW enabled on a virtual interface of a hypervisor, first the traffic is load balanced based on the configuration in DLB and then DFW rules are applied on traffic flowing from a VM to the hypervisor. DLB rules are applied on traffic originating from downlinks of a Tier-0 or Tier-1 logical routers going to the destination hypervisor. DLB rules cannot be applied on traffic flowing in the reverse direction - originating from outside the host going to a destination VM.

For example, if the DLB instance is load balancing traffic from Web-VMs to App-VMs, then to allow such traffic to pass through DFW, ensure that the DFW rule is set to value "Source=Web-VMs, Destination=App-VMs, Action=Allow".

Create and Attach a Distributed Load Balancer Instance

Unlike a central load balancer, a Distributed Load Balancer (DLB) instance is attached to virtual interfaces of a VM group.

At the end of the procedure a DLB instance is attached to the virtual interfaces of a VM group.

It is only possible to create and attach a DLB instance through API commands.

Prerequisites

- Add a policy group consisting of VMs. For example, such a VM group can be related to the App tier that receives requests from a VM on the Web-tier.

Procedure

- ◆ Run `Put /policy/api/v1/infra/lb-services/<mydlb>`.

```
{
  "connectivity_path" : "/infra/domains/default/groups/<clientVMGroup>",
  "enabled" : true,
  "size" : "DLB",
  "error_log_level" : "INFO",
  "access_log_enabled" : false,
  "resource_type" : "LBService",
  "display_name" : "mydlb"
}
```

Where,

- **connectivity_path:**
 - If the connectivity path is set to Null or Empty, the DLB instance is not applied to any transport nodes.
 - If the connectivity path is set ALL, all virtual interfaces of all transport nodes are bound to the DLB instance. One DLB instance is applied to all the virtual interfaces of the policy group.
- **size:** Set to value **DLB**. As each application or virtual interface gets an instance of DLB, there is just a single size form factor of the DLB instance.
- **enabled:** By default, the created DLB instance is enabled.

A DLB instance is created and attached to the VM group. The DLB instance created on the Web-tier is attached to all the virtual interfaces of the App-tier VM group.

What to do next

After creating a DLB instance, log in to the NSX Manager, go to **Networking -> Load Balancing -> Load Balancers**. View details of the DLB instance.

Next, [Create a Server Pool for Distributed Load Balancer](#).

Create a Server Pool for Distributed Load Balancer

Create a load balancer pool to include virtual machines that consume DLB services.

This task can be done both from the NSX-T UI and NSX-T API.

The API command to create a DLB pool, PUT `https://<NSXManager_IPaddress>/policy/api/v1/infra/lb-pools/<lb-pool-id>`

Prerequisites

- Create a VM group that consumes DLB service.
- Create and attach a DLB instance to a VM group.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Go to **Networking → Load Balancing → Server Pools**.
- 3 Click **Add Server Pool**.
- 4 Enter values in these fields.

Field	Description
Name	Enter name of the DLB pool.
Algorithm	Only ROUND_ROBIN is supported for a Distributed Load Balancer.
Members	Click Select Members and add individual members to the group. When adding individual members, only enter values to the following text boxes: ■ Name ■ IP address ■ Port Note This list includes the only supported fields to add members to a DLB pool.

Field	Description
Member Group	Click Select Members and add a member group. When adding individual members, enter values to the following fields: ■ Name ■ Compute Members: Click Set Members to add a group that includes all the pool members. ■ IP Revision Filter: Only IPv4 is supported. ■ Port: Default port for all the dynamic pool members. Note This list includes the only supported fields to add members to a DLB pool.
SNAT Translation Mode	Set this field to Disabled state. SNAT translation is not supported in a Distributed Load Balancer.

5 Click **Save**.

Results

Server pool members are added for the Distributed Load Balancer.

What to do next

See [Create a Virtual Server with a Fast TCP or UDP Profile](#).

Create a Virtual Server with a Fast TCP or UDP Profile

Create a virtual server and bind it to a Distributed Load Balancer service.

This task can be performed both from the NSX-T UI and NSX-T APIs.

The API command to create a virtual server is `PUT https://<NSXManager_IPAddress>/policy/api/v1/infra/lb-virtual-servers/<lb-virtual-server-id>`.

Prerequisites

- Create a server pool for the Distributed Load Balancer.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Go to **Networking** → **Load Balancing** → **Virtual Servers**.
- 3 Click **Add Virtual Server** -> **L4 TCP**.

- 4 To configure a virtual server for a Distributed Load Balancer, only the following fields are supported.

Field	Description
Name	Enter a name for the virtual server.
IP Address	IP address of the Distributed Load Balancer virtual server. Configures the IP address of the Distributed Load Balancer virtual server where it receives all client connections and distributes them among the backend servers.
Ports	Virtual server port number. Multiple ports or port ranges are not supported in the virtual server of a Distributed Load Balancer.
Load Balancer	Attach the Distributed Load Balancer instance that is associated to the virtual server. The virtual server then knows which policy group the load balancer is servicing.
Server Pool	Select the server pool. The server pool contains backend servers. Server pool consists of one or more servers that are similarly configured and are running the same application. It is also referred to as pool members.
Application Profile	Select the application profile for the virtual server. The application profile defines the application protocol characteristics. It is used to influence how load balancing is performed. The supported application profiles are: ■ Load Balancer Fast TCP Profile ■ Load Balancer Fast UDP Profile
Default Pool Member Ports	Optional field. Enter one port number to be used when member ports are not defined. Multiple ports or port ranges for default pool member ports are not supported in the virtual server of a Distributed Load Balancer.
Persistence	Optional field. Select Source IP or Disabled .

The Distributed Load Balancer configuration is complete.

Results

Verify whether the DLB is distributing traffic to all the servers in the pool based on the algorithm defined in the configuration. If you choose the Round_Robin algorithm, then DLB must be able to choose servers from the pool in a round robin fashion.

In the ESXi host, verify whether the DLB configuration is complete.

What to do next

See [Verifying Distributed Load Balancer Configuration on ESXi Hosts](#).

Verifying Distributed Load Balancer Configuration on ESXi Hosts

Verify whether the Distributed Load Balancer was configured completely on ESXi hosts.

After you securely connect to the ESXi host, run `/opt/vmware/nsx-nestdb/bin/nestdb-cli`. From the `nestdb-cli` prompt, run the following commands.

Command	Sample Response
To view the configured DLB service, run <code>get LbServiceMsg</code> .	<pre>{'id': {'left': 13946864992859343551, 'right': 10845263561610880178}, 'virtual_server_id': [{'left': 13384746951958284821, 'right': 11316502527836868364}], 'display_name': 'mydlb', 'size': 'DLB', 'enabled': True, 'access_log_enabled': False, 'log_level': 'LB_LOG_LEVEL_INFO', 'applied_to': {'type': 'CONTAINER', 'attachment_id': {'left': 2826732686997341216, 'right': 10792930437485655035}}}</pre>
To view the virtual server configured for DLB, run <code>get LbVirtualServerMsg</code> .	<pre>{'port': '80', 'revision': 0, 'display_name': 'mytcpvip', 'pool_id': {'left': 4370937730160476541, 'right': 13181758910457427118}, 'enabled': True, 'access_log_enabled': False, 'id': {'left': 13384746951958284821, 'right': 11316502527836868364}, 'ip_protocol': 'TCP', 'ip_address': {'ipv4': 2071690107}, 'application_profile_id': {'left': 1527034089224553657, 'right': 10785436903467108397}}}</pre>
To view configuration of the DLB pool members, run <code>get LbPoolMsg</code> .	<pre>{'tcp_multiplexing_number': 6, 'display_name': 'mylbpool', 'tcp_multiplexing_enabled': False, 'member': [{'port': '80', 'weight': 1, 'display_name': 'Member_VM30', 'admin_state': 'ENABLED', 'ip_address': {'ipv4': 3232261280}, 'backup_member': False}, {'port': '80', 'weight': 1, 'display_name': 'Member_VM31', 'admin_state': 'ENABLED', 'ip_address': {'ipv4': 3232261281}, 'backup_member': False}, {'port': '80', 'weight': 1, 'display_name': 'Member_VM32', 'admin_state': 'ENABLED', 'ip_address': {'ipv4': 3232261282}, 'backup_member': False}], 'id': {'left': 4370937730160476541, 'right': 13181758910457427118}, 'min_active_members': 1, 'algorithm': 'ROUND_ROBIN'}</pre>

Command	Sample Response
To view NSX controller configuration pushed to the ESXi host, run <code>get ContainerMsg</code> .	<pre>{'container_type': 'CONTAINER', 'id': {'left': 2826732686997341216, 'right': 10792930437485655035}, 'vif': ['cd2e482b-2998-480f-beba-65fbd7able62', 'f8aa2a58-5662-4c6b-8090-d1bd19174205', '83a1f709- e675-4e42-b677-ff501fd0f4ec', 'b8366b39-4c81-41fc- b89e-de7716462b2f'], 'name': 'default.clientVMGroup', 'mac_address': [{'mac': 52237218275}, {'mac': 52243694681}, {'mac': 52233233291}, {'mac': 52239463383}], 'ip_address': [{'ipv4': 16844388}, {'ipv4': 16844644}, {'ipv4': 16844132}, {'ipv4': 3232261283}, {'ipv4': 16844298}, {'ipv4': 16844554}, {'ipv4': 16844042}]}</pre>
To view application profile configuration on the ESXi host, run <code>get LbApplicationProfileMsg</code> .	<pre>{'display_name': 'default-tcp-lb-app-profile', 'id': {'left': 1527034089224553657, 'right': 10785436903467108397}, 'application_type': 'FAST_TCP', 'fast_tcp_profile': {'close_timeout': 8, 'flow_mirroring_enabled': False, 'idle_timeout': 1800}}</pre>

Monitoring Distributed Load Balancer Statistics

NSX-T Data Center CLI commands to monitor statistics for Distributed Load Balancer instances.

Action	Command
Display all load balancers.	<code>get load-balancers</code>
Display a specific load balancer.	<code>get load-balancer <UUID_LoadBalancer></code>
Show load balancer virtual-server configuration.	<code>get load-balancer <UUID_LoadBalancer> virtual-servers</code>
Show statistics of all pools of the specified load balancer	<code>get load-balancer <UUID_LoadBalancer> pools stats</code>
Show statistics of the specified load balancer and pool	<code>get load-balancer <UUID_LoadBalancer> pool <UUID_Pool> stats</code>
Show persistence-tables entry	<code>get load-balancer <UUID_LoadBalancer> persistence-tables</code>
Show load balancer pools configuration	<code>get load-balancer <UUID_LoadBalancer> pools</code>
Show statistics of all virtual servers of the specified load balancer	<code>get load-balancer <UUID_LoadBalancer> virtual-servers stats</code>
Show statistics of the specified load balancer and virtual server	<code>get load-balancer <UUID_LoadBalancer> virtual-server <UUID_VirtualServer> stat</code>
Clear statistics of the specified load balancer and pool	<code>clear load-balancer <UUID_LoadBalancer> pool <UUID_Pool> stats</code>

Action	Command
Clear statistics of all pools of the specified load balancer	<code>clear load-balancer <UUID_LoadBalancer> pools stats</code>
Clear statistics of the specified load balancer	<code>clear load-balancer <UUID_LoadBalancer> stats</code>
Clear statistics of the specified load balancer and virtual server	<code>clear load-balancer <UUID_LoadBalancer> virtual-server <UUID_VirtualServer> stats</code>
Clear statistics of all virtual servers of the specified load balancer	<code>clear load-balancer <UUID_LoadBalancer> virtual-servers stats</code>

Forwarding Policies

10

This feature pertains to NSX Cloud.

Forwarding Policies or Policy-Based Routing (PBR) rules define how NSX-T handles traffic from an NSX-managed VM. This traffic can be steered to NSX-T overlay or it can be routed through the cloud provider's (underlay) network.

Note See [Chapter 23 Using NSX Cloud](#) for details on how to manage your public cloud workload VMs with NSX-T Data Center.

Three default forwarding policies are set up automatically after you either deploy a PCG on a Transit VPC/VNet or link a Compute VPC/VNet to the Transit.

- 1 One **Route to Underlay** for all traffic that is addressed within the Transit/Compute VPC/VNet
- 2 Another **Route to Underlay** for all traffic destined to the metadata services of the public cloud.
- 3 One **Route to Overlay** for all other traffic, for example, traffic that is headed outside the Transit/Compute VPC/VNet. Such traffic is routed over the NSX-T overlay tunnel to the PCG and further to its destination.

Note **For traffic destined to another VPC/VNET managed by the same PCG:** Traffic is routed from the source NSX-managed VPC/VNet via the NSX-T overlay tunnel to the PCG and then routed to the destination VPC/VNet.

For traffic destined to another VPC/VNet managed by a different PCG: Traffic is routed from one NSX-managed VPC/VNet over the NSX overlay tunnel to the PCG of the source VPC/VNet and forwarded to the PCG of the destination NSX-managed VPC/VNet.

If traffic is headed to the internet, the PCG routes it to the destination in the internet.

Micro-segmentation while Routing to Underlay

Micro-segmentation is enforced even for workload VMs whose traffic is routed to the underlay network.

If you have direct connectivity from an NSX-managed workload VM to a destination outside the managed VPC/VNet and want to bypass the PCG, set up a forwarding policy to route traffic from this VM via underlay.

When traffic is routed through the underlay network, the PCG is bypassed and therefore the north-south firewall is not encountered by traffic. However, you still have to manage rules for east-west or distributed firewall (DFW) because those rules are applied at the VM-level before reaching the PCG.

Supported Forwarding Policies and Common Use Cases

You may see a list of forwarding policies in the drop-down menu but in this release only the following forwarding policies are supported:

- Route to Underlay
- Route from Underlay
- Route to Overlay

These are the common scenarios where forwarding policies are useful:

- **Route to Underlay:** Access a service on underlay from an NSX-managed VM. For example, access to the AWS S3 service on the AWS underlay network.
- **Route from Underlay:** Access a service hosted on an NSX-managed VM from the underlay network. For example, access from AWS ELB to the NSX-managed VM.

Read the following topics next:

- [Add or Edit Forwarding Policies](#)

Add or Edit Forwarding Policies

You can edit the auto-created forwarding policies or add new ones.

For example, to use services provided by the public cloud, such as S3 by AWS, you can manually create a policy to allow a set of IP addresses to access this service by being routed through underlay.

Prerequisites

You must have a VPC or VNet with a PCG deployed on it.

Procedure

- 1 Click **Add Section**. Name the section appropriately, for example, **AWS Services**.
- 2 Select the check box next to the section and click **Add Rule**. Name the rule, for example, **S3 Rules**.

- 3 In the **Sources** tab, select the VPC or VNet where you have the workload VMs to which you want to provide the service access, for example, the AWS VPC. You can also create a **Group** here to include multiple VMs matching one or more criteria.
- 4 In the **Destinations** tab, select the VPC or VNet where the service is hosted, for example, a **Group** that contains the IP address of the S3 service in AWS.
- 5 In the **Services** tab, select the service from the drop-down menu. If the service does not exist, you can add it. You can also leave the selection to **Any** because you can provide the routing details under **Destinations**.
- 6 In the **Action** tab, select how you want the routing to work, for example, select **Route to Underlay** if setting up this policy for the AWS S3 service.
- 7 Click **Publish** to finish setting up the Forwarding Policy.

IP Address Management (IPAM)

11

To manage IP addresses, you can configure DNS (Domain Name System), DHCP (Dynamic Host Configuration Protocol), IP address pools, and IP address blocks.

Note IP blocks are used by NSX Container Plug-in (NCP). For more info about NCP, see the *NSX Container Plug-in for Kubernetes and Cloud Foundry - Installation and Administration Guide*.

Read the following topics next:

- [Add a DNS Zone](#)
- [Add a DNS Forwarder Service](#)
- [Add a DHCP Profile](#)
- [Attach a DHCP Profile to a Tier-0 or Tier-1 Gateway](#)
- [Scenarios: Selection of Edge Cluster for DHCP Service](#)
- [Scenarios: Impact of Changing Segment Connectivity on DHCP](#)
- [Add an IP Address Pool](#)
- [Add an IP Address Block](#)

Add a DNS Zone

You can configure DNS zones for your DNS service. A DNS zone is a distinct portion of the domain name space in DNS.

When you configure a DNS zone, you can specify a source IP for a DNS forwarder to use when forwarding DNS queries to an upstream DNS server. If you do not specify a source IP, the DNS query packet's source IP will be the DNS forwarder's listener IP. Specifying a source IP is needed if the listener IP is an internal address that is not reachable from the external upstream DNS server. To ensure that the DNS response packets are routed back to the forwarder, a dedicated source IP is needed. Alternatively, you can configure SNAT on the logical router to translate the listener IP to a public IP. In this case, you do not need to specify a source IP.

Procedure

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

- 2 Select **Networking > IP Management > DNS**.
- 3 Click the **DNS Zones** tab.
- 4 To add a default zone, select **Add DNS Zone > Add Default Zone**
 - a Enter a name and optionally a description.
 - b Enter the IP address of up to three DNS servers.
 - c (Optional) Enter an IP address in the **Source IP** field.
- 5 To add an FQDN zone, select **Add DNS Zone > Add FQDN Zone**
 - a Enter a name and optionally a description.
 - b Enter a FQDN for the domain.
 - c Enter the IP address of up to three DNS servers.
 - d (Optional) Enter an IP address in the **Source IP** field.
- 6 Click **Save**.

Add a DNS Forwarder Service

You can configure a DNS forwarder to forward DNS queries to external DNS servers.

Before you configure a DNS forwarder, you must configure a default DNS zone. Optionally, you can configure one or more FQDN DNS zones. Each DNS zone is associated with up to 3 DNS servers. When you configure a FQDN DNS zone, you specify one or more domain names. A DNS forwarder is associated with a default DNS zone and up to 5 FQDN DNS zones. When a DNS query is received, 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. If a match is not found, the query is forwarded to the DNS servers specified in the default DNS zone.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > IP Management > DNS**.
- 3 Click **Add DNS Service**.
- 4 Enter a name and optionally a description.
- 5 Select a tier-0 or tier-1 gateway.
- 6 Enter the IP address of the DNS service.

Clients send DNS queries to this IP address, which is also known as the DNS forwarder's listener IP.
- 7 Select a default DNS zone.

- 8 Select a log level.
- 9 Select up to five FQDN zones.
- 10 Click the **Admin Status** toggle to enable or disable the DNS service.
- 11 Click **Save**.

Add a DHCP Profile

Before you can configure DHCP on a segment, you must add a DHCP profile in your network. You can create two types of DHCP profiles: DHCP server profile and DHCP relay profile.

A DHCP profile can be used simultaneously by multiple segments and gateways in your network. The following conditions apply when you attach a DHCP profile to a segment or a gateway:

- On a tier-0 or tier-1 gateway or a gateway-connected segment, you can attach either a DHCP server profile or a DHCP relay profile.
- On a standalone segment that is not connected to a gateway, you can attach only a DHCP server profile. Standalone segment supports only a local DHCP server.

- [Add a DHCP Server Profile](#)

You can add multiple DHCP server profiles in your network. Further, you can attach a single DHCP server profile to multiple DHCP servers.

- [Add a DHCP Relay Profile](#)

You can add a DHCP relay profile to relay the DHCP traffic to remote DHCP servers. The remote or external DHCP servers can be in any overlay segment, outside the SDDC, or in the physical network.

Add a DHCP Server Profile

You can add multiple DHCP server profiles in your network. Further, you can attach a single DHCP server profile to multiple DHCP servers.

Prerequisites

- Edge nodes are deployed in the network.
- Edge cluster is added in the network.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > IP Management > DHCP**.
- 3 Click **Add DHCP Profile**.
- 4 Enter a unique name to identify the DHCP server profile.
- 5 In the **Profile Type** drop-down menu, select **DHCP Server**.

- 6 (Optional) Enter the IP address of the DHCP server in a CIDR format.

Note A maximum of two DHCP server IP addresses are supported. You can enter 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 DHCP server IP address must not overlap with the addresses used in DHCP ranges and DHCP static binding.

If no server IP address is specified, 100.96.0.1/30 is autoassigned to the DHCP server.

The server IP address cannot be any of the following:

- Multicast IP address
 - Broadcast IP address
 - Loopback IP address
 - Unspecified IP address (address with all zeroes)
- 7 (Optional) Edit the lease time in seconds. The default value is 86400.

Valid range of values is 60–4294967295.

- 8 Select an Edge cluster.

Follow these guidelines:

- If you are using a local DHCP server on a segment, you must select an edge cluster in the DHCP server profile. If an edge cluster is unavailable in the profile, an error message is displayed when you save the segment.
- If you are using a Gateway DHCP server on the segment, select an edge cluster either in the gateway, or DHCP server profile, or both. If an edge cluster is unavailable in either the profile or the gateway, an error message is displayed when you save the segment.

Caution You can change the edge cluster in the profile after the DHCP server is created. However, this action causes all the existing DHCP leases that are assigned to the DHCP clients to be lost.

When a DHCP server profile is attached to a segment that uses a DHCP local server, the DHCP service is created in the edge cluster that you specified in the DHCP profile. However, if the segment uses a Gateway DHCP server, the edge cluster in which the DHCP service is created depends on a combination of several factors. For a detailed information about how an edge cluster is selected for DHCP service, see [Scenarios: Selection of Edge Cluster for DHCP Service](#).

- 9 (Optional) Next to **Edges**, click **Set** and select the preferred edge nodes where you want the DHCP service to run.

To select the preferred edge nodes, edge cluster must be selected. You can select a maximum of two preferred edge nodes. The following table explains the scenarios when DHCP HA is configured.

Scenario	DHCP HA
No preferred edge node is selected from the edge cluster.	DHCP HA is configured. A pair of active and standby edge nodes are selected automatically from the available nodes in the edge cluster.
Only one preferred edge node is selected from the edge cluster.	DHCP server runs without the HA support.
Two preferred edge nodes are selected from the edge cluster.	DHCP HA is configured. The first edge node that you add becomes the active edge, and the second edge node becomes the standby edge. The active edge is denoted with a sequence number 1, and the standby edge is denoted with a sequence number 2. You can interchange the active and standby edges. For example to change the current active edge to standby, select the active edge and click the Down arrow. Alternatively, you can select the passive edge and click the Up arrow to make it active. The sequence numbers are reversed in both situations.

After the DHCP server is created, you can change the preferred edge nodes in the DHCP server profile. However, this flexibility includes certain caveats.

For example, let us assume that the edge cluster in the DHCP profile has four edge nodes N1, N2, N3, and N4, and you have set N1 and N2 as the preferred edge nodes. N1 is the active edge and N2 is the standby edge. The DHCP service is running on the active edge node N1, and the DHCP server has started assigning leases to the DHCP clients on the segment.

Scenario	Impact on DHCP Service
Delete existing preferred edge nodes N1 and N2, and add N3 and N4 as the new preferred edge nodes.	A warning message informs you that the current DHCP leases will be lost due to the replacement of existing preferred edges. This action can cause a loss of network connectivity. You can prevent loss of connectivity by replacing one edge node at a time.
Delete existing preferred edges N1 and N2, and keep the preferred edge nodes list empty.	The DHCP servers remain on the edge nodes N1 and N2. The DHCP leases are retained and the DHCP clients do not lose network connectivity.
Delete any one of the preferred edges, either N1 or N2.	When any one of the preferred edges N1 or N2 is deleted, the other edge continues to provide IP addresses to the DHCP clients. The DHCP leases are retained and the DHCP clients do not experience a loss of network connectivity. However, DHCP HA support is lost. To retain DHCP HA, you must replace the deleted edge with another edge node, either N3 or N4, in the edge cluster.

- 10 (Optional) 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 DHCP profile.

- 11 Click **Save**.

What to do next

Attach the DHCP server profile either to a segment or a gateway, and configure the DHCP server settings at the level of each segment.

- [Attach a DHCP Profile to a Tier-0 or Tier-1 Gateway.](#)
- [Configure DHCP on a Segment.](#)

Add a DHCP Relay Profile

You can add a DHCP relay profile to relay the DHCP traffic to remote DHCP servers. The remote or external DHCP servers can be in any overlay segment, outside the SDDC, or in the physical network.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > IP Management > DHCP**.
- 3 Click **Add DHCP Profile**.
- 4 Enter a unique name to identify the relay profile.
- 5 In the **Profile Type** drop-down menu, select **DHCP Relay**.
- 6 (Required) Enter the IP addresses of the remote DHCP servers.

Both DHCPv4 and DHCPv6 servers are supported. You can enter multiple IP addresses. The server IP addresses of the remote DHCP servers must not overlap with the addresses that are used in DHCP ranges and DHCP static binding.

The server IP address cannot be any of the following:

- Multicast IP address
 - Broadcast IP address
 - Loopback IP address
 - Unspecified IP address (address with all zeroes)
- 7 (Optional) 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 DHCP profile.

8 Click **Save**.

What to do next

Attach the DHCP relay profile either to a gateway, or use the profile to configure a local DHCP relay on the segment.

- [Attach a DHCP Profile to a Tier-0 or Tier-1 Gateway.](#)
- [Configure DHCP on a Segment.](#)

Attach a DHCP Profile to a Tier-0 or Tier-1 Gateway

To use Gateway DHCP for a dynamic IP assignment, you must attach a DHCP server profile to a tier-0 or tier-1 gateway.

You can attach a DHCP profile to a gateway only when the segments connected to that gateway do not have a local DHCP server or DHCP relay configured on them. If a local DHCP server or DHCP relay exists on the segment, the UI throws an error when you try to attach a DHCP profile to the gateway. You must disconnect the segments from the gateway, and then attach a DHCP profile to the gateway.

Prerequisites

A DHCP server profile is added in the network.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Go to **Networking > Tier-0 Gateways** or **Networking > Tier-1 Gateways**.
- 3 Edit the appropriate gateway.
- 4 Do any of one of the following depending on the version of NSX-T Data Center that you are using:
 - In version 3.0.2, next to **DHCP**, click **Set DHCP Configuration**.
 - In version 3.0 and 3.0.1, next to **IP Address Management**, click **No Dynamic IP Allocation**.

- 5 In the **Type** drop-down menu, select **DHCP Server** or **DHCP Relay**.

Note If you select the profile type as DHCP Relay, the configuration does not take any effect. You must assign the DHCP relay profile to the segments, which are connected to the gateway. Attaching a DHCP relay profile to the gateway is a redundant configuration. This functional behavior is a known issue. For information about assigning a DHCP relay profile to the segment, click the [Configure DHCP on a Segment](#) link in the What to do next section of this topic.

- 6 Select a DHCP server profile to attach to this gateway.

- 7 Click **Save**.

What to do next

Navigate to **Networking > Segments**. On each segment that is connected to this gateway, configure the DHCP settings, static bindings, and other DHCP options.

For detailed steps, see:

- [Configure DHCP on a Segment](#).
- [Configure DHCP Static Bindings on a Segment](#).

After a Gateway DHCP server is in use, you can view the DHCP server statistics on the gateway. On the gateway, next to **DHCP** or **IP Address Management**, click the Servers link. On the **Set DHCP Configuration** page, click **Statistics**.

The Gateway DHCP server statistics are displayed in a pop-up window.

Note If you have configured a local DHCP server on a gateway-connected segment, the **Statistics** link on the **Set DHCP Configuration** page does not display the local DHCP server statistics. Only Gateway DHCP statistics are shown on this page.

Scenarios: Selection of Edge Cluster for DHCP Service

DHCP server runs as a service (service router) in the edge nodes of an NSX Edge cluster.

Standalone segments that are not connected to a gateway can use only a DHCP local server. Segments that are connected to a gateway can use either a DHCP local server, DHCP relay, or Gateway DHCP server.

Regardless of whether a segment uses a DHCP local server or a Gateway DHCP server, DHCP server always runs as a service router in the edge transport nodes of an edge cluster. If the segment uses a DHCP local server, the DHCP service is created in the edge cluster that you specified in the DHCP profile. However, if the segment uses a Gateway DHCP server, the edge cluster in which the DHCP service is created depends on the combination of the following factors:

- Is an edge cluster specified in the gateway?
- Is an edge cluster specified in the DHCP profile of the gateway?

- Is the edge cluster in the gateway and in the DHCP profile same or different?
- Is the tier-1 routed segment connected to a tier-0 gateway?

The following scenarios explain how the edge cluster is selected for creating the DHCP service.

Scenario 1: Standalone Segment Uses DHCP Local Server

Scenario Description:

- An edge cluster (Cluster1) is created with four edge nodes: N1, N2, N3, N4.
- A segment with None connectivity is added in the overlay transport zone.
- Segment uses a DHCP local server, by default.

The DHCP server profile configuration is as follows:

- Profile Type: **DHCP Server**
- Edge Cluster: **Cluster1**
- Preferred Edges: **None**

In this scenario, any two edge nodes from Cluster1 are autoallocated to create the DHCP service, and DHCP high availability (HA) is automatically configured. One of the edge nodes in Cluster1 runs in active mode and the other edge runs in passive mode.

Note

- If you select two preferred edge nodes in the DHCP profile, the edge node that is added first becomes the active edge. The second edge node takes the passive role.
 - If you select only one preferred edge node in the DHCP profile, DHCP HA is not configured.
-

Scenario 2: Tier-1 Routed Segment Uses Gateway DHCP and Different Edge Clusters in Gateway and DHCP Profile

Consider that you have two edge clusters in your network (Cluster1 and Cluster2). Both clusters have four edge nodes each:

- Cluster1 edge nodes: N1, N2, N3, N4
- Cluster2 edge nodes: N5, N6, N7, N8

Scenario Description:

- Segment is connected to a tier-1 gateway.
- Tier-1 gateway is not connected to a tier-0 gateway.
- DHCP server profile in the tier-1 gateway uses Cluster1.
- Tier-1 gateway uses Cluster2.
- Segment is configured to use the Gateway DHCP server.

The DHCP server profile in the tier-1 gateway has the following configuration:

- Profile Type: **DHCP Server**
- Edge Cluster: **Cluster1**
- Preferred Edges: **N1,N2** (added in the given sequence)

The tier-1 gateway configuration is as follows:

- Edge Cluster: **Cluster2**
- Preferred Edges: **N5,N6** (added in the given sequence)

In this scenario, DHCP service runs on the edge nodes of Cluster2. As Cluster2 contains multiple edge nodes, DHCP HA is autoconfigured. However, the preferred edges N5 and N6 on the gateway are ignored for DHCP HA. Any two nodes from Cluster2 are randomly autoallocated for DHCP HA.

This scenario also applies when the segment is directly connected to a tier-0 gateway, and there is no tier-1 gateway in your network topology.

Caution Starting in NSX-T Data Center 3.0.2, you can change the edge cluster on the Gateway DHCP server after the DHCP server is created. However, this action causes all the existing DHCP leases that are assigned to the DHCP clients to be lost.

To summarize, the main points of this scenario are as follows:

- When you use a Gateway DHCP server and set different edge clusters in the gateway DHCP profile and tier-1 gateway, then DHCP service is always created in the edge cluster of the gateway.
- The edge nodes are randomly allocated from the edge cluster of the tier-1 gateway for DHCP HA configuration.
- If no edge cluster is specified in the tier-1 gateway, the edge cluster in the DHCP profile of the tier-1 gateway (Cluster1) is used to create the DHCP service.

Scenario 3: Tier-1 Routed Segment Uses Local DHCP Server and Different Edge Clusters in Gateway and DHCP Profile

In this scenario, a segment is connected to a tier-1 gateway, but you use a local DHCP server on the segment. Consider that you have three edge clusters in your network (Cluster1, Cluster2, Cluster 3). Each cluster has two edges nodes each.

- Cluster1 edge nodes: N1, N2
- Cluster2 edge nodes: N3, N4
- Cluster3 edge nodes: N5, N6

Scenario Description:

- Segment is connected to a tier-1 gateway.

- Tier-1 gateway is connected to a tier-0 gateway (optional).
- DHCP profile on the gateway uses Cluster1.
- Gateway uses Cluster2.
- Segment is configured to use DHCP local server.
- Local DHCP server profile uses Cluster3.

The DHCP server profile on the gateway is as follows:

- Profile Name: **ProfileX**
- Profile Type: **DHCP Server**
- Edge Cluster: **Cluster1**
- Preferred Edges: **N1,N2** (added in the given sequence)

The tier-1 gateway configuration is as follows:

- Edge Cluster: **Cluster2**
- Preferred Edges: **N3,N4** (added in the given sequence)

The profile of the local DHCP server is as follows:

- Profile Name: **ProfileY**
- Profile Type: **DHCP Server**
- Edge Cluster: **Cluster3**
- Preferred Edges: **N5,N6** (added in the given sequence)

In this scenario, because the segment is configured to use a local DHCP server, the edge cluster (Cluster2) in the connected tier-1 gateway is ignored to create the DHCP service. DHCP service runs in the edge nodes of Cluster3 (N5, N6). DHCP HA is also configured. N5 becomes the active edge node and N6 becomes the standby edge.

If no preferred nodes are set in Cluster3, any two nodes from this cluster are autoallocated for creating the DHCP service and configuring DHCP HA. One of the edge nodes becomes an active edge and the other node becomes the standby edge. If only one preferred edge node is set in Cluster3, DHCP HA is not configured.

This scenario also applies when the segment is directly connected to a tier-0 gateway, and there is no tier-1 gateway in your network topology.

Scenario 4: Tier-1 Routed Segment Uses Gateway DHCP and Same Edge Clusters in Gateway and DHCP Profile

Consider that you have a single edge cluster (Cluster1) in your network with four edge nodes: N1, N2, N3, N4.

Scenario Description:

- Segment is connected to a tier-1 gateway.
- Tier-1 gateway is connected to a tier-0 gateway (optional)
- Gateway and DHCP profile on the gateway use the same edge cluster (Cluster1).
- Segment is configured to use Gateway DHCP server.

The DHCP server profile on the gateway is as follows:

- Profile Type: **DHCP Server**
- Edge Cluster: **Cluster1**
- Preferred Edges: **N1 ,N2** (added in the given sequence)

The tier-1 gateway configuration is as follows:

- Edge Cluster: **Cluster1**
- Preferred Edges: **N3 ,N4** (added in the given sequence)

In this scenario, as the gateway DHCP profile and gateway use a similar edge cluster (Cluster1), DHCP service is created in the preferred edge nodes N1 and N2 of the gateway DHCP profile. The preferred edge nodes N3 and N4 that you specified in the connected tier-1 gateway are ignored for creating the DHCP service.

If no preferred edges are set in the DHCP profile, any two nodes from Cluster1 are autoallocated for creating the DHCP service and configuring DHCP HA. One of the edge nodes becomes an active edge and the other edge becomes the standby edge.

To summarize, the main points of this scenario are as follows:

- When you use a Gateway DHCP server and specify similar edge clusters in the DHCP profile and connected gateway, then DHCP service is created in the preferred edge nodes of the DHCP profile.
- The preferred edges nodes specified in the connected gateway are ignored.

Scenario 5: Tier-1 Routed Segment is Connected to Tier-0 Gateway and No Edge Cluster is Set in Tier-1 Gateway

In this scenario, a segment is connected to a tier-1 gateway, and the tier-1 gateway is connected to a tier-0 gateway. Consider that you have three edge clusters in your network (Cluster1, Cluster2, Cluster 3). Each cluster has two edges nodes each.

- Cluster1 edge nodes: N1, N2
- Cluster2 edge nodes: N3, N4
- Cluster3 edge nodes: N5, N6

Scenario Description:

- Segment is directly connected to a tier-1 gateway.

- Tier-1 gateway is connected to a tier-0 gateway.
- DHCP server profile is specified on both tier-1 and tier-0 gateways.
- DHCP profile on tier-1 gateway uses Cluster1.
- DHCP profile on tier-0 gateway uses Cluster2.
- No edge cluster is selected in tier-1 gateway.
- Tier-0 gateway uses Cluster3.
- Segment is configured to use a Gateway DHCP server.

In this scenario, because the tier-1 gateway has no edge cluster specified, NSX-T Data Center falls back on the edge cluster of the connected tier-0 gateway. DHCP service is created in the edge cluster of tier-0 gateway (Cluster3). Any two edge nodes from this edge cluster are autoallocated for creating the DHCP service and configuring DHCP HA.

To summarize, the main points of this scenario are as follows:

- When a tier-1 gateway has no edge cluster specified, NSX-T Data Center falls back on the edge cluster of the connected tier-0 gateway to create the DHCP service.
- If no edge cluster is detected in the tier-0 gateway, DHCP service is created in the edge cluster of the tier-1 gateway DHCP profile.

Scenarios: Impact of Changing Segment Connectivity on DHCP

After you save a segment with DHCP configuration, you must be careful about changing the connectivity of the segment.

Segment connectivity changes are allowed only when the segments and gateways belong to the same transport zone.

The following scenarios explain the segment connectivity changes that are allowed or disallowed, and whether DHCP is impacted in each of these scenarios.

Scenario 1: Move a Routed Segment with Gateway DHCP Server to a Different Gateway

Consider that you have added a segment and connected it either to a tier-0 or tier-1 gateway. You configured Gateway DHCP server on this segment, saved the segment, and connected workloads to this segment. DHCP service is now used by the workloads on this segment.

Later, you decide to change the connectivity of this segment to another tier-0 or tier-1 gateway, which is in the same transport zone.

- Starting in NSX-T Data Center 3.0.2, this change is allowed. However, when you save the segment, an information message alerts you that changing the gateway connectivity impacts the existing DHCP leases, which are assigned to the workloads.

- In NSX-T Data Center 3.0 and 3.0.1, you cannot change the connectivity of the segment from one gateway to another gateway when the segment uses a Gateway DHCP server. Use the following steps in the workaround:

Workaround (only for versions 3.0 and 3.0.1):

- 1 Temporarily disconnect the existing segment from the gateway, or delete the segment. Temporary disconnection of the segment is supported only with the API. Follow these steps:
 - a Retrieve the segment details by running the following GET API:


```
GET https://{NSXManager_IP}/policy/api/v1/infra/segments/{segment-id}
```

Replace *segment-id* with the actual ID of the segment that you want to disconnect from the gateway.
 - b Observe that the `advanced_config` section in the API output shows `connectivity:"ON"`.
 - c Copy the GET API output in a text file and edit `connectivity` to OFF. Paste the complete API output in the request body of the following PATCH API:


```
PATCH https://{NSXManager_IP}/policy/api/v1/infra/segments/{segment-id}
```
 - d Run the PATCH API to disconnect the segment.
- 2 Add a new segment.
- 3 Connect this new segment to the gateway of your choice.

Scenario 2: Move a Routed Segment with Local DHCP Server or Relay to a Different Gateway

Consider that you have added a segment and connected it either to a tier-0 or tier-1 gateway. You configured local DHCP server or DHCP relay on this segment, saved the segment, and connected workloads to this segment. DHCP service is now used by the workloads on this segment.

Later, you decide to change the connectivity of this segment to another tier-0 or tier-1 gateway, which is in the same transport zone. This change is allowed. As the DHCP server is local to the segment, the DHCP configuration settings, including ranges, static bindings, and DHCP options are retained on the segment. The DHCP leases of the workloads are retained and there is no loss of network connectivity.

After the segment is moved to a new gateway, you can continue to update the DHCP configuration settings, and other segment properties.

- If you are using NSX-T Data Center 3.0 or 3.0.1, you cannot change the DHCP type and DHCP profile of a routed segment after moving the segment to a different gateway. For example, you cannot change the DHCP type from a local DHCP server or a DHCP relay to a Gateway DHCP server. In addition, you cannot select a different DHCP server profile or relay profile in the segment. But, you can edit the properties of the DHCP profile, if needed.
- Starting in version 3.0.2, you can change the DHCP type and DHCP profile of a routed segment after moving the segment to a different gateway.

Scenario 3: Move a Standalone Segment with Local DHCP Server to a Tier-0 or Tier-1 Gateway

Consider that you have added a segment with None connectivity in your network. You have configured local DHCP server on this segment, saved the segment, and connected workloads to this segment. DHCP service is now used by the workloads on this segment.

Later, you decide to connect this segment either to a tier-0 or tier-1 gateway, which is in the same transport zone. This change is allowed. As a local DHCP server existed on the segment, the DHCP configuration settings, including ranges, static bindings, and DHCP options are retained on the segment. The DHCP leases of the workloads are retained and there is no loss of network connectivity.

After the segment is connected to the gateway, you can continue to update the DHCP configuration settings, and other segment properties. However, you cannot select a different DHCP type and the DHCP profile in the segment. For example, you cannot change the DHCP type from a local DHCP server to a Gateway DHCP server or a DHCP relay. In addition, you cannot change the DHCP server profile in the segment. But, you can edit the properties of the DHCP profile, if needed.

Scenario 4: Move a Standalone Segment Without DHCP Configuration to a Tier-0 or Tier-1 Gateway

Consider that you have added a segment with None connectivity in your network. You have not configured DHCP on this segment, saved the segment, and connected workloads to this segment.

Later, you decide to connect this segment either to a tier-0 or tier-1 gateway, which is in the same transport zone. This change is allowed. As no DHCP configuration existed on the segment, the segment automatically uses the Gateway DHCP server after it is connected to the gateway. The DHCP profile attached to this gateway gets autoselected in the segment.

Now, you can specify the DHCP configuration settings, including ranges, static bindings, and DHCP options on the segment. You can also edit the other segment properties, if necessary. However, you cannot change the DHCP type from a Gateway DHCP server to a local DHCP server or a DHCP relay.

Remember, you can configure only a Gateway DHCPv4 server on the segment. In NSX-T Data Center 3.0, Gateway DHCPv6 server is not supported.

Scenario 5: Move a Segment with Tier-0 or Tier-1 Connectivity to None Connectivity

Consider that you have added a segment to a tier-0 or tier-1 gateway in your network. You have configured Gateway DHCP server or DHCP relay on this segment, saved the segment, and connected workloads to this segment. DHCP service is now used by the workloads on this segment.

Later, you decide to change the connectivity of this segment to None. This change is not allowed.

In this scenario, the following workaround can help:

- 1 Temporarily disconnect the existing segment from the gateway or delete the segment.
For information about temporarily disconnecting a segment from the gateway, see Scenario 1.
- 2 Add a new segment with a None connectivity.
- 3 Configure a local DHCP server on this standalone segment, if needed.

Add an IP Address Pool

You can configure IP address pools for use by components such as DHCP.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > IP Management > IP Address Pools**.
- 3 Click **Add IP Address Pool**.
- 4 Enter a name and optionally a description.
- 5 Click **Set** in the **Subnets** column to add subnets.
- 6 To specify an address block, select **Add Subnet > IP Block**.
 - a Select an IP block.
 - b Specify a size.
 - c Click the **Auto Assign Gateway** toggle to enable or disable automatic gateway IP assignment.
 - d Click **Add**.

7 To specify IP ranges, select **Add Subnet > IP Ranges**.

- a Enter IPv4 or IPv6 IP ranges.
- b Enter IP ranges in CIDR format.
- c Enter an address for **Gateway IP**.
- d Click **Add**.

8 Click **Save**.

Note: After you add an IP address pool and IP addresses have been allocated from the pool, you will not be able to delete the pool. If you want to expand the pool, you must add new IP ranges. For example, if the existing range is 192.168.1.11 - 192.168.1.20 and you want to expand it to 192.168.1.10 - 192.168.1.30, add the following two ranges:

- 192.168.1.10 - 192.168.1.10
- 192.168.1.21 - 192.168.1.30

Add an IP Address Block

You can configure IP address blocks for use by other components.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > IP Management > IP Address Pools**.
- 3 Click the **IP Address Blocks** tab.
- 4 Click **Add IP Address Block**.
- 5 Enter a name and optionally a description.
- 6 Enter an IP block in CIDR format.
- 7 Click **Save**.

Networking Settings

12

You can configure networking settings for IPv6, VNI (Virtual Network Identifier) pools, gateways, multicast, and BFD (Bidirectional Forwarding Detection).

Read the following topics next:

- [Configuring Multicast](#)
- [Add a VNI Pool](#)
- [Configure Gateway Settings](#)
- [Add a Gateway QoS Profile](#)
- [Add a BFD Profile](#)

Configuring Multicast

You can configure multicast on a tier-0 gateway for an IPv4 network to send the same multicast data to a group of recipients. In a multicast environment, any host, regardless of whether it is a member of a group, can send to a group. However, only the members of a group will receive packets sent to that group.

The multicast feature has the following capabilities and limitations:

- PIM Sparse Mode with IGMPv2.
- No Rendezvous Point (RP) or Bootstrap Router (BSR) functionality on NSX-T. However, RP information can be learned via PIM Bootstrap Messages (BSMs). In addition, a Static RP can be configured.

When a Static RP is configured, it serves as the RP for all multicast groups (224/4). If candidate RPs learned from BSMs advertise candidacy for the same group range, the Static RP is preferred. However, if candidate RPs advertise candidacy for a specific group or range of groups, they are preferred as the RP for those groups.

- The Reverse Path Forwarding (RPF) check for all multicast-specific IPs (senders of data traffic, BSRs, RPs) requires that a route to each of them exists. In NSX-T Data Center 3.0.0, reachability via the default route is not supported. Starting with NSX-T Data Center 3.0.1, reachability via the default route is also supported.

- The RPF check requires a route to each multicast-specific IP with an IP address as the next hop. Reachability via device routes, where the next hop is an interface index, is not supported.
- Tier-0 gateway only.
- Supported on only one uplink on a tier-0 gateway.
- Active-Cold Standby only is supported.
- The NSX Edge cluster can be in active-active or active-standby mode. When the cluster is in active-active mode, two of the cluster members will run multicast in active-cold standby mode. You can run the CLI command `get mcast high-availability role` on each Edge to identify the two nodes participating in multicast. Also note that since unicast reachability to NSX-T in an active-active cluster is via ECMP, it is imperative that the northbound PIM router selects the ECMP path that matches a PIM neighbor to send PIM Join/Prune messages to NSX-T. In this way it will select the active Edge which is running PIM.
- East-west multicast replication: up to 4 VTEP segments for maximum replication efficiency.
- ESXi host and NSX Edge only (KVM not supported).
- Layer 2 bridge attached to a downlink segment not supported.
- Edge Firewall services are not supported for multicast.
- Multi-site (federation) not supported.
- Multi-VRF not supported.

Multicast Configuration Prerequisites

Underlay network configurations:

- Acquire a multicast address range from your network administrator. This will be used to configure the Multicast Replication Range when you configure multicast on a tier-0 gateway (see [Configure Multicast](#)).
- Enable IGMP snooping on the layer 2 switches to which GENEVE participating hosts are attached. If IGMP snooping is enabled on layer 2, IGMP querier must be enabled on the router or layer 3 switch with connectivity to multicast enabled networks.

Multicast Configuration Steps

- 1 Create an IGMP profile. See [Create an IGMP Profile](#).
- 2 Optionally create a PIM profile to configure a Static Rendezvous Point (RP). See [Create a PIM Profile](#).
- 3 Configure a tier-0 gateway to support multicast. See [Add a Tier-0 Gateway and Configure Multicast](#)

Create an IGMP Profile

Internet Group Management Protocol (IGMP) is a multicast protocol used in IPv4 networks.

Note that the IGMP snooping timeout for reports is 2 times the general query timeout. By default, the IGMP snooping timeout value is 120 seconds. On ESXi, the default IGMP snooping timeout value is 60 seconds.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Networking Settings**.
- 3 Click the **Multicast Profiles** tab.
- 4 Click **Add IGMP Profile**.
- 5 Enter a profile name and the following profile details.

Option	Description
Query Interval (seconds)	Interval between general query messages. A larger value causes IGMP queries to be sent less often. Default: 30. Range: 1 - 1800.
Query Max Response Time (seconds)	Maximum allowed time before sending a response to a membership query message. Default: 10. Range: 1 - 25.
Last Member Query Interval (seconds)	Maximum amount of time between group-specific query messages, including those sent in response to leave-group messages. Default: 10. Range: 1 - 25.
Robustness Variable	Number of IGMP query messages sent. This helps alleviate the risk of loss of packets in a busy network. A larger number is recommended in a network with high traffic. Default: 2. Range: 1 - 255.

Create a PIM Profile

Protocol Independent Multicast (PIM) is a collection of multicast routing protocols for IP networks. It is not dependent on a specific unicast routing protocol and can leverage whichever unicast routing protocols are used to populate the unicast routing table.

This step is optional. It is needed only if you want to configure a Static Rendezvous Point (RP). A Rendezvous Point is a router in a multicast network domain that acts as a shared root for a multicast shared tree. If a Static RP is configured, it is preferred over the RPs that are learned from the elected Bootstrap Router (BSR).

Procedure

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

- 3 Click the **Multicast Profiles** tab.
- 4 In the **Select Profile type** drop-down menu, select **PIM Profiles**.
- 5 Click **Add PIM Profile** tab.
- 6 Enter a profile name.
- 7 Enter a Static Rendezvous Point (RP) address

Add a VNI Pool

You can create a VNI pool to be used when you configure EVPN for a tier-0 gateway. VNI pools cannot have values that overlap.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Networking Settings**.
- 3 Click the **VNI Pool** tab.
- 4 Click **Add VNI Pool**.
- 5 Enter a name for the pool.
- 6 Enter a start value.
The value must be from 75001 to 16777215.
- 7 Enter an end value.
The value must be from 75001 to 16777215.
- 8 Click **Save**.

Configure Gateway Settings

Set a global configuration for the layer 3 forwarding mode and the Maximum Transmission unit (MTU). IPv4 layer 3 forwarding is enabled by default. You can also configure IPv6 layer 3 forwarding.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Networking Settings**.
- 3 Click the **Global Networking Config** tab.

4 Click **Edit** next to **Global Gateway Configuration**.

a Enter a value for the **Gateway Interface MTU**.

The default value is 1500.

b Select the layer 3 forwarding mode.

5 Click **Save**.

Add a Gateway QoS Profile

Create a QoS profile for your tier-1 gateways to define limits on the traffic rates. You can specify the permitted information rate and the burst size to set the limitations. Any traffic that does not conform to the QoS policy, is dropped. QoS profiles can be set for both ingress and egress traffic, for all traffic types (unicast, BUM, IPv4/IPv6). You can choose to create a different profile for each tier-1 gateway.

Note Gateway QoS profile is supported only on tier-1 gateways.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Networking Settings**.
- 3 Click the **Gateway QoS Profiles** tab.
- 4 Click **Add Gateway QoS Profiles**.
- 5 Enter a name for the profile.
- 6 Enter the committed bandwidth limit that you want to set for the traffic.
- 7 Enter the burst size. Use the following guidelines for burst size.
 - B is the burst size in bytes.
 - R is the committed rate (or bandwidth) in Mbps.
 - I is the time interval in milliseconds, to refill or withdraw tokens(in bytes) from the token bucket. Use the `get dataplane` command from the NSX Edge CLI to retrieve the time interval, `Qos_wakeup_interval_ms`. The default value for `Qos_wakeup_interval_ms` is 50ms. However, this value is automatically adjusted by the dataplane based on the QoS configuration.

The constraints for burst size are:

- $B \geq R * 1000,000 * I / 1000 / 8$ because burst size is the maximum amount of tokens that can be refilled in each interval.
- $B \geq R * 1000,000 * 1 / 1000 / 8$ because the minimum value for I is 1 ms, taking into account dataplane CPU usage among other constraints.

- $B \geq \text{MTU of SR port}$ because at least the MTU-size amount of tokens need to be present in the token bucket for an MTU-size packet to pass rate-limiting check.

Since the burst size needs to satisfy all three constraints, the configured value of burst size would be:

$$\text{Max } (R * 1000,000 * I / 1000 / 8, R * 1000,000 * 1 / 1000 / 8, \text{MTU})$$

For example, if $R = 100$ Mbps, $I = 50$ ms, and $\text{MTU} = 1500$, then

$$B \geq \text{max } (100 * 1000,000 * 50 / 1000 / 8, 100 * 1000,000 * 1 / 1000 / 8, 1500) = 625000 \text{ in bytes}$$

- 8 Click **Save**.

Add a BFD Profile

BFD (Bidirectional Forwarding Detection) is a protocol that can detect forwarding path failures. You can create a BFD profile for your Tier-0 static routes.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Networking Settings**.
- 3 Click the **BFD Profile** tab.
- 4 Click **Add BFD Profile**.
- 5 Enter a name for the profile.
- 6 Enter values for the heartbeat **Interval** and **Declare Dead Multiple**.
- 7 Click **Save**.

The topics in this section cover north-south and east-west security for distributed firewall rules, identity firewall, network introspection, gateway firewall, and endpoint protection policies.

Read the following topics next:

- [Security Configuration Overview](#)
- [Security Overview](#)
- [Security Terminology](#)
- [Identity Firewall](#)
- [Layer 7 Context Profile](#)
- [Distributed Firewall](#)
- [Distributed IDS](#)
- [East-West Network Security - Chaining Third-party Services](#)
- [Gateway Firewall](#)
- [North-South Network Security - Inserting Third-party Service](#)
- [Endpoint Protection](#)
- [Security Profiles](#)
- [Time-Based Firewall Policy](#)
- [Network Introspection Settings](#)
- [Troubleshooting Firewall](#)
- [Bare Metal Server Security](#)

Security Configuration Overview

Configure east-west and north-south firewall policies under predefined categories for your environment.

Distributed Firewall (east-west) and Gateway Firewall (north-south) offer multiple sets of configurable rules divided by categories. You can configure an exclusion list that contains logical switches, logical ports, or groups, to be excluded from firewall enforcement.

Security policies are enforced as follows:

- Rules are processed in categories, left to right.
- Rules are processed in top-to-bottom ordering.
- Each packet is checked against the top rule in the rule table before moving down the subsequent rules in the table.
- The first rule in the table that matches the traffic parameters is enforced.

No subsequent rules can be enforced as the search is then terminated for that packet. Because of this behavior, it is always recommended to put the most granular policies at the top of the rule table. This ensures they will be enforced before more generic rules.

Whether an east-west or north-south firewall fails close or fails open upon failure depends on the last rule in the firewall. To ensure that a firewall fails close upon failure, configure the last rule to reject or drop all packets.

Security Overview

The security overview dashboard has three tabs: Insights, Configuration, and Capacity.

The **Insights** tab shows details for:

- URL Analysis:
 - The number of gateways with URL Filtering Enabled.
 - The number of gateways connected to the cloud service, and if connectivity to the cloud service is up.
 - The latest signature pack available on the cloud service, and what gateways are up to date.
 - Information for the top five URL categories, and the URLs accessed in each category.

- Intrusion Detection Summary:

Entry	Description
Total Intrusion Events	Displays the total number of intrusion events as a clickable link, and number in each severity category. For more information, see Distributed IDS Events .
Trending by Severity	Displays a graph with the number of intrusion events by time.
Top VMs by Intrusion Events or by Vulnerability Severity	Click the arrow to select the shown data.

- Distributed Firewall Rule Utilization:
 - Number of identity firewall rules.

- Number of Layer 7 rules.
- Number of compute rules.
- Number of rules with a combination of Layer 7 and IDFW.
- A summary of the configuration of endpoint protection for virtual machines. You can view components having issues, and virtual machine distribution by service profile.

The **Configuration** tab has clickable links with the number of:

- Distributed FW Policies
- Gateway Policies
- Endpoint Policies
- Network Introspection EW Policies
- Network Introspection NS Policies
- Distributed IDS Policies

You can also view details of your distributed firewall policies, along with the count per category.

The **Capacity** tab is not available in Policy view.

Security Terminology

The following terms are used throughout distributed firewall.

Table 13-1. Security-Related Terminology

Construct	Definition
Policy	A security policy includes various security elements including firewall rules and service configurations. Policy was previously called a firewall section.
Rule	A set of parameters with which flows are evaluated against, and define which actions will be taken upon a match. Rules include parameters such as source and destination, service, context profile, logging, and tags.
Group	Groups include different objects that are added both statically and dynamically, and can be used as the source and destination field of a firewall rule. Groups can be configured to contain a combination of virtual machines, IP sets, MAC sets, logical ports, logical switches, AD user groups, and other nested groups. Dynamic inclusion of groups can be based on tag, machine name, OS name, or computer name. When you create a group, you must include a domain that it belongs to, and by default this is the default domain. Groups were previously called NSGroup or security group.
Service	Defines a combination of port and protocol. Used to classify traffic based on port and protocol. Pre-defined services and user-defined services can be used in firewall rules.
Context Profile	Defines context aware attributes including APP-ID and domain name. Also includes sub attributes such as application version, or cipher set. Firewall rules can include a context profile to enable Layer-7 firewall rules.

Identity Firewall

With Identity Firewall (IDFW) features an NSX administrator can create Active Directory user-based Distributed Firewall (DFW) rules.

IDFW can be used for Virtual Desktops (VDI) or Remote desktop sessions (RDSH support), enabling simultaneous log ins by multiple users, user application access based on requirements, and the ability to maintain independent user environments. VDI management systems control what users are granted access to the VDI virtual machines. NSX-T controls access to the destination servers from the source virtual machine (VM), which has IDFW enabled. With RDSH, administrators create security groups with different users in Active Directory (AD), and allow or deny those users access to an application server based on their role. For example, Human Resources and Engineering can connect to the same RDSH server, and have access to different applications from that server.

IDFW can also be used on VMs that have supported operating systems. See [Identity Firewall Supported Configurations](#).

A high level overview of the IDFW configuration workflow begins with preparing the infrastructure. Preparation includes the administrator installing the host preparation components on each protected cluster, and setting up Active Directory synchronization so that NSX can consume AD users and groups. Next, IDFW must know which desktop an Active Directory user logs in to and apply IDFW rules. When network events are generated by a user, the thin agent installed with VMware Tools on the VM gathers and forwards the information, and sends it to the context engine. This information is used to provide enforcement for the distributed firewall.

IDFW processes the user identity at the source only in distributed firewall rules. Identity-based groups cannot be used as the destination in DFW rules.

Note IDFW relies on the security and integrity of the guest operating system. There are multiple methods for a malicious local administrator to spoof their identity to bypass firewall rules. User identity information is provided by the NSX Guest Introspection Thin Agent inside guest VMs. Security administrators must ensure that thin agent is installed and running in each guest VM. Logged-in users should not have the privilege to remove or stop the agent.

For supported IDFW configurations see [Identity Firewall Supported Configurations](#).

IDFW workflow:

- 1 A user logs in to a VM and starts a network connection, by opening Skype or Outlook.
- 2 A user login event is detected by the Thin Agent, which gathers connection information and identity information and sends it to the context engine.
- 3 The context engine forwards the connection and the identity information to Distributed Firewall rule for any applicable rule enforcement.

Identity Firewall Workflow

IDFW enhances traditional firewall by allowing firewall rules based on user identity. For example, administrators can allow or disallow customer support staff to access an HR database with a single firewall policy.

Identity based firewall rules are determined by membership in an Active Directory (AD) group membership. See [Identity Firewall Supported Configurations](#).

IDFW processes the user identity at the source only in distributed firewall rules. Identity-based groups cannot be used as the destination in DFW rules.

Note For Identity Firewall rule enforcement, Windows Time service should be **on** for all VMs using Active Directory. This ensures that the date and time is synchronized between Active Directory and VMs. AD group membership changes, including enabling and deleting users, do not immediately take effect for logged in users. For changes to take effect, users must log out and then log back in. AD administrator's should force a logout when group membership is modified. This behavior is a limitation of Active Directory.

Prerequisites

If Windows auto-logon is enabled on VMs, go to **Local Computer Policy > Computer configuration > Administrative Templates > System > Logon** and enable **Always wait for the network at computer startup and logon**.

For supported IDFW configurations see [Identity Firewall Supported Configurations](#).

Procedure

- 1 Enable NSX File Introspection driver and NSX Network Introspection driver. VMware Tools full installation adds these by default.
- 2 Enable IDFW on cluster or standalone host: [Enable Identity Firewall](#).
- 3 Configure Active Directory domain: [Add an Active Directory](#).
- 4 Configure Active Directory sync operations: [Synchronize Active Directory](#).
- 5 Create security groups (SG) with Active Directory group members: [Add a Group](#).
- 6 Assign SG with AD group members to a distributed firewall rule: [Add a Distributed Firewall](#).

Enable Identity Firewall

Identity Firewall must be enabled for IDFW firewall rules to take effect.

Procedure

- 1 Select **Security > Distributed Firewall**.
- 2 In the right corner, click **Actions > General Setting**.
- 3 Toggle the status button to enable IDFW.

Distributed firewall must also be enabled for IDFW to work.

- 4 To enable IDFW on standalone hosts or clusters, select the **Identity Firewall Settings** tab.
- 5 Toggle the **Enable** bar, and select the standalone hosts, or select the cluster where the IDFW host must be enabled.
- 6 Click **Save**.

Identity Firewall Best Practices

The following best practices will help maximize the success of identity firewall rules.

- IDFW supports the following protocols:
 - Single user (VDI, or Non-RDSH Server) use case support - TCP, UDP, ICMP
 - Multi-User (RDSH) use case support - TCP, UDP
- A single ID-based group can be used as the source only within a distributed firewall rule. If IP and ID-based groups are needed at the source, create two separate firewall rules.
- Any change on a domain, including a domain name change, will trigger a full sync with Active Directory. Because a full sync can take a long time, we recommend syncing during off-peak or non-business hours.
- For local domain controllers, the default LDAP port 389 and LDAPS port 636 are used for the Active Directory sync, and should not be edited from the default values.

Identity Firewall Supported Configurations

The following configurations are supported for IDFW on virtual machines (VMs). IDFW for physical devices is not supported.

IDFW supports the following protocols:

- Single user (VDI, or Non-RDSH Server) use case support - TCP, UDP, ICMP
- Multi-User (RDSH) use case support - TCP, UDP

Guest Operating Systems	Enforcement Type
Windows 8	Desktop - supports desktop users use case
Windows 10	Desktop - supports desktop users use case
Windows 2012	Server - supports server users use case
Windows 2012R2	Server - supports server users use case
Windows 2016	Server - supports server users use case
Windows 2012R2	RDSH - supports Remote Desktop Session Host
Windows 2016	RDSH - supports Remote Desktop Session Host

Active Directory Domain Controllers:

- Windows Server 2012
- Windows Server 2012R2
- Windows Server 2016
- Windows Server 2019

Host operating system: ESXi

VMware Tools - For supported versions of VMware Tools, see the [VMware Product Interoperability Matrices](#).

- VMCI Driver
- NSX File Introspection Driver
- NSX Network Introspection Driver

Layer 7 Context Profile

Layer 7 App IDs are configured as part of a context profile.

A context profile can specify one or more [Attributes](#), and can also include sub-attributes, for use in distributed firewall (DFW) rules and gateway firewall rules. When a sub-attribute, such as TLS version 1.2 is defined, multiple application identity attributes are not supported. In addition to attributes, DFW also supports a Fully Qualified Domain Name (FQDN) or URL that can be specified in a context profile for FQDN allowlisting or denylisting. See [Filtering Specific Domains \(FQDN/URLs\)](#) for more information. FQDNs can be configured with an attribute in a context profile, or each can be set in different context profiles. After a context profile has been defined, it can be applied to one or more distributed firewall rules.

Note

- Gateway firewall rules do not support the use of FQDN attributes or other sub attributes in context profiles.
 - Context profiles are not supported on tier-0 gateway firewall policy.
-

When a context profile is used in a rule, any traffic coming in from a virtual machine is matched against the rule table based on the 5-tuple. If the rule matching the flow also includes a Layer 7 context profile, the packet is redirected to a user-space component called the vDPI engine. Subsequent packets are sent to the vDPI engine for each flow. After the App ID has been determined, the information is stored in the in-kernel context table. When the next packet for the flow comes in the information in the context table is compared with the rule table again, and is matched on 5-tuple and the Layer 7 App ID. The appropriate action as defined in the fully matched rule is taken, and if there is an ALLOW rule, all subsequent packets for the flow are processed in the kernel, and matched against the connection table. For fully matched DROP rules, a reject packet is generated. Logs generated by the firewall will include the Layer 7 App ID and applicable URL, if that flow was sent to the vDPI engine.

Rule processing for an incoming packet:

- 1 Upon entering a DFW or Gateway filter, packets are looked up in the flow table based on 5-tuple.
- 2 If no flow/state is found, the flow is matched against the rule-table based on 5-tuple and an entry is created in the flow table.
- 3 If the flow matches a rule with a Layer 7 service object, the flow table state is marked as “DPI In Progress.”
- 4 The traffic is then punted to the DPI engine. The DPI Engine determines the App ID.
- 5 After the App ID has been determined, the DPI Engine sends down the attribute which is inserted into the context table for this flow. The "DPI In Progress" flag is removed, and traffic is no longer punted to the DPI engine.
- 6 The flow (now with App ID) is reevaluated against all rules that match the App ID, starting with the original rule that was matched based on 5-tuple, and the first fully matched L4/L7 rule is picked up. The appropriate action is taken (allow/deny/reject) and the flow table entry is updated accordingly.

Layer 7 Firewall Rule Workflow

Layer 7 App Ids are used in creating context profiles, which are used in distributed firewall rules or gateway firewall rules. Rule enforcement based on attributes enables users to allow or deny applications to run on any port.

NSX-T provides built in [Attributes](#) for common infrastructure and enterprise applications. App Ids include versions (SSL/TLS and CIFS/SMB) and Cipher Suite (SSL/TLS). For distributed firewall, App Ids are used in rules through context profiles, and can be combined with FQDN allowlisting and denylisting. App Ids are supported on ESXi and KVM hosts.

Note

- Gateway firewall rules do not support the use of FQDN attributes or other sub attributes in context profiles.
 - Context profiles are not supported on tier-0 gateway firewall policy.
-

Supported App IDs and FQDNs:

- For FQDN, users need to configure a high priority rule with a DNS App ID for the specified DNS servers on port 53.
- ALG App IDs (FTP, ORACLE, DCERPC, TFTP), require the corresponding ALG service for the firewall rule.
- SYSLOG App ID is detected only on standard ports.

KVM Supported App IDs and FQDNs:

- Sub attributes are not supported on KVM.

- FTP and TFTP ALG App IDs are supported on KVM.

Note that if you are using a combination of Layer 7 and ICMP, or any other protocols you need to put the Layer 7 firewall rules last. Any rules after a Layer 7 any/any rule will not be executed.

Procedure

- 1 Create a custom context profile: [Add a Context Profile](#).
- 2 Use the context profile in a distributed firewall rule, or a gateway firewall rule: [Add a Distributed Firewall](#) or [Add a Gateway Firewall Policy and Rule](#).

Multiple App Id context profiles can be used in a firewall rule with services set to **Any**. For ALG profiles (FTP, ORACLE, DCERPC, TFTP), one context profile is supported per rule.

Attributes

Layer 7 attributes (App Ids) identify which application a particular packet or flow is generated by, independent of the port that is being used.

Enforcement based on App Ids enable users to allow or deny applications to run on any port, or to force applications to run on their standard port. vDPI enables matching packet payload against defined patterns, commonly referred to as signatures. Signature-based identification and enforcement, enables customers not just to match the particular application/protocol a flow belongs to, but also the version of that protocol, for example TLS version 1.0 version TLS version 1.2 or different versions of CIFS traffic. This allows customers to get visibility into or restrict the use of protocols that have known vulnerabilities for all deployed applications and their E-W flows within the datacenter.

Layer 7 App Ids are used in context profiles in distributed firewall and gateway firewall rules, and are supported on ESXi and KVM hosts.

Note NFS version 4 is not a supported attribute.

Gateway firewall rules do not support the use of FQDN attributes or other sub attributes in context profiles.

Context profiles are not supported on tier-0 gateway firewall policy. Gateway firewall rules do not support the use of FQDN attributes, or other sub attributes.

Supported App IDs and FQDNs:

- For FQDN, users need to configure a high priority rule with a DNS App ID for the specified DNS servers on port 53.
- ALG App IDs (FTP, ORACLE, DCERPC, TFTP), require the corresponding ALG service for the firewall rule.
- SYSLOG App ID is detected only on standard ports.

KVM Supported App IDs and FQDNs:

- Sub attributes are not supported on KVM.

- FTP and TFTP ALG App IDs are supported on KVM.

Attribute (App Id)	Description	Type
360ANTIV	360 Safeguard is a program developed by Qihoo 360, an IT company based in China	Web Services
ACTIVDIR	Microsoft Active Directory	Networking
AMQP	Advanced Messaging Queuing Protocol is application layer protocol which supports business message communication between applications or organizations	Networking
AVAST	Traffic generated by browsing Avast.com official website of Avast! Antivirus downloads	Web Services
AVG	AVG Antivirus/Security software download and updates	File Transfer
AVIRA	Avira Antivirus/Security software download and updates	File Transfer
BLAST	A remote access protocol that compresses, encrypts, and encodes a computing experiences at a data center and transmits it across any standard IP network for VMware Horizon desktops.	Remote Access
BDEFENDER	BitDefender Antivirus/Security software download and updates.	File Transfer
CA_CERT	Certification authority (CA) issues digital certificates which certifies the ownership of a public key for message encryption	Networking
CIFS	CIFS (Common Internet File System) is used to provide shared access to directories, files, printers, serial ports, and miscellaneous communications between nodes on a network	File Transfer
CLDAP	Connectionless Lightweight Directory Access Protocol is an application protocol for accessing and maintaining distributed directory information services over an Internet Protocol (IP) network using UDP.	Networking
CTRXCGP	Citrix Common Gateway Protocol is an application protocol for accessing and maintaining distributed directory information services over an Internet Protocol (IP) network using UDP.	Database
CTRKGOTO	Hosting Citrix GoToMeeting, or similar sessions based on the GoToMeeting platform. Includes voice, video, and limited crowd management functions	Collaboration
CTRXICA	ICA (Independent Computing Architecture) is a proprietary protocol for an application server system, designed by Citrix Systems	Remote Access
DCERPC	Distributed Computing Environment / Remote Procedure Calls, is the remote procedure call system developed for the Distributed Computing Environment (DCE)	Networking
DIAMETER	An authentication, authorization, and accounting protocol for computer networks	Networking
DHCP	Dynamic Host Configuration Protocol is a protocol used management for the distribution of IP addresses within a network	Networking

Attribute (App Id)	Description	Type
DNS	Querying a DNS server over TCP or UDP	Networking
EPIC	Epic EMR is an electronic medical records application that provides patient care and healthcare information.	Client Server
ESET	Eset Antivirus/Security software download and updates	File Transfer
FPROT	F-Prot Antivirus/Security software download and updates	File Transfer
FTP	FTP (File Transfer Protocol) is used to transfer files from a file server to a local machine	File Transfer
GITHUB	Web-based Git or version control repository and Internet hosting service	Collaboration
HTTP	(HyperText Transfer Protocol) the principal transport protocol for the World Wide Web	Web Services
HTTP2	Traffic generated by browsing websites that support the HTTP 2.0 protocol	Web Services
IMAP	IMAP (Internet Message Access Protocol) is an Internet standard protocol for accessing email on a remote server	Mail
KASPRSKY	Kaspersky Antivirus/Security software download and updates	File Transfer
KERBEROS	Kerberos is a network authentication protocol designed to provide strong authentication for client/server applications by using secret-key cryptography	Networking
LDAP	LDAP (Lightweight Directory Access Protocol) is a protocol for reading and editing directories over an IP network	Database
MAXDB	SQL connections and queries made to a MaxDB SQL server	Database
MCAFEE	McAfee Antivirus/Security software download and updates	File Transfer
MSSQL	Microsoft SQL Server is a relational database.	Database
NFS	Allows a user on a client computer to access files over a network in a manner similar to how local storage is accessed. Note NFS version 4 is not a supported attribute.	File Transfer
NNTP	An Internet application protocol used for transporting Usenet news articles (netnews) between news servers, and for reading and posting articles by end user client applications.	File Transfer
NTBIOSNS	NetBIOS Name Service. In order to start sessions or distribute datagrams, an application must register its NetBIOS name using the name service	Networking
NTP	NTP (Network Time Protocol) is used for synchronizing the clocks of computer systems over the network	Networking
OCSP	An OCSP Responder verifying that a user's private key has not been compromised or revoked	Networking
ORACLE	An object-relational database management system (ORDBMS) produced and marketed by Oracle Corporation.	Database

Attribute (App Id)	Description	Type
PANDA	Panda Security Antivirus/Security software download and updates.	File Transfer
PCOIP	A remote access protocol that compresses, encrypts, and encodes a computing experiences at a data center and transmits it across any standard IP network.	Remote Access
POP2	POP (Post Office Protocol) is a protocol used by local e-mail clients to retrieve e-mail from a remote server.	Mail
POP3	Microsoft's implementation of NetBIOS Name Service (NBNS), a name server and service for NetBIOS computer names.	Mail
RADIUS	Provides centralized Authentication, Authorization, and Accounting (AAA) management for computers to connect and use a network service	Networking
RDP	RDP (Remote Desktop Protocol) provides users with a graphical interface to another computer	Remote Access
RTCP	RTCP (Real-Time Transport Control Protocol) is a sister protocol of the Real-time Transport Protocol (RTP). RTCP provides out-of-band control information for an RTP flow.	Streaming Media
RTP	RTP (Real-Time Transport Protocol) is primarily used to deliver real-time audio and video	Streaming Media
RTSP	RTSP (Real Time Streaming Protocol) is used for establishing and controlling media sessions between end points	Streaming Media
SIP	SIP (Session Initiation Protocol) is a common control protocol for setting up and controlling voice and video calls	Streaming Media
SMTP	SMTP (Simple Mail Transfer Protocol) An Internet standard for electronic mail (e-mail) transmission across Internet Protocol (IP) networks.	Mail
SNMP	SNMP (Simple Network Management Protocol) is an Internet-standard protocol for managing devices on IP networks.	Network Monitoring
SSH	SSH (Secure Shell) is a network protocol that allows data to be exchanged using a secure channel between two networked devices.	Remote Access
SSL	SSL (Secure Sockets Layer) is a cryptographic protocol that provides security over the Internet.	Web Services
SYMUPDAT	Symantec LiveUpdate traffic, this includes spyware definitions, firewall rules, antivirus signature files, and software updates.	File Transfer
SYSLOG	SYSLOG is a protocol that allows network devices to send event messages to a logging server.	Network Monitoring
TELNET	A network protocol used on the Internet or local area networks to provide a bidirectional interactive text-oriented communications facility using a virtual terminal connection.	Remote Access
TFTP	TFTP (Trivial File Transfer Protocol) being used to list, download, and upload files to a TFTP server like SolarWinds TFTP Server, using a client like WinAgents TFTP client.	File Transfer

Attribute (App Id)	Description	Type
VNC	Traffic for Virtual Network Computing.	Remote Access
WINS	Microsoft's implementation of NetBIOS Name Service (NBNS), a name server and service for NetBIOS computer names.	Networking

Distributed Firewall

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 13-2. Distributed Firewall Rule Categories

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

Firewall Drafts

A draft is a complete distributed firewall configuration with policy sections and rules. Drafts can be auto saved or manually saved, and immediately published or saved for publishing at a later date.

To save a manual draft firewall configuration, go to the upper right of the distributed firewall screen and click **Actions > Save**. After saving, the configuration can be viewed by selecting **Actions > View**. Auto drafts are enabled by default. Auto drafts can be disabled by going to **Actions > General Settings**. When auto drafts are enabled, any changes to a firewall configuration results in a system generated autodraft. A maximum of 100 auto drafts and 10 manual drafts can be saved. Auto drafts can be edited and saved as a manual draft, for publishing now or later. To prevent multiple users from opening and editing the draft, manual drafts can be locked. When a draft is published, the current configuration is replaced by the configuration in the draft.

Save or View a Firewall Draft

A draft is a distributed firewall configuration that has been published, or saved for publishing at a later date. Drafts are created automatically, and manually.

Manual drafts can be edited and saved. Auto drafts can be cloned, and saved as manual drafts, and then edited. The maximum number of drafts that can be saved is 100 autodrafts and 10 manual drafts.

Procedure

- 1 Click **Security > Distributed Firewall**.

- 2 To save a firewall configuration manually, go to **Actions > Save**.

A manual draft can be saved, or edited and then saved. After saving, you can revert to the original configuration.

- 3 **Name** the configuration.

- 4 To prevent multiple users from opening and editing a manual draft, **Lock** the configuration, and add a comment.

- 5 Click **Save**.

- 6 To view the saved configuration, click **Actions > View**.

A timeline opens up showing all saved configurations. To see details such as draft name, date, time and who saved it, point to the dot or star icon of any draft. Saved configurations can be filtered by time, showing all drafts in the last one day, one week, 30 days, or the last three months. They can be filtered by autodraft and saved by me. They can also be filtered by name, by using the search tool on the top right.

- 7 Hover over a draft to view name, date and time details of the saved configuration. Click the name to view draft details.

The detailed draft view shows the required changes to be made to the current firewall configuration, in order to be in sync with this draft. If this draft is published, all of the changes visible in this view will be applied to the current configuration.

Clicking the downward arrow expands each section, and displays the added, modified, and deleted changes in each section. The comparison shows added rules with a green bar on the left side of the box, modified elements (such as a name change) have a yellow bar, and deleted elements have a red bar.

- 8 To edit the name or description of a selected draft, click the menu icon (three dots) from the **View Draft Details** window, and select **Edit**.

Manual drafts can be locked. If locked, a comment for the draft must be provided.

Some roles, such as enterprise administrator have full access credentials, and cannot be locked out. See [Role-Based Access Control](#).

- 9 Auto drafts and manual drafts can also be cloned and saved by clicking **Clone**.

In the Saved Configurations window, you can accept the default name, or edit it. You can also lock the configuration. If locked, a comment for the draft must be provided.

- 10 To save the cloned version of the draft configuration, click **Save**. The draft is now present in the Saved Configurations section.

What to do next

After viewing a draft, you can load and publish it. It is then the active firewall configuration.

Publish or Revert a Firewall Draft

Both auto drafts and saved manual drafts can be loaded and published to become the active configuration.

During publishing, a new auto draft is created. This auto draft can be published to revert to the previous configuration.

Procedure

- 1 To view the saved configuration, click **Actions > View**.

A timeline opens up showing all saved configurations. To see details such as draft name, date, time and who saved it, point to the dot icon of any draft. Saved configurations are filtered by time, showing all drafts created in 1 day, 1 week, 30 days, or the last 3 months.

- 2 Click a draft name and the View Draft Details window appears.

- 3 Click **Load**. The new firewall configuration appears on the main window.

Note A draft cannot be loaded if firewall filters are being used, or if there are unsaved changes in the current configuration.

- 4 To commit the draft configuration and make it active, click **Publish**. To return to the previous published configuration, click **Revert**.

After publishing, the changes in the draft will be present in the active configuration.

- 5 To edit the contents of the selected draft before publishing, after clicking **Load**, edit the configuration.

- 6 To save the edited version of the draft configuration, click **Actions > Save**.

Manual drafts can be saved as a new configuration, or an update to the existing configuration. Auto drafts can only be saved as a new configuration.

- 7 Enter a **Name**, and optional **Description**. You can also **Lock** the draft. If locked, a comment for the draft must be provided.

- 8 Click **Save**.

- 9 To commit the draft configuration and make it active, click **Publish**, or to return to the previous published configuration, click **Revert**.

Add a Distributed Firewall

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

Prerequisites

To be DFW-protected, VMs must have their vNIC connected to an NSX overlay or VLAN segment.

If you are creating rules for Identity Firewall, first create a group with Active Directory members. To view supported protocols for IDFW, see [Identity Firewall Supported Configurations](#).

Note For Identity Firewall rule enforcement, Windows Time service should be **on** for all VMs using Active Directory. This ensures that the date and time is synchronized between Active Directory and VMs. AD group membership changes, including enabling and deleting users, do not immediately take effect for logged in users. For changes to take effect, users must log out and then log back in. AD administrator's should force a logout when group membership is modified. This behavior is a limitation of Active Directory.

Note that if you are using a combination of Layer 7 and ICMP, or any other protocols you need to put the Layer 7 firewall rules last. Any rules after a Layer 7 any/any rule will not be executed.

For Federation-specific details on distributed firewall policy and rule creation, see [Create DFW Policies and Rules from Global Manager](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Security > Distributed Firewall** from the navigation panel.
- 3 Ensure that you are in the correct pre-defined category, and click **Add Policy**. For more about categories, see [Distributed Firewall](#).
- 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. 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 (Optional) To configure the following policy settings, click the gear icon:

Option	Description
TCP Strict	A TCP connection begins with a three-way handshake (SYN, SYN-ACK, ACK) and typically ends with a two-way exchange (FIN, ACK). In certain circumstances, the distributed firewall (DFW) might not see the three-way handshake for a particular flow (due to asymmetric traffic or the distributed firewall being enabled while a flow exists). By default, the distributed firewall does not enforce the need to see a three-way handshake, and picks up sessions that are already established. TCP strict can be enabled on a per section basis to turn off mid-session pick-up and enforce the requirement for a three-way handshake. When enabling TCP strict mode for a particular DFW policy, and using a default ANY-ANY Block rule, packets that do not complete the three-way handshake connection requirements and that match a TCP-based rule in this section are dropped. Strict is only applied to stateful TCP rules, and is enabled at the distributed firewall policy level. TCP strict is not enforced for packets that match a default ANY-ANY Allow which has no TCP service specified.
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. See Role-Based Access Control.

- 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. See [Add a Group](#) for more information.
- IPv4, IPv6, and multicast addresses are supported.
- Note: IPv6 firewall must have IP Discovery for IPv6 enabled on a connected segment. For more information, see [Understanding IP Discovery Segment Profile](#).
- 10 In the **column, click the edit icon and select the destination of the rule. If not defined, the destination matches any. See [Add a Group](#) for more information.**
- IPv4, IPv6, and multicast addresses are supported.

- 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 [Add a Context Profile](#).

Context profiles use layer 7 APP ID attributes for use in distributed firewall rules and gateway firewall rules. Multiple App ID context profiles can be used in a firewall rule with services set to **Any**. For ALG profiles (FTP, or TFTP), one context profile is supported per rule.

Context profiles are not supported when creating IDS rules.

- 13 Click **Apply** to apply the context profile to the rule.
- 14 Use **Applied to** to apply the rule to a selected group. By default, the **Applied To** column is set to DFW, and the rule is applied to all workloads. If both policy and the rules within have **Applied to** set to a group, then 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.
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.

- 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.

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.

Distributed Firewall Packet Logs

If logging is enabled for firewall rules, you can look at the firewall packet logs to troubleshoot issues.

The log file is /var/log/dfwpktlogs.log for both ESXi and KVM hosts.

The following is a regular log sample for distributed firewall rules:

```
2018-07-03T19:44:09.749Z b6507827 INET match PASS mainrs/1024 IN 52 TCP 192.168.4.3/49627->192.168.4.4/49153 SEW

2018-07-03T19:46:02.338Z 7396c504 INET match DROP mainrs/1024 OUT 52 TCP 192.168.4.3/49676->192.168.4.4/135 SEW

2018-07-06T18:15:49.647Z 028cd586 INET match DROP mainrs/1027 IN 36 PROTO 2 0.0.0.0->224.0.0.1

2018-07-06T18:19:54.764Z 028cd586 INET6 match DROP mainrs/1027 OUT 143 UDP fe80:0:0:0:68c2:8472:2364:9be/546->ff02:0:0:0:0:1:2/547
```

The elements of a DFW log file format include the following, separated by a space:

- timestamp:
- last eight digits of the VIF ID of the interface
- INET type (v4 or v6)
- reason (match)
- action (PASS, DROP, REJECT)
- rule set name/ rule ID
- packet direction (IN/OUT)
- packet size

- protocol (TCP, UDP, or PROTO #)
- SVM direction for netx rule hit
- source IP address/source port>destination IP address/destination port
- TCP flags (SEW)

For passed TCP packets there is a termination log when the session has ended:

```
2018-07-03T19:44:30.585Z 7396c504 INET TERM mainrs/1024 OUT TCP RST 192.168.4.3/49627-
>192.168.4.4/49153 20/16 1718/76308
```

The elements of a TCP termination log include the following, separated by a space:

- timestamp:
- last 8 digits of the VIF ID of the interface
- INET type (v4 or v6)
- action (TERM)
- ruleset name/ rule ID
- packet direction (IN/OUT)
- protocol (TCP, UDP, or PROTO #)
- TCP RST flag
- SVM direction for netx rule hit
- source IP address/source port>destination IP address/destination port
- IN packet count/OUT packet count (all accumulated)
- IN packet size/OUT packet size

The following is a sample of FQDN log file for distributed firewall rules:

```
2019-01-15T00:34:45.903Z 7c607b29 INET match PASS 1031 OUT 48 TCP 10.172.178.226/32808-
>23.72.199.234/80 S www.sway.com(034fe78d-5857-0680-81e4-d8da6b28d1b4)
```

The elements of an FQDN log include the following, separated by a space:

- timestamp:
- last eight digits of the VIF ID of the interface
- INET type (v4 or v6)
- reason (match)
- action (PASS, DROP, REJECT)
- ruleset name/ rule ID
- packet direction (IN/OUT)
- packet size

- protocol (TCP, UDP, or PROTO #) - for TCP connections, the actual reason that a connection is terminated is indicated after the following IP address
- source IP address/source port>destination IP address/destination port
- TCP flags - S (SYN), SA (SYN-ACK), A (ACK), P (PUSH), U (URGENT), F (FIN), R (RESET)
- domain name/UUID where UUID is the binary internal representation of the domain name

The following is a sample of Layer 7 log file for distributed firewall rules:

```
2019-01-15T00:35:07.221Z 82f365ae INET match REJECT 1034 OUT 48 TCP 10.172.179.6/49818-
>23.214.173.202/80 S APP_HTTP

2019-01-15T00:34:46.486Z 7c607b29 INET match PASS 1030 OUT 48 UDP 10.172.178.226/42035-
>10.172.40.1/53 APP_DNS
```

The elements of a Layer 7 log include the following, separated by a space:

- timestamp:
- last eight digits of the VIF ID of the interface
- INET type (v4 or v6)
- reason (match)
- action (PASS, DROP, REJECT)
- ruleset name/ rule ID
- packet direction (IN/OUT)
- packet size
- protocol (TCP, UDP, or PROTO #) - for TCP connections, the actual reason that a connection is terminated is indicated after the following IP address
- source IP address/source port>destination IP address/destination port
- TCP flags - S (SYN), SA (SYN-ACK), A (ACK), P (PUSH), U (URGENT), F (FIN), R (RESET)
- APP_XXX is the discovered application

Manage a Firewall Exclusion List

Firewall exclusion lists are made of groups that can be excluded from a firewall rule based on group membership.

User excluded groups are managed by the user, and empty by default.

Virtual machines such as load balancers, firewalls, virtual network functions (routing, switching, etc.), and any virtual machines that require promiscuous mode must be in a DFW Exclusion list. VMware does not support adding those virtual machines to DFW; they must be manually added to user excluded groups.

In NSX Manager cluster, the first node must be manually added to the Distributed Firewall Exclude List.

User-defined groups can be excluded from firewall rules, and there are a maximum of 100 groups that can be on the list. IP sets, MAC sets, and Active Directory groups cannot be included as members in a group that is used in a firewall exclusion list.

Note NSX-T Data Center automatically adds NSX Edge node virtual machines to the firewall exclusion list.

Procedure

- 1 Navigate to **Security > Distributed Firewall > Actions > Exclusion List**.

A window appears listing available groups.

- 2 To add a group to the exclusion list, click the check box next to any group. Then click **Apply**.
- 3 To create a group, click **Add Group**. See [Add a Group](#).
- 4 To edit a group, click the three dot menu next to a group and select **Edit**.
- 5 To delete a group, click the three dot menu and select **Delete**.
- 6 To display group details, click **Expand All**.

Filtering Specific Domains (FQDN/URLs)

Set up a distributed firewall rule to filter specific domains identified with FQDN/URLs, for example, **.office365.com*.

Currently, a predefined list of domains is supported. You can see the list of FQDNs when you add a new context profile of attribute type *Domain (FQDN) Name*. You can also see a list of FQDNs by running the API call `/policy/api/v1/infra/context-profiles/attributes?attribute_key=DOMAIN_NAME`.

Note FQDN filtering does not support CNAME records in DNS as the FQDN attribute type entry in context profiles.

You must set up a DNS rule first, and then the FQDN allowlist or denylist rule below it. NSX-T Data Center uses time to live (TTL) in the DNS response (coming from DNS server to the virtual machine), for keeping the DNS to IP mapping cache entry for the virtual machine (VM). To override the DNS TTL using a DNS security profile, see [Configure DNS Security](#). For FQDN filtering to be effective, virtual machines need to use a DNS server for domain resolution (no static DNS entries), and also need to honor the TTL received in the DNS response. NSX-T Data Center uses DNS Snooping to obtain a mapping between the IP address and the FQDN. SpoofGuard should be enabled across the switch on all logical ports to protect against the risk of DNS spoofing attacks. A DNS spoofing attack is when a malicious VM can inject spoofed DNS responses to redirect traffic to malicious endpoints or bypass the firewall. For more information about SpoofGuard, see [Understanding SpoofGuard Segment Profile](#).

This feature works at layer 7 and does not cover ICMP. If a user creates a denylist rule for all services on `example.com` the feature is working as intended if `ping example.com` responds, but `curl example.com` does not.

Selecting a wild card FQDN is a best practice because it includes sub domains. For example, selecting `*example.com`, would include sub domains such as `americas.example.com` and `emea.example.com`. Using `example.com` would not include any sub domains.

FQDN-based rules are retained during vMotion for ESXi hosts.

Note ESXi and KVM hosts are supported. KVM hosts support the FQDN allowlist only. FQDN filtering is available only with TCP and UDP traffic.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Navigate to **Security > Distributed Firewall**.
- 3 Add a firewall policy section by following the steps in [Add a Distributed Firewall](#). An existing firewall policy section can also be used.
- 4 Select the new or existing firewall policy section and click **Add Rule** to create the DNS firewall rule first.
- 5 Provide a name for the firewall rule, such as **DNS rule**, and provide the following details:

Option	Description
Services	Click the edit icon and select the DNS or DNS-UDP service as applicable to your environment.
Profile	Click the edit icon and select the DNS context profile. This is precreated and is available in your deployment by default.
Applied To	Select a group as required.
Action	Select Allow .

- 6 Click **Add Rule** again to set up the FQDN allowlist or denylist rule.
- 7 Name the rule appropriately, such as, **FQDN/URL Allowlist**. Drag the rule under the DNS rule under this policy section.
- 8 Provide the following details:

Option	Description
Services	Click the edit icon and select the service you want to associate with this rule, for example, HTTP.
Profile	Click the edit icon and click Add New Context Profile . Click in the column titled Attribute , and select Domain (FQDN) Name . Select the list of Attribute Name/Values from the predefined list. Click Add . See Add a Context Profile for details.

Option	Description
Applied To	Select DFW or a group as required.
Action	Select Allow , Drop , or Reject .


9 Click **Publish**.

Extending Security Policies to Physical Workloads

NSX-T Data Center can act as a single point of administration for both virtual and physical workloads.

NSX-T Data Center supports integration with Arista CloudVision eXchange (CVX). This integration facilitates consistent networking and security services across virtual and physical workloads, independent of your application frameworks or physical network infrastructure. NSX-T Data Center does not directly program the physical network switch or router but integrates at the physical SDN controller level, therefore preserving the autonomy of security administrators and physical network administrators.

NSX-T Data Center supports integration with Arista EOS 4.22.1FX-PCS and later.



Limitations

- Arista switches require ARP traffic to exist before firewall rules are applied to an end host that is connected to an Arista switch. Packets can therefore pass through the switch before firewall rules are configured to block traffic.
- Allowed traffic does not resume when a switch crashes or is reloaded. The ARP tables need to be populated again, after the switch comes up, for the firewall rules to be enforced on the switch.
- Firewall rules cannot be applied on the Arista Physical Switch, for FTP passive clients that connect to FTP Server connected to the Arista Physical Switch.
- In CVX HA setup that uses Virtual IP for the CVX cluster, the CVX VM's dvpg's Promiscuous mode, and Forged transmits must be set to Accept. In case they are set to default (Reject), the CVX HA Virtual IP will not be reachable from NSX Manager.

Configure NSX-T Data Center to interact with Arista CVX

Complete the configuration procedure on NSX-T Data Center so that CVX can be added as an enforcement point in NSX-T Data Center and NSX-T Data Center can interact with CVX.

Prerequisites

Obtain the virtual IP address for the Arista CVX cluster.

Procedure

- 1 Log in to NSX Manager as a root user and run the following command to retrieve the thumbprint for CVX:

```
openssl s_client -connect <virtual IP address of CVX cluster> | openssl x509 -noout
-fingerprint -sha256
```

Sample output:

```
depth=0 CN = self.signed
verify error:num=18:self signed certificate
verify return:1
depth=0 CN = self.signed
verify return:1
SHA256
Fingerprint=35:C1:42:BC:7A:2A:57:46:E8:72:F4:C8:B8:31:E3:13:5F:41:95:EF:D8:1E:E9:3D:F0:CC:3
B:09:A2:FE:22:DE
```

- 2 Edit the retrieved thumbprint to use only lower case characters and exclude any colons in the thumbprint.

Sample of edited thumbprint for CVX:

```
35c142bc7a2a5746e872f4c8b831e3135f4195efd81ee93df0cc3b09a2fe22de
```

- 3 Call the `PATCH /policy/api/v1/infra/sites/default/enforcement-points` API and use the CVX thumbprint to create an enforcement endpoint for CVX. For example:

```
PATCH https://<nsx-manager>/policy/api/v1/infra/sites/default/enforcement-points/cvx-
default-ep
{
  "auto_enforce": "false",
  "connection_info": {
    "enforcement_point_address": "<IP address of CVX>",
    "resource_type": "CvxConnectionInfo",
    "username": "cvpadmin",
    "password": "1q2w3e4rT",
    "thumbprint": "65a9785e88b784f54269e908175ada662be55f156a2dc5f3a1b0c339cea5e343"
  }
}
```

- 4 Call the `GET /policy/api/v1/infra/sites/default/enforcement-points` API to retrieve the endpoint information. For example:

```
https://<nsx-manager>/policy/api/v1/infra/sites/default/enforcement-points/cvx-default-ep
{
  "auto_enforce": "false",
  "connection_info": {
    "enforcement_point_address": "<IP address of CVX>",
    "resource_type": "CvxConnectionInfo",
    "username": "admin",
    "password": "1q2w3e4rT",
    "thumbprint": "35c142bc7a2a5746e872f4c8b831e3135f4195efd81ee93df0cc3b09a2fe22de"
  }
}
```

Sample output:

```
{
  "connection_info": {
    "thumbprint": "35c142bc7a2a5746e872f4c8b831e3135f4195efd81ee93df0cc3b09a2fe22de",
    "enforcement_point_address": "192.168.2.198",
    "resource_type": "CvxConnectionInfo"
  },
  "auto_enforce": false,
  "resource_type": "EnforcementPoint",
  "id": "cvx-default-ep",
  "display_name": "cvx-default-ep",
  "path": "/infra/sites/default/enforcement-points/cvx-default-ep",
  "relative_path": "cvx-default-ep",
  "parent_path": "/infra/sites/default",
  "marked_for_delete": false,
  "_system_owned": false,
  "_create_user": "admin",
  "_create_time": 1564036461953,
  "_last_modified_user": "admin",
}
```

```

    "_last_modified_time": 1564036461953,
    "_protection": "NOT_PROTECTED",
    "_revision": 0
}

```

- 5 Call the POST `/api/v1/notification-watchers/` API and use the CVX thumbprint to create a notification ID. For example:

```

POST https://<nsx-manager>/api/v1/notification-watchers/
{
  "server": "<virtual IP address of CVX cluster>",
  "method": "POST",
  "uri": "/pcs/v1/nsgroup/notification",
  "use_https": true,
  "certificate_sha256_thumbprint":
  "35c142bc7a2a5746e872f4c8b831e3135f4195efd81ee93df0cc3b09a2fe22de",
  "authentication_scheme": {
    "scheme_name": "BASIC_AUTH",
    "username": "cvpadmin",
    "password": "1q2w3e4rT"
  }
}

```

- 6 Call the GET `/api/v1/notification-watchers/` to retrieve the notification ID.

Sample output:

```

{
  "id": "a0286cb6-de4d-41de-99a0-294465345b80",
  "server": "192.168.2.198",
  "port": 443,
  "use_https": true,
  "certificate_sha256_thumbprint":
  "35c142bc7a2a5746e872f4c8b831e3135f4195efd81ee93df0cc3b09a2fe22de",
  "method": "POST",
  "uri": "/pcs/v1/nsgroup/notification",
  "authentication_scheme": {
    "scheme_name": "BASIC_AUTH",
    "username": "cvpadmin"
  },
  "send_timeout": 30,
  "max_send_uri_count": 5000,
  "resource_type": "NotificationWatcher",
  "display_name": "a0286cb6-de4d-41de-99a0-294465345b80",
  "_create_user": "admin",
  "_create_time": 1564038044780,
  "_last_modified_user": "admin",
  "_last_modified_time": 1564038044780,
  "_system_owned": false,
  "_protection": "NOT_PROTECTED",
  "_revision": 0
}

```


- 7 Call the PATCH `/policy/api/v1/infra/domains/default/domain-deployment-maps/cvx-default-dmap` API to create a CVX domain deployment map. For example:

```
PATCH https://<nsx-manager>/policy/api/v1/infra/domains/default/domain-deployment-maps/cvx-
default-dmap
{

  "display_name": "cvx-deployment-map",

  "id": "cvx-default-dmap",

  "enforcement_point_path": "/infra/sites/default/enforcement-points/cvx-default-ep"

}
```

- 8 Call the GET `/policy/api/v1/infra/domains/default/domain-deployment-maps` API to retrieve the deployment map information.

Configure Arista CVX to interact with NSX-T Data Center

After configuring NSX-T Data Center, complete the configuration procedure on Arista CloudVision eXchange (CVX) to enable CVX to interact with NSX-T Data Center.

Prerequisites

NSX-T Data Center has registered the CVX as an enforcement point.

Procedure

- 1 Log in to NSX Manager as a root user and run the following command to create a thumbprint for CVX to communicate with NSX Manager:

```
openssl s_client -connect <IP address of nsx-manager>:443 | openssl x509 -pubkey -noout |
openssl rsa -pubin -outform der | openssl dgst -sha256 -binary | openssl base64
```

Sample output:

```
depth=0 C = US, ST = CA, L = Palo Alto, O = VMware Inc., OU = NSX, CN = nsx-mgr
verify error:num=18:self signed certificate
verify return:1
depth=0 C = US, ST = CA, L = Palo Alto, O = VMware Inc., OU = NSX, CN = nsx-mgr
verify return:1
writing RSA key
S+zwADluzeNf+dnffDpYvgs4YrS6QBgyeDry40bPgms=
```

- 2 Run the following commands from the CVX CLI:

```
cvx
no shutdown
service pcs
no shutdown
controller <IP address of nsx-manager>
username <NSX administrator user name>
```

```
password <NSX administrator password>
enforcement-point cvx-default-ep
pinned-public-key <thumbprint for CVX to communicate with NSX
                    Manager>
notification-id <notification ID created while registering CVX with NSX>
end
```

- 3 Run the following command from the CVX CLI to check the configuration:

```
show running-config
```

Sample output:

```
cvx
  no shutdown
  source-interface Management1
  !
  service hsc
    no shutdown

  !
  service pcs
    no shutdown
    controller 192.168.2.80
    username admin
    password 7 046D26110E33491F482F2800131909556B
    enforcement-point cvx-default-ep
    pinned-public-key sha256//S+zwADluzeNf+dnffDpYvgs4YrS6QBgyeDry40bPgms=
    notification-id a0286cb6-de4d-41de-99a0-294465345b80
```

- 4 Configure tag on the ethernet interface of the physical switch that connects to the physical server. Run the following commands on the physical switch managed by CVX.

```
configure terminal
interface ethernet 4
tag phy_app_server
end
copy running-config startup-config
Copy completed successfully.
```

- 5 Run the following command to verify tag configuration for the switch:

```
show running-config section tag
```

Sample output:

```
interface Ethernet4
  description connected-to-7150s-3
  switchport trunk allowed vlan 1-4093
  switchport mode trunk
  tag sx4_app_server
```

IP addresses that are learnt on the tagged interfaces, using ARP, are shared with NSX-T Data Center.

- 6 Log in to NSX Manager to create and publish firewall rules for the physical workloads managed by CVX. See [Chapter 13 Security](#) for more information on creating rules. For example:

	Name	Sources	Destinations	Services	Profiles	Applied To	Action
⋮	Firewall_Services	(2)	Applied To	DFW			● Up
⋮	vm_to_phy_server	vm	phy_server	Any	None	DFW	● Allow
⋮	phy_server_to_vm	phy_server	vm	Any	None	DFW	● Allow

NSX-T Data Center policies and rules published in NSX-T Data Center appear as dynamic ACLs on the physical switch managed by CVX.

```
prmh-nsx-tor-7050sx-4#show ip access-lists dynamic
IP Access List et4.v4.in [dynamic]
 10 permit ip host 71.1.1.3 host 27.1.1.11

IP Access List et4.v4.out [dynamic]
 10 permit ip host 27.1.1.11 host 71.1.1.3
```

For more information, see [CVX HA set up](#), [CVX HA Virtual IP setup](#), and [Physical Switch Mlag Setup](#)

Shared Address Sets

Security groups based on dynamic or logical objects can be created and used in the **Applied to** text box of distributed firewall rules.

Because address sets are dynamically populated based on virtual machine name or tags, and must be updated on each filter, they can exhaust the available amount of heap memory on hosts to store DFW rules and IP address sets.

In NSX-T Data Center version 2.5 and later, a feature called Global or Shared Address Sets, makes address sets shared across all the filters. While each filter can have different rules, based on **Applied To**, the address sets members are constant across all the filters. This feature is enabled by default, reducing heap memory use. It cannot be disabled.

In NSX-T Data Center version 2.4 and earlier, Global or Shared Address Sets is disabled, and environments with heavy distributed firewall rules might experience VSIP heap exhaustion.

Distributed IDS

Distributed Intrusion Detection Service (IDS) monitors network traffic on the host for suspicious activity.

IDS detects intrusion attempts based on already known malicious instruction sequences. The detected patterns in the IDS are known as signatures. Specific signatures can be excluded from intrusion detection.

Note Do not enable Distributed Intrusion Detection Service (IDS) in an environment that is using Distributed Load Balancer. NSX-T Data Center does not support using IDS with a Distributed Load Balancer.

Distributed IDS Configuration:

- 1 Enable IDS on hosts, download latest signature set, and configure signature settings.
[Distributed IDS Settings and Signatures](#)
- 2 Create IDS profiles. [Distributed IDS Profiles](#)
- 3 Create IDS rules. [Distributed IDS Rules](#)
- 4 Verify IDS status on hosts. [Verify Distributed IDS Status on Host](#)

Distributed IDS Settings and Signatures

NSX-T can automatically apply signatures to your hosts, and update intrusion detection signatures by checking our cloud-based service.

Distributed firewall (DFW) must be enabled for IDS to work. If traffic is blocked by a DFW rule, then IDS will not see the traffic.

Intrusion detection can be enabled on standalone hosts by toggling the **enabled** bar. If VC clusters are detected, IDS can also be enabled on a cluster basis by selecting the cluster and clicking **enable**.

Signatures

Signatures are applied to IDS rules through profiles. A single profile is applied to matching traffic. By default, NSX Manager checks for new signatures once per day. New signature update versions are published every two weeks (with additional non-scheduled 0-day updates). When a new update is available, there is a banner across the page with an **Update Now** link.

If **Auto update new versions** is selected, signatures are automatically applied to your hosts after they are downloaded from the cloud. If auto update is disabled, the signatures are stopped at the listed version. Click **view and change versions** to add another version, in addition to the default. Currently, two versions of signatures are maintained. Whenever there is a change in the version commit identification number, a new version is downloaded.

If a proxy server is configured for NSX Manager to access the Internet, click **Proxy Settings** and complete the configuration.

Offline Downloading and Uploading Signatures

To download and upload a signature bundle, when NSX Manager does not have Internet access:

- 1 This API is the first one to be called before any communication with the cloud service is started. It registers the client using the client's license key, and generates credentials for the client to use. Send all licenses and you will be provided with necessary permission. The `client_id` is the name given by the user. The `client_secret` is generated and used as the request for the Authentication API. If the client has previously registered, but does not have access to the `client_id` and `client_secret`, the client has to re-register using the same API.

```
POST https://api.nsx-sec-prod.com/1.0/auth/register
```

Body:

```
{
  "license_keys":["xxxxx-xxxxx-xxxxx-xxxxx-xxxxx"],
  "device_type":"NSX-Idps-Offline-Download",
  "client_id": "client_username"
}
```

Response:

```
{
  "client_id":"client_username",
  "client_secret": "Y54+V/
rCpEm50x5HAUIzH6aXtTq7s97wCA2QqZ8VyrFQjrJih7h0a1ItDQn02T46EJVnSMZWTseragTFScrtIwsiPSX7APQI
C7MxAYZ0BoAWvW2akMxyZKyzbYZjeROb/C2QchehC8GFIFNpwqiAcQjrQHwHGdtX4zTQ="
}
```

- 2 This API call authenticates the client using the `client_id` and `client_secret`, and generates an authorization token to use in the headers of requests to IDS Signatures APIs. The token is valid for 60 minutes. If the token is expired, the client has to reauthenticate using the `client_id` and `client_secret`.

```
POST https://api.nsx-sec-prod.com/1.0/auth/authenticate
```

Body:

```
{
  "client_id":"client_username",
  "client_secret": "Y54+V/
rCpEm50x5HAUIzH6aXtTq7s97wCA2QqZ8VyrFQjrJih7h0a1ItDQn02T46EJVnSMZWTseragTFScrtIwsiPSX7APQI
C7MxAYZ0BoAWvW2akMxyZKyzbYZjeROb/C2QchehC8GFIFNpwqiAcQjrQHwHGdtX4zTQ="
}
```

Response:

```
{
    "access_token":
        "eyJhbGciOiJIUzUxMiJ9.eyJqdGkiOiI3ZjMwNzVhMmQWNTYyZjJkZWMyNmZhNDhmMGZlMTk0YjMyMzU4OGU5NGU5NEZBMnNmNWZkOWYUWiJkdjLTJKYWYzMmE3LTYxMzc0TGJiNS05NzJlLTE0bjZhMGVkYmU3MCIsInN1YiI6IjdmMZA3ZWlyeDA3YjpmNmNiJmZlmZTk2MmM2ZmEOMWEvZWUxOTRtIGF0LnA3dCJpcXQiOiJlelNUUlNDNgD0TR9.x4U75GSshDLMyhiyu02B9HiIAdonzxzSmo01qRhvxUErQSPE_Kxq3rzg1_IIvyoy3SJwwDhSh8KEctGW50ecPg",
    "token_type": "bearer",
    "expires_in": 3600,
    "scope": "[idps_scope]"
}
```

- 3 The response to this command has the link for the ZIP file. NSXCloud downloads the signatures from the git hub repo every 24 hours, and saves the signatures in a ZIP file. Copy and paste the signatures URL into your browser, and the ZIP file will download.

```
GET https://api.nsx-sec-prod.com/1.0/intrusion-services/signatures
```

In the Headers tab, the Authorization key will have the `access_token` value from the authenticate API response.

Authorization: eyJhbGciOiJIUzUxMiJ9.eyJqdGkiOiI3ZjMwN2VhMmQwN2IyZjJjYzU5ZmU5NjJjNmZhNDZhMGZlMTk4YjMyMzU4OGU5NGU5NzE3NmNmNzkyYUWU1YjYjLTJkYWY2MmE3LTlyXmZctNGJiNS05NzJlLTE0NjZhMGNkYmU3MCIsInN1YiI6IjdmMzA3ZWVhZDA3YjJmMmNjMzlmZTk2MmM2ZmE0MWEwZmUxOThiMzIzNTg4ZTk0ZTk3MTEyY2Y3OTRhZTViNDM2MmRhZjYyYTctNjEzNy00YmI1LTk3MmUtMTQ2NmEwY2RlZTcwIiwiaXhwaWJoXzNTU1NTUyMjk0LCJpYXQiOiJlNTU1NDg2OTR9.x4U75GShDLmhiyUO2B9Hi1Adonzx3Smo01qRhvXuErQSPe Kxq3rzgl IIyvoy3SJwDhSh8KECTGW50eCPg

Response:

```
{
"signatures_url": "https://ncs-idps-us-west-2-prod-signatures.s3.us-west-2.amazonaws.com/a07fe284ff70dc67194f2e7cfla8178d69570528.zip?X-Amz-Security-Token=IQoJb3JpZ2luX2VjENF%2F%2F%2F%2F%2F%2F%2F%2F%2F%2FwEaCXvZLXdlc3QtMSJHMEUCIGlUYbzfbXOsm1lvdj1k36LPyoPota0L4CSOBMXgKGhmAiEA%2BQC1K4Gr7VCRiBM4ZH2WbP2rvIp0qfhfGLOx0ChGc4q6wEIHXABGgw1MTAwMTM3MTElNTMiDA4H4ir7eJl779wwWirIAdLiXuAukLwnhmlgLmydZhW7ZExe%2BamDkRU7KT46ZS93mC1CQeL00D2rjBybCBiGlmgNILPuQ2EyxmqxhEOzFYimXDDBER4pmv8%2BbKnDWPg08RNTqpD%2BAMicYNP7WlpxeZWYxcoBFruCDA2l3eXS6XNv3ot6T2a%2Bk4rMKhtZyFkzREIIcQlPg7Ej5q62EvvmMFQdo8TyZxFpMJBC4IEg0hlk6QZU1Jlkrq2RYKit5WwLD%2BQKJREdf4A0YctLbMCDBNbprUcCADMKyclu8FOuABuK90a%2BvnA%2FJFYiJ32eJl%2Bdt0YrbTnRyv1MuSUHxjNAdyrfXnkPyF80%2FQLYLVDVRUDatyAo10s3C0pzYN%2FvMKsumExy6FIcv%2FOLo08Y9RaMoTnUfeugpr6YsqMCH0pUR4dIVDYoi1hldNCf1XD74xmJSdnviayY4vXD4bBDKPnrFFhOxLTRFAVWLMDNYggLh3pV3rxDPnIwgFTFr7CmZGGJAQBKqaxzPMVZ2TQBAbmjxoRqCBip8Y662Tbjth7im2V522LMvonM6Tysf16ls6QU9IC6Wqjd0dei5yazK%2Fr9g%3D&X-Amz-Algorithm=AWS4-HMAC-SHA256&X-Amz-Date=20191202T222034Z&X-Amz-SignedHeaders=host&X-Amz-Expires=3599&X-Amz-Credential=ASIAXPZPUTA6A7V7P4X%2F20191202%2Fus-west-1%2Fs3%2Faws4_request&X-Amz-Signature=d85ca4aef6abe22062e2693acacf823f0a4fc51d1dc07cda8dec93d619050f5e"
}
```

- 4 Navigate to **Security > Distributed IDS > Settings**. Click **Upload IDS Signatures** in the right corner. Navigate to the saved signature ZIP file and upload the file. You can also upload the signature ZIP using the API call:

```
POST https://<mgr-ip>/policy/api/v1/infra/settings/firewall/security/intrusion-services/signatures?action=upload_signatures
```

Distributed IDS Profiles

IDS Profiles are used to group signatures, which can then be applied to select applications.

Signatures can be enabled based on the severity rating of the signature. A higher score indicates an increased risk associated with the intrusion event. Severity is determined based on the following:

- Severity specified in the signature itself
- CVSS (Common Vulnerability Scoring System) score specified in the signature
- Type-rating associated with the classification type

Exclusions are set per severity level and are used to disable signatures, reducing noise and improving performance. Exclusions are used to disable signatures:

- That cause false positives
- That are noisy
- That are irrelevant to the protected workloads

The default IDS profile includes critical severities and cannot be edited.

Procedure

- 1 Navigate to **Security > Distributed IDS > Profiles**.
- 2 Enter a profile name and description.
- 3 Click one or more of the severities you want to include.
See [IDS Severity Ratings](#) for more information.
- 4 To exclude a severity, click **select** under **Signatures to Exclude**. You can now view and exclude the signatures included in that severity level. Click **Add** to add a signature to the exclusion list. The following information is provided for each signature:

Variable	Description
Signature ID	Identification number that references individual signatures.
Details	Describes the threat.
Product Affected	Shows what product is vulnerable to the exploit.
Attack Target	Target of the attack.
IDS Severity	Indicates the severity of the signature. For more details, see IDS Severity Ratings .

Variable	Description
CVSS (Common Vulnerability Scoring System)	CVSS is a framework for rating the severity of security vulnerabilities in software. A CVSS base score of 0.0-3.9 is considered low severity. A CVSS base score of 4.0-6.9 is medium severity. A CVSS base score of 7.0-10.0 is high severity.
CVE (Common Vulnerability Enumeration)	Common Vulnerability Enumeration (CVE), is a dictionary of publicly known information security vulnerabilities and exposures.
Category	Type of attack.

5 Click **Save**.

What to do next

Create IDS rules.

IDS Severity Ratings

Signature severity helps security teams prioritize incidents.

A higher score indicates an increased risk associated with the intrusion event.

NSX IDS Severity Level	Classification Type-Rating	Classification Types
CRITICAL	1	■ Attempted User Privilege Gain ■ Unsuccessful User Privilege Gain ■ Successful User Privilege Gain ■ Attempted administrator Privilege Gain ■ Successful Administrator Privilege Gain ■ Executable Code was Detected ■ A Network Trojan was Detected ■ Web Application Attack ■ Inappropriate Content was Detected ■ Potential Corporate Privacy Violation ■ Targeted Malicious Activity was Detected ■ Exploit Kit Activity Detected ■ Domain Observed Used for C2 Detected ■ Successful Credential Theft Detected ■ Emerging Threat alert from SpiderLabs Research ■ RedAlert from SpiderLabs Research
High	2	■ Potentially Bad Traffic ■ Information Leak ■ Large Scale Information Leak ■ Attempted Denial of Service ■ Decode of an RPC Query ■ Suspicious Filename Detected ■ Attempted Login Using a Suspicious Username ■ System Call Detected ■ Client Using an Unusual Port ■ Detection of a Denial of Service Attack ■ Detection of a Non-Standard Protocol or Event ■ Access to a Potential Vulnerable Web Application Attack ■ Attempt to Log in By a Default Username and Password ■ Device Retrieving External IP Address Detected ■ Possibly Unwanted Program Detected

NSX IDS Severity Level	Classification Type-Rating	Classification Types
		■ Possible Social Engineering Attempted ■ Crypto Currency Mining Activity Detected
Medium	3	■ Not Suspicious Traffic ■ Unknown Traffic ■ Suspicious String was Detected ■ Detection of a Network Scan ■ Generic Protocol Command Decode ■ Misc Activity ■ Generic ICMP event
Low	4-9	■ TCP Connection Detected ■ Non-specific Potential Attack ■ Attempt to Exploit Client-side Web Application Vulnerability ■ Non-specific Potential Web App Attack ■ Traffic Which is Likely a Bad Idea or Misconfiguration ■ Attempt to Exploit Administrative-level Vulnerability ■ Attempt to Exploit user-level Vulnerability ■ IP Based Alert From SpiderLabs Research ■ Successful Exploitation of a Root-level Vulnerability ■ Indication of an Active Backdoor Channel ■ Worm Propagation ■ Specific Virus Detected

Distributed IDS Rules

IDS rules are used to apply a previously created profile to select applications and traffic.

IDS rules are created in the same manner as distributed firewall (DFW) rules. First, an IDS policy or section is created, and then rules are created. DFW must be enabled, and traffic must be allowed by DFW to be passed through to IDS rules.

IDS rules must:

- specify one IDS profile per rule
- stateful
- use of Layer 7 attributes (APP IDs) is not supported

One or more policy sections with rules must be created, because there are no default rules. Before creating rules, create a group that needs a similar rule policy. See [Add a Group](#).

- 1 Navigate to **Security > IDS > Rules**.
- 2 Click **Add Policy** to create a policy section, and give the section a name.
- 3 Click the gear icon to configure the following policy section options:

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. See Role-Based Access Control .

- 4 Click **Add Rule** to add a new rule, and give the rule a name.
- 5 Configure source/destination/services to determine which traffic needs IDS inspection. IDS supports any type of group for source and destination.
- 6 Select the **IDS Profile** to be used for the matching traffic. For more information, see [Distributed IDS Profiles](#).
- 7 Configure **Applied To**, to limit the scope of the rules. Groups consisting of only IP addresses, MAC addresses, or Active Directory groups cannot be used in the **Applied To** text box.
- 8 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.

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

For more information about creating policy sections and rules, see [Add a Distributed Firewall](#).

Distributed IDS Events

The events window contains the last 14 days of data.

Navigate to **Security > Distributed IDS > Events** to view time intrusion events.

There are three event log files in the `/var/log/nsx-idps` folder on ESXi hosts:

- fast log - contains internal logging of nsx-idps process events, with limited information and is used only for debugging purposes
- nsx-idps-log - contains general nsx-idps process logs with basic information and errors about the process workflow
- nsx-idps-events.log - contains detailed information about events (all alerts/drops/rejects) with NSX metadata

Colored dots indicate the unique type of intrusion events and can be clicked for details. The size of the dot indicates the number of times an intrusion event has been seen. A blinking dot indicates that an attack is ongoing. Point to a dot to see the attack name, number of attempts, first occurrence, and other details.

- Red dots - represent critical severity signature events.
- Orange dots - represent high severity signature events.
- Yellow dots - represent medium severity signature events.
- Gray dots - represent low severity signature events.

All the intrusion attempts for a particular signature are grouped and plotted at their first occurrence.

- Select the timeline by clicking the arrow in the upper right corner. The time line can be between 24 hours and 14 days.
- Filter events by:

Filter Criteria	Description
Attack Target	Target of the attack.
Attack Type	Type of attack, such as trojan horse, or denial of service (DoS).
CVSS (Common Vulnerability Score)	Common Vulnerability Score (filter based on a score above a set threshold).
Product Affected	Vulnerable product or (version) i.e Windows XP or Web_Browsers
VM Name	The VM (based on logical port) where exploit traffic originated from or was received by.

- Click the arrow next to an event to view details.

Detail	Description
Last Detected	This is the last time the signature was fired.
Details	The name of the signature that was fired.
Product Affected	Illustrates what product is vulnerable to the exploit.

Detail	Description
VM Affected	Lists of VMs involved in the intrusion attempt.
Vulnerability Details	If available, this shows a link to the CVE and the CVSS score associated with the vulnerability.
Source	IP address of the attacker and source port used.
Destination	IP address of the victim and destination port used.
Attack Direction	Client-Server or Server-Client.
Associated IDS Rule	Clickable link to the configured IDS Rule which resulted in this event.
Revision	The revision number of the IDS signature.
Activity	Displays the total number of times this particular IDS signature was triggered, the most recent occurrence, and the first occurrence.

- To view intrusion history, click the arrow next to an event, then click **View Intrusion History**. A window opens with the following details:

Detail	Description
Source IP	IP address of the attacker.
Source Port	Source port used in the attack.
Destination IP	IP address of the victim.
Destination Port	Destination port used in the attack.
Protocol	Traffic protocol of the detected intrusion.
Time Detected	This is the last time the signature was fired.

- The graph present under the chart represents events that occurred over a selected time span. You can zoom in to the specific time window on this graph to view details of signatures of the related events that happened during the time window.

Verify Distributed IDS Status on Host

To use the NSX virtual appliance CLI, you must have SSH access to an NSX virtual appliance. Each NSX virtual appliance contains a command-line interface (CLI).

The viewable modes in the CLI can differ based on the assigned role and rights of a user. If you are unable to access an interface mode or issue a particular command, consult your NSX administrator.

Procedure

- 1 Open an SSH session to a compute host running the work loads that were previously deployed. Log in as root.
- 2 Enter the `nsxcli` command to open the NSX-T Data Center CLI.

- 3 To confirm that IDS is enabled on this host, run the command: `get ids status`.

Sample Output:

```
localhost> get ids status
NSX IDS Status
-----
status: enabled
uptime: 793756 (9 days 04:29:16)
```

- 4 To confirm both of the IDS profiles have been applied to this host, run the command `get ids profile`.

```
localhost> get ids profiles
NSX IDS Profiles
-----
Profile count: 2
1. 31c1f26d-1f26-46db-b5ff-e6d3451efd71
2. 65776dba-9906-4207-9eb1-8e7d7fdf3de
```

- 5 To review IDS profile (engine) statistics, including the number of rules loaded, and the number of packets and sessions evaluated, run the command `get ids engine stats`.

The output is on a per profile basis, and shows the number of signatures loaded for each profile, and the number of packets that were evaluated.

```
localhost> get ids engine stats
NSX IDS Engine Statistics
-----
uptime: 18 (0 days 00:00:18)

app_layer:
-----
flow:
http: 10713
tx:
http: 25911
detect:
-----
engines:
alerts: 11129
id: 3
last_reload: 2020-03-17T21:29:39.387087+0000
packets_incoming: 572083
packets_outgoing: 571066
prof-uuid: 53ef4dba-0291-4ea3-96ef-d01259dca2fe
rules_failed: 0
rules_loaded: 11906

tcp:
---
memuse: 20872880
overlap: 50006
```

```
reassembly_memuse: 155439408
rst: 23797
sessions: 58811
syn: 89615
synack: 41635
```

East-West Network Security - Chaining Third-party Services

After partners register network services such as Intrusion Detection System or Intrusion Protection System (IDS/IPS) with NSX-T Data Center, as an administrator you can configure network services to introspect east-west traffic moving between VMs on an on-premises data center.

Prerequisites

- Partners must register services with NSX-T Data Center.
- ESXi hosts must be prepared as NSX-T Data Center transport nodes by using transport node profiles.

Note

- Service VMs are only supported on ESXi hosts and not supported on KVM hosts.
 - NSX-T Data Center only protects guest VMs running on ESXi hosts.
 - NSX-T Data Center does not protect guest VMs running on KVM hosts.
-

Key Concepts of Network Protection East-West

Traffic flowing between Guest VMs on an on-premises data center is protected by third-party services provided by partners. There are a few concepts that aid your understanding of the workflow.

- **Service:** Partners register services with NSX-T Data Center . A service represents the security functionality offered by the partner, service deployment details such as OVF URL of service VMs, point to attach the service, state of the service. When a notification is generated for a service, NSX-T Data Center notifies the partner after a time interval of 30 seconds.
- **Vendor Template:** It consists of functionality that a service can perform on a network traffic. Partners define vendor templates. For example, a vendor template can provide a network operation service such as tunneling with IPSec service.
- **Service Profile:** Is an instance of a vendor template. An NSX-T Data Center administrator can create a service profile to be consumed by service VMs.
- **Guest VM:** a source or destination of traffic in the network. The incoming or outgoing traffic is introspected by a service chain defined for a rule running east-west network services.
- **Service VM:** A VM that runs the OVA or OVF appliance specified by a service. It is connected over the service plane to receive redirected traffic.

- **Service Instance:** Is created when a service is deployed on a host. Each service instance has a corresponding service VM.
- **Service Segment:** A segment of a service plane that is associated to a transport zone. Each service attachment is segregated from other service attachments and from the regular L2 or L3 network segments provided by NSX-T. The service plane manages service attachments.
- **Service Manager:** Is the partner service manager that points to a set of services.
- **Service Chain:** Is a logical sequence of service profiles defined by an administrator. Service profiles introspect network traffic in the order defined in the service chain. For example, the first service profile is firewall, second service profile is monitor, and so on. Service chains can specify different sequence of service profiles for different directions of traffic (egress/ingress).
- **Redirection Policy:** Ensures that traffic classified for a specific service chain is redirected to that service chain. It is based on traffic patterns that match NSX-T Data Center security group and a service chain. All traffic matching the pattern is redirected along the service chain.
- **Service Path:** Is a sequence of service VMs that implement the service profiles of a service chain. An administrator defines the service chain, which consists of a pre-defined order of service profiles. NSX-T Data Center generates multiple service paths from a service chain based on the number, and locations of guest VMs and service VMs. It selects the optimum service path for the traffic flow to be introspected. Each service path is identified by a Service Path Index (SPI) and each hop along a path has a unique Service Index (SI).

NSX-T Data Center Requirements for East-West Traffic

In the NSX-T Data Center deployment, you need to ensure an overlay transport zone and overlay-backed logical switches exists.

East-West service insertion is applied to an entire NSX-T deployment. You cannot deploy the service at a cluster-level or a host-level.

All transport nodes must be of the type Overlay because the service sends traffic on GENEVE or overlay-backed logical switches. A overlay-backed (GENEVE-backed) logical switch is provisioned internally and not visible on the user interface.

Even if you plan a deployment using only VLAN-backed logical switches, East-West traffic passes through overlay transport zones and overlay-backed logical switches. So, ensure that you create an overlay transport zone and GENEVE-backed logical switches. Without these requirements, during a vMotion, the guestVM on a host cannot be migrated to another transport node. The guestVM goes into Disconnected state causing configuration errors in the East-West service.

High-Level Tasks for East-West Network Security

Follow these steps to set up network security for east-west traffic.

Table 13-3. List of Tasks to Configure East-West Network Introspection

Workflow Tasks	Persona	Implementation
Register Service	Partner	Only API
Register Vendor Template	Partner	Only API
Register Service Manager	Partner	Only API
Deploy a Service for East-West Traffic Introspection	Administrator	API and NSX Manager UI
Add a Service Profile	Administrator	API and NSX Manager UI
Add a Service Chain	Administrator	API and NSX Manager UI
Add Redirection Rules for East-West Traffic	Administrator	API and NSX Manager UI

Deploy a Service for East-West Traffic Introspection

After partners register services, as an administrator, you must deploy an instance of the service on member hosts of a cluster.

Deploy partner service VMs that run the partner security engine on all the NSX-T Data Center hosts in a cluster. After you deploy the SVMs, you can create policy rules used by SVM to protect guest VMs.

Prerequisites

- All hosts are managed by a vCenter Server.
- Partner services must be registered with NSX-T Data Center and are ready for deployment.
- NSX-T Data Center administrators can access partner services and vendor templates.
- Both the service VM and the partner service manager (console) must be able to communicate with each other at the management network level.
- Host-based service deployment: Before you deploy service VMs on each host, configure each host of the cluster with NSX-T Data Center by applying a transport node profile.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **System > Service Deployments > Deployment > Deploy Service**.
- 3 From the Partner Service field, select the partner service.
- 4 Enter the service deployment name.
- 5 In the Compute Manager field, select the vCenter Server to deploy the service.
- 6 In the Cluster field, select the cluster where the services need to be deployed.

- 7 In the Data Store drop-down menu, select a data store as the repository for the service virtual machine.
- 8 In the Network column, click **Set** and enter the Management Network interface by choosing DHCP or static IP address type, and data network.
- 9 In the Service Segments field, select a service segment from the list or click the Action icon to add or edit a service segment.

Guest VMs connected to a service segment are provided east-west network traffic protection.

To create a service segment:

- a Click the **+** icon next to the Service Segment field.
 - b In the Service Segment dialog box, click **Add Service Segment**.
 - c Enter a name, select a Transport Zone Overlay from the drop-down menu, and if applicable, select a gateway under Applied to Gateway.
 - d Click **Save**.
- 10 In the Deployment Type field, select from one of the following deployment options. Depending upon the services registered by the partner, multiple services can be deployed as part of a single service VM.
 - **Clustered:** Deploys the service on a host or hosts belonging to a cluster that is dedicated to host service VMs.
 - **Host Based:** Deploys the service on all the hosts within a cluster.
 - 11 In the Deployment Template field, select the template that provides attributes to protect the workload you want to run on guest VMs groups.
 - 12 (Cluster-based deployment only) In the Clustered Deployment Count, enter the number of service VMs to deploy on the cluster.

The vCenter Server decides on which host to deploy the service VMs.

- 13 Click **Save**.

Results

After service deployment, the partner Service Manager is notified about the update.

What to do next

Know deployment details and health status about service instances deployed on hosts. See [Add a Service Profile](#).

Add Redirection Rules for East-West Traffic

Add rules to redirect an east-west traffic for network introspection.

Rules are defined in a policy. Policy as a concept is similar to the concept of sections in firewalls. When you add a policy, select the service chain to redirect the traffic for introspection by service profiles of the service chain.

A rule definition consists of source and destination of the traffic, introspection service, the NSX-T Data Center object to apply the rule to, and traffic redirection policy. After you publish the rule, NSX Manager triggers the rule when a matching traffic pattern is found. The rule begins to introspect the traffic. For example, when NSX Manager classifies a traffic flow that must be introspected, it forwards the traffic to the regular distributed firewall and then to the specified service chain in the policy. The service profiles defined in the service chain introspect the traffic for network services the partner offers. If a service profile finishes introspection without detecting any security issues in the traffic, the traffic is forwarded to the next service profile in the service chain. At the end of the service chain, the traffic is forwarded to the destination target.

All notifications are sent to the partner Service Manager and NSX-T Data Center.

Prerequisites

A service chain is available to redirect the traffic for a network introspection.

Procedure

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

- 2 Verify the NSX Manager is in **Policy** mode.

- 3 Select **Security > East West Security > Network Introspection (E-W) > Add Policy**.

A policy section is similar to a firewall section where you define rules that determine how traffics flows.

- 4 Select a service chain.

- 5 To add a policy, click **Publish**.

- 6 Click the  vertical ellipsis on a section and click **Add Rule**.

- 7 In the **Sources** column, click the edit icon, and select the source of the rule. See [Add a Group](#) for more information.

IPv4, IPv6, and multicast addresses are supported.

- 8 Click **Save**.

- 9 In the **Destinations** column, click the edit icon, and select the destination of the rule. If not defined, the destination matches any. See [Add a Group](#) for more information.

IPv4, IPv6, and multicast addresses are supported.

- 10 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 or resources on ESXi and KVM hosts. It helps in defining a targeted policy for specific zones and tenants, without interfering with other policy defined for other tenants and zones.

Groups consisting of only IP addresses, MAC Addresses, or Active Directory groups cannot be used in the **Applied-to** text box.
- 11 In the Action text box, select **Redirect** to redirect traffic along the service chain or **Do Not Redirect** not to apply network introspection on the traffic.
- 12 Click **Publish**.
- 13 To revert a published rule, select a rule and click **Revert**.
- 14 To add a policy, click **+ Add Policy**.
- 15 To clone a policy or a rule, select the policy or rule and click **Clone**.
- 16 To enable a rule, enable the Enable/Disable icon or select the rule and from the menu click **Enable > Enable Rule**.
- 17 After enabling or disabling a rule, click **Publish** to enforce the rule.

Results

Traffic going to the source is redirected to the service chain for network introspection. After service profiles in the chain introspect the traffic, it is delivered to the destination.

During deployment, it is possible that the VM group membership for a particular policy changes. NSX-T Data Center notifies the partner Service Manager about these updates.

Uninstall an East-West Traffic Introspection Service

Uninstall an east-west traffic introspection service.

As part of uninstalling an east-west service, you need to delete the east-west policy, partner service deployed, service chain, service profile, and service segment.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Verify the NSX Manager is in **Policy** mode.
- 3 To delete a policy, select **Security** → **East West Security** → **Network Introspection (E-W)**.
- 4 Select the east-west policy, click the vertical ellipses, and click **Delete Policy**.
- 5 Click **Publish**.
- 6 To delete a partner service, select **System** → **Service Deployments**.
- 7 Select the partner service, click the vertical ellipses and click **Delete**.

- 8 Click **Delete** to complete the process.
- 9 To delete an east-west service chain, select **Security** → **Settings** → **Network Introspection Settings** → **Service Chain**.
- 10 Select the service chain, click the vertical ellipses and click **Delete**.
- 11 To delete an east-west service profile, select **Security** → **Settings** → **Network Introspection Settings** → **Service Profile**.
- 12 Select the service profile, click the vertical ellipses and click **Delete**.
- 13 To delete an east-west service segment, select **Security** → **Settings** → **Network Introspection Settings** → **Service Segment**.
- 14 Select the service segment, click the vertical ellipses and click **Delete**.

Gateway Firewall

Gateway firewall represents rules applied at the perimeter firewall.

There are predefined categories under the **All Shared Rules** view, where rules across all gateways are visible. Rules are evaluated top down, and left to right. The category names can be changed using the API.

Table 13-4. Categories for Gateway Firewall Rules

Rule Category	Purpose
Emergency	Used for Quarantine. Can also be used for Allow rules.
System	These rules are automatically generated by NSX-T Data Center and are specific to internal control plane traffic, such as, BFD rules, VPN rules and so on. Note Do not edit System rules.
Shared Pre Rules	These rules are globally applied across gateways.
Local Gateway	These rules are specific to a particular gateway.
Auto Service Rules	These are auto-plumbed rules applied to the data plane. You can edit these rules as required.
Default	These rules define the default gateway firewall behavior.

Add a Gateway Firewall Policy and Rule

Implement gateway firewall rules by adding them under a firewall policy section that belongs to a predefined category.

Procedure

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

- 2 Select **Security > North South Security > Gateway Firewall**.
- 3 To enable Gateway Firewall select **Actions > General Settings**, and toggle the status button. Click **Save**.
- 4 Click **Add Policy**, for more about categories see [Gateway Firewall](#).
- 5 Enter a **Name** for the new policy section.
- 6 Select the policy **Destination**.
- 7 Click the gear icon to configure the following policy settings:

Settings	Description
TCP Strict	A TCP connection begins with a three-way handshake (SYN, SYN-ACK, ACK), and typically ends with a two-way exchange (FIN, ACK). In certain circumstances, the firewall may not see the three-way handshake for a particular flow (i.e. due to asymmetric traffic). By default, the firewall does not enforce the need to see a three-way handshake, and will pick-up sessions that are already established. TCP strict can be enabled on a per section basis to turn off mid-session pick-up, and enforce the requirement for a three-way handshake. When enabling TCP strict mode for a particular firewall policy and using a default ANY-ANY Block rule, packets that do not complete the three-way handshake connection requirements and that match a TCP-based rule in this policy section are dropped. Strict is only applied to stateful TCP rules, and is enabled at the gateway firewall policy level. TCP strict is not enforced for packets that match a default ANY-ANY Allow which has no TCP service specified.
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 making changes to the same sections. When locking a section, you must include a comment.

- 8 Click **Publish**. Multiple Policies can be added, and then published together at one time.
The new policy is shown on the screen.
- 9 Select a policy section and click **Add Rule**.
- 10 Enter a name for the rule. IPv4, IPv6, and multicast addresses are supported.
- 11 In the **Sources** column, click the edit icon and select the source of the rule. See [Add a Group](#) for more information.
- 12 In the **Destinations** column, click the edit icon and select the destination of the rule. If not defined, the destination matches any. See [Add a Group](#) for more information.

13 In the **Services** column, click the pencil icon and select services. The service matches any if not defined.

14 In the **Profiles** column, click the edit icon and select a context profile, or click **Add New Context Profile**. See [Add a Context Profile](#).

- Context profiles are not supported on tier-0 gateway firewall policy.
- Gateway firewall rules do not support context profiles with FQDN attributes or other sub attributes.

Context profiles use layer 7 APP ID attributes for use in distributed firewall rules and gateway firewall rules. Multiple App Id context profiles can be used in a firewall rule with services set to **Any**. For ALG profiles (FTP, and TFTP), one context profile is supported per rule.

15 Click **Apply**.

16 The **Applied to** column defines the scope of enforcement per rule and allows users to selectively apply rules to one or more of an uplink interface or a service interface. By default, gateway firewall rules are applied to all the available uplinks and service interfaces on a selected gateway.

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

Option	Description
Allow	Allows all 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.
Reject	Rejects packets with the specified source, destination, and protocol. Rejecting a packet 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. The sending application is notified after one attempt that the connection cannot be established.

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

- 19 Click the gear icon to set logging, direction, IP protocol, and comments.

Option	Description
Logging	Logging can be turned off or on. Logs are stored at /var/log/syslog on the Edge.
Direction	The options are In , Out , and In/Out . The default is In/Out . This field 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	The options are IPv4 , IPv6 , and IPv4_IPv6 . The default is IPv4_IPv6 .

Note Click the graph icon to view the flow statistics of the firewall rule. You can see information such as the byte, packet count, and sessions.

- 20 Click **Publish**. Multiple rules can be added and then published together at one time.
- 21 On each policy section, click the **Info** icon to view the current status of edge firewall rules that are pushed to edge nodes. Any alarms generated when rules were pushed to edge nodes are also displayed.
- 22 To view consolidated status of policy rules that are applied to edge nodes, make the API call.

```
GET https://<policy-mgr>/policy/api/v1/infra/
realized-state/status?intent_path=/infra/domains/default/gateway-policies/
<GatewayPolicy_ID>&include_enforced_status=true
```

URL Analysis Workflow

URL analysis allows administrators to gain insight into the type of websites accessed within the organization, and understand the reputation and risk of the accessed websites.

- URL analysis is available on gateway firewall.
- The management interface IP of edge nodes must have Internet connectivity.
- DNS server must be configured on edge nodes.
- After enabling URL analysis, check that the connection is **UP**, and that the URL database is not 0.0.0.0.

Configure URL analysis and analyze traffic for external websites:

- 1 Ensure that DNS is configured on an edge node. See *Create an NSX Edge Transport Node* in the *NSX-T Data Center Installation Guide*.
- 2 Enable URL Analysis. [URL Settings](#)
- 3 (Optional) Configure custom URL Analysis profiles. [Add a Context Profile](#)
- 4 [Create a Layer 7 DNS Rule](#)
- 5 Generate traffic for external websites.

6 Review the URL Analysis dashboard. [URL Analysis Dashboard](#)

URL Analysis Dashboard

URL Analysis classifies websites into categories, and assigns a reputation score based on their domain.

There are more than 80 pre-defined URL categories. A site or domain can belong to multiple categories. For example, [www.vmware.com](#) belongs to both the **Business and Economy** category, and the **Computer and Internet Info** category. It is not possible to automatically drop or allow traffic based on URL analysis.

The URL Analysis dashboard shows a summary of all analyzed URLs, classified by reputation score and category. Only flows initiated from a Tier-1 gateway are analyzed. Based on their reputation score, URLs are classified into the following severities:

Severity Level	Description
High Risks (1-20) are red	Sites with a high probability of containing malicious links or payloads.
Suspicious (21-40) are orange	Sites with a higher than average probability of containing malicious links or payloads.
Moderate Risks (41-60) are yellow	Generally benign sites that exhibit some characteristics that suggest security risks.
Low Risks (61-80) are gray	Benign sites, that rarely exhibit characteristics that expose the user to security risks.
Trustworthy Sites (81-100) are green	Well-known sites with strong security practices.

On the left is a diagram of the URL distribution. And at the bottom of the page there are additional details about each URI including reputation score, URL, category, and session count.

URL Settings

You enable URL Analysis at the NSX Edge cluster level.

- 1 Navigate to **Security > URL Analysis > Settings**.
- 2 Toggle the **Enable** bar on the edge cluster you want to begin URL analysis on. NSX Edges require access to the Internet to download category and reputation definitions.
- 3 Click **Set** to add a context profile, with a **URL Category** attribute. URL Analysis Profiles specify the categories of traffic to be analyzed. If no profiles are created, all traffic is analyzed. See [Add a Context Profile](#) for more details.
- 4 Configure a Layer 7 gateway firewall rule for DNS traffic, so that URL Analysis can analyze domain information. See [Create a Layer 7 DNS Rule](#).

- Once enabled, you can check the connection status to the cloud service for each NSX Edge Node. You can also verify the version of the URL data used.

Note Fetching the URL data version is not supported if a proxy server is enabled in your environment. NSX Edge must have a direct connection with the cloud provider to fetch the URL data version.

Create a Layer 7 DNS Rule

URL Analysis relies on the configuration of a Layer 7 rule to capture the DNS traffic traversing the NSX Edge cluster.

A Layer 7 rule must be configured on all Tier-1 gateways, backed by the NSX Edge cluster for which you want to analyze traffic. The DNS traffic is analyzed to extract the hostname and IP information from the DNS packets. The extracted information is then used to categorize, and score the traffic.

Prerequisites

A medium-sized edge node (or higher), or a physical form factor edge.

Procedure

- Navigate to **Security > Gateway Firewall** and check that you are on the **All Shared Rules** tab.
- Click **Add Policy** to create a policy section, and give the section a name.
- Select the check box next to the policy and click **Add Rule**.
- Configure the following options:

Option	Description
Name	Name of rule.
Source	Any
Destinations	Any
Services	■ DNS-UDP ■ DNS
Profiles	DNS
Applied to	Select all of the tier-1 gateways backed by the NSX Edge cluster where URL Analysis is enabled.
Action	Allow

- Click **Publish**.

Gateway Firewall Packet Logs

If logging is enabled for a gateway firewall, gateway firewall packets will be logged.

The log file is `/var/log/syslog`. Each log message conforms to the syslog format, and consists of a syslog header and firewall-specific information. For more information about syslog, see [Log Messages and Error Codes](#).

The firewall-specific portion of a log message has the following fields:

Field	Notes
<VRF ID and interface UUID>	You can get this information about an interface by running a CLI command. For example: <pre> edge-1> get firewall interfaces Interface : 55f1af2f-4875-44e9-b0e0-59132ad7753d Type : UPLINK Sync enabled : true Name : Uplink_40_1 VRF ID : 1 ... </pre>
Address family	Possible values: INET, INET6
Reason	Possible values: ■ match: Packet matches a rule. ■ fragment: A fragment that comes after the first fragment. ■ short: Packet too short (for example, no IP header, or TCP/UDP header). ■ normalize: Malformed packets that do not have a correct header or a payload. ■ memory: Datapath out of memory. ■ ip-option: Invalid IP options are present. ■ TERM: A connection is terminated.
Action	Possible values: ■ PASS: Accept the packet. ■ DROP: Drop the packet. ■ NAT: SNAT ■ RDR: DNAT ■ PBR: Service insertion. ■ LB: Load balancer.
Rule ID	The firewall rule ID.
Direction	Possible values: IN, OUT
Packet length	Length in bytes.
Protocol	Possible values: TCP, UDP, or PROTO (protocol number)
Source IP address and port	For SNAT, this is the address before translation.
Destination IP address and port	For DNAT, this is the address before translation.

Examples of gateway firewall log messages for TCP:

```

<181>1 2020-09-21T22:14:12.080427+00:00 lur-svc.nsxedge-ob-16404613-1-gdefw NSX 2802
FIREWALL [nsx@6876 comp="nsx-edge" subcomp="datapathd.firewallpkt" level="INFO"] <1
55f1af2f487544e9:b0e059132ad7753d> INET reason-match PASS 1005 OUT 60 TCP 1.1.1.10/45120-
>91.189.92.38/443 S

```

```
<181>1 2020-09-21T22:14:19.963758+00:00 lur-svc.nsxedge-ob-16404613-1-gdefw NSX 2802
FIREWALL [nsx@6876 comp="nsx-edge" subcomp="datapathd.firewallpkt" level="INFO"] <1
55f1af2f487544e9:b0e059132ad7753d> INET TERM PASS 1005 OUT TCP 1.1.1.10/45120-
>91.189.92.38/443
```

Examples of gateway firewall log messages for UDP:

```
<181>1 2020-09-21T22:05:05.686346+00:00 lur-svc.nsxedge-ob-16404613-1-gdefw NSX 2802
FIREWALL [nsx@6876 comp="nsx-edge" subcomp="datapathd.firewallpkt" level="INFO"] <1
55f1af2f487544e9:b0e059132ad7753d> INET reason-match PASS 1005 IN 328 UDP 40.40.40.10/60613-
>1.1.1.10/42917
```

```
<181>1 2020-09-21T22:05:48.301116+00:00 lur-svc.nsxedge-ob-16404613-1-gdefw NSX 2802
FIREWALL [nsx@6876 comp="nsx-edge" subcomp="datapathd.firewallpkt" level="INFO"] <1
55f1af2f487544e9:b0e059132ad7753d> INET TERM PASS 1005 IN UDP 40.40.40.10/60613-
>1.1.1.10/42917
```

Examples of gateway firewall log messages for PROTO:

```
<181>1 2020-09-21T21:54:38.047682+00:00 lur-svc.nsxedge-ob-16404613-1-gdefw NSX 2802
FIREWALL [nsx@6876 comp="nsx-edge" subcomp="datapathd.firewallpkt" level="INFO"] <1
55f1af2f487544e9:b0e059132ad7753d> INET reason-match PASS 1005 IN 84 PROTO 1 40.40.40.10-
>1.1.1.10
```

```
<181>1 2020-09-21T21:54:45.036957+00:00 lur-svc.nsxedge-ob-16404613-1-gdefw NSX 2802
FIREWALL [nsx@6876 comp="nsx-edge" subcomp="datapathd.firewallpkt" level="INFO"] <1
55f1af2f487544e9:b0e059132ad7753d> INET TERM PASS 1005 IN PROTO 1 40.40.40.10->1.1.1.10
```

Examples of gateway firewall log messages for SNAT:

```
<181>1 2020-09-21T22:57:24.203037+00:00 lur-svc.nsxedge-ob-16404613-1-gdefw NSX 2802
FIREWALL [nsx@6876 comp="nsx-edge" subcomp="datapathd.firewallpkt" level="INFO"] <1
55f1af2f487544e9:b0e059132ad7753d> INET reason-match PASS 1005 OUT 60 TCP 1.1.2.10/49974-
>40.40.40.10/22 S
```

```
<181>1 2020-09-21T22:57:24.203615+00:00 lur-svc.nsxedge-ob-16404613-1-gdefw NSX 2802
FIREWALL [nsx@6876 comp="nsx-edge" subcomp="datapathd.firewallpkt" level="INFO"] <1
55f1af2f487544e9:b0e059132ad7753d> INET reason-match NAT 536870914 OUT 60 TCP 2.2.2.10/37305-
OR 1.1.2.10/49974->40.40.40.10/22 S
```

```
<181>1 2020-09-21T22:57:32.125757+00:00 lur-svc.nsxedge-ob-16404613-1-gdefw NSX 2802
FIREWALL [nsx@6876 comp="nsx-edge" subcomp="datapathd.firewallpkt" level="INFO"] <1
55f1af2f487544e9:b0e059132ad7753d> INET TERM NAT 536870914 OUT TCP 2.2.2.10/37305-OR
40.40.40.10/22->1.1.2.10/49974
```

Examples of gateway firewall log messages for DNAT:

```
<181>1 2020-09-21T22:49:00.978192+00:00 lur-svc.nsxedge-ob-16404613-1-gdefw NSX 2802
FIREWALL [nsx@6876 comp="nsx-edge" subcomp="datapathd.firewallpkt" level="INFO"] <1
55f1af2f487544e9:b0e059132ad7753d> INET reason-match RDR 536870913 IN 60 TCP
40.40.40.10/40082->10.10.10.1/22-OR 1.1.1.10/22 S
```

```
<181>1 2020-09-21T22:50:01.915154+00:00 lur-svc.nsxedge-ob-16404613-1-gdefw NSX 2802
FIREWALL [nsx@6876 comp="nsx-edge" subcomp="datapathd.firewallpkt" level="INFO"] <1
55f1af2f487544e9:b0e059132ad7753d> INET TERM RDR 536870913 IN TCP 40.40.40.10/40082-
>10.10.10.1/22-OR 1.1.1.10/22
```

North-South Network Security - Inserting Third-party Service

NSX-T Data Center provides the functionality to insert third-party services at tier-0 or tier-1 router in the data center to redirect traffic to the third-party service for introspection. Only ESXi hosts are supported to deploy north-south service VMs. KVM hosts are not supported.

High-Level Tasks for North-South Network Security

Follow these steps to set up network security for north-south traffic.

Table 13-5. List of Tasks to Configure North-South Network Introspection

Workflow Tasks	Persona	Implementation
Register Service with NSX-T Data Center	Partner	Only API
Deploy a Service for North-South Traffic Introspection	Administrator	API and NSX-T Data Center UI
Add Redirection Rules for North-South Traffic	Administrator	API and NSX-T Data Center UI

Deploy a Service for North-South Traffic Introspection

After you register a service, you must deploy an instance of the service on an NSX-T transport node for the service to start processing network traffic.

Deploy partner service VM at tier-0 or tier-1 logical router that acts as a gateway between the physical world and the logical network on vCenter Server. After you deploy the SVM as a standalone service instance or an active-standby service instance, you can create redirection rules to redirect traffic to the SVM for network introspection.

Prerequisites

- All hosts are managed by a vCenter Server.
- Partner services are registered with NSX-T Data Center and are ready for deployment.
- NSX-T Data Center administrators can access partner services.
- High Availability mode for logical router must be in active-standby mode.
- Turn on the Distributed Resource Scheduler utility.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **System > Service Deployments > Deployment**.
- 3 From the Partner Service drop-down menu, select the service that must be deployed.
- 4 Click **Deploy Service**. Enter details to deploy the service.

Table 13-6. Partner Service Details

Field	Description
Service Deployment Name	Enter a name to identify the service instance.
Deployment Specification	Select the form factor to deploy.
Attachment Points	Select the tier-0 or tier-1 logical router where the service instance must be deployed.
Failure Policy	Select Allow or Block .
Network	For a deployment of the type Active Standby, set value to the following fields: ■ Primary Interface Network: The interface to be used by the deployed service. ■ Primary Interface IP: Enter the IP address to be used by the service instance. ■ Primary Gateway Address: Enter the gateway address. ■ Primary Subnet Mask: Enter the subnet mask. ■ Secondary Interface Network: The standby interface that is used if the primary interface is unavailable. ■ Secondary Interface IP: Enter the IP address for the standby IP that is used if the primary IP is unavailable. ■ Secondart Gateway Address: Enter the standby gateway address that is used if the primary gateway is unavailable. ■ Secondary Subnet Mask: Enter the standby subnet mask that is used if the primary subnet mask is unavailable. For a deployment of the type Standalone, set values to the primary interfaces.
Compute Manager	Select the registered vCenter Server.
Datastore	Select the repository to store service instance data.

Table 13-6. Partner Service Details (continued)

Field	Description
Deployment Mode	Select Standalone to deploy a single service instance at the tier-0 or tier-1 logical router. Select Active Standby to deploy a couple of service instances in active-standby mode at the tier-0 or tier-1 logical router.
Deployment Template	Select the template to be used during deployment of the service instance.

5 Click **Save**.

Results

The Service Instances tab displays the deployment progress. It might take a few minutes for deployment to finish. Verify the deployment state to ensure that the service instance is successfully deployed at the tier-0 or tier-1 logical router.

Alternatively, go to the vCenter Server and verify the deployment status.

What to do next

Add redirection rules for north-south traffic. See [Add Redirection Rules for North-South Traffic](#).

Add Redirection Rules for North-South Traffic

Set up redirection rules to send traffic to third-party services inserted at a Tier-0 or Tier-1 router.

Prerequisites

- Register and deploy third-party services on NSX-T.
- Configure Tier-0 or Tier-1 router.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Security > North South Security > Network Introspection (N-S) > Add Policy**.
A policy section is similar to a firewall section where you define rules that determine how traffics flows.
- 3 Set **Redirection To** field for a service instance or a service chain to a Tier-0 or Tier-1 logical router to perform network introspection of traffic flowing between source and destination entities.
- 4 To add a policy, click **Publish**.
- 5 Click the  vertical ellipsis on a section and click **Add Rule**.

- 6 Edit the **Source** field to add a group by defining membership criteria, static members, IP/MAC addresses, or active directory groups. Membership criteria can be defined from one of these types: Virtual Machine, Logical Switch, Logical Port, IP Set. You can select static members from one of these categories: Group, Segment, Segment Port, Virtual Network Interface, or Virtual Machine.
- 7 Click **Save**.
- 8 To add a destination group, edit the **Destination** field.
- 9 In the **Applied To** field, you can do one of the following:
 - For a service inserted at Tier-0 logical router, select the uplink of Tier-0 router.
 - For a service inserted at Tier-1 logical router, you do not need to select any uplinks.
- 10 Each rule can be enabled individually. After you enable a rule, it is applied to the traffic that matches the rule.
- 11 Click **Advanced Settings** to configure the traffic direction and to enable logging.
- 12 In the Action field, select **Redirect** to redirect traffic along the service instance or **Do Not Redirect** not to apply network introspection on the traffic.
- 13 Click **Publish**.
- 14 To revert a published rule, select a rule and click **Revert**.
- 15 To add a policy, click **+ Add Policy**.
- 16 To clone a policy or a rule, select the policy or rule and click **Clone**.
- 17 To enable a rule, enable the Enable/Disable icon or select the rule and from the menu click **Enable > Enable Rule**.
- 18 After enabling or disabling a rule, click **Publish** to enforce the rule.

Results

Based on the actions set, north-south traffic is redirected to the service instance for network introspection.

Uninstall a North-South Traffic Introspection Service

Uninstall a north-south traffic introspection service.

Delete a policy and a partner service that was deployed for north-south introspection service.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Verify the NSX Manager is in **Policy** mode.
- 3 To delete a policy, select **Security** → **North South Security** → **Network Introspection (N-S)**.

- 4 Select the policy, click the vertical ellipses, and click **Delete Policy**.
- 5 Click **Publish**.
- 6 To delete a partner service, **System** → **Service Deployments**.
- 7 Select the service, click the vertical ellipses and click **Delete**.

Endpoint Protection

NSX-T Data Center allows you to insert third-party partner services as a separate service VM that provides Endpoint Protection services. A partner Service VM processes file, process, and registry events from the guest VM based on the endpoint protection policy rules applied by the NSX-T Data Center administrator.

Understand Endpoint Protection

Know the use case, workflow, and key concepts of endpoint protection.

Endpoint Protection Use Case

In a virtual environment, use the guest introspection platform to provide antivirus and antimalware protection to guest VMs.

As an NSX administrator, you implement an antivirus and antimalware solution that is deployed as a Service Virtual Machine (Service VM, or SVM) to monitor a file, network, or process activity on a guest VM. Whenever a file is accessed, such as a file open attempt, the antimalware Service VM is notified of the event. The Service VM then determines how to respond to the event. For example, to inspect the file for virus signatures.

- If the Service VM determines that the file contains no viruses, then it allows the file open operation to succeed.
- If the Service VM detects a virus in the file, it requests the Thin Agent on the guest VM to act in one of the following ways:
 - Delete the infected file or deny access to the file.
 - Infected VMs can be assigned a tag by NSX. Moreover, you can define a rule that automatically moves such tagged guest VMs to a security group that quarantines the infected VM for additional scan and isolation from the network until the infection is completely removed.

The benefits of using the guest introspection platform to protect guest VM endpoints:


- Reduced consumption of compute resources: Guest introspection offloads virus signatures and security scanning logic from each endpoint on a host to a third-party partner Service VM on the host. As virus scanning happens only on the Service VM, there is no need to spend compute resources on guest VMs to run virus scans.

- Better management: As virus scans are offloaded to a Service VM, virus signatures need to be updated to only one object per host. Such a mechanism works better than agent-based solution where same virus signatures need updates on all guest VMs.
- Continuous antivirus and antimalware protection: As the Service VM runs continuously, a guest VM is not mandated to run the latest virus signatures. For example, a snapshot VM might run some older version of the virus signature making it vulnerable in the traditional way of protecting endpoints. With the guest introspection platform, the Service VM is continuously running the latest virus and malware signatures thereby ensuring that any newly added VM is also protected with the latest virus signatures.
- Offloaded virus signatures to a Service VM: Virus database lifecycle is outside of guest VM lifecycle and so the Service VM is not affected by guest VM outages.

Endpoint Protection Architecture

Understand the architecture of service insertion and endpoint protection components (guest introspection) in NSX-T Data Center.

Figure 13-1. Endpoint Protection Architecture



Key Concepts:

- Partner console: It is the web application provided by the security vendor to work with the guest introspection platform.
- NSX Manager: It is the management plane appliance for NSX that provides API and graphical user interface to customers and partners for configuration of Network and Security policies. For guest introspection, the NSX Manager also provides API and GUI to deploy and manage partner appliances.
- Guest Introspection SDK: VMware provided library consumed by the security vendor.

- **Service VM:** Is the security vendor provided VM that consumes the guest introspection SDK provided by VMware. It contains the logic to scan file or process events to detect virus or malware on the guest. After scanning a request, it sends back a verdict or notification about the action taken by the guest VM on the request.
- **Guest Introspection host agent (Context Multiplexer):** It processes configuration of endpoint protection policies. It also multiplexes and forwards messages from protected VMs to the Service VM. It reports the health status of the guest introspection platform and maintains records of the Service VM configuration in the `muxconfig.xml` file.
- **Ops agent (Context engine and Guest Introspection client):** It forwards the guest introspection configuration to the guest introspection host agent (Context Multiplexer). It also relays the health status of the solution to NSX Manager.
- **EAM:** NSX Manager uses the ESXi agent manager to deploy a partner Service VM on every host on the cluster configured for protection.
- **Thin agent:** It is the file or network introspection agent running inside the guest VMs. It also intercepts file and network activities that are forwarded to the Service VM through the host agent. This agent is part of VMware Tools. It replaces the traditional agent provided by antivirus or antimalware security vendors. It is a generic and lightweight agent that facilitates offloading files and processes for scanning to the Service VM provided by the vendor.

Key Concepts of Endpoint Protection

The endpoint protection workflow needs partners to register their services with NSX-T Data Center and an administrator to consume these services. There are a few concepts that aid your understanding of the workflow.

- **Service Definition:** Partners define services with these attributes: name, description, supported form factors, deployment attributes that include network interfaces, and appliance OVF package location to be used by the SVM.
- **Service Insertion:** NSX provides the service insertion framework that allows partners to integrate networking and security solutions with the NSX platform. Guest introspection solution is one such form of service insertion.
- **Service Profiles and Vendor Templates:** Partners register vendor templates which expose protection levels for policies. For example, protection levels can be Gold, Silver, or Platinum. Service Profiles can be created from Vendor Templates, which allow the NSX administrators to name the Vendor Templates according to their preference. For services other than those of Guest Introspection, the Service Profiles allow further customization using attributes. The Service Profiles can then be used in the Endpoint Protection policy rules to configure protection for virtual machine groups defined in NSX. As an administrator, you can create groups based on VM name, tags, or identifiers. Multiple Service Profiles can optionally be created from a single Vendor Template.

- **Endpoint Protection Policy:** A policy is a collection of rules. When you have multiple policies, arrange them in the order to run them. The same applies for rules defined within a policy. For example, policy A has three rules, and policy B has four rules, and they are arranged in a sequence such that policy A precedes policy B. When guest introspection begins running policies, rules from policy A are run first before rules from policy B.
- **Endpoint Protection Rule:** As a NSX administrator, you can create rules that specify the virtual machine groups that are to be protected, and choose the protection level for those groups by specifying the Service Profile for each rule.
- **Service Instance:** It refers to the service VM on a host. The service VMs are treated as special VMs by vCenter and they are started before any of the guest VMs are powered on and stopped after all the guest VMs are powered off. There is one service instance per service per host.

Important Number of service instances is equal to the number of hosts on which the service is running host. For example, if you have eight hosts in a cluster, and the partner service was deployed on two clusters, the total number of service instances running are 16 SVMs.

- **Service Deployment:** As an admin you deploy partner Service VMs through NSX-T on a per cluster basis. Deployments are managed at a cluster level, so that when any host is added to the cluster, EAM automatically deploys the service VM on them.

Automatically deploying the SVM is important because if distributed resource scheduler (DRS) service is configured on a vCenter Cluster, then vCenter can rebalance or distribute existing VMs to any new host that got added to the cluster after the SVM is deployed and started on the new host. Since partner Service VMs need NSX-T platform to provide security to guest VMs, the host must be prepared as a transport node.

Important One service deployment refers to one cluster on the vCenter Server that is managed for deploying and configuring one partner service.

- **File Introspection driver:** Is installed on the guest VM, intercepts the file activity on the guest VM.
- **Network Introspection driver:** Is installed on the guest VM, intercepts the network traffic, process, and user activity on the guest VM.

High-level Tasks for Endpoint Protection

Third-party partners services containing security scanning logic, are registered with NSX-T Data Center for guest VM protection. The partner service is enforced when the NSX admin deploys the registered services and applies end point protection policies to guest VM groups.

The guest introspection workflow for the endpoint protection use case is as follows:

Figure 13-2. Endpoint Protection Workflow

Workflow Tasks	Role/Persona	Implementation
Register a Service with NSX-T Data Center	Partner Admin	Partner Console
Register a Service with NSX-T Data Center	Partner Admin	Partner Console
Register a Service with NSX-T Data Center	Partner Admin	Partner Console
Deploy a Service	NSX Admin	API and NSX Manager UI
View Service Instance Details	NSX Admin	API and NSX Manager UI
Bring up Service Instance	NSX Admin	API and NSX Manager UI
Add a Service Profile	NSX Admin	API and NSX Manager UI
Consume Guest Introspection Policy	NSX Admin	API and NSX Manager UI
Add and Publish Endpoint Protection Rules	NSX Admin	API and NSX Manager UI
Monitor Endpoint Protection Status	NSX Admin	API and NSX Manager UI

Configure Endpoint Protection

Protect guest VMs running in an NSX-T Data Center environment using third-party partner security services.

The high-level steps to configure endpoint protection policies:

- 1 Ensure [Prerequisites to Configure Endpoint Protection](#) are met before you configure endpoint protection on guest VMs.
- 2 Supported software. See [Supported Software](#).
- 3 Install File Introspection Driver for Linux VMs. See [Install the Guest Introspection Thin Agent on Linux Virtual Machines](#).
- 4 Install File Introspection Driver for Windows VMs. See [Install the Guest Introspection Thin Agent on Windows Virtual Machines](#).
- 5 Install Network Introspection Driver for Linux VMs. See [Install the Linux Thin Agent for Network Introspection](#).
- 6 Create a User with Guest Introspection Partner Admin Role. See [Create a User with Guest Introspection Partner Admin Role](#).
- 7 Register partner service with NSX-T Data Center. Refer to Partner documentation.
- 8 Deploy a service. See [Deploy a Service](#).
- 9 Consume Guest Introspection Policy. See [Consume Guest Introspection Policy](#).
- 10 Add and Publish Endpoint Protection Rules. See [Add and Publish Endpoint Protection Rules](#).
- 11 Monitor endpoint protection rules. See [Monitor Endpoint Protection Status](#).

Prerequisites to Configure Endpoint Protection

Before you configure endpoint protection for guest VMs, ensure that the prerequisites are met.

Prerequisites

- NSX Manager is installed on all the hosts.
- Prepare and configure NSX-T Data Center cluster as transport nodes by applying transport node profiles. After the host is configured as the transport node, guest introspection components are installed. See *NSX-T Data Center Installation Guide*.
- Partner console is installed and configured to register services with NSX-T Data Center.
- Ensure that the guest VMs run VM Hardware Configuration file version 9 or higher.
- Configure VMware Tools and install thin agents.
 - See [Install the Guest Introspection Thin Agent on Linux Virtual Machines](#).
 - See [Install the Guest Introspection Thin Agent on Windows Virtual Machines](#).
 - See [Install the Linux Thin Agent for Network Introspection](#).

Install the Guest Introspection Thin Agent on Linux Virtual Machines

Guest Introspection supports File Introspection in Linux for anti-virus only. To protect Linux VMs using a Guest Introspection security solution, you must install the Guest Introspection thin agent.

The Linux thin agent is available as part of the operating system specific packages (OSPs). The packages are hosted on VMware packages portal. Enterprise or Security Administrator (non-NSX Administrator) can install the agent on guest VMs outside of NSX.

Installing VMware Tools is not required.

Based on your Linux operating system, perform the following steps with root privilege:

Prerequisites

- Ensure that the guest virtual machine has a supported version of Linux installed:
 - Red Hat Enterprise Linux (RHEL) 7.4 (64 bit) GA
 - SUSE Linux Enterprise Server (SLES) 12 (64 bit) GA
 - Ubuntu 16.04.5 LTS (64 bit) GA
 - CentOS 7.4 GA
- Verify GLib 2.0 is installed on the Linux VM.

Procedure**1 For Ubuntu systems**

- a Obtain and import the VMware packaging public keys using the following commands.

```
curl -O https://packages.vmware.com/packages/nsx-gi/keys/VMWARE-PACKAGING-NSX-GI-GPG-
RSA-KEY.pub
apt-key add VMWARE-PACKAGING-NSX-GI-GPG-RSA-KEY.pub
```

- b Create a new file named `vmware.list` file under `/etc/apt/sources.list.d`
- c Edit the file with the following content:

```
deb [arch=amd64] https://packages.vmware.com/packages/nsx-gi/latest/ubuntu/ xenial main
```

- d Install the package.

```
apt-get update
apt-get install vmware-nsx-gi-file
```

2 For RHEL7 systems

- a Obtain and import the VMware packaging public keys using the following commands.

```
curl -O https://packages.vmware.com/packages/nsx-gi/keys/VMWARE-PACKAGING-NSX-GI-GPG-
RSA-KEY.pub
rpm --import VMWARE-PACKAGING-NSX-GI-GPG-RSA-KEY.pub
```

- b Create a new file named `vmware.repo` file under `/etc/yum.repos.d`.
- c Edit the file with the following content:

```
[vmware]
name = VMware
baseurl = https://packages.vmware.com/packages/nsx-gi/latest/rhel/x86_64
enabled = 1
gpgcheck = 1
metadata_expire = 86400
ui_repoid_vars = basearch
```

- d Install the package.

```
yum install vmware-nsx-gi-file
```

3 For SLES systems

- a Obtain and import the VMware packaging public keys using the following commands.

```
curl -O https://packages.vmware.com/packages/nsx-gi/keys/VMWARE-PACKAGING-NSX-GI-GPG-
RSA-KEY.pub
rpm --import VMWARE-PACKAGING-NSX-GI-GPG-RSA-KEY.pub
```

- b Add the following repository:

```
zypper ar -f "https://packages.vmware.com/packages/nsx-gi/latest/sle12/x86_64/" VMware
```

- c Install the package.

```
zypper install vmware-nsx-gi-file
```

4 For CentOS systems

- a Obtain and import the VMware packaging public keys using the following commands.

```
curl -O https://packages.vmware.com/packages/nsx-gi/keys/VMWARE-PACKAGING-NSX-GI-GPG-
RSA-KEY.pub
rpm --import VMWARE-PACKAGING-NSX-GI-GPG-RSA-KEY.pub
```

- b Create a new file named `vmware.repo` file under `/etc/yum.repos.d`.

- c Edit the file with the following content:

```
[vmware]
name = VMware
baseurl = https://packages.vmware.com/packages/nsx-gi/latest/centos/x86_64
enabled = 1
gpgcheck = 1
metadata_expire = 86400
ui_repoid_vars = basearch
```

- d Install the package.

```
yum install vmware-nsx-gi-file
```

What to do next

Verify whether the thin agent is running using the service `vsepd` `status` command with the administrative privileges. The status must be running.

Install the Linux Thin Agent for Network Introspection

Install the Linux thin agent to introspect network traffic.

Important To protect guest VMs against antivirus, you do not need to install the Linux thin agent for network introspection.

The Linux thin agent driver that is used to introspect network traffic depends on an open-source driver.

Prerequisites

Install the following packages:

- glib2
- libnetfilter-contrack3/ libnetfilter-contrack
- libnetfilter-queue1/ libnetfilter-queue
- iptables

Procedure

- 1 To install the open-source driver provided by guest introspection.

- a Add following URL as the base URL for your operating system.

```
deb [arch=amd64] https://packages.vmware.com/guest-introspection-for-vmware-nsx/latest/
ubuntu xenial main
```

- b Import the VMware packaging key.

```
https://packages.vmware.com/packages/nsx-gi/keys/VMWARE-PACKAGING-NSX-GI-GPG-RSA-
KEY.pub
```

- c Update the repository and install the open-source driver.

```
apt-get install Guest-Introspection-for-VMware-NSX
```

- 2 To install the Linux thin agent that is used to introspect file and or network traffic.

- To install file and network introspection packages, select `vmware-nsx-gi` package in step c.
- To install network introspection packages, select the `vmware-nsx-gi-net` package in step c.

- a Add following URL as the base URL for your operating system.

```
deb [arch=amd64] https://packages.vmware.com/packages/nsx-gi/latest
```

- b Import the VMware packaging key.

```
https://packages.vmware.com/packages/nsx-gi/keys/VMWARE-PACKAGING-NSX-GI-GPG-RSA-
KEY.pub
```

- c Install one of the drivers.

```
vmware-nsx-gi
vmware-nsx-gi-net
```

Install the Guest Introspection Thin Agent on Windows Virtual Machines

To protect VMs using a Guest Introspection security solution, you must install Guest Introspection thin agent, also called Guest Introspection drivers, on the VM. Guest Introspection drivers are included with VMware Tools for Windows, but are not part of the default installation. To install Guest Introspection on a Windows VM, you must perform a custom install and select the drivers.

Windows virtual machines with the Guest Introspection drivers installed are automatically protected whenever they are started up on an ESXi host that has the security solution installed. Protected virtual machines retain the security protection through shut downs and restarts, and even after a vMotion move to another ESXi host with the security solution installed.

- If you are using vSphere 6.0, see these instructions for installing VMware Tools, see [Manually Install or Upgrade VMware Tools in a Windows Virtual Machine](#).
- If you are using vSphere 6.5, see these instructions for installing VMware Tools: <https://www.vmware.com/support/pubs/vmware-tools-pubs.html>.

Prerequisites

Ensure that the guest virtual machine has a supported version of Windows installed. The following Windows operating systems are supported for NSX Guest Introspection:

- Windows XP SP3 and above (32 bit)
- Windows Vista (32 bit)
- Windows 7 (32/64 bit)
- Windows 8 (32/64 bit)
- Windows 8.1 (32/64) (vSphere 6.0 and later)
- Windows 10
- Windows 2003 SP2 and above (32/64 bit)
- Windows 2003 R2 (32/64 bit)
- Windows 2008 (32/64 bit)
- Windows 2008 R2 (64 bit)
- Win2012 (64)
- Win2012 R2 (64) (vSphere 6.0 and later)
- Windows Server 2016
- Windows Server 2019

Procedure

- 1 Start the VMware Tools installation, following the instructions for your version of vSphere. Select **Custom install**.

2 Expand the VMCI Driver section.

The options available vary depending on the version of VMware Tools.

3 Select the driver to be installed on the VM.

Driver	Description
vShield Endpoint Drivers	Installs File Introspection (<code>vsepflt</code>) and Network Introspection (<code>vnetflt</code>) drivers.
Guest Introspection Drivers	Installs File Introspection (<code>vsepflt</code>) and Network Introspection (<code>vnetflt</code>) drivers.
NSX File Introspection Driver and NSX Network Introspection Driver	Select NSX File Introspection Driver to install <code>vsepflt</code>. Optionally select NSX Network Introspection Driver to install <code>vnetflt</code> (<code>vnetWFP</code> on Windows 10 or later). Note Select NSX Network Introspection Driver only if you are using the Identity Firewall or Endpoint Monitoring features.

4 In the drop-down menu next to the drivers you want to add, select This feature is installed on the local hard drive.

5 Follow the remaining steps in the procedure.

What to do next

Verify whether the thin agent is running using the `fltmc` command with the administrative privileges. The Filter Name column in the output lists the thin agent with an entry `vsepflt`.

Supported Software

Guest Introspection is interoperable with specific versions of software.

VMware Tools

VMware Tool 10.3.10 version is supported.

Check out interoperability between VMware Tools and NSX-T. See [VMware Product Interoperability Matrices](#).

Supported OS

- Windows 7
- Windows 8/8.1
- Windows 10
- Windows 2008 server R2
- Windows 2012 server R2
- Windows 2016 Server
- CentOS 7.4 GA
- RHEL 7.4 GA

- Ubuntu 16.04.5 LTS (64 bit)
- SLES 12 GA

Supported Hosts

For supported ESXi hosts, see the [VMware Product Interoperability Matrices](#).

Create a User with Guest Introspection Partner Admin Role

Assign a user with the Guest Introspection Partner Admin role that is available in NSX-T Data Center.

Note: It is recommended to register partner services by a user that is associated with the Guest Introspection Partner Admin role to avoid any security issues.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **System** → **User** → **Role Assignments**.
- 3 Click **Add**.
- 4 Select the user and assign that user the **GI Partner Admin** role.

What to do next

Register services with NSX-T Data Center. See [Register a Service with NSX-T Data Center](#).

Register a Service with NSX-T Data Center

Register third-party security services with NSX-T Data Center.

Prerequisites

- Ensure that prerequisites are met. See [Prerequisites to Configure Endpoint Protection](#).
- Ensure that a vIDM user is assigned the GI Partner Admin role. This role is used to register services with NSX-T Data Center.

Procedure

- 1 Log in with the GI Partner Admin privileges to the partner console.
- 2 Register a service, vendor template, and configure the partner solution with NSX-T Data Center. See partner documentation.

What to do next

View catalog of partner services. See [View Catalog of Partner Services](#).

View Catalog of Partner Services

The catalog page displays all the partners and their services that are registered with NSX-T Data Center.

Prerequisites

- Partners register services with NSX-T Data Center.
- Services are deployed on a cluster.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **System > Service Deployments > Catalog**.
- 3 Click **View** on a service. The Deployment page displays the details about the service, such as status of deployment, network details, cluster details, and so on.

What to do next

Upgrade a partner service VM.

Deploy a Service

After you register a service, you must deploy an instance of the service for the service to start processing network traffic.

Deploy partner service VMs that run the partner security engine on all the NSX-T Data Center hosts in a cluster. The vSphere ESX Agency Manager (EAM) service is used to deploy the partner service VMs on each host. After you deploy the SVMs, you can create policy rules used by SVM to protect guest VMs.

Prerequisites

- All hosts are managed by a vCenter Server.
- Partner services are registered with NSX-T Data Center and are ready for deployment.
- NSX-T Data Center administrators can access partner services and vendor templates.
- Both the service VM and the partner Service Manager (console) must be able to communicate with each other at the management network level.
- Prepare hosts as NSX-T Data Center transport nodes:
 - Create a transport zone.
 - Create an IP pool for tunnel endpoint IP addresses.
 - Create an uplink profile.
 - Add a transport node profile to prepare a cluster for auto deployment of NSX-T Data Center transport nodes.

- Configure a standalone or managed host.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Go to the **System** tab and click **Service Deployment**.
- 3 From the Partner Service drop-down, select the service to be deployed.
- 4 Click **Deployment** and click **Deploy Service**.
- 5 Enter the service deployment name.
- 6 In the Compute Manager field, select the compute resource on the vCenter Server to deploy the service.
- 7 In the Cluster field, select the cluster where the services need to be deployed.
- 8 In the Data Store drop-down menu, you can:
 - a Select a datastore as the repository for the service virtual machine.
 - b Select **Specified on Host**. This setting means that you do not need to select a datastore and port group on this wizard. You can directly configure agent settings on EAM in vCenter Server to point to a specific datastore and port group to be used for service deployment.

To know how to configure EAM, refer to the vSphere documentation.
- 9 In the Network column, click **Set**.
- 10 Set the Management Network interface to **Specified on Host** or **DVPG**.
- 11 Set the network type to DHCP or Static IP pool. If you set the network type to Static IP pool, select from the list of available IP pools.
- 12 In the Deployment Specification field, select host-based deployment to deploy service on all hosts. Depending upon the services registered by the partner, multiple services can be deployed as part of a single service VM.
- 13 In the Deployment Template field, select the registered deployment template.
- 14 Click **Save**.

Results

When a new host is added to the cluster, EAM automatically deploys the service VM on the new host. The deployment process might take some time, depending on the vendor's implementation. You can view the status in the NSX Manager user interface. The service is successfully deployed on the host when the status turns `Deployment Successful`.

To remove host from a cluster, first move it into maintenance mode. Then, select the option to migrate the guest VMs to another host to complete migration.

What to do next

Know deployment details and health status about service instances deployed on hosts. See [View Service Instance Details](#).

View Service Instance Details

Know deployment details and health status of service instance deployed on member hosts of a cluster.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **System > Service Deployments > Service Instances**.
- 3 From the Partner Service drop-down menu, select the partner service to view details related to service instances.

Table 13-7.

Field	Description
Service Instance Name	A unique ID identifying the service instance on a particular host.
Service Deployment Name	The name you entered when deploying the service.
Deployed To	Host IP address or FQDN
Deployment Mode	Cluster or Standalone
Deployment Status	Up status to determine a successful deployment
Health Status	When the service instance is deployed, the health status is Ready. To bring the health status from Ready to Up, make the required configuration changes. See Bring up Service Instance. After the following parameters are successfully realized by NSX-T Data Center, the health status changes from Ready to Up. ■ Solution status: Up ■ Connectivity between NSX-T Data Center Guest Introspection agent and NSX-T Data Center Ops Agent: Up ■ Health Status received at: <Day, Date, Time>

What to do next

Bring up Service Instance. See [Bring up Service Instance](#).

Bring up Service Instance

After deploying the service instance, certain parameters need to be realized in NSX-T Data Center for the health status to be Up.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **System > Service Deployments > Service Instances**.
- 3 From the Partner Service drop-down menu, select the partner service to view details related to service instances.
- 4 The Health Status column displays state of the service instance as `Ready`. It indicates that the service instance is ready to be configured with endpoint protection policy rules to protect VMs.
- 5 The following parameters must be realized in NSX-T Data Center for the health status to change to `Up`.
 - Guest virtual machines must be available on the host.
 - Guest virtual machines must be powered on.
 - Endpoint protection rules must be applied to the guest virtual machines.
 - Guest virtual machines must be configured with the supported version of VMtools and file introspection drivers.

What to do next

Add a service profile. See [Add a Service Profile](#).

Add a Service Profile

Guest introspection policies can be implemented only when a service profile is available in NSX-T Data Center. Service profiles are created from a template provided by the partner. Service Profiles are a way for the administrator to choose protection levels (Gold, Silver, Platinum policy) for a VM by choosing the vendor templates provided by the vendor.

For example, a vendor can provide Gold, Platinum, and Silver policy levels. Each profile created might serve a different type of workload. A Gold service profile provides complete antimalware to a PCI-type workload, while a silver service profile only provides basic antimalware protection to a regular workload.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Security > Endpoint Protection > Endpoint Protection Rules > Service Profiles**.

- 3 From the Partner Service field, select the service for which you want to create a service profile.
- 4 Click **Add Service Profile**.
- 5 Enter the service profile name and select the vendor template. Optionally, add description and tags.
- 6 Click **Save**.

The vendor template ID used to create the service profile is passed on to the partner console. Partners store the vendor template ID to track usage of which guest VMs are protected by these vendor template.

Results

After creating service profile, an NSX admin creates rules to associate a service profile to a group of VMs before publishing the policy rule.

What to do next

Apply endpoint protection policy on guest VM groups that need to be protected from malware. See [Consume Guest Introspection Policy](#).

Consume Guest Introspection Policy

Policy can be enforced on VM groups by creating rules that associate service profiles with VM groups. Protection begins immediately after rules are applied to a VM group.

The endpoint protection policy is a protection service offered by partners to protect guest VMs from malware by implementing service profiles on guest VMs. With a rule applied to a VM group, all guest VMs within that group are protected by that service profile. When a file access event on a guest VM occurs, the GI thin agent (running on each guest VM) collects context of the file (file attributes, file handle, and other context details) and notifies the event to SVM. If the SVM wants to scan the file content, it request for details using the EPSec API library. Upon a clean verdict from SVM, the GI thin agent allows the user to access the file. In case SVM reports the file as infected, the GI thin agent denies user access to the file.

To execute an security service on a VM group, you need to:

Procedure

- 1 Define policy and rules.
- 2 Define membership criteria to form VM group.
- 3 Define rules for VM groups.
- 4 Publish the rule.

Add and Publish Endpoint Protection Rules

Publishing policy rules to VM groups means associating VM groups that need to be protected with a specific service profile.

Procedure

- 1 In the policy section, select the policy section.
- 2 Click **Add** -> **Add Rule**.
- 3 In the new rule, enter the rule name.
- 4 In the Select Groups field, click the Edit icon.
- 5 In the Set Groups window, select from the existing list of groups or add a new group.
 - a To add a new group, click **Add Group**, enter details and click **Save**.
See [Add a Group](#).
- 6 In the Group column, select the VM group.
- 7 In the Service Profiles column, select the service profile that provides the desired protection level to the guest VMs in the group.
 - a To add a new service profile, click **Add Service Profile**, enter details and click **Save**.
See [Add a Service Profile](#).
- 8 Click **Publish**.

Results

Endpoint protection policies protect VM groups.

What to do next

You might want to change the sequence of rules depending on the type of protection required for different VM groups. See [How Guest Introspection Runs Endpoint Protection Policy](#)

Monitor Endpoint Protection Status

Monitor the configuration status of protected and unprotected VMs, issues with Host agent and service VMs, and VMs configured with the file introspection driver that was installed as part of the VMtools installation.

You can view:

- View Service Deployment Status.
- View Configuration Status of Endpoint Protection.
- View Capacity Status Set for Endpoint Protection.

View Service Deployment Status

View service deployment details on the Monitoring Dashboard.

View the system-wide status of EPP policy.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Navigate to **Home > Monitoring - Dashboards**.
- 3 From the drop-down menu, click **Monitoring - System**.
- 4 To view the deployment status across clusters in the system, navigate to the Endpoint Protection widget, click the doughnut chart to view successful or unsuccessful deployments.

The Service Deployments page displays the deployment details.

View Configuration Status of Endpoint Protection

View configuration status of the endpoint protection service.

View the system-wide status of EPP policy.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Navigate to **Home > Security > Security Overview**.
- 3 To view status of EPP on clusters, click the Security widget.
- 4 In the Security Overview page, click **Configuration**.



- 5 In the Endpoint Protection section, view:

a VM Distribution by Service Profile widget displays:

- 1 Number of VMs protected by top profile. Top profile represents a profile that protects the maximum number of VMs on a cluster.
- 2 VMs protected by remaining service profiles categorized under Other Profiles.
- 3 VMs not protected categorized under No Profile.

The Endpoint Protection Rules page displays VMs protected by Endpoint Protection policies.

b Components having issues widget displays:

- 1 Host: Issues related to the context multiplexer.
- 2 SVM: Issues related to service VMs. For example, the SVM state is down, SVM connection with guest VM is down.

The Status column on the Deployment page displays health issues.

c Configure VMs running File Introspection widget displays:

- 1 VMs protected by File Introspection driver.
- 2 VMs where the File Introspection driver status is unknown.

ESXi Agency Manager (EAM) attempts to resolve a few issues related to hosts, SVMs, and configuration errors. See [Resolve Partner Services Issues](#).

View Capacity Status Set for Endpoint Protection

View capacity status of the endpoint protection service.

View the capacity status of EPP policy.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Navigate to **Home > Monitoring - Dashboards**.
- 3 From the drop-down menu, click **Monitoring - Networking and Security**.
- 4 To view status of EPP on clusters, click the Security widget.
- 5 In the Security Overview page, click **Capacity** and view capacity status of these parameters.

SECURITY OVERVIEW					
Configuration		Capacity			
REFRESH		Last updated on: Jun 26, 2019 at 4:30:00 PM			
		RESET VALUES		REVERT	SAVE
Limit	Maximum Capacity	Current Inventory (realized)		Warning Alert	Critical Alert
System Wide Endpoint Protection Enabled Hosts	1,024	5	0.49%	70%	100%
System Wide Endpoint Protection Enabled Virtual Machines	10,000	5	0.05%	70%	100%

- System Wide Endpoint Protection Enabled Hosts:** If the number of host numbers protected reaches the threshold limit, NSX Manager notifies a warning alert or critical alert when corresponding threshold limits are reached.
- System Wide Endpoint Protection Enabled Virtual Machines:** If the number of virtual machine numbers protected reaches the threshold limit, NSX Manager notifies a warning alert or critical alert when corresponding threshold limits are reached.

Note You can set threshold limits for these parameters, view status and receive alerts when these parameters reach the set threshold limit.

Change the Third Party Service Virtual Machine

An NSX-T Data Center administrator can change or deploy a new form factor or version of the service VM (SVM).

This task can be done both from UI or API.

The API command to change or upgrade a SVM is /


```
POST https://<NSX_Manager_IPaddress>/api/v1/serviceinsertion/services/{{service_id}}/
service-deployments /<service-deployment-id>?action=upgrade.
```

Prerequisites

- Ensure that partners have registered multiple service VMs differentiated by versions and or form factor (disk, vCPU, or RAM).
- Ensure that the SVM deployment status is `Deployment Successful` before changing the appliance. If the SVM is in a different state, go to **Home** → **Alarms**, and search for any open alarms of the event type EAM. Resolve them before trying to change to a newer SVM.
- Ensure that all prerequisites required to deploy an endpoint service or a combined partner service (For example, endpoint protection service and network gateway firewall) are met before proceeding to change the appliance.
- Ensure that storage is available before you change the existing SVM with a new one.
- If there are workloads that are protected by the existing SVM, first perform vMotion to migrate the workload and then change or deploy the new SVM.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Go to **System** → **Service Deployments** → **Deployment**.
- 3 Go to the service deployment and click **Change Appliance**.



- 4 In the **Change Appliance** window, select the specification of SVM you want to deploy and click **Update**.
- 5 If the service deployed is Upgrade Failed, click **Resolve** if the issue cannot be auto-resolved.
- 6 Click **Resolve All**.

The new SVM is deployed in NSX-T.

Results**Example:****What to do next**

To know the Runtime status of a combined partner service, Health Status of an endpoint protection service, or Deployment Status, go to the **Service Instances** tab.

After changing the existing SVM to a new SVM, enable guest VMs that are to be protected.

Manage Endpoint Protection

Resolve policy conflicts, health issues with service VMs, and know how endpoint protection policy works.

Resolve Partner Services Issues

Without partner service virtual machine functional, guest VMs are not protected against malware.

On each host, verify that the following services or process are up and running:

- ESXi Agency Manager (EAM) service must be up and running. The following URL must be accessible.

```
https://<vCenter_Server_IP_Address>/eam/mob
```

Verify the ESXi Agency Manager is online.

```
root> service-control --status vmware-eam
```

- Port groups of SVMs must not be deleted because these port groups are required to ensure that SVM continues to protect guest VMs.

```
https://<vCenter_Server_IP_Address>/ui
```

- In vCenter Server, go to the virtual machine, click the **Networks** tab, and check whether **vmervice-vshield-pg** is listed.
- Context Multiplexer (MUX) service is up and running. Check `nsx-context-mux` VIB is UP and running on the host.
- The management interface on which NSX-T Data Center communicates with the partner service console must be up.
- The control interface enabling communication between MUX and SVM must be up. Port group connecting MUX with SVM must be created. Both interface and port group are required for the partner service to be functional.

ESXi Agency Manager Issues

The table lists the ESXi Agency Manager issues that can be resolved using the Resolve button on the NSX Manager user interface. It notifies NSX Manager with error details.

Table 13-8. ESXi Agency Manager Issues

Issue	Category	Description	Resolution
Cannot Access Agent OVF	VM Not Deployed	An agent virtual machine is expected to be deployed on a host, but the agent virtual machine cannot be deployed because the ESXi Agent Manager is unable to access the OVF package for the agent. It might happen because the web server providing the OVF package is down. The web server is often internal to the solution that created the Agency.	ESXi Agency Manager (EAM) service retries the OVF download operation. Check the partner management console status. Click Resolve .

Table 13-8. ESXi Agency Manager Issues (continued)

Incompatible Host Version	VM Not Deployed	An agent virtual machine is expected to be deployed on a host. However, because of compatibility issues the agent was not deployed on the host.	Upgrade either the host or the solution to make the agent compatible with the host. Check the compatibility of the SVM. Click Resolve .
Insufficient Resources	VM Not Deployed	An agent virtual machine is expected to be deployed on a host. However, ESXi Agency Manager (EAM) service did not deploy the agent virtual machine because the host has less CPU or memory resources.	ESXi Agency Manager (EAM) service attempts to redeploy the virtual machine. Ensure that CPU and memory resources are available. Check the host and free up some resources. Click Resolve .
Insufficient Space	VM Not Deployed	An agent virtual machine is expected to be deployed on a host. However, the agent virtual machine was not deployed because the agent datastore on the host did not have enough free space.	ESXi Agency Manager (EAM) service attempts to redeploy the virtual machine. Free up some space on the datastore. Click Resolve .
No Agent VM Network	VM Not Deployed	An agent virtual machine is expected to be deployed on a host, but the agent cannot be deployed because the agent network has not been configured on the host.	Add one of the networks listed in customAgentVmNetwork to the host. The issue resolves automatically after the datastore is available.
Ovf Invalid Format	VM Not Deployed	An Agent virtual machine is expected to be provisioned on a host, but it failed to do so because the provisioning of the OVF package failed. The provisioning is unlikely to succeed until the solution that provides the OVF package has been upgraded or patched to provide a valid OVF package for the agent virtual machine.	ESXi Agency Manager (EAM) service attempts to redeploy the SVM. Check the partner solution documentation or upgrade the partner solution to get the valid OVF package. Click Resolve .
Missing Agent IP Pool	VM Powered Off	An agent virtual machine is expected to be powered on, but the agent virtual machine is powered off because there are no IP addresses defined on the agent's virtual machine network.	Define the IP address on the virtual machine network. Click Resolve .

Table 13-8. ESXi Agency Manager Issues (continued)

No Agent VM Datastore	VM Powered Off	An agent virtual machine is expected to be deployed on a host, but the agent cannot be deployed because the agent datastore has not been configured on the host.	Add one of the datastores listed in customAgentVmDatastore to the host. The issue resolves automatically after the datastore is available.
No Custom Agent VM Network	No Agent VM Network	An agent virtual machine is expected to be deployed on a host, but the agent cannot be deployed because the agent network has not been configured on the host.	Add the host to one of the networks listed in a custom agent VM network. The issue resolves automatically after a custom VM network is available.
No Custom Agent VM Datastore	No Agent VM Datastore	An agent virtual machine is expected to be deployed on a host, but the agent cannot be deployed because the agent datastore has not been configured on the host.	Add the host to one of the datastores listed in a custom agent VM datastore. The issue resolves automatically.
Orphaned Agency	Agency Issue	The solution that created the agency is no longer registered with the vCenter Server.	Register the solution with the vCenter Server.
Orphaned DvFilter Switch	Host Issue	A dvFilter switch exists on a host but no agents on the host depend on dvFilter. It happens if a host is disconnected when an agency configuration changed.	Click Resolve . ESXi Agency Manager (EAM) service attempts to connect the host before the agency configuration is updated.
Unknown Agent VM	Host Issue	An agent virtual machine has been found in the vCenter Server inventory that does not belong to any agency in this vSphere ESX Agent Manager server instance.	Click Resolve . ESXi Agency Manager (EAM) service attempts to place the virtual machine to the inventory it belongs to.
Ovf Invalid Property	VM Issue	An agent virtual machine must be powered on, but an OVF property is either missing or has an invalid value.	Click Resolve . ESXi Agency Manager (EAM) service attempts to reconfigure the correct OVF property.
VM Corrupted	VM Issue	An agent virtual machine is corrupt.	Click Resolve . ESXi Agency Manager (EAM) service attempts to repair the virtual machine.

Table 13-8. ESXi Agency Manager Issues (continued)

VM Orphaned	VM Issue	An agent virtual machine exists on a host, but the host is no longer part of scope for the agency. It happens if a host is disconnected when the agency configuration is changed.	Click Resolve . ESXi Agency Manager (EAM) service attempts to connect the host back to the agency configuration.
VM Deployed	VM Issue	An agent virtual machine is expected to be removed from a host, but the agent virtual machine has not been removed. The specific reason why vSphere ESX Agent Manager was unable to remove the agent virtual machine, such as the host is in maintenance mode, powered off or in standby mode.	Click Resolve . ESXi Agency Manager (EAM) service attempts to remove the agent virtual machine from the host.
VM Powered Off	VM Issue	An agent virtual machine is expected to be powered on, but the agent virtual machine is powered off.	Click Resolve . ESXi Agency Manager (EAM) service attempts to power on the virtual machine.
VM Powered On	VM Issue	An agent virtual machine is expected to be powered off, but the agent virtual machine is powered off.	Click Resolve . ESXi Agency Manager (EAM) service attempts to power off the virtual machine.
VM Suspended	VM Issue	An agent virtual machine is expected to be powered on, but the agent virtual machine is suspended.	Click Resolve . ESXi Agency Manager (EAM) service attempts to power on the virtual machine.
VM Wrong Folder	VM Issue	An agent virtual machine is expected to be located in a designated agent virtual machine folder, but is found in a different folder.	Click Resolve . ESXi Agency Manager (EAM) service attempts to place the agent virtual machine to the designated folder.

Table 13-8. ESXi Agency Manager Issues (continued)



VM Wrong Resource Pool	VM Issue	An agent virtual machine is expected to be located in a designated agent virtual machine resource pool, but is found in a different resource pool.	Click Resolve . ESXi Agency Manager (EAM) service attempts to place the agent virtual machine to a designated resource pool.
VM Not Deployed	Agent Issue	An agent virtual machine is expected to be deployed on a host, but the agent virtual machine has not been deployed. Specific reasons why ESXi Agent Manager was unable to deploy the agent, such as being unable to access the OVF package for the agent or a missing host configuration. This issue can also happen if the agent virtual machine is explicitly deleted from the host.	Click Resolve to deploy the agent virtual machine.

Next, configure the Endpoint Protection for VM groups. See [Endpoint Protection](#).

How Guest Introspection Runs Endpoint Protection Policy

Endpoint protection policies are enforced in a specific order. When you design policies, consider the sequence number associated to rules and the domains that host the rules.

Scenario: Out of the many workloads that run in your organization, for the purposes of illustration we consider two kinds of workloads - VMs running Virtual Desktop Infrastructure (VDI), and VMs running Payments Cards Industry Data Security Standards (PCI-DSS) workloads. A section of employees in the organization requires remote desktop access, which makes up the virtual desktop infrastructure (VDI) workload. These VDI workloads might require a Gold protection policy level based on the compliance rules set up by the organization. Whereas a PCI-DSS workload needs the highest level of protection, Platinum level protection.



As there are two workload types, create two policies one each for VDI workloads and server workloads. Within each policy or section, define a domain to reflect the workload type and within that section define rules for that workload. Publish the rules to start GI services on guest VMs. GI internally uses the two sequence numbers: Policy sequence number and rule sequence number to determine the complete sequence of rules to run. Each rule serves two purposes: determines which VMs to protect and the protection policy that must be applied to protect the VMs.

To change the sequence order, drag a rule in the NSX-T Data Center Policy Manager UI to change its sequence order. Alternatively, you can explicitly assign sequence number for rules using API.

Alternatively make an NSX-T Data Center API call to manually define a rule by associating a service profile with a VM group and declare the sequence number of the rules. The API and parameter details are detailed in the NSX-T Data Center *API guide*. Make Service configuration APIs calls to apply profiles to entities such as VM groups and so on.

Table 13-9. NSX-T Data Center APIs used to define rule that apply service profile to VM groups

API	Details
Get all service configuration details.	<pre>GET /api/v1/service-configs</pre> The service configuration API returns details of the service profile applied to a VM group, the VM group protected, and the sequence or precedence number that decides priority of the rule.
Create a service configuration.	<pre>POST /api/v1/service-configs</pre> The service configuration API takes input parameters of a service profile, VM group to be protected, and sequence or precedence number that must be applied to the rule.
Delete a service configuration.	<pre>DELETE /api/v1/service-configs/ <config-set-id></pre> The service configuration API deletes the configuration applied to the VM group.
Get details of a specific configuration.	<pre>GET /api/v1/service-configs/ <config-set-id></pre> Get details of a specific configuration

Table 13-9. NSX-T Data Center APIs used to define rule that apply service profile to VM groups (continued)

API	Details
Update a service configuration.	<pre>PUT /api/v1/service-configs/ <config-set-id></pre> Update a service configuration.
Get effective profiles.	<pre>GET /api/v1/service-configs/ effective-profiles?resource_id=<resource-id> &resource_type=<resource-type></pre> The service configuration API returns only that profile which is applied to a particular VM group.

Efficiently manage rules by following these recommendations:

- Set a higher sequence number for a policy for which rules must be ran first. From the UI, you can drag policies to change their priority.
- Similarly, set a higher sequence number for rules within each policy.
- Depending on how many rules you need, you can position rules apart in multiples of 2, 3, 4, or even 10. So, two consecutive rules that are 10 positions apart give you more flexibility to resequence rules without having to change the sequence order of all the rules. For example, if you do not plan to define many rules, you can select to position rules 10 positions apart. So, rule 1 gets a sequence number of 1, rule 2 gets a sequence number of 10, rule 3 gets a sequence number of 20, and so on. This recommendation provides flexibility to efficiently manage rules so that you do not need to resequence all the rules.

Internally, guest introspection sequences these policy rules in the following way.

```
Policy 1 ↔ Sequence Number 1 (1000)

- Rule 1 : Group 1↔ Service Profile ↔ Sequence Number 1 (1001)
- Rule 2 : Group 1↔ Service Profile ↔ Sequence Number 10 (1010)
- Rule 3 : Group 1↔ Service Profile ↔ Sequence Number 20 (1020)
- Rule 4 : Group 1↔ Service Profile ↔ Sequence Number 30 (1030)

Policy 2 ↔ Sequence Number 2 (2000)

- Rule 1 : Group 1↔ Service Profile ↔ Sequence Number 1 (2001)
- Rule 2 : Group 1↔ Service Profile ↔ Sequence Number 10 (2010)
- Rule 3 : Group 1↔ Service Profile ↔ Sequence Number 20 (2020)
```

```
- Rule 4 : Group 1↔ Service Profile ↔ Sequence Number 30 (2030)
```

Based on the above sequence numbers, G1 runs rules of Policy 1 before it runs rules of Policy 2.

But there are situations when the intended rules are not applied to a VM group or a VM. These conflicts need to be resolved to apply the desired policy protection levels.

Endpoint Policy Conflict Resolution

Consider a scenario where two policy domains exist, each consisting of multiple rules. As an admin you are not always certain of which VMs can end up getting membership of a group because VMs get associated to a group based on dynamic membership criteria, such as OS Name, Computer Name, User, Tagging.

Conflicts arise in the following scenarios:

- A VM is part of two groups, where each group is protected by a different profile.
- A partner service VM is associated with more than one service profile.
- An unexpected rule ran on a guest VM, or when a rule does not run on a VM group.
- Sequence number is not assigned to policy rules or domains.

Table 13-10. Resolve policy conflicts

Scenario	Expected Endpoint Protection Flow	Resolution
When a VM gets membership to multiple groups. And each group is protected by a different type of service profile. Expected protection was not applied to the VM.	A VM group created with a membership criteria means that VMs are added to the group dynamically. In such a case, the same VM can be part of multiple groups. There is no way to pre-determine which group that VM is going to be part of because the membership criteria dynamically populates VM into the group. Consider VM 1 is part of Group 1 and Group 2. ■ Rule 1: Group 1 (by OS name) is applied Gold (Service Profile) with Sequence Number 1 ■ Rule 2: Group 2 (by tag) is applied Platinum with Sequence Number 10 Endpoint protection policy runs the Gold service profile on VM 1 but does not run Platinum service profile on VM1.	Change the Sequence Number of Rule 2 such that it runs before Rule 1. ■ On the NSX-T Data Center Policy Manager UI, drag the Rule 2 before Rule 1 on the rule list. ■ Using NSX-T Data Center Policy Manager API, manually add a higher sequence number for Rule 2.
When a rule associates the same service profile to protect two VM groups. Endpoint protection does not run the rule on the second VM group.	Endpoint protection only runs the first service profile on the VM because the same service profile cannot be applied again to any other rule across policies or domain. Consider VM 1 is part of Group 1 and Group 2. Rule 1: Group 1 (by OS name) is applied Gold (service profile) Rule 2: Group 2 (by tag) is applied Gold (service profile)	■ Add Group 2 to Rule 1. (Rule 1: Group 1, Group 2 is applied Profile 1)

Quarantine VMs

After rules are applied to VM groups, based on the protection level and tag set by partners, there might be VMs that are identified as infected that need to be quarantined.

Partners use the API with tag `virus_found=true` to tag VMs that are infected. Affected VMs are attached with the `virus_found=true` tag.

As an administrator, you can create a pre-defined quarantine group based on tag with `virus_found=true` value, such that the group gets populated with infected VMs as and when they are tagged. As an admin, you might choose to set specific firewall rules for the quarantine group. You can set firewall rules for the quarantine group. For example, you might choose to block all traffic incoming and outgoing from the quarantine group.

Verify Health Status of Service Instances

Health status of a service instance depends on many factors: status of the partner solution, connectivity between Guest Introspection Agent (Context Multiplexer) and Context Engine (Ops

Agent), and availability of Guest Introspection Agent information, SVM protocol information with NSX Manager.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **System > Service Deployments > Service Instances**.
- 3 In the Health Status column, click  to know the health of the service instance.

Table 13-11. Health Status of Third-Party Service Instance

Parameter	Description
Health Status received at	The latest timestamp when NSX Manager received the health status details of the service instance.
Solution Status	Status of partner solution running on an SVM. Status UP indicates that the partner solution is correctly running.
Connectivity between NSX-T Data Center Guest Introspection Agent and NSX-T Data Center Ops Agent	Status is UP when NSX-T Data Center Guest Introspection agent (context multiplexer) is connected with the Ops agent (includes the context engine). The context multiplexer forwards health information of SVMs to the context engine. They also share SVM-VM configuration between each other to know which guest VMs are protected by the SVM.
Service VM Protocol Version	Transport protocol version used internally for troubleshooting issues.
NSX-T Data Center Guest Introspection Agent Information	Represents protocol version compatibility between NSX-T Data Center Guest Introspection agent and SVM.

- 4 If the Health Status is **Up** (status displayed in green) and the partner console displays all guest VMs as protected, the health status of the service instance is **Up**.
- 5 If the Health Status is **Up** (status displayed in green) but the partner console displays guest VMs in unprotected state, perform the following step:
 - a Contact VMware support to resolve the issue. The health status of the service instance might be down not correctly reflected by the NSX Manager user interface.

- 6 If the Health Status is **Down** (status displayed in red), then one or more factors that determine the service instance health are down.

Table 13-12. Troubleshoot Health Status

Health Status Attribute	Resolution
Solution Status is Down or Not available .	1 Verify that service deployment status is Up (green). If you encounter errors, see Resolve Partner Services Issues. 2 Ensure that at least one guest VM in the affected host is protected with an endpoint protection policy. 3 From the partner console, verify whether the solution service is running on the SVM on the host. See the Partner documentation for more details. 4 If none of the above steps resolve the issue, contact VMware support.
Connectivity between NSX-T Data Center Guest Introspection Agent and NSX-T Data Center Ops Agent is Down .	1 Verify that service deployment status is Up (green). If you encounter errors, see Resolve Partner Services Issues. 2 Ensure that at least one guest VM in the affected host is protected with an endpoint protection policy. 3 From the partner console, verify whether the solution service is running on the SVM on the host. See the Partner documentation for more details. 4 If none of the above steps resolve the issue, contact VMware support.
Service VM Protocol Version is Unavailable .	1 Verify that service deployment status is Up (green). If you encounter errors, see Resolve Partner Services Issues. 2 Ensure that at least one guest VM in the affected host is protected with an endpoint protection policy. 3 From the partner console, verify whether the solution service is running on the SVM on the host. See the Partner documentation for more details. 4 If none of the above steps resolve the issue, contact VMware support.
NSX-T Data Center Guest Introspection Agent Information is Unavailable .	Contact VMware support.

Delete Partner Services

To delete partner services, make an API call . Before you make the API call to delete partner services or SVMs deployed on a host, you need to do the following from the NSX Manager user interface.

To delete partner services:

Procedure

- 1 Remove EPP rules applied to VM groups running on the host.

- 2 Remove service profile protection applied to VM groups.
- 3 To remove solution binding SVMs with partner service manager, make the following API call.

```
/DEL https://<NSX_Manager_IPAddress>/api/v1/serviceinsertion/services/{{service_id}}/
solution-configs/<solution-config-id>
```

- 4 To delete the service deployment, make the following API call.

```
/DEL https://<NSX_Manager_IPAddress>/api/v1/serviceinsertion/services/<service-id>/service-
deployments/<service-deployment-id>
```

Refer to the *NSX-T Data Center API guide* for more information on API parameters.

Security Profiles

This section contains profiles that fine tune Firewall Operations: Session Timers, Flood Protection, and DNS Security

Create a Session Timer

Session Timers define how long a session is maintained on the firewall after inactivity in the session.

When the session timeout for the protocol expires, the session closes. On the firewall, several timeouts for TCP, UDP, and ICMP sessions can be specified to apply to a user-defined group or a Tier-0 or Tier-1 gateway. Default session values can be modified depending on your network needs. Note that setting a value too low might cause frequent timeouts, and setting a value too high might delay failure detection. See [Default Session Timer Values](#) for more information.

Session timers are supported on ESXi and KVM hosts.

Procedure

- 1 Navigate to **Security > Settings > Security Profiles > Session Timer**.

- 2 Click **Add Profile**.

The **Profile** screen appears, populated with the default values.

- 3 Enter a **name** and a **description** (optional) for the timer profile.
- 4 Click **Set** to select the Tier-0 or Tier-1 gateway or group to apply the timer profile.
- 5 Select the protocol. Accept the default values or enter your own values.

TCP Variables	Description
First Packet	The timeout value for the connection after the first packet has been sent. The default is 120 seconds.
Opening	The timeout value for the connection after a second packet has been transferred. The default is 30 seconds.

TCP Variables	Description
Established	The timeout value for the connection once the connection has become fully established.
Closing	The timeout value for the connection after the first FIN has been sent. The default is 120 seconds.
Fin Wait	The timeout value for the connection after both FINs have been exchanged and the connection is closed. The default is 45 seconds.
Closed	The timeout value for the connection after one endpoint sends an RST. The default is 20 seconds.
UDP Variables	Description
First Packet	The timeout value for the connection after the first packet is sent. This is the initial timeout for the new UDP flow. The default is 60 seconds.
Single	The timeout value for the connection if the source host sends more than one packet and the destination host has not sent one back. The default is 30 seconds. ESXi hosts only. KVM hosts use the UDP first packet.
Multiple	The timeout value for the connection if both hosts have sent packets. The default is 60 seconds.
ICMP Variables	Description
First Packet	The timeout value for the connection after the first packet is sent. This is the initial timeout for the new ICMP flow. The default is 20 seconds.
Error reply	The timeout value for the connection after an ICMP error is returned in response to an ICMP packet. The default is 10 seconds. ESXi hosts only. KVM hosts use the ICMP first packet.

6 Click **Save**.

What to do next

After saving, click [Manage Group to Profile Precedence](#) to manage group to profile binding precedence.

Default Session Timer Values

The session timer profile applies the timeout values to Tier-0 or Tier-1 router interfaces or groups containing segments, segment-ports, tags, or any other non-IP based groups. The timeout values decide how long a protocol session remains active after the session closes.

Session Timer Values

- Default Timer Profile shown with API and UI applies only to distributed firewall (DFW).
- Gateway Firewall (GFW) default session timers are different than the default timer profile seen when using API and UI. GFW default session timers are optimized for North-South traffic, and some of them are lower than minimum configurable values by default.
- Firewall session timers can be changed for both DFW and GFW by using the API and UI.
- The same non-default timer profile can be applied to both DFW and GFW, if needed.

If you do not customize timer values, the gateway takes default values. Gateway firewall default timer values:

Timer Property	Edge Default (secs)	Minimum (secs)	Maximum (secs)
ICMP Error Reply	6	10	4320000
ICMP First Packet	6	10	4320000
TCP Closed	2	10	4320000
TCP Closing	900	10	4320000
TCP Established	7200	120	4320000
TCP Fin-wait	4	10	4320000
TCP First Packet	120	10	4320000
TCP Opening	30	10	4320000
UDP First Packet	30	10	4320000
UDP Multiple	30	10	4320000
UDP Single	30	10	4320000

Distributed firewall default session timer values:

Timer Property	DFW Default (secs)	Minimum (secs)	Maximum (secs)
ICMP Error Reply	10	10	4320000
ICMP First Packet	20	10	4320000
TCP Closed	20	10	4320000
TCP Closing	120	10	4320000
TCP Established	43200	120	4320000
TCP Fin-wait	45	10	4320000
TCP First Packet	120	10	4320000
TCP Opening	30	10	4320000
UDP First Packet	60	10	4320000
UDP Multiple	60	10	4320000
UDP Single	30	10	4320000

Flood Protection

Flood protection helps to protect against Denial of Service (DDoS) attacks.

DDoS attacks aim to make a server unavailable to legitimate traffic by consuming all the available server resources - the server is flooded with requests. Creating a flood protection profile imposes active session limits for ICMP, UDP, and half-open TCP flows. Distributed firewall can cache flow entries which are in SYN_SENT and SYN_RECEIVED states, and promote each entry to a TCP state after an ACK is received from the initiator, completing the three-way handshake.

Procedure

- 1 Navigate to **Security > Security Profiles > Flood Protection**.
- 2 Click **Add Profile**, and select **Add Edge Gateway Profile** or **Add Firewall Profile**.
- 3 Fill out the flood protection profile parameters:

Table 13-13. Parameters for Firewall and Edge Gateway Profiles

Parameter	Minimum and maximum values	Default	
TCP Half Open Connection Limit - TCP SYN flood attacks are prevented by limiting the number of active, not-fully-established TCP flows which are allowed by the firewall.	1-1,000,000	Firewall - None Edge Gateway - 1,000,000	Set this text box to limit the number of active TCP half open connections. If this text box is empty, this limit is disabled on ESX nodes and set to the default on value of Edge Gateways.
UDP Active Flow Limit -UDP flood attacks are prevented by limiting the number of active UDP flows which are allowed by the firewall. Once the set UDP flow limit is reached, subsequent UDP packets which can establish a new flow are dropped.	1-1,000,000	Firewall - None Edge Gateway - 1,000,000	Set this text box to limit the number of active UDP connections. If this text box is empty, this limit is disabled on ESX nodes and set to the default on value of Edge Gateways.
ICMP Active Flow Limit - ICMP flood attacks are prevented by limiting the number of active ICMP flows which are allowed by the firewall. After the set flow limit is reached, subsequent ICMP packets which can establish a new flow are dropped.	1-1,000,000	Firewall - None Edge Gateway - 10,000	Set this text box to limit the number of active ICMP open connections. If this text box is empty, this limit is disabled on ESX nodes and set to the default on value of Edge Gateways.

Table 13-13. Parameters for Firewall and Edge Gateway Profiles (continued)

Parameter	Minimum and maximum values	Default	
Other Active Connection Limit	1-1,000,000	Firewall - None Edge Gateway - 10,000	Set this text box to limit the number of active connections other than ICMP, TCP, and UDP half open connections. If this text box is empty, this limit is disabled on ESX nodes, and set to the default on value of Edge Gateways.
SYN Cache - Syn cache is used when a TCP half open connection limit has also been configured. The number of active half-open connections are enforced by maintaining a syncache of the not-fully-established TCP sessions. This cache maintains the flow entries which are in SYN_SENT and SYN_RECEIVED states. Each syncache entry will be promoted to a full TCP state entry after an ACK is received from the initiator, completing the three-way handshake.		Only available for firewall profiles.	Toggle on and off. Enabling SYN cache is effective only when a TCP half open connection limit is configured.
RST Spoofing - Generates spoofed RST to server when purging half-open states from SYN cache. Allows server to clean up states associated with SYN flood (half open).		Only available for firewall profiles.	Toggle on and off. SYN Cache must be selected for this option to be available

4 To apply the profile to edge gateways and firewall groups, click **Set**.

5 Click **Save**.

What to do next

After saving, click [Manage Group to Profile Precedence](#) to manage group to profile binding precedence.

Configure DNS Security

Creating a DNS Security Profile helps to guard against DNS-related attacks.

You can do the following after you set up the DNS Security Profile:

- Snoop on DNS responses for a VM, or a group of VMs on the transport node to associate FQDN with IP addresses.
- Add global and default DNS server information and apply it to all VMs that are using DFW rules.
- Specify selected DNS server information for selected VMs.
- Apply DNS profiles to groups.

Note Only ESXi is supported in the current release.

Procedure

- 1 Navigate to **Security > Settings > Security Profiles > DNS Security** .
- 2 Click **Add Profile**.
- 3 Enter the following values:

Option	Description
Profile Name	Provide a profile name.
TTL	This field captures the Time to live for the DNS cache entry in seconds. You have the following options: TTL 0 - cached entry never expires. TTL 1 to 3599 - invalid TTL 3600 to 864000 – valid TTL left empty – automatic TTL, set from the DNS response packet. Note DNS Security Profile has a default DNS cache timeout of 24 hours.
Applied To	You can select a group based on any criteria to apply the DNS security profile to. Note Only one DNS server profile is applied to a VM.
Tags	Optional. Assign a tag and scope to the DNS profile to make it easy to search. See Add Tags to an Object for more information.

- 4 Click **Save**.

What to do next

After saving, click [Manage Group to Profile Precedence](#) to manage group to profile binding precedence.

Manage Group to Profile Precedence

You can bind multiple groups to a security profile. NSX-T Data Center applies the security profile to the group with highest precedence level.

If you bind a security profile to multiple groups, NSX-T Data Center assigns highest precedence to the newest group from that list. However, you can change the precedence level for groups.

To assign precedence to groups:

Prerequisites

- Session timer groups must only contain segments, segment ports and VMs as members. Other category types are not supported.
- DNS security groups must contain only VMs as members. Other category types are not supported.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Navigate to **Security > Security Profiles**.
- 3 Click **Manage Group to Profile Precedence**.
- 4 To assign a group highest level of precedence, move it to the top of the list.
- 5 Click **Close**.

Results

The security profile is applied to the group with highest precedence level.

Time-Based Firewall Policy

With time windows, security administrators can restrict traffic from a source or to a destination, for a specific time period.

Time windows apply to a firewall policy section and all the rules in it. Each firewall policy section can have one time window. The same time window can be applied to more than one policy section. If you want the same rule applied on different days or different times for different sites, you must create more than one policy section. Time-based rules are available for distributed and gateway firewalls.

Prerequisites

Network Time Protocol (NTP) service must be running on each transport node when using time-based rule publishing. See [Configuring Appliances](#).

If a time-zone is changed on the edge transport node after the node is deployed, reload the edge node or restart the data plane for time-based gateway firewall policy to take effect.

Create a firewall policy.

Procedure

- 1 Click the clock icon on the firewall policy you want to have a time window.
A time window appears.
- 2 Click **Add New Time Window** and enter a **name**.
- 3 Select a time zone: UTC (Coordinated Universal Time), or the local time of the transport node. Distributed firewall only supports UTC with NTP service enabled, a change of time zone configuration is not supported on ESXi hosts.
- 4 Select the frequency of the time window - **Weekly** or **One time**.
- 5 Select the days of the week that the time window takes effect.

NSX-T Data Center supports configuring weekly UTC time-windows for the local time-zone, when the entire time-window for the local time-zone is within the same day as the UTC time-zone. For example, you cannot configure a time window in UTC for a 7am-7pm PDT, which maps to UTC 2pm-2am of the next day.
- 6 Select the beginning and ending dates for the time window, and the times the window will be in effect.
- 7 Click **Save**.
- 8 Click the check box next to the policy section you want to have a time window. Then click the clock icon.
- 9 Select the time window you want to apply, and click **Apply**.
- 10 Click **Publish**. The clock icon for the section turns green.

For the first publication of a time-based rule, the time is taken, and rule enforcement begins at less than 2 minutes. After the rules are deployed, enforcement as per time window, is instantaneous.

Network Introspection Settings

This section contains settings to configure network introspection.

Add a Service Segment

When you configure east-west network introspection or you want to redirect packets from the uplink of an NSX Edge to the service chain, create a service segment.

Prerequisites

- If you are configuring east-west service chaining to redirect packets from the uplink of an NSX Edge to the service chain, create a Tier-0 and or Tier-1 gateway. The segment is later connected to the Tier-0 and or Tier-1 gateway.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Click **Security > Settings > Network Introspection Settings > Service Segment > Add Service Segment**.
- 3 Click **Add Service Segment**.
- 4 In the **Name** field, enter a name for the segment.
- 5 In the **Transport Zone (Overlay)** field, select an overlay transport zone that is associated to the segment.
- 6 In the **Connected To** field, do one of the following:
 - Leave the field blank if you are configuring east-west network introspection to protect guest VMs by third-party security vendors.
 - Select a Tier-0 or Tier-1 gateway if you are configuring an east-west service chaining to redirect packets from the uplink of an NSX Edge to the service chain.
- 7 Click **Save**.

Results

The **Status** column displays the status of the service segment.

Add a Service Profile

A service profile is an instance of a partner vendor template. Administrators can customize attributes of a vendor template to create an instance of the template.

Note You can create multiple service profile for a single vendor. For example, the service profile set for the forward path provides IDS protection, whereas the service profile set for the reverse path supports IPS protection. However, a single service profile can be set for both forward and reverse path.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Navigate to **Security > Network Introspection Settings > Service Profiles**.
- 3 From the Partner Service drop-down field, select a service. You can create a service profile for the selected service.
- 4 Enter the service profile name and select the vendor template.
- 5 The Redirection Action field inherits functionality from the vendor template. For example, if COPY is the functionality provided by the vendor template, then by default the redirection action when you create a service profile is COPY.

- 6 (Optional) Define any tags to filter out and manage service profiles.
- 7 Click **Save**.

Results

A new service profile is created for the partner service.

What to do next

Add a service chain. See [Add a Service Chain](#).

Add a Service Chain

A service chain is a logical sequence of service profiles defined by the network administrator.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Verify the NSX Manager is in **Policy** mode.
- 3 Select **Security > Network Introspection Settings > Service Chain > Add Chain**.
- 4 Enter the service chain name.
- 5 In the Service Segments field, select the service segment to which you want to apply the service chain.

A service segment is a segment of service plane that connects multiple service VMs of an overlay transport zone. Each service VM in the service chain is separate from another service VM and L2 and L3 network segments run by NSX-T Data Center. The service plane controls access to service VMs.
- 6 To set the forward path, click the **Set Forward Path** field and click **Add Profile in Sequence**.
- 7 Add the first profile in the service chain and click **Add**.
- 8 To specify the next service profile, click **Add Profile in Sequence** and enter details.

You can also rearrange the profile order by using the Up and Down arrow icons.
- 9 Click **Save** to finish adding a forward path for the service chain.
- 10 In the Reverse Path column, select **Inverse Forward Path** for the service plane to use the service profile you set for the forward path.
- 11 To set a new service profile for the reverse path, click **Set Reverse Path** and add a service profile.
- 12 Click **Save** to finish adding a reverse path for the service chain.

13 In the Failure Policy field,

- Select **Allow** to send traffic to the destination VM when the service VM fails. Service VM failure is detected by the liveness detection mechanism which can be enabled only by partners.
- Select **Block** to not send traffic to the destination VM when the service VM fails.

14 Click **Save**.

Results

After adding a service chain, the partner Service Manager is notified about the update.

What to do next

Create a redirection rule to introspect east-west network traffic. See [Add Redirection Rules for East-West Traffic](#).

Troubleshooting Firewall

This section provides information about troubleshooting firewall issues.

Monitor and Troubleshoot Firewall on NSX Manager

There are several steps to take when troubleshooting firewall.

- 1 Check the Firewall policy realization status. See [Check Rule Realization Status](#).
- 2 Check the rule hits statistics by navigating to **Security > Distributed Firewall** or **Security > Gateway Firewall**, and clicking the graph icon. Rule level statistics are aggregated every 15 minutes from all the transport nodes. Rule statistics can be reset using **Reset All Rules Stats** from the three dot menu icon .
- 3 Check for Capacity Dashboard to make sure configuration is within the supported limit of NSX-T Data Center. The Capacity dashboard can be accessed from **Security > Security Overview > Capacity** , see [View the Usage and Capacity of Categories of Objects](#).
- 4 Check for supported configuration max limit for the given release by checking the [Configuration Limits](#).
- 5 Check for per VM level Firewall Rules pushed to datapath in Manager Mode by navigating **Logical Switches > Ports > Related Firewall Rules**.

You can also use the following NSX DFW helper script from github to get the total firewall rules configured and per VM firewall rules. <https://github.com/vmware-samples/nsx-t/blob/master/helper-scripts/DFW/nsx-get-dfw-rules-per-vm.py>

Troubleshooting Distributed Firewall on ESX Hosts

On ESX hosts, follow these steps to troubleshoot the NSX distributed firewall (DFW) data path issues.

Get the list of VMs on the ESXi host and associated Filter Name

This lists all VM's on this ESXi host. Note down the value of "name" field and use that in subsequent commands to get relevant output for a given VM.

```
[root@esxcomp-2a:~] summarize-dvfilter | grep -A 3 vmm
world 1371516 vmm0:PROD-MRS-DB-01 vcUuid:'50 20 92 e1 11 b7 10 d3-56 c5 e0 da 46 87 b5 d2'
port 67108881 PROD-MRS-DB-01.eth0
vNic slot 2
name: nic-1371516-eth0-vmware-sfw.2
--
world 1622816 vmm0:DEV-MRS-DB-01 vcUuid:'50 2d f3 a3 96 a4 f4 94-6e 55 84 85 c1 bd 05 2c'
port 67108883 DEV-MRS-DB-01.eth0
vNic slot 2
name: nic-1622816-eth0-vmware-sfw.2
--
world 7014985 vmm0:PROD-MRS-APP-01 vcUuid:'50 20 9b 5f cd b7 43 de-ab bb 8d 0e f5 bb ca 99'
port 67108895 PROD-MRS-APP-01.eth0
vNic slot 2
name: nic-7014985-eth0-vmware-sfw.2
--
world 7022287 vmm0:PROD-MRS-APP-02 vcUuid:'50 20 4a 44 17 fb 21 cf-fb 62 1e a3 d0 3c 7d cf'
port 67108896 PROD-MRS-APP-02.eth0
vNic slot 2
name: nic-7022287-eth0-vmware-sfw.2
[root@esxcomp-2a:~]
```

Get the firewall rules applied to a VM

Use Filter name associated with the VM from above output to get all the firewall rules applied to that VM's vNIC

```
[root@esxcomp-2a:~] vsipioctl getrules -f nic-7014985-eth0-vmware-sfw.2
ruleset mainrs {
  # generation number: 0
  # realization time : 2020-12-16T23:41:30
  # PRE_FILTER rules
  rule 5134 at 1 inout protocol any from addrset d8e7adac-af3b-4f22-9785-0cc30f0e81b1 to
addrset d8e7adac-af3b-4f22-9785-0cc30f0e81b1 accept with log tag 'ipv6-app-allow';
  rule 5133 at 2 inout protocol any from any to any accept with log tag 'ipv6-app-deny-
default';
  rule 5132 at 3 inout inet protocol icmp from any to addrset
9b14a216-4318-4bb1-94b0-56dfedec6f24 accept with log tag 'icmp-test';
  rule 5132 at 4 inout inet protocol tcp strict from any to addrset
9b14a216-4318-4bb1-94b0-56dfedec6f24 port 22 accept with log tag 'icmp-test';
  rule 5132 at 5 inout inet protocol ipv6-icmp from any to addrset
9b14a216-4318-4bb1-94b0-56dfedec6f24 accept with log tag 'icmp-test';
  rule 5130 at 6 inout inet protocol icmp from any to addrset rdst5130 accept with log tag
'icmp-test-gb-default';
  rule 5130 at 7 inout inet protocol ipv6-icmp from any to addrset rdst5130 accept with log
tag 'icmp-test-gb-default';
  # FILTER (APP Category) rules
  rule 5102 at 1 inout protocol any from addrset rsrc5102 to addrset d19f38e1-c13e-4fbb-9d6b-
b6971f251e2d accept;
```

```

    rule 5126 at 2 in protocol any from addrset rsrc5127 to addrset d19f38e1-c13e-4fbb-9d6b-
b6971f251e2d accept;
    rule 5127 at 3 out protocol any from addrset rsrc5127 to addrset d19f38e1-c13e-4fbb-9d6b-
b6971f251e2d accept;
    rule 5128 at 4 out protocol any from addrset rsrc5128 to addrset rdst5128 accept;
    rule 5129 at 5 in protocol any from addrset rsrc5128 to addrset
98abd76f-351b-4a4a-857f-1d91416b0798 accept;
    rule 5103 at 6 in protocol any from addrset rsrc5128 to addrset bled4d3d-ab4c-4bab-999b-
a50642cad495 accept;
    rule 5135 at 7 inout protocol any from any to any with attribute profile
acf76e7d-400b-438b-966f-8d5c10bebbda accept;
    rule 5135 at 8 inout protocol any from any to any with attribute profile 88dc6bf0-808e-49f6-
a692-dd0e5cee6ab3 accept;
    rule 5124 at 9 inout protocol any from any to any with attribute profile 8774c654-0f9e-43ad-
a803-4aa720e590cf accept;
    rule 5123 at 10 inout protocol any from any to any with attribute profile 13e599b5-
dd2d-420f-8473-9d45f0d324ac accept;
    rule 5125 at 11 inout protocol any from any to any with attribute profile e4be8d7e-
e4ab-4466-8f2e-998445ead95d accept;
    rule 2 at 12 inout protocol any from any to any drop with log tag 'icmp-default-rule';
}

ruleset mainrs_L2 {
    # generation number: 0
    # realization time : 2020-12-16T23:41:30
    # FILTER rules
    rule 1 at 1 inout ethertype any stateless from any to any accept;
}

[root@esxcomp-2a:~]

```

Get stats per FW rule per VM VNIC

Use "-s" with the above command to get the firewall stats associated with the VM firewall rules.

```

[root@esxcomp-2a:~] vsipioctl getrules -f nic-7014985-eth0-vmware-sfw.2 -s
ruleset mainrs {
    # PRE_FILTER rules
    rule 5134 at 1, 68 evals, 68 hits, 68 sessions, in 1120 out 1120 pkts, in 113952 out 114184
bytes
    rule 5133 at 2, 24 evals, 24 hits, 24 sessions, in 16 out 8 pkts, in 896 out 768 bytes
    rule 5132 at 3, 0 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
    rule 5132 at 4, 0 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
    rule 5132 at 5, 0 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
    rule 5130 at 6, 0 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
    rule 5130 at 7, 0 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
    # FILTER (APP Category) rules
    rule 5102 at 1, 0 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
    rule 5126 at 2, 0 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
    rule 5127 at 3, 0 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
    rule 5128 at 4, 0 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
    rule 5129 at 5, 0 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
    rule 5103 at 6, 0 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
    rule 5135 at 7, 92 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes

```

```

rule 5135 at 8, 92 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
rule 5124 at 9, 92 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
rule 5123 at 10, 92 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
rule 5125 at 11, 92 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
rule      2 at 12, 92 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
}

ruleset mainrs_L2 {
    # FILTER rules
rule      1 at 1, 0 evals, 0 hits, 0 sessions, in 0 out 0 pkts, in 0 out 0 bytes
}

[root@esxcomp-2a:~]

```

Get the addrset/groups used in the VM's Firewall rules

The firewall rule uses groups/addrset in the Source or destination. This output gets the all the addrset used in the rules based on the grouping configuration.

```

[root@esxcomp-2a:~] vsipioctl getaddrset -f nic-1371516-eth0-vmware-sfw.2
addrset is shared for this filter
global addrset
addrset 98abd76f-351b-4a4a-857f-1d91416b0798 {
ip 7.7.7.7,
ip 8.8.8.8,
}
addrset 9b14a216-4318-4bb1-94b0-56dfedec6f24 {
ip 10.1.0.0,
ip 10.2.0.2,
ip 10.114.217.26,
ip 172.16.202.2,
ip 172.16.202.22,
ip 192.168.202.2,
ip 2001::172:16:202:2,
ip 2001::172:16:202:22,
mac 00:50:56:a0:0e:25,
mac 00:50:56:a0:26:dc,
mac 00:50:56:a0:2d:c0,
mac 00:50:56:a0:8d:90,
}
addrset b1ed4d3d-ab4c-4bab-999b-a50642cad495 {
ip 7.7.7.7,
ip 8.8.8.8,
}
addrset d19f38e1-c13e-4fbb-9d6b-b6971f251e2d {
ip 3.3.3.3,
ip 4.4.4.4,
}
addrset d8e7adac-af3b-4f22-9785-0cc30f0e81b1 {
ip 172.16.202.2,
ip 172.16.202.22,
ip 2001::172:16:202:2,
ip 2001::172:16:202:22,
mac 00:50:56:a0:26:dc,
}

```



```

mac 00:50:56:a0:8d:90,
}
addrset rdst5128 {
ip 3.3.3.3,
ip 4.4.4.4,
ip 7.7.7.7,
ip 8.8.8.8,
}
addrset rdst5130 {
ip 1.1.1.1,
ip 1.1.1.2,
ip 100.100.100.100,
}
addrset rsrc5102 {
ip 1.1.1.1,
ip 1.1.1.2,
}
addrset rsrc5127 {
ip 1.1.1.1,
ip 1.1.1.2,
ip 3.3.3.3,
ip 4.4.4.4,
}
addrset rsrc5128 {
ip 1.1.1.1,
ip 1.1.1.2,
ip 3.3.3.3,
ip 4.4.4.4,
ip 7.7.7.7,
ip 8.8.8.8,
}
local addrset
No address sets.
[root@esxcomp-2a:~]

```

Get the active Firewall flow per VM

NSX DFW maintains active flow per VNIC. This output gets the all the active flows over that VNIC.

```

[root@esxcomp-2a:~] vsipioctl getflows -f nic-7014985-eth0-vmware-sfw.2
Count retrieved from kernel active=6, inactive=0, drop=0
ecbd448200000001 Active ipv6-icmp 86dd IN 5134 0 0 2001::172:16:202:22 ->
2001::172:16:202:2 128 0 1039376 1039376 9994 9994 tmo 9
ecbd4482000000b9 Active tcp 0800 OUT 5134* 0 0 (est) 172.16.202.2:Unknown(39914) ->
172.16.202.22:ssh(22) 305 EST:EST rtt 21020 retrans 0/0 4409 3725 23 25 tmo 43195
ecbd4482000000ba Active ipv6-icmp 86dd OUT 5134* 0 0 fe80::250:56ff:fea0:8d90 ->
2001::172:16:202:22 135 0 64 72 1 1
ecbd4482000000bb Active igmp 0800 IN 5133* 0 0 (D) 0.0.0.0 -> 224.0.0.1 36 0 1 0 tmo 51
ecbd4482000000bc Active ipv6-icmp 86dd IN 5133* 0 0 (D) fe80::ffff:ffff:ffff:ffff ->
ff02::1 130 0 76 0 1 0 tmo 11
ecbd4482000000bd Active ipv6-icmp 86dd OUT 5133* 0 0 (D) fe80::250:56ff:fea0:8d90 ->
ff02::16 143 0 0 96 0 1 tmo 11
[root@esxcomp-2a:~]

```

Get the active Full Firewall config per VM

This output provides full firewall config per VNIC- Rules, Addrset & Profiles used.

```
[root@esxcomp-2a:~] vsipioctl getfwconfig -f nic-7014985-eth0-vmware-sfw.2
ruleset mainrs {
  # generation number: 0
  # realization time : 2020-12-16T23:41:30
  # PRE_FILTER rules
  rule 5134 at 1 inout protocol any from addrset d8e7adac-af3b-4f22-9785-0cc30f0e81b1 to
addrset d8e7adac-af3b-4f22-9785-0cc30f0e81b1 accept with log tag 'ipv6-app-allow';
  rule 5133 at 2 inout protocol any from any to any accept with log tag 'ipv6-app-deny-
default';
  rule 5132 at 3 inout inet protocol icmp from any to addrset
9b14a216-4318-4bb1-94b0-56dfedec6f24 accept with log tag 'icmp-test';
  rule 5132 at 4 inout inet protocol tcp strict from any to addrset
9b14a216-4318-4bb1-94b0-56dfedec6f24 port 22 accept with log tag 'icmp-test';
  rule 5132 at 5 inout inet protocol ipv6-icmp from any to addrset
9b14a216-4318-4bb1-94b0-56dfedec6f24 accept with log tag 'icmp-test';
  rule 5130 at 6 inout inet protocol icmp from any to addrset rdst5130 accept with log tag
'icmp-test-gb-default';
  rule 5130 at 7 inout inet protocol ipv6-icmp from any to addrset rdst5130 accept with log
tag 'icmp-test-gb-default';
  # FILTER (APP Category) rules
  rule 5102 at 1 inout protocol any from addrset rsrc5102 to addrset d19f38e1-c13e-4fbb-9d6b-
b6971f251e2d accept;
  rule 5126 at 2 in protocol any from addrset rsrc5127 to addrset d19f38e1-c13e-4fbb-9d6b-
b6971f251e2d accept;
  rule 5127 at 3 out protocol any from addrset rsrc5127 to addrset d19f38e1-c13e-4fbb-9d6b-
b6971f251e2d accept;
  rule 5128 at 4 out protocol any from addrset rsrc5128 to addrset rdst5128 accept;
  rule 5129 at 5 in protocol any from addrset rsrc5128 to addrset
98abd76f-351b-4a4a-857f-1d91416b0798 accept;
  rule 5103 at 6 in protocol any from addrset rsrc5128 to addrset bled4d3d-ab4c-4bab-999b-
a50642cad495 accept;
  rule 5135 at 7 inout protocol any from any to any with attribute profile
acf76e7d-400b-438b-966f-8d5c10bebbda accept;
  rule 5135 at 8 inout protocol any from any to any with attribute profile 88dc6bf0-808e-49f6-
a692-dd0e5cee6ab3 accept;
  rule 5124 at 9 inout protocol any from any to any with attribute profile 8774c654-0f9e-43ad-
a803-4aa720e590cf accept;
  rule 5123 at 10 inout protocol any from any to any with attribute profile 13e599b5-
dd2d-420f-8473-9d45f0d324ac accept;
  rule 5125 at 11 inout protocol any from any to any with attribute profile e4be8d7e-
e4ab-4466-8f2e-998445ead95d accept;
  rule 2 at 12 inout protocol any from any to any drop with log tag 'icmp-default-rule';
}

ruleset mainrs_L2 {
  # generation number: 0
  # realization time : 2020-12-16T23:41:30
  # FILTER rules
  rule 1 at 1 inout ethertype any stateless from any to any accept;
}
```

```

addrset is shared for this filter
global addrset
addrset 98abd76f-351b-4a4a-857f-1d91416b0798 {
ip 7.7.7.7,
ip 8.8.8.8,
}
addrset 9b14a216-4318-4bb1-94b0-56dfedec6f24 {
ip 10.1.0.0,
ip 10.2.0.2,
ip 10.114.217.26,
ip 172.16.202.2,
ip 172.16.202.22,
ip 192.168.202.2,
ip 2001::172:16:202:2,
ip 2001::172:16:202:22,
ip fe80::250:56ff:fea0:26dc,
ip fe80::250:56ff:fea0:8d90,
mac 00:50:56:a0:0e:25,
mac 00:50:56:a0:26:dc,
mac 00:50:56:a0:2d:c0,
mac 00:50:56:a0:8d:90,
}
addrset b1ed4d3d-ab4c-4bab-999b-a50642cad495 {
ip 7.7.7.7,
ip 8.8.8.8,
}
addrset d19f38e1-c13e-4fbb-9d6b-b6971f251e2d {
ip 3.3.3.3,
ip 4.4.4.4,
}
addrset d8e7adac-af3b-4f22-9785-0cc30f0e81b1 {
ip 172.16.202.2,
ip 172.16.202.22,
ip 2001::172:16:202:2,
ip 2001::172:16:202:22,
ip fe80::250:56ff:fea0:26dc,
ip fe80::250:56ff:fea0:8d90,
mac 00:50:56:a0:26:dc,
mac 00:50:56:a0:8d:90,
}
addrset rdst5128 {
ip 3.3.3.3,
ip 4.4.4.4,
ip 7.7.7.7,
ip 8.8.8.8,
}
addrset rdst5130 {
ip 1.1.1.1,
ip 1.1.1.2,
ip 100.100.100.100,
}
addrset rsrc5102 {
ip 1.1.1.1,
ip 1.1.1.2,
}

```

```

addrset rsrc5127 {
ip 1.1.1.1,
ip 1.1.1.2,
ip 3.3.3.3,
ip 4.4.4.4,
}
addrset rsrc5128 {
ip 1.1.1.1,
ip 1.1.1.2,
ip 3.3.3.3,
ip 4.4.4.4,
ip 7.7.7.7,
ip 8.8.8.8,
}
local addrset
No address sets.
containers are shared for this filter
global containers
container 13e599b5-dd2d-420f-8473-9d45f0d324ac {
# generation number: 21208
# realization time : 2020-12-16T23:41:30
FQDN : login\.microsoft\.com(3940c0d7-cbfc-abbb-35b4-786fc4199684),
}
container 8774c654-0f9e-43ad-a803-4aa720e590cf {
# generation number: 21208
# realization time : 2020-12-16T23:41:30
FQDN : outlook\.office365\.com(6e465c1d-7d81-9672-00e1-76ddfc280b8b),
}

container 88dc6bf0-808e-49f6-a692-dd0e5cee6ab3 {
# generation number: 21208
# realization time : 2020-12-16T23:41:30
APP_ID : APP_360ANTIV,
}

container acf76e7d-400b-438b-966f-8d5c10bebbda {
# generation number: 21208
# realization time : 2020-12-16T23:41:30
APP_ID : APP_ACTIVDIR,
}

container e4be8d7e-e4ab-4466-8f2e-998445ead95d {
# generation number: 21208
# realization time : 2020-12-16T23:41:30
FQDN : play\.google\.com(c44ef0fc-a922-eb1b-f155-4f0625271198),
}
local containers
No containers.
[root@esxcomp-2a:~]

```

Other output for FW troubleshooting

In addition to above command option NSX allows other options to debug the NSX FW datapath on ESX. Use the help menu as below.

```
[root@esxcomp-2a:~] vsipioctl -h
Usage: help <cmd> <options>
below is a list of available cmd:
  getfilters      : get list of filters
  getfwconfig     : get rules, addrsets and containers of a filter
  getrules       : get rules of a filter
  getaddrsets    : get addrsets of a filter
  getcontainers  : get containers of a filter
  getspoofguard  : get spoofguard setting of a filter
  getflows       : get flows of a filter
  getconncount   : get active connection count
  getconnections : get active connections
  getsismstats   : get service insertion service VM stats
  getsisvctable  : dump service insertion service table
  getsinshtable  : display service insertion nsh table
  getsiproxytable : display service insertion proxy table
  getsifailedspis : get service insertion failed spi table
  getsiflowprohtable : get service insertion flow programming table
  getsislottedid : get service insertion slot id
  getsilbenablestatus : get service insertion load balance enable status
  getmeminfo     : get meminfo data
  initvsiplogging : init vsip logger
  getfqdnentries : get fqdn entries
  getdnsconfigprofile : get dns config profile for a filter
  getfilterstat  : get statistics of a filter
  gettimeout     : get connection timeout setting of a filter
  getfloodstat   : get flood protection status
  getsidcache    : get sid cache of a filter
  help          : this help message
run `vsipioctl <cmd> -h` to find out available options of a cmd.
[root@esxcomp-2a:~]
```

NSX CLI for FW troubleshooting

On ESXi, nsxcli option can be used as an alternative option to ESX cli, by typing "nsxcli" and user can use "get firewall" command tree to get the similar output as above.

```
[root@esxcomp-2a:~] nsxcli
esxcomp-2a.dg.vsphere.local>
esxcomp-2a.dg.vsphere.local> get firewall
% Command not found: get firewall

Possible alternatives:
  get firewall <vifuuid> addrsets
  get firewall <vifuuid> profile
  get firewall <vifuuid> ruleset rules
  get firewall exclusion
  get firewall ipfix-containers
  get firewall ipfix-filters
```

```

get firewall ipfix-profiles
get firewall ipfix-stats
get firewall packetlog
get firewall packetlog last <lines>
get firewall rule-stats
get firewall rule-stats total
get firewall status
get firewall thresholds
get firewall vifs

```

```

esxcomp-2a.dg.vsphere.local> get firewall packetlog last 10
Wed Dec 16 2020 UTC 23:53:55.693
2020-12-16T23:53:23.878Z fd2e9266 INET6 match PASS 5134 OUT 72 ICMP fe80::250:56ff:fea0:8d90-
>fe80::250:56ff:fea0:26dc ipv6-app-allow
2020-12-16T23:53:23.878Z 5f46e9b1 INET6 match PASS 5134 IN 72 ICMP fe80::250:56ff:fea0:8d90-
>fe80::250:56ff:fea0:26dc ipv6-app-allow
2020-12-16T23:53:29.234Z fd2e9266 INET6 TERM 5134 OUT ICMP 135 0 fe80::250:56ff:fea0:8d90-
>2001::172:16:202:22 1/1 72/64 ipv6-app-allow
2020-12-16T23:53:29.234Z 5f46e9b1 INET6 TERM 5134 IN ICMP 135 0 fe80::250:56ff:fea0:8d90-
>2001::172:16:202:22 1/1 72/64 ipv6-app-allow
2020-12-16T23:53:30.234Z fd2e9266 INET6 TERM 5134 IN ICMP 135 0 fe80::250:56ff:fea0:26dc-
>fe80::250:56ff:fea0:8d90 1/1 72/64 ipv6-app-allow
2020-12-16T23:53:30.234Z 5f46e9b1 INET6 TERM 5134 OUT ICMP 135 0 fe80::250:56ff:fea0:26dc-
>fe80::250:56ff:fea0:8d90 1/1 72/64 ipv6-app-allow
2020-12-16T23:53:35.239Z fd2e9266 INET6 TERM 5134 OUT ICMP 135 0 fe80::250:56ff:fea0:8d90-
>fe80::250:56ff:fea0:26dc 1/1 72/64 ipv6-app-allow
2020-12-16T23:53:35.241Z 5f46e9b1 INET6 TERM 5134 IN ICMP 135 0 fe80::250:56ff:fea0:8d90-
>fe80::250:56ff:fea0:26dc 1/1 72/64 ipv6-app-allow
2020-12-16T23:53:51.876Z fd2e9266 INET6 match PASS 5134 OUT 72 ICMP fe80::250:56ff:fea0:8d90-
>2001::172:16:202:22 ipv6-app-allow
2020-12-16T23:53:51.876Z 5f46e9b1 INET6 match PASS 5134 IN 72 ICMP fe80::250:56ff:fea0:8d90-
>2001::172:16:202:22 ipv6-app-allow

```

```

esxcomp-2a.dg.vsphere.local> get firewall exclusion
Wed Dec 16 2020 UTC 23:53:57.731

```

Firewall Exclusion

Exclusion count: 7

```

00894e3c-8948-4b6b-a4cd-acd3a2c21205
15f077e9-4492-4391-9f63-a99b6c978003
2936443e-128c-4b6d-9fcf-3b2fad778b08
3602f84a-8333-44f3-a3c2-e04fbf5e848f
8149b7ec-553d-48e1-af04-1ee2f5ae266e
d615679c-092e-4bfe-8c17-803fe8b3315d
da619e9d-48a0-4c82-a831-bf580d3bec05

```

```

esxcomp-2a.dg.vsphere.local> get firewall thresholds
Wed Dec 16 2020 UTC 23:53:59.905

```

Firewall Threshold Monitors

#	Name	Raised	Threshold	CurrValue	CurrSize	MaxSize	PeakEver	EverTime (ago)
1	dfw-cpu	False	60	0	--	--	0	--:--:--
2	vsip-attr	False	60	3	4 MB	128 MB	3	4d 23:35:06

3	vsip-flow	False	60	0	0 MB	312 MB	0	--:--:--
4	vsip-fprules	False	60	0	0 MB	128 MB	0	--:--:--
5	vsip-fqdn	False	60	0	0 MB	128 MB	0	--:--:--
6	vsip-module	False	60	15	153 MB	1024 MB	15	4d 23:35:06
7	vsip-rules	False	60	0	0 MB	512 MB	0	--:--:--
8	vsip-si	False	60	0	0 MB	128 MB	0	--:--:--
9	vsip-state	False	60	0	0 MB	384 MB	0	--:--:--

esxcomp-2a.dg.vsphere.local>

DFW L2 Rules Show Unknown MAC Address

After configuring a layer-2 firewall rule with one MAC set as source and another MAC set as destination, the `getrules` command on the host shows the destination MAC set as 01:00:00:00:00:00/01:00:00:00:00:00. For example,

```
[root@host1:~] vsipioctl getrules -f nic-1000052822-eth1-vmware-sfw.2
ruleset mainrs {
  # generation number: 0
  # realization time : 2018-07-26T12:42:28
  rule 1039 at 1 inout protocol tcp from any to any port 1521 accept as oracle;
  # internal # rule 1039 at 2 inout protocol tcp from any to any port 1521 accept;
  rule 1039 at 3 inout protocol icmp from any to any accept;
  rule 2 at 4 inout protocol any from any to any accept with log;
}

ruleset mainrs_L2 {
  # generation number: 0
  # realization time : 2018-07-26T12:42:28
  rule 1040 at 1 inout ethertype any stateless from addrset
d83a1523-0d07-4b18-8a5b-77a634540b57 to addrset 9ad9c6ef-c7dd-4682-833d-57097b415e41 accept;
  # internal # rule 1040 at 2 in ethertype any stateless from addrset
d83a1523-0d07-4b18-8a5b-77a634540b57 to addrset 9ad9c6ef-c7dd-4682-833d-57097b415e41 accept;
  # internal # rule 1040 at 3 out ethertype any stateless from addrset
d83a1523-0d07-4b18-8a5b-77a634540b57 to mac 01:00:00:00:00:00/01:00:00:00:00:00 accept;
  rule 1 at 4 inout ethertype any stateless from any to any accept;
}
```

The internal OUT rule with the address 01:00:00:00:00:00/01:00:00:00:00:00 is created by design to handle outbound broadcasting packets and does not indicate a problem. The firewall rule will work as configured.

Troubleshooting Distributed Firewall on KVM Hosts

To troubleshoot firewall issues with a KVM host, you can look at the firewall rules that apply on the host.

Get the list of VIFs on the KVM host

```
localhost.localdomain> get firewall vifs
***sample output***

                          Firewall VIFs
-----VIF count: 3

1   239140cf-6c6c-464f-96eb-dfb13203171e
2   eb277d27-0d28-4fb0-82ce-f59d86ea5bee
3   afb2aa98-85ee-4bb4-8318-d699fa84c7f0
```

Discover firewall rules that apply to a specific VIF

Specify the VIF by UUID (in this example: 239140cf-6c6c-464f-96eb-dfb13203171e).

```
localhost.localdomain> get firewall 239140cf-6c6c-464f-96eb-dfb13203171e ruleset rules
***sample output***

                          Firewall Rules
-----

VIF UUID : 239140cf-6c6c-464f-96eb-dfb13203171e
Ruleset UUID : 7c5838e5-ab75-427d-b4dd-9452e5607805
Rule count : 5345
rule 3073 inout protocol any from any to any profile fbb4b84f-f6c1-40c5-a509-f7c6f81fe7d9
accept with log tag dns;
rule 3072 inout protocol any from any to any profile 6bc09f62-a188-4e36-9708-291af7237039
accept with log tag youtube.com;
rule 3072 inout protocol any from any to any profile 27b9a15b-8071-4d09-a7e8-71eecfca0779
accept with log tag youtube.com;
rule 3075 inout protocol tcp from addrset 81d95211-ab77-4f2d-beaf-3e15b045fb5e to addrset
3d41a802-a899-4464-ba2b-da9240598552 port 5000 accept with log tag portlist1;
rule 3075 inout protocol tcp from addrset 81d95211-ab77-4f2d-beaf-3e15b045fb5e to addrset
3d41a802-a899-4464-ba2b-da9240598552 port 4992/0xffff8 accept with log tag portlist1;
rule 3075 inout protocol tcp from addrset 81d95211-ab77-4f2d-beaf-3e15b045fb5e to addrset
3d41a802-a899-4464-ba2b-da9240598552 port 4864/0xff80 accept with log tag portlist1;
```

Get the list of address sets used in a specific VIF

Specify the VIF by UUID (in this example: 239140cf-6c6c-464f-96eb-dfb13203171e).

```
localhost.localdomain> get firewall 239140cf-6c6c-464f-96eb-dfb13203171e addrsets
***sample output***

                          Firewall Address Sets
-----

Address set count : 11
UUID : 09f6da50-bcf2-4347-91a7-df00dca003a6
Address count : 7
mac 00:50:56:81:9b:2e
mac 00:0c:29:03:4d:0d
mac 00:0c:29:03:4d:03
ip 10.172.177.231
ip 10.172.177.111
ip 192.168.1.11
ip 192.168.2.11
```


Get the list of APPIDs and FQDNs used in a specific VIF

To check the FQDN profiles on the hypervisor run the command `localhost.localdomain> get firewall <vif-id> profile`

Look for the URL which was configured on the policy UI.

Specify the VIF by UUID (in this example: 239140cf-6c6c-464f-96eb-dfb13203171e).

```
localhost.localdomain> get firewall 239140cf-6c6c-464f-96eb-dfb13203171e profile
***sample output***

                          Firewall Profiles
-----

Profiles count : 9
  UUID : 87de2b6b-bdf5-49b6-bae2-824f455a21a4
  Attribute count : 2
    FQDN : www\.youtube\.com
    FQDN : .*\.microsoft\.com

  UUID : 68dc8321-5cb5-4cd4-b1d1-14961d71c05e
  Attribute count : 1
    APP_ID : APP_SSL
```

Get the list of APPIDs and FQDNs used in a specific VIF

Specify the VIF by UUID (in this example: 239140cf-6c6c-464f-96eb-dfb13203171e).

```
localhost.localdomain> get firewall 239140cf-6c6c-464f-96eb-dfb13203171e profile
***sample output***

                          Firewall Profiles
-----

Profiles count : 9
  UUID : 87de2b6b-bdf5-49b6-bae2-824f455a21a4
  Attribute count : 2
    FQDN : www\.youtube\.com
    FQDN : .*\.microsoft\.com

  UUID : 68dc8321-5cb5-4cd4-b1d1-14961d71c05e
  Attribute count : 1
    APP_ID : APP_SSL
```

Discover FQDN of specific VIF

```
localhost.localdomain> get firewall 239140cf-6c6c-464f-96eb-dfb13203171e fqdn

                          Firewall Profile FQDN
-----

Profiles count   : 3
  Profile UUID   : 87de2b6b-bdf5-49b6-bae2-824f455a21a4
  FQDN count     : 2
    FQDN UUID    : 37efd4dd-961c-4756-afdd-ec04f44b6c10
    Value        : www\.youtube\.com
    IP set       : 172.217.6.46
```

Check connections through the Linux Conntrack module.

In this example, we look for flows between two specific IP addresses.

```
ovs-appctl dpctl/dump-conntrack -m | grep 192.168.1.15 | grep 192.168.1.16
icmp,orig=(src=192.168.1.15,dst=192.168.1.16,id=7972,type=8,code=0),
reply=(src=192.168.1.16,dst=192.168.1.15,id=7972,type=0,code=0),
id=2901517888,zone=61437,status=SEEN_REPLY|CONFIRMED,mark=2083,labels=0x1f
```

Check FQDN Roles and Profiles On KVM Hosts

Firewall rules can be created to filter specific domains with FQDN/URLs. To check the FQDN profiles on the hypervisor run the command `localhost.localdomain> get firewall <vif-id> profile`. Look for the URL which was configured on the policy UI.

Sample `nsxcli` output for published context profile with APP_ID and FQDN entries:

```
localhost.localdomain> get firewall 989bdcf6-c6fc-47cd-86a3-367e552dba32 profile
Firewall Profiles
-----
Profiles count : 3
UUID : b34b868e-f113-4463-84a6-14736e50168e
Attribute count : 1
APP_ID : APP_HTTP
UUID : c4689750-d5e1-41f5-ba2c-0bfc846ed494
Attribute count : 1
FQDN : www\.youtube\.com
UUID : 77a599db-b2d3-4510-bbff-fa2bb3laceae
Attribute count : 1
APP_ID : APP_DNS
localhost.localdomain> get firewall 989bdcf6-c6fc-47cd-86a3-367e552dba32 fqdn
Firewall Profile FQDN
-----
Profiles count : 1
Profile UUID : c4689750-d5e1-41f5-ba2c-0bfc846ed494
FQDN count : 1
FQDN UUID : 1c9d612c-c398-409e-b6f0-f1ec49b778fe
Value : www\.youtube\.com
IP set : 172.217.6.46, 172.217.164.110, 172.217.5.110, 216.58.194.206, 172.217.6.78,
172.217.0.46, 216.58.195.78, 216.58.194.174
```

Troubleshooting Gateway Firewall

Use the user interface and API to troubleshoot gateway firewall.

Use NSX Manager UI and API to check the following:

- Gateway Firewall is enabled for the given Gateway.
- Check the realization state for a given gateway firewall policy. The UI shows the realization status next to the top right side of the FW Policy header.
- Check rule stats to see any traffic is hitting the FW policy.
- Enable logging for the rule for troubleshooting the policy.

Gateway firewall is implemented on NSX Edge transport node. As a next step, use datapath troubleshooting as below using nsxcli commands on the NSX Edge node command prompt.

Get UUID of the Gateway on which Firewall is enabled

```
EDGE-VM-A01> get logical-router
Logical Router
UUID                                VRF    LR-ID  Name
Type                                Ports
736a80e3-23f6-5a2d-81d6-bbefb2786666  0      0
TUNNEL                                4
8ccc0151-82bd-43d3-a2dd-6a31bf0cd29b  1      1      DR-DC-Tier-0-GW
DISTRIBUTED_ROUTER_TIER0             5
5a914d04-305f-402e-9d59-e443482c0e15  2      1025   SR-DC-Tier-0-GW
SERVICE_ROUTER_TIER0                 7
495f69d7-c46e-4044-8b40-b053a86d157b  4      2050   SR-PROD-Tier-1
SERVICE_ROUTER_TIER1                 5
```

Get all Gateway interfaces using UUID

Gateway firewall is implemented per Uplink interface of a Gateway. Identify the uplink interface and get the interface ID from the output below.

```
dc02-nsx-edgevm-1> get logical-router 16f04a64-ef71-4c03-bb5c-253a61752222 interfaces
Wed Dec 16 2020 PST 17:24:13.134
Logical Router
UUID                                VRF    LR-ID  Name                                Type
16f04a64-ef71-4c03-bb5c-253a61752222  5      2059   SR-PROD-ZONE-GW
SERVICE_ROUTER_TIER1
Interfaces (IPv6 DAD Status A-DAD_Success, F-DAD_Duplicate, T-DAD_Tentative, U-
DAD_Unavailable)
  Interface      : 748d1f17-34d0-555e-8984-3ef9f9367a6c
  Ifuid          : 274
  Mode           : cpu
  Port-type      : cpu

  Interface      : 1bd7ef7f-4f3e-517a-adf0-846d7dff4e24
  Ifuid          : 275
  Mode           : blackhole
  Port-type      : blackhole

  Interface      : 2403a3a4-1bc8-4c9f-bfb0-c16c0b37680f
  Ifuid          : 300
  Mode           : loopback
  Port-type      : loopback
  IP/Mask        : 127.0.0.1/8;::1/128 (NA)

  Interface      : 16cea0ab-c977-4ceb-b00f-3772436ad972      <<<<<<<<<< INTERFACE ID
  Ifuid          : 289
  Name           : DC-02-Tier0-A-DC-02-PROD-Tier-1-t1_lrp
  Fwd-mode       : IPV4_ONLY
  Mode           : lif
  Port-type      : uplink      <<<<<<<<<< Port-type Uplink
Interface
```

```

    IP/Mask      :
100.64.96.1/31;fe80::50:56ff:fe56:4455/64 (NA);fc9f:aea3:1afb:d800::2/64 (NA)
    MAC         : 02:50:56:56:44:55
    VNI         : 69633
    Access-VLAN  : untagged
    LS_port     : be42fb2e-b10b-499e-a6a9-221da47a4bcc
    Urpf-mode   : NONE
    DAD-mode    : LOOSE
    RA-mode     : SLAAC_DNS_THROUGH_RA (M=0, O=0)
    Admin       : up
    Op_state    : up
    MTU         : 1500
    arp_proxy   :

```

Get Gateway Firewall Rules on a GW Interface

Use Interface ID to get firewall rules programmed on a gateway interface.

```

dc02-nsx-edgevm-2> get firewall 16cea0ab-c977-4ceb-b00f-3772436ad972 ruleset rules
Wed Dec 16 2020 PST 17:43:53.047
DNAT rule count: 0

SNAT rule count: 0

Firewall rule count: 6
  Rule ID   : 5137
  Rule      : inout protocol tcp from any to any port {22, 443} accept with log

  Rule ID   : 3113
  Rule      : inout protocol icmp from any to any accept with log

  Rule ID   : 3113
  Rule      : inout protocol ipv6-icmp from any to any accept with log

  Rule ID   : 5136
  Rule      : inout protocol any from any to any accept with log

  Rule ID   : 1002
  Rule      : inout protocol any from any to any accept

  Rule ID   : 1002
  Rule      : inout protocol any stateless from any to any accept

dc02-nsx-edgevm-2>

```

Check Gateway Firewall Sync status

Gateway Firewall sync flow status between Edge Nodes for high availability. Gateway firewall sync config can be seen using the output below.

```
dc02-nsx-edgevm-1> get firewall 16cea0ab-c977-4ceb-b00f-3772436ad972 sync config
Wed Dec 16 2020 PST 17:30:55.686
HA mode           : secondary-active
Firewall enabled  : true
Sync pending      : false
Bulk sync pending : true           Last status: ok
Failover mode     : non-preemptive
Local VTEP IP     : 172.16.213.125
Peer VTEP IP      : 172.16.213.123
Local context     : 16f04a64-ef71-4c03-bb5c-253a61752222
Peer context      : 16f04a64-ef71-4c03-bb5c-253a61752222

dc02-nsx-edgevm-1>

dc02-nsx-edgevm-2> get firewall 16cea0ab-c977-4ceb-b00f-3772436ad972 sync config
Wed Dec 16 2020 PST 17:47:43.683
HA mode           : primary-passive
Firewall enabled  : true
Sync pending      : false
Bulk sync pending : true           Last status: ok
Failover mode     : non-preemptive
Local VTEP IP     : 172.16.213.123
Peer VTEP IP      : 172.16.213.125
Local context     : 16f04a64-ef71-4c03-bb5c-253a61752222
Peer context      : 16f04a64-ef71-4c03-bb5c-253a61752222

dc02-nsx-edgevm-2>
```

Check Gateway Firewall Active Flows

Gateway firewall active flows can be seen using the command below. The flow states are synced between active and standby edge nodes for that gateway. The example below shows output from both edge-node-1 and edge-node-2.

```
dc02-nsx-edgevm-2> get firewall 16cea0ab-c977-4ceb-b00f-3772436ad972 connection
Wed Dec 16 2020 PST 17:45:55.889
Connection count: 2
0x0000000330000598: 10.166.130.107:57113 -> 10.114.217.26:22  dir in protocol tcp state
ESTABLISHED:ESTABLISHED fn 5137:0
0x040000033000058f1: 10.166.130.107 -> 10.114.217.26  dir in protocol icmp  fn 5136:0

dc02-nsx-edgevm-2>

dc02-nsx-edgevm-1> get firewall 16cea0ab-c977-4ceb-b00f-3772436ad972 connection
Wed Dec 16 2020 PST 17:47:09.980
Connection count: 2
0x0000000330000598: 10.166.130.107:57113 -> 10.114.217.26:22  dir in protocol tcp state
```

```
ESTABLISHED:ESTABLISHED fn 5137:0
0x04000003300058f1: 10.166.130.107 -> 10.114.217.26 dir in protocol icmp fn 3113:0

dc02-nsx-edgevm-1>
```

Check Gateway Firewall Logs

Gateway firewall logs provide the gateway VRF and GW Interface information, along with flow details. Gateway firewall logs can be accessed on the edge, or can be sent to Syslog Server. Firewall logs provide the logical router VRF, firewall interface ID, FW rule ID & flow details.

```
dc02-nsx-edgevm-1> get log-file syslog | find datapathd.firewallpkt

<181>1 2020-08-04T21:18:25.633996+00:00 dc02-nsx-edgevm-1 NSX 26581 FIREWALL [nsx@6876
comp="nsx-edge" subcomp="datapathd.firewallpkt"
level="INFO"] <8 16cea0abc9774ceb:b00f3772436ad972> INET reason-match PASS 3061 OUT 48 TCP
10.114.217.26/33646->10.114.208.136/22 S

<181>1 2020-08-04T21:18:41.182424+00:00 dc02-nsx-edgevm-1 NSX 26581 FIREWALL [nsx@6876
comp="nsx-edge" subcomp="datapathd.firewallpkt"
level="INFO"] <2 460b362ce1254ebd:98498057bc3b18df> INET TERM PASS 3053 IN TCP
10.166.56.254/60291->10.114.217.26/22

dc02-nsx-edgevm-1>
```

Other Command Line Options for debugging Gateway Firewall

```
dc02-nsx-edgevm-2> get firewall 16cea0ab-c977-4ceb-b00f-3772436ad972

Possible alternatives:
  get firewall <uuid> addrset name <string>
  get firewall <uuid> addrset sets
  get firewall <uuid> attrset name <string>
  get firewall <uuid> attrset sets
  get firewall <uuid> connection
  get firewall <uuid> connection count
  get firewall <uuid> connection raw
  get firewall <uuid> connection state
  get firewall <uuid> ike policy [<rule-id>]
  get firewall <uuid> interface stats
  get firewall <uuid> ruleset [type <rule-type>] rules [<ruleset-detail>]
  get firewall <uuid> ruleset [type <rule-type>] stats
  get firewall <uuid> sync config
  get firewall <uuid> sync stats
  get firewall <uuid> timeouts
  get firewall [logical-switch <uuid>] interfaces
  get firewall interfaces sync

dc02-nsx-edgevm-2>
```

Check Rule Realization Status

DFW rules can be created, updated, and deleted using both the UI and API.

Rule Realization Status on UI

You can see the rule realization status for DFW and Gateway firewall policies by navigating to **Security > Distributed Firewall** or **Security Gateway Firewall**, and checking the rule realization status reported by transport nodes.

There are four possible values for the rule realization status:

- Success
- Error
- In Progress
- Unknown

Rule Realization Status Through APIs

If the rule was created and enforced at relevant nodes, the realization status can be checked by following Policy Manager APIs.

To check realization status for all the entities created in policy manager run the command: `GET: https://<Policy Appliance IP>/policy/api/v1/infra/realized-state/realized-entities` The realized state of the object should be "REALIZED" and 'runtime_status' should be "SUCCESS"

For example, the query to check the realized state of <e2d4c010-96c8-11e9-8c0a-f7581ab92530> of security policy at the Policy manager level is <f96f27c0-92b8-11e9-96af-b5e746a259e7>

is `GET https://10.172.121.219/policy/api/v1/infra/realized-state/realized-entities?intent_path=/infra/domains/default/security-policies/f96f27c0-92b8-11e9-96af-b5e746a259e7/rules/e2d4c010-96c8-11e9-8c0a-f7581ab92530`

```
{
  "results": [
    {
      "extended_attributes": [],
      "entity_type": "RealizedFirewallRule",
      "intent_paths": [
        "/infra/domains/default/security-policies/1-communication-560"
      ],
      "resource_type": "GenericPolicyRealizedResource",
      "id": "default.1-communication-560.3-communication-110",
      "display_name": "default.1-communication-560.3-communication-110",
      "description": "default.1-communication-560.3-communication-110",
      "path": "/infra/realized-state/enforcement-points/default/firewalls/firewall-sections/default.1-communication-560/firewall-rules/default.1-communication-560.3-communication-110",
      "relative_path": "default.1-communication-560.3-communication-110",
      "parent_path": "/infra/realized-state/enforcement-points/default/firewalls/firewall-sections/default.1-communication-560",
      "intent_reference": [],
      "realization_specific_identifier": "1028",
      "state": "REALIZED",
      "alarms": [],
      "runtime_status": "IN_PROGRESS",
      "_create_user": "system",
    }
  ]
}
```

```

"_create_time": 1561673625030,
"_last_modified_user": "system",
"_last_modified_time": 1561674044534,
"_system_owned": false,
"_protection": "NOT_PROTECTED",
"_revision": 6
},
"result_count": 1
}

```

To check the overall realized status of section of every rule in a section on the hypervisor run the command: `GET https://<policy-mgr>/policy/api/v1/infra/realized-state/status?include_enforced_status=true&intent_path=<Security_policy_path>`.

There are four possible values for the consolidated status:

- Success
- Error
- In Progress
- Unknown

Table 13-14. Consolidated Status

Transport Node 1 Overall Status	Transport Node 2 Overall Status	Consolidated Status
ERROR	ERROR	ERROR
ERROR	IN_PROGRESS	ERROR
ERROR	UNKNOWN	ERROR
IN_PROGRESS	IN_PROGRESS	IN_PROGRESS
IN_PROGRESS	UNKNOWN	IN_PROGRESS
SUCCESS	SUCCESS	SUCCESS
SUCCESS	ERROR	ERROR
SUCCESS	IN_PROGRESS	IN_PROGRESS
SUCCESS	UNKNOWN	UNKNOWN
UNKNOWN	UNKNOWN	UNKNOWN

Distributed Firewall Packet Logs

If logging is enabled for firewall rules, you can look at the firewall packet logs to troubleshoot issues.

The log file is `/var/log/dfwpktlogs.log` for both ESXi and KVM hosts.

The following is a regular log sample for distributed firewall rules:

```
2018-07-03T19:44:09.749Z b6507827 INET match PASS mainrs/1024 IN 52 TCP 192.168.4.3/49627->192.168.4.4/49153 SEW

2018-07-03T19:46:02.338Z 7396c504 INET match DROP mainrs/1024 OUT 52 TCP 192.168.4.3/49676->192.168.4.4/135 SEW

2018-07-06T18:15:49.647Z 028cd586 INET match DROP mainrs/1027 IN 36 PROTO 2 0.0.0.0->224.0.0.1

2018-07-06T18:19:54.764Z 028cd586 INET6 match DROP mainrs/1027 OUT 143 UDP
fe80:0:0:0:68c2:8472:2364:9be/546->ff02:0:0:0:0:1:2/547
```

The elements of a DFW log file format include the following, separated by a space:

- timestamp:
- last eight digits of the VIF ID of the interface
- INET type (v4 or v6)
- reason (match)
- action (PASS, DROP, REJECT)
- rule set name/ rule ID
- packet direction (IN/OUT)
- packet size
- protocol (TCP, UDP, or PROTO #)
- SVM direction for netx rule hit
- source IP address/source port>destination IP address/destination port
- TCP flags (SEW)

For passed TCP packets there is a termination log when the session has ended:

```
2018-07-03T19:44:30.585Z 7396c504 INET TERM mainrs/1024 OUT TCP RST 192.168.4.3/49627->192.168.4.4/49153 20/16 1718/76308
```

The elements of a TCP termination log include the following, separated by a space:

- timestamp:
- last 8 digits of the VIF ID of the interface
- INET type (v4 or v6)
- action (TERM)
- ruleset name/ rule ID
- packet direction (IN/OUT)
- protocol (TCP, UDP, or PROTO #)

- TCP RST flag
- SVM direction for netx rule hit
- source IP address/source port>destination IP address/destination port
- IN packet count/OUT packet count (all accumulated)
- IN packet size/OUT packet size

The following is a sample of FQDN log file for distributed firewall rules:

```
2019-01-15T00:34:45.903Z 7c607b29 INET match PASS 1031 OUT 48 TCP 10.172.178.226/32808-
>23.72.199.234/80 S www.sway.com(034fe78d-5857-0680-81e4-d8da6b28d1b4)
```

The elements of an FQDN log include the following, separated by a space:

- timestamp:
- last eight digits of the VIF ID of the interface
- INET type (v4 or v6)
- reason (match)
- action (PASS, DROP, REJECT)
- ruleset name/ rule ID
- packet direction (IN/OUT)
- packet size
- protocol (TCP, UDP, or PROTO #) - for TCP connections, the actual reason that a connection is terminated is indicated after the following IP address
- source IP address/source port>destination IP address/destination port
- TCP flags - S (SYN), SA (SYN-ACK), A (ACK), P (PUSH), U (URGENT), F (FIN), R (RESET)
- domain name/UUID where UUID is the binary internal representation of the domain name

The following is a sample of Layer 7 log file for distributed firewall rules:

```
2019-01-15T00:35:07.221Z 82f365ae INET match REJECT 1034 OUT 48 TCP 10.172.179.6/49818-
>23.214.173.202/80 S APP_HTTP

2019-01-15T00:34:46.486Z 7c607b29 INET match PASS 1030 OUT 48 UDP 10.172.178.226/42035-
>10.172.40.1/53 APP_DNS
```

The elements of a Layer 7 log include the following, separated by a space:

- timestamp:
- last eight digits of the VIF ID of the interface
- INET type (v4 or v6)
- reason (match)

- action (PASS, DROP, REJECT)
- ruleset name/ rule ID
- packet direction (IN/OUT)
- packet size
- protocol (TCP, UDP, or PROTO #) - for TCP connections, the actual reason that a connection is terminated is indicated after the following IP address
- source IP address/source port>destination IP address/destination port
- TCP flags - S (SYN), SA (SYN-ACK), A (ACK), P (PUSH), U (URGENT), F (FIN), R (RESET)
- APP_XXX is the discovered application

Bare Metal Server Security

Secure workloads running on Windows Server 2016 bare metal servers.

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


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

The workloads can be on overlay or VLAN-backed networks and the workloads must not be outside of the perimeter of a Windows Server 2016 bare metal server. As part of the configuration, the NSX Agent is installed on the bare metal host. Network connectivity between the application IP address of the Windows bare metal server, NSX Agent and NSX Manager needs to be established before applying DFW rules to protect workloads.


Apply DFW rules to secure ingress and egress traffic flowing through the L2 and L3 networks between workloads on a Windows Server 2016 bare metal server and virtual or physical workloads.

A couple of use cases where ingress and egress traffic is filtered at the Windows bare metal server.

Traffic between Virtual and Physical Bare Metal Workloads



Traffic between Physical Bare Metal Workloads



Before you apply DFW rules to the Windows bare metal workloads, integrate NSX-T on the Windows Server using Ansible scripts. To install and integrate NSX-T on a Windows bare metal server, refer to the Secure Workloads on Windows Server 2016 Bare Metal Server topic in the *NSX-T Data Center Installation Guide*.

You can configure services, groups, context profiles, and virtual machines for the NSX-T Data Center inventory.

When you click the **Inventory** tab, an overview of the inventory objects is displayed, showing the number of groups, services, virtual machines, and context profiles that are in the inventory. In addition, the following information about groups is shown:

- the number of groups used in policies
- the number of groups not used in policies
- the number of groups with members
- the number of groups without members
- the number of identity groups
- the number of identity groups used in policies
- the number of identity groups not used in policies

Read the following topics next:

- [Add a Service](#)
- [Add a Group](#)
- [Add a Context Profile](#)
- [Containers](#)
- [Public Cloud Services](#)
- [Physical Servers](#)
- [Tags](#)

Add a Service

You can configure a service, and specify parameters for matching network traffic such as a port and protocol pairing.

You can also use a service to allow or block certain types of traffic in firewall rules. You cannot change the type after you create a service. Some services are predefined and cannot be modified or deleted.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Inventory > Services**.
- 3 Click **Add Service**.
- 4 Enter a name.
- 5 Click **Set Service Entries**.
- 6 Select a type.

The choices are **Layer 2** and **Layer 3 and above**.

- 7 Under **Port-Protocol**, click **Add Service Entry** to add one or more service entries.

For layer 2, the only available service type is **Ether**.

For layer 3 and above, the available service types are **IP, IGMP, ICMPv4, ICMPv6, ALG, TCP, and UDP**.

- 8 Click the **Services** tab to add one or more services.

Any service that you add is considered a nested service because it is included in the service that you are creating. The recommended maximum level of nesting is 3. An example of three levels of nesting: service A includes service B, service B includes service C, and service C includes service D. In addition, cyclic nesting is not allowed. In the previous example, service C cannot include service A or B.

- 9 Click **Apply**.
- 10 (Optional) Add one or more tags.
- 11 (Optional) Enter a description.
- 12 Click **Save**.

Add a Group

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

Groups can be configured to contain a combination of virtual machines, IP sets, MAC sets, segment ports, segments, AD user groups, and other groups. Dynamic inclusion of groups can be based on tag, machine name, OS name, or computer name.

Note If you create a group in the API using LogicalPort based criteria, you cannot edit the group in the UI using the AND operator between SegmentPort criteria.

Groups can also be excluded from firewall rules, and there are a maximum of 100 groups that can be on the list. IP sets, MAC sets, and AD groups cannot be included as members in a group that is used in a firewall exclusion list. See [Manage a Firewall Exclusion List](#) for more information.

A single IP or AD group can be used as the source only within a distributed firewall rule. If IP and AD groups are needed at the source, create two separate firewall rules.

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

Note When a host is added to or removed from a vCenter Server, the external ID of the VMs on the host changes. If a VM is a static member of a group and the VM's external ID changes, the NSX Manager UI will no longer show the VM as a member of the group. However, the API that lists the groups will still show that the group contains the VM with its original external ID. If you add a VM as a static member of a group and the VM's external ID changes, you must add the VM again using its new external ID. You can also use dynamic membership criteria to avoid this issue.

Tags in NSX are case-sensitive, but a group that is based on tags is "case- insensitive." For example, if the dynamic grouping membership criterion is `VM Tag Equals 'quarantine'`, the group includes all VMs that contain either the tags 'quarantine' or 'QUARANTINE'.

If you are using NSX Cloud, see [Group VMs using NSX-T Data Center and Public Cloud Tags](#) for information on the how to use public cloud tags to group your workload VMs in NSX Manager.

Prerequisites

If you are using Federation, see [Security in NSX Federation](#) for details on configuration options.

Note If you are using NSX Federation, you cannot create groups from the Global Manager that include AD user groups.

Procedure

- 1 Select **Inventory > Groups** from the navigation panel.
- 2 Click **Add Group**.
- 3 Enter a group name.
- 4 If you are adding a group from a Global Manager for Federation, either accept the default region selection, or select a region from the drop-down menu. Once you create a group with a region, you cannot edit the region selection. However, you can change the span of the region itself by adding or removing locations from it. You can create customized regions before you create the group. See [Create a Region from Global Manager](#).

Note For groups added from a Global Manager in a Federation environment, selecting a region is mandatory. This text box is not available if you are not using the Global Manager.

5 (Optional) Click **Set Members**.

For each membership criterion, you can specify up to five rules, which are combined with the logical AND operator. The available member criterion can apply to the following:

- **Segment Port** - specify a tag, scope, or both.
- **Segment** - specify a tag, scope, or both.
- **Virtual Machine** - specify a name, tag, computer OS name, or computer name that equals, contains, starts with, ends with, or does not equal a particular string.
- **IP Set** - specify a tag, scope, or both.

6 (Optional) Click **Members** to select members.

The available member types are:

■ **Groups**

Note If you are using Federation, you can add a group as a member that has an equal or smaller span than the region you selected for the group you are creating from the Global Manager, see [Security in NSX Federation](#) .

■ **Segments**

Note IP addresses assigned to a gateway interface, and NSX load balancer virtual IP addresses are not included as segment group members.

■ **Segment Ports**

■ **VIFs**

■ **Virtual Machines**

■ **Physical Servers**

■ **Cloud Native Service Instances**

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 (Optional) Click **AD Groups** to add Active Directory Groups. Groups with Active Directory members can be used in the source text box of a distributed firewall rule for Identity Firewall. Groups can contain both AD and compute members.

If you are using NSX Federation, you cannot create groups from the Global Manager to include AD groups.

9 (Optional) Enter a description and tag.

10 Click **Apply**

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

Add a Context Profile

Context profiles enable creating attributes key value pairs such as layer 7 App Id, and Domain Names. After a context profile has been defined, it can be used in one or more distributed firewall rules and gateway firewall rules.

There are two attributes for use in context profiles: App Id and Domain (FQDN) Name. Select App Ids can have one or more sub attributes, such the TLS_Version and CIPHER_SUITE. Both App Id and domain name can be used in a single context profile. Multiple App Ids can be used in the same profile. One App Id with sub attributes can be used - sub attributes are cleared when multiple App Id attributes are used in a single profile.

Currently, a predefined list of domains is supported. You can see the list of FQDNs when you add a new context profile of attribute type *Domain (FQDN) Name*. You can also see a list of FQDNs by running the API call `/policy/api/v1/infra/context-profiles/attributes?attribute_key=DOMAIN_NAME`.

Procedure

- 1 Select **Inventory > Context Profiles**.
- 2 Click **Add New Context Profile**.
- 3 Enter a **Profile Name**.
- 4 In the Attributes column, click **Set**.
- 5 Select an attribute, or click **Add Attribute**, and select **App Id**, **URL Category**, or **Domain (FQDN) Name**.
- 6 Select one or more attributes.
- 7 (Optional) If you have selected an attribute with sub attributes such as SSL or CIFS, click **Set** in the Sub Attributes/Values column.
 - a Click **Add Sub Attribute** and select a sub attribute category from the drop-down menu.
 - b Select one or more sub attributes.
 - c Click **Add**. Another sub attribute can be added by clicking **Add Sub Attribute**.
 - d Click **Apply**.
- 8 Click **Add**.
- 9 (Optional) To add another type of attribute, click **Add Attribute** again.
- 10 Click **Apply**.
- 11 (Optional) Enter a description.
- 12 (Optional) Enter a tag.
- 13 Click **Save**.

What to do next

Apply this context profile to a layer 7 distributed firewall rule (for layer 7 or Domain name) or gateway firewall rule (for layer 7).

Containers

You can view the inventory of containers-related objects by navigating to **Inventory > Containers**.

Containers-related objects are configured through NSX Container Plugin (NCP). For NCP documentation, go to <https://docs.vmware.com/en/VMware-NSX-T-Data-Center/index.html>.

The **Namespaces** tab shows the namespaces that you have configured. The following information is displayed:

- Namespace
- Type
- Cluster
- IP Address
- Pods
- Services
- Networking
- Networking Status

You can expand each row to see more details such as Ingress Rules, Labels, Network and Security Policies.

You can click the value in the **Pods**, **Services**, and **Networking** fields to get more information.

The **Clusters** tab shows the clusters that you have configured. The following information is displayed:

- Container Cluster Name
- Infrastructure Type
- Nodes
- Namespaces
- Pods
- Services
- Networking
- Networking Status

You can click the value in the **Nodes**, **Namespaces**, **Pods**, **Services**, and **Networking** fields to get more information.

The **Pods** screen displays the following information:

- Pod Name
- Container Node
- Transport Node
- IP Address (Only IP addresses from the pod's corresponding segment ports are shown. For the pod's network interfaces that are not attached to segments, their IP addresses are not shown.)
- Segment
- Segment Port
- Services
- Labels
- Status
- Networking Status

The **Services** screen displays the following information:

- Service Name
- Namespace
- Networking
- Labels
- Status
- Networking Status

The **Networking** screen displays the following information:

- Entity Name
- Entity Type
- Connectivity
- Tags
- Status

The **Nodes** screen displays the following information:

- Node Name
- External ID
- IP Address
- Labels
- Networking Status

Public Cloud Services

You can see a list of public cloud services that are available for your public cloud workload VMs.

Public cloud services that can be protected using cloud native security constructs, can be on-boarded with NSX Cloud.

Note In the current release, only the following AWS services are supported:

- RDS
 - Application ELB (network ELB not supported)
-

How to onboard public cloud services using NSX Cloud

You onboard public cloud services in the same way as you onboard workload VMs in the Native Cloud Enforced Mode.

Once onboarded, the public cloud services are available from **Inventory > Public Cloud Services**.

You can create firewall rules for these services in the same way as for workload VMs.

See [Managing VMs in the Native Cloud Enforced Mode](#).

Note You must enable the ports used by these services, for example, port 80 for ELB, in the firewall rules in NSX Manager.

Physical Servers

You can view the inventory of physical servers by navigating to **Inventory > Physical Servers**. These are transport nodes running on bare-metal servers.

For each physical server, the following information is displayed:

- Name
- OS type
- IP address
- Tags

Tags

Tags help you to label NSX-T Data Center objects so that you can quickly search or filter objects, troubleshoot and trace, and do other related tasks.

You can create tags using both the UI and APIs. Each tag has the following two attributes:

- Tag (refers to the tag name. It is required, must be unique and case-sensitive.)
- Scope (optional)

Tag scope is analogous to a key and tag name is analogous to a value. For example, let us say, you want to label all virtual machines based on their operating system (Windows, Mac, Linux). You can create three tags, such as Windows, Linux, and Mac, and set the scope of each tag to OS. Other examples of tag scope can be tenant, owner, name, and so on.

After you save a tag, you cannot update the name and scope. However, you can unassign or remove tags from objects.

For information about the maximum number of tags supported in NSX-T Data Center objects, see the VMware Configuration Maximums tool at <https://configmax.vmware.com/home>.

Following are some of the operations that you can do with tags:

- Assign or unassign tags to an object.
- Assign or unassign a single tag to multiple objects simultaneously (supported only for VMs).
- View a list of all tags in the inventory.
- Filter the list of tags by tag name, tag source, and tag scope.
- View a list of objects that are assigned a specific tag.

Use Cases of Tags

The following table describes some use cases of using tags.

Use Case	Description
Manageability	■ Simplify searching of objects in a large-scale inventory management. ■ Provide more information to differentiate objects that share similar or unclear names.
Third-party sharing and context sharing	■ Annotate objects with custom information. ■ Allow third-party non-NSX systems to add metadata information in an automated fashion. For example, metadata from partners, cloud management providers, container platforms, and so on. ■ Capture attributes or properties and relationships that are learned using NSX discovery agent, inventory collection, public cloud agent, Guest Introspection, VM Tools, and so on.
Security	■ Create grouping membership criteria. ■ Specify the firewall source and destination.
Troubleshooting (Traceability)	■ Trace a firewall rule into the logs (Rule tags) ■ Trace and correlate objects back to an OpenStack network.

System Tags

System tags are tags that are system-defined, and you cannot add, edit, or delete them.

Table 14-1. System Tags in Public Cloud Manager Objects

Objects	System Tags
Logical Switch	■ CrossCloud
Node	■ CloudType
Logical Router	■ CloudScope
Logical Router Uplink Port	■ CloudRegion
Static Route	■ CloudVpcId
DHCP Profile	■ PcId
Firewall Section Rule List	■ EntityType
NAT Rule	■ CrossCloud ■ CloudType ■ CloudScope ■ CloudRegion ■ CloudVpcId ■ PcId ■ EntityType ■ DefaultSnatRule ■ DefaultLinkLocalSNatRule/Cloud-Public-IP ■ DefaultSiNatRule

Table 14-2. System Tags in Cloud Service Manager (CSM) Objects

Objects	System Tags
BFD Health Monitoring Profile	■ CrossCloud
Transport Zone	■ CloudType
Uplink Host Switch Profile	■ CloudScope
Transport Node	■ CloudRegion
Edge Cluster	■ CloudVpcId ■ PcId ■ EntityType

Table 14-3. System Tags in NSX Cloud VMs

Tag Source	System Tags
Amazon	■ aws:account ■ aws:availabilityzone ■ aws:region ■ aws:vpc ■ aws:subnet ■ aws:transit_vpc
Microsoft Azure	■ azure:subscription_id ■ azure:region ■ azure:vm_rg ■ azure:vnet_name ■ azure:vnet_rg ■ azure:transit_vnet_name ■ azure:transit_vnet_rg

Table 14-4. System Tags in Other NSX-T Data Center Objects

Objects	System Tags
Group	■ autoPlumbing ■ abstractionPath ■ NLB-VIP_ID ■ NLB-Lb-ID ■ NLB-Pool_ID
Segment	■ subnet-cidr
IP Address Pool	■ abstractionPath
IP Address Block	

Discovered Tags

NSX-T Data Center can discover and synchronize tags from Amazon and Microsoft Azure.

Discovered tags are tags that you have added to your VMs in the public cloud and are automatically discovered by NSX Cloud. The discovered tags are displayed for your workload VMs in the NSX Manager inventory. You cannot edit these tags in the UI.

The prefix for discovered AWS tags is "dis:aws", and the prefix for discovered Azure tags is "dis:azure". When you make changes to the tags in the public cloud, the changes are reflected in NSX Manager. By default, this feature is enabled.

You can enable or disable the discovery of AWS tags at the time of adding the AWS account. Similarly, you can enable or disable Microsoft Azure tags at the time of adding the Microsoft Azure subscription.

Add Tags to an Object

You can select existing tags that are available in the NSX-T Data Center inventory or create new tags to add to an object.

The following procedure explains the steps for adding tags to a single object. For this procedure, the virtual machine object is considered. The steps for adding tags to other objects remain the same. You can navigate to the specific object page, and follow similar steps to add tags to that object.

For information about the maximum number of tags supported in NSX-T Data Center objects, see the VMware Configuration Maximums tool at <https://configmax.vmware.com/home>.

Procedure

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

- 2 Edit an object.

For example, click **Inventory > Virtual Machines**. Next to the virtual machine that you want to edit, click the vertical ellipses, 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 a Tag to Multiple Objects

Starting in NSX-T Data Center 3.0, you can add a tag to multiple objects simultaneously. However, in v3.0, this feature is available only for the virtual machine object.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Click **Inventory > Tags**.
- 3 Click **Add Tag**.
- 4 Enter a tag name.
The maximum length of the tag name is 256 characters.
- 5 (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.
- 6 In **Assigned To**, click **Set Virtual Machines**.
- 7 (Required) Select one or more virtual machines to which you want to assign the tag, and click **Apply**.

You must assign a tag to at least one virtual machine before you can save the tag.

Note You can do a bulk assignment of a tag on a maximum of 1000 virtual machines at one go.

- 8 Click **Save**.

Results

- If the tag is assigned to many virtual machines, the assignment might take some time. When the assignment is in progress, the **Last Assignment Status** shows *Running*. After the tag is assigned successfully to all the selected virtual machines, the **Last Assignment Status** column changes to *Successful*.
- If a partial assignment occurs, NSX-T Data Center does not roll back the tag assignment from the VMs on which the tag is applied. For example, assume that you selected 100 VMs for a bulk tag assignment, and the assignment fails for 10 VMs. The tag that is assigned on the remaining 90 VMs is not rolled back.

In such partial assignment situations, run the following API to retrieve the status of the tag operation:

```
GET /api/v1/infra/tags/tag-operations/<tag-operation-id>/status
```

You can also retrieve the realized status of the tag operation with the following API:

```
GET /api/v1/infra/realized-state/realized-entities?intent_path=/infra/
tags/tag-operations/<operation-id>
```

For more details about these APIs, see the *NSX-T Data Center API Guide*.

What to do next

If you have a long list of tags in the inventory, you can filter or search tags to find the tags of your interest quickly. You can filter on source, scope, and tag (name of the tag). You can also sort tags in the UI. However, due to the case-sensitive nature of tags, tags are sorted only in a lexical order.

The following limitations apply to searching or filtering tags:

- You cannot filter tags on the source and scope attributes simultaneously because both work on the scope attribute of the tag.
- The API does not support filtering tags with special characters, such as *, &, /, \, and so on. However, you can use special characters to filter tags in the UI.

Unassign Tags from an Object

You can remove tags that you had assigned previously to an object.

The following procedure explains the steps for unassigning tags from a single NSX-T Data Center object. For this procedure, the virtual machine object is considered. The steps for unassigning tags from other objects remain the same. You can navigate to the specific object page, and follow similar steps to unassign tags from that object.

Procedure

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

For example, click **Inventory > Virtual Machines**. Next to the virtual machine that you want to edit, click the vertical ellipses, and click **Edit**.
- 3 Click the **X** icon for each tag that you want to unassign from the virtual machine.
- 4 Click **Save**.

Unassign a Tag from Multiple Objects

Starting in NSX-T Data Center 3.0, you can unassign a tag from multiple objects simultaneously. However, in v3.0, this feature is available only for virtual machines.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Click **Inventory > Tags**.
- 3 Next to that tag that you want to edit, click the vertical ellipses, and click **Edit**.
- 4 In the **Assigned To** column, click the number of virtual machines that are assigned this tag.

- 5 Click the **X** icon for each virtual machine that you want to unassign this tag.

Note

- You can do a bulk unassignment of a tag on a maximum of 1000 virtual machines at one go.
 - When a tag is unassigned from all objects, the tag is deleted automatically from the inventory after five days.
-

- 6 Click **Apply**, and then click **Save**.

Multisite and Federation

15

There are two options for managing NSX-T Data Center across multiple locations.

Table 15-1. Comparison of Multisite and Federation

	Multisite	Federation
Availability	NSX-T Data Center 2.3	NSX-T Data Center 3.0
Environments	Small and medium enterprises	Large enterprises
Environment Profile	■ Small and medium scale deployments ■ No specific site management or policy requirements	■ Large scale deployments ■ Specific site management or policy requirements
Number of NSX Manager clusters	1	1 per location
VMware Site Recovery Manager support for disaster recovery of Management Plane	Fully supported starting in NSX-T Data Center 3.0.2.	Fully supported for the recovery of Global Manager and Local Manager appliances starting in NSX-T Data Center 3.0.2.
VMware Site Recovery Manager support for disaster recovery of compute VMs	Fully supported starting in NSX-T Data Center 3.0.2.	Not Supported.
vSphere High Availability (vSphere HA) support for disaster recovery of Management Plane	Fully supported.	Fully supported for the recovery of Global Manager and Local Manager appliances starting in NSX-T Data Center 3.0.2.

Read the following topics next:

- [NSX-T Data Center Multisite](#)
- [NSX Federation](#)


NSX-T Data Center Multisite

NSX-T Data Center supports multisite deployments where you can manage all the sites from one NSX Manager cluster.

Two types of multisite deployments are supported:


- Disaster recovery
- Active-active

The following diagram illustrates a disaster recovery deployment.



In an active-active deployment, all the sites are active and layer 2 traffic crosses the site boundaries. In a disaster recovery deployment, NSX-T Data Center at the primary site handles networking for the enterprise. The secondary site is standing by to take over if a catastrophic failure occurs at the primary site.

The following diagram illustrates an active-active deployment.



You can deploy two sites for automatic or manual/scripted recovery of the management plane and the data plane.

Automatic Recovery of the Management Plane


Requirements:

- A stretched vCenter cluster with HA across sites configured.
- A stretched management VLAN.


The NSX Manager cluster is deployed on the management VLAN and is physically in the primary site. If there is a primary site failure, vSphere HA will restart the NSX Managers in the secondary site. All the transport nodes will reconnect to the restarted NSX Managers automatically. This process takes about 10 minutes. During this time, the management plane is not available but the data plane is not impacted.

The following diagrams illustrates automatic recovery of the management plane.

Before the disaster:



After disaster recovery:



Automatic Recovery of the Data Plane

You can configure failure domains for Edge nodes to achieve automatic recovery of the data plane. You can group Edge nodes within an Edge cluster in different failure domains. NSX Manager will automatically place any new active tier-1 gateway in the preferred failure domain, and the standby tier-1 gateway in the other domain.

Requirements:

- The maximum latency between Edge nodes is 10 ms.
- The HA mode for the tier-0 gateway must be active-standby, and the failover mode must be preemptive.
- If asymmetric routing is possible (for example, the two locations are two buildings without any physical firewall between them), then the HA mode for the tier-0 gateway can be active-active.

Note: The failover mode of the tier-1 gateway can be preemptive or non-preemptive, but preemptive is recommended to guarantee that the tier-0 and tier-1 gateways are in the same location.

Configuration steps:

- Using the API, create failure domains for the two sites, for example, FD1A-Preferred_Site1 and FD2A-Preferred_Site1. Set the parameter `preferred_active_edge_services` to `true` for the primary site and set it to `false` for the secondary site.

```
POST /api/v1/failure-domains
{
  "display_name": "FD1A-Preferred_Site1",
  "preferred_active_edge_services": "true"
}

POST /api/v1/failure-domains
{
  "display_name": "FD2A-Preferred_Site1",
  "preferred_active_edge_services": "false"
}
```

- Using the API, configure an Edge cluster that is stretched across the two sites. For example, the cluster has Edge nodes `EdgeNode1A` and `EdgeNode1B` in the primary site, and Edge nodes `EdgeNode2A` and `EdgeNode2B` in the secondary site. The active tier-0 and tier-1 gateways will run on `EdgeNode1A` and `EdgeNode1B`. The standby tier-0 and tier-1 gateways will run on `EdgeNode2A` and `EdgeNode2B`.
- Using the API, associate each Edge node with the failure domain for the site. First call the `GET /api/v1/transport-nodes/<transport-node-id>` API to get the data about the Edge node. Use the result of the GET API as the input for the `PUT /api/v1/transport-nodes/<transport-node-id>` API, with the additional property, `failure_domain_id`, set appropriately. For example,

```
GET /api/v1/transport-nodes/<transport-node-id>
Response:
{
  "resource_type": "TransportNode",
  "description": "Updated NSX configured Test Transport Node",
  "id": "77816de2-39c3-436c-b891-54d31f580961",
  ...
}

PUT /api/v1/transport-nodes/<transport-node-id>
{
  "resource_type": "TransportNode",
  "description": "Updated NSX configured Test Transport Node",
  "id": "77816de2-39c3-436c-b891-54d31f580961",
  ...
  "failure_domain_id": "<UUID>",
}
```

- Using the API, configure the Edge cluster to allocate nodes based on failure domain. First call the GET `/api/v1/edge-clusters/<edge-cluster-id>` API to get the data about the Edge cluster. Use the result of the GET API as the input for the PUT `/api/v1/edge-clusters/<edge-cluster-id>` API, with the additional property, `allocation_rules`, set appropriately. For example,

```
GET /api/v1/edge-clusters/<edge-cluster-id>
Response:
{
  "_revision": 0,
  "id": "bf8d4daf-93f6-4c23-af38-63f6d372e14e",
  "resource_type": "EdgeCluster",
  ...
}


PUT /api/v1/edge-clusters/<edge-cluster-id>
{
  "_revision": 0,
  "id": "bf8d4daf-93f6-4c23-af38-63f6d372e14e",
  "resource_type": "EdgeCluster",
  ...
  "allocation_rules": [
    {
      "action": {
        "enabled": true,
        "action_type": "AllocationBasedOnFailureDomain"
      }
    }
  ],
}
```

- Create tier-0 and tier-1 gateways using the API or NSX Manager UI.


In case of a full primary site failure, the tier-0 standby and tier-1 standby in the secondary site automatically take over and become the new active gateways. In case of a failure of one of the Edge nodes in the primary site, the same principle applies. For example, in the diagram below, assume that Edge node 1B hosts Tier-0-Test and Tier-1-Test, Edge node 2A hosts the Tier-0-Test standby and Edge node 2B hosts the Tier-1-Test standby. If Edge node 1B fails, the standby Tier-0-Test on Edge node 2A and standby Tier-1-Test on Edge node 2B take over and become the new active gateways.

The following diagrams illustrates automatic recovery of the data plane.

Before the disaster:



After disaster recovery:



Manual/Scripted Recovery of the Management Plane

Requirements:


- DNS for NSX Managers with a short TTL (for example, 5 minutes).
- Continuous backup.

Neither vSphere HA, nor a stretched management VLAN, is required. NSX-T Managers must be associated with a DNS name with a short TTL. All transport nodes (Edge nodes and hypervisors) must connect to the NSX Manager using their DNS name. To save time, you can optionally pre-install an NSX Manager cluster in the secondary site.

The recovery steps are:

- 1 Change the DNS record so that the NSX Manager cluster has different IP addresses.
- 2 Restore the NSX Manager cluster from a backup.
- 3 Connect the transport nodes to the new NSX Manager cluster.

The following diagram illustrates manual/scripted recovery of the management plane.



Manual/Scripted Recovery of the Data Plane

Requirement:


- The maximum latency between Edge nodes is 150 ms.

The Edge nodes can be VMs or bare metal. The tier-0 gateways in each location can be active-standby or active-active. Edge node VMs can be installed in different vCenter Servers. No vSphere HA is required.

The recovery steps are:

- 1 Using the API, move the tier-1 gateways that are connected to the primary tier-0 gateway (blue in the diagram below) to the secondary tier-0 gateway (green).
- 2 Using the API, move the standalone tier-1 gateways to the secondary site.
- 3 Using the API, move the layer-2 bridges to the secondary site.

The following diagram illustrates manual/scripted recovery of the data plane.



Requirements for Multisite Deployments

Inter-site Communication

- The bandwidth must be at least 1 Gbps and the latency (RTT) must be less than 150 ms.
- MTU must be at least 1600. 9000 is recommended.

NSX Manager

- With Automatic Recovery of Management Plane
 - VLAN management stretched between sites.
 - vSphere HA across sites for NSX Manager VMs.
- With Manual/Scripted Recovery of the Management Plane
 - Continuous backup.
 - NSX Manager must be set up to use FQDN.

Data Plane

- The same internet provider must be used if public IP addresses are exposed through services such as NAT or load balancer.

- With Automatic Recovery of Management Plane
 - Maximum latency between locations is 10 ms.
 - The HA mode for the tier-0 gateway must be active-standby and the failover mode must be preemptive to guarantee no asymmetric routing.
 - The HA mode for the tier-0 gateway can be active-active if asymmetric routing is acceptable (such as different buildings in a metropolitan region).
- With Manual/Scripted Recovery of Management Plane
 - Maximum latency between locations is 150 ms.

Cloud Management System

- The cloud management system (CMS) must support an NSX-T Data Center plug-in. In this release, VMware Integrated OpenStack (VIO) and vRealize Automation (vRA) satisfy this requirement.

Limitations

- No local-egress capabilities. All north-south traffic must occur within one site.
- The compute disaster recovery software must support NSX-T Data Center, for example, VMware SRM 8.1.2 or later.

Working with VMware Site Recovery Manager

Starting in NSX-T Data Center 3.0.2, you can use VMware Site Recovery Manager with NSX-T Data Center Multisite.

For detailed instructions on using VMware Site Recovery Manager, see: [VMware Site Recovery Manager Documentation](#).

The following Site Recovery Manager workflows are supported with NSX-T Data Center Multisite:

- NSX-T Data Center Management VMs:
 - Full recovery as well as test recovery of management VMs.
 - Recovery of management cluster using a VIP.
 - Recovery of management cluster using individual node IP addresses, instead of the VIP.
- Compute VMs:
 - Full recovery as well as test recovery of compute VMs.
 - Recovery of network services for compute VMs.
 - Recovery of VM tags, security groups and firewall rules.

NSX Federation

With NSX Federation, you can manage multiple NSX-T Data Center environments with a single pane of glass view, create gateways and segments that span one or more locations, and configure and enforce firewall rules consistently across locations.

Once you have installed the Global Manager and have added locations, you can configure networking and security from Global Manager.

For information about the initial NSX Federation configuration, including installing Global Manager and adding locations, see *Getting Started with Federation* in the *NSX-T Data Center Installation Guide*.

Overview of NSX Federation

Before you configure your NSX Federation environment, understand what features are supported, how NSX Federation shares information across locations, and how the user interface works.

Federation Key Concepts

NSX Federation introduces some new terms and concepts, such as remote tunnel endpoint (RTEP), span, and region.

Federation Systems: Global Manager and Local Manager

A Federation environment includes two types of management systems:

- Global Manager: a system similar to NSX Manager that federates multiple Local Managers.
- Local Manager: an NSX Manager system in charge of network and security services for a location.

Federation Span: Local and Stretched

When you create a networking object from Global Manager, it can span one or more locations.

- Local: the object spans only one location.
- Stretched: the object spans more than one location.

You do not directly configure the span of a segment. A segment has the same span as the gateway it is attached to.

Federation Regions

Security objects have a region. The region can be one of the following:

- Location: a region is automatically created for each location. This region has the span of that location.
- Global: a region that has the span of all available locations.
- Custom Region: you can create regions that include a subset of the available locations.

Federation Tunnel Endpoints

In a Federation environment, there are two types of tunnel endpoints.

- Tunnel End Point (TEP): the IP address of a transport node (Edge node or Host) used for Geneve encapsulation within a location.
- Remote Tunnel End Points (RTEP): the IP address of a transport node (Edge node only) used for Geneve encapsulation across locations.

Features and Configurations Supported in Federation

All configurations made from the Global Manager are made in Policy mode. Manager mode is not available in Federation.

See [Chapter 1 NSX Manager](#) for more information about the two modes.

Configuration Maximums

A Federation environment has the following configuration maximums:

- For most configurations, the Local Manager cluster has the same configuration maximums as an NSX Manager cluster. Go to [VMware Configuration Maximums tool](#) and select NSX-T Data Center.

Select the Federation category for NSX-T Data Center in the [VMware Configuration Maximums tool](#) for exceptions and other Federation-specific values.

- For a given location, the following configurations contribute to the configuration maximum:
 - Objects that were created on the Local Manager.
 - Objects that were created on the Global Manager and include the location in its span.

You can view the capacity and usage on each Local Manager. See [View the Usage and Capacity of Categories of Objects](#).

Feature Support

Table 15-2. Features Supported in Federation

Feature	Details	Related Links
Tier-0 Gateway	■ 3.0.1 and later: Active Active and Active Standby ■ 3.0.0: Active Active only	Add a Tier-0 Gateway from Global Manager
Tier-1 Gateway		Add a Tier-1 Gateway from Global Manager
Segments	Layer 2 Bridge is not supported.	Add a Segment from Global Manager
Groups	Some limitations. See Security in NSX Federation .	Create Groups from Global Manager
Distributed Firewall		Create DFW Policies and Rules from Global Manager

Table 15-2. Features Supported in Federation (continued)

Feature	Details	Related Links
Gateway Firewall		Create Gateway Policies and Rules from Global Manager
Network Address Translation (NAT)	1 Tier-0 gateway: ■ Active Active: You can configure stateless NAT only, that is, with action type Reflexive. ■ Active Standby: You can create stateful or stateless NAT rules. 2 Tier-1 gateway: ■ You can create stateful or stateless NAT rules. ■ Stateless NAT rules are pushed to all locations in the gateway's span unless scoped to one or more locations specifically. ■ Stateful NAT rules are also pushed to all locations in the gateway's span or to the specific location selected. However, stateful NAT rules are realized and enforced only on the primary location.	Configure NAT on a Gateway
DNS		See Add a DNS Forwarder Service
DHCP and SLAAC	■ DHCP Relay is supported on segments and gateways. ■ DHCPv4 server is supported on gateways with DHCP static bindings configured on segments. ■ IPv6 addresses can be assigned using SLAAC with DNS Through RA (DAD detects duplicates within a location only).	■ DHCP Relay: Add a DHCP Relay Profile ■ DHCP Server (supported on gateway only): ■ Add a DHCP Server Profile ■ Attach a DHCP Profile to a Tier-0 or Tier-1 Gateway ■ Configure DHCP Static Bindings on a Segment ■ IPv6 address assignment: Create SLAAC and DAD Profiles for IPv6 Address Assignment

Table 15-2. Features Supported in Federation (continued)

Feature	Details	Related Links
Using objects created on Global Manager in a Local Manager configuration	- Most configurations are supported. For example: ■ Connecting a Local Manager tier-1 gateway to a Global Manager tier-0 gateway. ■ Using a Global Manager group in a Local Manager distributed firewall rule. - These configurations are not supported: ■ Connecting a Local Manager segment to a Global Manager tier-0 or tier-1 gateway. ■ Connecting a load balancer to a Global Manager tier-1 gateway.	
Backup and Restore	■ 3.0.1 and later: backup with FQDN or IP is supported. ■ 3.0.0: backup with FQDN is not supported.	Backup and Restore in NSX Federation
vMotion between locations	Cold migration between locations is not supported.	

Understanding Federation


In Federation, you make configuration changes on the Global Manager. The changes are synced with the relevant Local Managers. Local Managers also sync some information with each other.

Making Changes on Global Manager


The Global Manager provides a user interface similar to the NSX Manager interface.

Configurations that are created on the Global Manager are read-only on the Local Managers. Configurations on the Local Managers are not synced with the Global Manager.

The Global Manager syncs a configuration with a Local Manager only if the configuration is relevant to that location. For example, if you create a tier-0 gateway and add it to Location 1, Location 2, and Location 3, the configuration is synced with all three Local Managers.




If the tier-0 gateway is added only to Location 1 and Location 2, the configuration is not synced with Location 3.



Making Changes on Local Managers


You can use the Local Manager to create objects on that specific Local Manager. These objects are not synced with the Global Manager or any other Local Manager.



Realizing Global Manager Changes on Local Managers

The Global Manager validates changes against the Global Manager configuration only. When a Local Manager receives a configuration from the Global Manager, the configuration is realized in the fabric nodes of that Local Manager. During this realization, errors or conflicts might be detected.

For example, you can create a tier-0 gateway from Global Manager, and then from a Local Manager you can create and attach a tier-1 gateway to the tier-0 gateway.



Because Local Managers do not sync their configurations to the Global Manager, from the Global Manager context the tier-0 gateway does not appear to be connected to anything. You can delete the tier-0 gateway from the Global Manager, and this change is synced to the Local Managers. When the changes are realized in each location, the following occurs:

- The tier-0 gateway can be deleted from the Local Manager in Location 2.
- The tier-0 gateway cannot be deleted from the Local Manager in Location 1.
- The tier-0 gateway is marked for deletion on the Global Manager.

When the tier-0 is disconnected from the tier-1 in Location 1, the tier-0 is deleted from Global Manager.

Most problems are displayed on the user interface. Additional problems can be displayed using these API calls.

- On Global Manager:

```
GET /global-manager/api/v1/global-infra/realized-state/alarms
```

- On Local Manager:


```
GET /policy/api/v1/infra/realized-state/alarms
```

Using the Global and Local Manager Web Interfaces

You can use the Global Manager to create objects that are limited to one location, or span multiple locations.

Location Drop-Down Menu on Global Manager

When you log into the Global Manager web interface, you see a Location drop-down menu in the top navigation bar. Using this menu, you can switch between the Global Manager and any associated Local Managers.



Local and Global Objects


Objects created on a Local Manager are local objects. They are specific to that Local Manager and are not viewable from the Global Manager web interface.

Objects created from the Global Manager are global objects, though their span might not include all available locations.

On a Local Manager, you can see local objects, and any global objects that apply to that location.

The global objects have an icon next to them:  .

This screenshot from the Local Manager web interface shows two segments. The segment `segment-01` has the  icon next to it, which indicates that it was created on the Global Manager. The segment `segment-02` has no  icon, which indicates that it was created on the Local Manager.




Because all objects on the Global Manager are global, there is no icon displayed when you are logged into the Global Manager.

Status of Local and Global Objects

Local Managers display the status of both global and local objects.

The Global Manager displays only global objects, but does not automatically receive the status of the objects.

To retrieve the latest status from the Local Managers, click **Check Status** for the object. To refresh the status, click the **Refresh** icon.



Overriding Global Manager Configurations on Local Manager

When you create an object from Global Manager, the same configuration is propagated to all relevant locations. Starting in NSX-T Data Center 3.0.1 you can override some Global Manager configurations on a Local Manager.

To override a configuration, click the three dots menu () next to the configuration, and click **Edit**. If the **Edit** menu item is dimmed, you cannot override this configuration.

If a configuration is overridden, you see this icon in the status column on both Global Manager and Local Manager: .

To remove an override, click the three dots menu (⋮) next to the configuration, and click **Revert**. The configuration from Global Manager is restored.

If you override a configuration from Global Manager on a Local Manager, and then you delete the configuration from the Global Manager, the configuration persists on the Local Manager. When you revert the configuration, the configuration is deleted from Local Manager.

You can get a list of all configurations that have been overridden. Make this API request to the Global Manager: GET <https://<global-mgr>/global-manager/api/v1/global-infra/overridden-resources>.

Gateway Configurations

Gateway configurations are found in **Networking > Tier-0 Gateways** and **Networking > Tier-1 Gateways**.

You can modify the following gateway configurations:

- Tier-0 Gateway BGP Configuration
- Tier-0 Gateway Interfaces

Profile Configurations

Profile configurations on Global Manager are used in all Local Managers. There is no span setting for a profile configuration.

You can override the following global profile configurations from Local Manager:

- Segment Profiles: **Networking > Segments > Segment Profiles**
 - IP Discovery Profiles
 - MAC Discovery Profiles
 - Segment Security Profiles
 - SpoofGuard Profiles
- Networking Profiles: **Networking > Networking Settings**
 - IPv6 DAD Profiles
 - IPv6 ND Profiles
 - Gateway QoS Profiles
 - BFD Profiles
- Context Profiles: **Inventory > Context Profiles**
- Security Profiles: **Security > Security Profiles**
 - Firewall Session Timer Profile
 - Edge Gateway Flood Protection Profiles
 - Firewall Flood Protection Profiles

- DNS Security Profiles
- CPU and Memory Threshold Profiles are API only:
 - Override with PUT/PATCH `https://<local-manager>/policy/api/v1/global-infra/settings/firewall/cpu-mem-thresholds-profiles/<id>?action=override`.
 - Revert with DELETE `https://<local-manager>/policy/api/v1/global-infra/settings/firewall/cpu-mem-thresholds-profiles/<id>`.
- Troubleshooting Profiles: **Plan & Troubleshoot**
 - Firewall IPFIX Profiles
 - Switch IPFIX Profiles
 - IPFIX Firewall Collector
 - IPFIX Switch Collector
 - Remote L3 Span Port Mirroring Profile
 - Logical Span Port Mirroring Profile
 - QoS Profile

Networking in NSX Federation

Tier-0 gateways, tier-1 gateways, and segments can span one or more locations in the NSX Federation environment.

When you plan your network topology, keep these requirements in mind:

- Tier-0 and tier-1 gateways can have a span of one or more locations.
- The span of a tier-1 gateway must be equal to, or a subset of, the span of the tier-0 gateway it is attached to.
- A segment has the same span as the tier-0 or tier-1 gateway it is attached to. Isolated segments are not realized until they are connected to a gateway.
- NSX Edge nodes in the Edge Cluster selected on the Global Manager for tier-0 and tier-1 gateways must be configured with the Default TZ Overlay.

You can create different topologies to achieve different goals.

- You can create segments and gateways that are specific to a given location. Each site has its own configuration, but you can manage everything from the Global Manager interface.
- You can create segments and gateways that span locations. These stretched networks provide consistent networking across sites.

Tier-0 Gateway Configurations in NSX Federation

With NSX Federation, you can deploy a tier-0 gateway that is limited to a single location, or you can stretch it across multiple locations.


Tier-0 gateways can have one of the following configurations:

- Non-stretched tier-0 gateway.
- Stretched active-active with primary and secondary locations.
- Stretched active-active with all primary locations.
- Stretched active-standby with primary and secondary locations.

Note Active-standby tier-0 gateways are supported starting in NSX-T Data Center 3.0.1.

Non-Stretched Tier-0 Gateway

You can create a tier-0 gateway from Global Manager that spans only one location. This is similar to creating the tier-0 gateway on the Local Manager directly, but has the advantage that you can manage it from Global Manager.



Stretched Active-Active Tier-0 Gateway with Primary and Secondary Locations

In an active-active tier-0 gateway with primary and secondary locations, the following applies:


- All Edge nodes are active at the same time, therefore the tier-0 cannot run stateful services.
- All traffic enters and leaves through the Edge nodes in the primary location.

If both the tier-0 gateway and the linked tier-1 gateway have primary and secondary locations, configure the same location to be primary for both gateways to reduce cross-location traffic.

Important In this topology, NSX-T Data Center ensures that all egress traffic leaves through the primary location.

If your environment has stateful services, such as external firewall, on the physical network, you must ensure that the return traffic enters through the primary location. For example, you can add AS path prepending on the BGP peers in your secondary locations.

If you do not have stateful services on your physical network, and you choose to have asymmetric routing, you must disable Unicast Reverse Path Forwarding (uRPF) on all externally tier-0 interfaces.




Stretched Active-Active Tier-0 Gateway with All Primary Locations

In an active-active tier-0 gateway with all primary locations, the following applies:

- All Edge nodes are active at the same time, therefore the tier-0 cannot run stateful services.
- All traffic enters and leaves through Edge nodes in the same location as the workloads.


Important This topology allows traffic to egress locally from each location. You must ensure that return traffic enters the same location to allow stateful services such as firewall. For example, you can configure a location-specific NAT IP so that return traffic is always routed back to the same location that it left.



Stretched Active-Standby Tier-0 Gateway with Primary and Secondary Locations

In an active-standby tier-0 gateway with primary and secondary locations, the following applies:

- Only one Edge node is active at a time, therefore the tier-0 can run stateful services.
- All traffic enters and leaves through the active Edge node in the primary location.



For Active Standby tier-0 gateways, the following services are supported:

- Network Address Translation (NAT)
- Gateway Firewall
- DNS
- DHCP


See [Features and Configurations Supported in Federation](#) for more information.

Tier-1 Gateway Configurations in Federation

With NSX Federation, you can deploy a tier-1 gateway to provide distributed routing only, or you can configure services on it.

Tier-1 Gateway for Distributed Routing Only

You can create a tier-1 gateway in Federation for distributed routing only. This gateway has the same span as the tier-0 gateway it is linked to. The tier-1 does not use Edge nodes for routing. All traffic is routed from host transport nodes to the tier-0 gateway. However, to enable cross-location forwarding, the tier-1 allocates two Edge nodes from the Edge cluster configured on the linked tier-0 to use for that traffic.



Tier-1 Gateway with Services or Custom Span


You configure the tier-1 gateway with Edge clusters if you need one of the following configurations:

- You want to run services on the tier-1 gateway.
- You want to deploy a tier-1 gateway that has a different span than the linked tier-0 gateway.

You can remove locations, but you cannot add locations that are not already included the span of the tier-0 gateway.

You select one of the locations to be the primary location. All other locations are secondary. The HA mode for the tier-1 gateway is Active Standby. All traffic passing through this tier-1 gateway passes through the active edge node in the primary location.

If both the tier-1 gateway and the linked tier-0 gateway have primary and secondary locations, configure the same location to be primary for both gateways to reduce cross-location traffic.



Configure Edge Nodes for Stretched Networking

If you want to create gateways and segments that span more than one location, you must configure a remote tunnel endpoint (RTEP) on Edge nodes in each location to handle the cross-location traffic.

When you configure an RTEP, do it on an Edge cluster basis. All Edge nodes in the cluster must have an RTEP configured. You do not need to configure all Edge clusters with RTEP. RTEPs are required only if the Edge cluster is used to configure a gateway that spans more than one location.

You can configure the TEP and RTEP to use the same physical NIC on the Edge node or use separate physical NICs.

You can also configure RTEPs from each Local Manager. Select **System > Get Started > Configure Remote Tunnel Endpoint**.

You can edit RTEPs on an Edge node. Log into the Local Manager and select **System > Fabric > Nodes > Edge Transport Nodes**. Select an Edge node, and click **Tunnels**. If an RTEP is configured, it is displayed in the **Remote Tunnel Endpoint** section. Click **Edit** to modify the RTEP configuration.

Prerequisites

- Verify that each location participating in the stretched network has at least one Edge cluster.
- Determine which layer 3 networks and VLANs to use for RTEP networks.
 - Intra-location tunnel endpoints (TEP) and inter-location tunnel endpoints (RTEP) must use separate VLANs and layer 3 subnets.

- Verify that all RTEP networks used in a given Federation environment have IP connectivity to each other.
- Verify that external firewalls allow cross-location RTEP tunnels, and BGP sessions between Edges. See VMware Ports and Protocols at <https://ports.vmware.com/home/NSX-T-Data-Center>.
- Configure the MTU for RTEP on each Local Manager. The default is 1500. Set the RTEP MTU to be as high as your physical network supports. On each Local Manager, select **System > Fabric > Settings**. Click **Edit** next to **Remote Tunnel Endpoint**.

Procedure

- 1 From your browser, log in with admin privileges to the active Global Manager at <https://<global-manager-ip-address>>.
- 2 Go to **System > Location Manager** and click **Networking** from the location you want to configure for stretched networking.
- 3 Click **Configure** next to the Edge cluster for which you want to set up the RTEP.

The **Configure Edge Nodes for Stretched Networking** screen opens in the Local Manager with that Edge cluster selected.

- 4 You can select all Edge Nodes in this cluster or one node at a time. Provide the following details for the RTEP configuration:

Option	Description
Host Switch	Select a host switch from the drop-down menu.
Teaming Policy	Select a teaming policy if you have one configured.
RTEP VLAN	Enter the VLAN ID for the RTEP network. Valid values are between 1 and 4094.
IP Pool for all nodes	Select an IP pool for all nodes in this Edge Cluster. If you want to assign an IP address to an individual node, you can edit the RTEP configuration later.
Inter Location MTU	The default is 1500.

- 5 Click **Save**.

You can click each of the Edge Nodes that are marked as Configured to see the Edge node configuration details. Select the **Tunnels** tab to view and edit the RTEP configuration.

Add a Tier-0 Gateway from Global Manager

You can add a tier-0 gateway from the Global Manager. This gateway can have a span of one or more locations. This span affects the span of the tier-1 gateways and segments attached to it.

See [Tier-0 Gateway Configurations in NSX Federation](#) for details about tier-0 gateway configurations in NSX Federation.

The following settings must be kept consistent across locations. If you change these settings from the Global Manager web interface, those changes are automatically applied on all locations. However, if you change these settings using the API, you must manually make the same changes in each location.

- Local AS
- ECMP settings
- Multipath Relax settings
- Graceful Restart

Important When you create a tier-0 gateway from Global Manager, you must configure an external interface in each location that the tier-0 is stretched to. Each external interface must be connected to a segment that was created from Global Manager, with the **Connectivity** set to None and the **Traffic type** set to VLAN. See [Add a Segment from Global Manager](#). The Edge nodes configured with those external interfaces are used for inter-location communication, even if northbound communication is not needed.

Prerequisites

- If you are creating a tier-0 gateway that spans more than one location, verify that each location has Edge nodes configured with RTEPs for stretched networking. See [Configure Edge Nodes for Stretched Networking](#).

Procedure

- 1 From your browser, log in with admin privileges to the active Global Manager at `https://<global-manager-ip-address>`.
- 2 Select **Networking > Tier-0 Gateways**.
- 3 Enter a name for the gateway.
- 4 Select an HA (high availability) mode to configure within each location.

The default mode is active-active. In the active-active mode, traffic is load balanced across edge nodes in all locations. In the active-standby mode, an elected Edge node processes traffic in each location. If the active node fails, the standby node becomes active.

Note Active-standby tier-0 gateways are supported starting in NSX-T Data Center 3.0.1.

- 5 If the HA mode is active-standby, select a failover mode.

Option	Description
Preemptive	If the preferred node fails and recovers, it will preempt its peer and become the active node. The peer will change its state to standby.
Non-preemptive	If the preferred node fails and recovers, it will check if its peer is the active node. If so, the preferred node will not preempt its peer and will be the standby node.

- 6 Specify the span of this tier-0 gateway by providing the following details for each location. To add additional locations, click **Add Location**.

Option	Description
Location	Select the location from the drop-down menu.
Edge Cluster	Select an Edge cluster from this location. If you are configuring a stretched tier-0, you must select an Edge cluster that contains Edge nodes that are configured with an RTEP.
Mode	Each location of the tier-0 gateway can have a mode of Primary or Secondary . ■ If the HA mode is Active Active, you can configure the tier-0 gateway with all locations mode set to primary. 1 Select the Mark all locations as Primary toggle to mark all locations as primary. ■ If the HA mode is Active Active or Active Standby, you can configure the tier-0 gateway with one location set to Primary, and all others set to Secondary. 1 Select Primary mode for one location. In all other locations, set mode to Secondary. 2 For secondary locations, you must select a fallback preference.

- 7 Click **Additional Settings**.

- a In the **Internal Transit Subnet** field, enter a subnet.

This is the subnet used for communication between components within this gateway. The default is 169.254.0.0/24.

- b In the **T0-T1 Transit Subnets** field, enter one or more subnets.

These subnets are used for communication between this gateway and all tier-1 gateways that are linked to it. After you create this gateway and link a tier-1 gateway to it, you will see the actual IP address assigned to the link on the tier-0 gateway side and on the tier-1 gateway side. The address is displayed in **Additional Settings > Router Links** on the tier-0 gateway page and the tier-1 gateway page. The default is 100.64.0.0/16.

- c In the **Intersite Transit Subnet** field, enter a subnet. This subnet is used for cross-location communication between gateway components. The default is 169.254.32.0/20.

- 8 Click **Save**.

- 9 To configure interfaces, click **Interfaces** and **Set**. Configure an external interface for each location that the tier-0 gateway spans.

- Click **Add Interface**.
- Enter a name.
- Select a location.

- d Select a type.

If the HA mode is active-standby, the choices are **External**, **Service**, and **Loopback**. If the HA mode is active-active, the choices are **External** and **Loopback**.

Service interfaces are supported only on gateways that span one location. If the gateway is stretched, service interfaces are not supported.

- e Enter an IP address in CIDR format.

- f Select a segment.

The segment must be created from the Global Manager, with the **Connectivity** set to None and the **Traffic type** set to VLAN. See [Add a Segment from Global Manager](#).

- g If the interface type is not **Service**, select an NSX Edge node.

- h (Optional) If the interface type is not **Loopback**, enter an MTU value.

- i Skip **PIM** configuration.

Multicast is not supported in Federation.

- j (Optional) Add tags and select an ND profile.

- k (Optional) If the interface type is **External**, for **URPF Mode**, you can select **Strict** or **None**.

URPF (Unicast Reverse Path Forwarding) is a security feature.

- l After you create an interface, you can download the ARP table by clicking the menu icon (three dots) for the interface and selecting **Download ARP table**.

10 Click **Routing** to add IP prefix lists, community lists, static routes, and route maps.

When you add a static route on a tier-0 gateway, the default behavior is that the static routes are pushed to all locations configured on the gateway. However, the routes are enabled only on the primary locations. This ensures that on the secondary locations, the routes that are learned from the primary location are preferred.

If you want to change this behavior, you can use the **Enabled on Secondary** setting and the **Scope** setting.

If you select **Enabled on Secondary**, the static route is also enabled on the secondary locations.

When you add a next hop for a static route, you can set the **Scope**. The scope can be an interface, a gateway, or a segment. On a tier-0 gateway created from Global Manager, the scope can also be a location. You can use the scope setting to configure different next hops for each location.

11 Click **BGP** to configure BGP.

When you configure BGP on a tier-0 gateway from the Global Manager, most settings apply to all locations.

Some of the settings within the BGP configuration, such as **Route Aggregation** and **BGP Neighbors** prompt you to provide separate values for each location.

See [Configure BGP](#) for more information about configuring BGP.

- 12 To configure route redistribution, click **Route Redistribution**, and for each location, click **Set**.

Select one or more of the sources:

- Tier-0 subnets: **Static Routes, NAT IP, IPsec Local IP, DNS Forwarder IP, EVPN TEP IP, Connected Interfaces & Segments.**

Under **Connected Interfaces & Segments**, you can select one or more of the following: **Service Interface Subnet, External Interface Subnet, Loopback Interface Subnet, Connected Segment.**

- Advertised tier-1 subnets: **DNS Forwarder IP, Static Routes, LB VIP, NAT IP, LB SNAT IP, IPsec Local Endpoint, Connected Interfaces & Segments.**

Under **Connected Interfaces & Segments**, you can select **Service Interface Subnet** and/or **Connected Segment**.

What to do next

Set up a tier-1 gateway from Global Manager.

Add a Tier-1 Gateway from Global Manager

A gateway can be configured in one or more locations. These locations are the span of the gateway. A tier-1 gateway cannot have a greater span than the tier-0 gateway it is connected to.

See [Tier-1 Gateway Configurations in Federation](#) for details about tier-1 gateway configuration options in NSX Federation.

Prerequisites

Verify you have a tier-0 gateway configured.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<global-manager-ip-address>>.
- 2 Select **Networking > Tier-1 Gateways**.
- 3 Click **Add Tier-1 Gateway**.
- 4 Enter a name for the gateway.
- 5 Select a tier-0 gateway to connect to this tier-1 gateway to create a multi-tier topology.
 - If you select a tier-0 gateway, the Locations configuration is populated with the same locations that are configured on the tier-0. If needed, you can modify the locations configuration in the Locations section.
 - If you do not select a tier-0 gateway, you can select locations. However, if you later connect the tier-1 gateway to a tier-0 gateway, you might need to update the locations to create a valid configuration.

- 6 In **Locations**, you can change the **Enable Edge Clusters for Services or Custom Span** setting. It is disabled by default.

- Leave **Enable Edge Clusters for Services or Custom Span** disabled if you want the tier-1 gateway to have the same span as the tier-0 gateway, and you do not need to enable services on the tier-1 gateway. The tier-1 gateway will perform distributed routing only.
- Enable **Enable Edge Clusters for Services or Custom Span** if you want to choose a subset of locations for the tier-1 gateway, or if you want to enable services on the tier-1 gateway.

If you enable **Enable Edge Clusters for Services or Custom Span**, enter the location, cluster, and mode information.

- a Select a location from the drop-down menu. If you linked this tier-1 gateway to a tier-0 gateway, the locations of that tier-0 gateway are automatically listed. If needed, you can delete a location.
- b Select an NSX Edge cluster for each location. If the tier-1 gateway spans more than one location, the Edge clusters must already be configured with an RTEP for each of its Edge Nodes.
- c (Optional) To select specific Edge nodes, click **Set** next the Edge cluster.
Edge nodes are automatically allocated if you do not select Edge nodes.
- d Select a mode for each location. Mode can be Primary or Secondary.

Only one location can be configured with Primary mode. All northbound traffic from this tier-1 gateway is sent through this location.

- 7 If you have enabled Edge clusters, select a failover mode.

Option	Description
Preemptive	If the preferred NSX Edge node fails and recovers, it will preempt its peer and become the active node. The peer will change its state to standby.
Non-preemptive	If the preferred NSX Edge node fails and recovers, it will check if its peer is the active node. If so, the preferred node will not preempt its peer and will be the standby node. This is the default option.

- 8 Skip selecting a size from the **Edge Pool Allocation Size** drop-down menu.
- 9 If you have enabled Edge clusters, select a setting for **Enable StandBy Relocation**.

Standby relocation means that if the Edge node where the active or standby logical router is running fails, a new standby logical router is created on another Edge node to maintain high availability. If the Edge node that fails is running the active logical router, the original standby logical router becomes the active logical router and a new standby logical router is created. If the Edge node that fails is running the standby logical router, the new standby logical router replaces it.

10 (Optional) Click **Route Advertisement.**

Select one or more of the following:

- **All Static Routes**
- **All NAT IP's**
- **All DNS Forwarder Routes**
- **All LB VIP Routes**
- **All Connected Segments and Service Ports**
- **All LB SNAT IP Routes**
- **All IPsec Local Endpoints**

11 Click **Save.****12 (Optional) Click **Route Advertisement**.**

- a In the **Set Route Advertisement Rules** field, click **Set** to add route advertisement rules.

13 (Optional) Click **Additional Settings.**

- a For IPv6, you can select or create an **ND Profile** and a **DAD Profile**.

These profiles are used to configure Stateless Address Autoconfiguration (SLAAC) and Duplicate Address Detection (DAD) for IPv6 addresses.

- b Select an **Ingress QoS Profile** and an **Egress QoS Profile** for traffic limitations.

These profiles are used to set information rate and burst size for permitted traffic. See [Add a Gateway QoS Profile](#) for more information on creating QoS profiles.

If this gateway is linked to a tier-0 gateway, the **Router Links** field shows the link addresses.

14 (Optional) Click **Service Interfaces and **Set** to configure connections to segments. Required in some topologies such as VLAN-backed segments or one-arm load balancing.**

Service interfaces are supported only on gateways that span one location. If the gateway is stretched, service interfaces are not supported.

- a Click **Add Interface**.
- b Enter a name and IP address in CIDR format.
- c Select a segment.
- d In the **MTU** field, enter a value between 64 and 9000.
- e For **URPF Mode**, you can select **Strict** or **None**.

URPF (Unicast Reverse Path Forwarding) is a security feature.

- f Add one or more tags.
- g In the **ND Profile** field, select or create a profile.
- h Click **Save**.

15 (Optional) Click **Static Routes** and **Set** to configure static routes.

- a Click **Add Static Route**.
- b Enter a name and a network address in the CIDR or IPv6 CIDR format.
- c Click **Set Next Hops** to add next hop information.
- d Click **Save**.

Add a Segment from Global Manager

You can add two kinds of segments: overlay-backed segments and VLAN-backed segments. When you create segments from Global Manager, only overlay-backed segments can span multiple locations.

You can view segments ports from Global Manager, but you cannot create or modify them. If you need to create or modify a segment port, you must do it from the Local Manager.

Important Do not change the gateway connectivity of a segment in Federation. Changing the gateway affects the span of the segment. If the span changes in such a way that it excludes a location, the segment is deleted on the excluded location. You must disconnect all VMs before you shrink the span of a segment.

Prerequisites

Verify that each location has a default overlay transport zone configured. The default overlay transport zone is used to create global overlay segments. From each Local Manager, select **System > Fabric > Transport Zones**. Select an overlay transport zone, and click **Actions > Set as Default Transport Zone**.

Procedure

- 1 From your browser, log in with admin privileges to a Global Manager at <https://<global-manager-ip-address>>.
- 2 Select **Networking > Segments**.
- 3 Click **Add Segment**.
- 4 Enter a name for the segment.

5 Select the Connectivity, Traffic Type, and Locations for this segment.

Table 15-3. Segment Configurations

Connectivity	Traffic Type	Location and Transport Zone	Details
A global tier-0 or tier-1 gateway	Overlay	The Location section is populated with the following configurations: ■ the same locations that are configured on the attached gateway. ■ the default overlay transport zone for each location.	Use this configuration to create a global overlay-backed segment connected to the selected global gateway.
None	VLAN	You must select one location for this segment. You must also select a transport zone from that location.	Use this configuration to create a global VLAN-backed segment to use for a tier-0 external interface.
None	Overlay	No locations or transport zones can be selected.	This segment is created on the Global Manager but is not realized in any Local Managers. You can attach it to a gateway later.

Creating a VLAN-backed segment that is attached to a gateway is not supported.

6 Enter the Gateway IP address of the subnet in a CIDR format. A segment can contain an IPv4 subnet, or an IPv6 subnet, or both.

- If a segment is not connected to a gateway, subnet is optional.
- If a segment is connected either to a tier-1 or tier-0 gateway, subnet is required.

Subnets of one segment must not overlap with the subnets of other segments in your network. A segment is always associated with a single virtual network identifier (VNI) regardless of whether it is configured with one subnet, two subnets, or no subnet.

7 Skip **Set DHCP Config**.

Only static bindings are supported on a segment created from Global Manager. See [Features and Configurations Supported in Federation](#).

8 If the transport zone is of type VLAN, specify a list of VLAN IDs. If the transport zone is of type Overlay, and you want to support layer 2 bridging or guest VLAN tagging, specify a list of VLAN IDs or VLAN ranges

9 (Optional) Select an uplink teaming policy for the segment.

This drop-down menu displays the named teaming policies, if you have added them in the VLAN transport zone. If no uplink teaming policy is selected, the default teaming policy is used.

- Named teaming policies are not applicable to overlay segments. Overlay segments always follow the default teaming policy.
- For VLAN-backed segments, you have the flexibility to override the default teaming policy with a selected named teaming policy. This capability is provided so that you can steer the infrastructure traffic from the host to specific VLAN segments in the VLAN transport zone. Before adding the VLAN segment, ensure that the named teaming policy names are added in the VLAN transport zone.

10 Click **Save**.

11 To continue configuring the segment, click **Yes** when prompted.

12 To select segment profiles, click **Segment Profiles**.

13 To bind a static IP address to the MAC address of a VM on the segment, expand **DHCP Static Bindings**, and then click **Set**.

14 Click **Save**.


Security in NSX Federation

You can create distributed and gateway firewall rules from the Global Manager with global, regional or local spans.

Distributed and gateway firewall policies and rules created from the Global Manager are synced to Local Managers and appear in the Local Managers with a  icon. You can edit rules created from the Global Manager only from the Global Manager. They cannot be edited from Local Managers.

Federation of Distributed Firewall (DFW) Policies and Rules

Use this example to understand the supported firewall workflows:




- In the example, the Global Manager has three Local Managers registered with it, named: *Location1*, *Location2* and *Location3*.
- The Global Manager auto-creates the following regions:
 - *Global*
 - *Location1*
 - *Location2*
 - *Location3*
- You create a customized region named: **Region1** that includes Local Managers *Location2* and *Location3*.

- You create the following groups:
 - **Group1:** Region *Global*.
 - **Group2:** Region *Location1*.
 - **Group3:** Region *Location2*.
 - **Group4:** Region *Location3*.
 - **Group5:** Region **Region1**.

DFW Policies and Rules in NSX-T Data Center 3.0.1

The following use cases are supported:

- **Group Span:** You can create groups in the Global Manager with a global, local or regional span. See [Create Groups from Global Manager](#) .
- **Dynamic Groups:** You can create groups based on dynamic criteria, such as tags.
- **DFW Policy Span:** DFW policies can be applied to a global, regional or local span.
- **DFW Rule's Source and Destination Groups:** Either all the groups in the source field or all the groups in the destination field must match the DFW policy's span. The system auto-creates groups in locations that are outside the policy's span.



Refer to the table for examples of valid and invalid source and destination groups in DFW rules:

Table 15-4. Valid Source and Destination for a DFW rule based on the DFW Policy's Span in 3.0.1

DFW Policy Span (Applied To)	Scenarios supported in DFW rules in version 3.0.1.
<i>Global</i> From the example, this region contains the following groups: ■ Group1	For a DFW policy with the span of <i>Global</i> region, all groups are allowed in the DFW rule's source and destination. Following are some typical scenarios that are supported, using our example: ■ Source: Group2; Destination Group3 ■ Source: Group3; Destination Group4 ■ Source: Group4; Destination: Any ■ Source: Group1; Destination Group2.
<i>Location1</i>: auto-created region for the Local Manager in location 1. From the example, this region contains the following groups: ■ Group2	For a DFW policy with the span of one location: <i>Location1</i> in this example, either the source or the destination group for the DFW rule must belong to <i>Location1</i>. The following scenarios are supported: ■ Source: Group2; Destination Group2 ■ Source: Group3; Destination Group2. ■ Source: Group2; Destination Group4. ■ Source Group1; Destination Group2. The following is an example of unsupported group selections for this policy span. Both the source and the destination groups are outside the policy's span: ■ Source Group5; Destination Group3. ■ Source Group1; Destination Group3.
Region1 : user-created region that spans <i>Location2</i> and <i>Location3</i>. From the example, this region contains the following groups: ■ Group5	For a DFW policy with the span of a user-created region: Region1 in this example, either the source or the destination group for the DFW rule must contain locations that belong to Region1. The following scenarios are supported: ■ Source: Group5; Destination Group2. ■ Source: Group2; Destination Group5. ■ Source: Group2; Destination Group3. ■ Source: Group3; Destination Group4. ■ Source: Any; Destination: Group5 ■ Source Group4; Destination Any The following is an example of unsupported group selections for this policy span. Both the source and the destination groups are outside the policy's span: ■ Source Group2; Destination Group2. ■ Source Group1; Destination Group2. ■ Source Group1; Destination Group1.

- If a group contains segments, the span of the DFW policy must be greater than or equal to the span of the segment. For example, if you have a group containing a segment whose span is *Location1*, the DFW policy cannot be applied to region **Region1** because it only contains *Location2* and *Location3*.

DFW Policies and Rules in NSX-T Data Center 3.0.0

- **Group Span:** You can create groups in the Global Manager with a global, local or regional span. See [Create Groups from Global Manager](#) .
- **Dynamic Groups:** You can create groups based on dynamic criteria, such as tags.
- **DFW Policy Span:** DFW policies can be applied to a global, regional, or local span as well.
- **DFW Rule's Source and Destination Groups:** All the groups in the source field and all the groups in the destination field must match the DFW policy's span.

Refer to the table to understand how the span of the policy determines what source and destination groups are valid in a DFW rule.

Table 15-5. Valid Source and Destination for a DFW rule based on the DFW Policy's Span in 3.0.0

DFW Policy Span (Applied To)	Source and Destination Groups supported in DFW Rules in version 3.0.0.
<i>Global.</i> From the example, this region contains the following groups: ■ Group1	For a DFW policy spanned to the <i>Global</i> region, you can select either the keyword <i>Any</i> or a <i>Global</i> group in the source and destination for a DFW rule: For example, ■ Source: Group1; Destination: Group1. ■ Source: Group1; Destination: Any ■ Source: Any ; Destination: Group1. ■ Source: Any ; Destination Any Caution Other rule configurations can be created, but are not supported, for example: . ■ Source: Group2; Destination: Group3 ■ Source: Group4; Destination: Group1
<i>Location1</i>: auto-created region for the Local Manager in location 1. From the example, this region contains the following groups: ■ Group2	For a DFW policy spanned to the region for one location: <i>Location1</i> in this example, both the source and destination groups must belong to this region. For example, these rules are supported: ■ Source: Group2; Destination: Group2 Caution Other rule configurations can be created, but are not supported, for example: . ■ Source: Group2; Destination: Group3 ■ Source: Group4; Destination: Group2
Region1 : customized region that spans <i>Location2</i> and <i>Location3</i>. From the example, this region contains the following groups: ■ Group5	For a DFW policy with a span to a customized region: Region1 in this example, both the source and destination groups must belong to this region. For example, these rule are supported: ■ Source: Group5; Destination: Group5. ■ Source: Group5; Destination: Any. ■ Source: Any ;Destination: Group5. Caution Other rule configurations can be created, but are not supported, for example: . ■ Source: Group2 and Destination: Group3 ■ Source: Group2; Destination: Group4 ■ Source: Group3; Destination: Group2 ■ Source: Group4; Destination: Group2

- If a group contains segments, the span of the DFW policy must be greater than or equal to the span of the segment. For example, if you have a group containing a segment whose span is *Location1*, the DFW policy cannot be applied to region **Region1** because it only contains *Location2* and *Location3*.

Federation of Gateway Firewall Policies and Rules

Gateway firewall rules can be applied to all the locations included in the gateway's span, or all interfaces of a particular location, or specific interfaces of one or more locations.

Note The span of the source and destination groups for gateway firewall rules must be the same as or a subset of the gateway's span on which you are creating the rule.

Table 15-6. Span Options for Gateway Firewall Rules

Gateway Firewall Rule's Span (Applied To)	Applies to
Apply rule to gateway	The rule applies to all interfaces attached to this gateway, in all locations that this gateway is stretched to.
Select a location and then select Apply rule to all Entities.	The rule applies only to the selected location.
Select a location and then select interfaces from that location. Repeat for other locations, selecting interfaces for each location that you want to apply the rule to.	The rule applies only to the selected interfaces.

Create a Region from Global Manager

Each location added to the Global Manager automatically becomes a region. You can also create customized regions.

Use regions to create focused groups for security and networking policies. Some regions are created automatically after you onboard locations in Global Manager. You can add more regions as necessary.

Note Each location can be a part of only one customized region.

The following regions are added by default:

- A Global region including all the locations added to the Global Manager.
- One region for each location added to the Global Manager.

For existing regions, you can view the following information:

- Name of the region.
- Locations included in the region.
- Groups the region belongs to.
- Security/Network policies the region is a part of.

Prerequisites

Refer to [Security in NSX Federation](#) for details on the implication of the span of regions and groups in creating and maintaining security policies and rules.

Procedure

- 1 Select **Inventory > Regions**.

2 Click **Add Region**.

3 Provide the following information:

Option	Description
Name	Provide a name for the region, for example, EMEA, or APAC.
Locations	Select the locations that you want to include in this region.

4 Click **Save**.

The region with the specified locations is created.

What to do next

[Create Groups from Global Manager](#) .


Create Groups from Global Manager

Create Groups from Global Manager that apply globally across your NSX-T Data Center deployments or cover selected locations or regions.

Group Span

When you create a group from the Global Manager, you select a region for the group. The group is synced with all locations in that region. A global region containing all locations, and a region for each location that has been added to the Global Manager are available automatically as regions you can select for a group's span. You can create customized regions before you create groups. See [Create a Region from Global Manager](#).

In this example, **Group1** is created in the Global region, and is therefore synced with all Local Managers.




Dynamic Groups

If a group that spans more than one location has dynamic membership, you need information from each location to list the group membership.

In this example, **Group1** has the following members:

- **VM1** in *Location1*
- **VM2** in *Location2*
- **VM3** in *Location3*

Each Local Manager syncs its dynamic group membership with the other Local Managers. As a result, each Local Manager has a complete list of group members.



Nested Groups

For groups created from the Global Manager, you can add another group as a member if it has an equal or smaller span than the group's region.

Note If you are using NSX-T Data Center version 3.0.0, you can add a group as a member of another group only if the span of both the groups is exactly the same.

Extending the example using **Region1** that contains *Location2* and *Location3*, note the following additional configurations:

Task	Effect
From Global Manager, create Group-Loc2 with region <i>Location2</i> .	■ Group-Loc2 is created in Global Manager. ■ Group-Loc2 is created in the Local Manager <i>Location2</i>.

From Global Manager, create group **Group-Region1** with region **Region1**. Add **Group-Loc2** as a member.
This is a nested group.

- **Group-Region1** is created in Global Manager.
- **Group-Region1** is created in *Location2* and *Location3*.
- **Group-Loc2** is created in Local Manager *Location3*.

From Global Manager, navigate to **Inventory > Regions** and edit **Region1** to remove *Location2*.

This action is not allowed because of the nested group **Group-Region1**.

See [Add a Group](#) for detailed steps for creating groups.

Create DFW Policies and Rules from Global Manager

You can create security policies and DFW rules to apply to multiple locations registered with the Global Manager.

Prerequisites

Ensure that you have already created any customized regions that you want to use for firewall rules. See [Create a Region from Global Manager](#).

Procedure

- 1 From your browser, log in with Enterprise Admin or Security Admin privileges to a Global Manager at <https://<global-manager-ip-address>>.
- 2 Select **Security > Distributed Firewall**
- 3 Ensure that you are in the correct pre-defined category, and click **Add Policy**. For more about categories, see [Distributed Firewall](#).

Note Ethernet, Emergency categories and Default Policy are not supported on Global Manager.

- 4 Click **Add Policy**.
- 5 Enter a **Name** for the new policy section.
- 6 Click the pencil icon next to **Applied To** to set the span of this policy.
- 7 In the **Set Applied To** dialog box, you can make the following selections:
 - **Region:** select which Local Managers to apply the policy to. Each Local Manager is automatically added as a region. You can also create customized regions. See [Create a Region from Global Manager](#).
 - **Select Applied To:** By default, policy is applied to **DFW**, that is, the policy is applied to all the workloads on the Local Managers based on the selected region for this policy. You can also apply a policy to selected groups. Applied to defines the scope of enforcement per policy, and is used mainly for resource optimization on ESXi and KVM hosts. It helps in defining a targeted policy for specific zones, tenants and application without interfering with other policy defined for other tenants, zones & applications.

See [DFW Policies and Rules in NSX-T Data Center 3.0.1](#) to understand how the span of the policy determines whether your DFW rule is valid or invalid.

- 8 To configure the following policy settings, click the gear icon:

Option	Description
TCP Strict	A TCP connection begins with a three-way handshake (SYN, SYN-ACK, ACK) and typically ends with a two-way exchange (FIN, ACK). In certain circumstances, the distributed firewall (DFW) might not see the three-way handshake for a particular flow (due to asymmetric traffic or the distributed firewall being enabled while a flow exists). By default, the distributed firewall does not enforce the need to see a three-way handshake, and picks up sessions that are already established. TCP strict can be enabled on a per section basis to turn off mid-session pick-up and enforce the requirement for a three-way handshake. When enabling TCP strict mode for a particular DFW policy, and using a default ANY-ANY Block rule, packets that do not complete the three-way handshake connection requirements and that match a TCP-based rule in this section are dropped. Strict is only applied to stateful TCP rules, and is enabled at the distributed firewall policy level. TCP strict is not enforced for packets that match a default ANY-ANY Allow which has no TCP service specified.
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. See Role-Based Access Control.

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

The new policy is shown on the screen.

- 10 Select a policy section and click **Add Rule**.

- 11 Enter a name for the rule.

- 12 The Source and Destination are validated based on the DFW policy's span. See [DFW Policies and Rules in NSX-T Data Center 3.0.1](#) for more information.

- If the DFW policy is applied to a location, for example, **Loc1**, source or destination can be either the keyword **ANY** or a group that belongs to **Loc1**.
- If DFW policy is applied to a user-created region, for example, **Region1** source or destination can be either the keyword **ANY** or a group that has the same span as **Region1** or spans a location in **Region1**.

- If DFW policy is applied to **Global**, source or destination can be anything.

Note Active Directory and IDFW are not supported for NSX Federation, that is, you cannot use these features from the Global Manager.

- a In the **Sources** column, click the pencil icon and select the source of the rule.
- b In the **Destinations** column, click the pencil icon and select the destination of the rule. If not defined, the destination matches any.
- 13 In the **Services** column, click the pencil icon and select services. The service matches any if not defined.
- 14 In the **Profiles** column, click the edit icon and select a context profile, or click **Add New Context Profile**. See [Add a Context Profile](#).
- 15 Click **Apply** to apply the context profile to the rule.
- 16 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 a selected group. **Applied to** defines the scope of enforcement per rule, and is used mainly for optimization of resources on ESXi and KVM hosts. It helps in defining a targeted policy for specific zones, tenants, and applications without interfering with other policy defined for other tenants and zones and applications.

Note You cannot select the following types of groups in **Applied-to**:

- a group with IP or MAC addresses
 - an Active Directory user group
-

- 17 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.
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.

- 18 Click the toggle button to enable or disable the rule.

- 19 Click the gear icon to configure the following rule options:

Option	Description
Logging	Logging is turned off by default. Logs are stored at <code>/var/log/dfwpktlogs.log</code> 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 appears in the Firewall Log when logging is enabled.

- 20 Click **Publish**. Multiple rules can be added and then published together at one time.
- 21 On each policy, click **Check Status** to view the status of rules it contains, per location. You can click **Success** or **Failed** to open the policy status window.
- 22 Click **Check Status** to check the realization status of policies that are applied to Transport Nodes on different locations.

Create Gateway Policies and Rules from Global Manager

You can create gateway firewall policies and rules to be applied to multiple locations or selected interfaces for particular locations, from the Global Manager.

Tier-0 or tier-1 gateways created from the Global Manager span all or a set of locations. You have a few options when applying gateway firewall rules created from the Global Manager: Gateway firewall rules can be applied to all the locations included in the gateway's span, or all interfaces of a particular location, or specific interfaces of one or more locations.

On the Local Manager rules are enforced in the following order:

- 1 Any rules you create from the Global Manager, that get successfully realized on the Local Manager, are enforced first.
- 2 Any rules that you create from the Local Manager are enforced next.
- 3 The last rule enforced is the default gateway firewall rule. This is the allow-all or deny-all rule applicable to all locations and all workloads. You can edit the behavior for this default rule from the Global Manager.

Procedure

- 1 From your browser, log in with Enterprise Admin or Security Admin privileges to the Global Manager at `https://<global-manager-ip-address>`.
- 2 Select **Security > Gateway Firewall**.

- 3 Ensure that you are in the correct pre-defined category. Only **Pre Rules**, **Local Gateway** and **Default** categories are supported on Global Manager. To define policy under the **Local Gateway** category, click the category name from the **All Shared Rules** tab or directly click the **Gateway Specific Rules** tab.

Select a tier-0 or tier-1 gateway from the drop-down menu next to **Gateway**. The span of the tier-0 or tier-1 gateway you selected becomes the default span of the Gateway Firewall policy and rule. You can reduce the span but not expand it.

- 4 Click **Add Policy**.
- 5 Enter a **Name** for the new policy section.
- 6 (Optional) Click the gear icon to configure the following policy settings:

Settings	Description
TCP Strict	A TCP connection begins with a three-way handshake (SYN, SYN-ACK, ACK), and typically ends with a two-way exchange (FIN, ACK). In certain circumstances, the firewall may not see the three-way handshake for a particular flow (i.e. due to asymmetric traffic). By default, the firewall does not enforce the need to see a three-way handshake, and will pick up sessions that are already established. TCP strict can be enabled on a per section basis to turn off mid-session pick-up, and enforce the requirement for a three-way handshake. When enabling TCP strict mode for a particular firewall policy and using a default ANY-ANY Block rule, packets that do not complete the three-way handshake connection requirements and that match a TCP-based rule in this policy section are dropped. Strict is only applied to stateful TCP rules, and is enabled at the gateway firewall policy level. TCP strict is not enforced for packets that match a default ANY-ANY Allow which has no TCP service specified.
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 making changes to the same sections. When locking a section, you must include a comment.

- 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**.
- 9 Enter a name for the rule.
- 10 In the **Sources** column, click the edit icon and select the source of the rule. The source group must have the same or a subset of the gateway's span.

- 11 In the **Destinations** column, click the edit icon and select the destination of the rule. If not defined, the destination matches any. The destination group must have the same or a subset of the gateway's span.
- 12 In the **Services** column, click the pencil icon and select services. The service matches any if not defined. Click **Apply** to save.
- 13 In the **Profiles** column, click the edit icon and select a context profile, or click **Add New Context Profile**. See [Add a Context Profile](#).

Note Context profiles are not supported for tier-0 gateways. You can apply L7 context profiles to tier-1 gateways.

- 14 Click the pencil icon in the **Applied to** column. In the **Applied To** dialog box:

<i>Applied To</i> Selection	Result
Select Apply rule to gateway	The gateway firewall rule is applied to all locations covered by the gateway's span. If you add another location to the gateway, this gateway firewall rule automatically gets applied to the location.
Select a location and then select Apply rules to all Entities	Apply this rule to all interfaces in the selected location.
Select a location and then select interfaces for that location	Apply the rule only to selected interfaces in one or more locations.

Note There is no default selection for **Applied To**. You must make a selection to be able to publish this rule.

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

Option	Description
Allow	Allows all 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.
Reject	Rejects packets with the specified source, destination, and protocol. Rejecting a packet 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. The sending application is notified after one attempt that the connection cannot be established.

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

- 17 Click the gear icon to set logging, direction, IP protocol, tag, and notes.

Option	Description
Logging	Logging can be turned off or on. You can access logs using the following NSX CLI command on NSX Edge: <pre>get log-file syslog find datapathd.firewallpkt</pre> Logs can also be sent to an external syslog server.
Direction	The options are In, Out, and In/Out. The default is In/Out. This field 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	The options are IPv4, IPv6, and IPv4_IPv6. The default is IPv4_IPv6.
Log Label	Log label that has been added to the rule.

Note Click the graph icon to view the flow statistics of the firewall rule. You can see information such as the byte, packet count, and sessions.

- 18 Click **Publish**. Multiple rules can be added and then published together at one time.
- 19 Click **Check Status** to view the realization status of policy applied to gateways through edge nodes in different locations. You can click **Success** or **Failed** to open the policy status window.

Backup and Restore in NSX Federation

You can configure and start backups for Global Manager and each Local Manager from within the Global Manager.

Important Starting in NSX-T Data Center 3.0.1, restoring Global Manager to an FQDN is supported. If you are using NSX-T Data Center 3.0.0, do not use FQDN for the Global Manager. Only IP address backups are supported for the Global Manager appliance in NSX-T Data Center 3.0.0.

- Log in to the active Global Manager and select **System > Backup & Restore**. Each Global Manager and Local Manager in the environment is listed. See [Configure Backups](#) for instructions.
- You cannot restore Local Manager from within the Global Manager. To restore a Local Manager backup, log in to the Local Manager to restore. See [Restore a Backup](#) for instructions.
- The system treats backup and restore operations as specific to each appliance, whether it is the Global Manager or the Local Manager you are backing up or restoring. The Global Manager's backup contains a backup of the database of that appliance only. The Local Manager contains a backup of the database and inventory of that appliance only.

- If you are restoring a Global Manager and a Local Manager, select backup timestamps of each appliance as close to each other as possible.
- After each appliance is restored, the `async replicator` service restores communication between the Global Manager and each Local Manager.

Restore Scenarios in Federation

Scenario	Restore Workflow
Global Manager is lost.	Restore the Global Manager. When restored, the Global Manager pushes configurations to the Local Managers registered with it.
A Local Manager is lost.	Restore the Local Manager. When restored, configurations from the Global Manager are synchronized with the Local Manager.
Both the Global Manager and the Local Manager are lost.	If you are restoring both the Global Manager and the Local Manager, use the latest backups of each appliance. When the Global Manager and the Local Manager are restored, the Global Manager pushes the configurations to the Local Manager. You must manually resolve any discrepancies in inventory and fabric related changes between the Local Manager and the Global Manager.

You can monitor the health and performance of the NSX-T Data Center environment.

Read the following topics next:

- [Monitor NSX Edge Nodes](#)
- [Working with Events and Alarms](#)
- [Using vRealize Log Insight for System Monitoring](#)
- [Using vRealize Operations Manager for System Monitoring](#)
- [Using vRealize Network Insight Cloud for System Monitoring](#)

Monitor NSX Edge Nodes

You can monitor an NSX Edge node's usage of resources such as CPU, memory, and storage.

Starting with NSX-T Data Center 3.0.1, the following additional information is displayed:

- Alarms
 - Edge node - Overall alarm count
 - CPU - Alarms for Datapath CPU and Service CPUs
 - Disk - Overall disk alarms and alarms for each partition
 - Memory - Overall memory alarms and alarms for each memory pool
- CPU
 - Datapath CPU - Number of datapath CPU cores and their usage details, which include the average usage of all cores and the highest usage among cores.
 - Service CPU - Number of service CPU cores and their usage details, which include the average usage of all cores and highest usage among cores.
- Disk
 - Total disk usage for all ext4 disk partitions and the list of RAM disk and disk partitions. Also the available free space of each partition.

■ Memory

- Datapath Memory - Includes heap memory, memory pool, and resident memory.
- Memory Pools - List of all memory pools along with their description and usage values except for QAT memory pool (of Bare Metal Edge) whose usage is always around 100%. The memory pools are:

Name	Description
jumbo_mbuf_pool	Packet Pool for Jumbo Frame Used by Ipsec Crypto Device
common_mbuf_pool	Datapath Common Packet Pool
sp_pktmbuf_pool	Datapath Slowpath Packet Pool
fw_mon_msg	Stateful Service Sync Message Pool
vxstt4_frag_q	Vxstt Fragment Pool for Reassembly
pfstatepl3	Stateful Service State Pool
pffqdnippl	Stateful Service FQDN to IP Map Pool
pffqdnsyncpl	Stateful Service FQDN SYNC Pool
pffqdndnpl	Stateful Service FQDN Internal Pool
pfdnsdnpl	Stateful Service FQDN Internal Pool
pfpktpl3	Stateful Service Fragmented Packet Pool
pfsyncmbufpl3	Stateful Service SYNC Pool
pf_fp_rule_node	Stateful Service Rule Node Pool
pf_fp_root_rule_node	Stateful Service Root Rule Node Pool
pf_tb_root_rule_node	Stateful Service Fastpath Root Table Node Pool
pfa_intattr_pl3	Stateful Service Integer Attribute Pool
pfa_attrconn_pl3	Stateful Service Attribute Connection Pool
pfa_ctx_pl3	Stateful Service Context Pool
pfa_key_ace_pl3	Stateful Service Integer Attribute Key Pool
pfa_value_ace_pl3	Stateful Service Integer Attribute Value Pool
lb_pkt_pl3	Load Balancer Temp Packet Cache Pool

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **System > Fabric > Nodes**.
- 3 Click the **Edge Transport Nodes** tab.
- 4 Click the name of an Edge node.

5 Click the **Monitor** tab.

Usage information for CPU, memory, and disk is displayed, as well as the node status, network interfaces, and NAT rule statistics.

Working with Events and Alarms

NSX-T Data Center provides alarms to call your attention to events that can potentially affect performance and system operation. Alarms provide detailed event information such as which component is affected, the type of event, and then recommends a corrective action.

For example, one of the NSX Edge nodes can be experiencing unusually high CPU usage or low available disk space.

Note Alarms are system events with a severity level greater than `LOW`.

If an alarm (for example, Certificate About to Expire) is raised and later a higher-severity alarm (for example, Certificate Expired) is raised about the same issue, the lower-severity alarm is not automatically resolved. You must take the recommended action to resolve the alarm.

Alarm information displays in several locations within the NSX Manager interface.

About Events and Alarms

All alarms are events that have a severity level greater than `LOW`. However, they are treated and reported differently. This section describes those differences.

Event Catalog

The following tables describe the events that trigger alarms, including alarm messages and recommended actions to resolve them. Any event with a severity greater than `LOW` triggers an alarm.

Alarm Management Events

Alarm management events arise from the NSX Manager and Global Manager nodes.

Event Name	Severity	Alert Message	Recommended Action
Alarm Service Overloaded	Critical	The alarm service is overloaded. When event detected: "Due to heavy volume of alarms reported, the alarm service is temporarily overloaded. The NSX UI and GET /api/v1/alarms NSX API have stopped reporting new alarms. Syslog entries and SNMP traps (if enabled) are still being emitted reporting the underlying event details. When the underlying issues causing the heavy volume of alarms are addressed, the alarm service starts reporting new alarms again." When event resolved: "The heavy volume of alarms has subsided and new alarms are being reported again."	Review all active alarms using the Alarms page in the NSX UI or using the GET /api/v1/alarms? status=OPEN,ACKNOWLEDGED,SUPPRESSED NSX API. For each active alarm, investigate the root cause by following the recommended action for the alarm. When sufficient alarms are resolved, the alarm service will start reporting new alarms again.
Heavy Volume of Alarms	Critical	Heavy volume of a specific alarm type detected. When event detected: "Due to heavy volume of {event_id} alarms, the alarm service has temporarily stopped reporting alarms of this type. The NSX UI and GET /api/v1/alarms NSX API are not reporting new instances of these alarms. Syslog entries and SNMP traps (if enabled) are still being emitted reporting the underlying event details. When the underlying issues causing the heavy volume of {event_id} alarms are addressed, the alarm service starts reporting new {event_id} alarms when new issues are detected again." When event resolved: "The heavy volume of {event_id} alarms has subsided and new alarms of this type are being reported again."	Review all active alarms using the Alarms page in the NSX UI or using the GET /api/v1/alarms? status=OPEN,ACKNOWLEDGED,SUPPRESSED NSX API. For each active alarm, investigate the root cause by following the recommended action for the alarm. When sufficient alarms are resolved, the alarm service will start reporting new {event_id} alarms again.

Certificates Events

Certificate events arise from the NSX Manager node.

Event Name	Severity	Alert Message	Recommended Action
Certificate Expired	Critical	A certificate has expired. When event detected: "Certificate <i>{entity-id}</i> has expired." When event resolved: "The expired certificate <i>{entity-id}</i> has been removed or is no longer expired."	Ensure services that are currently using the certificate are updated to use a new, non-expired certificate. For example, to apply a new certificate to the HTTP service, invoke the following API call: <pre>POST /api/v1/node/services/http? action=apply_certificate&certificate_ id=<cert-id></pre> where <i><cert-id></i> is the ID of a valid certificate reported by the API call <code>GET /api/v1/trust-management/certificates</code>. After the expired certificate is no longer in use, it should be deleted with the following API call: <pre>DELETE /api/v1/trust-management/ certificates/{entity_id}</pre>
Certificate About to Expire	High	A certificate is about to expire. When event detected: "Certificate <i>{entity-id}</i> is about to expire." When event resolved: "The expiring certificate <i>{entity-id}</i> or is no longer about to expire."	Ensure services that are currently using the certificate are updated to use a new, non-expiring certificate. For example, to apply a new certificate to the HTTP service, invoke the following API call: <pre>POST /api/v1/node/services/http? action=apply_certificate&certificate_ id=<cert-id></pre> where <i><cert-id></i> is the ID of a valid certificate reported by the API call <code>GET /api/v1/trust-management/certificates</code>. After the expiring certificate is no longer in use, it should be deleted using the API call: <pre>DELETE /api/v1/trust-management/ certificates/{entity_id}</pre>
Certificate Expiration Approaching	Medium	A certificate is approaching expiration. When event detected: "Certificate <i>{entity-id}</i> is approaching expiration." When event resolved: "The expiring certificate <i>{entity-id}</i> or is no longer approaching expiration."	Ensure services that are currently using the certificate are updated to use a new, non-expiring certificate. For example, to apply a new certificate to the HTTP service, invoke the following API call: <pre>POST /api/v1/node/services/http? action=apply_certificate&certificate_ id=<cert-id></pre> where <i><cert-id></i> is the ID of a valid certificate reported by the API call <code>GET /api/v1/trust-management/certificates</code>. After the expiring certificate is no longer in use, it should be deleted using the API call:

Event Name	Severity	Alert Message	Recommended Action
			DELETE /api/v1/trust-management/certificates/{entity_id}

CNI Health Events

CNI health events arise from the ESXi and KVM nodes.

Event Name	Severity	Alert Message	Recommended Action
Hyperbus Manager Connection Down	Medium	Hyperbus cannot communicate with the Manager node. When event detected: "Hyperbus cannot communicate with the Manager node." When event resolved: "Hyperbus can communicate with the Manager node."	The hyperbus vmkernel interface (vmk50) may be missing. See Knowledge Base article 67432 .

DHCP Events

DHCP events arise from the NSX Edge and public gateway nodes.

Event Name	Severity	Alert Message	Recommended Action
Pool Lease Allocation Failed	High	IP addresses in an IP Pool have been exhausted. When event detected: "The addresses in IP Pool <i>{entity_id}</i> of DHCP Server <i>{dhcp_server_id}</i> have been exhausted. The last DHCP request has failed and future requests will fail." When event resolved: "IP Pool <i>{entity_id}</i> of DHCP Server <i>{dhcp_server_id}</i> is no longer exhausted. A lease is successfully allocated to the last DHCP request."	Review the DHCP pool configuration in the NSX UI or on the Edge node where the DHCP server is running by invoking the NSX CLI command <code>get dhcp ip-pool</code>. Also review the current active leases on the Edge node by invoking the NSX CLI command <code>get dhcp lease</code>. Compare the leases to the number of active VMs. Consider reducing the lease time on the DHCP server configuration if the number of VMs are low compared to the number of active leases. Also consider expanding the pool range for the DHCP server by visiting the Networking > Segments > Segment page in the NSX UI.
Pool Overloaded	Medium	An IP Pool is overloaded. When event detected: "DHCP Server <i>{dhcp_server_id}</i> IP Pool <i>{entity_id}</i> usage is approaching exhaustion with <i>{dhcp_pool_usage}%</i> IPs allocated." When event resolved: "The DHCP Server <i>{dhcp_server_id}</i> IP Pool <i>{entity_id}</i> has fallen below the high usage threshold."	Review the DHCP pool configuration in the NSX UI or on the Edge node where the DHCP server is running by invoking the NSX CLI command <code>get dhcp ip-pool</code>. Also review the current active leases on the Edge node by invoking the NSX CLI command <code>get dhcp lease</code>. Compare the leases to the number of active VMs. Consider reducing the lease time on the DHCP server configuration if the number of VMs are low compared to the number of active leases. Also consider expanding the pool range for the DHCP server by visiting the Networking > Segments > Segment page in the NSX UI.

Distributed Firewall Events

Distributed firewall events arise from the NSX Manager or ESXi nodes.

Event Name	Severity	Alert Message	Recommended Action
Distributed Firewall CPU Usage Very High	Critical	Distributed firewall CPU usage is very high. When event detected: "The DFW CPU usage on Transport node <i>{entity_id}</i> has reached <i>{system_resource_usage}</i> % which is at or above the very high threshold value of <i>{system_usage_threshold}</i> %." When event resolved: "DNS forwarder <i>{entity_id}</i> is running again."	Consider re-balancing the VM workloads on this host to other hosts. Please review the security design for optimization. For example, use the apply-to configuration if the rules are not applicable to the entire datacenter.
Distributed Firewall Memory Usage Very High	Critical	Distributed firewall memory usage is very high. When event detected: "The DFW memory usage <i>{heap_type}</i> on Transport Node <i>{entity_id}</i> has reached <i>{system_resource_usage}</i> % which is at or above the very high threshold value of <i>{system_usage_threshold}</i> %." When event resolved: "The DFW memory usage <i>{heap_type}</i> on Transport Node <i>{entity_id}</i> has reached <i>{system_resource_usage}</i> % which is below the very high threshold value of <i>{system_usage_threshold}</i> %."	View the current DFW memory usage by invoking the NSX CLI command <code>get firewall thresholds</code> on the host. Consider re-balancing the workloads on this host to other hosts.

DNS Events

DNS events arise from the NSX Edge and public gateway nodes.

Event Name	Severity	Alert Message	Recommended Action
Forwarder Down	High	A DNS forwarder is down. When event detected: "DNS forwarder <i>{entity_id}</i> is not running. This is impacting all configured DNS Forwarders that are currently enabled." When event resolved: "DNS forwarder <i>{entity_id}</i> is running again."	1 Invoke the NSX CLI command <code>get dns-forwarders status</code> to verify if the DNS forwarder is in down state. 2 Check <code>/var/log/syslog</code> to see if there are errors reported. 3 Collect a support bundle and contact the NSX support team.
Forwarder Disabled	High	A DNS forwarder is disabled. When event detected: "DNS forwarder <i>{entity_id}</i> is disabled." When event resolved: ""DNS forwarder <i>{entity_id}</i> is enabled."	1 Invoke the NSX CLI command <code>get dns-forwarders status</code> to verify if the DNS forwarder is in a disabled state. 2 Use NSX Policy API or Manager API to enable the DNS forwarder it should not be in the disabled state.

Edge Health Events

Edge health events arise from the NSX Edge and public gateway nodes.

Event Name	Severity	Alert Message	Recommended Action
Edge CPU Usage Very High	Critical	Edge node CPU usage is very high. When event detected: "The CPU usage on the Edge node <i>{entity-id}</i> has reached <i>{system_resource_usage}</i> %, which is at or above the very high threshold value of <i>{system_usage_threshold}</i>%." When event resolved: "The CPU usage on the Edge node <i>{entity-id}</i> has reached <i>{system_resource_usage}</i>%, which is below the very high threshold value of <i>{system_usage_threshold}</i>%."	Please review the configuration, running services and sizing of this Edge node. Consider adjusting the Edge appliance form factor size or rebalancing services to other Edge nodes for the applicable workload.
Edge CPU Usage High	Medium	Edge node CPU usage is high. When event detected: "The CPU usage on the Edge node <i>{entity-id}</i> has reached <i>{system_resource_usage}</i>%, which is at or above the high threshold value of <i>{system_usage_threshold}</i>%." When event resolved: "The CPU usage on the Edge node <i>{entity-id}</i> has reached <i>{system_resource_usage}</i> %, which is below the high threshold value of <i>{system_usage_threshold}</i>%."	Please review the configuration, running services and sizing of this Edge node. Consider adjusting the Edge appliance form factor size or rebalancing services to other Edge nodes for the applicable workload.
Edge Datapath Configuration Failure	High	Edge node datapath configuration has failed. When event detected: "Failed to enable the datapath on the Edge node after three attempts." When event resolved: "Datapath on the Edge node has been successfully enabled."	Ensure the Edge node connection to the Manager node is healthy. From the Edge node NSX CLI, invoke the command <code>get services</code> to check the health of services. If the dataplane service is stopped, invoke the command <code>start service dataplane</code> to restart it.
Edge Datapath CPU Usage Very High	Critical	Edge node datapath CPU usage is very high. When event detected: "The datapath CPU usage on Edge node <i>{entity-id}</i> has reached <i>{datapath_resource_usage}</i>% which is at or above the very high threshold for at least two minutes." When event resolved: "Datapath CPU usage on Edge node <i>{entity-id}</i> has reduced below the maximum threshold."	Review the CPU statistics on the Edge node by invoking the NSX CLI command <code>get dataplane cpu stats</code> to show packet rates per CPU core. Higher CPU usage is expected with higher packet rates. Consider increasing the Edge appliance form factor size and rebalancing services on this Edge node to other Edge nodes in the same cluster or other Edge clusters.

Event Name	Severity	Alert Message	Recommended Action
Edge Datapath CPU Usage High	Medium	Edge node datapath CPU usage is high. When event detected: "The datapath CPU usage on Edge node <i>{entity-id}</i> has reached <i>{datapath_resource_usage}%</i> which is at or above the high threshold for at least two minutes." When event resolved: "The CPU usage on Edge node <i>{entity-id}</i> has reached below the high threshold."	Review the CPU statistics on the Edge node by invoking the NSX CLI command <code>get dataplane cpu stats</code> to show packet rates per CPU core. Higher CPU usage is expected with higher packet rates. Consider increasing the Edge appliance form factor size and rebalancing services on this Edge node to other Edge nodes in the same cluster or other Edge clusters.
Edge Datapath Crypto Driver Down	Critical	The Edge node datapath crypto driver is down. When event detected: "Edge node crypto driver is down." When event resolved: "Edge node crypto driver is up."	Upgrade the Edge node as needed.
Edge Datapath Memory Pool is High	Medium	The Edge node datapath memory pool is high. When event detected: "The datapath mempool usage for <i>{mempool_name}</i> on Edge node <i>{entity-id}</i> has reached <i>{system_resource_usage}%</i> which is at or above the high threshold value of <i>{system_usage_threshold}%</i>." When event resolved: "The datapath mempool usage for <i>{mempool_name}</i> on Edge node <i>{entity-id}</i> has reached <i>{system_resource_usage}%</i> which is below the high threshold value of <i>{system_usage_threshold}%</i>."	Log in as the root user and invoke the commands <code>edge-appctl -t /var/run/vmware/edge/dpdctl mempool/show</code> and <code>edge-appctl -t /var/run/vmware/edge/dpdctl memory/show malloc_heap</code> to check DPDK memory usage.
Edge Disk Usage Very High	Critical	Edge node disk usage is very high. When event detected: "The disk usage for the Edge node disk partition <i>{disk_partition_name}</i> has reached <i>{system_resource_usage}%</i>, which is at or above the very high threshold value of <i>{system_usage_threshold}%</i>." When event resolved: "The disk usage for the Edge node disk partition <i>{disk_partition_name}</i> has reached <i>{system_resource_usage}%</i>, which is below the very high threshold value of <i>{system_usage_threshold}%</i>."	Examine the partition with high usage and see if there are any unexpected large files that can be removed.

Event Name	Severity	Alert Message	Recommended Action
Edge Disk Usage High	Medium	Edge node disk usage is high. When event detected: "The disk usage for the Edge node disk partition <i>{disk_partition_name}</i> has reached <i>{system_resource_usage}%</i>, which is at or above the high threshold value of <i>{system_usage_threshold}%</i>." When event resolved: "The disk usage for the Edge node disk partition <i>{disk_partition_name}</i> has reached <i>{system_resource_usage}%</i>, which is below the high threshold value of <i>{system_usage_threshold}%</i>."	Examine the partition with high usage and see if there are any unexpected large files that can be removed.
Edge Global ARP Table Usage High	Medium	The Edge node global ARP table usage is high. When event detected: "Global ARP table usage on Edge node <i>{entity-id}</i> has reached <i>{datapath_resource_usage}%</i> which is above the high threshold for over two minutes." When event resolved: "Global arp table usage on Edge node <i>{entity-id}</i> has reached below the high threshold."	Increase the ARP table size: 1 Log in as the root user. 2 Invoke the command <code>edge-appctl -t /var/run/vmware/edge/dpdctl neigh/show</code>. 3 Check if neigh cache usage is normal. - a If it is normal, invoke the command <code>edge-appctl -t /var/run/vmware/edge/dpdctl neigh/set_param max_entries</code> to increase the ARP table size.
Edge Memory Usage Very High	Critical	Edge node memory usage is very high. When event detected: "The memory usage on the Edge node <i>{entity-id}</i> has reached <i>{system_resource_usage}%</i>, which is at or above the very high threshold value of <i>{system_usage_threshold}%</i>." When event resolved: "The memory usage on the Edge node <i>{entity-id}</i> has reached <i>{system_resource_usage}%</i>, which is below the very high threshold value of <i>{system_usage_threshold}%</i>."	Please review the configuration, running services and sizing of this Edge node. Consider adjusting the Edge appliance form factor size or rebalancing services to other Edge nodes for the applicable workload.
Edge Memory Usage High	Medium	Edge node memory usage is high. When event detected: "The memory usage on the Edge node <i>{entity-id}</i> has reached <i>{system_resource_usage}%</i>, which is at or above the high threshold value of <i>{system_usage_threshold}%</i>." When event resolved: "The memory usage on the Edge node <i>{entity-id}</i> has reached <i>{system_resource_usage}%</i>, which is below the high threshold value of <i>{system_usage_threshold}%</i>."	Please review the configuration, running services and sizing of this Edge node. Consider adjusting the Edge appliance form factor size or rebalancing services to other Edge nodes for the applicable workload.

Event Name	Severity	Alert Message	Recommended Action
Edge NIC Link Status Down	Critical	Edge node NIC link is down. When event detected: "Edge node NIC {edge_nic_name} link is down." When event detected: "Edge node NIC {edge_nic_name} link is up."	On the Edge node, confirm if the NIC link is physically down by invoking the NSX CLI command <code>get interfaces</code> . If it is down, verify the cable connection.
Edge NIC Out of Receive Buffer	Critical	Edge node NIC receive descriptor ring buffer has no space left. When event detected: "Edge NIC {edge_nic_name} receive ring buffer has overflowed by {rx_ring_buffer_overflow_percentage}% on Edge node {entity-id} for over 60 seconds." When event resolved: "Edge NIC {edge_nic_name} receive ring buffer usage on Edge node {entity-id} is no longer overflowing."	Invoke the NSX CLI command <code>get dataplane</code> , and check the following: 1 If PPS and CPU usage is high and check rx ring size by invoking using <code>get dataplane find ring-size rx</code> . ■ If PPS and CPU is high and rx ring size is low, invoke <code>set dataplane ring-size rx <ring-size></code> , and set <code><ring-size></code> to a higher value to accommodate incoming packets. ■ If the above condition is not satisfied, and ring size is high and CPU usage is still high, the cause may be due to a dataplane processing overhead delay.
Edge NIC Out of Transmit Buffer	Critical	Edge node NIC transmit descriptor ring buffer has no space left. When event detected: "Edge node NIC {edge_nic_name} transmit ring buffer has overflowed by {tx_ring_buffer_overflow_percentage}% on Edge node {entity-id} for over 60 seconds." When event resolved: "Edge node NIC {edge_nic_name} transmit ring buffer usage on Edge node {entity-id} is no longer overflowing."	Invoke the NSX CLI command <code>get dataplane</code> , and check the following: 1 If PPS and CPU usage is high and check rx ring size by invoking using <code>get dataplane find ring-size tx</code> . ■ If If PPS and CPU is high and tx ring size is low, invoke <code>set dataplane ring-size tx <ring-size></code> , and set <code><ring-size></code> to a higher value to accommodate outgoing packets. ■ If the above condition is not satisfied, and ring size is high and CPU usage is low or nominal, the cause may be due to the transmit ring size setting on the hypervisor.
Storage Error	Critical	Starting in NSX-T Data Center 3.0.1. The following disk partitions on the Edge node are in read-only mode: {disk_partition_name} .	Examine the read-only partition to see if reboot resolves the issue, or the disk needs to be replaced. Refer to KB article https://kb.vmware.com/s/article/2146870 .

Endpoint Protection Events

Endpoint protection events arise from the NSX Manager or ESXi nodes.

Event Name	Severity	Alert Message	Recommended Action
EAM Status Down	Critical	ESX Agent Manager (EAM) service on a compute manager is down. When event detected: "ESX Agent Manager (EAM) service on compute manager <i>{entity_id}</i> is down." When event resolved: "ESX Agent Manager (EAM) service on compute manager <i>{entity_id}</i> is either up or compute manager <i>{entity_id}</i> has been removed."	Restart the ESX Agent Manager (EAM) service: ■ SSH into the vCenter node and run: <pre>service vmware-eam start</pre>
Partner Channel Down	Critical	Host module and Partner SVM connection is down. When event detected: "The connection between host module and Partner SVM <i>{entity_id}</i> is down." When event resolved: "The connection between host module and Partner SVM <i>{entity_id}</i> is up."	See Knowledge Base article 2148821 Troubleshooting NSX Guest Introspection and make sure that the Partner SVM identified by <i>{entity_id}</i> is reconnected to the host module.

Federation Events

Federation events arise from the NSX Manager, NSX Edge, and the public gateway nodes.

Event Name	Severity	Alert Message	Recommended Action
LM to LM Synchronization Error	High	Starting in NSX-T Data Center 3.0.1. The synchronization between <i>{site_name}</i> <i>{{site_id}}</i> and <i>{remote_site_name}</i> <i>{{remote_site_id}}</i> failed for more than 5 minutes.	1 Invoke the NSX CLI command <code>get site-replicator remote-sites</code> to get the connection state between remote locations. If a remote location is connected but not synchronized, it is possible that the location is still in the process of master resolution. In this case, wait for approximately 10 seconds and try invoking the CLI again to check for the state of the remote location. If a location is disconnected, try the next step. 2 Check the connectivity from Local Manager (LM) in location <i>{site_name}</i><i>{site_id}</i> to the LMs in location <i>{remote_site_name}</i> <i>{remote_site_id}</i> via ping. If they are not pingable, check for flakiness in WAN connectivity. If there are no physical network connectivity issues, try the next step. 3 Check the <code>/var/log/cloudnet/nsx-ccp.log</code> file on the Manager nodes in the local cluster in location <i>{site_name}</i> <i>{{site_id}}</i> that triggered the alarm to see if there are any cross-site communication errors. In addition, also look for errors being logged by the <code>nsx-appl-proxy</code> subcomponent within <code>/var/log/syslog</code>.
LM to LM Synchronization Warning	Medium	Starting in NSX-T Data Center 3.0.1.	1 Invoke the NSX CLI command <code>get site-replicator remote-sites</code> to get

Event Name	Severity	Alert Message	Recommended Action
		The synchronization between <i>{site_name}</i> (<i>{site_id}</i>) and <i>{remote_site_name}</i> (<i>{remote_site_id}</i>) failed.	the connection state between remote locations. If a remote location is connected but not synchronized, it is possible that the location is still in the process of master resolution. In this case, wait for approximately 10 seconds and try invoking the CLI again to check for the state of the remote location. If a location is disconnected, try the next step. 2 Check the connectivity from Local Manager (LM) in location <i>{site_name}</i>/<i>{site_id}</i> to the LMs in location <i>{remote_site_name}</i> (<i>{remote_site_id}</i>) via ping. If they are not pingable, check for flakiness in WAN connectivity. If there are no physical network connectivity issues, try the next step. 3 Check the <code>/var/log/cloudnet/nsx-ccp.log</code> file on the Manager nodes in the local cluster in location <i>{site_name}</i> (<i>{site_id}</i>) that triggered the alarm to see if there are any cross-site communication errors. In addition, also look for errors being logged by the nsx-appl-proxy subcomponent within <code>/var/log/syslog</code>.
RTEP BGP Down	High	Starting in NSX-T Data Center 3.0.1.	1 Invoke the NSX CLI command <code>get logical-routers</code> on the affected edge node. 2 Switch to <code>REMOTE_TUNNEL_VRF</code> context

Event Name	Severity	Alert Message	Recommended Action
		RTEP BGP session from source IP <i>{bgp_source_ip}</i> to remote location <i>{remote_site_name}</i> neighbor IP <i>{bgp_neighbor_ip}</i> is down. Reason: <i>{failure_reason}</i> .	3 Invoke the NSX CLI command <code>get bgp neighbor</code> to check the BGP neighbor. 4 Alternatively, invoke the NSX API <code>GET /api/v1/transport-nodes/<transport-node-id>/inter-site/bgp/summary</code> to get the BGP neighbor status. 5 Invoke the NSX CLI command <code>get interfaces</code> and check if the correct RTEP IP address is assigned to the interface with name <code>remote-tunnel-endpoint</code>. 6 . Check if the ping is working successfully between assigned RTEP IP address <i>{bgp_source_ip}</i> and the remote location <i>{remote_site_name}</i> neighbor IP <i>{bgp_neighbor_ip}</i>. 7 Check <code>/var/log/syslog</code> for any errors related to BGP. 8 Invoke the API <code>GET</code> or <code>PUT /api/v1/transport-nodes/<transport-node-id></code> to get/update <code>remote_tunnel_endpoint</code> configuration on the edge node. This will update the RTEP IP assigned to the affected edge node.

High Availability Events

High availability events arise from the NSX Edge and public cloud gateway nodes.

Event Name	Severity	Alert Message	Recommended Action
Tier0 Gateway Failover	High	A tier0 gateway has failed over. When event detected: "The tier0 gateway <i>{entity-id}</i> failover from <i>{previous_gateway_state}</i> to <i>{current_gateway_state}</i> ." When event resolved: "The tier0 gateway <i>{entity-id}</i> is now up."	Determine the service that is down and restart it. 1 Identify the tier0 VRF ID by running the NSX CLI command <code>get logical-routers</code> . 2 Switch to the VRF context by running <code>vrf <vrf-id></code> . 3 View which service is down by running <code>get high-availability status</code> .
Tier1 Gateway Failover	High	A tier1 gateway has failed over. When event detected: "The tier1 gateway <i>{entity-id}</i> failover from <i>{previous_gateway_state}</i> to <i>{current_gateway_state}</i> ." When event resolved: "The tier1 gateway <i>{entity-id}</i> is now up."	Determine the service that is down and restart it. 1 Identify the tier1 VRF ID by running the NSX CLI command <code>get logical-routers</code> . 2 Switch to the VRF context by running <code>vrf <vrf-id></code> . 3 View which service is down by running <code>get high-availability status</code> .

Infrastructure Communication Events

Infrastructure communication events arise from the NSX Edge, KVM, ESXi, and public gateway nodes.

Event Name	Severity	Alert Message	Recommended Action
Edge Tunnels Down	Critical	An Edge node's tunnel status is down. When event detected: "Overall tunnel status of Edge node <i>{entity_id}</i> is down." When event resolved: "The tunnels of Edge node <i>{entity_id}</i> have been restored."	1 Using SSH, log into the Edge node. 2 Get the status. <pre>nsxcli get tunnel-ports</pre> 3 On each tunnel, check the stats for any drops. <pre>get tunnel-port <UUID> stats</pre> 4 Check the <code>syslog</code> file for any tunnel related errors.

Infrastructure Service Events

Infrastructure service events arise from the NSX Edge and public gateway nodes.

Event Name	Severity	Alert Message	Recommended Action
Edge Service Status Down	Critical	Edge service is down for at least one minute. When event detected: "The service <i>{edge_service_name}</i> is down for at least one minute." When event resolved: "The service <i>{edge_service_name}</i> is up."	On the Edge node, verify the service hasn't exited due to an error by looking for core dump files in the <code>/var/log/core</code> directory. To confirm whether the service is stopped, invoke the NSX CLI command <code>get services</code>. If so, run <code>start service <service-name></code> to restart the service.
Edge Service Status Changed	Low	Edge service status has changed. When event detected: "The service <i>{edge_service_name}</i> changed from <i>{previous_service_state}</i> to <i>{current_service_state}</i>." When event resolved: "The service <i>{edge_service_name}</i> changed from <i>{previous_service_state}</i> to <i>{current_service_state}</i>."	On the Edge node, verify the service hasn't exited due to an error by looking for core dump files in the <code>/var/log/core</code> directory. To confirm whether the service is stopped, invoke the NSX CLI command <code>get services</code>. If so, run <code>start service <service-name></code> to restart the service.

Intelligence Communication Events

NSX Intelligence communication events arise from the NSX Manager node, ESXi node, and NSX Intelligence appliance.

Event Name	Severity	Alert Message	Recommended Action
Transport node flow exporter disconnected	High	A Transport node is disconnected from its Intelligence node's messaging broker. Data collection is affected. When event detected: "The flow exporter on Transport node <i>{entity-id}</i> is disconnected from the Intelligence node's messaging broker. Data collection is affected." When event resolved: "The flow exporter on Transport node <i>{entity-id}</i> has reconnected to the Intelligence node's messaging broker."	1 Restart messaging service if it is not running in the NSX Intelligence node. 2 Resolve the network connection failure between the transport node and the NSX Intelligence node.
Control Channel to Transport Node Down	Critical	Control channel to Transport Node Down. When event detected: Controller service <i>central_control_plane_id</i> to Transport node <i>{entity-id}</i> down for atleast three minutes from Controller services point of view. When event resolved: Controller service <i>central_control_plane_id</i> restores connection to Transport node <i>{entity-id}</i>.	1 Check the connectivity from the Controller service <i>central_control_plane_id</i> and Transport node <i>{entity-id}</i> interface by using the ping command. If they are not pingable, check for flakiness in network connectivity. 2 Check to see if the TCP connections are established using the netstat output to see if the Controller service <i>{central_control_plane_id}</i> is listening for connections on port 1235. If not, check firewall (or) iptables rules to see if port 1235 is blocking Transport node <i>{entity-id}</i> connection requests. Ensure that there are no host firewalls or network firewalls in the underlay are blocking the required IP ports between Manager nodes and Transport nodes. This is documented in our ports and protocols tool which is here: https://ports.vmware.com/. 3 It is possible that the Transport node <i>{entity-id}</i> may still be in maintenance mode. You can check whether the Transport node is in maintenance mode via the following API: <pre>GET https://<nsx-mgr>/api/v1/transport-nodes/<tn-uuid></pre> When maintenance mode is set, the Transport node will not be connected

Event Name	Severity	Alert Message	Recommended Action
			to the Controller service. This is usually the case when host upgrade is in progress. Wait for a few minutes and check connectivity again. Note This alarm is not critical and should be resolved. GSS need not be contacted for the notification of this alarm unless the alarm remains unresolved over an extended period of time.
Control Channel to Transport Node Down for too long	Warning	Control channel to Transport Node Down for too long. When event detected: Controller service <i>central_control_plane_id</i> to Transport node <i>{entity-id}</i> down for atleast 15 minutes from Controller services point of view. When event resolved: Controller service <i>central_control_plane_id</i> restores connection to Transport node <i>{entity-id}</i>.	1 Check the connectivity from the Controller service <i>central_control_plane_id</i> and Transport node <i>{entity-id}</i> interface by using the ping command. If they are not pingable, check for flakiness in network connectivity. 2 Check to see if the TCP connections are established using the netstat output to see if the Controller service <i>{central_control_plane_id}</i> is listening for connections on port 1235. If not, check firewall (or) iptables rules to see if port 1235 is blocking Transport node <i>{entity-id}</i> connection requests. Ensure that there are no host firewalls or network firewalls in the underlay are blocking the required IP ports between Manager nodes and Transport nodes. This is documented in our ports and protocols tool which is here: https://ports.vmware.com/. 3 It is possible that the Transport node <i>{entity-id}</i> may still be in maintenance mode. You can check whether the Transport node is in maintenance mode via the following API: <pre>GET https://<nsx-mgr>/api/v1/transport-nodes/<tn-uuid></pre> When maintenance mode is set, the Transport node will not be connected to the Controller service. This is usually the case when host upgrade is in progress. Wait for a few minutes and check connectivity again.

Event Name	Severity	Alert Message	Recommended Action
Management Channel To Transport Node Down	Critical	Disconnection from Manager node to transport node. When event detected: When event resolved	1 Ensure that there is a network connectivity between Manager nodes and the Transport node <i>nodename (IP)</i> and there are no firewalls blocking the traffic between these nodes. 2 Ensure the nsx-proxy service is running on the Transport node by invoking the following command. <pre>/etc/init.d/nsx-prxy status</pre> If the nsx-proxy service is not running, restart it by invoking the following command. <pre>/etc/init.d/nsx-proxy restart</pre>
Manager Control Channel Down	Critical	Manager to controller channel is down. When event detected: When event resolved:	On the Manager node <i>managernode (IP)</i>, invoke the following two NSX CLI commands: ■ <code>restart service mgmt-plane-bus</code> ■ <code>restart service manage</code>

Intelligence Health Events

NSX Intelligence health events arise from the NSX Manager node and NSX Intelligence appliance.

Event Name	Severity	Alert Message	Recommended Action
CPU Usage Very High	Critical	Intelligence node CPU usage is very high. When event detected: "The CPU usage on NSX Intelligence node <i>{intelligence_node_id}</i> is above the very high threshold value of <i>{system_usage_threshold}%</i> ." When event resolved: "The CPU usage on NSX Intelligence node <i>{intelligence_node_id}</i> is below the very high threshold value of <i>{system_usage_threshold}%</i> ."	Use the <code>top</code> command to check which processes have the most memory usages, and then check <code>/var/log/syslog</code> and these processes' local logs to see if there are any outstanding errors to be resolved.
CPU Usage High	Medium	Intelligence node CPU usage is high. When event detected: "The CPU usage on NSX Intelligence node <i>{intelligence_node_id}</i> is above the high threshold value of <i>{system_usage_threshold}%</i> ." When event resolved: "The CPU usage on NSX Intelligence node <i>{intelligence_node_id}</i> is below the high threshold value of <i>{system_usage_threshold}%</i> ."	Use the <code>top</code> command to check which processes have the most memory usages, and then check <code>/var/log/syslog</code> and these processes' local logs to see if there are any outstanding errors to be resolved.

Event Name	Severity	Alert Message	Recommended Action
Memory Usage Very High	Critical	Intelligence node memory usage is very high. When event detected: "The memory usage on NSX Intelligence node <i>{intelligence_node_id}</i> is above the very high threshold value of <i>{system_usage_threshold}%</i>." When event resolved: "The memory usage on NSX Intelligence node <i>{intelligence_node_id}</i> is below the very high threshold value of <i>{system_usage_threshold}%</i>."	Use the <code>top</code> command to check which processes have the most memory usages, and then check <code>/var/log/syslog</code> and these processes' local logs to see if there are any outstanding errors to be resolved.
Memory Usage High	Medium	Intelligence node memory usage is high. When event detected: "The memory usage on NSX Intelligence node <i>{intelligence_node_id}</i> is above the high threshold value of <i>{system_usage_threshold}%</i>." When event resolved: "The memory usage on NSX Intelligence node <i>{intelligence_node_id}</i> is below the high threshold value of <i>{system_usage_threshold}%</i>."	Use the <code>top</code> command to check which processes have the most memory usages, and then check <code>/var/log/syslog</code> and these processes' local logs to see if there are any outstanding errors to be resolved.
Disk Usage Very High	Critical	Intelligence node disk usage is very high. When event detected: "The disk usage of disk partition <i>{disk_partition_name}</i> on the NSX Intelligence node <i>{intelligence_node_id}</i> is above the very high threshold value of <i>{system_usage_threshold}%</i>." When event resolved: "The disk usage of disk partition <i>{disk_partition_name}</i> on the NSX Intelligence node <i>{intelligence_node_id}</i> is below the very high threshold value of <i>{system_usage_threshold}%</i>."	Examine disk partition <i>{disk_partition_name}</i> and see if there are any unexpected large files that can be removed.
Disk Usage High	Medium	Intelligence node disk usage is high. When event detected: "The disk usage of disk partition <i>{disk_partition_name}</i> on the NSX Intelligence node <i>{intelligence_node_id}</i> is above the high threshold value of <i>{system_usage_threshold}%</i>." When event resolved: "The disk usage of disk partition <i>{disk_partition_name}</i> on the NSX Intelligence node <i>{intelligence_node_id}</i> is below the high threshold value of <i>{system_usage_threshold}%</i>."	Examine disk partition <i>{disk_partition_name}</i> and see if there are any unexpected large files that can be removed.

Event Name	Severity	Alert Message	Recommended Action
Data disk partition usage very high	Critical	Intelligence node data disk partition usage is very high. When event detected: "The disk usage of disk partition /data on NSX Intelligence node <i>{intelligence_node_id}</i> is above the very high threshold value of <i>{system_usage_threshold}%</i>." When event resolved: "The disk usage of disk partition /data on NSX Intelligence node <i>{intelligence_node_id}</i> is below the very high threshold value of <i>{system_usage_threshold}%</i>."	Stop NSX Intelligence data collection until the disk usage is below the threshold. In the NSX UI, navigate to System Appliances NSX Intelligence Appliance. Then select ACTIONS > Stop Collecting Data.
Data disk partition usage high	Medium	Intelligence node data disk partition usage is high. When event detected: "The disk usage of disk partition /data on NSX Intelligence node <i>{intelligence_node_id}</i> is above the high threshold value of <i>{system_usage_threshold}%</i>." When event resolved: "The disk usage of disk partition /data on NSX Intelligence node <i>{intelligence_node_id}</i> is below the high threshold value of <i>{system_usage_threshold}%</i>."	Stop NSX Intelligence data collection until the disk usage is below the threshold. Examine the /data partition and see if there are any unexpected large files that can be removed.
Node status degraded	High	Intelligence node status is degraded. When event detected: "Service <i>{service_name}</i> on NSX Intelligence node <i>{intelligence_node_id}</i> is not running." When event resolved: "Service <i>{service_name}</i> on NSX Intelligence node <i>{intelligence_node_id}</i> is running properly."	Examine service status and health information with NSX CLI command <code>get services</code> in the NSX Intelligence node. Restart unexpected stopped services with NSX CLI command <code>restart service <service-name></code>.

License Events

License events arise from the NSX Manager node.

Event Name	Severity	Alert Message	Recommended Action
License Expired	Critical	A license has expired. When event detected: "The license of type <i>{license_edition_type}</i> has expired." When event resolved: "The expired license of type <i>{license_edition_type}</i> has been removed, updated, or is no longer expired."	Add a new, non-expired license: 1 In the NSX UI, by navigate to System > Licenses . 2 Click Add and specify the key of the new license. 3 Delete the expired license by selecting the check box and clicking Unassign .
License About to Expire	Medium	When event detected: "The license of type <i>{license_edition_type}</i> is about to expire." When event resolved: "The expiring license identified by <i>{license_edition_type}</i> has been removed, updated, or is no longer about to expire."	Add a new, non-expired license: 1 In the NSX UI, by navigate to System > Licenses . 2 Click Add and specify the key of the new license. 3 Delete the expired license by selecting the check box and clicking Unassign .

Load Balancer Events

Load balancer events arise from the NSX Edge node.

Event Name	Severity	Alert Message	Recommended Action
Load Balancer CPU Very High	Medium	Load balancer CPU usage is very high. When event detected: "The CPU usage of load balancer <i>{entity_id}</i> is <i>{system_resource_usage}%</i> , which is higher than the very high threshold of <i>{system_usage_threshold}%</i> ." When event resolved: "The CPU utilization of load balancer <i>{entity_id}</i> is <i>{system_resource_usage}%</i> , which is lower than the very high threshold of <i>{system_usage_threshold}%</i> ."	If the load balancer CPU utilization of is higher than <i>{system_usage_threshold}%</i> , the workload is too high for this load balancer. Rescale the load balancer service by changing the load balancer size from small to medium or from medium to large. If the CPU utilization of this load balancer is still high, consider adjusting the Edge appliance form factor size or moving load balancer services to other Edge nodes for the applicable workload.
Load Balancer Status Down	Medium	Load balancer service is down. When event detected: "The load balancer service <i>{entity_id}</i> is down." When event resolved: "The load balancer service <i>{entity_id}</i> is up."	Verify whether the load balancer service in the Edge node is running. If the status of the load balancer service is not ready, move the Edge node into maintenance mode, then exit maintenance mode. If the status of the load balancer service is still not recovered, please check whether there are any error log in syslog.

Event Name	Severity	Alert Message	Recommended Action
Virtual Server Status Down	Medium	Load balancer virtual service is down. When event detected: "The load balancer virtual server <i>{entity_id}</i> is down." When event resolved: "The load balancer virtual server <i>{entity_id}</i> is up."	Consult the load balancer pool to determine its status and verify its configuration. If incorrectly configured, reconfigure it and remove the load balancer pool from the virtual server then re-add it to the virtual server again.
Pool Status Down	Medium	When event detected: "The load balancer pool <i>{entity_id}</i> status is down." When event resolved: "The load balancer pool <i>{entity_id}</i> status is up."	1 Consult the load balancer pool to determine which members are down. 2 Check network connectivity from the load balancer to the impacted pool members. 3 Validate application health of each pool member. 4 Validate the health of each pool member using the configured monitor. When the health of the member is established, the pool member status is updated to healthy based on the Rise Count.

Manager Health Events

NSX Manager health events arise from the NSX Manager node cluster.

Event Name	Severity	Alert Message	Recommended Action
Duplicate IP Address	Medium	Manager node's IP address is in use by another device. When event detected: "Manager node <i>{entity_id}</i> IP address <i>{duplicate_ip_address}</i> is currently being used by another device in the network." When event detected: "Manager node <i>{entity_id}</i> appears to no longer be using <i>{duplicate_ip_address}</i> ."	1 Determine which device is using the Manager's IP address and assign the device a new IP address. Note Reconfiguring the Manager to use a new IP address is not supported. 2 Verify if the static IP address pool/ DHCP server is configured correctly. 3 Correct the IP address of the device if it is manually assigned.
Manager CPU Usage Very High	Critical	Manager node CPU usage is very high. When event detected: "The CPU usage on the Manager node <i>{entity_id}</i> has reached <i>{system_resource_usage}</i> %, which is at or above the very high threshold value of <i>{system_usage_threshold}</i> %." When event resolved: "The CPU usage on the Manager node <i>{entity_id}</i> has reached <i>{system_resource_usage}</i> %, which is below the very high threshold value of <i>{system_usage_threshold}</i> %."	Please review the configuration, running services and sizing of this Manager node. Consider adjusting the Manager appliance form factor size.

Event Name	Severity	Alert Message	Recommended Action
Manager CPU Usage High	Medium	Starting in NSX-T Data Center 3.0.1. Manager node CPU usage is high. When event detected: "The CPU usage on the Manager node <i>{entity_id}</i> has reached <i>{system_resource_usage}%</i>, which is at or above the high threshold value of <i>{system_usage_threshold}%</i>." When event resolved: "The CPU usage on the Manager node <i>{entity_id}</i> has reached <i>{system_resource_usage}%</i>, which is below the high threshold value of <i>{system_usage_threshold}%</i>."	Please review the configuration, running services and sizing of this Manager node. Consider adjusting the Manager appliance form factor size.
Manager Memory Usage Very High	Critical	Starting in NSX-T Data Center 3.0.1. Manager node memory usage is very high. When event detected: "The memory usage on the Manager node <i>{entity_id}</i> has reached <i>{system_resource_usage}%</i>, which is at or above the very high threshold value of <i>{system_usage_threshold}%</i>." When event resolved: "The memory usage on the Manager node <i>{entity_id}</i> has reached <i>{system_resource_usage}%</i>, which is below the very high threshold value of <i>{system_usage_threshold}%</i>."	Please review the configuration, running services and sizing of this Manager node. Consider adjusting the Manager appliance form factor size.
Manager Memory Usage High	Medium	Manager node memory usage is high. When event detected: "The memory usage on the Manager node <i>{entity_id}</i> has reached <i>{system_resource_usage}%</i>, which is at or above the high threshold value of <i>{system_usage_threshold}%</i>." When event resolved: "The memory usage on the Manager node <i>{entity_id}</i> has reached <i>{system_resource_usage}%</i>, which is below the high threshold value of <i>{system_usage_threshold}%</i>."	Please review the configuration, running services and sizing of this Manager node. Consider adjusting the Manager appliance form factor size.

Event Name	Severity	Alert Message	Recommended Action
Manager Disk Usage Very High	Critical	Manager node disk usage is very high. When event detected: "The disk usage for the Manager node disk partition <i>{disk_partition_name}</i> has reached <i>{system_resource_usage}%</i>, which is at or above the very high threshold value of <i>{system_usage_threshold}%</i>." When event resolved: "The disk usage for the Manager node disk partition <i>{disk_partition_name}</i> has reached <i>{system_resource_usage}%</i>, which is below the very high threshold value of <i>{system_usage_threshold}%</i>."	Examine the partition with high usage and see if there are any unexpected large files that can be removed.
Manager Disk Usage High	Medium	Manager node disk usage is high. When event detected: "The disk usage for the Manager node disk partition <i>{disk_partition_name}</i> has reached <i>{system_resource_usage}%</i>, which is at or above the high threshold value of <i>{system_usage_threshold}%</i>." When event resolved: "The disk usage for the Manager node disk partition <i>{disk_partition_name}</i> has reached <i>{system_resource_usage}%</i>, which is below the high threshold value of <i>{system_usage_threshold}%</i>."	Examine the partition with high usage and see if there are any unexpected large files that can be removed.
Manager Configuration Disk Usage Very High	Critical	Manager node config disk usage is very high. When event detected: "The disk usage for the Manager node disk partition <code>/config</code> has reached <i>{system_resource_usage}%</i>, which is at or above the very high threshold value of <i>{system_usage_threshold}%</i>. This can be an indication of high disk usage by the NSX Datastore service under the <code>/config/corfu</code> directory." When event resolved: "The disk usage for the Manager node disk partition <code>/config</code> has reached <i>{system_resource_usage}%</i>, which is below the very high threshold value of <i>{system_usage_threshold}%</i>."	Examine the <code>/config</code> partition and see if there are any unexpected large files that can be removed.

Event Name	Severity	Alert Message	Recommended Action
Manager Configuration Disk Usage High	Medium	Manager node config disk usage is high. When event detected: "The disk usage for the Manager node disk partition <code>/config</code> has reached <code>{system_resource_usage}%</code>, which is at or above the high threshold value of <code>{system_usage_threshold}%</code>. This can be an indication of rising disk usage by the NSX Datastore service under the <code>/config/corfu</code> directory." When event resolved: "The disk usage for the Manager node disk partition <code>/config</code> has reached <code>{system_resource_usage}%</code>, which is below the high threshold value of <code>{system_usage_threshold}%</code>."	Examine the <code>/config</code> partition and see if there are any unexpected large files that can be removed.
Operations DB Disk Usage High	Medium	The disk usage for the Manager node disk partition <code>/nonconfig</code> has reached <code>{system_resource_usage}%</code> which is at or above the high threshold value of <code>{system_usage_threshold}%</code> . This can be an indication of rising disk useage by the NSX Datastore service under the <code>/nonconfig/corfu</code> directory.	Please run the following tool, and contact GSS if any issues are reported <code>/opt/vmware/tools/support/inspect_checkpoint_issues.py --nonconfig</code> .
Operations DB Disk Usage Very High	Critical	The disk usage for the Manager node disk partition <code>/nonconfig</code> has reached <code>{system_resource_usage}%</code> which is at or above the very high threshold value of <code>{system_usage_threshold}%</code> . This can be an indication of rising disk useage by the NSX Datastore service under the <code>/nonconfig/corfu</code> directory.	Please run the following tool, and contact GSS if any issues are reported <code>/opt/vmware/tools/support/inspect_checkpoint_issues.py --nonconfig</code> .

NCP Events

NSX Container Plug-in (NCP) events arise from the ESXi and KVM nodes.

Event Name	Severity	Alert Message	Recommended Action
NCP Plugin Down	Critical	Manager Node has detected the NCP is down or unhealthy. When event detected: "Manager node has detected the NCP is down or unhealthy." When event resolved: "Manager Node has detected the NCP is up or healthy again."	To find the clusters which are having issues, invoke the NSX API: <code>GET /api/v1/systemhealth/container-cluster/ncp/status</code> to fetch all cluster statuses and determine the name of any clusters that report DOWN or UNKNOWN. Go to the NSX UI Inventory > Container > Clusters page to find the names of clusters that reported DOWN or UNKNOWN status and click the Nodes tab which lists all Kubernetes and PAS cluster members. For Kubernetes cluster: 1 Check NCP Pod liveness by finding the K8s master node from all the cluster members and log onto the master node. Then invoke the kubectl command <code>kubectl get pods --all-namespaces</code>. If there is an issue with the NCP Pod, please use <code>kubectl logs</code> command to check the issue and fix the error. 2 Check the connection between NCP and Kubernetes API server. The NSX CLI can be used inside the NCP Pod to check this connection status by invoking the following commands from the master VM. <pre>kubectl exec -it <NCP-Pod-Name> -n nsx-system bash nsxcli get ncp-k8s-api-server status</pre> If there is an issue with the connection, please check both the network and NCP configurations. 3 Check the connection between NCP and NSX Manager. The NSX CLI can be used inside the NCP Pod to check this connection status by invoking the following command from the master VM. <pre>kubectl exec -it <NCP-Pod-Name> -n nsx-system bash nsxcli get ncp-nsx status</pre>

Event Name	Severity	Alert Message	Recommended Action
			If there is an issue with the connection, please check both the network and NCP configurations. For PAS cluster: 1 Check the network connections between virtual machines and fix any network issues. 2 Check the status of both nodes and services and fix crashed nodes or services. Invoke the commands <code>bosh vms</code> and <code>bosh instances -p</code> to check the status of nodes and services.

Node Agents Health Events

Node agent health events arise from the ESXi and KVM nodes.

Event Name	Severity	Alert Message	Recommended Action
Node Agents Down	High	The agents running inside the Node VM appear to be down. When event detected: "The agents running inside the node VM appear to be down." When event resolved: "The agents inside the Node VM are running."	For ESX: 1 If Vmk50 is missing, see Knowledge Base article 67432. 2 If Hyperbus 4094 is missing: restarting <code>nsx-cfgagent</code> or restarting the container host VM may help. 3 If container host VIF is blocked, check the connection to the controller make sure all configurations are sent down. 4 If <code>nsx-cfgagent</code> has stopped, please restart <code>nsx-cfgagent</code>. For KVM: 1 If the Hyperbus namespace is missing, restarting the <code>nsx-opsagent</code> may help recreate the namespace. 2 If Hyperbus interface is missing inside the hyperbus namespace, restarting the <code>nsx-opsagent</code> may help. 3 If the <code>nsx-agent</code> has stopped, restart <code>nsx-agent</code>. For both ESX and KVM: 1 If the <code>node-agent</code> package is missing: check whether the <code>node-agent</code> package has been successfully installed in the container host VM. 2 If the interface for the <code>node-agent</code> in the container host VM is down: check the <code>eth1</code> interface status inside the container host VM.

Password Management Events

Password management events arise from the NSX Manager, NSX Edge, and the public gateway nodes.

Event Name	Severity	Alert Message	Recommended Action
Password expired	Critical	User password has expired. When event detected: "The password for user <i>{username}</i> has expired." When event resolved: "The password for the user <i>{username}</i> has been changed successfully or is no longer expired."	The password for the user <i>{username}</i> must be changed now to access the system. For example, to apply a new password to a user, invoke the following NSX API with a valid password in the request body: <pre>PUT /api/v1/node/users/<userid></pre> where <i><userid></i> is the ID of the user. If the admin user (with <i><userid></i> 10000) password has expired, admin must login to the system via SSH (if enabled) or console in order to change the password. Upon entering the current expired password, admin will be prompted to enter a new password.
Password about to expire	High	User password is about to expire. When event detected: "The password for user <i>{username}</i> is about to expire in <i>{password_expiration_days}</i> days."" When event resolved: "The password for the user <i>{username}</i> has been changed successfully or is no longer about to expire."	Ensure the password for the user identified by <i>{username}</i> is changed immediately. For example, to apply a new password to a user, invoke the following NSX API with a valid password in the request body: <pre>PUT /api/v1/node/users/<userid></pre> where <i><userid></i> is the ID of the user.
Password expiration approaching	Medium	User password is approaching expiration. When event detected: "The password for user <i>{username}</i> is about to expire in <i>{password_expiration_days}</i> days." When event resolved: "The password for the user <i>{username}</i> has been changed successfully or is no longer about to expire."	The password for the user identified by <i>{username}</i> needs to be changed soon. For example, to apply a new password to a user, invoke the following NSX API with a valid password in the request body: <pre>PUT /api/v1/node/users/<userid></pre> where <i><userid></i> is the ID of the user.

Routing Events

Event Name	Severity	Alert Message	Recommended Action
BGP Down	High	BGP neighbor down. When event detected: "In Router <i>{entity_id}</i>, BGP neighbor <i>{bgp_neighbor_ip}</i> is down, reason: <i>{failure_reason}</i>." When event resolved: "In Router <i>{entity_id}</i>, BGP neighbor <i>{bgp_neighbor_ip}</i> is up."	1 SSH into the Edge node. 2 Invoke the NSX CLI command: <code>get logical-routers</code> 3 Switch to the service router <i>{sr_id}</i>. 4 Check <code>/var/log/syslog</code> to see if there are any errors related to BGP connectivity.
Bidirectional Forwarding Detection Down (BFD) on External Interface	High	BFD session is down. When event detected: "In router <i>{entity_id}</i>, BFD session for peer <i>{peer_address}</i> is down." When event resolved: "In router <i>{entity_id}</i>, BFD session for peer <i>{peer_address}</i> is up."	1 SSH into the Edge node. 2 Invoke the NSX CLI command: <code>get logical-routers</code> 3 Switch to the service router <i>{sr_id}</i>. 4 Verify the connectivity by invoking the NSX CLI command: <code>ping <peer_address></code>.
Routing Down	High	All BGP/BFD sessions are down. When event detected: "All BGP/BFD sessions are down." When event resolved: "At least one BGP/BFD sessions up."	1 Invoke the NSX CLI command <code>get logical-routers</code> to get the Tier0 service router. 2 Switch to the Tier0 service router VRF, then invoke the following NSX CLI commands: ■ Verify connectivity: <code>ping <BFD peer IP address></code> ■ Check BFD health: <pre>get bfd-config get bfd-sessions</pre> ■ Check BGP health: <code>get bgp neighbor summary</code> <pre>get bfd neconfig get bfd-sessions</pre> Check <code>/var/log/syslog</code> to see if there are any errors related to BGP connectivity.
Static Routing Removed	High	Static route removed. When event detected: "In router <i>{entity_id}</i>, static route <i>{static_address}</i> was removed because BFD was down." When event resolved: "In router <i>{entity_id}</i>, static route <i>{static_address}</i> was re-added as BFD recovered."	1 SSH into the Edge node. 2 Invoke the NSX CLI command: <code>get logical-routers</code> 3 Switch to the service router <i>{sr_id}</i>. 4 Verify the connectivity by invoking the NSX CLI command: <pre>get bgp neighbor summary</pre> 5 Also, verify the configuration in both NSX and the BFD peer to ensure that timers have not been changed.

Transport Node Health

Transport node health events arise from the KVM and ESXi nodes.

Event Name	Severity	Alert Message	Recommended Action
LAG Member Down	Medium	LACP reporting member down. When event detected: "LACP reporting member down." When event resolved: "LACP reporting member up."	Check the connection status of LAG members on hosts. 1 In the NSX UI, navigate to Fabric > Nodes > Transport Nodes > Host Transport Nodes. 2 In the Host Transport Nodes list, check the Node Status column. Find the Transport node with the degraded or down Node Status. 3 Select <transport node> > Monitor. Find the bond (uplink) which is reporting degraded or down. 4 Check the LACP member status details by logging into the failed host and running the appropriate command: ■ ESXi: <code>esxcli network vswitch dvs vmware lacp status get</code> ■ KVM: <code>ovs-appctl bond/show</code> and <code>ovs-appctl lacp/show</code>
N-VDS Uplink Down	Medium	Uplink is going down. When event detected: "Uplink is going down." When event resolved: "Uplink is going up."	Check the physical NICs status of uplinks on hosts. 1 In the NSX UI, navigate to Fabric > Nodes > Transport Nodes > Host Transport Nodes. 2 In the Host Transport Nodes list, check the Node Status column. Find the Transport node with the degraded or down Node Status. 3 Select <transport node> > Monitor. Check the status details of the bond (uplink) which is reporting degraded or down. To avoid a degraded state, ensure all uplink interfaces are connected and up regardless of whether they are in use or not.

VPN Events

VPN events arise from the NSX Edge and public gateway nodes.

Event Name	Severity	Alert Message	Recommended Action
IPsec Policy-Based Session Down	Medium	Policy-based IPsec VPN session is down. When event detected: "The policy-based IPsec VPN session <i>{entity_id}</i> is down. Reason: <i>{session_down_reason}</i>." When event resolved: "The policy-based IPsec VPN session <i>{entity_id}</i> is up."	Check IPsec VPN session configuration and resolve errors based on the session down reason.
IPsec Route-Based Session Down	Medium	Route-based IPsec VPN session is down. When event detected: "The route-based IPsec VPN session <i>{entity_id}</i> is down. Reason: <i>{session_down_reason}</i>." When event resolved: "The route-based IPsec VPN session <i>{entity_id}</i> is up."	Check IPsec VPN session configuration and resolve errors based on the session down reason.
IPsec Policy-Based Tunnel Down	Medium	Policy-based IPsec VPN tunnels are down. When event detected: "One or more policy-based IPsec VPN tunnels in session <i>{entity_id}</i> are down." When event resolved: "All policy-based IPsec VPN tunnels in session <i>{entity_id}</i> are up."	Check IPsec VPN session configuration and resolve errors based on the tunnel down reason.
IPsec Route-Based Tunnel Down	Medium	Route-based IPsec VPN tunnels are down. When event detected: "One or more route-based IPsec VPN tunnels in session <i>{entity_id}</i> are down." When event resolved: "All route-based IPsec VPN tunnels in session <i>{entity_id}</i> are up."	Check IPsec VPN session configuration and resolve errors based on the tunnel down reason.
L2VPN Session Down	Medium	L2VPN session is down. When event detected: "The L2VPN session <i>{entity_id}</i> is down." When event resolved: "The L2VPN session <i>{entity_id}</i> is up."	Check IPsec VPN session configuration and resolve errors based on the reason.

Identity Firewall Events

Event Name	Severity	Alert Message	Recommended Action
Connectivity to AD server	Critical	Connectivity to AD server is lost. When event detected: Connectivity to Identity Firewall AD server is down. When event detected: Connectivity to Identity Firewall AD server is up .	1 The AD server is reachable from NSX nodes. 2 The AD server details are configured correctly in NSX. 3 The AD server is running correctly. 4 There are no firewalls blocking access between the AD server and NSX nodes. After fixing the connection issue, use the "TEST CONNECTION" in LDAP server UI to test the connection to AD server.
Errors during Delta Sync	Critical	Failure to Sync AD server <i>error description</i> When event detected: Failure during selective sync of Identity Firewall AD server: <i>error details</i> . When event detected: Selective sync errors of Identity Firewall AD server fixed.	1 Verify whether the load balancer service in the Edge node is running. 2 If the status of the load balancer service is not ready, move the Edge node into maintenance mode, then exit maintenance mode. 3 If the status of the load balancer service is still not recovered, please check whether there are any error log in syslog.

View Alarm Information

Alarms information is displayed in several locations within the NSX Manager interface. Alarm and event information is also included with other notifications in the Notifications drop-down menu in the title bar.

An alarm can be in one of the following states:

State	Description
Open	Alarm is in an active, unacknowledged state.
Acknowledged	Alarm has been acknowledged by a user. The alarm remains open but no longer appears in the NSX Manager notifications.
Suppressed	Status reporting for this alarm has been disabled by the user for a user-specified duration.
Resolved	Alarm has been resolved, whether by the system or through user action. The alarm will continue to appear in the alarm table in the Resolved state for up to eight days, after which it automatically deletes. (The system may delete resolved alarms earlier to accommodate resource needs.) Note If a user changes an alarm state to Resolved but the condition that triggered the alarm is not resolved, a new alarm instance will be instantiated. Also, an event may be resolved for several minutes before the reported state updates in the interface.

Note The following steps show how to view alarms from the Home page. However, you can also view alarms from other pages, such as the Tier-0, Tier-1, and Load Balancing pages, among others. See the Alarms columns in the tables on these pages.

Procedure

- 1 Navigate to the Home page and click **Alarms**.

Note A red exclamation mark (!) next to the **Alarms** panel label indicates at least one open alarm with a severity of Critical. .

The Alarms panel appears, displaying along the top graphic dashboards such as Active Alarms, Top Features with the Most Alarms, and Top Events by Occurrence. Below the dashboards is a sortable, filterable list of the current alarms. The table details the following information about each active alarm:

- Feature affected
- Event Type
- Node
- Entity
- Severity (Critical, High, Medium)
- Last Reported Time
- Alarm State (Open, Suppressed, Resolved, Acknowledged)

Each row in the Alarms table can be expanded to show more details.

- 2 Filter the results displayed in the dashboards by clicking the funnel icon in the upper-right corner of the dashboards.

You can filter by the last 24 hours, last 48 hours, or custom time range, or all open alarms.

- 3 Filter the results displayed in the table by clicking the filter text box above the table.

You are prompted to specify a filter: Alarm State, Description, Entity Name, Entity Type, Event Type, Node, and so on.

What to do next

After viewing an alarm, you can decide on how to respond. See [Managing Alarm States](#).

View Alarm Definitions

Detailed alarm definitions are provided on a separate panel in the Alarms tab. You can open the panel directly or arrive there by clicking the value in the Event Type column in the Alarms table.

Alarms details are displayed in several locations in the NSX Manager. See [View Alarm Information](#).

Procedure

1 From the Alarms tab.

- a Navigate to the Home page and click **Alarms**.

The Alarms panel has two modes, as shown at the top of the panel: **Alarms** and **Alarm Definitions**.

- b Click **Alarm Definitions**.

The Alarms tab redisplay to show the table of Alarm definitions.

2 From the Tier-0 Gateways page.

- a Go to **Networking > Connectivity > Tier-0 Gateways**.

The Open Alarms column of the gateway table displays the number of open alarms.

- b Click the number in the Open Alarms column.

A dialog opens to display the open alarms in table format.

- c Click the value in the Event Type column.

This action moves you to the **Home > Alarms > Alarm Definitions** panel described above.

3 From the Load Balancing page.

- a Go to **Networking > Network Services > Load Balancing**.

The Open Alarms column of the gateway table displays the number of open alarms.

- b Click the number in the Open Alarms column.

A dialog opens to display the open alarms in table format.

- c Click the value in the Event Type column.

This action moves you to the **Home > Alarms > Alarm Definitions** panel described above.

4 After you access the **Alarm Definitions**, expand any definition to view details and user-definable settings.

Alarm definition details include:

Column	Description
Feature	Displays the component where the alarm is originating, for example: Transport Node.
Event Type	Displays the specific type of error, for example: CPU Usage High.
Severity	Displays the level of alarm: Critical, High, or Medium.
Enabled	Displays whether detection of the Alarm is enabled.
Create Alarms	Displays whether to report the alarm in the interface or API.
Create SNMP Traps	Displays whether the system emits an SNMP trap when the alarm is detected or resolved.

The panel also displays the following:

Item	Description
Description	Describes the condition that triggers the alarm.
Recommended Action	Describes steps you can take to correct the condition.
SNMP OID for Event true	Displays the SNMP Object Identifier for the Event when status is true.
SNMP OID for Event false	Displays the SNMP Object Identifier for the Event when status is false.
Threshold	User-configured threshold for triggering the alarm.
Sensitivity (%)	User-configured sensitivity for triggering the alarm.

What to do next

Some of fields in an alarm definition can be modified. See [Configuring Alarm Definition Settings](#).

Configuring Alarm Definition Settings

Several settings in an alarm definition can be customized. From the Alarm Definitions page, you can enable or disable an alarm, configure if an event (when true) creates an alarm, create an SNMP trap, set alarm threshold, and set alarm sensitivity. From the Alarm Definitions page, you can enable or disable detection of an alarm, whether an alarm is reported in the API/user interface, and whether a SNMP trap is emitted when an alarm is detected or resolved.

You can configure the following alarm definition settings:

Setting	Control Type	Description
Enabled	Toggle	Enables or disables detection of the alarm.
Create Alarms	Toggle	Enables or disables whether the alarm is reported in the API/UI.
Create SNMP Traps	Toggle	Enables or disables whether an SNMP trap is emitted when an alarm is detected or resolved.
Threshold	Numerical value	Configures the threshold for triggering the event. This value determines if a single sample is true and triggers an event. ■ For CPU, disk, and memory alarms, threshold is the percentage usage value to indicate an alarming condition. ■ For certificate or license expiration alarms, this is the number of days before expiration, including local password expiration.
Sensitivity (%)	Numerical value (percentage)	Configures the sensitivity for triggering the alarm. Sensitivity defines the conditions that trigger an alarm. (The sample size is internally defined and cannot be modified.) If the sample size is ten and sensitivity is set to 80%, then eight or more occurrences in the sample of ten raises the alarms. See the NSX-T Data Center REST API documentation .

Procedure

- 1 Navigate to the Home page and click **Alarms**.

The Alarms panel has two modes, as shown at the top of the panel: **Alarms** and **Alarm Definitions**.

2 Click **Alarm Definitions**.

The Alarms tab redisplay to show the Alarm Definitions panel.

3 Right-click three vertical dots icon in the leftmost column of an alarm, and select **Edit**.

The selected alarm definition expands to show the definition details, and puts the configurable settings into edit mode.

4 Modify the settings as desired.

5 Click **Save**.

What to do next

For details about alarm definitions, see [View Alarm Definitions](#). For details about SNMP traps, see [Simple Network Management Protocol \(SNMP\)](#).

Managing Alarm States

In addition to correcting the underlying causes, you can manage alarms by modifying their states as reported in the Alarms list.

Triggered alarms can be in one of the following states: Open, Acknowledged, Suppressed, or Resolved.

Procedure

1 Navigate to the Home page and click **Alarms**.

The Alarms panel has two modes, as shown at the top of the panel: **Alarms** and **Alarm Definitions**.

2 Click the **Alarms** mode, if the panel is not already displayed.

The **Alarms** panel displays a list of all alarms, including Resolved alarms.

Note Resolved alarms continue to be listed for up to eight days after their resolution.

3 Locate the alarm in the table on the page, and select the checkbox in the leftmost column.

4 Click **Action** and select the desired action.

- If you change the state of an alarm to Acknowledged, this indicates that you are aware of, and have acknowledged, the alarm.
- If you move an alarm into a Suppressed state, you are prompted to specify the duration in hours. After the specified duration passes, the alarm state reverts to Open. However, if the system determines the condition has been corrected, the alarm state changes to Resolved.
- You can restore the state of an Acknowledged or Suppressed alarm to Open.
- You cannot change the state of a Resolved alarm.

The value in the Alarm State column updates accordingly.

Using vRealize Log Insight for System Monitoring

You can monitor your NSX-T Data Center environment using the Log Insight NSX-T content pack.

This content pack has the following alerts:

Alert Name	Description
SysCpuUsage	CPU usage is above 95% for more than 10 minutes.
SysMemUsage	Memory usage is above 95% for more than 10 minutes.
SysDiskUsage	Disk usage for one or more partitions is above 89% for more than 10 minutes.
PasswordExpiry	Password for appliance user account is about to expire or expired.
CertificateExpiry	One or more CA signed certificate is expired.
ClusterNodeStatus	Local edge cluster node is down.
BackupFailure	NSX scheduled backup operation failed.
VipLeadership	NSX Management cluster VIP is down.
ApiRateLimit	Client API reached configured threshold.
CorfuQuorumLost	Two nodes went down in the cluster and lost corfu quorum.
DfwHeapMem	DFW heap memory exceeded configured threshold.
ProcessStatus	Critical process status changed.
ClusterFailoverStatus	SR high availability state changed or active/standby services failover.
DhcpPoolUsageOverloadedEvent	DHCP pool reached configured usage threshold.
FabricCryptoStatus	Edge crypto mux driver is down for failing Known_Answer_Tests (KAT).
VpnTunnelState	VPN tunnel is down.
BfdTunnelStatus	BFD Tunnel status changed.
RoutingBgpNeighborStatus	BGP neighbor status is down.
VpnL2SessionStatus	L2 VPN session is down.
VpnIkeSessionStatus	IKE session is down.
RoutingStatus	Routing(BGP/BFD) is down.
DnsForwarderStatus	DNS forwarder running status is DOWN.
TnConnDown_15min	Transport Node connection to a controller/Manager is down for at least 15 minutes.
TnConnDown_5min	Transport Node connection to controller/Manager is down for at least 5 minutes.
ServiceDown	One or more services are down.
IpNotAvailableInPool	There is no IP available in the Pool or reaches configured threshold.
LoadBalancerError	NSX Load Balancer Service status is ERROR.
LoadBalancerDown	NSX Load Balancer Service status is DOWN.
LoadBalancerVsDown	VS status: all pool members are down.

Alert Name	Description
LoadBalancerPoolDown	Pool status: all pool members are down.
ProcessCrash	Process or daemon crashes in the datapath or other LB process like dispatcher, etc..

Using vRealize Operations Manager for System Monitoring

You can monitor your NSX-T Data Center environment using vRealize Operations Manager.

Table 16-1. Alerts in the Management Pack for NSX-T Data Center

Alert	Description	Recommendation
NSX-T Data Center Management service has failed	Triggered when the management service on the NSX-T Data Center host is not running.	Log in to the NSX Manager and restart the failed management service.
Logical Switch's admin state is not UP	Triggered when the admin state is disabled on the logical switch.	Log in to NSX-T Data Center and enable the admin state if it is intended so.
Edge Node Controller/ Manager Connectivity is not UP	Triggered when the edge node connectivity status is down in NSX-T Data Center.	Check the Edge node connection status with Controller Cluster and Manager Cluster and fix the broken connection.
Edge Host node is in Failed/Error state	Triggered when the host node in NSX-T Data Center is in error or failed state due to one of the following reasons: ■ Edge configuration error ■ Installation failure ■ Uninstallation failure ■ Upgrade failure ■ Virtual Machine deployment failure ■ Virtual Machine power off failure ■ Virtual Machine power on failure ■ Virtual Machine undeployment failure	Edge host node is in failed/error state, check the host node state and fix the issue.
BFD service is disabled	Triggered when the BFD service is not enabled on the logical router.	BFD Service for a TIER0 router is not enabled even though neighbors are configured. Enable the BFD service if required.
NAT rule not configured	Triggered when the NAT rule on the logical router is not configured.	Log in to the NSX Manager and add the NAT rules for the Logical Router.
Static Route not configured	Triggered when the static route on the logical router is not configured.	Log in to the NSX Manager and add the static routes for the Logical Router if required.

Table 16-1. Alerts in the Management Pack for NSX-T Data Center (continued)

Alert	Description	Recommendation
Route Advertisement service is disabled	Triggered when the route advertisement service is not enabled on the logical router.	Route Advertisement service for a TIER1 router is not enabled even though route advertisements are configured, log in to NSX Manager and enable the service.
Route Redistribution service is disabled	Triggered when the route redistribution service is not enabled on the logical router.	Route Redistribution service for a TIER0 router is not enabled even though route redistribution rules are configured, log in to NSX Manager and enable the service.
ECMP service is disabled for Logical Router	Triggered when the ECMP service is not enabled on the logical router.	BGP ECMP service for a TIER0 router is not enabled even though neighbors are configured, log in to NSX Manager and enable the service.
Controller Node Connectivity is broken	Triggered when the controller node connection status is down in NSX-T Data Center	Log in to NSX Manager and check the connectivity of the controller node with Management Node and Controller cluster and resolve the disconnected state.
Less than 3 controller nodes are deployed	Triggered when the NSX-T Data Center server has less than three controller nodes.	Deploy at least 3 controller nodes in the cluster.
Controller Cluster Status is not stable	Triggered when all the controller nodes are down in NSX-T Data Center.	Check the status of controller cluster.
Management Status is not stable	Triggered when the status of any node on the management cluster is down.	Check the status of management cluster.
File System usage is more than 85 percent	Triggered when the guest file systems usage of the Controller Virtual Machine is more than 85 percent.	File system usage is more than 85, check and clean the File System to make more space.
File System usage is more than 75 percent	Triggered when the guest file systems usage of the Controller Virtual Machine is more than 75 percent.	File system usage is more than 75, check and clean the File System to make more space.
File System usage is higher than 70 percent	Triggered when the guest file systems usage of the Controller Virtual Machine is more than 70 percent.	File system usage is more than 70, check and clean the File System to make more space.

Table 16-1. Alerts in the Management Pack for NSX-T Data Center (continued)

Alert	Description	Recommendation
Edge Cluster Status is down	Triggered when edge cluster status is down.	Check the edge cluster status and if required follow standard troubleshooting steps recommended by NSX-T Data Center documentation and VMware documentation.
Logical Switch State has failed	Triggered when the state of logical switch has failed.	Check the logical switch state and if necessary follow standard troubleshooting steps recommended by NSX-T Data Center documentation and VMware documentation.
Load Balancer Service operational status down	Triggered when the operational status of load balancer service is down.	Check the operational status of load balancer service and if necessary follow standard troubleshooting steps recommended by NSX-T Data Center documentation and VMware documentation.
Load balancer service operational status error	Triggered when the operational status of load balancer service contains error.	Check the operational status of load balancer service and if necessary follow standard troubleshooting steps recommended by NSX-T Data Center documentation and VMware documentation.
Load Balancer virtual server operational state down	Triggered when the operational state of load balancer virtual server is down.	Check the operational state of load balancer virtual server and if necessary follow standard troubleshooting steps recommended by NSX-T Data Center documentation and VMware documentation.

Table 16-1. Alerts in the Management Pack for NSX-T Data Center (continued)

Alert	Description	Recommendation
Load Balancer virtual server operational state detached	Triggered when the operational state of load balancer virtual server is detached.	Check the operational state of load balancer virtual server and if necessary follow standard troubleshooting steps recommended by NSX-T Data Center documentation and VMware documentation.
Edge node configuration state has failed	Triggered when the configuration state of edge node has failed.	Check the configuration state of the edge node and if necessary follow standard troubleshooting steps recommended by NSX-T Data Center documentation and VMware documentation.
Management service monitor runtime state has failed	Triggered when the monitor runtime state of the management service stops running.	Log in to the NSX Manager VA and restart the failed management service.
Management cluster's management status is not stable	Triggered when the management status of a management cluster is not stable.	Check the status of management cluster.
Less than 3 manager nodes are deployed	Triggered when the NSX-T Data Center server has less than three manager nodes deployed.	Deploy at least 3 manager nodes in the cluster.
Manager node connectivity is broken	Triggered when the manager connection status of manager node is down.	Log in to NSX Manager and check the manager connectivity of manager node and follow standard troubleshooting steps recommended by NSX-T Data Center documentation and VMware documentation.
File System usage of manager node is more than 85 percent	Triggered when the guest file systems usage of the manager node is more than 85 percent.	File system usage is more than 85, check and clean the File System to make more space.
File System usage of manager node is more than 75 percent	Triggered when the guest file systems usage of the manager node is more than 75 percent.	File system usage is more than 75, check and clean the File System to make more space.
File System usage of manager node is more than 70 percent	Triggered when the guest file systems usage of the manager node is more than 70 percent.	File system usage is more than 70, check and clean the File System to make more space.

Using vRealize Network Insight Cloud for System Monitoring

You can monitor your NSX-T Data Center environment using vRealize Network Insight Cloud.

Table 16-2. vRealize Network Insight Computed NSX-T Events

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80205	NSXTNoUplinkConnectivityEvent	Warning	NSX-T Tier-1 logical router disconnect event	NSX-T Tier-1 logical router is disconnected from Tier-0 router. Networks under this router are not reachable from outside and vice versa.
1.3.6.1.4.1.6876.100.1.0.80206	NSXTRoutingAdvertisementEvent	Warning	Routing advertisement disabled	Routing advertisement is disabled for NSX-T Tier-1 logical router. Networks under this router are not reachable from outside.
1.3.6.1.4.1.6876.100.1.0.80207	NSXTManagerConnectivityDownEvent	Critical	NSX-T Edge Node has no manager connectivity	NSX-T Edge Node has lost manager connectivity.
1.3.6.1.4.1.6876.100.1.0.80208	NSXTControllerConnectivityDegradedEvent	Warning	Controller connectivity degraded for NSX-T Edge Node	NSX-T Edge Node is not able to communicate with one or more controllers.
1.3.6.1.4.1.6876.100.1.0.80209	NSXTControllerConnectivityDownEvent	Critical	NSX-T Edge Node has no controller connectivity	NSX-T Edge Node is not able to communicate with any of the controllers.
1.3.6.1.4.1.6876.100.1.0.80210	NSXTMTuMismatchEvent	Warning	MTU mismatch between NSX-T Tier-0 and uplink switch/router	The MTU configured on interfaces of Tier-0 logical router does not match with the interfaces of uplink switch/router from same L2 network. This can impact the network performance.

Table 16-2. vRealize Network Insight Computed NSX-T Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80211	NSXTExcludedVmFlowEvent	Info	One or More VMs excluded from NSX-T DFW Firewall.	One or more VMs are not protected by NSX-T DFW firewall. vRealize Network Insight will not receive IPFIX flows for these VMs.
1.3.6.1.4.1.6876.100.1.0.80212	NSXTDoubleVlanTaggingEvent	Warning	Uplink Vlan misconfiguration	Communication is disrupted because VLAN on uplink port of Tier 0 router is different than VLAN on the external gateway.
1.3.6.1.4.1.6876.100.1.0.80213	NSXTNoTzAttachedOnTnEvent	Warning	No transport zone is attached to the transport node.	No transport zone attached to the transport node. VMs might lose connectivity because of this.
1.3.6.1.4.1.6876.100.1.0.80214	NSXTVtepDeleteEvent	Warning	No VTEP available on the transport node.	All vteps are deleted from the transport node. VMs might lose connectivity because of this.
1.3.6.1.4.1.6876.100.1.0.80225	NSXTControllerNodeToControlClusterConnectivityEvent	Critical	NSX-T controller node has no control cluster connectivity	NSX-T controller node has lost control cluster connectivity.
1.3.6.1.4.1.6876.100.1.0.80226	NSXTControllerNodeToMgmtPlaneConnectivityEvent	Critical	NSX-T controller node has no management plane connectivity	NSX-T controller node has lost management plane connectivity.
1.3.6.1.4.1.6876.100.1.0.80227	NSXTMPNodeToMgmtClusterConnectivityEvent	Critical	NSX-T management node has no management cluster connectivity	NSX-T management node has lost management cluster connectivity.
1.3.6.1.4.1.6876.100.1.0.80246	NSXTHostNodeMgmtConnectivityStatusDownEvent	Warning	NSX-T Host Node has no manager connectivity	Desynchronization between NSX Manager's State of connectivity with Host Transport Nodes
1.3.6.1.4.1.6876.100.1.0.80247	NSXTEdgeNodeCtrlConnectivityStatusUnknownEvent	Critical	Controller connectivity for NSX-T Edge Node is Unknown.	NSX-T Edge Node Controller connectivity is Unknown.

Table 16-2. vRealize Network Insight Computed NSX-T Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80248	NSXTHostNodeCtrlConnectivityStatusDownEvent	Warning	NSX-T Host Node has no controller connectivity	NSX-T Host Node is not able to communicate with any of the controllers.
1.3.6.1.4.1.6876.100.1.0.80249	NSXTHostNodeCtrlConnectivityStatusDegradedEvent	Warning	Controller connectivity degraded for NSX-T Host Node	NSX-T Host Node is not able to communicate with one or more controllers.
1.3.6.1.4.1.6876.100.1.0.80250	NSXTHostNodeCtrlConnectivityStatusUnknownEvent	Warning	Controller connectivity for NSX-T Host Node is Unknown.	NSX-T Host Node Controller connectivity is Unknown.
1.3.6.1.4.1.6876.100.1.0.80228	NSXTHostNodePnicStatusDownEvent	Warning	NSX-T Host Transport Node Pnic Status is 'Down'.	NSX-T Host Transport Node Pnic Status is 'Down'.
1.3.6.1.4.1.6876.100.1.0.80229	NSXTHostNodePnicStatusDegradedEvent	Warning	NSX-T Host Transport Node Pnic Status is 'Degraded'	NSX-T Host Transport Node Pnic Status is 'Degraded'.
1.3.6.1.4.1.6876.100.1.0.80230	NSXTHostNodePnicStatusUnknownEvent	Warning	NSX-T Host Transport Node Pnic Status is 'Unknown'.	NSX-T Host Transport Node Pnic Status is 'Unknown'.
1.3.6.1.4.1.6876.100.1.0.80237	NSXTEdgeNodePnicStatusDownEvent	Critical	NSX-T Edge Transport Node Pnic Status is 'Down'.	NSX-T Edge Transport Node Pnic Status is 'Down'.
1.3.6.1.4.1.6876.100.1.0.80238	NSXTEdgeNodePnicStatusDegradedEvent	Critical	NSX-T Edge Transport Node Pnic Status is 'Degraded'.	NSX-T Edge Transport Node Pnic Status is 'Degraded'.
1.3.6.1.4.1.6876.100.1.0.80239	NSXTEdgeNodePnicStatusUnknownEvent	Critical	NSX-T Edge Transport Node Pnic Status is 'Unknown'.	NSX-T Edge Transport Node Pnic Status is 'Unknown'.
1.3.6.1.4.1.6876.100.1.0.80231	NSXTHostNodeTunnelStatusDownEvent	Warning	NSX-T Host Transport Node Tunnel Status is 'Down'.	NSX-T Host Transport Node Tunnel Status is 'Down'.
1.3.6.1.4.1.6876.100.1.0.80232	NSXTHostNodeTunnelStatusDegradedEvent	Warning	NSX-T Host Transport Node Tunnel Status is 'Degraded'.	NSX-T Host Transport Node Tunnel Status is 'Degraded'.
1.3.6.1.4.1.6876.100.1.0.80233	NSXTHostNodeTunnelStatusUnknownEvent	Warning	NSX-T Host Transport Node Tunnel Status is 'Unknown'.	NSX-T Host Transport Node Tunnel Status is 'Unknown'.

Table 16-2. vRealize Network Insight Computed NSX-T Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80240	NSXTEdgeNodeTunnelStatusDownEvent	Critical	NSX-T Edge Transport Node Tunnel Status is 'Down'.	NSX-T Edge Transport Node Tunnel Status is 'Down'.
1.3.6.1.4.1.6876.100.1.0.80241	NSXTEdgeNodeTunnelStatusDegradeEvent	Critical	NSX-T Edge Transport Node Tunnel Status is 'Degraded'.	NSX-T Edge Transport Node Tunnel Status is 'Degraded'.
1.3.6.1.4.1.6876.100.1.0.80242	NSXTEdgeNodeTunnelStatusUnknownEvent	Critical	NSX-T Edge Transport Node Tunnel Status is 'Unknown'.	NSX-T Edge Transport Node Tunnel Status is 'Unknown'.
1.3.6.1.4.1.6876.100.1.0.80234	NSXTHostNodeStatusDownEvent	Warning	NSX-T Host Transport Node Status is 'Down'.	NSX-T Host Transport Node Status is 'Down'.
1.3.6.1.4.1.6876.100.1.0.80235	NSXTHostNodeStatusDegradeEvent	Warning	NSX-T Host Transport Node Status is 'Degraded'.	NSX-T Host Transport Node Status is 'Degraded'.
1.3.6.1.4.1.6876.100.1.0.80236	NSXTHostNodeStatusUnknownEvent	Warning	NSX-T Host Transport Node Status is 'Unknown'.	NSX-T Host Transport Node Status is 'Unknown'.
1.3.6.1.4.1.6876.100.1.0.80243	NSXTEdgeNodeStatusDownEvent	Critical	NSX-T Edge Transport Node Status is 'Down'.	NSX-T Edge Transport Node Status is 'Down'.
1.3.6.1.4.1.6876.100.1.0.80244	NSXTEdgeNodeStatusDegradeEvent	Critical	NSX-T Edge Transport Node Status is 'Degraded'.	NSX-T Edge Transport Node Status is 'Degraded'.
1.3.6.1.4.1.6876.100.1.0.80245	NSXTEdgeNodeStatusUnknownEvent	Critical	NSX-T Edge Transport Node Status is 'Unknown'.	NSX-T Edge Transport Node Status is 'Unknown'.
1.3.6.1.4.1.6876.100.1.0.80252	NSXTLogicalSwitchAdminStatusDownEvent	Warning	NSX-T Logical Switch Admin Status is 'Down'	NSX-T Logical Switch Admin Status is 'Down'

Table 16-2. vRealize Network Insight Computed NSX-T Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80253	NSXTLogicalPortOperationalStatusDownEvent	Critical	NSX-T Logical Port Operational Status is 'Down'	NSX-T Logical Port Operational Status is 'Down'. This could cause a communication failure between two virtual interfaces (VIFs) that are connected to the same logical switch, for example, you cannot ping one VM from another.
1.3.6.1.4.1.6876.100.1.0.80254	NSXTLogicalPortOperationalStatusUnknownEvent	Warning	NSX-T Logical Port Operational Status is 'Unknown'	NSX-T Logical Port Operational Status is 'Unknown'. This could cause a communication failure between two virtual interfaces (VIFs) that are connected to the same logical switch, for example, you cannot ping one VM from another.
1.3.6.1.4.1.6876.100.1.0.80255	NSXTComputeManagerConnectionStatusNotUpEvent	Warning	NSX-T Compute Manager Connection Status in not up	NSX-T Compute Manager Connection status is not up
1.3.6.1.4.1.6876.100.1.0.80256	NSXTClusterBackupDisabledEvent	Warning	NSX-T Manager backup is not scheduled.	NSX-T Manager backup is not scheduled
1.3.6.1.4.1.6876.100.1.0.80257	NSXTDFWFirewallDisabledEvent	Critical	NSX-T DFW Firewall is disabled.	Distributed Firewall is disabled in the NSX-T Manager
1.3.6.1.4.1.6876.100.1.0.80258	NSXTLogicalPortReceivedPacketDropEvent	Warning	NSX-T Logical Port Received Packets are getting dropped.	Received packets are getting dropped on the NSX-T Logical Port and associated entities might get affected
1.3.6.1.4.1.6876.100.1.0.80259	NSXTLogicalPortTransmittedPacketDropEvent	Warning	NSX-T Logical Port Transmitted Packets are getting dropped.	Transmitted packets are getting dropped on the NSX-T Logical Port and associated entities might get affected

Table 16-2. vRealize Network Insight Computed NSX-T Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80260	NSXTLogicalSwitchReceivedPacketDropEvent	Warning	NSX-T Logical Switch Received Packets are getting dropped	Received packets are getting dropped on the NSX-T Logical Switch and associated entities might get affected
1.3.6.1.4.1.6876.100.1.0.80261	NSXTLogicalSwitchTransmittedPacketDropEvent	Warning	NSX-T Logical Switch Transmitted Packets are getting dropped	Transmitted packets are getting dropped on the NSX-T Logical Switch and associated entities might get affected
1.3.6.1.4.1.6876.100.1.0.80262	NSXTRxPacketDropOnMPNicEvent	Warning	Received packets are dropping on NSX-T Management Node's network interface	Received packets are getting dropped on NSX-T Management Node's network interface. This may impact the network traffic related to NSX-T management cluster.
1.3.6.1.4.1.6876.100.1.0.80263	NSXTRxPacketDropOnEdgeTnNicEvent	Critical	Received packets are dropping on NSX-T Edge Node's network interface	Received packets are getting dropped on NSX-T Edge Node's network interface. This may impact the network traffic of edge cluster.
1.3.6.1.4.1.6876.100.1.0.80264	NSXTRxPacketDropOnHostTnNicEvent	Warning	Received packets are dropping on NSX-T Host Node's network interface	Received packets are getting dropped on NSX-T Host Node's network interface. This may impact the network traffic on ESXi Host.
1.3.6.1.4.1.6876.100.1.0.80265	NSXTTxPacketDropOnMPNicEvent	Warning	Transmitted packets are dropping on NSX-T Management Node's network interface	Transmitted packets are getting dropped on NSX-T Management Node's network interface. This may impact the network traffic related to NSX-T management cluster.

Table 16-2. vRealize Network Insight Computed NSX-T Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80266	NSXTTxPacketDropOnEdgeTnNicEvent	Critical	Transmitted packets are dropping on NSX-T Edge Node's network interface	Transmitted packets are getting dropped on NSX-T Edge Node's network interface. This may impact the network traffic of edge cluster.
1.3.6.1.4.1.6876.100.1.0.80267	NSXTTxPacketDropOnHostTnNicEvent	Warning	Transmitted packets are dropping on NSX-T Host Node's network interface	Transmitted packets are getting dropped on NSX-T Host Node's network interface. This may impact the network traffic on ESXi Host.
1.3.6.1.4.1.6876.100.1.0.80402	NSXTMPNodeServiceCmInventoryStatusEvent	Warning	CM Inventory Service has stopped running	CM Inventory Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80403	NSXTMPNodeServiceControllerStatusEvent	Warning	Controller Service has stopped running.	Controller Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80404	NSXTMPNodeServiceDataStoreStatusEvent	Warning	DataStore Service has stopped running.	DataStore Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80405	NSXTMPNodeServiceHttpStatusEvent	Warning	HTTP Service has stopped running.	HTTP Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80406	NSXTMPNodeServiceInstallUpgradeEvent	Warning	Install Upgrade Service has stopped running.	Install Upgrade Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80407	NSXTMPNodeServiceLiagentStatusEvent	Warning	Liagent service has stopped running.	Liagent Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80408	NSXTMPNodeServiceManagerStatusEvent	Warning	Manager Service has stopped running.	Manager Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80409	NSXTMPNodeServiceMgmtPlaneBusStatusEvent	Warning	Management Plane Service has stopped running.	Management Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80410	NSXTMPNodeServiceMigrationCoordinatorStatusEvent	Warning	Migration Coordinator Service has stopped running.	Migration Coordinator Service status has turned to stopped.

Table 16-2. vRealize Network Insight Computed NSX-T Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80411	NSXTMPNodeService NodeMgmtStatusEvent	Warning	Node Management Service has stopped running.	Node Management Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80412	NSXTMPNodeService NodeStatsStatusEvent	Warning	Node Statistics Service has stopped running.	Node Statistics Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80413	NSXTMPNodeService NSXMessageBusStatusEvent	Warning	Message Bus Service has stopped running.	Message Bus Client Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80414	NSXTMPNodeService NSXPlatformClientStatusEvent	Warning	Platform Client Service has stopped running.	Platform Client Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80415	NSXTMPNodeService NSXUpgradeAgentStatusEvent	Warning	Upgrade Agent Service has stopped running.	Upgrade Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80416	NSXTMPNodeService NTPStatusEvent	Warning	NTP Service has stopped running.	NTP Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80417	NSXTMPNodeService PolicyStatusEvent	Warning	Policy Service has stopped running.	Policy Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80418	NSXTMPNodeService SearchStatusEvent	Warning	Search Service has stopped running.	Search Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80419	NSXTMPNodeService SNMPStatusEvent	Warning	SNMP Service has stopped running.	SNMP Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80420	NSXTMPNodeService SSHStatusEvent	Warning	SSH Service has stopped running.	SSH Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80421	NSXTMPNodeService SyslogStatusEvent	Warning	Syslog Service has stopped running.	Syslog Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80422	NSXTMPNodeService TelemetryStatusEvent	Warning	Telemetry Service has stopped running.	Telemetry Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80423	NSXTMPNodeService UIServiceStatusEvent	Warning	UI Service has stopped running.	UI Service status has turned to stopped.

Table 16-2. vRealize Network Insight Computed NSX-T Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0 .80402	NSXTMPNodeService CmlInventoryStatusEvent	Critical	CM Inventory Service has stopped	One of the Services of the NSX-T Management Node, namely CM Inventory Service has stopped running.
1.3.6.1.4.1.6876.100.1.0 .80403	NSXTMPNodeService ControllerStatusEvent	Critical	Controller Service has stopped	One of the Services of the NSX-T Management Node, namely Controller Service has stopped running.
1.3.6.1.4.1.6876.100.1.0 .80404	NSXTMPNodeService DataStoreStatusEvent	Critical	DataStore Service has stopped	One of the Services of the NSX-T Management Node, namely DataStore Service has stopped running.
1.3.6.1.4.1.6876.100.1.0 .80405	NSXTMPNodeService HttpStatusEvent	Critical	HTTP Service has stopped	One of the Services of the NSX-T Management Node, namely HTTP Service has stopped running.
1.3.6.1.4.1.6876.100.1.0 .80406	NSXTMPNodeService InstallUpgradeEvent	Warning	Install Upgrade Service has stopped	One of the Services of the NSX-T Management Node, namely Install Upgrade Service has stopped running.
1.3.6.1.4.1.6876.100.1.0 .80407	NSXTMPNodeService LiagentStatusEvent	Warning	Liagent service has stopped	One of the Services of the NSX-T Management Node, namely LI Agent Service has stopped running.
1.3.6.1.4.1.6876.100.1.0 .80408	NSXTMPNodeService ManagerStatusEvent	Critical	Manager Service has stopped	One of the Services of the NSX-T Management Node, namely Manager Service has stopped running.
1.3.6.1.4.1.6876.100.1.0 .80409	NSXTMPNodeService MgmtPlaneBusStatus Event	Warning	Management Plane Service has stopped	One of the Services of the NSX-T Management Node, namely Management Plane Bus Service has stopped running.

Table 16-2. vRealize Network Insight Computed NSX-T Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80410	NSXTMPNodeService MigrationCoordinator StatusEvent	Warning	Migration Co- ordinator Service has stopped	One of the Services of the NSX-T Management Node, namely Migration Co- ordinator Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80411	NSXTMPNodeService NodeMgmtStatusEve nt	Critical	Node Management Service has stopped	One of the Services of the NSX-T Management Node, namely Node Management Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80412	NSXTMPNodeService NodeStatsStatusEven t	Critical	Node Statistics Service has stopped	One of the Services of the NSX-T Management Node, namely Node Statistics has stopped running.
1.3.6.1.4.1.6876.100.1.0.80413	NSXTMPNodeService NSXMessageBusStat usEvent	Warning	Message Bus Service has stopped	One of the Services of the NSX-T Management Node, namely Message Bus Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80414	NSXTMPNodeService NSXPlatformClientSta tusEvent	Critical	Platform Client Service has stopped	One of the Services of the NSX-T Management Node, namely Platform Client Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80415	NSXTMPNodeService NSXUpgradeAgentSt atusEvent	Warning	Upgrade Agent Service has stopped	One of the Services of the NSX-T Management Node, namely Upgrade Agent Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80416	NSXTMPNodeService NTPStatusEvent	Critical	NTP Service has stopped	One of the Services of the NSX-T Management Node, namely NTP Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80417	NSXTMPNodeService PolicyStatusEvent	Critical	Policy Service has stopped	One of the Services of the NSX-T Management Node, namely Policy Service has stopped running.

Table 16-2. vRealize Network Insight Computed NSX-T Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80418	NSXTMPNodeService SearchStatusEvent	Critical	Search Service has stopped	One of the Services of the NSX-T Management Node, namely Search Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80419	NSXTMPNodeService SNMPStatusEvent	Warning	SNMP Service has stopped	One of the Services of the NSX-T Management Node, namely SNMP Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80420	NSXTMPNodeService SSHStatusEvent	Critical	SSH Service has stopped	One of the Services of the NSX-T Management Node, namely SSH Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80421	NSXTMPNodeService SyslogStatusEvent	Critical	Syslog Service has stopped	One of the Services of the NSX-T Management Node, namely Syslog Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80422	NSXTMPNodeService TelemetryStatusEvent	Warning	Telemetry Service has stopped	One of the Services of the NSX-T Management Node, namely Telemetry Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80423	NSXTMPNodeService UIServiceStatusEvent	Critical	UI Service has stopped	One of the Services of the NSX-T Management Node, namely UI Service has stopped running.
1.3.6.1.4.1.6876.100.1.0.80424	NSXTMPNodeService ClusterManagerStatusEvent	Critical	Cluster Manager Service has stopped	One of the Services of the NSX-T Management Node, namely Cluster Manager Service has stopped running.

Network Monitoring

17

The topics in this section show you how to configure monitoring using Internet Protocol Flow Information Export (IPFIX) profiles for the firewall and switches, as well as how to configure an IPFIX collector.

Read the following topics next:

- [Add an IPFIX Collector](#)
- [Add a Firewall IPFIX Profile](#)
- [Add a Switch IPFIX Profile](#)
- [IPFIX Monitoring on a vSphere Distributed Switch](#)
- [Add a Port Mirroring Profile](#)
- [Port Mirroring on a vSphere Distributed Switch](#)
- [Perform a Traceflow](#)
- [Simple Network Management Protocol \(SNMP\)](#)
- [Monitor Fabric Nodes](#)
- [Network Latency Statistics](#)
- [Monitoring Tools in Manager Mode](#)

Add an IPFIX Collector

You can configure IPFIX collectors for firewalls and switches.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Plan & Troubleshoot > IPFIX**.
- 3 Click the **Collectors** tab.
- 4 Select **Add New Collector > IPFIX Switch** or **Add New Collector > IPFIX Firewall**.
- 5 Enter a name.

- 6 Enter the IP address and port of up to four collectors. Both IPv4 and IPv6 addresses are supported.
- 7 Click **Save**.

Add a Firewall IPFIX Profile

You can configure IPFIX profiles for firewalls.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Plan & Troubleshoot > IPFIX**.
- 3 Click the **Firewall IPFIX Profiles** tab.
- 4 Click **Add Firewall IPFIX Profile**.
- 5 Complete the following details.

Setting	Description
Name and Description	Enter a name and optionally a description. Note If you want to create a global profile, name the profile Global . A global profile cannot be edited or deleted from the UI, but you can do so using NSX-T Data Center APIs.
Active Flow Export Timeout (Minutes)	The length of time after which a flow will time out, even if more packets associated with the flow are received. Default is 1.
Observation Domain ID	This parameter identifies which observation domain the network flows originate from. The default is 0 and indicates no specific observation domain.
Collector Configuration	Select a collector from the drop-down menu.
Applied To	Click Set and select a group to apply the filter to, or create a new group.
Priority	This parameter resolves conflicts when multiple profiles apply. The IPFIX exporter will use the profile with the highest priority only. A lower value means a higher priority.

- 6 Click **Save** and then **Yes** to continue configuring the profile.
- 7 Click **Save**.

Add a Switch IPFIX Profile

You can configure IPFIX profiles for switches, also known as segments.

Flow-based network monitoring enable network administrators to gain insight into traffic traversing a network.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Plan & Troubleshoot > IPFIX**.
- 3 Click the **Switch IPFIX Profiles** tab.
- 4 Click **Add Switch IPFIX Profile**.
- 5 Enter the following details:

Setting	Description
Name and Description	Enter a name and optionally a description. Note If you want to create a global profile, name the profile Global . A global profile cannot be edited or deleted from the UI, but you can do so using NSX-T Data Center APIs.
Active Timeout (seconds)	The length of time after which a flow times out, even if more packets associated with the flow are received. Default is 300.
Idle Timeout (seconds)	The length of time after which a flow times out, if no more packets associated with the flow are received (ESXi only, KVM times out all flows based on the active timeout). Default is 300.
Packet Sampling Probability (%)	The percentage of packets that will be sampled (approximately). Increasing this setting can have a performance impact on the hypervisors and collectors. If all hypervisors are sending more IPFIX packets to the collector, the collector might not be able to collect all packets. Setting the probability at the default value of 0.1% keeps the performance impact low.
Collector Configuration	Select a collector from the drop-down menu .
Applied To	Select a category: Segment, Segment Port, or Groups. The IPFIX profile is applied to the selected object.
Priority	This parameter resolves conflicts when multiple profiles apply. The IPFIX exporter uses the profile with the highest priority only. A lower value means a higher priority.
Max Flows	The maximum flows cached on a bridge (KVM only, not configurable on ESXi). Default is 16384.
Observation Domain ID	The observation domain ID identifies which observation domain the network flows originate from. Enter 0 to indicate no specific observation domain.
Export Overlay Flow	This parameter defines whether to sample and export the overlay flows on uplink and tunnel ports. Both the vNIC flow and overlay flow are included in the sample. The default is enabled . When disabled, only vNIC flows are sampled and exported.
Tags	Enter a tag to make searching easier.

- 6 Click **Save** and then **Yes** to continue configuring the profile.
- 7 Click **Applied To** to apply the profile to objects.
Select one or more of the objects.
- 8 Click **Save**.

IPFIX Monitoring on a vSphere Distributed Switch

Configure IPFIX monitoring for NSX Distributed Virtual port groups, and vSphere Distributed Virtual port groups that are connected to a VDS switch enabled to support NSX-T networking.

From vSphere, enable IPFIX for Distributed Virtual port groups (vSphere) and from NSX Manager, enable IPFIX for segments (NSX-T) created on a VDS switch.

To enable IPFIX monitoring for Distributed Virtual port groups, see the *vSphere Networking* documentation.

To enable IPFIX monitoring for NSX-T port groups, see [Add a Switch IPFIX Profile](#).

A VDS switch enabled for NSX-T displays the following behavior:

- Both non-uplink and uplink ports support bidirectional traffic incoming and outgoing on:
 - Ports, port groups, VMs on vSphere.
 - Segments, segment ports, and groups on NSX-T
- IPFIX profile samples packets on uplink ports when are coming from or going to non-uplink ports that are IPFIX enabled. For example, consider that *VM-A* and *VM-B* are connected to non-uplink ports (port-1, port-2), where port-1 connected to *VM-A* is IPFIX enabled, and port-2 connected to *VM-B* is not IPFIX enabled. When you send traffic from *VM-A* and *VM-B* traffic to port-1, only packets from *VM-A* are sampled because IPFIX is enabled only on the port that *VM-A* is connected to. IPFIX does not sample packets coming from the port-2 associated to *VM-B* because IPFIX is not enabled on that port.
- Packet count exported to the IPFIX collector is the total count based on a sampling rate, not the sampled packets. For example, IPFIX calculates the count of total packets and exports the info. For 100 incoming packets, IPFIX might sample 9–11 packets. It exports 90 or 110 packets to the IPFIX collector.

Add a Port Mirroring Profile

You can configure port mirroring profiles for port mirroring sessions.

Note that logical SPAN is supported for overlay segments only and not VLAN segments.

Note Port Mirroring is not recommended for monitoring because when used for longer durations performance is impacted.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Plan & Troubleshoot > Port Mirroring**
- 3 Select **Add Profile > Remote L3 Span** or **Add Profile > Logical Span**.
- 4 Enter a name and optionally a description.

5 Complete the following profile details.

Session Type	Parameters
Remote L3 SPAN	■ Direction - Select Bidirectional, Ingress, or Egress. ■ Snap Length - Specify the number of bytes to capture from a packet. ■ Encapsulation Type - Select GRE, ERSPAN TWO, or ERSPAN THREE. ■ GRE Key - Specify a GRE key if encapsulation type is GRE. ■ ERSPAN ID - Specify an ERSPAN ID if encapsulation type is ERSPAN TWO or ERSPAN THREE.
Logical SPAN	■ Direction - Select Bidirectional, Ingress, or Egress. ■ Snap Length - Specify the number of bytes to capture from a packet.

6 Click **Set** in the **Source** column to set a source.

For Logical SPAN, the available sources are **Segment Port**, **Group of Virtual Machines**, and **Group of Virtual Network Interfaces**.

For Remote L3 SPAN, the available sources are **Segment**, **Segment Port**, **Group of Virtual Machines**, and **Group of Virtual Network Interfaces**.

7 Click **Set** in the **Destination** column to set a destination.

8 Click **Save**.

Port Mirroring on a vSphere Distributed Switch

You can configure port mirroring for port groups, virtual NICs of VMs, and VMs created in NSX-T and vSphere Distributed Virtual port groups created in vSphere that are connected to a vSphere Distributed Switch (VDS) switch.

In vCenter Server, configure port mirroring for vSphere Distributed Virtual port groups on a VDS switch.

In NSX Manager, configure port mirroring for segments (in NSX-T) on a VDS switch.

Note You cannot edit configuration for a segment created in NSX-T in vCenter Server. As an admin, you can view the properties of a port mirroring session to know on which switch it is created.

To enable port mirroring on vSphere Distributed Virtual port groups, see the *vSphere Networking* documentation.

To enable port mirroring on segments, ports, groups in NSX-T from both Policy and Manager modes in NSX Manager, see:

- [Add a Port Mirroring Profile](#)
- [Monitor Port Mirroring Sessions in Manager Mode](#)

Uplink Conflict Between Teaming and Remote SPAN

In vSphere, by default, the **Remote SPAN** in a teaming policy is set to `Disallowed`. If you use all the available physical NICs to configure remote SPAN, there are no free uplinks available for the teaming policy to consume. The unavailability of any free uplink means that uplink traffic is not allowed on destination ports, resulting in configuration errors.

However, in NSX-T, by default, **Normal I/O on Destination Ports** is set to `Allowed`. In NSX-T, port mirroring configured for NSX port groups on an N-VDS switch allows teaming and port mirroring on destination ports. So, uplink configuration errors do not occur in NSX-T.

To resolve a uplink conflict when configuring teaming and remote SPAN:

- Ensure that a free uplink is available. For example, on an ESXi host with 2 physical NICs, do not assign both these uplinks as destination IP addresses in the remote span port mirroring profile to avoid uplink conflicts in configuration. There must be at least one available uplink that can be configured in teaming profile.
- In vCenter Server, edit the port mirror configuration profile and set the **Normal I/O on Destination Ports** to `Allowed`.

Perform a Traceflow

Use Traceflow to inspect the path of a packet. Traceflow traces the transport node-level path of a packet. The trace packet traverses the logical switch overlay, but is not visible to interfaces attached to the logical switch. In other words, no packet is actually delivered to the test packet's intended recipients.

Note Traceflow is not supported for a VLAN-backed logical switch or segment.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Plan & Troubleshoot > Troubleshooting Tools > Traceflow**.
- 3 Select an IPv4 or IPv6 address type.
- 4 Select a traffic type.

For IPv4 addresses the traffic type choices are Unicast, Multicast, and Broadcast. For IPv6 address the traffic type choices are Unicast or Multicast.

Note: Multicast and broadcast are not supported in a VMware Cloud (VMC) environment.

5 (Optional) Select a protocol and provide related information.

Protocol	Parameters
DHCP	Select a DHCP OP code: Boot Request or Boot Reply .
DHCPv6	Select a DHCP message type: Solicit , Advertise , Request , or Reply . Note This option is available only when IPv6 is selected for IP Address.
DNS	Specify an address and select a message type: Query or Response .
ICMP	Specify an ICMP ID and a sequence.
ICMPv6	Specify an ICMP ID and a sequence. Note This option is available only when IPv6 is selected for IP Address.
TCP	Specify a source port, a destination port, and TCP flags.
UDP	Specify a source port and a destination port.

For the TCP protocol, note the following:

- The default flag is SYN.
- SYN cannot be combined with RST or FIN.
- If SYN is not selected, you must select ACK or RST.
- ACK cannot be combined with FIN, PSH, or URG.

6 Specify the source and destination information according to the traffic type.

Traffic Type	Source	Destination
Unicast	Select a VM or a logical port. For a VM: ■ Select a VM from the drop-down list. ■ Select a virtual interface. ■ The IP address and MAC address are displayed if VMtools is installed in the VM, or if the VM is deployed using OpenStack plug-in (address bindings will be used in this case). If the VM has more than one IP address, select one from the drop-down list. ■ If the IP address and MAC address are not displayed, enter the IP address and MAC address in the text boxes. For a logical port: ■ Select an attachment type: VIF, DHCP, Edge Uplink, or Edge Centralized Service. ■ Select a port.	Select a VM, a logical port, or IP-MAC. For a VM: ■ Select a VM from the drop-down list. ■ Select a virtual interface. ■ The IP address and MAC address are displayed if VMtools is installed in the VM or if the VM is deployed using OpenStack plug-in (address bindings will be used in this case). If the VM has more than one IP address, select one from the drop-down list. ■ If the IP address and MAC address are not displayed, enter the IP address and MAC address in the text boxes. For a logical port: ■ Select an attachment type: VIF, DHCP, Edge Uplink, or Edge Centralized Service. ■ Select a port. For IP-MAC: ■ Select the trace type (layer 2 or layer 3). For layer 2, enter an IP address and a MAC address. For layer 3, enter an IP address.
Multicast	Same as above.	Enter an IP Address. It must be a multicast address from 224.0.0.0 - 239.255.255.255.
Broadcast	Same as above.	Enter a subnet prefix length.

7 (Optional) Click **Advanced Settings** to see the advanced options.

In the left column, enter the desired values or input for the following fields:

Option	Description
Frame Size	The default is 128.
TTL	The default is 64.
Timeout (ms)	The default is 10000.
Ethertype	The default is 2048.
Payload Type	Select Base64, Hex, Plaintext, Binary, or Decimal.
Payload Data	Payload formatted based on selected type.

8 Click **Trace**.

The output includes a graphical map of the topology and a table listing the observed packets. The first packet listed has the observation type *Injected* and shows the packet that is injected at the injection point.

You can apply a filter (**All**, **Delivered**, **Dropped**) on the observations that are displayed. If there are dropped observations, the **Dropped** filter is applied by default. Otherwise, the **All** filter is applied.

The graphical map shows the backplane and router links. Note that bridging information is not displayed.

Simple Network Management Protocol (SNMP)

You can use Simple Network Management Protocol (SNMP) to monitor your NSX-T Data Center components. The SNMP service is not started by default after installation.

The SNMP Framework in NSX-T Data Center enables you to monitor various system entities (such as disk on NSX Edge) and logical entities (such as NSX Edge VPN tunnel) using their SNMP managers. This framework enables NSX-T Data Center verticals and platform to define SNMP MIB objects to be monitored and which can be used to enable their SMNP managers to interact with NSX-T Data Center.

To download the SNMP MIB files, see [Knowledge Base article 1013445: SNMP MIB module file download](#). Download and use the file **VMWARE-NSX-MIB.mib**.

For SNMP configuration, see *Configure SNMP for ESXi* in the VMware vSphere product documentation.

Procedure

- 1 Log in to the NSX Manager CLI or the NSX Edge CLI.
- 2 Run the following commands

- For SNMPv1/SNMPv2:

```
set snmp community <community-string>
start service snmp
```

The maximum character limit for **community-string** is 64.

- For SNMPv3

```
set snmp v3-users <user_name> auth-password <auth_password> priv-password
<priv_password>

start service snmp
```

The maximum character limit for **user_name** is 32. Ensure that your passwords meet PAM constraints. If you want to change the default engine id, use the following command:

```
set snmp v3-engine-id <v3-engine-id>

start service snmp
```

v3-engine-id is a hexadecimal string that is 10 to 64 characters long.

NSX-T Data Center supports SHA1 and AES128 as the authentication and privacy protocols. You can also use API calls to set up SNMPv3. For more information, see the *NSX-T Data Center API Guide*.

Monitor Fabric Nodes

You can monitor fabric nodes such as hosts, edges, NSX Edge clusters, bridges, and transport nodes from the NSX Manager UI.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **System > Fabric > Nodes** from the navigation panel.
- 3 Select one of the following tabs.
 - Hosts
 - Edges
 - Edge Clusters
 - Bridges
 - Transport Nodes

Results

Note On the Hosts screen, if the MPA Connectivity status is Down or Unknown for a host, ignore the LCP Connectivity status because it might be inaccurate.

Network Latency Statistics

In a network, latency can accumulate at multiple endpoints in the data path. As a network administrator, you need the ability to monitor the latency of a network to diagnose and troubleshoot performance bottlenecks in the network.

The following network latency statistics can be measured on host transport nodes:

- pNIC to vNIC
- vNIC to pNIC
- vNIC to vNIC
- VTEP to VTEP

In NSX-T Data Center, the following limitations apply to measuring latency statistics:

- Only ESXi host transport nodes are supported for measuring network latency in the data plane.

- KVM hosts and edge transport nodes are not supported.
- On VLAN segments, network latency is measured only when the two vNICs belong to VMs on the same ESXi host.
- When the VMs are attached to separate segments, network latency is measured only when the data traffic is routed through the distributed router (DR) instance on the ESXi host transport nodes. If the data traffic is routed through the DR instance on the edge transport nodes, network latency is not measured.
- Enhanced networking stack (ENS) does not support vNIC to pNIC, pNIC to vNIC, and vNIC to vNIC latency.
- Latency measurement is not supported when an east-west network traffic protection is configured using partner service VMs. Latency monitoring is disabled on the ports of service virtual machine (SVMs) and guest VMs.

You can export the latency data to external network performance monitoring tools and run analytics on the data. The external monitoring tools are also called collectors. By using a collector, you can achieve greater network visibility, optimize network performance, and identify the endpoints in the data path that cause a significant latency in the network.

After the hosts are configured to measure network latency statistics, the network operations agent (netopa) on the hosts periodically polls the data plane. When latency data is available, the agent exports the data at preconfigured intervals to the external collectors.

Note

- The netopa agent can export the network latency statistics only to vRealize Network Insight (vRNI). Other collector tools are not supported currently.
- You can configure ESXi hosts to measure network latency statistics only by using the NSX REST APIs.

The following support matrix summarizes the transport nodes and collectors that are supported for various network latency statistics.


Table 17-1. Support Matrix

Network Latency Statistics	Starting in NSX-T Data Center Version	Supported Transport Nodes	Supported Collectors	Comments
VTEP to VTEP	2.5	ESXi hosts	vRNI 5.0 or later	
pNIC to vNIC vNIC to pNIC vNIC to vNIC	3.0	ESXi hosts	vRNI 5.3	Support for exporting statistics to vRNI 5.3 is available starting in NSX-T Data Center 3.0.2.

You can measure network latency statistics for both standalone ESXi hosts and ESXi hosts that are a part of the vCenter Server cluster. However, network latency statistics from only vCenter-managed ESXi hosts can be exported to vRNI. vRNI does not support collecting latency statistics from standalone ESXi hosts that are not managed by a vCenter Server.

pNIC to vNIC and vNIC to pNIC Latency

When pNIC latency measurement is enabled on a host transport node, vNIC to pNIC latency and pNIC to vNIC latency are computed for each vNIC on the host transport node.



pNIC to vNIC and vNIC to pNIC latency statistics are exported to the external collector in the following format:


Endpoint1, Endpoint2, Max, Min, Avg

Where:

- *Endpoint1, Endpoint2* can either be the virtual interface ID (VIF ID) or the name of the physical adapter on an ESXi host (vmnic).
- *Max, Min, and Avg* represent the maximum, minimum, and average latency values between the two endpoints in microseconds.

vNIC to vNIC Latency

This latency represents the time taken by the data packet to travel from the source vNIC to the destination vNIC either on the same ESXi host or different ESXi hosts. If the vNICs are on different ESXi hosts, only GENEVE encapsulation protocol is supported in the overlay tunnel between the hosts.



vNIC to vNIC network latency is computed as follows:

- When the source vNIC1 on VM1 and the destination vNIC2 on VM2 are on the same host, a single-trip latency is calculated for each trip and exported to the collector. In other words, latency for each trip vNIC1 to vNIC2 and vNIC2 to vNIC1 is computed separately.
- When the source vNIC1 on VM1 and the destination vNIC2 on VM2 are on different hosts, total round-trip latency is calculated, and only a single latency value is exported to the collector. If there is no return traffic from vNIC2 to vNIC1, no network latency is exported to the collector.

Note NSX-T Data Center calculates the vNIC to vNIC latency between hosts directly by using the timestamps in the GENEVE encapsulated packets. You do not have to enable pNIC latency measurement on the host and the VTEP to VTEP latency. The pNIC to vNIC, vNIC to pNIC, and VTEP to VTEP statistics are independent of the vNIC to vNIC statistic.

vNIC to vNIC latency statistics are exported to the external collector in the following format:


VIF1, VIF2, Max, Min, Avg

Where:

- *VIF1, VIF2* represent the virtual interfaces or the vNICs.
- *Max, Min, and Avg* represent the maximum, minimum, and average latency values between the two vNICs in microseconds.

VTEP to VTEP Latency

This latency represents the total round-trip time taken by the data packet to travel from the source VTEP to the destination VTEP. To measure VTEP to VTEP latency, you must enable latency in the transport zone profile.



To calculate the VTEP to VTEP latency between ESXi hosts, Bidirectional Flow Detection (BFD) protocol is used. NSX-T Data Center extends the BFD protocol with timestamps in the payload to support latency computation between the VTEPs. BFD packets are transmitted at regular intervals in each tunnel between the hosts to compute the VTEP to VTEP latency.

What to read next

- [Measure Network Latency Statistics](#)

You can configure ESXi hosts in your network to measure: pNIC to vNIC, vNIC to pNIC, vNIC to vNIC, and VTEP to VTEP network latency statistics.

- [Export Network Latency Statistics](#)

You can export network latency statistics to external collectors and run analytics on the data. The netopa agent that is running in the ESXi hosts can export the network latency statistics only to vRealize Network Insight (vRNI). Other collector tools are not supported currently.

Measure Network Latency Statistics

You can configure ESXi hosts in your network to measure: pNIC to vNIC, vNIC to pNIC, vNIC to vNIC, and VTEP to VTEP network latency statistics.

Configuration is supported only using the NSX REST APIs. The steps in the following procedure list the management plane APIs that you must run to configure the calculation of various network latency statistics. For a detailed information about the API schema, example request, example response, and error messages of all the APIs, you must read the *NSX-T Data Center API Guide*.

Prerequisites

Both vCenter-managed hosts and standalone ESXi hosts that you want to configure for measuring network latency statistics must be prepared for NSX-T Data Center. That is, NSX-T Data Center components must be installed on all the ESXi hosts in your network.

Procedure

- 1 To compute vNIC to vNIC, pNIC to vNIC, and vNIC to pNIC network latency statistics, do these steps:

- a Create a latency profile with the following POST API:

```
POST https://<NSX-Manager-IP>/api/v1/latency-profiles
```

By default, vNIC to vNIC latency is measured for all the vNICs on the host transport node.

In the request body of this API, configure the following information:

- Enable or disable pNIC latency on the host. When it is enabled, pNIC to vNIC and vNIC to pNIC latency are calculated for each vNIC on the host transport node.
- Specify either the sampling rate or the sampling interval.

If you configure both, sampling interval takes a precedence.

- b Create a NSGroup with transport nodes as the target type in the NSGroupsSimpleExpression with the following POST API:

```
POST https://<NSX-Manager-IP>/api/v1/ns-groups
```

If you have **Manager** mode enabled in your UI, you can use the UI to create the NSGroup, and specify the transport nodes in the membership criteria.

- c Create a Service Configuration profile with the following POST API:

```
POST https://<NSX-Manager-IP>/api/v1/service-configs
```

This API combines the latency profile and the NSGroup that you created in the previous steps.

- 2 To measure VTEP to VTEP latency statistics, enable latency in the BFD health monitoring profile, which is a resource type in the transport zone profile. Run either the following PUT or the POST API:

- POST `https://<NSX-Manager-IP>/api/v1/transportzone-profiles`
- PUT `https://<NSX-Manager-IP>/api/v1/transportzone-profiles/<transportzone-profile-id>`

What to do next

Export the statistics to an external collector for a deeper network insight and troubleshooting network-specific latency problems.

Export Network Latency Statistics

You can export network latency statistics to external collectors and run analytics on the data. The netopa agent that is running in the ESXi hosts can export the network latency statistics only to vRealize Network Insight (vRNI). Other collector tools are not supported currently.

In vRNI, you can collect network latency statistics from only the vCenter-managed ESXi hosts. vRNI does not support collecting latency statistics from standalone ESXi hosts that are not managed by a vCenter Server.

You can export network latency statistics by using any one of the following methods:

- Method 1: Use the management plane APIs in NSX-T Data Center.
- Method 2: Enable an optional setting in the vRNI UI to collect latency statistics.

Prerequisites

- In the vRNI UI, complete the following tasks in the given order:
 - a Add vCenter Server as the data source. If you have multiple vCenter Servers added as Compute Managers in your NSX-T Data Center environment, you can add all vCenter Servers as the data source.
 - b Add NSX Manager as the data source.

For a detailed explanation about adding data sources in vRNI, see the *Using vRealize Network Insight* documentation at <https://docs.vmware.com/en/VMware-vRealize-Network-Insight/index.html>.

- Ensure that port 1991 is open on the collector to receive network latency data from the ESXi hosts.

Procedure

1 Method 1: Use the NSX-T Data Center REST APIs.

- a Ensure that you have configured the ESXi hosts to measure network latency statistics. For detailed steps, see [Measure Network Latency Statistics](#).

- b Export the network latency statistics to the collector with the following PUT API:

```
PUT https://<manager-ip>/api/v1/global-configs/OperationCollectorGlobalConfig
-d '<content>'
```

In the request body of this API, configure the following information:

- Details of external collectors, such as collector IP address, collector port.
- Report interval that controls the frequency at which the netopa agent sends statistics to the collector.

2 Method 2: Enable an optional setting in the vRNI UI to collect latency statistics.

When you add NSX Manager as the data source in vRNI, select the **Enable latency metric collection** check box. This option enables vRNI to collect latency statistics from the ESXi hosts.

For a detailed information about adding NSX Manager as the data source in vRNI, see the *Using vRealize Network Insight* documentation.

Results

vNIC to vNIC latency statistics are exported to the external collector in the following format:

VIF1, VIF2, Max, Min, Avg

Where:

- *VIF1, VIF2* represent the virtual interfaces or the vNICs.
- *Max, Min, and Avg* represent the maximum, minimum, and average time between the two vNICs in microseconds.

pNIC to vNIC and vNIC to pNIC latency statistics are exported to the external collector in the following format:

Endpoint1, Endpoint2, Max, Min, Avg

Where:

- *Endpoint1, Endpoint2* can either be the virtual interface ID (VIF ID) or the name of the physical adapter on an ESXi host (vmnic).
- *Max, Min, and Avg* represent the maximum, minimum, and average time between the two endpoints in microseconds.

Monitoring Tools in Manager Mode

NSX-T support monitoring methods in **Manager** mode, including viewing port connections, traceflow, port mirroring, and activity monitoring.

View Port Connection Information in Manager Mode

You can use the port connection tool to quickly visualize and troubleshoot the connection between two VMs.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Plan and Troubleshoot > Port Connection** from the navigation panel.
- 3 Select a VM from the **Source Virtual Machine** drop-down menu.
- 4 Select a VM from the **Destination Virtual Machine** drop-down menu.

5 Click **Go**.

A visual map of the port connection topology is displayed. You can click on any of the components in the visual output to reveal more information about that component.


Traceflow

Traceflow allows you to inject a packet into the network and monitor its flow across the network. This flow allows you to monitor your network and identify issues such as bottlenecks or disruptions.

Traceflow allows you to identify the path (or paths) a packet takes to reach its destination or, conversely, where a packet is dropped along the way. Each entity reports the packet handling on input and output, so you can determine whether issues occur when receiving a packet or when forwarding the packet.

The NSX Manager interface graphically displays the trace route based on the parameters you set (IP address type, traffic type, source, and destination). This display page also enables you to edit the parameters, retrace the traceflow, or create a new one.

Figure 17-1. Sample traceflow diagram



What is Traceflow?

Traceflow is not the same as a ping request/response that goes from guest-VM stack to guest-VM stack. Traceflow observes a marked packet as it traverses the overlay network, and each packet is monitored as it crosses the overlay network until it reaches a destination guest VM or an Edge uplink. Note that the injected marked packet is never actually delivered to the destination guest VM.

Traceflow can be used on transport nodes and supports both IPV4 and IPv6 protocols including: ICMP, TCP, UDP, DHCP, DNS and ARP/NDP.

Traceflow Parameters

You can construct packets with custom header fields and packet sizes. The source or destination for the traceflow can be a logical switch port, logical router uplink port, CSP or DHCP port. The destination endpoint can be any device in the NSX-T Data Center overlay or in the underlay. However, you cannot select a destination that is north of an NSX Edge node. The destination must be on the same subnet, or must be reachable through NSX-T Data Center distributed logical routers.

If NSX-T Data Center bridging is configured, packets with unknown destination MAC addresses are always sent to the bridge. Typically, the bridge forwards these packets to a VLAN and reports the traceflow packet as delivered. A packet reported as delivered does not necessarily mean that the trace packet was delivered to the specified destination.

For a unicast traceflow packet, you can observe packet replication and/or flooding in traceflow observations.

- A traceflow packet is replicated if the logical switch does not know the TEP(s) to which the packet is destined.
- A traceflow packet is flooded if N-VDS or VDS does not know the virtual switch port(s) to which the packet is destined.

You can specify multicast and broadcast packets as traceflow packets.

- For multicast traffic, the source is a VM vNIC or a logical port, and the destination is a multicast IP address.
- For broadcast traffic, the source is a VM vNIC or a logical port, and the Layer 2 destination MAC address is FF:FF:FF:FF:FF:FF.

To create a valid packet for firewall inspection, the broadcast traceflow operation requires a subnet prefix length. The subnet mask enables NSX-T Data Center to calculate an IP network address for the packet. A multicast or broadcast traceflow packet can be delivered to multiple VM vNICs or Edge uplinks, resulting in the generation of multiple delivered observations.

Trace the Path of a Packet with Traceflow in Manager Mode

Use Traceflow to inspect the path of a packet. Traceflow traces the transport node-level path of a packet. The trace packet traverses the logical switch overlay, but is not visible to interfaces attached to the logical switch. In other words, no packet is actually delivered to the test packet's intended recipients.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Plan & Troubleshoot > Traceflow**.
- 3 Select an IPv4 or IPv6 address type.
- 4 Select a traffic type.

For IPv4 addresses the traffic type choices are Unicast, Multicast, and Broadcast. For IPv6 address the traffic type choices are Unicast or Multicast.

Note: Multicast and broadcast are not supported in a VMware Cloud environment.

- 5 Specify the source and destination information according to the traffic type.

Traffic Type	Source	Destination
Unicast	Select a VM or a logical port. For a VM: ■ Select a VM from the drop-down list. ■ Select a virtual interface. ■ The IP address and MAC address are displayed if VMtools is installed in the VM, or if the VM is deployed using OpenStack plug-in (address bindings will be used in this case). If the VM has more than one IP address, select one from the drop-down list. ■ If the IP address and MAC address are not displayed, enter the IP address and MAC address in the text boxes. For a logical port: ■ Select an attachment type: VIF, DHCP, Edge Uplink, or Edge Centralized Service. ■ Select a port.	Select a VM, a logical port, or IP-MAC. For a VM: ■ Select a VM from the drop-down list. ■ Select a virtual interface. ■ The IP address and MAC address are displayed if VMtools is installed in the VM or if the VM is deployed using OpenStack plug-in (address bindings will be used in this case). If the VM has more than one IP address, select one from the drop-down list. ■ If the IP address and MAC address are not displayed, enter the IP address and MAC address in the text boxes. For a logical port: ■ Select an attachment type: VIF, DHCP, Edge Uplink, or Edge Centralized Service. ■ Select a port. For IP-MAC: ■ Select the trace type (layer 2 or layer 3). For layer 2, enter an IP address and a MAC address. For layer 3, enter an IP address.
Multicast	Same as above.	Enter an IP Address. It must be a multicast address from 224.0.0.0 - 239.255.255.255.
Broadcast	Same as above.	Enter a subnet prefix length.

- 6 (Optional) Click **Advanced** to see the advanced options.

- 7 (Optional) In the left column, enter the desired values or input for the following fields:

Option	Description
Frame Size	The default is 128.
TTL	The default is 64.
Timeout (ms)	The default is 10000.
Ethertype	The default is 2048.
Payload Type	Select Base64 , Hex , Plaintext , Binary , or Decimal .
Payload Data	Payload formatted based on selected type.

- 8 (Optional) Select a protocol and provide related information.

Protocol	Parameters
TCP	Specify a source port, a destination port, and TCP flags.
UDP	Specify a source port and a destination port.
ICMPv6	Specify an ICMP ID and a sequence.
ICMP	Specify an ICMP ID and a sequence.
DHCPv6	Select a DHCP message type: Solicit , Advertise , Request , or Reply .
DHCP	Select a DHCP OP code: Boot Request or Boot Reply .
DNS	Specify an address and select a message type: Query or Response .

- 9 Click **Trace**.

Information about the connections, components, and layers is displayed. The output includes a table listing Observation Type (Delivered, Dropped, Received, Forwarded), Transport Node, and Component, and a graphical map of the topology if unicast and logical switch as a destination are selected. You can apply a filter (**All**, **Delivered**, **Dropped**) on the observations that are displayed. If there are dropped observations, the **Dropped** filter is applied by default. Otherwise, the **All** filter is applied. The graphical map shows the backplane and router links. Note that bridging information is not displayed.

Monitor Port Mirroring Sessions in Manager Mode

You can monitor port mirroring sessions for troubleshooting and other purposes.

Note that logical SPAN is supported for overlay logical switches only and not VLAN logical switches.

NSX Cloud Note If using NSX Cloud, see [NSX-T Data Center Features Supported with NSX Cloud](#) for a list of auto-generated logical entities, supported features, and configurations required for NSX Cloud.

This feature has the following restrictions:

- A source mirror port cannot be in more than one mirror session.
- With KVM, multiple NICs can be attached to the same OVS port. The mirroring happens at the OVS uplink port, meaning that traffic on all the pNICs attached to the OVS port is mirrored.
- For a local SPAN session, the mirror session source and destination ports must be on the same host vSwitch. Therefore, if you vMotion the VM that has the source or destination port to another host, traffic on that port can no longer be mirrored.
- On ESXi, when mirroring is enabled on the uplink, raw production TCP packets are encapsulated using the Geneve protocol by VDL2 into UDP packets. A physical NIC that supports TSO (TCP segmentation offload) can change the packets and mark the packets with the MUST_TSO flag. On a monitor VM with VMXNET3 or E1000 vNICs, the driver treats the packets as regular UDP packets and cannot handle the MUST_TSO flag, and will drop the packets.

If a lot of traffic is mirrored to a monitor VM, there is a potential for the driver's buffer ring to become full and packets to be dropped. To alleviate the problem, you can take one or more of the following actions:

- Increase the rx buffer ring size.
- Assign more CPU resources to the VM.
- Use the Data Plane Development Kit (DPDK) to improve packet processing performance.

Note Make sure that the monitor VM's MTU setting (in the case of KVM, the hypervisor's virtual NIC device's MTU setting also) is large enough to handle the packets. This is especially important for encapsulated packets because encapsulation increases the size of packets. Otherwise, packets might be dropped. This is not an issue with ESXi VMs with VMXNET3 NICs, but is a potential issue with other types of NICs on both ESXi and KVM VMs.

Note In an L3 port mirroring session involving VMs on KVM hosts, you must set the MTU size to be large enough to handle the extra bytes required by encapsulation. The mirror traffic goes through an OVS interface and OVS uplink. You must set the OVS interface's MTU to be at least 100 bytes larger than the size of the original packet (before encapsulation and mirroring). If you see dropped packets, increase the MTU setting for the host's virtual NIC and the OVS interface. Use the following command to set the MTU for an OVS interface:

```
ovs-vsctl -- set interface <ovs_Interface> mtu_request=<MTU>
```

Note When you monitor the logical port of a VM and the uplink port of a host where the VM resides, you will see different behaviors depending on whether the host is ESXi or KVM. For ESXi, the logical-port mirror packets and the uplink mirror packets are tagged with the same VLAN ID and appear the same to the monitor VM. For KVM, the logical-port mirror packets are not tagged with a VLAN ID but the uplink mirror packets are tagged, and they appear different to the monitor VM.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

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

- 2 Select **Plan & Troubleshoot > Port Mirroring > Port Mirroring Session**.

- 3 Click **Add** and select a session type.

The available types are **Local SPAN**, **Remote SPAN**, **Remote L3 SPAN**, and **Logical SPAN**.

- 4 Enter a session name and optionally a description.

- 5 Provide additional parameters.

Session Type	Parameters
Local SPAN	■ Transport Node - Select a transport node. ■ Direction - Select Bidirectional, Ingress, or Egress. ■ Packet Truncation - Select a packet truncation value.
Remote SPAN	■ Session Type - Select RSPAN Source session or RSPAN Destination session. ■ Transport Node - Select a transport node. ■ Direction - Select Bidirectional, Ingress, or Egress. ■ Packet Truncation - Select a packet truncation value. ■ Encap. VLAN ID - Specify an encapsulation VLAN ID. ■ Preserve Orig. VLAN - Select whether to preserve the original VLAN ID.
Remote L3 SPAN	■ Encapsulation - Select GRE, ERSPAN TWO, or ERSPAN THREE. ■ GRE Key - Specify a GRE key if encapsulation is GRE. ERSPAN ID - Specify an ERSPAN ID if encapsulation is ERSPAN TWO or ERSPAN THREE. ■ Direction - Select Bidirectional, Ingress, or Egress. ■ Packet Truncation - Select a packet truncation value.
Logical SPAN	■ Logical Switch - Select a logical switch. ■ Direction - Select Bidirectional, Ingress, or Egress. ■ Packet Truncation - Select a packet truncation value.

- 6 Click **Next**.

7 Provide source information.

Session Type	Parameters
Local SPAN	■ Select an N-VDS. ■ Select physical interfaces. ■ Enable or disable encapsulated packet. ■ Select virtual machines. ■ Select virtual interfaces.
Remote SPAN	■ Select virtual machines. ■ Select virtual interfaces.
Remote L3 SPAN	■ Select virtual machines. ■ Select virtual interfaces. ■ Select a logical switch.
Logical SPAN	■ Select logical ports.

8 Click **Next**.

9 Provide destination information.

Session Type	Parameters
Local SPAN	■ Select virtual machines. ■ Select virtual interfaces.
Remote SPAN	■ Select an N-VDS. ■ Select physical interfaces.
Remote L3 SPAN	■ Specify an IPv4 address.
Logical SPAN	■ Select logical ports.

10 Click **Save**.

You cannot change the source or destination after saving the port mirroring session.

Configure Filters for a Port Mirroring Session

You can configure filters for port mirroring sessions to limit the amount of data that is mirrored.

This feature has the following capabilities and restrictions:

- Only ESXi and KVM host transport nodes are supported.
- IP address, IP prefix, and IP ranges are supported for source and destination.
- IPSet for source or destination is not supported.
- Mirror statistics on ESXi or KVM are not supported.

You must configure filters using the API. Using the NSX Manager UI is not supported. For more information about the port mirroring API and the `PortMirroringFilter` schema, see the *NSX-T Data Center API Reference*.

Procedure

- 1 Configure a port mirroring session using the NSX Manager UI or API.
- 2 Call the GET /api/v1/mirror-sessions API to get information about the port mirroring session.
- 3 Call the GET /api/v1/mirror-sessions/<mirror-session-id> API to add one or more filters. For example,

```
PUT https://<nsx-mgr>/api/v1/mirror-sessions/e57e8b2d-3047-4550-b230-dd1ee0e10b49
{
  "resource_type": "PortMirroringSession",
  "id": "e57e8b2d-3047-4550-b230-dd1ee0e10b49",
  "display_name": "port-mirror-session-1",
  "description": "Pnic port mirror session 1",
  "mirror_sources": [
    {
      "resource_type": "LogicalPortMirrorSource",
      "port_ids": [
        "6a361832-43e4-430d-a48a-b84a6cba73c3"
      ]
    }
  ],
  "mirror_destination": {
    "resource_type": "LogicalPortMirrorDestination",
    "port_ids": [
      "3e42e8b2d-3047-4550-b230-dd1ee0e10b34"
    ]
  },
  "port_mirroring_filters": [
    {
      "filter_action": "MIRROR",
      "src_ips": {
        "ip-addresses": [
          "192.168.175.250",
          "2001:bd6::c:2957:160:126"
        ]
      },
      "dst_ips": {
        "ip-addresses": [
          "192.168.160.126",
          "2001:bd6::c:2957:175:250"
        ]
      }
    }
  ],
  "session_type": "LogicalPortMirrorSession",
  "preserve_original_vlan": false,
  "direction": "BIDIRECTIONAL",
  "_revision": 0
}
```


- 4 (Optional) You can call the `get mirroring-session <session-number>` CLI command to show the properties of the port mirroring session, including the filters.

Configure IPFIX in Manager Mode

IPFIX (Internet Protocol Flow Information Export) is a standard for the format and export of network flow information. You can configure IPFIX for switches and firewalls. For switches, network flow at VIFs (virtual interfaces) and pNICs (physical NICs) is exported. For firewalls, network flow that is managed by the distributed firewall component is exported.

NSX Cloud Note If using NSX Cloud, see [NSX-T Data Center Features Supported with NSX Cloud](#) for a list of auto-generated logical entities, supported features, and configurations required for NSX Cloud.

This feature is compliant with the standards specified in RFC 7011 and RFC 7012.

When you enable IPFIX, all configured host transport nodes will send IPFIX messages to the IPFIX collectors using port 4739. In the case of ESXi, NSX-T Data Center automatically opens port 4739. In the case of KVM, if firewall is not enabled, port 4739 is open, but if firewall is enabled, you must ensure that the port is open because NSX-T Data Center does not automatically open the port.

IPFIX on ESXi and KVM sample tunnel packets in different ways. On ESXi the tunnel packet is sampled as two records:

- Outer packet record with some inner packet information
 - SrcAddr, DstAddr, SrcPort, DstPort, and Protocol refer to the outer packet.
 - Contains some enterprise entries to describe the inner packet.
- Inner packet record
 - SrcAddr, DstAddr, SrcPort, DstPort, and Protocol refer to the inner packet.

On KVM the tunnel packet is sampled as one record:

- Inner packet record with some outer tunnel information
 - SrcAddr, DstAddr, SrcPort, DstPort, and Protocol refer to the inner packet.
 - Contains some enterprise entries to describe the outer packet.

Configure Switch IPFIX Collectors in Manager Mode

You can configure IPFIX collectors for switches.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Plan & Troubleshoot > IPFIX**.
- 3 Click the **Switch IPFIX Collectors** tab.
- 4 Click **Add** to add a collector.
- 5 Enter a name and optionally a description.
- 6 Click **Add** and enter the IP address and port of a collector.
You can add up to 4 collectors.
- 7 Click **Add**.

Configure Switch IPFIX Profiles in Manager Mode

You can configure IPFIX profiles for switches.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Plan & Troubleshoot > IPFIX**.
- 3 Click the **Switch IPFIX Profiles** tab.
- 4 Click **Add** to add a profile.

Setting	Description
Name and Description	Enter a name and optionally a description. Note If you want to create a global profile, name the profile Global . A global profile cannot be edited or deleted from the UI, but you can do so using NSX-T Data Center APIs.
Active Timeout (seconds)	The length of time after which a flow will time out, even if more packets associated with the flow are received. Default is 300.
Idle Timeout (seconds)	The length of time after which a flow will time out, if no more packets associated with the flow are received (ESXi only, KVM times out all flows based on active timeout). Default is 300.
Max Flows	The maximum flows cached on a bridge (KVM only, not configurable on ESXi). Default is 16384.
Export Overlay Flow	Setting that controls whether the sample result includes overlay flow information.

Setting	Description
Sampling Probability (%)	The percentage of packets that will be sampled (approximately). Increasing this setting may have a performance impact on the hypervisors and collectors. If all hypervisors are sending more IPFIX packets to the collector, the collector may not be able to collect all packets. Setting the probability at the default value of 0.1% will keep the performance impact low.
Observation Domain ID	The observation domain ID identifies which observation domain the network flows originate from. Enter 0 to indicate no specific observation domain.
Collector Profile	Select a switch IPFIX collector that you configure in the previous step.
Priority	This parameter resolves conflicts when multiple profiles apply. The IPFIX exporter will use the profile with the highest priority only. A lower value means a higher priority.

5 Click **Add**.

Configure Firewall IPFIX Collectors in Manager Mode

You can configure IPFIX collectors for firewalls.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Plan & Troubleshoot > IPFIX**.
- 3 Click the **Firewall IPFIX Collectors** tab.
- 4 Click **Add** to add a collector.
- 5 Enter a name and optionally a description.
- 6 Click **Add** and enter the IP address and port of a collector.
You can add up to 4 collectors.
- 7 Click **Add**.

Configure Firewall IPFIX Profiles in Manager Mode

You can configure IPFIX profiles for firewalls.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Plan & Troubleshoot > IPFIX**.
- 3 Click the **Firewall IPFIX Profiles** tab.
- 4 Click **Add** to add a profile.

Setting	Description
Name and Description	Enter a name and optionally a description. Note If you want to create a global profile, name the profile Global . A global profile cannot be edited or deleted from the UI, but you can do so using NSX-T Data Center APIs.
Collector Configuration	Select a collector from the drop-down list.
Active Flow Export Timeout (Minutes)	The length of time after which a flow will time out, even if more packets associated with the flow are received. Default is 1.
Priority	This parameter resolves conflicts when multiple profiles apply. The IPFIX exporter will use the profile with the highest priority only. A lower value means a higher priority.
Observation Domain ID	This parameter identifies which observation domain the network flows originate from. The default is 0 and indicates no specific observation domain.

- 5 Click **Add**.

ESXi IPFIX Templates

An ESXi host transport node supports eight logical switch IPFIX flow templates and two distributed firewall IPFIX flow templates.

The following table lists VMware-specific elements in logical switch IPFIX packets.

Element ID	Parameter Name	Data Type	Unit
880	tenantProtocol	unsigned8	1 byte
881	tenantSourceIPv4	ipv4Address	4 bytes
882	tenantDestIPv4	ipv4Address	4 bytes
883	tenantSourceIPv6	ipv6Address	16 bytes
884	tenantDestIPv6	ipv6Address	16 bytes
886	tenantSourcePort	unsigned16	2 bytes
887	tenantDestPort	unsigned16	2 bytes
888	egressInterfaceAttr	unsigned16	2 bytes

Element ID	Parameter Name	Data Type	Unit
889	vxlanExportRole	unsigned8	1 byte
890	ingressInterfaceAttr	unsigned16	2 bytes
898	virtualObsID	string	variable length


The following table lists VMware-specific elements in distributed firewall IPFIX packets.

Element ID	Parameter Name	Data Type	Unit
950	ruleId	unsigned32	4 bytes
951	vmUuid	string	16 bytes
952	vnidIndex	unsigned32	4 bytes
953	sessionFlags	unsigned8	1 byte
954	flowDirection	unsigned8	1 byte
955	algControlFlowId	unsigned64	8 bytes
956	algType	unsigned8	1 byte
957	algFlowType	unsigned8	1 byte
958	averageLatency	unsigned32	4 bytes
959	retransmissionCount	unsigned32	4 bytes
960	vifUuid	octetArray	16 bytes
961	vifId	string	variable length

ESXi Logical Switch IPFIX Templates

An ESXi host transport node supports eight logical switch IPFIX flow templates.

The following diagram shows the flow of traffic between VMs attached to ESXi hosts monitored by the IPFIX feature:



The IPv4 Encapsulated template will have the following elements:

- standard elements
- SrcAddr: VTEP1
- DstAddr: VTEP2
- tenantSourceIPv4: IP1
- tenantDestIPv4: IP2
- tenantSourcePort: 10000
- tenantDestPort: 80
- tenantProtocol: TCP
- ingressInterfaceAttr: 0x03 (tunnel port)
- egressInterfaceAttr: 0x01
- encapExportRole: 01
- virtualObsID: 89fd5032-2dc9-4fc3-993a-9bb4b616de54 (logical port ID)

IPv4 Template

Template ID: 256

```
IPFIX_TEMPLATE_START(IPFIX_FLOW_TYPE_IPv4)
IPFIX_TEMPLATE_FIELD(sourceIPv4Address, 4)
IPFIX_TEMPLATE_FIELD(destinationIPv4Address, 4)
IPFIX_TEMPLATE_FIELD(octetDeltaCount, 8)
```

```

IPFIX_TEMPLATE_FIELD(packetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(flowStartSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(flowEndSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(sourceTransportPort, 2)
IPFIX_TEMPLATE_FIELD(destinationTransportPort, 2)
IPFIX_TEMPLATE_FIELD(ingressInterface, 4)
IPFIX_TEMPLATE_FIELD(egressInterface, 4)
IPFIX_TEMPLATE_FIELD(encapId, 8)
IPFIX_TEMPLATE_FIELD(protocolIdentifier, 1)
IPFIX_TEMPLATE_FIELD(flowEndReason, 1)
IPFIX_TEMPLATE_FIELD(tcpFlags, 1)
IPFIX_TEMPLATE_FIELD(IPv4TOS, 1)
IPFIX_TEMPLATE_FIELD(maxTTL, 1)
IPFIX_TEMPLATE_FIELD(flowDir, 1)
// Specify the Interface port- Uplink Port, Access port,N.A
IPFIX_VMW_TEMPLATE_FIELD(ingressInterfaceAttr, 2)
IPFIX_VMW_TEMPLATE_FIELD(egressInterfaceAttr, 2)
IPFIX_VMW_TEMPLATE_FIELD(encapExportRole, 1)
IPFIX_VMW_TEMPLATE_VAR_LEN_FIELD(virtualObsID, virtualObsDataLen)
IPFIX_TEMPLATE_PADDING(paddingOctets, 1)
IPFIX_TEMPLATE_END()

```

IPv4 Encapsulated Template

Template ID: 257

```

IPFIX_TEMPLATE_START(IPFIX_FLOW_TYPE_IPv4_ENCAP)
IPFIX_TEMPLATE_FIELD(sourceIPv4Address, 4)
IPFIX_TEMPLATE_FIELD(destinationIPv4Address, 4)
IPFIX_TEMPLATE_FIELD(octetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(packetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(flowStartSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(flowEndSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(sourceTransportPort, 2)
IPFIX_TEMPLATE_FIELD(destinationTransportPort, 2)
IPFIX_TEMPLATE_FIELD(ingressInterface, 4)
IPFIX_TEMPLATE_FIELD(egressInterface, 4)
IPFIX_TEMPLATE_FIELD(protocolIdentifier, 1)
IPFIX_TEMPLATE_FIELD(flowEndReason, 1)
IPFIX_TEMPLATE_FIELD(tcpFlags, 1)
IPFIX_TEMPLATE_FIELD(IPv4TOS, 1)
IPFIX_TEMPLATE_FIELD(maxTTL, 1)
IPFIX_TEMPLATE_FIELD(flowDir, 1)
IPFIX_TEMPLATE_FIELD(encapId, 8)
IPFIX_VMW_TEMPLATE_FIELD(tenantSourceIPv4, 4)
IPFIX_VMW_TEMPLATE_FIELD(tenantDestIPv4, 4)
IPFIX_VMW_TEMPLATE_FIELD(tenantSourcePort, 2)
IPFIX_VMW_TEMPLATE_FIELD(tenantDestPort, 2)
IPFIX_VMW_TEMPLATE_FIELD(tenantProtocol, 1)
// Specify the Interface port - Uplink Port, Access port, N.A
IPFIX_VMW_TEMPLATE_FIELD(ingressInterfaceAttr, 2)
IPFIX_VMW_TEMPLATE_FIELD(egressInterfaceAttr, 2)

```

```
// TUNNEL-GW or no.
IPFIX_VMW_TEMPLATE_FIELD(encapExportRole, 1)
IPFIX_VMW_TEMPLATE_VAR_LEN_FIELD(virtualObsID, virtualObsDataLen)
IPFIX_TEMPLATE_END()
```

IPv4 ICMP Template

Template ID: 258

```
IPFIX_TEMPLATE_START(IPFIX_FLOW_TYPE_IPv4_ICMP)
IPFIX_TEMPLATE_FIELD(sourceIPv4Address, 4)
IPFIX_TEMPLATE_FIELD(destinationIPv4Address, 4)
IPFIX_TEMPLATE_FIELD(octetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(packetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(flowStartSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(flowEndSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(ingressInterface, 4)
IPFIX_TEMPLATE_FIELD(egressInterface, 4)
IPFIX_TEMPLATE_FIELD(protocolIdentifier, 1)
IPFIX_TEMPLATE_FIELD(flowEndReason, 1)
IPFIX_TEMPLATE_FIELD(IPv4TOS, 1)
IPFIX_TEMPLATE_FIELD(maxTTL, 1)
IPFIX_TEMPLATE_FIELD(flowDir, 1)
IPFIX_TEMPLATE_FIELD(encapId, 8)
// Specify the Interface port - Uplink Port, Access Port, or NA.
IPFIX_VMW_TEMPLATE_FIELD(ingressInterfaceAttr, 2)
IPFIX_VMW_TEMPLATE_FIELD(egressInterfaceAttr, 2)
IPFIX_VMW_TEMPLATE_FIELD(encapExportRole, 1)
IPFIX_VMW_TEMPLATE_VAR_LEN_FIELD(virtualObsID, virtualObsDataLen)
IPFIX_TEMPLATE_PADDING(paddingOctets, 2)
IPFIX_TEMPLATE_END()
```

IPv4 ICMP Encapsulated Template

Template ID: 259

```
IPFIX_TEMPLATE_START(IPFIX_FLOW_TYPE_IPv4_ICMP_ENCAP)
IPFIX_TEMPLATE_FIELD(sourceIPv4Address, 4)
IPFIX_TEMPLATE_FIELD(destinationIPv4Address, 4)
IPFIX_TEMPLATE_FIELD(octetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(packetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(flowStartSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(flowEndSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(sourceTransportPort, 2)
IPFIX_TEMPLATE_FIELD(destinationTransportPort, 2)
IPFIX_TEMPLATE_FIELD(ingressInterface, 4)
IPFIX_TEMPLATE_FIELD(egressInterface, 4)
IPFIX_TEMPLATE_FIELD(protocolIdentifier, 1)
IPFIX_TEMPLATE_FIELD(flowEndReason, 1)
IPFIX_TEMPLATE_FIELD(IPv4TOS, 1)
IPFIX_TEMPLATE_FIELD(maxTTL, 1)
IPFIX_TEMPLATE_FIELD(flowDir, 1)
IPFIX_TEMPLATE_FIELD(encapId, 8)
IPFIX_VMW_TEMPLATE_FIELD(tenantSourceIPv4, 4)
IPFIX_VMW_TEMPLATE_FIELD(tenantDestIPv4, 4)
IPFIX_VMW_TEMPLATE_FIELD(tenantProtocol, 1)
```



```
// Specify the Interface port- Uplink Port, Access port,N.A
IPFIX_VMW_TEMPLATE_FIELD(ingressInterfaceAttr, 2)
IPFIX_VMW_TEMPLATE_FIELD(egressInterfaceAttr, 2)
// TUNNEL-GW or no.
IPFIX_VMW_TEMPLATE_FIELD(encapExportRole, 1)
IPFIX_VMW_TEMPLATE_VAR_LEN_FIELD(virtualObsID, virtualObsDataLen)
IPFIX_TEMPLATE_PADDING(paddingOctets, 1)
IPFIX_TEMPLATE_END()
```

IPv6 Template

Template ID: 260

```
IPFIX_TEMPLATE_START(IPFIX_FLOW_TYPE_IPv6)
IPFIX_TEMPLATE_FIELD(sourceIPv6Address, 16)
IPFIX_TEMPLATE_FIELD(destinationIPv6Address, 16)
IPFIX_TEMPLATE_FIELD(octetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(packetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(flowStartSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(flowEndSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(sourceTransportPort, 2)
IPFIX_TEMPLATE_FIELD(destinationTransportPort, 2)
IPFIX_TEMPLATE_FIELD(ingressInterface, 4)
IPFIX_TEMPLATE_FIELD(egressInterface, 4)
IPFIX_TEMPLATE_FIELD(encapId, 8)
IPFIX_TEMPLATE_FIELD(protocolIdentifier, 1)
IPFIX_TEMPLATE_FIELD(flowEndReason, 1)
IPFIX_TEMPLATE_FIELD(tcpFlags, 1)
IPFIX_TEMPLATE_FIELD(IPv6TOS,1)
IPFIX_TEMPLATE_FIELD(maxTTL, 1)
IPFIX_TEMPLATE_FIELD(flowDir, 1)
// Specify the Interface port - Uplink Port, Access Port, or NA.
IPFIX_VMW_TEMPLATE_FIELD(ingressInterfaceAttr, 2)
IPFIX_VMW_TEMPLATE_FIELD(egressInterfaceAttr, 2)
IPFIX_VMW_TEMPLATE_FIELD(encapExportRole, 1)
IPFIX_VMW_TEMPLATE_VAR_LEN_FIELD(virtualObsID, virtualObsDataLen)
IPFIX_TEMPLATE_PADDING(paddingOctets, 1)
IPFIX_TEMPLATE_END()
```

IPv6 Encapsulated Template

Template ID: 261

```
IPFIX_TEMPLATE_START(IPFIX_FLOW_TYPE_IPv6_ENCAP)
IPFIX_TEMPLATE_FIELD(sourceIPv4Address, 4)
IPFIX_TEMPLATE_FIELD(destinationIPv4Address, 4)
IPFIX_TEMPLATE_FIELD(octetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(packetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(flowStartSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(flowEndSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(sourceTransportPort, 2)
IPFIX_TEMPLATE_FIELD(destinationTransportPort, 2)
IPFIX_TEMPLATE_FIELD(ingressInterface, 4)
IPFIX_TEMPLATE_FIELD(egressInterface, 4)
IPFIX_TEMPLATE_FIELD(protocolIdentifier, 1)
IPFIX_TEMPLATE_FIELD(flowEndReason, 1)
```

```

IPFIX_TEMPLATE_FIELD(tcpFlags, 1)
IPFIX_TEMPLATE_FIELD(IPv6TOS, 1)
IPFIX_TEMPLATE_FIELD(maxTTL, 1)
IPFIX_TEMPLATE_FIELD(flowDir, 1)
//ENCAP specific
IPFIX_TEMPLATE_FIELD(encapId, 8)
IPFIX_VMW_TEMPLATE_FIELD(tenantSourceIPv6, 16)
IPFIX_VMW_TEMPLATE_FIELD(tenantDestIPv6, 16)
IPFIX_VMW_TEMPLATE_FIELD(tenantSourcePort, 2)
IPFIX_VMW_TEMPLATE_FIELD(tenantDestPort, 2)
IPFIX_VMW_TEMPLATE_FIELD(tenantProtocol, 1)
// Specify the Interface port - Uplink Port, Access Port, or NA
IPFIX_VMW_TEMPLATE_FIELD(ingressInterfaceAttr, 2)
IPFIX_VMW_TEMPLATE_FIELD(egressInterfaceAttr, 2)
// TUNNEL-GW or no.
IPFIX_VMW_TEMPLATE_FIELD(encapExportRole, 1)
IPFIX_VMW_TEMPLATE_VAR_LEN_FIELD(virtualObsID, virtualObsDataLen)
IPFIX_TEMPLATE_END()

```

IPv6 ICMP Template

Template ID: 262

```

IPFIX_TEMPLATE_START(IPFIX_FLOW_TYPE_IPv6_ICMP)
IPFIX_TEMPLATE_FIELD(sourceIPv6Address, 16)
IPFIX_TEMPLATE_FIELD(destinationIPv6Address, 16)
IPFIX_TEMPLATE_FIELD(octetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(packetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(flowStartSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(flowEndSysUpTime, 8)
IPFIX_TEMPLATE_FIELD(ingressInterface, 4)
IPFIX_TEMPLATE_FIELD(egressInterface, 4)
IPFIX_TEMPLATE_FIELD(protocolIdentifier, 1)
IPFIX_TEMPLATE_FIELD(flowEndReason, 1)
IPFIX_TEMPLATE_FIELD(IPv6TOS, 1)
IPFIX_TEMPLATE_FIELD(maxTTL, 1)
IPFIX_TEMPLATE_FIELD(flowDir, 1)
IPFIX_TEMPLATE_FIELD(encapId, 8)
// Specify the Interface port - Uplink Port, Access Port, or NA.
IPFIX_VMW_TEMPLATE_FIELD(ingressInterfaceAttr, 2)
IPFIX_VMW_TEMPLATE_FIELD(egressInterfaceAttr, 2)
IPFIX_VMW_TEMPLATE_FIELD(encapExportRole, 1)
IPFIX_VMW_TEMPLATE_VAR_LEN_FIELD(virtualObsID, virtualObsDataLen)
IPFIX_TEMPLATE_PADDING(paddingOctets, 2)
IPFIX_TEMPLATE_END()

```

IPv6 ICMP Encapsulated Template

Template ID: 263

```

IPFIX_TEMPLATE_START(IPFIX_FLOW_TYPE_IPv6_ICMP_ENCAP)
IPFIX_TEMPLATE_FIELD(sourceIPv4Address, 4)
IPFIX_TEMPLATE_FIELD(destinationIPv4Address, 4)
IPFIX_TEMPLATE_FIELD(octetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(packetDeltaCount, 8)
IPFIX_TEMPLATE_FIELD(flowStartSysUpTime, 8)

```

```

IPFIX_TEMPLATE_FIELD(flowEndSysUpTime, 8)
IPFIX_VMW_TEMPLATE_FIELD(sourceTransportPort, 2)
IPFIX_VMW_TEMPLATE_FIELD(destinationTransportPort, 2)
IPFIX_TEMPLATE_FIELD(ingressInterface, 4)
IPFIX_TEMPLATE_FIELD(egressInterface, 4)
IPFIX_TEMPLATE_FIELD(protocolIdentifier, 1)
IPFIX_TEMPLATE_FIELD(IPv6TOS, 1)
IPFIX_TEMPLATE_FIELD(maxTTL, 1)
IPFIX_TEMPLATE_FIELD(flowDir, 1)
IPFIX_TEMPLATE_FIELD(flowEndReason, 1)
//ENCAP Specific
IPFIX_TEMPLATE_FIELD(encapId, 8)
IPFIX_VMW_TEMPLATE_FIELD(tenantSourceIPv6, 16)
IPFIX_VMW_TEMPLATE_FIELD(tenantDestIPv6, 16)
IPFIX_VMW_TEMPLATE_FIELD(tenantProtocol, 1)
// Specify the Interface port - Uplink Port, Access Port, or NA
IPFIX_VMW_TEMPLATE_FIELD(ingressInterfaceAttr, 2)
IPFIX_VMW_TEMPLATE_FIELD(egressInterfaceAttr, 2)
// TUNNEL-GW or no.
IPFIX_VMW_TEMPLATE_FIELD(encapExportRole, 1)
IPFIX_VMW_TEMPLATE_VAR_LEN_FIELD(virtualObsID, virtualObsDataLen)
IPFIX_TEMPLATE_PADDING(paddingOctets, 1)
IPFIX_TEMPLATE_END()

```

ESXi Distributed Firewall IPFIX Templates

An ESXi host transport node supports two distributed firewall IPFIX flow templates.

IPv4 Template

Template ID: 288

```

IPFIX_TEMPLATE_FIELD(sourceIPv4Address,4)
IPFIX_TEMPLATE_FIELD(destinationIPv4Address,4)
IPFIX_TEMPLATE_FIELD(sourceTransportPort,2)
IPFIX_TEMPLATE_FIELD(destinationTransportPort,2)
IPFIX_TEMPLATE_FIELD(protocolIdentifier,1)
IPFIX_TEMPLATE_FIELD(icmpTypeIPv4,1)
IPFIX_TEMPLATE_FIELD(icmpCodeIPv4,1)
IPFIX_TEMPLATE_FIELD(flowStartSeconds,4)
IPFIX_TEMPLATE_FIELD(flowEndSeconds,4)
IPFIX_TEMPLATE_FIELD(octetDeltaCount,8)
IPFIX_TEMPLATE_FIELD(packetDeltaCount,8)
IPFIX_TEMPLATE_FIELD(firewallEvent,1)
IPFIX_TEMPLATE_FIELD(direction,1)
IPFIX_TEMPLATE_FIELD(ruleId,4)
IPFIX_TEMPLATE_FIELD(vifUuid,16)
IPFIX_TEMPLATE_FIELD(sessionFlags,1)
IPFIX_TEMPLATE_FIELD(flowDirection,1)
IPFIX_TEMPLATE_FIELD(flowId,8)
IPFIX_TEMPLATE_FIELD(algControlFlowId,8)
IPFIX_TEMPLATE_FIELD(algType,1)
IPFIX_TEMPLATE_FIELD(algFlowType,1)
IPFIX_TEMPLATE_FIELD(averageLatency,4)
IPFIX_TEMPLATE_FIELD(retransmissionCount,4)

```

IPv6 Template

Template ID: 289

```
IPFIX_TEMPLATE_FIELD(sourceIPv6Address,16)
IPFIX_TEMPLATE_FIELD(destinationIPv6Address,16)
IPFIX_TEMPLATE_FIELD(sourceTransportPort,2)
IPFIX_TEMPLATE_FIELD(destinationTransportPort,2)
IPFIX_TEMPLATE_FIELD(protocolIdentifier,1)
IPFIX_TEMPLATE_FIELD icmpTypeIPv6,1)
IPFIX_TEMPLATE_FIELD icmpCodeIPv6,1)
IPFIX_TEMPLATE_FIELD(flowStartSeconds,4)
IPFIX_TEMPLATE_FIELD(flowEndSeconds,4)
IPFIX_TEMPLATE_FIELD(octetDeltaCount,8)
IPFIX_TEMPLATE_FIELD(packetDeltaCount,8)
IPFIX_TEMPLATE_FIELD(firewallEvent,1)
IPFIX_TEMPLATE_FIELD(direction,1)
IPFIX_TEMPLATE_FIELD(ruleId,4)
IPFIX_TEMPLATE_FIELD(vifUuid,16)
IPFIX_TEMPLATE_FIELD(sessionFlags,1)
IPFIX_TEMPLATE_FIELD(flowDirection,1)
IPFIX_TEMPLATE_FIELD(flowId,8)
IPFIX_TEMPLATE_FIELD(algControlFlowId,8)
IPFIX_TEMPLATE_FIELD(algType,1)
IPFIX_TEMPLATE_FIELD(algFlowType,1)
IPFIX_TEMPLATE_FIELD(averageLatency,4)
IPFIX_TEMPLATE_FIELD(retransmissionCount,4)
```


KVM IPFIX Templates

A KVM host transport node supports 88 IPFIX flow templates and one options template.

The following table lists VMware-specific elements in the KVM IPFIX packets.

Element ID	Parameter Name	Data Type	Unit
891	tunnelType	unsigned8	1 byte
892	tunnelKey	bytes	variable length
893	tunnelSourceIPv4Address	unsigned32	4 bytes
894	tunnelDestinationIPv4Address	unsigned32	4 bytes
895	tunnelProtocolIdentifier	unsigned8	1 byte
896	tunnelSourceTransportPort	unsigned16	2 bytes
897	tunnelDestinationTransportPort	unsigned16	2 bytes
898	virtualObsID	string	variable length

The following diagram shows the flow of traffic between VMs attached to KVM hosts monitored by the IPFIX feature:



The KVM IPv4 IPFIX ingress template will have the following elements:

- standard elements
- virtualObsID: 6d876a1c-e0ac-4bcf-85ee-bdd42fa7ba34 (logical port ID)

KVM Ethernet IPFIX Templates

There are four KVM Ethernet IPFIX templates: ingress, egress, ingress with tunnel, and egress with tunnel.

Ethernet Ingress

Template ID: 256. Field count: 27.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)

- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)

Ethernet Egress

Template ID: 257. Field count: 31.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)

- IF_DESC (length: variable)
- OUTPUT_SNMP (length: 4)
- Unknown(369) (length: 8)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMcastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)

Ethernet Ingress with Tunnel

Template ID: 258. Field count: 34.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)

- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)

Ethernet Egress with Tunnel

Template ID: 259. Field count: 38.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)

- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (length: 4)
- Unknown(369) (length: 8)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)

- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)

KVM IPv4 IPFIX Templates

There are four KVM IPv4 IPFIX templates: ingress, egress, ingress with tunnel, and egress with tunnel.

IPv4 Ingress

Template ID: 276. Field count: 45.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)

- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

IPv4 Egress

Template ID: 277. Field count: 49.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)

- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)

- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMcastOctetTotalCount (Length: 8)

IPv4 Ingress with Tunnel

Template ID: 278. Field count: 52.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)

- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)

- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

IPv4 Egress with Tunnel

Template ID: 279. Field count: 56.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))

- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

KVM TCP over IPv4 IPFIX Templates

There are four KVM TCP over IPv4 IPFIX templates: ingress, egress, ingress with tunnel, and egress with tunnel.

TCP over IPv4 Ingress

Template ID: 280. Field count: 53.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)

- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)
- tcpAckTotalCount (Length: 8)
- tcpFinTotalCount (Length: 8)
- tcpPshTotalCount (Length: 8)
- tcpRstTotalCount (Length: 8)
- tcpSynTotalCount (Length: 8)
- tcpUrgTotalCount (Length: 8)

TCP over IPv4 Egress

Template ID: 281. Field count: 57.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)

- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)

- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)
- tcpAckTotalCount (Length: 8)
- tcpFinTotalCount (Length: 8)
- tcpPshTotalCount (Length: 8)
- tcpRstTotalCount (Length: 8)
- tcpSynTotalCount (Length: 8)
- tcpUrgTotalCount (Length: 8)

TCP over IPv4 Ingress with Tunnel

Template ID: 282. Field count: 60.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)

- IF_DESC (length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)

- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)
- tcpAckTotalCount (Length: 8)
- tcpFinTotalCount (Length: 8)
- tcpPshTotalCount (Length: 8)
- tcpRstTotalCount (Length: 8)
- tcpSynTotalCount (Length: 8)
- tcpUrgTotalCount (Length: 8)

TCP over IPv4 Egress with Tunnel

Template ID: 283. Field count: 64.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)

- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)

- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMcastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMcastOctetTotalCount (Length: 8)
- tcpAckTotalCount (Length: 8)
- tcpFinTotalCount (Length: 8)
- tcpPshTotalCount (Length: 8)
- tcpRstTotalCount (Length: 8)
- tcpSynTotalCount (Length: 8)
- tcpUrgTotalCount (Length: 8)

KVM UDP over IPv4 IPFIX Templates

There are four KVM UDP over IPv4 IPFIX templates: ingress, egress, ingress with tunnel, and egress with tunnel.

UDP over IPv4 Ingress

Template ID: 284. Field count: 47.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)

- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)

- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMcastOctetTotalCount (Length: 8)

UDP over IPv4 Egress

Template ID: 285. Field count: 51.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)

- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)

- postMCastOctetTotalCount (Length: 8)

UDP over IPv4 Ingress with Tunnel

Template ID: 286. Field count: 54.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))

- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

UDP over IPv4 Egress with Tunnel

Template ID: 287. Field count: 58.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)

- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)

- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

KVM SCTP over IPv4 IPFIX Templates

There are four KVM SCTP over IPv4 IPFIX templates: ingress, egress, ingress with tunnel, and egress with tunnel.

SCTP over IPv4 Ingress

Template ID: 288. Field count: 47.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)

- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)

- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMcastOctetTotalCount (Length: 8)

SCTP over IPv4 Egress

Template ID: 289. Field count: 51.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)

- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMcastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)

- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

SCTP over IPv4 Ingress with Tunnel

Template ID: 290. Field count: 54.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))

- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

SCTP over IPv4 Egress with Tunnel

Template ID: 291. Field count: 58.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)

- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)

- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMcastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMcastOctetTotalCount (Length: 8)

KVM ICMPv4 IPFIX Templates

There are four KVM ICMPv4 IPFIX templates: ingress, egress, ingress with tunnel, and egress with tunnel.

ICMPv4 Ingress

Template ID: 292. Field count: 47.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)

- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- ICMP_IPv4_TYPE (Length: 1)
- ICMP_IPv4_CODE (Length: 1)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)

- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

ICMPv4 Egress

Template ID: 293. Field count: 51.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)

- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- ICMP_IPv4_TYPE (Length: 1)
- ICMP_IPv4_CODE (Length: 1)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)

- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

ICMPv4 Ingress with Tunnel

Template ID: 294. Field count: 54.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- ICMP_IPv4_TYPE (Length: 1)
- ICMP_IPv4_CODE (Length: 1)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))

- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMcastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMcastOctetTotalCount (Length: 8)

ICMPv4 Egress with Tunnel

Template ID: 295. Field count: 58.

The fields are:

- observationPointId (length: 4)

- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- ICMP_IPv4_TYPE (Length: 1)
- ICMP_IPv4_CODE (Length: 1)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)

- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

KVM IPv6 IPFIX Templates

There are four KVM IPv6 IPFIX templates: ingress, egress, ingress with tunnel, and egress with tunnel.

IPv6 Ingress

Template ID: 296. Field count: 46.

The fields are:

- observationPointId (length: 4)

- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)

- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

IPv6 Egress

Template ID: 297. Field count: 50.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)

- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMcastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)

- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

IPv6 Ingress with Tunnel

Template ID: 298. Field count: 53.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))

- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

IPv6 Egress with Tunnel

Template ID: 299. Field count: 57.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)

- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)

- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

KVM TCP over IPv6 IPFIX Templates

There are four KVM TCP over IPv6 IPFIX templates: ingress, egress, ingress with tunnel, and egress with tunnel.

TCP over IPv6 Ingress

Template ID: 300. Field count: 54.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)

- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)

- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)
- tcpAckTotalCount (Length: 8)
- tcpFinTotalCount (Length: 8)
- tcpPshTotalCount (Length: 8)
- tcpRstTotalCount (Length: 8)
- tcpSynTotalCount (Length: 8)
- tcpUrgTotalCount (Length: 8)

TCP over IPv6 Egress

Template ID: 301. Field count: 58.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)

- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)

- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)
- tcpAckTotalCount (Length: 8)
- tcpFinTotalCount (Length: 8)
- tcpPshTotalCount (Length: 8)
- tcpRstTotalCount (Length: 8)
- tcpSynTotalCount (Length: 8)
- tcpUrgTotalCount (Length: 8)

TCP over IPv6 Ingress with Tunnel

Template ID: 302. Field count: 61.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)

- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)

- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)
- tcpAckTotalCount (Length: 8)
- tcpFinTotalCount (Length: 8)
- tcpPshTotalCount (Length: 8)
- tcpRstTotalCount (Length: 8)
- tcpSynTotalCount (Length: 8)
- tcpUrgTotalCount (Length: 8)

TCP over IPv6 Egress with Tunnel

Template ID: 303. Field count: 65.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)

- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)

- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMcastOctetTotalCount (Length: 8)
- tcpAckTotalCount (Length: 8)
- tcpFinTotalCount (Length: 8)
- tcpPshTotalCount (Length: 8)
- tcpRstTotalCount (Length: 8)
- tcpSynTotalCount (Length: 8)
- tcpUrgTotalCount (Length: 8)

KVM UDP over IPv6 IPFIX Templates

There are four KVM UDP over IPv6 IPFIX templates: ingress, egress, ingress with tunnel, and egress with tunnel.

UDP over IPv6 Ingress

Template ID: 304. Field count: 48.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)

- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)

- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

UDP over IPv6 Egress

Template ID: 305. Field count: 52.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)

- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMcastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMcastOctetTotalCount (Length: 8)

UDP over IPv6 Ingress with Tunnel

Template ID: 306. Field count: 55.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))

- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

UDP over IPv6 Egress with Tunnel

Template ID: 307. Field count: 59.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)

- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)

- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMcastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMcastOctetTotalCount (Length: 8)

KVM SCTP over IPv6 IPFIX Templates

There are four KVM SCTP over IPv6 IPFIX templates: ingress, egress, ingress with tunnel, and egress with tunnel.

SCTP over IPv6 Ingress

Template ID: 308. Field count: 48.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)

- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)

- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

SCTP over IPv6 Egress

Template ID: 309. Field count: 52.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- IP_PROTOCOL_VERSION (Length: 1)

- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)

- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

SCTP over IPv6 Ingress with Tunnel

Template ID: 310. Field count: 55.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))

- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

SCTP over IPv6 Egress with Tunnel

Template ID: 311. Field count: 59.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))

- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

KVM ICMPv6 IPFIX Templates

There are four KVM ICMPv6 IPFIX templates: ingress, egress, ingress with tunnel, and egress with tunnel.

ICMPv6 Ingress

Template ID: 312. Field count: 48.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- ICMP_IPv6_TYPE (Length: 1)
- ICMP_IPv6_CODE (Length: 1)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)

- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

ICMPv6 Egress

Template ID: 313. Field count: 52.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)

- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- ICMP_IPv6_TYPE (Length: 1)
- ICMP_IPv6_CODE (Length: 1)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)

- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

ICMPv6 Ingress with Tunnel

Template ID: 314. Field count: 55.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)

- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- ICMP_IPv6_TYPE (Length: 1)
- ICMP_IPv6_CODE (Length: 1)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)

- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMcastOctetTotalCount (Length: 8)

ICMPv6 Egress with Tunnel

Template ID: 315. Field count: 59.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)

- FLOW_LABEL (Length: 4)
- ICMP_IPv6_TYPE (Length: 1)
- ICMP_IPv6_CODE (Length: 1)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)

- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

KVM Ethernet VLAN IPFIX Templates

There are four KVM Ethernet VLAN IPFIX templates: ingress, egress, ingress with tunnel, and egress with tunnel.

Ethernet VLAN Ingress

Template ID: 316. Field count: 30.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)

- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)

Ethernet VLAN Egress

Template ID: 317. Field count: 34.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (length: 4)
- Unknown(369) (length: 8)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)

- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMcastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)

Ethernet VLAN Ingress with Tunnel

Template ID: 318. Field count: 37.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))

- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMcastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)

Ethernet VLAN Egress with Tunnel

Template ID: 319. Field count: 41.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)

- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (length: 4)
- Unknown(369) (length: 8)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)

- Unknown(353) (length: 8)
- flowEndReason (length: 1)

KVM IPv4 VLAN IPFIX Templates

There are four KVM IPv4 VLAN IPFIX templates: ingress, egress, ingress with tunnel, and egress with tunnel.

IPv4 VLAN Ingress

Template ID: 336. Field count: 48.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)

- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMcastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMcastOctetTotalCount (Length: 8)

IPv4 VLAN Egress

Template ID: 337. Field count: 52.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)

- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)

- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

IPv4 VLAN Ingress with Tunnel

Template ID: 338. Field count: 55.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)

- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)

- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

IPv4 VLAN Egress with Tunnel

Template ID: 339. Field count: 59.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)

- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)

- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

KVM TCP over IPv4 VLAN IPFIX Templates

There are four KVM TCP over IPv4 VLAN IPFIX templates: ingress, egress, ingress with tunnel, and egress with tunnel.

TCP over IPv4 VLAN Ingress

Template ID: 340. Field count: 56.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)

- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMcastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMcastOctetTotalCount (Length: 8)
- tcpAckTotalCount (Length: 8)

- tcpFinTotalCount (Length: 8)
- tcpPshTotalCount (Length: 8)
- tcpRstTotalCount (Length: 8)
- tcpSynTotalCount (Length: 8)
- tcpUrgTotalCount (Length: 8)

TCP over IPv4 VLAN Egress

Template ID: 341. Field count: 60.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)

- IP_DST_ADDR (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)
- tcpAckTotalCount (Length: 8)
- tcpFinTotalCount (Length: 8)

- tcpPshTotalCount (Length: 8)
- tcpRstTotalCount (Length: 8)
- tcpSynTotalCount (Length: 8)
- tcpUrgTotalCount (Length: 8)

TCP over IPv4 VLAN Ingress with Tunnel

Template ID: 342. Field count: 63.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))

- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMcastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMcastOctetTotalCount (Length: 8)

- tcpAckTotalCount (Length: 8)
- tcpFinTotalCount (Length: 8)
- tcpPshTotalCount (Length: 8)
- tcpRstTotalCount (Length: 8)
- tcpSynTotalCount (Length: 8)
- tcpUrgTotalCount (Length: 8)

TCP over IPv4 VLAN Egress with Tunnel

Template ID: 343. Field count: 67.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)

- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)

- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)
- tcpAckTotalCount (Length: 8)
- tcpFinTotalCount (Length: 8)
- tcpPshTotalCount (Length: 8)
- tcpRstTotalCount (Length: 8)
- tcpSynTotalCount (Length: 8)
- tcpUrgTotalCount (Length: 8)

KVM UDP over IPv4 VLAN IPFIX Templates

There are four KVM UDP over IPv4 VLAN IPFIX templates: ingress, egress, ingress with tunnel, and egress with tunnel.

UDP over IPv4 VLAN Ingress

Template ID: 344. Field count: 50.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)

- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMcastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)

- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

UDP over IPv4 VLAN Egress

Template ID: 345. Field count: 54.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)

- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

UDP over IPv4 VLAN Ingress with Tunnel

Template ID: 346. Field count: 57.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))

- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

UDP over IPv4 VLAN Egress with Tunnel

Template ID: 347. Field count: 61.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)

- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))

- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

KVM SCTP over IPv4 VLAN IPFIX Templates

There are four KVM SCTP over IPv4 VLAN IPFIX templates: ingress, egress, ingress with tunnel, and egress with tunnel.

SCTP over IPv4 VLAN Ingress

Template ID: 348. Field count: 50.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)

- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

SCTP over IPv4 VLAN Egress

Template ID: 349. Field count: 54.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)

- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)

- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMcastOctetTotalCount (Length: 8)

SCTP over IPv4 VLAN Ingress with Tunnel

Template ID: 350. Field count: 57.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)

- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)

- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

SCTP over IPv4 VLAN Egress with Tunnel

Template ID: 351. Field count: 61.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)

- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)

- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

KVM ICMPv4 VLAN IPFIX Templates

There are four KVM ICMPv4 VLAN IPFIX templates: ingress, egress, ingress with tunnel, and egress with tunnel.

ICMPv4 VLAN Ingress

Template ID: 352. Field count: 50.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)

- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- ICMP_IPv4_TYPE (Length: 1)
- ICMP_IPv4_CODE (Length: 1)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

ICMPv4 VLAN Egress

Template ID: 353. Field count: 54.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- ICMP_IPv4_TYPE (Length: 1)
- ICMP_IPv4_CODE (Length: 1)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)

- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

ICMPv4 VLAN Ingress with Tunnel

Template ID: 354. Field count: 57.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)

- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- ICMP_IPv4_TYPE (Length: 1)
- ICMP_IPv4_CODE (Length: 1)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)

- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

ICMPv4 VLAN Egress with Tunnel

Template ID: 355. Field count: 61.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)

- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IP_SRC_ADDR (Length: 4)
- IP_DST_ADDR (Length: 4)
- ICMP_IPv4_TYPE (Length: 1)
- ICMP_IPv4_CODE (Length: 1)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)

- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

KVM IPv6 VLAN IPFIX Templates

There are four KVM IPv6 VLAN IPFIX templates: ingress, egress, ingress with tunnel, and egress with tunnel.

IPv6 VLAN Ingress

Template ID: 356. Field count: 49.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)

- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)

- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

IPv6 VLAN Egress

Template ID: 357. Field count: 53.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)

- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMcastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)

- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

IPv6 VLAN Ingress with Tunnel

Template ID: 358. Field count: 56.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))

- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

IPv6 VLAN Egress with Tunnel

Template ID: 359. Field count: 60.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))

- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

KVM TCP over IPv6 VLAN IPFIX Templates

There are four KVM TCP over IPv6 VLAN IPFIX templates: ingress, egress, ingress with tunnel, and egress with tunnel.

TCP over IPv6 VLAN Ingress

Template ID: 360. Field count: 57.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)

- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)
- tcpAckTotalCount (Length: 8)
- tcpFinTotalCount (Length: 8)
- tcpPshTotalCount (Length: 8)
- tcpRstTotalCount (Length: 8)
- tcpSynTotalCount (Length: 8)
- tcpUrgTotalCount (Length: 8)

TCP over IPv6 VLAN Egress

Template ID: 361. Field count: 61.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 898 (length: variable, PEN: VMware Inc. (6876))

- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)
- tcpAckTotalCount (Length: 8)
- tcpFinTotalCount (Length: 8)
- tcpPshTotalCount (Length: 8)
- tcpRstTotalCount (Length: 8)
- tcpSynTotalCount (Length: 8)
- tcpUrgTotalCount (Length: 8)

TCP over IPv6 VLAN Ingress with Tunnel

Template ID: 362. Field count: 64.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))

- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)
- tcpAckTotalCount (Length: 8)
- tcpFinTotalCount (Length: 8)
- tcpPshTotalCount (Length: 8)

- tcpRstTotalCount (Length: 8)
- tcpSynTotalCount (Length: 8)
- tcpUrgTotalCount (Length: 8)

TCP over IPv6 VLAN Egress with Tunnel

Template ID: 363. Field count: 68.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)

- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMcastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)

- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)
- tcpAckTotalCount (Length: 8)
- tcpFinTotalCount (Length: 8)
- tcpPshTotalCount (Length: 8)
- tcpRstTotalCount (Length: 8)
- tcpSynTotalCount (Length: 8)
- tcpUrgTotalCount (Length: 8)

KVM UDP over IPv6 VLAN IPFIX Templates

There are four KVM UDP over IPv6 VLAN IPFIX templates: ingress, egress, ingress with tunnel, and egress with tunnel.

UDP over IPv6 VLAN Ingress

Template ID: 364. Field count: 51.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)

- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)

- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

UDP over IPv6 VLAN Egress

Template ID: 365. Field count: 55.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)

- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

UDP over IPv6 VLAN Ingress with Tunnel

Template ID: 366. Field count: 58.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))

- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

UDP over IPv6 VLAN Egress with Tunnel

Template ID: 367. Field count: 62.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)

- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))

- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

KVM SCTP over IPv6 VLAN IPFIX Templates

There are four KVM SCTP over IPv6 VLAN IPFIX templates: ingress, egress, ingress with tunnel, and egress with tunnel.

SCTP over IPv6 VLAN Ingress

Template ID: 368. Field count: 51.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)

- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

SCTP over IPv6 VLAN Egress

Template ID: 369. Field count: 55.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)

- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)

- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

SCTP over IPv6 VLAN Ingress with Tunnel

Template ID: 370. Field count: 58.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)

- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)

- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

SCTP over IPv6 VLAN Egress with Tunnel

Template ID: 371. Field count: 62.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)

- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- L4_SRC_PORT (Length: 2)
- L4_DST_PORT (Length: 2)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)

- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

KVM ICMPv6 VLAN IPFIX Templates

There are four KVM ICMPv6 IPFIX templates: ingress, egress, ingress with tunnel, and egress with tunnel.

ICMPv6 Ingress

Template ID: 372. Field count: 51.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)

- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- ICMP_IPv6_TYPE (Length: 1)
- ICMP_IPv6_CODE (Length: 1)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)

- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

ICMPv6 Egress

Template ID: 373. Field count: 55.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)

- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- ICMP_IPv6_TYPE (Length: 1)
- ICMP_IPv6_CODE (Length: 1)
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP_LENGTH_MINIMUM (Length: 8)
- IP_LENGTH_MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)

- postMCastOctetTotalCount (Length: 8)

ICMPv6 Ingress with Tunnel

Template ID: 374. Field count: 58.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- ICMP_IPv6_TYPE (Length: 1)
- ICMP_IPv6_CODE (Length: 1)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))
- 896 (length: 2, PEN: VMware Inc. (6876))

- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

ICMPv6 Egress with Tunnel

Template ID: 375. Field count: 62.

The fields are:

- observationPointId (length: 4)
- DIRECTION (length: 1)
- SRC_MAC (length: 6)
- DESTINATION_MAC (length: 6)
- ethernetType (length: 2)
- ethernetHeaderLength (length: 1)
- INPUT_SNMP (length: 4)
- Unknown(368) (length: 4)
- IF_NAME (length: variable)
- IF_DESC (length: variable)
- OUTPUT_SNMP (Length: 4)
- Unknown(369) (Length: 4)
- IF_NAME (Length: variable)
- IF_DESC (Length: variable)
- SRC_VLAN (Length: 2)
- dot1qVlanId (Length: 2)
- dot1qPriority (Length: 1)
- IP_PROTOCOL_VERSION (Length: 1)
- IP_TTL (Length: 1)
- PROTOCOL (Length: 1)
- IP_DSCP (Length: 1)
- IP_PRECEDENCE (Length: 1)
- IP_TOS (Length: 1)
- IPV6_SRC_ADDR (Length: 4)
- IPV6_DST_ADDR (Length: 4)
- FLOW_LABEL (Length: 4)
- ICMP_IPv6_TYPE (Length: 1)
- ICMP_IPv6_CODE (Length: 1)
- 893 (length: 4, PEN: VMware Inc. (6876))
- 894 (length: 4, PEN: VMware Inc. (6876))
- 895 (length: 1, PEN: VMware Inc. (6876))

- 896 (length: 2, PEN: VMware Inc. (6876))
- 897 (length: 2, PEN: VMware Inc. (6876))
- 891 (length: 1, PEN: VMware Inc. (6876))
- 892 (length: variable, PEN: VMware Inc. (6876))
- 898 (length: variable, PEN: VMware Inc. (6876))
- flowStartDeltaMicroseconds (length: 4)
- flowEndDeltaMicroseconds (length: 4)
- DROPPED_PACKETS (length: 8)
- DROPPED_PACKETS_TOTAL (length: 8)
- PKTS (length: 8)
- PACKETS_TOTAL (length: 8)
- Unknown(354) (length: 8)
- Unknown(355) (length: 8)
- Unknown(356) (length: 8)
- Unknown(357) (length: 8)
- Unknown(358) (length: 8)
- MUL_DPKTS (length: 8)
- postMCastPacketTotalCount (length: 8)
- Unknown(352) (length: 8)
- Unknown(353) (length: 8)
- flowEndReason (length: 1)
- DROPPED_BYTES (Length: 8)
- DROPPED_BYTES_TOTAL (Length: 8)
- BYTES (Length: 8)
- BYTES_TOTAL (Length: 8)
- BYTES_SQUARED (Length: 8)
- BYTES_SQUARED_PERMANENT (Length: 8)
- IP LENGTH MINIMUM (Length: 8)
- IP LENGTH MAXIMUM (Length: 8)
- MUL_DOCTETS (Length: 8)
- postMCastOctetTotalCount (Length: 8)

KVM Options IPFIX Templates

There is one KVM options template, based on IETF RFC 7011, section 3.4.2.

Options Template

Template ID: 462. Scope count: 1. Data count: 1.

Monitor a Logical Switch Port Activity in Manager Mode

You can monitor the logical port activity for example, to troubleshoot network congestion and packets being dropped

Prerequisites

- Verify that a logical switch port is configured. See [Connecting a VM to a Logical Switch in Manager Mode](#).
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Verify that a logical switch port is configured. See [Connecting a VM to a Logical Switch in Manager Mode](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Logical Switches > Ports**
- 3 Click the name of a port.
- 4 Click the **Monitor** tab.
The port status and statistics are displayed.
- 5 To download a CSV file of the MAC addresses that has been learned by the host, click **Download MAC Table**.
- 6 To monitor activity on the port, click **Begin Tracking**.

A port tracking page opens. You can view the bidirectional port traffic and identify dropped packets. The port tracker page also lists the switching profiles attached to the logical switch port.

Results

If you notice dropped packets because of network congestion, you can configure a QoS switching profile for the logical switch port to prevent data loss on preferred packets. See [Understanding QoS Switching Profile](#).

Authentication and Authorization

18

You can log in to NSX Manager using a local user account, a user account managed by VMware Identity Manager (vIDM), or a user account managed by a directory service such as Active Directory over LDAP or OpenLDAP. You can also assign roles to user accounts managed by vIDM or a directory service to implement role-based access control.

NSX Manager recognizes only system-generated session identifiers and invalidates session identifiers upon administrator logout or other session termination. Upon successful login, the NSX Manager uses a random number generator to create a random session ID and stores that ID in memory. When clients make requests to the NSX Manager, it only allows clients to authenticate if the session ID they present matches one of the IDs generated by the server. When any user logs out of NSX Manager, the session identifier is immediately destroyed and cannot be reused.

Access to NSX Manager via UI, API and CLI is subject to authentication and authorization. In addition, such access will generate audit logs. This logging is enabled by default and cannot be disabled. Auditing of sessions is initiated at system startup. An audit log message is distinguished by having the text `audit="true"` in the structured data part of the log message.

Local user passwords on NSX appliances are secured using the default Linux/PAM libraries which store the hashed and salted representation in `/etc/shadow`. During authentication, the password entered by the user is obfuscated. Other passwords are encrypted using a random key that is stored in the local file system.

Read the following topics next:

- [Local User Accounts](#)
- [Integration with VMware Identity Manager/Workspace ONE Access](#)
- [Integration with LDAP](#)
- [Add a Role Assignment or Principal Identity](#)
- [Configuring Both vIDM and LDAP or Transitioning from vIDM to LDAP](#)
- [Role-Based Access Control](#)
- [Logging User Account Changes](#)

Local User Accounts

Each NSX-T appliance has three local accounts, `admin`, `audit`, and `root`. To administer NSX-T, you must log in as `admin`.

The `root` user has special privileges. You must not log in as `root` and make changes that are not documented in this guide, except when under the guidance of VMware. Changes made by the `root` user can cause catastrophic failures. In a production environment, the `root` password should be secured and made available for privileged access only.

For additional security-related information about the NSX Manager, see the section "Security" in [Chapter 1 NSX Manager](#).

Manage a User's Password or Name

You can manage the `admin` and `audit` user's account through an NSX-T appliance's CLI.

The `admin` user can manage the password and change the name of the `admin` and `audit` users, but cannot add, delete, or disable users. Any change to the `admin` or `audit` user's account will be audited.

The `audit` user has read privileges to the NSX-T environment and is not active by default. To activate it, log in as `admin` and run the `set user audit password` command and provide a new password. When prompted for the current password, press the **Enter** key.

By default, user passwords expire after 90 days. You can change or disable the password expiration for each user.

When the password of `admin` or `audit` on the NSX Manager will expire within 30 days, the NSX Manager web interface displays a password expiration notification. If you set the password expiration to 30 days or less the notification is always present. The notification includes a **Change Password** link. Click the link to change the user's password.

Prerequisites

Familiarize yourself with the password complexity requirements for NSX Manager and NSX Edge. See "NSX Manager Installation" and "NSX Edge Installation" in the *NSX-T Data Center Installation Guide*.

Procedure

- 1 Log in to the appliance's CLI as `admin`.
- 2 To change the password, run the `set user <username> password` command. For example:

```
nsx> set user admin password
Current password:
New password:
Confirm new password:
nsx>
```

- 3 To change the name of the `admin` or `audit` user, run the `set user <username> username <new username>` command. For example:

```
nsx> set user admin username admin1
nsx>
```

- 4 To get the password expiration information, run the `get user <username> password-expiration` command. For example:

```
nsx> get user admin password-expiration
Password expires 90 days after last change
nsx>
```

- 5 To set the password expiration time in days, run the `set user <username> password-expiration <number of days>` command. For example:

```
nsx> set user admin password-expiration 120
nsx>
```

- 6 To disable password expiration, run the `clear user <username> password-expiration` command. For example:

```
nsx> clear user admin password-expiration
nsx>
```

Resetting the Passwords of an Appliance

The following procedure applies to NSX Manager, NSX Edge, Cloud Service Manager, and NSX Intelligence appliances.

Note If you have an NSX Manager cluster, resetting the password for the `admin` or `audit` user on one NSX Manager will automatically reset the password for the other NSX Managers in the cluster. Note that the synchronization of the password can take several minutes or more.

If you have renamed the user `admin` or `audit`, use the new name in the following procedures.

When you reboot an appliance, the GRUB boot menu does not appear by default. The following procedure requires that you have configured GRUB to display the GRUB boot menu. For more information about configuring GRUB and changing the GRUB `root` password, see "Configure NSX-T Data Center to Display the GRUB Menu at Boot Time" in the *NSX-T Data Center Installation Guide*.

If you know the password for `root` but have forgotten the password for `admin` or `audit`, you can reset it using the following procedure:

- 1 Log in to the appliance as `root`.

- 2 For an NSX Intelligence appliance or a Cloud Service Manager, skip this step. For NSX Edge, run the command `/etc/init.d/nsx-edge-api-server stop`. Otherwise, run the command `/etc/init.d/nsx-mp-api-server stop`.
- 3 To reset the password for `admin`, run the command `passwd admin`.
- 4 To reset the password for `audit`, run the command `passwd audit`.
- 5 Run the command `touch /var/vmware/nsx/reset_cluster_credentials`.
- 6 For NSX Edge, run the command `/etc/init.d/nsx-edge-api-server start`. Otherwise, run the command `/etc/init.d/nsx-mp-api-server start`.

If you have forgotten the `root` user's password, you can reset it using the following procedure. You can then use the above procedure to reset the password for `admin` or `audit`.

Procedure

- 1 Connect to the console of the appliance.
- 2 Reboot the system.
- 3 When the GRUB boot menu appears, press the left **SHIFT** or **ESC** key quickly. If you wait too long and the boot sequence does not pause, you must reboot the system again.
- 4 Press **e** to edit the menu.

Enter the user name `root` and the GRUB password for `root` (not the same as the appliance's user `root`).
- 5 Keep the cursor on the Ubuntu selection.
- 6 Press **e** to edit the selected option.
- 7 Search for the line starting with `linux` and add `systemd.wants=PasswordRecovery.service` to the end of the line.
- 8 Press **Ctrl-X** to boot.
- 9 When the log messages stop, enter the new password for `root`.
- 10 Enter the password again.

The boot process will continue.
- 11 After the reboot, you can verify the password change by logging in as `root` with the new password.

Authentication Policy Settings

You can view or change the authentication policy settings through the CLI.

You can view or set the minimum password length with the following commands:

```
get auth-policy minimum-password-length
set auth-policy minimum-password-length <password-length>
```

The following commands apply to logging in to the NSX Manager UI, or making an API call:

```
get auth-policy api lockout-period
get auth-policy api lockout-reset-period
get auth-policy api max-auth-failures
set auth-policy api lockout-period <lockout-period>
set auth-policy api lockout-reset-period <lockout-reset-period>
set auth-policy api max-auth-failures <auth-failures>
```

The following commands apply to logging in to the CLI on an NSX Manager or an NSX Edge node:

```
get auth-policy cli lockout-period
get auth-policy cli max-auth-failures
set auth-policy cli lockout-period <lockout-period>
set auth-policy cli max-auth-failures <auth-failures>
```

For more information about the CLI commands, see the *NSX-T Command-Line Interface Reference*.

By default, after five consecutive failed attempts to log in to the NSX Manager UI, the administrator account is locked for 15 minutes. You can disable account lockout with the following command:

```
set auth-policy api lockout-period 0
```

Similarly, you can disable account lockout for the CLI with the following command:

```
set auth-policy cli lockout-period 0
```

Integration with VMware Identity Manager/Workspace ONE Access

You can configure NSX Manager to authenticate users using VMware Identity Manager (vIDM).

Note: The new product name for VMware Identity Manager is VMware Workspace ONE Access.

Time Synchronization between NSX Manager, vIDM, and Related Components

For authentication to work correctly, NSX Manager, vIDM and other service providers such as Active Directory must all be time synchronized. This section describes how to time synchronize these components.

VMware Infrastructure

Follow the instructions in the following KB articles to synchronize ESXi hosts.

- <https://kb.vmware.com/kb/1003736>
- <https://kb.vmware.com/kb/2012069>

Third-Party Infrastructure

Follow the vendor's documentation on how synchronize VMs and hosts.

Configuring NTP on the vIDM Server (Not Recommended)

If you are not able to synchronize time across the hosts, you can disable synchronizing to host and configure NTP on the vIDM server. This method is not recommend because it requires the opening of UDP port 123 on the vIDM server

- Check the clock on the vIDM server and make sure it is correct.

```
# hwclock
Tue May 9 12:08:43 2017 -0.739213 seconds
```

- Edit `/etc/ntp.conf` and add the following entries if they don't exist.

```
server time.nist.gov
server pool.ntp.org
server time.is dynamic
restrict 192.168.100.0 netmask 255.255.255.0 nomodify notrap
```

- Open UDP port 123.

```
# iptables -A INPUT -p udp --dport 123 -j ACCEPT
```

Run the following command to check that the port is open.

```
# iptables -L -n
```

- Start the NTP service.

```
/etc/init.d/ntp start
```

- Make NTP run automatically after a reboot.

```
# chkconfig --add ntp
# chkconfig ntp on
```

- Check that the NTP server can be reached.

```
# ntpq -p
```

The `reach` column should not show 0. The `st` column should show some number other than 16..

Obtain the Certificate Thumbprint from a vIDM Host

Before you configure the integration of vIDM with NSX-T Data Center, you must get the certificate thumbprint from the vIDM host.

You must use OpenSSL version 1.x or higher for the thumbprint. On a vIDM host of version 3.3.2 or earlier, the command `openssl` might be running an older version of OpenSSL. In that case, you must use the command `openssl1`. This command is only available on a vIDM host.

You can check your version of OpenSSL with the following command:

```
openssl version
```

On a server that is not the vIDM host, you can use the `openssl` command that is running OpenSSL version 1.x or later.

Procedure

- 1 Log in at the vIDM host's console, or SSH to the vIDM host as the user **sshuser**, or log in to any server that can ping the vIDM host.
- 2 Run one of the following commands to get the thumbprint of the vIDM host.
 - If you are logged in to a server that can ping the vIDM host, run the `openssl` command to get the thumbprint:

```
openssl s_client -connect <FQDN of vIDM host>:443 < /dev/null 2> /dev/null | openssl  
x509 -sha256 -fingerprint -noout -in /dev/stdin
```

- If you are logged in to the vIDM host, do one of the following:
 - If the OpenSSL version is 0.9.x or earlier, run the following command:

```
openssl1 s_client -connect <FQDN of vIDM host>:443 < /dev/null 2> /dev/null |  
openssl x509 -sha256 -fingerprint -noout -in /dev/stdin
```

If you get an error running the command, you might need to run `openssl1` with the `sudo` command, that is, `sudo openssl1 ....`

- If the OpenSSL version is 1.x or later, run the following command:

```
openssl s_client -connect <FQDN of vIDM host>:443 < /dev/null 2> /dev/null |  
openssl x509 -sha256 -fingerprint -noout -in /dev/stdin
```

If you get an error running the command, you might need to run `openssl` with the `sudo` command, that is, `sudo openssl ....`

Configure VMware Identity Manager/Workspace ONE Access Integration

You can integrate NSX-T Data Center with VMware Identity Manager (vIDM), which provides identity management services. The vIDM deployment can be a standalone vIDM host or a vIDM cluster.

Note: The new product name for VMware Identity Manager is VMware Workspace ONE Access.

The vIDM host or all the vIDM cluster components should have a certificate signed by a certificate authority (CA). Otherwise, logging in to vIDM from NSX Manager might not work with certain browsers, such as Microsoft Edge or Internet Explorer 11. For information about installing a CA-signed certificate on vIDM, see the VMware Identity Manager documentation at <https://docs.vmware.com/en/VMware-Identity-Manager/index.html>.

When you register NSX Manager with vIDM, you specify a redirect URI that points to NSX Manager. You can provide either the fully qualified domain name (FQDN) or the IP address. It is important to remember whether you use the FQDN or the IP address. When you try to log in to NSX Manager through vIDM, you must specify the host name in the URL the same way, that is, if you use the FQDN when registering the manager with vIDM, you must use the FQDN in the URL, and if you use the IP address when registering the manager with vIDM, you must use the IP address in the URL. Otherwise, login will fail.

If NSX-T API access is needed, one of the following configurations must be true:

- vIDM has a known CA-signed certificate.
- vIDM has the connector CA certificate trusted on the vIDM service side.

- vIDM uses outbound connector mode.

Note NSX Managers and vIDM must be in the same time zone. The recommended way is to use UTC.

You must configure your DNS servers to have PTR records if you are not using Virtual IP or an external load balancer (this means that the manager is configured using the physical IP or FQDN of the node).

If you configure vIDM to be integrated with an external load balancer, you must enable session persistence on the load balancer to avoid issues such as pages not loading or a user being unexpectedly logged out.

If the vIDM deployment is a vIDM cluster, the vIDM load balancer must be configured for SSL termination and re-encryption.

With vIDM enabled, you can still log in to NSX Manager with a local user account if you use the URL `https://<nsx-manager-ip-address>/login.jsp?local=true`.

If you use the UserPrincipalName (UPN) to log in to vIDM, authentication to NSX-T might fail. To avoid this issue, use a different type of credentials, for example, SAMAccountName.

If using NSX Cloud, you can log in to CSM separately using the URL `https://<csm-ip-address>/login.jsp?local=true`

Prerequisites

- Verify that you have the certificate thumbprint from the vIDM host or the vIDM load balancer, depending on the type of vIDM deployment (a standalone vIDM host or a vIDM cluster). The command to obtain the thumbprint is the same in both cases. See [Obtain the Certificate Thumbprint from a vIDM Host](#).
- Verify that NSX Manager is registered as an OAuth client to vIDM. During the registration process, note the client ID and the client secret. For more information, see the VMware Identity Manager documentation at <https://docs.vmware.com/en/VMware-Workspace-ONE-Access/3.3/idm-administrator/GUID-AD4B6F91-2D68-48F2-9212-5B69D40A1FAE.html>. When you create the client, you only need to do the following:
 - Set **Access Type** to **Service Client Token**.
 - Specify a client ID.
 - Expand the **Advanced** field and click **Generate Shared Secret**.
 - Click **Add**.

NSX Cloud Note If using NSX Cloud, also verify that CSM is registered as an OAuth client to vIDM.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **System > Users and Roles**.
- 3 Click the **VMware Identity Manager** tab.
- 4 Click **Edit**.
- 5 To enable external load balancer integration, click the **External Load Balancer Integration** toggle.

Note If you have Virtual IP (VIP) set up (check **System > Appliances > Virtual IP**), you cannot use the **External Load Balancer Integration** even if you enable it. This is because you can either have VIP or the External Load Balancer while configuring vIDM but not both. Disable VIP if you want to use the External Load Balancer. See [Configure a Virtual IP \(VIP\) Address for a Cluster](#) in the *NSX-T Data Center Installation Guide* for details.

- 6 To enable VMware Identity Manager integration, click the **VMware Identity Manager Integration** toggle.
- 7 Provide the following information.

Parameter	Description
VMware Identity Manager Appliance	The fully qualified domain name (FQDN) of the vIDM host or the vIDM load balancer, depending on the type of vIDM deployment (a standalone vIDM host or a vIDM cluster).
OAuth Client ID	The ID that is created when registering NSX Manager to vIDM.
OAuth Client Secret	The secret that is created when registering NSX Manager to vIDM.
SSL Thumbprint	The certificate thumbprint of the vIDM host. It must be an SHA-256 thumbprint.
NSX Appliance	The IP address or fully qualified domain name (FQDN) of NSX Manager. If you are using an NSX Manager cluster, use the load balancer FQDN or cluster VIP FQDN or IP address. If you specify a FQDN, you must access NSX Manager from a browser using the manager's FQDN in the URL, and if you specify an IP address, you must use the IP address in the URL. Alternatively, the vIDM administrator can configure the NSX Manager client so that you can connect using either the FQDN or the IP address.

- 8 Click **Save**.
- 9 If using NSX Cloud, repeat steps 1 through 8 from the CSM appliance by logging in to CSM instead of NSX Manager.

Validate VMware Identity Manager Functionality

After configuring VMware Identity Manager, validate the functionality. Unless VMware Identity Manager is properly configured and validated, some users may receive Not Authorized (Error Code 98) messages when trying to log in.

Unless VMware Identity Manager is properly configured and validated, some users may receive Not Authorized (Error Code 98) messages when trying to log in.

Procedure

- 1 Create a base64 encoding of the username and password.

Run the following command to get the encoding and remove the trailing '\n' character. For example:

```
echo -n 'sfadmin@ad.node.com:password1234!' | base64 | tr -d '\n'
c2ZhZG1pbkZhZC5ub2RlLmNvbTpwYXNzd29yZDEyMzQhCg==
```

- 2 Verify that each user can make API call to each node.

Use a Remote Authorization curl command: `curl -k -H 'Authorization: Remote <base64 encoding string>' https://<node FQDN>/api/v1/node/aaa/auth-policy`. For example:

```
curl -k -H 'Authorization: Remote c2ZhZG1pbkZhZC5ub2RlLmNvbTpwYXNzd29yZDEyMzQhCg==' /
https://tmgr1.cptroot.com/api/v1/node/aaa/auth-policy
```

This returns the authorization policy settings, such as:

```
{
  "_schema": "AuthenticationPolicyProperties",
  "_self": {
    "href": "/node/aaa/auth-policy",
    "rel": "self"
  },
  "api_failed_auth_lockout_period": 900,
  "api_failed_auth_reset_period": 900,
  "api_max_auth_failures": 5,
  "cli_failed_auth_lockout_period": 900,
  "cli_max_auth_failures": 5,
  "minimum_password_length": 12
}
```

If the command does not return an error, the VMware Identity Manager is working correctly. No further steps are required. If the curl command returns an error, the user may be locked out.

Note Account lockout policies are set and enforced on a per node basis. If one node in the cluster has locked out a user, other nodes may have not.

3 To reset a user lockout on a node:

- a Retrieve the authorization policy using the local NSX Manager admin user:

```
curl -k -u 'admin:<password>' https://nsxmgr/api/v1/node/aaa/auth-policy
```

- b Save the output to a JSON file in current working directory.
- c Modify the file to change lockout period settings.

For example, many of the default settings apply lockout and reset periods of 900 seconds. Change these values to enable immediate reset, such as:

```
{
  "_schema": "AuthenticationPolicyProperties",
  "_self": {
    "href": "/node/aaa/auth-policy",
    "rel": "self"
  },
  "api_failed_auth_lockout_period": 1,
  "api_failed_auth_reset_period": 1,
  "api_max_auth_failures": 5,
  "cli_failed_auth_lockout_period": 1,
  "cli_max_auth_failures": 5,
  "minimum_password_length": 12
}
```

- d Apply the change to the affected node.

```
curl -k -u 'admin:<password>' -H 'Content-Type: application/json' -d \
@<modified_policy_setting.json> https://nsxmgr/api/v1/node/aaa/auth-policy
```

- e (Optional) Return the authorization policy settings files to its previous settings.

This should resolve the lockout issue. If you can still make remote auth API calls, but are still unable to log in through the browser, the browser may have an invalid cache or cookie stored. Clear your cache and cookies, and try again.

Integration with LDAP

You can configure NSX Manager to authenticate users using a directory service such as Active Directory over LDAP or OpenLDAP.

If you are using Active Directory (AD), and your AD forest is comprised of multiple subdomains, you should point NSX-T Data Center at your AD Global Catalog (GC) and configure each subdomain as an alternative domain name in NSX. The Global Catalog service usually runs on your primary AD domain controllers, and is a read-only copy of the most important information from all the primary and secondary domains. The GC service runs on port 3268 (plaintext), and 3269 (LDAP over TLS, encrypted).

For example, if your primary domain is "example.com" and you have subdomains "americas.example.com" and "emea.example.com", you should:

- 1 Configure NSX to use either the LDAP protocol on port 3268 or the LDAPS protocol on port 3269.
- 2 Add alternative domain names "americas.example.com" and "emea.example.com" in the NSX LDAP configuration.

Users in one of the subdomains must log in using the appropriate domain in their login name. For example, user "john" in the emea.example.com domain, must log in with the username "john@emea.example.com".

Note LDAP integration is not supported on Global Manager (Federation).

LDAP Identity Source

NSX Manager acts as an LDAP client, and interfaces with LDAP servers.

Three identity sources can be configured for user authentication. When a user logs into NSX Manager, the user is authenticated against the appropriate LDAP server of the user's domain. The LDAP server responds back with the authentication results, and the user group information. Once successfully authenticated, the user is assigned the roles corresponding to the groups that they belong to.

NSX Manager does not support multiple LDAP servers behind a load balancer, and LDAPS or StartTLS. If LDAP servers are behind a load balancer, configure NSX to connect directly to one of the LDAP servers, and not the load balancer virtual IP address.

Note Nested Active Directory groups with the parent group mapped as an NSX role is not supported.

Procedure

- 1 Navigate to **System > Users and Roles > LDAP**.
- 2 Click **Add Identity Source**.
- 3 Enter a **name** for the identity source.
- 4 Enter the **domain name** This must correspond to the domain name of your Active Directory server, if using Active Directory.
- 5 Select the type: either **Active Directory over LDAP** or **Open LDAP**.
- 6 Click **set** to configure LDAP servers. One LDAP server is supported for each domain.

Hostname/IP	The hostname or IP address of your LDAP server.
LDAP Protocol	Select the protocol : LDAP (unsecured) or LDAPS (secured).

Port	The default port is populated based on the selected protocol. If your LDAP server is running on a non-standard port, you can edit this text box to give the port number.
Connection Status	After filling in the mandatory text boxes, including the LDAP server information, you can click this to test the connection.
Use StartTLS	If selected, the LDAPv3 StartTLS extension is used to upgrade the connection to use encryption. To determine if you should use this option, consult your LDAP server administrator. This option can only be used if LDAP protocol is selected.
Certificate	If you are using LDAPS or LDAP + StartTLS, this text box should contain the PEM-encoded X.509 certificate of the server. If you leave this text box blank and click the Check Status link, NSX connects to the LDAP server. NSX then retrieves the LDAP server's certificate, and asks if you want to trust that certificate. If you have verified that the certificate is correct, click OK , and the certificate text box will be populated with the retrieved certificate.
Bind Identity	The format is user@domainName, or you can specify the distinguished name. For Active Directory, you can use either the userPrincipalName (user@domainName) or the distinguished name. For OpenLDAP, you must supply a distinguished name. This text box is required unless your LDAP server supports anonymous bind, then it is optional. Consult your LDAP server administrator if you are not sure.
Password	Enter a password for the LDAP server. This text box is required unless your LDAP server supports anonymous bind, then it is optional. Consult your LDAP server administrator.

7 Click **Add**.

8 Enter the **Base Domain**.

A base distinguished name (Base DN) is needed to add an Active Directory domain. A Base DN is the starting point that an LDAP server uses when searching for users authentication within an Active Directory domain. For example, if your domain name is corp.local the DN for the Base DN for Active Directory would be "DC=corp,DC=local".

All of the user and group entries you intend to use to control access to NSX-T Data Center must be contained within the LDAP directory tree rooted at the specified Base DN. If the Base DN is set to something too specific, such as an Organizational Unit deeper in your LDAP tree, NSX may not be able to find the entries it needs to locate users and determine group membership. Selecting a broad Base DN is a best practice if you are unsure.

9 Your NSX-T Data Center end users can now log in using their login name followed by @ and the domain name of your LDAP server, **user_name@domain_name**.

What to do next

Assign roles to users and groups. See [Add a Role Assignment or Principal Identity](#).

Add a Role Assignment or Principal Identity

You can assign roles to users or user groups if VMware Identity Manager is integrated with NSX-T Data Center, or if you have LDAP as an authentication provider. You can also assign roles to principal identities.

A principal is an NSX-T Data Center component or a third-party application such as an OpenStack product. With a principal identity, a principal can use the identity name to create an object and ensure that only an entity with the same identity name can modify or delete the object. A principal identity has the following properties:

- Name
- Node ID - this can be any alphanumeric value assigned to a principal identity
- Certificate
- RBAC role indicating the access rights of this principal

Users (local, remote, or principal identity) with the Enterprise Administrator role can modify or delete objects owned by principal identities. Users (local, remote, or principal identity) without the Enterprise Administrator role cannot modify or delete protected objects owned by principal identities, but can modify or delete unprotected objects.

If a principal identity user's certificate expires, you must import a new certificate and make an API call to update the principal identity user's certificate (see the procedure below). For more information about the NSX-T Data Center API, a link to the API resource is available at <https://docs.vmware.com/en/VMware-NSX-T-Data-Center>.

A principal identity user's certificate must satisfy the following requirements:

- SHA256 based.
- RSA/DSA message algorithm with 2048 bits or above key size.
- Cannot be a root certificate.

You can delete a principal identity using the API. However, deleting a principal identity does not automatically delete the corresponding certificate. You must delete the certificate manually.

Steps to delete a principal identity and its certificate:

- 1 Get the details of the principal identity to delete and note the `certificate_id` value in the response.

```
GET /api/v1/trust-management/principal-identities/<principal-identity-id>
```

- 2 Delete the principal identity.

```
DELETE /api/v1/trust-management/principal-identities/<principal-identity-id>
```

- 3 Delete the certificate using the `certificate_id` value obtained in step 1.

```
DELETE /api/v1/trust-management/certificates/<certificate_id>
```

For LDAP, you configure user groups to user roles mapping information; the groups correspond to the user groups specified in the Active Directory (AD). To grant a user permissions on NSX, add that user to the mapped group in AD.

Prerequisites

You must have an authentication provider configured:

- For role assignment for vIDM, verify that a vIDM host is associated with NSX-T. For more information, see [Configure VMware Identity Manager/Workspace ONE Access Integration](#).
- For role assignment for LDAP, verify that you have an LDAP identity source. For more information, see [LDAP Identity Source](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **System > Users and Roles**.
- 3 To assign roles to users, select **Add > Role Assignment for vIDM**.
 - a Select a user or user group.
 - b Select a role.
 - c Click **Save**.
- 4 To add a principal identity, select **Add > Principal Identity with Role**.
 - a Enter a name for the principal identity.
 - b Select a role.
 - c Enter a node ID.
 - d Enter a certificate in PEM format.
 - e Click **Save**.
- 5 To add a role assignment for LDAP select **Add > Role Assignment for LDAP**.
 - a Select a domain.
 - b Enter the first few characters of the user's name, login ID, or a group name to search the LDAP directory, then select a user or group from the list that appears.
 - c Select a role.
 - d Click **Save**.
- 6 (Optional) If using NSX Cloud, log in to the CSM appliance instead of NSX Manager and repeat steps 1 through 4.

- 7 If the certificate for the principal identity expires, perform the following steps:
- Import a new certificate and note the certificate's ID. See [Import a Self-signed or CA-signed Certificate](#).

- Call the following API to get the ID of the principal identity.

```
GET https://<nsx-mgr>/api/v1/trust-management/principal-identities
```

- Call the following API to update the principal identity's certificate. You must provide the imported certificate's ID and the principal identity user's ID.

For example,

```
POST https://<nsx-mgr>/api/v1/trust-management/principal-identities?
action=update_certificate
{
  "principal_identity_id": "ebd3032d-728e-44d4-9914-d4f81c9972cb",
  "certificate_id" : "abd3032d-728e-44d4-9914-d4f81c9972cc"
}
```

Configuring Both vIDM and LDAP or Transitioning from vIDM to LDAP

If you have configured vIDM as the authentication server, you can add LDAP as an additional authentication server. You can also disable vIDM and use LDAP exclusively.

To configure vIDM integration, see [Integration with VMware Identity Manager/Workspace ONE Access](#). To configure LDAP integration, see [Integration with LDAP](#).

If you have both vIDM and LDAP integration configured, the URL for the login page for vIDM users is `https://<nsx-manager-ip-address>`. Users will be redirected to the vIDM login page. The URL for the login page for LDAP users is `https://<nsx-manager-ip-address>/login.jsp?local=true` and the login name must be in the format `user_name@domain_name`.

If you only have LDAP integration configured, the URL for the login page for vIDM users is `https://<nsx-manager-ip-address>` and the login name must be in the format `user_name@domain_name`.

If you have vIDM integration configured and want to transition to using LDAP only, first configure LDAP integration. The AD servers must be the same as the AD servers used in vIDM. Then disable vIDM on the vIDM configuration page. The roles, users, and role assignments created in vIDM will exist in LDAP.

Role-Based Access Control

With role-based access control (RBAC), you can restrict system access to authorized users. Users are assigned roles and each role has specific permissions.

There are four types of permissions:

- Full access (Create, Read, Update, and Delete)
- Execute (Read, Update)
- Read
- None

Full access gives the user all permissions.

NSX-T Data Center has the following built-in roles. You cannot add any new roles.

- Enterprise Administrator
- Auditor
- Network Engineer
- Network Operations
- Security Engineer
- Security Operations
- Load Balancer Administrator
- Load Balancer Auditor
- VPN Administrator
- Guest Introspection Administrator
- Network Introspection Administrator

To view the built-in roles and the associated permissions, navigate to **System > Users and Roles > Roles**.

After an Active Directory (AD) user is assigned a role, if the username is changed on the AD server, you need to assign the role again using the new username.

Roles and Permissions

[Table 18-1. Roles and Permissions](#) and [Table 18-2. Roles and Permissions for Manager Mode](#) show the permissions each role has for different operations. The following abbreviations are used:

- EA - Enterprise Administrator
- A - Auditor
- NE - Network Engineer
- NO - Network Operations
- SE - Security Engineer
- SO - Security Operations
- LB Adm - Load Balancer Administrator

- LB Aud - Load Balancer Auditor
- VPN Adm - VPN Administrator
- GI Adm - Guest Introspection Administrator
- NI Adm - Network Introspection Administrator
- FA - Full access
- E - Execute
- R - Read

Table 18-1. Roles and Permissions

Operation	EA	A	NE	NO	SE	SO	CS Adm	CS Aud	LB Adm	LB Aud	VPN Adm	GI Adm	NI Adm
Networking > Tier-0 Gateways	FA	R	FA	R	R	R	FA	R	R	R	R	R	R
Networking > Tier-1 Gateways	FA	R	FA	R	R	R	FA	R	R	R	R	R	R
Networking > Network Interface	FA	R	FA	FA	R	R	FA	R	R	R	R	R	R
Networking > Network Static Routes	FA	R	FA	FA	R	R	FA	R	R	R	R	R	R
Networking > Locale Services	FA	R	FA	FA	R	R	FA	R	R	R	R	R	R
Networking > Static ARP Configuration	FA	R	FA	FA	R	R	FA	R	R	R	R	R	R

Table 18-1. Roles and Permissions (continued)

Operati on	EA	A	NE	NO	SE	SO	CS Adm	CS Aud	LB Ad m	LB Au d	VPN Adm	GI Adm	NI Adm
Networ king > Segmen ts	FA	R	FA	R	R	R	FA	R	R	R	R	R	R
Networ king > Segmen ts > Segmen t Profiles	FA	R	FA	R	R	R	FA	R	R	R	R	R	R
Networ king > IP Address Pools	FA	R	FA	FA	R	R	FA	R	R	R	None	None	None
Networ king Forwar ding Policies	FA	R	FA	R	FA	R	FA	R	Non e	No ne	None	None	None
Networ king > DNS	FA	R	FA	FA	R	R	FA	R	R	R	None	None	None
Networ king > DHCP	FA	R	FA	R	R	R	FA	R	R	R	None	None	None
Networ king > Load Balancin g	FA	R	None	None	R	None	FA	R	FA	R	None	None	None
Networ king > NAT	FA	R	FA	R	FA	R	FA	R	R	R	None	None	None
Networ king > VPN	FA	R	FA	R	FA	R	FA	R	Non e	No ne	FA	None	None
Networ king > IPv6 Profiles	FA	R	FA	R	R	R	FA	R	R	R	None	None	None

Table 18-1. Roles and Permissions (continued)

Operati on	EA	A	NE	NO	SE	SO	CS Adm	CS Aud	LB Ad m	LB Au d	VPN Adm	GI Adm	NI Adm
Security > Distribu ted Firewall	FA	R	R	R	FA	R	FA	R	R	R	R	R	R
Security > Gatewa y Firewall	FA	R	R	R	FA	R	FA	R	Non e	No ne	None	None	FA
Security > Networ k Introspe ction	FA	R	R	R	R	R	FA	R	Non e	No ne	None	None	FA
Security > Endpoin t Protecti on Rules	FA	R	R	R	R	R	FA	R	Non e	No ne	None	FA	None
Inventor y > Context Profiles	FA	R	FA	R	FA	R	FA	R	R	R	R	R	R
Inventor y > Virtual Machine s	R	R	R	R	R	R	R	R	R	R	R	R	R
Inventor y > Virtual Machine s > Create & Assign Tags to VM	FA	R	R	R	FA	R	FA	R	R	R	R	FA	FA
Inventor y > Contain ers	FA	R	R	R	R	R	None	Non e	Non e	No ne	None	None	None

Table 18-1. Roles and Permissions (continued)

Operati on	EA	A	NE	NO	SE	SO	CS Adm	CS Aud	LB Ad m	LB Au d	VPN Adm	GI Adm	NI Adm
Inventor y > Physical Servers	FA	R	R	R	R	R	R	R	R	R	None	None	None
Plan & Trouble shoot > Port Mirrorin g	FA	R	FA	R	R	R	FA	R	Non e	No ne	None	None	None
Plan & Trouble shoot > Port Mirrorin g Binding	FA	R	FA	FA	R	R	FA	R	R	R	R	R	R
Plan & Trouble shoot > Monitori ng Profile Binding	FA	R	FA	FA	R	R	FA	R	R	R	R	R	R
Plan & Trouble shoot > IPFIX > Firewall IPFIX Profiles	FA	R	FA	R	FA	R	FA	R	R	R	R	R	R
Plan & Trouble shoot > IPFIX > Switch IPFIX Profiles	FA	R	FA	R	R	R	FA	R	R	R	R	R	R
System > Fabric > Nodes > Hosts	FA	R	R	R	R	R	R	R	Non e	No ne	None	None	None
System > Fabric > Nodes > Nodes	FA	R	FA	R	FA	R	R	R	R	R	None	None	None

Table 18-1. Roles and Permissions (continued)

Operation	EA	A	NE	NO	SE	SO	CS Adm	CS Aud	LB Adm	LB Aud	VPN Adm	GI Adm	NI Adm
System > Fabric > Nodes > Edges	FA	R	FA	R	R	R	R	R	None	None	None	None	None
System > Fabric > Nodes > Edge Clusters	FA	R	FA	R	R	R	R	R	None	None	None	None	None
System > Fabric > Nodes > Bridges	FA	R	FA	R	R	R	None	None	R	R	None	None	None
System > Fabric > Nodes > Transport Nodes	FA	R	R	R	R	R	R	R	R	R	None	None	None
System > Fabric > Nodes > Tunnels	R	R	R	R	R	R	R	R	R	R	None	None	None
System > Fabric > Profiles > Uplink Profiles	FA	R	R	R	R	R	R	R	R	R	None	None	None
System > Fabric > Profiles > Edge Cluster Profiles	FA	R	FA	R	R	R	R	R	R	R	None	None	None
System > Fabric > Profiles > Configuration	FA	R	None	None	None	None	R	R	None	None	None	None	None

Table 18-1. Roles and Permissions (continued)

Operation	EA	A	NE	NO	SE	SO	CS Adm	CS Aud	LB Adm	LB Aud	VPN Adm	GI Adm	NI Adm
System > Fabric > Transport Zones > Transport Zones	FA	R	R	R	R	R	R	R	R	R	None	None	None
System > Fabric > Transport Zones > Transport Zone Profiles	FA	R	R	R	R	R	R	R	None	None	None	None	None
System > Fabric > Compute Managers	FA	R	R	R	R	R	R	R	None	None	None	R	R
System > Certificates	FA	R	None	None	FA	R	None	None	FA	R	FA	None	None
System > Service Deployments > Service Instances	FA	R	R	R	FA	R	FA	R	None	None	None	FA	FA
System > Utilities > Support Bundle	FA	R	None	None	None	None	None	None	None	None	None	None	None
System > Utilities > Backup	FA	R	None	None	None	None	None	None	None	None	None	None	None

Table 18-1. Roles and Permissions (continued)

Operati on	EA	A	NE	NO	SE	SO	CS Adm	CS Aud	LB Ad m	LB Au d	VPN Adm	GI Adm	NI Adm
System > Utilities > Restore	FA	R	None	None	None	None	None	None	None	None	None	None	None
System > Utilities > Upgrad e	FA	R	R	R	R	R	None	None	None	None	None	None	None
System > Users > Role Assign ments	FA	R	None	None	None	None	None	None	None	None	None	None	None
System > Active Director y	FA	R	FA	R	FA	FA	R	R	R	R	R	R	R
System > Users > Configu ration	FA	R	None	None	None	None	None	None	None	None	None	None	None
System > License s	FA	R	R	R	R	R	None	None	None	None	None	None	None
System > System Adminis tration	FA	R	R	R	R	R	R	R	None	None	None	None	None

Table 18-1. Roles and Permissions (continued)

Operation	EA	A	NE	NO	SE	SO	CS Adm	CS Aud	LB Adm	LB Aud	VPN Adm	GI Adm	NI Adm
Custom Dashboard Configuration	FA	R	R	R	R	R	FA	R	R	R	R	R	R
System > Lifecycle Management > Migrate	FA	None	None	None	None	None	None	None	None	None	None	None	None

Table 18-2. Roles and Permissions for Manager Mode

Operation	EA	A	NE	NO	SE	SO	CS Adm	CS Aud	LB Adm	LB Aud	VPN Adm	GI Adm	NI Adm
Plan & Troubleshoot > Port Connection	E	R	E	E	E	E	E	R	E	E	None	None	None
Plan & Troubleshoot > Traceflow	E	R	E	E	E	E	E	R	E	E	None	None	None
Plan & Troubleshoot > Port Mirroring	FA	R	FA	R	R	R	FA	R	None	None	None	None	None
Plan & Troubleshoot > IPFIX	FA	R	FA	R	FA	R	FA	R	R	R	R	R	R
Security > Distributed Firewall > General	FA	R	R	R	FA	R	FA	R	None	None	None	None	R

Table 18-2. Roles and Permissions for Manager Mode (continued)

Operation	EA	A	NE	NO	SE	SO	CS Adm	CS Aud	LB Adm	LB Aud	VPN Adm	GI Adm	NI Adm
Security > Distributed Firewall > Configuration	FA	R	R	R	FA	R	FA	R	None	None	None	None	None
Security > Edge Firewall	FA	R	R	R	FA	R	FA	R	None	None	None	None	FA
Networking > Routers	FA	R	FA	FA	R	R	FA	R	R	R	R	None	R
Networking > NAT	FA	R	FA	R	FA	R	FA	R	R	R	None	None	None
Networking > DHCP > Server Profiles	FA	R	FA	R	None	None	FA	R	None	None	None	None	None
Networking > DHCP > Servers	FA	R	FA	R	None	None	FA	R	None	None	None	None	None
Networking > DHCP > Relay Profiles	FA	R	FA	R	None	None	FA	R	None	None	None	None	None
Networking > DHCP > Relay Services	FA	R	FA	R	None	None	FA	R	None	None	None	None	None
Networking > DHCP > Metadata Proxies	FA	R	FA	R	None	None	None	None	None	None	None	None	None
Networking > IPAM	FA	R	FA	FA	R	R	None	None	R	R	None	None	None

Table 18-2. Roles and Permissions for Manager Mode (continued)

Operati on	EA	A	NE	NO	SE	SO	CS Adm	CS Aud	LB Ad m	LB Au d	VPN Adm	GI Adm	NI Adm
Network king > Logical Switches > Switches	FA	R	FA	FA	R	R	FA	R	R	R	R	None	R
Network king > Logical Switches > Ports	FA	R	FA	FA	R	R	FA	R	R	R	R	None	R
Network king > Logical Switches > Switchin g Profiles	FA	R	FA	FA	R	R	FA	R	R	R	None	None	None
Network king > Load Balancin g > Load Balance rs	FA	R	None	None	R	None	FA	R	FA	R	None	None	None
Network king > Load Balancin g > Profiles > SSL Profiles	FA	R	None	None	FA	R	FA	R	FA	R	None	None	None
Inventor y > Groups	FA	R	FA	R	FA	R	FA	R	R	R	R	R	R
Inventor y > Groups > IP Sets	FA	R	FA	R	FA	R	FA	R	R	R	R	R	R

Table 18-2. Roles and Permissions for Manager Mode (continued)

Operation	EA	A	NE	NO	SE	SO	CS Adm	CS Aud	LB Adm	LB Aud	VPN Adm	GI Adm	NI Adm
Inventor y > Groups > IP Pools	FA	R	FA	R	None	None	None	None	R	R	R	R	R
Inventor y > Groups > MAC Sets	FA	R	FA	R	FA	R	FA	R	R	R	R	R	R
Inventor y > Services	FA	R	FA	R	FA	R	FA	R	R	R	R	R	R
Inventor y > Virtual Machine s	R	R	R	R	R	R	R	R	R	R	R	R	R
Inventor y > Virtual Machine s > Create & Assign Tags to VM	FA	R	R	R	FA	R	FA	R	R	R	R	FA	FA
Inventor y > Virtual Machine s > Configu re Tags	FA	None	None	None	None	None	None	None	None	None	None	None	None

Logging User Account Changes

Changes to a user's role assignment are automatically written to syslog and the audit log.

For more information about syslog and the audit log, see [Log Messages and Error Codes](#).

An example of a log message when assigning a role to a vIDM user:

```
2020-09-24T16:05:51.244Z nsxmanager-14663974-1-CertKB-FS NSX 5519 - [nsx@6876 audit="true"
comp="nsx-manager" entId="e3c2af75-9d0f-4020-90cc-f2f00d6af255" level="INFO"
```



```
reqId="b27711c6-0590-4b39-b8b6-f0980a0597f0" subcomp="policy" update="true" username="admin"]
UserName="admin", ModuleName="AAA", Operation="CreateRoleBinding", Operation
status="success", New
value=[{"name":"test_AU@idfw.local","type":"remote_user","identity_source_type":"VIDM","roles":
[{"role":"auditor"}]}, {"id":"bba634c9-cfbd-4806-a831-e63ec195e1f9","_protection":"UNKNOWN"}]
```

An example of a log message when updating the role of a vIDM user:

```
2020-09-24T16:12:51.217Z nsxmanager-14663974-1-CertKB-FS NSX 5519 -
[nsx@6876 audit="true" comp="nsx-manager" entId="e3c2af75-9d0f-4020-90cc-f2f00d6af255"
level="INFO" reqId="973faed4-f4b5-443d-bd79-7d995c027183" subcomp="policy" update="true"
username="admin"] UserName="admin", ModuleName="AAA", Operation="UpdateRoleBinding",
Operation status="success", New value=[{"e3c2af75-9d0f-4020-90cc-f2f00d6af255"
{"name":"test_AU@idfw.local","type":"remote_user","identity_source_type":"VIDM","roles":
[{"role":"security_engineer"}]}, {"_protection":"UNKNOWN"}]
```

An example of a log message when assigning a role to an LDAP user:

```
2020-09-24T16:06:28.663Z nsxmanager-14663974-1-CertKB-FS NSX 5519 - [nsx@6876 audit="true"
comp="nsx-manager" entId="35e45569-6da6-4dcd-b4a1-75747cdd6cf8" level="INFO"
reqId="db27f4ae-25a7-4482-b3f4-49228d12960b" subcomp="policy" update="true" username="admin"]
UserName="admin", ModuleName="AAA", Operation="CreateRoleBinding", Operation
status="success", New
value=[{"name":"skrasner@airius.com","type":"remote_user","identity_source_type":"LDAP","ident
ity_source_id":"ldap","roles":[{"role":"auditor"}]}, {"id":"dd8d3675-
c574-454b-975e-300b65462827","_protection":"UNKNOWN"}]
```

An example of a log message when updating the role of an LDAP user:

```
2020-09-24T16:12:37.449Z nsxmanager-14663974-1-CertKB-FS NSX 5519 - [nsx@6876 audit="true"
comp="nsx-manager" entId="35e45569-6da6-4dcd-b4a1-75747cdd6cf8" level="INFO"
reqId="d7cdd3de-75a1-4d29-9fea-27e1dda4b5e2" subcomp="policy" update="true" username="admin"]
UserName="admin", ModuleName="AAA", Operation="UpdateRoleBinding", Operation
status="success", New value=[{"35e45569-6da6-4dcd-b4a1-75747cdd6cf8"
{"name":"skrasner@airius.com","type":"remote_user","identity_source_type":"LDAP","identity_sou
rce_id":"ldap","roles":[{"role":"network_engineer"}]}, {"_protection":"UNKNOWN"}]
```

After you install NSX-T Data Center, the manager nodes and cluster have self-signed certificates.

If you are using Federation, additional certificates are set up to establish trust between the Local Managers and Global Manager.

You can import certificates, create a certificate signing request (CSR), generate self-signed certificates, and import a certificate revocation list (CRL). To improve security, it is recommended that you replace the self-signed certificates with CA-signed certificates.

Read the following topics next:

- [Types of Certificates](#)
- [Certificates for NSX Federation](#)
- [Create a Certificate Signing Request File](#)
- [Creating Self-signed Certificates](#)
- [Importing and Replacing Certificates](#)
- [Importing and Retrieving CRLs](#)
- [Storage of Public Certificates and Private Keys for Load Balancer or VPN service](#)
- [Alarm Notification for Certificate Expiration](#)

Types of Certificates

There are three categories of self-signed certificates in NSX-T Data Center.

- Platform Certificates
- NSX Services Certificates
- Principal Identity Certificates

Refer to the following sections for details on each certificate category.

Platform Certificates

After installing NSX-T Data Center, navigate to **System > Certificates** to view the platform certificates created by the system. By default these are self-signed X.509 RSA 2048/SHA256 certificates for internal communication within NSX-T Data Center and for external authentication when NSX Manager is accessed using APIs or the UI.

The internal certificates are not viewable or editable.

Table 19-1. Platform Certificates in NSX-T Data Center

Naming Convention in NSX Manager	Purpose	Replaceable?	Default Validity
tomcat	This is an API certificate used for external communication with individual NSX Manager nodes through UI/API.	Yes. See Replace Certificates	825 days
mp-cluster	This is an API certificate used for external communication with the NSX Manager cluster using the cluster VIP, through UI/API.	Yes. See Replace Certificates	825 days
Additional certificates	Certificates specifically for NSX Federation. If you are not using NSX Federation, these certificate are not used.	See Certificates for NSX Federation for details on self-signed certificates auto-configured for NSX Federation.	
Not visible in the UI	Certificates used for internal communication between different system components.	No	10 years

NSX Service Certificates

NSX service certificates are used for services such as load balancer and VPN.

NSX service certificates cannot be self signed. You must import them. See [Importing and Replacing Certificates](#) for instructions.

A certificate signing request (CSR) can be used as NSX service certificate if it is signed by CA (local CA or public CA like Verisign). Once the CSR is signed, you can import that signed certificate into NSX Manager. A CSR can be generated on NSX Manager or outside of NSX Manager. Note that the **Service Certificate** flag is disabled for CSRs generated on NSX Manager. Therefore, these signed CSRs cannot be used as service certificates, but only as platform certificates.

Platform and NSX service certificates are stored separately within the system and certificates imported as NSX service certificate cannot be used for platform or the reverse.

Principal Identity (PI) Certificates

PI certificates can be for services or for platform.

PI for Cloud Management Platforms (CMP), such as Openstack, uses X.509 certificates that are uploaded when onboarding a CMP as a client. For information on assigning roles to Principal Identity and replacing PI certificates, see [Add a Role Assignment or Principal Identity](#)

PI for NSX Federation uses X.509 platform certificates for the Local Manager and Global Manager appliances. See [Certificates for NSX Federation](#) for details on self-signed certificates auto-configured for NSX Federation.

Certificates for NSX Federation

The system creates certificates required for communication between NSX Federation appliances as well as for external communication.

By default, the Global Manager uses self-signed certificates for communicating with internal components and registered Local Managers, as well as for authentication for NSX Manager UI or APIs.

You can view the external (UI/API) and inter-site certificates in NSX Manager. The internal certificates are not viewable or editable.

Certificates for Global Manager and Local Managers

After you add a Local Manager into the Global Manager, all certificates that authenticate the Local Manager for external and internal communication are copied into the Global Manager and trust is established between the two systems. These certificates are also copied into each of the sites registered with the Global Manager.

See the following table for a list of all the certificates created for each appliance using NSX Federation, and the certificates these appliances exchange with each other:

Table 19-2. Certificates for the Global Manager and Local Managers

Naming Convention in the Global Manager or Local Manager	Purpose	Replaceable?	Default Validity
The following are certificates specific to each NSX Federation appliance.			
APH-AR certificate	■ For Global Manager and each Local Manager. ■ Used for inter-site communication using the AR channel (Async-Replicator channel).	No	10 years
GlobalManager	■ For the Global Manager. ■ PI certificate for the Global Manager.	Yes. See Replace Certificates .	825 days
mp-cluster certificate	■ For the Global Manager and each Local Manager. ■ Used for UI/API communication with the VIP of the Global Manager or Local Manager cluster.		

Table 19-2. Certificates for the Global Manager and Local Managers (continued)

Naming Convention in the Global Manager or Local Manager	Purpose	Replaceable?	Default Validity
tomcat certificate	■ For the Global Manager and each Local Manager. ■ Used for UI/API communication with individual Global Manager and Local Manager nodes for each of the locations added to the Global Manager.		
LocalManager	■ For Local Manager. ■ PI certificate for this specific Local Manager.		
The following are certificates exchanged between NSX Federation appliances.			
Naming Convention in the Global Manager or Local Manager	Purpose	Replaceable?	Default Validity
Hashed code, for example, 1729f966-67b7-4c17-bdf5-325affb79f4f	■ Exchanged between all the Local Managers registered with the Global Manager. ■ PI certificate for the Global Manager exchanged with Local Managers. ■ PI certificates for each of the locations exchanged with all registered Location Managers.	Not Applicable	
Site certificate CN=<>,O	■ Exchanged between all NSX Federation appliances: all registered Local Managers and the Global Manager. ■ All types of certificates.		

Principal Identity (PI) Users for NSX Federation

The following PI users with corresponding roles are created after you add a Local Manager to the Global Manager:

Table 19-3. Principal Identity (PI) Users Created for NSX Federation

NSX Federation Appliance	PI Username	PI User Role
Global Manager	LocalManagerIdentity One for each Local Manager registered with this Global Manager.	auditor
Local Manager	GlobalManagerIdentity	Enterprise Admin
	LocalManagerIdentity One for each Local Manager registered with the same Global Manager. Use the following API to get a list of all the Local Manager PI users because they are not visible in the UI: <pre>GET https://<local-mgr>/api/v1/trust-management/ principal-identities</pre>	auditor

Create a Certificate Signing Request File

Certificate signing request (CSR) is an encrypted text that contains specific information such as, organization name, common name, locality, and country. You send the CSR file to a certificate authority (CA) to apply for a digital identity certificate.

Prerequisites

Gather the information that you need to fill out the CSR file. You must know the FQDN of the server and the organizational unit, organization, city, state, and country.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **System > Certificates**.
- 3 Click the **CSRs** tab.
- 4 Click **Generate CSR**.
- 5 Complete the CSR file details.

Option	Description
Name	Assign a name for your certificate.
Common Name	Enter the fully qualified domain name (FQDN) of your server. For example, test.vmware.com.
Organization Name	Enter your organization name with applicable suffixes. For example, VMware Inc.
Organization Unit	Enter the department in your organization that is handling this certificate For example, IT department.

Option	Description
Locality	Add the city in which your organization is located. For example, Palo Alto.
State	Add the state in which your organization is located. For example, California.
Country	Add the country in which your organization is located. For example, United States (US).
Message Algorithm	Set the encryption algorithm for your certificate. RSA encryption - is used for digital signatures and encryption of the message. Therefore, it is slower than DSA when creating an encrypted token but faster to analyze and validate this token. This encryption is slower to decrypt and faster to encrypt.
Key Size	Set the key bits size of the encryption algorithm. The default value, 2048, is adequate unless you specifically need a different Key size. Other supported sizes are 3072 and 4096. Many CAs require a minimum value of 2048. Larger key sizes are more secure but have a greater impact on performance.
Description	Enter specific details to help you identify this certificate at a later date.

6 Click **Generate**.

A custom CSR appears as a link.

7 Select the CSR and click **Actions** to select one of the following options:

- **Self Sign Certificate**
- **Import Certificate for CSR**
- **Download the PEM**

If you selected **Download CSR PEM**, you can save the CSR PEM file for your records and CA submission. Use the contents of the CSR file to submit a certificate request to the CA in accordance with the CA enrollment process. For the other two options, refer the required topics.

Results

The CA creates a server certificate based on the information in the CSR file, signs it with its private key, and sends you the certificate. The CA also sends you a root CA certificate.

Creating Self-signed Certificates

Create a Self-Signed Certificate

You can create a self-signed certificate. However, using a self-signed certificate is less secure than using a trusted certificate.

When you use a self-signed certificate the client user receives a warning message such as, *Invalid Security Certificate*. The client user must then accept the self-signed certificate when first connecting to the server in order to proceed. Allowing client users to select this option provides reduced security than other authorization methods.

Prerequisites

Verify that a CSR is available. See [Create a Certificate Signing Request File](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **System > Certificates**.
- 3 Click the **CSRs** tab.
- 4 Select a CSR.
- 5 Select **Actions > Self Sign Certificate for CSR**.
- 6 Enter the number of days the self-sign certificate is valid.
The default is 825 days. Even if you change this value for previously generated self-signed certificate, the default value is displayed every time you generate a new certificate.
- 7 Click **Add**.

Results

The self-signed certificate appears in the **Certificates** tab.

Import a Certificate for a CSR

You can import a signed certificate of NSX-T Data Center generated CSR. This page provides the step to import a signed certificate of NSX generated CSR.

A self-signed certificate acts as a certificate as well as CA. It is not required to be signed from any external CA, whereas CSR is a certificate signing request that cannot act as CA and must be signed by external CA. Note that a self-signed certificate is not supported for LB.

When you use a self-signed certificate the client user receives a warning message such as, *Invalid Security Certificate*. The client user must then accept the self-signed certificate when first connecting to the server in order to proceed. Allowing client users to select this option provides reduced security than other authorization methods.

Prerequisites

- Verify that a CSR is available. See [Create a Certificate Signing Request File](#).
- NSX-T Data Center generated CSR was used as a CSR for signed certificate.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **System > Certificates**.
- 3 Click the **CSRs** tab.
- 4 Select a CSR.
- 5 Select **Actions > Import Certificate for CSR**.
- 6 Browse to the signed certificate file on your computer and add the file.
- 7 Click **Add**.

Results

The self-signed certificate appears in the **Certificates** tab.

Importing and Replacing Certificates

You can import self-signed or CA-signed certificates for platform or services. You can replace some of the self-signed certificates using APIs.

You can also import CA certificates for services such as Load Balancer.

Import a Self-signed or CA-signed Certificate

You can import a certificate with a private key to replace the default self-signed certificate, after activation.

You can import self-signed or CA-signed certificates for platform or services using this procedure. Note that a CSR generated on NSX Manager and which is self-signed cannot be used as a service certificate, such as the Load Balancer service. If you want to import a CA certificate for the Load Balancer service, see [Import a CA Certificate](#).

Prerequisites

- Verify that a certificate is available.
- The server certificate must contain the Basic Constraints extension `basicConstraints = cA:FALSE`.

Procedure

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

- 3 Select **Import > Import Certificate** and enter the certificate details.

Option	Description
Name	Assign a name to the certificate.
Certificate Contents	Browse to the certificate file on your computer and add the file. The certificate must not be encrypted. If it is a CA-signed certificate, be sure to include the whole chain in this order: certificate - intermediate - root.
Private Key	Browse to the private key file on your computer and add the file. This is an optional field if imported certificate is based on NSX Manager generated CSR as a private key already exists on the NSX Manager appliance.
Passphrase	Add a passphrase for this certificate if it is encrypted. In this release, this field is not used because encrypted certificate is not supported.
Description	Enter a description of what is included in this certificate.
Service Certificate	Set to Yes to use this certificate for services such as a load balancer and VPN. Set to No if this certificate is for the NSX Manager nodes.

- 4 Click **Import**.

Import a CA Certificate

You can import a CA certificate from a system external to NSX-T Data Center, for example, to use with the Load Balancer service.

If you want to import a self-signed or a CA-signed certificate, see instructions at [Import a Self-signed or CA-signed Certificate](#).

Prerequisites

Verify that a CA certificate is available.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **System > Certificates**.
- 3 Select **Import > Import CA Certificate** and enter the certificate details.

Option	Description
Name	Assign a name to the CA certificate.
Certificate Contents	Browse to the CA certificate file on your computer and add the file.
Description	Enter a summary of what is included in this CA certificate.
Service Certificate	Set to Yes to use this certificate for services such as a load balancer and VPN.

- 4 Click **Import**.

Set Checks for Certificate Imports

You can enable or disable the Extended Key Usage (EKU) Extension and the Certificate Revocation List Distribution Point (CDP) validation checks that NSX-T Data Center performs while importing a certificate.

Note: If you have CA-signed certificates without a CDP then you might have problems after upgrade. To avoid this problem you can turn CRL checking off or replace the certificates with certificates that include a CDP.

To set validation checks, use the following API with payload. For more information about the API, see the *NSX-T Data Center API Guide*.

```
PUT https://<manager>/api/v1/global-configs/SecurityGlobalConfig
{
  "crl_checking_enabled": false,
  "ca_signed_only": false,
  "eku_checking_enabled": false,
  "resource_type": "SecurityGlobalConfig"
}
```

Where:

- `crl_checking_enabled`: Enabled by default to check CDP specified in the imported CA-signed certificate. Support includes HTTP based CRL-DP only. File or LDAP-based options are not supported.
- `ca_signed_only`: Disabled by default. It allows checks signed by CA only.
- `eku_checking_enabled`: Disabled by default. It checks for EKU Extension in the imported certificate.
- `revision`: The current revision of the resource that must be included in a request. To obtain the value of this parameter issue a GET operation.

Replace Certificates

You can replace certificates for a manager node or the manager cluster virtual IP (VIP) by making an API call.

After you install NSX-T Data Center, the manager nodes and cluster have self-signed certificates. It is recommended that you replace the self-signed certificates with a CA-signed certificate and that you use a single common CA-signed certificate with a SAN (Subject Alternative Names) list that matches all the nodes and VIP for the cluster. See [Types of Certificates](#) for details on the default self-signed certificates configured by the system.

If you are using Federation, you can replace Global Manager nodes, Global Manager cluster, Local Manager nodes and Local Manager cluster certificates using the following APIs. You can also replace the platform Principal Identity certificates auto-created for the Global Manager and Local Manager appliances. See [Certificates for NSX Federation](#) for details on self-signed certificates auto-configured for Federation.

Prerequisites

- Verify that a certificate is available in the NSX Manager. See [Import a Self-signed or CA-signed Certificate](#).
- The server certificate must contain the Basic Constraints extension `basicConstraints = cA:FALSE`.
- Verify that the certificate is valid by making the following API call:

```
GET https://<nsx-mgr>/api/v1/trust-management/certificates/<certificate-id>?
action=validate
```

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **System > Certificates**.
- 3 In the ID column, click the ID of the certificate you want to use and copy the certificate ID from the pop-up window.

Make sure that when this certificate was imported, the option **Service Certificate** was set to **No**.

- 4 To replace the certificate of a manager node, use the `POST /api/v1/node/services/http?action=apply_certificate` API call. For example,

```
POST https://<nsx-mgr>/api/v1/node/services/http?
action=apply_certificate&certificate_id=e61c7537-3090-4149-b2b6-19915c20504f
```

Note: The certificate chain must be in the industry standard order of 'certificate - intermediate - root.'

For more information about the API, see the *NSX-T Data Center API Reference*.

- 5 To replace the certificate of the manager cluster VIP, use the `POST /api/v1/cluster/api-certificate?action=set_cluster_certificate` API call. For example,

```
POST https://<nsx-mgr>/api/v1/cluster/api-certificate?
action=set_cluster_certificate&certificate_id=d60c6a07-6e59-4873-8edb-339bf75711ac
```

Note: The certificate chain must be in the industry standard order of 'certificate - intermediate - root.'

For more information about the API, see the *NSX-T Data Center API Reference*. This step is not necessary if you did not configure VIP.

- 6 (Optional) To replace the Principal Identity certificates for Federation, use the API call: POST `https://<nsx-mgr>/api/v1/trust-management/certificates?action=set_pi_certificate_for_federation`. For example:

```
POST https://<nsx-mgr>/api/v1/trust-management/certificates?
action=set_pi_certificate_for_federation
{ "cert_id": "<id>",
  "service_type": "LOCAL_MANAGER" }
```

- 7 (Optional) If you currently have an NSX Intelligence appliance deployed with your NSX Manager cluster, you must update the NSX Manager node IP, certificate, and thumbprint information that is on the NSX Intelligence appliance. See VMware Knowledge Base article <https://kb.vmware.com/s/article/78505> for more information.

Importing and Retrieving CRLs

Import a Certificate Revocation List

A certificate revocation list (CRL) is a list of subscribers and their certificate status. When a potential user attempts to access a server, the server denies access based on the CRL entry for that particular user.

The list contains the following items:

- Revoked certificates and the reasons for revocation
- Dates the certificates are issued
- Entities that issued the certificates
- Proposed date for the next release

Prerequisites

Verify that a CRL is available.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **System > Certificates**.
- 3 Click the **CRLs** tab.

4 Click **Import** and add the CRL details.

Option	Description
Name	Assign a name to the CRL.
Certificate Contents	Copy all of the items in the CRL and paste them in this section. A sample CRL. <pre> -----BEGIN X509 CRL----- MIIBODCB4zANBgkqhkiG9w0BAQQFADBGMQswCQYDVQQGEwJBVTEEMMAoGA1 UECBMD UUxEMRkwFwYDVQQKEwBNaw5jb20gUHR5LiBMdGQuMQswCQYDVQQLEwJDUz EbMBkG A1UEAxMSU1NMZW55IGRlbW8gc2VydMVFw0wMTAxMTUxNjI2NTdaFw0wMT AyMTQx NjI2NTdaMFwiEgIBARcNOTUxMDA5MjMzMjA1WjASAgEDFw05NTEyMDEwMT AwMDBa MBMCAhI0Fw0wMTAxMTUxNjE5NDdaMBMCAhI1Fw0wMTAxMTUxNjIzNDZaMA OGCSqG SIB3DQEBBAUAA0EAHPjQ3M93QOj8Ufi+jZM7Y78TfAzG4jJn/ E6MYBPFVQFY0/Gp UZexfjSVo5CIyySOtYscz8o07avwBxTiMpDEQg== -----END X509 CRL-- </pre>
Description	Enter a summary of what is included in this CRL.

5 Click **Import**.

Results

The imported CRL appears as a link.

Configuring NSX Manager to Retrieve a Certificate Revocation List

Using the API, you can configure NSX Manager to retrieve a certificate revocation list (CRL). You can then check the CRL by making an API call to NSX Manager instead of to the certificate authority.

This feature provides the following benefits:

- It is more efficient to have the CRL cached on the server, that is, NSX Manager.
- The client does not need to create any outbound connection to the certificate authority.

The following APIs related to certificate revocation lists are available:

```

GET /api/v1/trust-management
GET /api/v1/trust-management/crl-distribution-points
POST /api/v1/trust-management/crl-distribution-points
DELETE /api/v1/trust-management/crl-distribution-points/<crl-distribution-point-id>
GET /api/v1/trust-management/crl-distribution-points/<crl-distribution-point-id>
PUT /api/v1/trust-management/crl-distribution-points/<crl-distribution-point-id>
GET /api/v1/trust-management/crl-distribution-points/<crl-distribution-point-id>/status
POST /api/v1/trust-management/crl-distribution-points/pem-file

```

You can manage CRL distribution points and retrieve the CRLs stored in NSX Manager. For more information, see the *NSX-T Data Center API Reference*.

Storage of Public Certificates and Private Keys for Load Balancer or VPN service

Public certificates and private keys are stored on the NSX Managers for load balancer or VPN service. When a load balancer or VPN service is created that requires a private key, NSX Manager sends a copy of the private key to the Edge node where the load balancer or VPN service is running.

Alarm Notification for Certificate Expiration

NSX-T Data Center generates alarms when a certificate is nearing its expiry or if a certificate has already expired.

NSX-T Data Center generates alarms under following events:

- Medium severity alarm starting 30 day before certificate expiry.
- High severity alarm starting 7 days prior to expiry.
- Critical severity alarm every day after certificate expires.

Certificate Expiry alarms contains details on certificate ID, severity, node, first/last report time, and recommended action.

As a remedial, you must replace the expiring External Platform certificate with a new valid certificate and delete expiring certificate.

Configuring NSX-T Data Center in Manager Mode

20

NSX-T Data Center has two user interface modes: Policy mode and Manager mode. If you have objects that were created in Manager mode, you should continue to use Manager mode to make changes.

For more information about the two modes, see [Chapter 1 NSX Manager](#).

If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Read the following topics next:

- [Logical Switches in Manager Mode](#)
- [Logical Routers in Manager Mode](#)
- [NAT in Manager Mode](#)
- [Grouping Objects in Manager Mode](#)
- [DHCP in Manager Mode](#)
- [IP Address Management in Manager Mode](#)
- [Load Balancing in Manager Mode](#)
- [Firewall in Manager Mode](#)

Logical Switches in Manager Mode

You can configure logical switches and related objects in **Manager** mode. A logical switch reproduces switching functionality, broadcast, unknown unicast, multicast (BUM) traffic, in a virtual environment decoupled from the underlying hardware.


Note If you use **Manager** mode to modify objects created in the **Policy** mode, some settings might not be configurable. These read-only settings have this icon next to them: . See [Chapter 1 NSX Manager](#) for more information.

Logical switches are similar to VLANs, in that they provide network connections to which you can attach virtual machines. The VMs can then communicate with each other over tunnels between hypervisors if the VMs are connected to the same logical switch. Each logical switch has a virtual network identifier (VNI), like a VLAN ID. Unlike VLAN, VNIs scale well beyond the limits of VLAN IDs.

To see and edit the VNI pool of values, log in to NSX Manager, navigate to **Fabric > Profiles**, and click the **Configuration** tab. Note that if you make the pool too small, creating a logical switch will fail if all the VNI values are in use. If you delete a logical switch, the VNI value will be re-used, but only after 6 hours.

When you add logical switches, it is important that you map out the topology that you are building.

Figure 20-1. Logical Switch Topology



For example, the topology above shows a single logical switch connected to two VMs. The two VMs can be on different hosts or the same host, in different host clusters or in the same host cluster. Because the VMs in the example are on the same virtual network, the underlying IP addresses configured on the VMs must be in the same subnet.

NSX Cloud Note If using NSX Cloud, see [NSX-T Data Center Features Supported with NSX Cloud](#) for a list of auto-generated logical entities, supported features, and configurations required for NSX Cloud.

Understanding BUM Frame Replication Modes

Each host transport node is a tunnel endpoint. Each tunnel endpoint has an IP address. These IP addresses can be in the same subnet or in different subnets, depending on your configuration of IP pools or DHCP for your transport nodes.

When two VMs on different hosts communicate directly, unicast-encapsulated traffic is exchanged between the two tunnel endpoint IP addresses associated with the two hypervisors without any need for flooding.

However, as with any Layer 2 network, sometimes traffic that is originated by a VM needs to be flooded, meaning that it needs to be sent to all of the other VMs belonging to the same logical switch. This is the case with Layer 2 broadcast, unknown unicast, and multicast traffic (BUM traffic). Recall that a single NSX-T Data Center logical switch can span multiple hypervisors. BUM traffic originated by a VM on a given hypervisor needs to be replicated to remote hypervisors that host other VMs that are connected to the same logical switch. To enable this flooding, NSX-T Data Center supports two different replication modes:

- Hierarchical two-tier (sometimes called MTEP)
- Head (sometimes called source)

Hierarchical two-tier replication mode is illustrated by the following example. Say you have Host A, which has VMs connected to virtual network identifiers (VNIs) 5000, 5001, and 5002. Think of VNIs as being similar to VLANs, but each logical switch has a single VNI associated with it. For this reason, sometimes the terms VNI and logical switch are used interchangeably. When we say a host is on a VNI, we mean that it has VMs that are connected to a logical switch with that VNI.

A tunnel endpoint table shows the host-VNI connections. Host A examines the tunnel endpoint table for VNI 5000 and determines the tunnel endpoint IP addresses for other hosts on VNI 5000.

Some of these VNI connections will be on the same IP subnet, also called an IP segment, as the tunnel endpoint on Host A. For each of these, Host A creates a separate copy of every BUM frame and sends the copy directly to each host.

Other hosts' tunnel endpoints are on different subnets or IP segments. For each segment where there is more than one tunnel endpoint, Host A nominates one of these endpoints to be the replicator.

The replicator receives from Host A one copy of each BUM frame for VNI 5000. This copy is flagged as Replicate locally in the encapsulation header. Host A does not send copies to the other hosts in the same IP segment as the replicator. It becomes the responsibility of the replicator to create a copy of the BUM frame for each host it knows about that is on VNI 5000 and in the same IP segment as that replicator host.

The process is replicated for VNI 5001 and 5002. The list of tunnel endpoints and the resulting replicators might be different for different VNIs.

With head replication also known as headend replication, there are no replicators. Host A simply creates a copy of each BUM frame for each tunnel endpoint it knows about on VNI 5000 and sends it.

If all the host tunnel endpoints are on the same subnet, the choice of replication mode does not make any difference because the behaviour will not differ. If the host tunnel endpoints are on different subnets, hierarchical two-tier replication helps distribute the load among multiple hosts. Hierarchical two-tier is the default mode.

Create a Logical Switch in Manager Mode

Logical switches attach to single or multiple VMs in the network. The VMs connected to a logical switch can communicate with each other using the tunnels between hypervisors.

Prerequisites

- Verify that a transport zone is configured. See the *NSX-T Data Center Installation Guide*.
- Verify that fabric nodes are successfully connected to NSX-T Data Center management plane agent (MPA) and NSX-T Data Center local control plane (LCP).

In the GET `https://<nsx-mgr>/api/v1/transport-nodes/<transport-node-id>/state` API call, the state must be `success`. See the *NSX-T Data Center Installation Guide*.

- Verify that transport nodes are added to the transport zone. See the *NSX-T Data Center Installation Guide*.
- Verify that the hypervisors are added to the NSX-T Data Center fabric and VMs are hosted on these hypervisors.
- Familiarize yourself with the logical switch topology and BUM frame replication concepts. See [Logical Switches in Manager Mode](#) and [Understanding BUM Frame Replication Modes](#).
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Logical Switches > Switches > Add**.
- 3 Enter a name for the logical switch and optionally a description.
- 4 Select a transport zone for the logical switch.

VMs that are attached to logical switches that are in the same transport zone can communicate with each other.

- 5 Enter the name of an uplink teaming policy.
- 6 Set **Admin Status** to either **Up** or **Down**.

7 Select a replication mode for the logical switch.

The replication mode (hierarchical two-tier or head) is required for overlay logical switches, but not for VLAN-based logical switches.

Replication Mode	Description
Hierarchical two-tier	The replicator is a host that performs replication of BUM traffic to other hosts within the same VNI. Each host nominates one host tunnel endpoint in every VNI to be the replicator. This is done for each VNI.
Head	Hosts create a copy of each BUM frame and send the copy to each tunnel endpoint it knows about for each VNI.

8 (Optional) Specify a VLAN ID or ranges of VLAN IDs for VLAN tagging.

To support guest VLAN tagging for VMs connected to this switch, you must specify VLAN ID ranges, also called trunk VLAN ID ranges. The logical port will filter packets based on the trunk VLAN ID ranges, and a guest VM can tag its packets with its own VLAN ID based on the trunk VLAN ID ranges.

9 (Optional) Click the **Switching Profiles** tab and select switching profiles.

10 Click **Save**.

In the NSX Manager UI, the new logical switch is a clickable link.

What to do next

Attach VMs to your logical switch. See [Connecting a VM to a Logical Switch in Manager Mode](#).

Connecting a VM to a Logical Switch in Manager Mode


Depending on your host, the configuration for connecting a VM to a logical switch can vary.

The supported hosts that can connect to a logical switch are; an ESXi host that is managed in vCenter Server, a standalone ESXi host, and a KVM host.

Attach a VM Hosted on vCenter Server to a Logical Switch in Manager Mode

If you have a ESXi host that is managed in vCenter Server, you can access the host VMs through the Web-based vSphere Web Client. In this case, you can use this procedure to attach VMs to NSX-T Data Center logical switches.

The example shown in this procedure shows how to attach a VM called app-vm to a logical switch called app-switch.



The installation-based vSphere Client application does not support attaching a VM to an NSX-T Data Center logical switch. If you do not have the (Web-based) vSphere Web Client, see [Attach a VM Hosted on Standalone ESXi to a Logical Switch in Manager Mode](#).


Prerequisites

- The VMs must be hosted on hypervisors that have been added to the NSX-T Data Center fabric.
- The fabric nodes must have NSX-T Data Center management plane (MPA) and NSX-T Data Center control plane (LCP) connectivity.
- The fabric nodes must be added to a transport zone.
- A logical switch must be created.

Procedure

- 1 In the vSphere Web Client, edit the VM settings, and attach the VM to the NSX-T Data Center logical switch.

For example:



T1-web-sv-01a - Edit Settings				
Virtual Hardware		VM Options	SDRS Rules	vApp Options
CPU	1			
Memory	512	MB		
Hard disk 1	750	MB		
SCSI controller 0	LSI Logic Parallel			
Network adapter 1	LS.ONE@0 (nsx.LogicalSwitch)		☒ Connect...	
CD/DVD drive 1	Client Device		☐ Connect...	
Floppy drive 1	Client Device		☐ Connect...	
Video card	Specify custom settings			
VMCI device				

2 Click **OK**.

Results

After attaching a VM to a logical switch, logical switch ports are added to the logical switch. You can view logical switch ports and the VIF attachment ID on the NSX Manager UI. In **Manager mode**, select **Networking > Logical Switches > Ports**.

Use the GET `https://<mgr-ip>/api/v1/logical-ports/` API call to view port details and Admin status for the corresponding VIF attachment ID. To view the Operational status, use the `https://<mgr-ip>/api/v1/logical-ports/<logical-port-id>/status` API call with the appropriate logical port ID.

If two VMs are attached to the same logical switch and have IP addresses configured in the same subnet, they should be able to ping each other.

What to do next

Add a logical router.

You can monitor the activity on the logical switch port to troubleshoot problems. See "Monitor a Logical Switch Port Activity" in the *NSX-T Data Center Administration Guide*.


Attach a VM Hosted on Standalone ESXi to a Logical Switch in Manager Mode

If you have a standalone ESXi host, you cannot access the host VMs through the web-based vSphere Web Client. In this case, you can use this procedure to attach VMs to NSX-T Data Center logical switches.

The example shown in this procedure shows how to attach a VM called app-vm to a logical switch called app-switch.

Switch's opaque network ID:
22b22448-38bc-419b-bea8-b51126bec7ad

VM's external ID:
50066bae-0f8a-386b-e62e-b0b9c6013a51




Prerequisites

- The VM must be hosted on hypervisors that have been added to the NSX-T Data Center fabric.
- The fabric nodes must have NSX-T Data Center management plane (MPA) and NSX-T Data Center control plane (LCP) connectivity.
- The fabric nodes must be added to a transport zone.
- A logical switch must be created.
- You must have access to the NSX Manager API.
- You must have write access to the VM's VMX file.

Procedure

- 1 Using the (install-based) vSphere Client application or some other VM management tool, edit the VM and add a VMXNET 3 Ethernet adapter.

Select any named network. You will change the network connection in a later step.



- 2 Use the NSX-T Data Center API to issue the GET `https://<nsx-mgr>/api/v1/fabric/virtual-machines/<VM-ID>` API call.

In the results, find the VM's externalId.

For example:

```
GET https://<nsx-mgr>/api/v1/fabric/virtual-machines/60a5a5d5-ea2b-407e-a806-4fdc8468f735

{
  "resource_type": "VirtualMachine",
  "id": "60a5a5d5-ea2b-407e-a806-4fdc8468f735",
  "display_name": "app-vm",
  "compute_ids": [
    "instanceUuid:50066bae-0f8a-386b-e62e-b0b9c6013a51",
    "moIdOnHost:5",
    "externalId:50066bae-0f8a-386b-e62e-b0b9c6013a51",
    "hostLocalId:5",
    "locationId:564dc020-1565-e3f4-f591-ee3953eef3ff",
    "biosUuid:4206f47d-fef7-08c5-5bf7-ea26a4c6b18d"
  ],
  "external_id": "50066bae-0f8a-386b-e62e-b0b9c6013a51",
  "type": "REGULAR",
  "host_id": "cb82b0fa-a8f1-11e5-92a9-6b7d1f8661fa",
```



```
"local_id_on_host": "5"
}
```

3 Power off and unregister the VM from the host.

You can use your VM management tool or the ESXi CLI, as shown here.

```
[user@host:~] vim-cmd /vmsvc/getallvms
Vmid    Name      File           Guest OS      Version  Annotation
5       app-vm    [ds2] app-vm/app-vm.vmx  ubuntuGuest   vmx-08
8       web-vm    [ds2] web-vm/web-vm.vmx  ubuntu64Guest vmx-08

[user@host:~] vim-cmd /vmsvc/power.off 5
Powering off VM:

[user@host:~] vim-cmd /vmsvc/unregister 5
```

4 From the NSX Manager UI, get the logical switch ID.

For example:

app-switch

[Overview](#)
[Monitor](#)
[Manage](#)
[Related](#)

[Summary](#)
[EDIT](#)

Name	app-switch
ID	9b2c8ead-f7b4-496c-bc22-6870bb44dd80
Location	
Description	
Admin Status	● Up
Replication Mode	Hierarchical Two-Tier replication
VLAN	N/A
VNI	65549
Logical Ports	0
Traffic Type	Overlay
Transport Zone	TZ_OVERLAY
Uplink Teaming Policy Name	[Use Default]
N-VDS Mode	STANDARD
Created	8/31/2018, 3:43:01 PM by admin
Last Updated	8/31/2018, 3:43:01 PM by admin

5 Modify the VM's VMX file.

Delete the **ethernet1.networkName = "<name>"** field and add the following fields:

- ethernet1.opaqueNetwork.id = "<logical switch's ID>"
- ethernet1.opaqueNetwork.type = "nsx.LogicalSwitch"
- ethernet1.externalId = "<VM's externalId>"
- ethernet1.connected = "TRUE"
- ethernet1.startConnected = "TRUE"

For example:

OLD

```
ethernet1.pciSlotNumber = "224"
ethernet1.virtualDev = "vmxnet3"
ethernet1.networkName = "VM Network"
ethernet1.addressType = "vpx"
```

```

ethernet1.generatedAddress = "00:50:56:86:7b:d7"
ethernet1.uptCompatibility = "true"
ethernet1.present = "TRUE"

```

NEW

```

ethernet1.pciSlotNumber = "224"
ethernet1.virtualDev = "vmxnet3"
ethernet1.addressType = "vpx"
ethernet1.generatedAddress = "00:50:56:86:7b:d7"
ethernet1.uptCompatibility = "true"
ethernet1.present = "TRUE"
ethernet1.opaqueNetwork.id = "22b22448-38bc-419b-bea8-b51126bec7ad"
ethernet1.opaqueNetwork.type = "nsx.LogicalSwitch"
ethernet1.externalId = "50066bae-0f8a-386b-e62e-b0b9c6013a51"
ethernet1.connected = "TRUE"
ethernet1.startConnected = "TRUE"

```

- 6 In the NSX Manager UI, add a logical switch port, and use the VM's externalId for the VIF attachment.
- 7 Reregister the VM and power it on.

You can use your VM management tool or the ESXi CLI, as shown here.

```

[user@host:~] vim-cmd /solo/register /path/to/file.vmx

For example:
[user@host:~] vim-cmd solo/registervm /vmfs/volumes/355f2049-6c704347/app-vm/app-vm.vmx
9

[user@host:~] vim-cmd /vmsvc/power.on 9
Powering on VM:

```

Results

In the NSX Manager UI in **Manager** mode, select **Networking > Logical Switches > Ports**. Find the VIF attachment ID matching the VM's externalId and make sure that the Admin and Operational status are Up/Up.

If two VMs are attached to the same logical switch and have IP addresses configured in the same subnet, they should be able to ping each other.

What to do next


Add a logical router.

You can monitor the activity on the logical switch port to troubleshoot problems. See "Monitor a Logical Switch Port Activity" in the *NSX-T Data Center Administration Guide*.

Attach a VM Hosted on KVM to a Logical Switch in Manager Mode

If you have a KVM host, you can use this procedure to attach VMs to NSX-T Data Center logical switches.

The example shown in this procedure shows how to attach a VM called app-vm to a logical switch called app-switch.



Prerequisites

- The VM must be hosted on hypervisors that have been added to the NSX-T Data Center fabric.
- The fabric nodes must have NSX-T Data Center management plane (MPA) and NSX-T Data Center control plane (LCP) connectivity.
- The fabric nodes must be added to a transport zone.
- A logical switch must be created.

Procedure

- 1 From the KVM CLI, run the `virsh dumpxml <your vm> | grep interfaceid` command.
- 2 In the NSX Manager UI, add a logical switch port, and use the VM's interface ID for the VIF attachment.

Results

In the NSX Manager UI in **Manager** mode, select **Networking > Logical Switches > Ports**. Find the VIF attachment ID and make sure that the Admin and Operational status are Up/Up.

If two VMs are attached to the same logical switch and have IP addresses configured in the same subnet, they should be able to ping each other.

What to do next

Add a logical router.

You can monitor the activity on the logical switch port to troubleshoot problems. See "Monitor a Logical Switch Port Activity" in the *NSX-T Data Center Administration Guide*.

Create a Logical Switch Port In Manager Mode

A logical switch has multiple switch ports. A logical switch port connects another network component, a VM, or a container to a logical switch.

If you connect a VM to a logical switch on an ESXi host that is managed by vCenter Server, a logical switch port is created automatically. For more information about connecting a VM to a logical switch, see [Connecting a VM to a Logical Switch in Manager Mode](#).

For more information about connecting a container to a logical switch, see the *NSX-T Container Plug-in for Kubernetes - Installation and Administration Guide*.

Note The IP address and MAC address bound to a logical switch port for a container are allocated by NSX Manager. Do not change the address binding manually.

To monitor activity on a logical switch port, see [Monitor a Logical Switch Port Activity in Manager Mode](#).

Prerequisites

- Verify that a logical switch is created. See [Logical Switches in Manager Mode](#).
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Logical Switches > Ports > Add**.
- 3 In the **General** tab, complete the port details.

Option	Description
Name and Description	Enter a name and optionally a description.
Logical Switch	Select a logical switch from the drop-down menu.
Admin Status	Select Up or Down .
Attachment Type	Select None or VIF . Select VIF if this port is for connecting to a VM.
Attachment ID	If the attachment type is VIF , enter the attachment ID.

Using the API, you can set the attachment type to additional values (LOGICALROUTER, BRIDGEENDPOINT, DHCP_SERVICE, METADATA_PROXY, L2VPN_SESSION). If the attachment type is DHCP service, metadata proxy, or L2 VPN session, the switching profiles for the port must be the default ones. You cannot use any user-defined profile.


- 4 (Optional) In the **Switching Profiles** tab, select switching profiles.
- 5 Click **Save**.

Test Layer 2 Connectivity in Manager Mode

After you successfully set up your logical switch and attach VMs to the logical switch, you can test the network connectivity of the attached VMs.

If your network environment is configured properly, based on the topology the App2 VM can ping the App1 VM.

Figure 20-2. Logical Switch Topology



Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 Log in to one of the VMs attached to the logical switch using SSH or the VM console.
For example, App2 VM 172.16.20.11.
- 2 Ping the second VM attached to the logical switch to test connectivity.

```
$ ping -c 2 172.16.20.10
PING 172.16.20.10 (172.16.20.10) 56(84) bytes of data.
64 bytes from 172.16.20.10: icmp_seq=1 ttl=63 time=0.982 ms
64 bytes from 172.16.20.10: icmp_seq=2 ttl=63 time=0.654 ms
64 bytes from 172.16.20.10: icmp_seq=3 ttl=63 time=0.791 ms

--- 172.16.20.10 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1990ms
rtt min/avg/max/mdev = 0.654/0.809/0.902/0.104 ms
```

- 3 (Optional) Identify the problem that causes the ping to fail.
 - a Verify that the VM network settings are correct.
 - b Verify that the VM network adapter is connected to the correct logical switch.
 - c Verify that the logical switch Admin status is UP.
 - d From the NSX Manager, select **Networking > Logical Switches > Switches**.

- e Click the logical switch and note the UUID and VNI information.
- f Run the following commands to troubleshoot the problem.

Command	Description
get logical-switch <vni-or-uuid> arp-table	Displays the ARP table for the specified logical switch. Sample output. <pre>nsx-manager1> get logical-switch 41866 arp-table VNI IP MAC Connection-ID 41866 172.16.20.11 00:50:56:b1:70:5e 295422</pre>
get logical-switch <vni-or-uuid> connection-table	Displays the connections for the specified logical switch. Sample output. <pre>nsx-manager1> get logical-switch 41866 connection-table Host-IP Port ID 192.168.110.37 36923 295420 192.168.210.53 37883 295421 192.168.210.54 57278 295422</pre>
get logical-switch <vni-or-uuid> mac-table	Displays the MAC table for the specified logical switch. Sample output. <pre>nsx-manager1> get logical-switch 41866 mac-table VNI MAC VTEP-IP Connection-ID 41866 00:50:56:86:f2:b2 192.168.250.102 295421 41866 00:50:56:b1:70:5e 192.168.250.101 295422</pre>
get logical-switch <vni-or-uuid> stats	Displays statistics information about the specified logical switch. Sample output. <pre>nsx-manager1> get logical-switch 41866 stats update.member 11 update.vtep 11 update.mac 4 update.mac.invalidate 0 update.arp 7 update.arp.duplicate 0 query.mac 2 query.mac.miss 0 query.arp 9 query.arp.miss 6</pre>
get logical-switch <vni-or-uuid> stats-sample	Displays a summary of all logical switch statistics over time. Sample output. <pre>nsx-manager1> get logical-switch 41866 stats-sample 21:00:00 21:10:00 21:20:00 21:30:00 21:40:00 update.member 0 0 0 0 0 update.vtep 0 0 0 0 0 update.mac 0 0 0 0 0 update.mac.invalidate 0 0 0 0 0 update.arp 0 0 0 0 0 update.arp.duplicate 0 0 0 0 0</pre>

Command	Description
	<pre>query.mac 0 0 0 0 0 query.mac.miss 0 0 0 0 0 query.arp 0 0 0 0 0 query.arp.miss 0 0 0 0 0</pre>
get logical-switch <vni-or-uuid> vtep	Displays all virtual tunnel end points related to the specified logical switch. Sample output. <pre>nsx-manager1> get logical-switch 41866 vtep VNI IP LABEL Segment MAC Connection-ID 41866 192.168.250.102 0x8801 192.168.250.0 00:50:56:65:f5:fc 295421 41866 192.168.250.100 0x1F801 192.168.250.0 02:50:56:00:00:00 295420 41866 192.168.250.101 0x16001 192.168.250.0 00:50:56:64:7c:28 295422</pre>


Results

The first VM attached to the logical switch is able to send packets to the second VM.

Create a VLAN Logical Switch for the NSX Edge Uplink in Manager Mode

Edge uplinks go out through VLAN logical switches.

When you are creating a VLAN logical switch, it is important to have in mind a particular topology that you are building. For example, the following simple topology shows a single VLAN logical switch inside of a VLAN transport zone. The VLAN logical switch has VLAN ID 100. This matches the VLAN ID on the TOR port connected to the hypervisor host port used for the Edge's VLAN uplink.



Prerequisites

- To create a VLAN logical switch, you must first create a VLAN transport zone.
- An NSX-T Data Center vSwitch must be added to the NSX Edge. To confirm on an Edge, run the `get host-switches` command. For example:

```
nsx-edge1> get host-switches

Host Switch      : c0a78378-1c20-432a-9e23-ddb34f1c80c9
Switch Name      : hs1
Transport Zone   : c46dcd72-808a-423d-b4cc-8752c33f6b2c
Transport Zone   : 73def985-d122-4b7b-ab6a-a58176dfc32d
Physical Port    : fp-eth0
Uplink Name      : uplink-1
Transport VLAN   : 4096
Default Gateway  : 192.168.150.1
Subnet Mask      : 255.255.255.0
Local VTEP Device : fp-eth0
Local VTEP IP    : 192.168.150.102
```

- Verify that fabric nodes are successfully connected to the NSX-T Data Center management plane agent (MPA) and the NSX-T Data Center local control plane (LCP).

In the GET `https://<nsx-mgr>/api/v1/transport-nodes/<transport-node-id>/state` API call, the state must be `success`. See the *NSX-T Data Center Installation Guide*.

- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From a browser, log in to an NSX Manager at `https://<nsx-mgr>`.
- 2 Select **Networking > Logical Switches > Switches > Add**.
- 3 Type a name for the logical switch.
- 4 Select a transport zone for the logical switch.
- 5 Select an uplink teaming policy.
- 6 For admin status, select **Up** or **Down**.
- 7 Type a VLAN ID.

Enter 0 in the VLAN field if there is no VLAN ID for the uplink to the physical TOR.

- 8 (Optional) Click the **Switching Profiles** tab and select switching profiles.

Results

Note If you have two VLAN logical switches that have the same VLAN ID, they cannot be connected to the same Edge N-VDS (previously known as hostswitch). If you have a VLAN logical switch and an overlay logical switch, and the VLAN ID of the VLAN logical switch is the same as the transport VLAN ID of the overlay logical switch, they also cannot be connected to the same Edge N-VDS.

What to do next

Add a logical router.

Switching Profiles for Logical Switches and Logical Ports

Switching profiles include Layer 2 networking configuration details for logical switches and logical ports. NSX Manager supports several types of switching profiles, and maintains one or more system-defined default switching profiles for each profile type.

The following types of switching profiles are available.

- QoS (Quality of Service)
- Port Mirroring
- IP Discovery
- SpoofGuard
- Switch Security

■ MAC Management

Note You cannot edit or delete the default switching profiles in the NSX Manager. You can create custom switching profiles instead.

Before using a default profile, make sure that the settings are what you need them to be. When you create a custom profile, some settings have default values. Do not assume that in the default profile, these settings will have the default values.

Each default or custom switching profile has a unique reserved identifier. You use this identifier to associate the switching profile to a logical switch or a logical port. For example, the default QoS switching profile ID is f313290b-eba8-4262-bd93-fab5026e9495.

A logical switch or logical port can be associated with one switching profile of each type. You cannot have for example, two QoS different switching profiles associated to a logical switch or logical port.

If you do not associate a switching profile type while creating or updating a logical switch, then the NSX Manager associates a corresponding default system-defined switching profile. The children logical ports inherit the default system-defined switching profile from the parent logical switch.

When you create or update a logical switch or logical port you can choose to associate either a default or a custom switching profile. When the switching profile is associated or disassociated from a logical switch the switching profile for the children logical ports is applied based on the following criteria.

- If the parent logical switch has a profile associated with it, the child logical port inherits the switching profile from the parent.
- If the parent logical switch does not have a switching profile associated with it, a default switching profile is assigned to the logical switch and the logical port inherits that default switching profile.
- If you explicitly associate a custom profile with a logical port, then this custom profile overrides the existing switching profile.

Note If you have associated a custom switching profile with a logical switch, but want to retain the default switching profile for one of the child logical port, then you must make a copy of the default switching profile and associate it with the specific logical port.

You cannot delete a custom switching profile if it is associated to a logical switch or a logical port. You can find out whether any logical switches and logical ports are associated with the custom switching profile by going to the Assigned To section of the Summary view and clicking on the listed logical switches and logical ports.

Understanding QoS Switching Profile

QoS provides high-quality and dedicated network performance for preferred traffic that requires high bandwidth. The QoS mechanism does this by prioritizing sufficient bandwidth, controlling

latency and jitter, and reducing data loss for preferred packets even when there is a network congestion. This level of network service is provided by using the existing network resources efficiently.

For this release, shaping and traffic marking namely, CoS and DSCP is supported. The Layer 2 Class of Service (CoS) allows you to specify priority for data packets when traffic is buffered in the logical switch due to congestion. The Layer 3 Differentiated Services Code Point (DSCP) detects packets based on their DSCP values. CoS is always applied to the data packet irrespective of the trusted mode.

NSX-T Data Center trusts the DSCP setting applied by a virtual machine or modifying and setting the DSCP value at the logical switch level. In each case, the DSCP value is propagated to the outer IP header of encapsulated frames. This enables the external physical network to prioritize the traffic based on the DSCP setting on the external header. When DSCP is in the trusted mode, the DSCP value is copied from the inner header. When in the untrusted mode, the DSCP value is not preserved for the inner header.

Note DSCP settings work only on tunneled traffic. These settings do not apply to traffic inside the same hypervisor.

You can use the QoS switching profile to configure the average ingress and egress bandwidth values to set the transmit limit rate. The peak bandwidth rate is used to support burst traffic a logical switch is allowed to prevent congestion on the northbound network links. These settings do not guarantee the bandwidth but help limit the use of network bandwidth. The actual bandwidth you will observe is determined by the link speed of the port or the values in the switching profile, whichever is lower.

The QoS switching profile settings are applied to the logical switch and inherited by the child logical switch port.

Configure a Custom QoS Switching Profile in Manager Mode

You can define the DSCP value and configure the ingress and egress settings to create a custom QoS switching profile.

Prerequisites

- Familiarize yourself with the QoS switching profile concept. See [Understanding QoS Switching Profile](#).
- Identify the network traffic you want to prioritize.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

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

2 Select **Networking > Logical Switches > Switching Profiles > Add**

3 Select **QoS** and complete the QoS switching profile details.

Option	Description
Name and Description	Assign a name to the custom QoS switching profile. You can optionally describe the setting that you modified in the profile.
Mode	Select either a Trusted or Untrusted option from the Mode drop-down menu. When you select the Trusted mode the inner header DSCP value is applied to the outer IP header for IP/IPv6 traffic. For non IP/IPv6 traffic, the outer IP header takes the default value. Trusted mode is supported on an overlay-based logical port. The default value is 0. Untrusted mode is supported on overlay-based and VLAN-based logical port. For the overlay-based logical port, the DSCP value of the outbound IP header is set to the configured value irrespective to the inner packet type for the logical port. For the VLAN-based logical port, the DSCP value of IP/IPv6 packet will be set to the configured value. The DSCP values range for untrusted mode is between 0 to 63. Note DSCP settings work only on tunneled traffic. These settings do not apply to traffic inside the same hypervisor.
Priority	Set the DSCP value. The priority values range from 0 to 63.
Class of Service	Set the CoS value. CoS is supported on VLAN-based logical port. CoS groups similar types of traffic in the network and each type of traffic is treated as a class with its own level of service priority. The lower priority traffic is slowed down or in some cases dropped to provide better throughput for higher priority traffic. CoS can also be configured for the VLAN ID with zero packet. The CoS values range from 0 to 7, where 0 is the best effort service.
Ingress	Set custom values for the outbound network traffic from the VM to the logical network. You can use the average bandwidth to reduce network congestion. The peak bandwidth rate is used to support burst traffic and the burst size is based on the duration with peak bandwidth. You set burst duration in the burst size setting. You cannot guarantee the bandwidth. However, you can use the Average, Peak, and Burst Size settings to limit network bandwidth. For example, if the average bandwidth is 30 Mbps, peak bandwidth is 60 Mbps, and the allowed duration is 0.1 second, then the burst size is $60 * 1000000 * 0.10 / 8 = 750000$ Bytes. The default value 0 disables rate limiting on the ingress traffic.

Option	Description
Ingress Broadcast	Set custom values for the outbound network traffic from the VM to the logical network based on broadcast. Set custom values for the outbound network traffic from the VM to the logical network based on broadcast. For example, when you set the average bandwidth for a logical switch to 3000 Kbps, peak bandwidth is 6000 Kbps, and the allowed duration is 0.1 second, then the burst size is $6000 * 1000 * 0.10 / 8 = 75000$ Bytes. The default value 0 disables rate limiting on the ingress broadcast traffic.
Egress	Set custom values for the inbound network traffic from the logical network to the VM. The default value 0 disables rate limiting on the egress traffic.

If the ingress, ingress broadcast, and egress options are not configured, the default values are used.

4 Click **Save**.

Results

A custom QoS switching profile appears as a link.

What to do next

Attach this QoS customized switching profile to a logical switch or logical port so that the modified parameters in the switching profile are applied to the network traffic. See [Associate a Custom Profile with a Logical Switch in Manager Mode](#) or [Associate a Custom Profile with a Logical Port in Manager Mode](#).

Understanding Port Mirroring Switching Profile

Logical port mirroring lets you replicate and redirect all of the traffic coming in or out of a logical switch port attached to a VM VIF port. 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.

We recommend you use port mirroring only for troubleshooting.

Note Port Mirroring is not recommended for monitoring because when used for longer durations performance is impacted.

Compared to the physical port mirroring, logical port mirroring ensures that all of the VM network traffic is captured. If you implement port mirroring only in the physical network, some of the VM network traffic fails to be mirrored. This happens because communication between VMs residing on the same host never enters the physical network and therefore does not get mirrored. With logical port mirroring you can continue to mirror VM traffic even when that VM is migrated to another host.

The port mirroring process is similar for both VM ports in the NSX-T Data Center domain and ports of physical applications. You can forward the traffic captured by a workload connected to a logical network and mirror that traffic to a collector. The IP address should be reachable from the guest IP address on which the VM is hosted. This process is also true for physical applications connected to gateway nodes.

Configure a Custom Port Mirroring Switching Profile in Manager Mode

You can create a custom port mirroring switching profile with a different destination and key value.

Prerequisites

- Familiarize yourself with the port mirroring switching profile concept. See [Understanding Port Mirroring Switching Profile](#).
- Identify the IP address of the destination logical port ID you want to redirect network traffic to.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Logical Switches > Switching Profiles > Add**
- 3 Select **Port Mirroring** and complete the port mirroring switching profile details.

Option	Description
Name and Description	Assign a name to the custom port mirroring switching profile. You can optionally describe the setting you modified to customize this profile.
Direction	Select an option from the drop-down menu to use this source for Ingress , Egress , or Bidirectional 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. Bidirectional is the two-way of traffic from the VM to the logical network and from the logical network to the VM. This is the default option.
Packet Truncation	Optional. The range is 60 - 65535.

Option	Description
Key	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.
Destinations	Enter the destination ID of the collector for the mirroring session. The destination IP address ID can only be an IPv4 address within the network or a remote IPv4 address not managed by NSX-T Data Center. You can add up to three destination IP addresses separated by a comma.

4 Click **Save**.

Results

A custom port mirroring switching profile appears as a link.

What to do next

Attach the switching profile to a logical switch or logical port. See [Associate a Custom Profile with a Logical Switch in Manager Mode](#) or [Associate a Custom Profile with a Logical Port in Manager Mode](#).

Verify that the customized port mirroring switching profile works. See [Verify Custom Port Mirroring Switching Profile](#).

Verify Custom Port Mirroring Switching Profile

Before you start using the custom port mirroring switching profile, verify that the customization works properly.

Prerequisites

- Verify that the custom port mirroring switching profile is configured. See [Configure a Custom Port Mirroring Switching Profile in Manager Mode](#).
- Verify that the customized port mirroring switching profile is attached to a logical switch. See [Associate a Custom Profile with a Logical Switch in Manager Mode](#).

Procedure

- 1 Locate two VMs with VIF attachments to the logical port configured for port mirroring.
For example, VM1 10.70.1.1 and VM2 10.70.1.2 have VIF attachments and they are located in the same logical network.
- 2 Run the `tcpdump` command on a destination IP address.

```
sudo tcpdump -n -i eth0 dst host destination_IP_address and proto gre
```


For example, the destination IP address is 10.24.123.196.
- 3 Log in to the first VM and ping the second VM to verify that the corresponding ECHO requests and replies are received at the destination address.

What to do next

Attach this port mirroring customized switching profile to a logical switch so that the modified parameters in the switching profile are applied to the network traffic. See [Associate a Custom Profile with a Logical Switch in Manager Mode](#).

Understanding IP Discovery Switching Profile

IP Discovery uses DHCP and DHCPv6 snooping, ARP (Address Resolution Protocol) snooping, ND (Neighbor Discovery) snooping, and VM Tools to learn MAC and IP addresses.

The discovered MAC and IP addresses are used to achieve ARP/ND suppression, which minimizes traffic between VMs connected to the same logical switch. The addresses are also used by the SpoofGuard and distributed firewall (DFW) components. DFW uses the address bindings to determine the IP address of objects in firewall rules.

DHCP/DHCPv6 snooping inspects the DHCP/DHCPv6 packets exchanged between the DHCP/DHCPv6 client and server to learn the IP and MAC addresses.

ARP snooping inspects the outgoing ARP and GARP (gratuitous ARP) packets of a VM to learn the IP and MAC addresses.

VM Tools is software that runs on an ESXi-hosted VM and can provide the VM's configuration information including MAC and IP or IPv6 addresses. This IP discovery method is available for VMs running on ESXi hosts only.

ND snooping is the IPv6 equivalent of ARP snooping. It inspects neighbor solicitation (NS) and neighbor advertisement (NA) messages to learn the IP and MAC addresses.

Duplicate address detection checks whether a newly discovered IP address is already present on the realized binding list for a different port. This check is performed for ports on the same segment. If a duplicate address is detected, the newly discovered address is added to the discovered list, but is not added to the realized binding list. All duplicate IPs have an associated discovery timestamp. If the IP that is on the realized binding list is removed, either by adding it to the ignore binding list or by disabling snooping, the duplicate IP with the oldest timestamp is moved to the realized binding list. The duplicate address information is available through an API call.

By default, the discovery methods ARP snooping and ND snooping operate in a mode called trust on first use (TOFU). In TOFU mode, when an address is discovered and added to the realized bindings list, that binding remains in the realized list forever. TOFU applies to the first 'n' unique <IP, MAC, VLAN> bindings discovered using ARP/ND snooping, where 'n' is the binding limit that you can configure. You can disable TOFU for ARP/ND snooping. The methods will then operate in trust on every use (TOEU) mode. In TOEU mode, when an address is discovered, it is added to the realized bindings list and when it is deleted or expired, it is removed from the realized bindings list. DHCP snooping and VM Tools always operate in TOEU mode.

For each port, NSX Manager maintains an ignore bindings list, which contains IP addresses that cannot be bound to the port. If you navigate to **Networking > Logical Switches > Ports** in **Manager** mode, and select a port, you can add discovered bindings to the ignore bindings list. You can also delete an existing discovered or realized binding by copying it to **Ignore Bindings**.

Note TOFU is not the same as SpoofGuard, and it does not block traffic in the same way as SpoofGuard. For more information, see [Understanding SpoofGuard Segment Profile](#).

For Linux VMs, the ARP flux problem might cause ARP snooping to obtain incorrect information. The problem can be prevented with an ARP filter. For more information, see <http://linux-ip.net/html/ether-arp.html#ether-arp-flux>.

Configure IP Discovery Switching Profile in Manager Mode

NSX-T Data Center has several default IP Discovery switching profiles. You can also create additional ones.

Prerequisites

- Familiarize yourself with the IP Discovery switching profile concepts. See [Understanding IP Discovery Switching Profile](#).
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Logical Switches > Switching Profiles > Add**.
- 3 Select **IP Discovering** and specify the IP Discovery switching profile details.

Option	Description
Name and Description	Enter a name and optionally a description.
ARP Snooping	For an IPv4 environment. Applicable if VMs have static IP addresses.
ARP Binding Limit	The maximum number of IPv4 IP addresses that can be bound to a port. The minimum value allowed is 1 (the default) and the maximum is 256.

Option	Description
ARP ND Binding Limit Timeout	The timeout value, in minutes, for IP addresses in the ARP/ND binding table if TOFU is disabled. If an address times out, a newly discovered address replaces it.
DHCP Snooping	For an IPv4 environment. Applicable if VMs have IPv4 addresses.
DHCP V6 Snooping	For an IPv6 environment. Applicable if VMs have IPv6 addresses.
VM Tools	Available for ESXi-hosted VMs only.
VM Tools for IPv6	Available for ESXi-hosted VMs only.
Neighbor Discovery Snooping	For an IPv6 environment. Applicable if VMs have static IP addresses.
Neighbor Discovery Binding Limit	The maximum number of IPv6 addresses that can be bound to a port.
Trust on First Use	Applicable to ARP and ND snooping.
Duplicate IP Detection	For all snooping methods and both IPv4 and IPv6 environments.

4 Click **Add**.

What to do next

Attach this IP Discovery customized switching profile to a logical switch or logical port so that the modified parameters in the switching profile are applied to the network traffic. See [Associate a Custom Profile with a Logical Switch in Manager Mode](#) or [Associate a Custom Profile with a Logical Port in Manager Mode](#).

Understanding SpoofGuard

SpoofGuard helps prevent a form of malicious attack called "web spoofing" or "phishing." A SpoofGuard policy blocks traffic determined to be spoofed.

SpoofGuard is a tool that is designed to prevent virtual machines in your environment from altering their existing IP address. In the instance that a virtual machine's IP address does not match the IP address on the corresponding logical port and switch address binding in SpoofGuard, the virtual machine's vNIC is prevented from accessing the network entirely. SpoofGuard can be configured at the port or switch level. There are several reasons SpoofGuard might be used in your environment:

- Preventing a rogue virtual machine from assuming the IP address of an existing VM.
- Ensuring the IP addresses of virtual machines cannot be altered without intervention – in some environments, it's preferable that virtual machines cannot alter their IP addresses without proper change control review. SpoofGuard facilitates this by ensuring that the virtual machine owner cannot simply alter the IP address and continue working unimpeded.
- Guaranteeing that distributed firewall (DFW) rules will not be inadvertently (or deliberately) bypassed – for DFW rules created utilizing IP sets as sources or destinations, the possibility always exists that a virtual machine could have its IP address forged in the packet header, thereby bypassing the rules in question.

NSX-T Data Center SpoofGuard configuration covers the following:

- MAC SpoofGuard - authenticates MAC address of packet
- IP SpoofGuard - authenticates MAC and IP addresses of packet
- Dynamic Address Resolution Protocol (ARP) inspection, that is, ARP and Gratuitous Address Resolution Protocol (GARP) SpoofGuard and Neighbor Discovery (ND) SpoofGuard validation are all against the MAC source, IP Source and IP-MAC source mapping in the ARP/GARP/ND payload.

At the port level, the allowed MAC/VLAN/IP allow-list is provided through the Address Bindings property of the port. When the virtual machine sends traffic, it is dropped if its IP/MAC/VLAN does not match the IP/MAC/VLAN properties of the port. The port level SpoofGuard deals with traffic authentication, i.e. is the traffic consistent with VIF configuration.

At the switch level, the allowed MAC/VLAN/IP allow-list is provided through the Address Bindings property of the switch. This is typically an allowed IP range/subnet for the switch and the switch level SpoofGuard deals with traffic authorization.

Traffic must be permitted by port level AND switch level SpoofGuard before it will be allowed into switch. Activating or deactivating port and switch level SpoofGuard, can be controlled using the SpoofGuard switch profile.

Configure Port Address Bindings in Manager Mode

Address bindings specify the IP and MAC address of a logical port and are used to specify the port whitelist in SpoofGuard.

With port address bindings you'll specify the IP and MAC address, and VLAN if applicable, of the logical port. When SpoofGuard is enabled, it ensures that the specified address bindings are enforced in the data path. In addition to SpoofGuard, port address bindings are used for DFW rule translations.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

1 In NSX Manager, select to **Networking > Logical Switches > Ports**.

2 Click the logical port to which you want apply address binding.

The logical port summary appears.

3 In the **Overview** tab, expand **Address Bindings > Manual Bindings**.

4 Click **Add**.

The Add Address Binding dialogue box appears.

- 5 Specify the IP (IPv4 address, IPv6 address, or IPv6 subnet) and MAC address of the logical port to which you want to apply address binding. For example, for IPv6, 2001::/64 is an IPv6 subnet, 2001::1 is a host IP, whereas 2001::1/64 is an invalid input. You can also specify a VLAN ID.
- 6 Click **Add**.

What to do next

Use the port address bindings when you [Configure a SpoofGuard Switching Profile in Manager Mode](#).

Configure a SpoofGuard Switching Profile in Manager Mode

When SpoofGuard is configured, if the IP address of a virtual machine changes, traffic from the virtual machine may be blocked until the corresponding configured port/switch address bindings are updated with the new IP address.

Enable SpoofGuard for the port group(s) containing the guests. When enabled for each network adapter, SpoofGuard inspects packets for the prescribed MAC and its corresponding IP address.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Logical Switches > Switching Profiles > Add**.
- 3 Select **Spoof Guard**.
- 4 Enter a name and optionally a description.
- 5 To enable port level SpoofGuard, set **Port Bindings** to **Enabled**.
- 6 Click **Add**.

Results

A new switching profile has been created with a SpoofGuard Profile.

What to do next

Associate the SpoofGuard profile with a logical switch or logical port. See [Associate a Custom Profile with a Logical Switch in Manager Mode](#) or [Associate a Custom Profile with a Logical Port in Manager Mode](#).

Understanding Switch Security Switching Profile

Switch security provides stateless Layer2 and Layer 3 security by checking the ingress traffic to the logical switch and dropping unauthorized packets sent from VMs by matching the IP address, MAC address, and protocols to a set of allowed addresses and protocols. You can use switch security to protect the logical switch integrity by filtering out malicious attacks from the VMs in the network.

You can configure the Bridge Protocol Data Unit (BPDU) filter, DHCP Snooping, DHCP server block, and rate limiting options to customize the switch security switching profile on a logical switch.

Configure a Custom Switch Security Switching Profile in Manager Mode

You can create a custom switch security switching profile with MAC destination addresses from the allowed BPDU list and configure rate limiting.

Prerequisites

- Familiarize yourself with the switch security switching profile concept. See [Understanding Switch Security Switching Profile](#).
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Logical Switches**.
- 3 Click the **Switching Profiles** tab.
- 4 Click **Add** and select **Switch Security**.
- 5 Complete the switch security profile details.

Option	Description
Name and Description	Assign a name to the custom switch security profile. You can optionally describe the setting that you modified in the profile.
BPDU Filter	Toggle the BPDU Filter button to enable BPDU filtering. Disabled by default. When the BPDU filter is enabled, all of the traffic to BPDU destination MAC address is blocked. The BPDU filter when enabled also disables STP on the logical switch ports because these ports are not expected to take part in STP.
BPDU Filter Allow List	Click the destination MAC address from the BPDU destination MAC addresses list to allow traffic to the permitted destination. You must enable BPDU Filter to be able to select from this list.

Option	Description
DHCP Filter	Toggle the Server Block button and Client Block button to enable DHCP filtering. Both are disabled by default. DHCP Server Block blocks traffic from a DHCP server to a DHCP client. Note that it does not block traffic from a DHCP server to a DHCP relay agent. DHCP Client Block prevents a VM from acquiring a DHCP IP address by blocking DHCP requests.
DHCPv6 Filter	Toggle the V6 Server Block button and V6 Client Block button to enable DHCP filtering. Both are disabled by default. DHCPv6 Server Block blocks traffic from a DHCPv6 server to a DHCPv6 client. Note that it does not block traffic from a DHCP server to a DHCP relay agent. Packets whose UDP source port number is 547 are filtered. DHCPv6 Client Block prevents a VM from acquiring a DHCP IP address by blocking DHCP requests. Packets whose UDP source port number is 546 are filtered.
Block Non-IP Traffic	Toggle the Block Non-IP Traffic button to allow only IPv4, IPv6, ARP, and BPDU traffic. The rest of the non-IP traffic is blocked. The permitted IPv4, IPv6, ARP, GARP and BPDU traffic is based on other policies set in address binding and SpoofGuard configuration. By default, this option is disabled to allow non-IP traffic to be handled as regular traffic.
RA Guard	Toggle the RA Guard button to filter out ingress IPv6 router advertisements. ICMPv6 type 134 packets are filtered out. This option is enabled by default.
Rate Limits	Set a rate limit for broadcast and multicast traffic. This option is enabled by default. Rate limits can be used to protect the logical switch or VMs from events such as broadcast storms. To avoid any connectivity problems, the minimum rate limit value must be ≥ 10 pps.

6 Click **Add**.

Results

A custom switch security profile appears as a link.

What to do next

Attach this switch security customized switching profile to a logical switch or logical port so that the modified parameters in the switching profile are applied to the network traffic. See [Associate a Custom Profile with a Logical Switch in Manager Mode](#) or [Associate a Custom Profile with a Logical Port in Manager Mode](#).

Understanding MAC Management Switching Profile

The MAC management switching profile supports two functionalities: MAC learning and MAC address change.

The MAC address change feature allows a VM to change its MAC address. A VM connected to a port can run an administrative command to change the MAC address of its vNIC and still send and receive traffic on that vNIC. This feature is supported on ESXi only and not on KVM. In the default MAC management switching profile, this property is enabled by default.

MAC learning provides network connectivity to deployments where multiple MAC addresses are configured behind one vNIC, for example, in a nested hypervisor deployment where an ESXi VM runs on an ESXi host and multiple VMs run inside the ESXi VM. Without MAC learning, when the ESXi VM's vNIC connects to a switch port, its MAC address is static. VMs running inside the ESXi VM do not have network connectivity because their packets have different source MAC addresses. With MAC learning, the vSwitch inspects the source MAC address of every packet coming from the vNIC, learns the MAC address and allows the packet to go through. If a MAC address that is learned is not used for a certain period of time, it is removed. This aging property is not configurable.

MAC learning also supports unknown unicast flooding. Normally, when a packet that is received by a port has an unknown destination MAC address, the packet is dropped. With unknown unicast flooding enabled, the port floods unknown unicast traffic to every port on the switch that has MAC learning and unknown unicast flooding enabled. This property is enabled by default, but only if MAC learning is enabled.

The number of MAC addresses that can be learned is configurable. The maximum value is 4096, which is the default. You can also set the policy for when the limit is reached. The options are:

- **Drop** - Packets from an unknown source MAC address are dropped. Packets inbound to this MAC address will be treated as unknown unicast. The port will receive the packets only if it has unknown unicast flooding enabled.
- **Allow** - Packets from an unknown source MAC address are forwarded although the address will not be learned. Packets inbound to this MAC address will be treated as unknown unicast. The port will receive the packets only if it has unknown unicast flooding enabled.

If you enable MAC learning or MAC address change, to improve security, configure SpoofGuard as well.

Configure MAC Management Switching Profile in Manager Mode

You can create a MAC management switching profile to manage MAC addresses.

Prerequisites

- Familiarize yourself with the MAC management switching profile concept. See [Understanding MAC Management Switching Profile](#).
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Logical Switches > Switching Profiles > Add**.
- 3 Select **MAC Management** and complete the MAC management profile details.

Option	Description
Name and Description	Assign a name to the MAC management profile. You can optionally describe the setting that you modified in the profile.
MAC Change	Enable or disable the MAC address change feature. The default is disabled.
Status	Enable or disable the MAC learning feature. The default is disabled.
Unknown Unicast Flooding	Enable or disable the unknown unicast flooding feature. The default is enabled. This option is available if you enable MAC learning
MAC Limit	Set the maximum number of MAC addresses. The default is 4096. This option is available if you enable MAC learning
MAC Limit Policy	Select Allow or Drop . The default is Allow . This option is available if you enable MAC learning

- 4 Click **Add**.

What to do next

Attach the switching profile to a logical switch or logical port. See [Associate a Custom Profile with a Logical Switch in Manager Mode](#) or [Associate a Custom Profile with a Logical Port in Manager Mode](#).

Associate a Custom Profile with a Logical Switch in Manager Mode

You can associate a custom switching profile to a logical switch so that the profile applies to all the ports on the switch.

When custom switching profiles are attached to a logical switch they override existing default switching profiles. The custom switching profile is inherited by children logical switch ports.

Note If you have associated a custom switching profile with a logical switch, but want to retain the default switching profile for one of the child logical switch port, then you must make a copy of the default switching profile and associate it with the specific logical switch port.

Prerequisites

- Verify that a logical switch is configured. See [Create a Logical Switch in Manager Mode](#).
- Verify that a custom switching profile is configured. See [Switching Profiles for Logical Switches and Logical Ports](#).

- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Logical Switches > Switches**.
- 3 Click the logical switch to apply the custom switching profile.
- 4 Click the **Manage** tab.
- 5 Select the custom switching profile type from the drop-down menu.
 - QoS
 - Port Mirroring
 - IP Discovering
 - SpoofGuard
 - Switch Security
 - MAC Management
- 6 Click **Change**.
- 7 Select the previously created custom switching profile from the drop-down menu.
- 8 Click **Save**.
The logical switch is now associated with the custom switching profile.
- 9 Verify that the new custom switching profile with the modified configuration appears under the **Manage** tab.
- 10 (Optional) Click the **Related** tab and select **Ports** from the drop-down menu to verify that the custom switching profile is applied to child logical ports.

What to do next

If you do not want to use the inherited switching profile from a logical switch, you can apply a custom switching profile to the child logical switch port. See [Associate a Custom Profile with a Logical Port in Manager Mode](#).

Associate a Custom Profile with a Logical Port in Manager Mode

A logical port provides a logical connection point for a VIF, a patch connection to a router, or a Layer 2 gateway connection to an external network. Logical ports also expose switching profiles, port statistics counters, and a logical link status.

You can change the inherited switching profile from the logical switch to a different custom switching profile for the child logical port.

Prerequisites

- Verify that a logical port is configured. See [Connecting a VM to a Logical Switch in Manager Mode](#).
- Verify that a custom switching profile is configured. See [Switching Profiles for Logical Switches and Logical Ports](#).
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Logical Switches > Ports**.
- 3 Click the logical port to apply the custom switching profile.
- 4 Click the **Manage** tab.
- 5 Select the custom switching profile type from the drop-down menu.
 - QoS
 - Port Mirroring
 - IP Discovering
 - SpoofGuard
 - Switch Security
 - MAC Management
- 6 Click **Change**.
- 7 Select the previously created custom switching profile from the drop-down menu.
- 8 Click **Save**.

The logical port is now associated with the custom switching profile.
- 9 Verify that the new custom switching profile with the modified configuration appears under the **Manage** tab.

What to do next

You can monitor the activity on the logical switch port to troubleshoot problems. See [Monitor a Logical Switch Port Activity in Manager Mode](#).

Layer 2 Bridging in Manager Mode

When an NSX-T Data Center logical switch requires a Layer 2 connection to a VLAN-backed port group or needs to reach another device, such as a gateway, that resides outside of an NSX-T

Data Center deployment, you can use an NSX-T Data Center Layer 2 bridge. This Layer 2 bridge is especially useful in a migration scenario, in which you need to split a subnet across physical and virtual workloads.

The NSX-T Data Center concepts involved in Layer 2 bridging are Edge Clusters and Edge Bridge profiles. You can configure layer 2 bridging using NSX Edge transport nodes. To use NSX Edge transport nodes for bridging, you create an Edge bridge profile. An Edge Bridge profile specifies which Edge Cluster to use for bridging and which Edge Transport node acts as the primary and backup bridge.

The Edge Bridge Profile is attached to a logical switch and the mapping specifies the physical uplink on the Edge used for bridging and the VLAN ID to be associated with the logical switch. A logical switch can be attached to several bridge profiles.

Create an Edge Bridge Profile in Manager Mode

An Edge bridge profile makes an NSX Edge cluster capable of providing layer 2 bridging to a logical switch.

When you create an edge bridge profile, if you set the failover mode to be preemptive and a failover occurs, the standby node becomes the active node. After the failed node recovers, it becomes the active node again. If you set the failover mode to be non-preemptive and a failover occurs, the standby node becomes the active node. After the failed node recovers, it becomes the standby node. You can manually set the standby edge node to be the active node by running the CLI command `set l2bridge-port <uuid> state active` on the standby edge node. The command can only be applied in non-preemptive mode. Otherwise, there will be an error. In non-preemptive mode, the command will trigger an HA failover when applied on a standby node, and it will be ignored when applied on an active node. For more information, see the *NSX-T Data Center Command-Line Interface Reference*.

Prerequisites

- Verify that you have an NSX Edge cluster with two NSX Edge transport nodes.
- Verify that you are in Manager mode.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Logical Switches > Edge Bridge Profiles > Add**.
- 3 Enter a name for the Edge bridge profile and optionally a description.
- 4 Select an NSX Edge cluster.
- 5 Select a primary node.
- 6 Select a backup node.

7 Select a failover mode.

The options are **Preemptive** and **Non-Preemptive**.

8 Click the **Add** button.

What to do next

Configure Edge-based bridging. See [Configure Edge-Based Bridging](#).

Configure Edge-Based Bridging

When you configure edge-based bridging, after creating an edge bridge profile for an edge cluster, some additional configurations are required.

Note that bridging a logical switch twice on the same Edge node is not supported. However, you can bridge two VLANs to the same logical switch on two different Edge nodes.

There are three configuration options.

Option 1: Configure Promiscuous Mode

- Set promiscuous mode on the portgroup.
- Allow forged transmit on the portgroup.
- Run the following command to enable reverse filter on the ESXi host where the Edge VM is running:

```
esxcli system settings advanced set -o /Net/ReversePathFwdCheckPromisc -i 1
```

Then disable and enable promiscuous mode on the portgroup with the following steps:

- Edit the portgroup's settings.
- Disable promiscuous mode and save the settings.
- Edit the portgroup's settings again.
- Enable promiscuous mode and save the settings.
- Do not have other port groups in promiscuous mode on the same host sharing the same set of VLANs.
- The active and standby Edge VMs should be on different hosts. If they are on the same host the throughput might be reduced because VLAN traffic needs to be forwarded to both VMs in promiscuous mode.

Option 2: Configure MAC Learning

If the Edge is deployed on a host with NSX-T installed, it can connect to a VLAN logical switch or segment. The logical switch must have a MAC Management profile with MAC Learning enabled. Similarly, the segment must have a MAC Discovery profile with MAC Learning enabled.

Option 3: Configure a Sink Port

- 1 Retrieve the port number for the trunk vNIC that you want to configure as a sink port.
 - a Log in to the vSphere Web Client, and navigate to **Home > Networking**.
 - b Click the distributed port group to which the NSX Edge trunk interface is connected, and click **Ports** to view the ports and connected VMs. Note the port number associated with the trunk interface. Use this port number when fetching and updating opaque data.
- 2 Retrieve the dvsUuid value for the vSphere Distributed Switch.
 - a Log in to the vCenter Mob UI at `https://<vc-ip>/mob`.
 - b Click **content**.
 - c Click the link associated with the **rootFolder** (for example: *group-d1 (Datacenters)*).
 - d Click the link associated with the **childEntity** (for example: *datacenter-1*).
 - e Click the link associated with the **networkFolder** (for example: *group-n6*).
 - f Click the DVS name link for the vSphere distributed switch associated with the NSX Edges (for example: *dvs-1 (Mgmt_VDS)*).
 - g Copy the value of the uuid string. Use this value for dvsUuid when fetching and updating opaque data.
- 3 Verify if opaque data exists for the specified port.
 - a Go to `https://<vc-ip>/mob/?moid=DVSManager&vmodl=1`.
 - b Click **fetchOpaqueDataEx**.
 - c In the **selectionSet** value box paste the following XML input:

```
<selectionSet xsi:type="DVPortSelection">
  <dvsUuid>c2 1d 11 50 6a 7c 77 68-e6 ba ce 6a 1d 96 2a 15</dvsUuid> <!-- example
dvsUuid -->
  <portKey>393</portKey> <!-- example port number -->
</selectionSet>
```

Use the port number and dvsUuid value that you retrieved for the NSX Edge trunk interface.

- d Set **isRuntime** to **false**.
 - e Click **Invoke Method**. If the result shows values for `vim.dvs.OpaqueData.ConfigInfo`, then there is already opaque data set, use the **edit** operation when you set the sink port. If the value for `vim.dvs.OpaqueData.ConfigInfo` is empty, use the **add** operation when you set the sink port.
- 4 Configure the sink port in the vCenter managed object browser (MOB).
 - a Go to `https://<vc-ip>/mob/?moid=DVSManager&vmodl=1`.
 - b Click **updateOpaqueDataEx**.

- c In the **selectionSet** value box paste the following XML input. For example,

```
<selectionSet xsi:type="DVPortSelection">
  <dvsUuid>c2 1d 11 50 6a 7c 77 68-e6 ba ce 6a 1d 96 2a 15</dvsUuid> <!-- example
dvsUuid -->
  <portKey>393</portKey> <!-- example port number -->
</selectionSet>
```

Use the dvsUuid value that you retrieved from the vCenter MOB.

- d On the opaqueDataSpec value box paste one of the following XML inputs.

Use this input to enable a SINK port if opaque data is not set (operation is set to add):

```
<opaqueDataSpec>
  <operation>add</operation>
  <opaqueData>
    <key>com.vmware.etherswitch.port.extraEthFRP</key>
    <opaqueData
xsi:type="vmobl.Binary">AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
    </opaqueData>
  </opaqueDataSpec>
```

Use this input to enable a SINK port if opaque data is already set (operation is set to edit):

```
<opaqueDataSpec>
  <operation>edit</operation>
  <opaqueData>
    <key>com.vmware.etherswitch.port.extraEthFRP</key>
    <opaqueData
xsi:type="vmobl.Binary">AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
    </opaqueData>
  </opaqueDataSpec>
```

Use this input to disable a SINK port:

```
<opaqueDataSpec>
  <operation>edit</operation>
  <opaqueData>
    <key>com.vmware.etherswitch.port.extraEthFRP</key>
    <opaqueData
xsi:type="vmobl.Binary">AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
    </opaqueData>
  </opaqueDataSpec>
```


- e Set `isRuntime` to `false`.
- f Click **Invoke Method**.

What to do next

Associate a logical switch with the bridge profile. See [Create a Layer 2 Bridge-Backed Logical Switch in Manager Mode](#).

Create a Layer 2 Bridge-Backed Logical Switch in Manager Mode

When you have VMs that are connected to the NSX-T Data Center overlay, you can configure a bridge-backed logical switch to provide layer 2 connectivity with other devices or VMs that are outside of your NSX-T Data Center deployment.

Prerequisites

- Verify that you have an Edge bridge profile.
- Configure one of the following options: promiscuous mode, MAC learning, or a sink port. See [Configure Edge-Based Bridging](#).
- At least one ESXi or KVM host to serve as a regular transport node. This node has hosted VMs that require connectivity with devices outside of a NSX-T Data Center deployment.
- A VM or another end device outside of the NSX-T Data Center deployment. This end device must be attached to a VLAN port matching the VLAN ID of the bridge-backed logical switch.
- One logical switch in an overlay transport zone to serve as the bridge-backed logical switch.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From a browser, log in to an NSX Manager at `https://<nsx-mgr>`.
- 2 Select **Networking > Logical Switches**.
- 3 Click the name of an overlay switch (traffic type: overlay).
- 4 Click **Related > Edge Bridge Profiles**.
- 5 Click **Attach**.
- 6 To attach to an Edge bridge profile,
 - a Select an Edge bridge profile.
 - b Select a transport zone.
 - c Enter a VLAN ID.
 - d Click **Save**.

7 Connect VMs to the logical switch if they are not already connected.

The VMs must be on transport nodes in the same transport zone as the Edge bridge profile.

Results

You can test the functionality of the bridge by sending a ping from the NSX-T Data Center-internal VM to a node that is external to NSX-T Data Center.

You can monitor traffic on the bridge switch by clicking the **Monitor** tab.

You can also view the bridge traffic with the `GET https://192.168.110.31/api/v1/bridge-endpoints/<endpoint-id>/statistics` API call:

```
{
  "tx_packets": {
    "total": 134416,
    "dropped": 0,
    "multicast_broadcast": 0
  },
  "rx_bytes": {
    "total": 22164,
    "multicast_broadcast": 0
  },
  "tx_bytes": {
    "total": 8610134,
    "multicast_broadcast": 0
  },
  "rx_packets": {
    "total": 230,
    "dropped": 0,
    "multicast_broadcast": 0
  },
  "last_update_timestamp": 1454979822860,
  "endpoint_id": "ba5ba59d-22f1-4a02-b6a0-18ef0e37ef31"
}
```

Logical Routers in Manager Mode

NSX-T Data Center supports a 2-tier routing model.

In the top tier is the tier-0 logical router. Northbound, the tier-0 logical router connects to one or more physical routers or layer 3 switches and serves as a gateway to the physical infrastructure. Southbound, the tier-0 logical router connects to one or more tier-1 logical routers or directly to one or more logical switches.

In the bottom tier is the tier-1 logical router. Northbound, the tier-1 logical router connects to a tier-0 logical router. Southbound, it connects to one or more logical switches.


Note If you use **Manager** mode to modify objects created in the **Policy** mode, some settings might not be configurable. These read-only settings have this icon next to them: . See [Chapter 1 NSX Manager](#) for more information.

Tier-1 Logical Router

Tier-1 logical routers have downlink ports to connect to logical switches and uplink ports to connect to tier-0 logical routers.

When you add a logical router, it is important that you plan the networking topology you are building.

Figure 20-3. Tier-1 Logical Router Topology



For example, this simple topology shows two logical switches connected to a tier-1 logical router. Each logical switch has a single VM connected. The two VMs can be on different hosts or the same host, in different host clusters or in the same host cluster. If a logical router does not separate the VMs, the underlying IP addresses configured on the VMs must be in the same subnet. If a logical router does separate them, the IP addresses on the VMs must be in different subnets.

In some scenarios, external clients send ARP queries for MAC addresses bound to LB VIP ports. However, LB VIP ports do not have MAC addresses and cannot handle such queries. Proxy ARP is implemented on the centralized service ports of a tier-1 logical router to handle ARP queries on behalf of the LB VIP ports.

When a tier-1 logical router is configured with DNAT, Edge firewall, and load balancer, traffic to and from another tier-1 logical router is processed in this order: DNAT first, then Edge firewall, and then load balancer. Traffic within the tier-1 logical router is processed through DNAT first and then load balancer. Edge firewall processing is skipped.

On a tier-0 or tier-1 logical router, you can configure different types of ports. One type is called centralized service port (CSP). You must configure a CSP on a tier-0 logical router in active-standby mode or a tier-1 logical router to connect to a VLAN-backed logical switch, or to create a standalone tier-1 logical router. A CSP supports the following services on a tier-0 logical router in active-standby mode or a tier-1 logical router:

- NAT
- Load balancing
- Stateful firewall
- VPN (IPsec and L2VPN)

Create a Tier-1 Logical Router in Manager Mode

The tier-1 logical router must be connected to the tier-0 logical router to get the northbound physical router access.

Prerequisites

- Verify that the logical switches are configured. See [Create a Logical Switch in Manager Mode](#).
- Verify that an NSX Edge cluster is deployed to perform network address translation (NAT) configuration. See the *NSX-T Data Center Installation Guide*.
- Familiarize yourself with the tier-1 logical router topology. See [Tier-1 Logical Router](#).
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Tier-1 Logical Routers > Add**.
- 3 Enter a name for the logical router and optionally a description.
- 4 (Optional) Select a tier-0 logical router to connect to this tier-1 logical router.

If you do not yet have any tier-0 logical routers configured, you can leave this field blank for now and edit the router configuration later.

5 (Optional) Select an NSX Edge cluster.

To deselect a cluster that you selected, click the **x** icon. If the tier-1 logical router is going to be used for NAT configuration, it must be connected to an NSX Edge cluster. If you do not yet have any NSX Edge clusters configured, you can leave this field blank for now and edit the router configuration later.

6 (Optional) Click the **StandBy Relocation** toggle to enable or disable standby relocation.

Standby relocation means that if the Edge node where the active or standby logical router is running fails, a new standby logical router is created on another Edge node to maintain high availability. If the Edge node that fails is running the active logical router, the original standby logical router becomes the active logical router and a new standby logical router is created. If the Edge node that fails is running the standby logical router, the new standby logical router replaces it.

7 (Optional) If you selected an NSX Edge cluster, select a failover mode.

Option	Description
Preemptive	If the preferred node fails and recovers, it will preempt its peer and become the active node. The peer will change its state to standby. This is the default option.
Non-preemptive	If the preferred node fails and recovers, it will check if its peer is the active node. If so, the preferred node will not preempt its peer and will be the standby node.

8 (Optional) Click the **Advanced** tab and enter a value for **Intra Tier-1 Transit Subnet**.

9 Click **Add**.

Results

After the logical router is created, if you want to remove the Edge cluster from the router's configuration, perform the following steps:

- Click the name of the router to see the configuration details.
- Select **Services > Edge Firewall**.
- Click **Disable Firewall**.
- Click the **Overview** tab and click **Edit**.
- In the **Edge Cluster** field, click the **x** icon.
- Click **Save**.

If this logical router supports more than 5000 VMs, you must run the following commands on each node of the NSX Edge cluster to increase the size of the ARP table.

```
set debug-mode
set dataplane neighbor max-arp-logical-router 10000
```

You must re-run the commands after a dataplane restart or a node reboot because the change is not persistent.

What to do next

Create downlink ports for your tier-1 logical router. See [Add a Downlink Port on a Tier-1 Logical Router in Manager Mode](#).

Add a Downlink Port on a Tier-1 Logical Router in Manager Mode

When you create a downlink port on a tier-1 logical router, the port serves as a default gateway for the VMs that are in the same subnet.

Prerequisites

- Verify that a tier-1 logical router is configured. See [Create a Tier-1 Logical Router in Manager Mode](#).
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Tier-1 Logical Routers**.
- 3 Click the name of a tier-1 router.
- 4 Click the **Configuration** tab and select **Router Ports**.
- 5 Click **Add**.
- 6 Enter a name for the router port and optionally a description.
- 7 In the **Type** field, select **Downlink**.
- 8 For **URPF Mode**, select **Strict** or **None**.
URPF (unicast Reverse Path Forwarding) is a security feature.
- 9 (Optional) Select a logical switch.
- 10 Select whether this attachment creates a switch port or updates an existing switch port.
If the attachment is for an existing switch port, select the port from the drop-down menu.
- 11 Enter an IP address and a prefix length for the router port.
- 12 (Optional) Select a DHCP relay service.
- 13 Click **Add**.

What to do next

Enable route advertisement to provide North-South connectivity between VMs and external physical networks or between different tier-1 logical routers that are connected to the same tier-0 logical router. See [Configure Route Advertisement on a Tier-1 Logical Router in Manager Mode](#).

Add a VLAN Port on a Tier-0 or Tier-1 Logical Router in Manager Mode

If you have only VLAN-backed logical switches, you can connect the switches to VLAN ports on a tier-0 or tier-1 router so that NSX-T Data Center can provide layer-3 services.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).


Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Locate the router in **Networking > Tier-0 Logical Routers** or **Networking > Tier-1 Logical Routers** and select it.
- 3 Click the **Configuration** tab and select **Router Ports**.
- 4 Click **Add**.
- 5 Enter a name for the router port and optionally a description.
- 6 In the **Type** field, select **Centralized**.
- 7 For **URPF Mode**, select **Strict** or **None**.
URPF (unicast Reverse Path Forwarding) is a security feature.
- 8 (Required) Select a logical switch.
- 9 Select whether this attachment creates a switch port or updates an existing switch port.
If the attachment is for an existing switch port, select the port from the drop-down menu.
- 10 Enter the router port IP address in CIDR notation.
- 11 Click **Add**.

Configure Route Advertisement on a Tier-1 Logical Router in Manager Mode

To provide Layer 3 connectivity between VMs connected to logical switches that are attached to different tier-1 logical routers, it is necessary to enable tier-1 route advertisement towards tier-0. You do not need to configure a routing protocol or static routes between tier-1 and tier-0 logical routers. NSX-T Data Center creates NSX-T Data Center static routes automatically when you enable route advertisement.

For example, to provide connectivity to and from the VMs through other peer routers, the tier-1 logical router must have route advertisement configured for connected routes. If you don't want to advertise all connected routes, you can specify which routes to advertise.



Prerequisites

- Verify that VMs are attached to logical switches. See [Logical Switches in Manager Mode](#).
- Verify that downlink ports for the tier-1 logical router are configured. See [Add a Downlink Port on a Tier-1 Logical Router in Manager Mode](#).
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Tier-1 Logical Routers**.
- 3 Click the name of a tier-1 router.
- 4 Select **Route Advertisement** from the **Routing** drop-down menu.

- 5 Click **Edit** to edit the route advertisement configuration.

You can toggle the following switches:

- **Status**
- **Advertise All NSX Connected Routes**
- **Advertise All NAT Routes**
- **Advertise All Static Routes**
- **Advertise All LB VIP Routes**
- **Advertise All LB SNAT IP Routes**
- **Advertise All DNS Forwarder Routes**

- a Click **Save**.

- 6 Click **Add** to advertise routes.

- a Enter a name and optionally a description.
- b Enter a route prefix in CIDR format.
- c Click **Apply Filter** to set the following options:

Action	Specify Allow or Deny .
Match route types	Select one or more of the following: ■ Any ■ NSX Connected ■ Tier-1 LB VIP ■ Static ■ Tier-1 NAT ■ Tier-1 LB SNAT
Prefix operator	Select GE (greater than or equal) or EQ (equal).

- d Click **Add**.

What to do next

Familiarize yourself with the tier-0 logical router topology and create the tier-0 logical router. See [Tier-0 Logical Router](#).


If you already have a tier-0 logical router connected to the tier-1 logical router, you can verify that the tier-0 router is learning the tier-1 router connected routes. See [Verify that a Tier-0 Router Has Learned Routes from a Tier-1 Router](#).

Configure a Tier-1 Logical Router Static Route in Manager Mode

You can configure a static route on a tier-1 logical router to provide connectivity from NSX-T Data Center to a set of networks that are accessible through a virtual router.

For example, in the following diagram, the tier-1 A logical router has a downlink port to an NSX-T Data Center logical switch. This downlink port (172.16.40.1) serves the default gateway for the virtual router VM. The virtual router VM and tier-1 A are connected through the same NSX-T Data Center logical switch. The tier-1 logical router has a static route 10.10.0.0/16 that summarizes the networks available through the virtual router. Tier-1 A then has route advertisement configured to advertise the static route to tier-1 B.

Figure 20-4. Tier-1 Logical Router Static Route Topology



Recursive static routes are supported.

Prerequisites

- Verify that a downlink port is configured. See [Add a Downlink Port on a Tier-1 Logical Router in Manager Mode](#).
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Verify that a downlink port is configured. See [Add a Downlink Port on a Tier-1 Logical Router in Manager Mode](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Tier-1 Logical Routers**.
- 3 Click the name of a tier-1 router.
- 4 Click the **Routing** tab and select **Static Routes** from the drop-down menu.
- 5 Click **Add**.
- 6 Enter a network address in the CIDR format.

Static route based on IPv6 is supported. IPv6 prefixes can only have an IPv6 next hop.
For example, 10.10.10.0/16 or an IPv6 address.
- 7 Click **Add** to add a next-hop IP address.

For example, 172.16.40.10. You can also specify a null route by clicking the pencil icon and selecting **NULL** from the drop-down. To add another next hop addresses, click **Add** again.
- 8 Click **Add** at the bottom of the dialog box.

The newly created static route network address appears in the row.
- 9 From the tier-1 logical router, select **Routing > Route Advertisement**.
- 10 Click **Edit** and select **Advertise All Static Routes**.
- 11 Click **Save**.

The static route is propagated across the NSX-T Data Center overlay.

Create a Standalone Tier-1 Logical Router in Manager Mode

A standalone tier-1 logical router has no downlink and no connection to a tier-0 router. It has a service router but no distributed router. The service router can be deployed on one NSX Edge node or two NSX Edge nodes in active-standby mode.

A standalone tier-1 logical router:

- Must not have a connection to a tier-0 logical router.
- Can have only one centralized service port (CSP) if it is used to attach a load balancer (LB) service.
- Can connect to an overlay logical switch or a VLAN logical switch.
- Supports any combination of the services IPsec, NAT, firewall, load balancer, and service insertion. For ingress, the order of processing is: IPsec – DNAT – firewall – load balancer - service insertion. For egress, the order of processing is: service insertion - load balancer - firewall - SNAT - IPsec.

Typically, a standalone tier-1 logical router is connected to a logical switch that a regular tier-1 logical router is also connected to. The standalone tier-1 logical router can communicate with other devices through the regular tier-1 logical router after static routes and route advertisements are configured.

Before using the standalone tier-1 logical router, note the following:

- To specify the default gateway for the standalone tier-1 logical router, you must add a static route. The subnet should be 0.0.0.0/0 and the next hop is the IP address of a regular tier-1 router connected to the same switch.
- ARP proxy on the standalone router is supported. You can configure an LB virtual server IP or LB SNAT IP in the CSP's subnet. For example, if the CSP IP is 1.1.1.1/24, the virtual IP can be 1.1.1.2. It can also be an IP in another subnet such as 2.2.2.2 if routing is properly configured so that traffic for 2.2.2.2 can reach the standalone router.
- For an NSX Edge VM, you cannot have more than one CSPs which are connected to the same VLAN-backed logical switch or different VLAN-backed logical switches that have the same VLAN ID.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Tier-1 Logical Routers > Add**.
- 3 Enter a name for the logical router, and optionally a description.
- 4 (Required) Select an NSX Edge cluster to connect to this tier-1 logical router.
- 5 (Required) Select a failover mode and cluster members.

Option	Description
Preemptive	If the preferred node fails and recovers, it will preempt its peer and become the active node. The peer will change its state to standby. This is the default option.
Non-preemptive	If the preferred node fails and recovers, it will check if its peer is the active node. If so, the preferred node will not preempt its peer and will be the standby node.

- 6 Click **Add**.
- 7 Click the name of the router that you just created.
- 8 Click the **Configuration** tab and select **Router Ports**.
- 9 Click **Add**.

10 Enter a name for the router port and optionally a description.

11 In the **Type** field, select **Centralized**.

12 For **URPF Mode**, select **Strict** or **None**.

URPF (Unicast Reverse Path Forwarding) is a security feature.

13 (Required) Select a logical switch.

14 Select whether this attachment creates a switch port or updates an existing switch port.

15 Enter the router port IP address in CIDR notation.

16 Click **Add**.

Tier-0 Logical Router


A tier-0 logical router provides a gateway service between the logical and physical network.

NSX Cloud Note If using NSX Cloud, see [NSX-T Data Center Features Supported with NSX Cloud](#) for a list of auto-generated logical entities, supported features, and configurations required for NSX Cloud.

An Edge node can support only one tier-0 gateway or logical router. When you create a tier-0 gateway or logical router, make sure you do not create more tier-0 gateways or logical routers than the number of Edge nodes in the NSX Edge cluster.

When you add a tier-0 logical router, it is important that you map out the networking topology you are building.

Figure 20-5. Tier-0 Logical Router Topology



For simplicity, the sample topology shows a single tier-1 logical router connected to a single tier-0 logical router hosted on a single NSX Edge node. Keep in mind that this is not a recommended topology. Ideally, you should have a minimum of two NSX Edge nodes to take full advantage of the logical router design.

The tier-1 logical router has a web logical switch and an app logical switch with respective VMs attached. The router-link switch between the tier-1 router and the tier-0 router is created automatically when you attach the tier-1 router to the tier-0 router. Thus, this switch is labeled as system generated.

In some scenarios, external clients send ARP queries for MAC addresses bound to loopback or IKE IP ports. However, loopback and IKE IP ports do not have MAC addresses and cannot handle such queries. Proxy ARP is implemented on the uplink and centralized service ports of a tier-0 logical router to handle ARP queries on behalf of the loopback and IKE IP ports.

When a tier-0 logical router is configured with DNAT, IPsec, and Edge firewall, traffic is processed in this order: IPsec first, then DNAT, and then Edge firewall.

On a tier-0 or tier-1 logical router, you can configure different types of ports. One type is called centralized service port (CSP). You must configure a CSP on a tier-0 logical router in active-standby mode or a tier-1 logical router to connect to a VLAN-backed logical switch, or to create a standalone tier-1 logical router. A CSP supports the following services on a tier-0 logical router in active-standby mode or a tier-1 logical router:

- NAT
- Load balancing
- Stateful firewall
- VPN (IPsec and L2VPN)

Create a Tier-0 Logical Router in Manager Mode

Tier-0 logical routers have downlink ports to connect to NSX-T Data Center tier-1 logical routers and uplink ports to connect to external networks.

Prerequisites

- Verify that at least one NSX Edge is installed. See the *NSX-T Data Center Installation Guide*
- Verify that an NSX Edge cluster is configured. See the *NSX-T Data Center Installation Guide*.
- Familiarize yourself with the networking topology of the tier-0 logical router. See [Tier-0 Logical Router](#).
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Tier-0 Logical Routers > Add**.
- 3 Enter a name for the tier-0 logical router.
- 4 Select an existing NSX Edge cluster from the drop-down menu to back this tier-0 logical router.
- 5 (Optional) Select a high-availability mode.

By default, the active-active mode is used. In the active-active mode, traffic is load balanced across all members. In active-standby mode, all traffic is processed by an elected active member. If the active member fails, a new member is elected to be active.

- 6
- 7 (Optional) Click the **Advanced** tab to enter a subnet for the intra-tier 0 transit subnet.

This is the subnet that connects to the tier-0 services router to its distributed router. If you leave this blank, the default 169.0.0.0/28 subnet is used.

- 8 (Optional) Click the **Advanced** tab to enter a subnet for the tier-0-to-tier-1 transit subnet.

This is the subnet that connects the tier-0 router to any tier-1 routers that connect to this tier-0 router. If you leave this blank, the default address space assigned for these tier-0-to-tier-1 connections is 100.64.0.0/16. Each tier-0-to-tier-1 peer connection is provided a /31 subnet within the 100.64.0.0/16 address space.

- 9 Click **Save**.

The new tier-0 logical router appears as a link.

- 10 (Optional) Click the tier-0 logical router link to review the summary.

What to do next

Attach tier-1 logical routers to this tier-0 logical router.


Configure the tier-0 logical router to connect it to a VLAN logical switch to create an uplink to an external network. See [Connect a Tier-0 Logical Router to a VLAN Logical Switch for the NSX Edge Uplink in Manager Mode](#).

Attach Tier-1 Router to a Tier-0 Router in Manager Mode

You can attach the tier-0 logical router to the tier-1 logical router so that the tier-1 logical router gets northbound and east-west network connectivity.

When you attach a tier-1 logical router to a tier-0 logical router, a router-link switch between the two routers is created. This switch is labeled as system-generated in the topology. The default address space assigned for these tier-0-to-tier-1 connections is 100.64.0.0/16. Each tier-0-to-tier-1 peer connection is provided a /31 subnet within the 100.64.0.0/16 address space. Optionally, you can configure the address space in the tier-0 **Summary > Advanced** configuration.

The following figure shows a sample topology.



Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Tier-1 Logical Routers**.
- 3 Select the tier-1 logical router.
- 4 In the Tier-0 Connection section, click **Connect**.
- 5 Select a tier-0 logical router from the drop-down menu.
- 6 (Optional) Select an NSX Edge cluster from the drop-down menu.

The tier-1 router needs to be backed by an edge device if the router is going to be used for services, such as NAT. If you do not select an NSX Edge cluster, the tier-1 router cannot perform NAT.

- 7 Specify members and a preferred member.

If you select an NSX Edge cluster and leave the members and preferred member fields blank, NSX-T Data Center sets the backing edge device from the specified cluster for you.

- 8 Click **Save**.
- 9 Click the **Configuration** tab of the tier-1 router to verify that a new point-to-point linked port IP address is created.

For example, the IP address of the linked port can be 100.64.1.1/31.

- 10 Select the tier-0 logical router from the navigation panel.
- 11 Click the **Configuration** tab of the tier-0 router to verify that a new point-to-point linked port IP address is created.

For example, the IP address of the linked port can be 100.64.1.1/31.

What to do next

Verify that the tier-0 router is learning routes that are advertised by the tier-1 routers.

Verify that a Tier-0 Router Has Learned Routes from a Tier-1 Router

When a tier-1 logical router advertises routes to a tier-0 logical router, the routes are listed in the tier-0 router's routing table as NSX-T Data Center static routes.

Procedure

- 1 On the NSX Edge, run the `get logical-routers` command to find the VRF number of the tier-0 service router.

```

nsx-edge-1> get logical-routers
Logical Router
UUID          : 736a80e3-23f6-5a2d-81d6-bbefb2786666
vrf           : 0
type          : TUNNEL

Logical Router
UUID          : 421a2d0d-f423-46f1-93a1-2f9e366176c8
vrf           : 5
type          : SERVICE_ROUTER_TIER0

Logical Router
UUID          : f3ce9d7d-7123-47d6-aba6-45cf1388ca7b
vrf           : 6
type          : DISTRIBUTED_ROUTER

Logical Router
UUID          : c8e64eff-02b2-4462-94ff-89f3788f1a61
vrf           : 7
type          : SERVICE_ROUTER_TIER1

Logical Router
UUID          : fb6c3f1f-599f-4421-af8a-99692dff3dd4
vrf           : 8
type          : DISTRIBUTED_ROUTER

```

- 2 Run the `vrf <number>` command to enter the tier-0 service router context.

```

nsx-edge-1> vrf 5
nsx-edge1(tier0_sr)>

```

- 3 On the tier-0 service router, run the `get route` command and make sure the expected routes appear in the routing table.

Notice that the NSX-T Data Center static routes (ns) are learned by the tier-0 router because the tier-1 router is advertising routes.

```

nsx-edge1(tier0_sr)> get route

Flags: c - connected, s - static, b - BGP, ns - nsx_static
nc - nsx_connected, rl - router_link, t0n: Tier0-NAT, t1n: Tier1-NAT

Total number of routes: 7

b    10.10.10.0/24      [20/0]      via 192.168.100.254


```

r1	100.91.176.0/31	[0/0]	via 169.254.0.1
c	169.254.0.0/28	[0/0]	via 169.254.0.2
ns	172.16.10.0/24	[3/3]	via 169.254.0.1
ns	172.16.20.0/24	[3/3]	via 169.254.0.1
c	192.168.100.0/24	[0/0]	via 192.168.100.2

Connect a Tier-0 Logical Router to a VLAN Logical Switch for the NSX Edge Uplink in Manager Mode

To create an NSX Edge uplink, you must connect a tier-0 router to a VLAN switch.

The following simple topology shows a VLAN logical switch inside of a VLAN transport zone. The VLAN logical switch has a VLAN ID that matches the VLAN ID on the TOR port for the Edge's VLAN uplink.



Prerequisites

- Create a VLAN logical switch. See [Create a VLAN Logical Switch for the NSX Edge Uplink in Manager Mode](#).
- Create a tier-0 router.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

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

- 2 Select **Networking > Tier-0 Logical Routers**.
- 3 Select the tier-0 logical router.
- 4 From the **Configuration** tab, add a new logical router port.
- 5 Type a name for the port, such as uplink.
- 6 Select the **Uplink** type.
- 7 Select an edge transport node.
- 8 Select a VLAN logical switch.
- 9 Select either **Attach to new switch port** or **Attach to existing switch port**.

If you select **Attach to new switch port**, the port will be automatically created.

- 10 Specify an IP address that is in the same subnet as the connected port on the TOR switch.

Results

A new uplink port is added for the tier-0 router.

What to do next

Configure BGP or a static route.

Verify the Tier-0 Logical Router and TOR Connection

For routing to work on the uplink from the tier-0 router, connectivity with the top-of-rack device must be in place.

Prerequisites

- Verify that the tier-0 logical router is connected to a VLAN logical switch. See [Connect a Tier-0 Logical Router to a VLAN Logical Switch for the NSX Edge Uplink in Manager Mode](#).

Procedure

- 1 Log in to the NSX Edge CLI.
- 2 On the NSX Edge, run the `get logical-routers` command to find the VRF number of the tier-0 service router.

```
nsx-edge-1> get logical-routers
Logical Router
UUID          : 736a80e3-23f6-5a2d-81d6-bbefb2786666
vrf           : 0
type          : TUNNEL

Logical Router
UUID          : 421a2d0d-f423-46f1-93a1-2f9e366176c8
vrf           : 5
type          : SERVICE_ROUTER_TIER0
```

```

Logical Router
UUID       : f3ce9d7d-7123-47d6-aba6-45cf1388ca7b
vrf        : 6
type       : DISTRIBUTED_ROUTER

Logical Router
UUID       : c8e64eff-02b2-4462-94ff-89f3788f1a61
vrf        : 7
type       : SERVICE_ROUTER_TIER1

Logical Router
UUID       : fb6c3f1f-599f-4421-af8a-99692dff3dd4
vrf        : 8
type       : DISTRIBUTED_ROUTER

```

- 3 Run the `vrf <number>` command to enter the tier-0 service router context.

```

nsx-edge-1> vrf 5
nsx-edgel(tier0_sr)>

```

- 4 On the tier-0 service router, run the `get route` command and make sure the expected route appears in the routing table.

Notice that the route to the TOR appears as connected (c).

```

nsx-edgel(tier0_sr)> get route
Flags: t0c - Tier0-Connected, t0s - Tier0-Static, b - BGP,
t0n - Tier0-NAT, t1s - Tier1-Static, t1c - Tier1-Connected,
t1n: Tier1-NAT, t1l: Tier1-LB VIP, t1ls: Tier1-LB SNAT,
t1ld: Tier1-DNS FORWARDER, t1lipsec: Tier1-IPSec, isr: Inter-SR,
> - selected route, * - FIB route

Total number of routes: 11

t1c> * 1.1.1.0/25 [3/0] via 100.64.1.1, downlink-282, 08w4d03h
t1c> * 1.1.2.0/24 [3/0] via 100.64.1.1, downlink-282, 08w4d03h
t0c> * 1.1.3.0/24 is directly connected, downlink-275, 08w4d03h
b > * 2.1.4.0/24 [20/0] via 40.40.40.10, uplink-273, 01w0d02h
b > * 10.182.48.0/20 [20/0] via 40.40.40.10, uplink-273, 01w0d02h
t0c> * 40.40.40.0/24 is directly connected, uplink-273, 08w4d03h
t0c> * 100.64.1.0/31 is directly connected, downlink-282, 08w4d03h
t0c> * 169.254.0.0/24 is directly connected, downlink-277, 01w0d02h
b > * 172.17.0.0/16 [20/0] via 40.40.40.10, uplink-273, 01w0d02h
t0c> * fc36:a750:db0d:7800::/64 is directly connected, downlink-282, 08w4d03h
t0c> * fe80::/64 is directly connected, downlink-282, 08w4d03h

```

5 Ping the TOR.

```

nsx-edge1(tier0_sr)> ping      192.168.100.254
PING 192.168.100.254 (192.168.100.254): 56 data bytes
64 bytes from 192.168.100.254: icmp_seq=0 ttl=64 time=2.822 ms
64 bytes from 192.168.100.254: icmp_seq=1 ttl=64 time=1.393 ms
^C
nsx-edge1>
--- 192.168.100.254 ping statistics ---
3 packets transmitted, 2 packets received, 33.3% packet loss
round-trip min/avg/max/stddev = 1.393/2.107/2.822/0.715 ms

```

Results

Packets are sent between the tier-0 logical router and physical router to verify a connection.

What to do next

Depending on your networking requirements, you can configure a static route or BGP. See [Configure a Static Route in Manager Mode](#) or [Configure BGP on a Tier-0 Logical Router in Manager Mode](#).

Add a Loopback Router Port in Manager Mode

You can add a loopback port to a tier-0 logical router.

The loopback port can be used for the following purposes:

- Router ID for routing protocols
- NAT
- BFD
- Source address for routing protocols

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Tier-0 Logical Routers**.
- 3 Select the tier-0 logical router.
- 4 Select **Configuration > Router Ports**

- 5 Click **Add**.
- 6 Enter a name and optionally a description.
- 7 Select the **Loopback** type.
- 8 Select an edge transport node.
- 9 Enter an IP address in CIDR format.

Results

A new port is added for the tier-0 router.

Add a VLAN Port on a Tier-0 or Tier-1 Logical Router in Manager Mode

If you have only VLAN-backed logical switches, you can connect the switches to VLAN ports on a tier-0 or tier-1 router so that NSX-T Data Center can provide layer-3 services.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Locate the router in **Networking > Tier-0 Logical Routers** or **Networking > Tier-1 Logical Routers** and select it.
- 3 Click the **Configuration** tab and select **Router Ports**.
- 4 Click **Add**.
- 5 Enter a name for the router port and optionally a description.
- 6 In the **Type** field, select **Centralized**.
- 7 For **URPF Mode**, select **Strict** or **None**.
URPF (unicast Reverse Path Forwarding) is a security feature.
- 8 (Required) Select a logical switch.
- 9 Select whether this attachment creates a switch port or updates an existing switch port.
If the attachment is for an existing switch port, select the port from the drop-down menu.
- 10 Enter the router port IP address in CIDR notation.
- 11 Click **Add**.

Configure High Availability VIP in Manager Mode

With HA VIP (high availability virtual IP) configured, a tier-0 logical router is operational even if one uplink is down. The physical router interacts with the HA VIP only.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure


- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Tier-0 Logical Routers**.
- 3 Click the tier-0 logical router name.
- 4 Click **Configuration > HA VIP**.
- 5 Click **Add**.
- 6 Enter an IP address in CIDR format.
- 7 To enable HA VIP, set the status to **Enabled**.
- 8 Select exactly two uplink ports.
- 9 Click **Add**.

Configure a Static Route in Manager Mode

You can configure a static route on the tier-0 router to external networks. After you configure a static route, there is no need to advertise the route from tier-0 to tier-1, because tier-1 routers automatically have a static default route towards their connected tier-0 router.

The static route topology shows a tier-0 logical router with a static route to the 10.10.10.0/24 prefix in the physical architecture. For test purposes, the 10.10.10.10/32 address is configured on the external router loopback interface. The external router has a static route to the 172.16.0.0/16 prefix to reach the app and web VMs.

Figure 20-6. Static Route Topology



Recursive static routes are supported.

Prerequisites

- Verify that the physical router and tier-0 logical router are connected. See [Verify the Tier-0 Logical Router and TOR Connection](#).
- Verify that the tier-1 router is configured to advertise connected routes. See [Create a Tier-1 Logical Router in Manager Mode](#).
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Tier-0 Logical Routers**.
- 3 Select the tier-0 logical router.
- 4 Click the **Routing** tab and select **Static Route** from the drop-down menu.
- 5 Select **Add**.

- 6 Enter a network address in the CIDR format.

For example, 10.10.10.0/24.

- 7 Click **+ Add** to add a next-hop IP address.

For example, 192.168.100.254. You can also specify a null route by clicking the pencil icon and selecting **NULL** from the drop-down.

- 8 Specify the administrative distance.

- 9 Select a logical router port from the dropdown list.

The list includes IPsec Virtual Tunnel Interface (VTI) ports.

- 10 Click the **Add** button.

What to do next

Check that the static route is configured properly. See [Verify the Static Route on a Tier-0 Router](#).

Verify the Static Route on a Tier-0 Router

Use the CLI to verify that the static route is connected. You must also verify the external router can ping the internal VMs and the internal VMs can ping the external router.

Prerequisites

Verify that a static route is configured. See [Configure a Static Route in Manager Mode](#).

Procedure

- 1 Log in to the NSX Manager CLI.

2 Confirm the static route.

a Get the service router UUID information.

```
get logical-routers
```

```
nsx-edge1> get logical-routers
Logical Router
UUID       : 736a80e3-23f6-5a2d-81d6-bbefb2786666
vrf        : 2
type       : TUNNEL

Logical Router
UUID       : d40bbfa4-3e3d-4178-8615-6f42ea335037
vrf        : 4
type       : SERVICE_ROUTER_TIER0

Logical Router
UUID       : d0289ba4-250e-41b4-8ffc-7cab4a46c3e4
vrf        : 5
type       : DISTRIBUTED_ROUTER

Logical Router
UUID       : a6ee6316-2212-4171-99cc-930c98bcad7f
vrf        : 6
type       : DISTRIBUTED_ROUTER
```

b Locate the UUID information from the output.

```
Logical Router
UUID       : d40bbfa4-3e3d-4178-8615-6f42ea335037
vrf        : 4
type       : SERVICE_ROUTER_TIER0
```

c Verify that the static route works.

```
get logical-router d40bbfa4-3e3d-4178-8615-6f42ea335037 route static
```

```
Flags: c - connected, s - static, b - BGP, ns - nsx_static
nc - nsx_connected, rl - router_link, t0n: Tier0-NAT, t1n: Tier1-NAT

s    10.10.10.0/24      [1/1]      via 192.168.100.254
rl   100.64.1.0/31     [0/0]      via 169.0.0.1
ns   172.16.10.0/24    [3/3]      via 169.0.0.1
ns   172.16.20.0/24    [3/3]      via 169.0.0.1
```

- 3 From the external router, ping the internal VMs to confirm that they are reachable through the NSX-T Data Center overlay.

- a Connect to the external router.

```
ping 172.16.10.10
```

```
PING 172.16.10.10 (172.16.10.10) 56(84) bytes of data.
64 bytes from 172.16.10.10: icmp_req=1 ttl=62 time=127 ms
64 bytes from 172.16.10.10: icmp_req=2 ttl=62 time=1.96 ms
^C
--- 172.16.10.10 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1001ms
rtt min/avg/max/mdev = 1.966/64.793/127.620/62.827 ms
```

- b Test the network connectivity.

```
tracert 172.16.10.10
```

```
tracert to 172.16.10.10 (172.16.10.10), 30 hops max, 60 byte packets
 1  192.168.100.3 (192.168.100.3)  0.640 ms  0.575 ms  0.696 ms
 2  100.64.1.1 (100.64.1.1)  0.656 ms  0.604 ms  0.578 ms
 3  172.16.10.10 (172.16.10.10)  3.397 ms  3.703 ms  3.790 ms
```

- 4 From the VMs, ping the external IP address.


```
ping 10.10.10.10
```

```
PING 10.10.10.10 (10.10.10.10) 56(84) bytes of data.
64 bytes from 10.10.10.10: icmp_req=1 ttl=62 time=119 ms
64 bytes from 10.10.10.10: icmp_req=2 ttl=62 time=1.93 ms
^C
--- 10.10.10.10 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1001ms
rtt min/avg/max/mdev = 1.936/60.865/119.795/58.930 ms
```

BGP Configuration Options

To take full advantage of the tier-0 logical router, the topology must be configured with redundancy and symmetry with BGP between the tier-0 routers and the external top-of-rack peers. This design helps to ensure connectivity in the event of link and node failures.

There are two modes of configuration: active-active and active-standby. The following diagram shows two options for symmetric configuration. There are two NSX Edge nodes shown in each topology. In the case of an active-active configuration, when you create tier-0 uplink ports, you can associate each uplink port with up to eight NSX Edge transport nodes. Each NSX Edge node can have two uplinks.



For option 1, when the physical leaf-node routers are configured, they should have BGP neighborships with the NSX Edges. Route redistribution should include the same network prefixes with equal BGP metrics to all of the BGP neighbors. In the tier-0 logical router configuration, all leaf-node routers should be configured as BGP neighbors.

When you are configuring the tier-0 router's BGP neighbors, if you do not specify a local address (the source IP address), the BGP neighbor configuration is sent to all NSX Edge nodes associated with the tier-0 logical router uplinks. If you do configure a local address, the configuration goes to the NSX Edge node with the uplink owning that IP address.

In the case of option1, if the uplinks are on the same subnet on the NSX Edge nodes, it makes sense to omit the local address. If the uplinks on the NSX Edge nodes are in different subnets, the local address should be specified in the tier-0 router's BGP neighbor configuration to prevent the configuration from going to all associated NSX Edge nodes.

For option 2, ensure that the tier-0 logical router configuration includes the tier-0 services router's local IP address. The leaf-node routers are configured with only the NSX Edges that they are directly connected to as the BGP neighbor.

Configure BGP on a Tier-0 Logical Router in Manager Mode

To enable access between your VMs and the outside world, you can configure an external or internal BGP (eBGP/iBGP) connection between a tier-0 logical router and a router in your physical infrastructure.

The iBGP feature has the following capabilities and restrictions:

- Redistribution, prefix lists, and routes maps are supported.
- Route reflectors are not supported.

- BGP confederation is not supported.

When configuring BGP, you must configure a local Autonomous System (AS) number for the tier-0 logical router. For example, the following topology shows the local AS number is 64510. You must also configure the remote AS number. EBGP neighbors must be directly connected and in the same subnet as the tier-0 uplink. If they are not in the same subnet, BGP multi-hop should be used.


A tier-0 logical router in active-active mode supports inter-SR (service router) routing. If router #1 is unable to communicate with a northbound physical router, traffic is re-routed to router #2 in the active-active cluster. If router #2 is able to communicate with the physical router, traffic between router #1 and the physical router will not be affected.

In a topology with a tier-0 logical router in active-active mode attached to a tier-1 logical router in active-standby mode, you must enable inter-SR routing to handle asymmetric routing. You have asymmetric routing if you configure a static route on one of the SRs, or if one SR needs to reach another SR's uplink. In addition, note the following:

- In the case of a static route configured on one SR (for example, SR #1 on Edge node #1), another SR (for example, SR #2 on Edge node #2) might learn the same route from an eBGP peer and prefer the learned route to the static route on SR #1, which might be more efficient. To ensure that SR #2 uses the static route configured on SR #1, configure the tier-1 logical router in pre-emptive mode and configure Edge node #1 as the preferred node.
- If the tier-0 logical router has an uplink port on Edge node #1 and another uplink port on Edge node #2, ping traffic from tenant VMs to the uplinks works if the two uplinks are in different subnets. Ping traffic will fail if the two uplinks are in the same subnet.

Note Router ID used for forming BGP sessions on an edge node is automatically selected from the IP addresses configured on the uplinks of a tier-0 logical router. BGP sessions on an edge node can flap when router ID changes. This can happen when the IP address auto-selected for router ID is deleted or the logical router port on which this IP is assigned is deleted.

Figure 20-7. BGP Connection Topology



Note the following scenarios when there are connection failures involving BGP or BFD:

- With only BGP configured, if all BGP neighbors go down, the service router's state will be down.
- With only BFD configured, if all BFD neighbors go down, the service router's state will be down.
- With BGP and BFD configured, if all BGP and BFD neighbors go down, the service router's state will be down.
- With BGP and static routes configured, if all BGP neighbors go down, the service router's state will be down.
- With only static routes configured, the service router's state will always be up unless the node is experiencing a failure or in a maintenance mode.

Prerequisites

- Verify that the tier-1 router is configured to advertise connected routes. See [Configure Route Advertisement on a Tier-1 Logical Router in Manager Mode](#). This is not strictly a prerequisite for BGP configuration, but if you have a two-tier topology and you plan to redistribute your tier-1 networks into BGP, this step is required.
- Verify that a tier-0 router is configured. See [Create a Tier-0 Logical Router in Manager Mode](#).
- Make sure the tier-0 logical router has learned routes from the tier-1 logical router. See [Verify that a Tier-0 Router Has Learned Routes from a Tier-1 Router](#).

- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Tier-0 Logical Routers**.
- 3 Select the tier-0 logical router.
- 4 Click the **Routing** tab and select **BGP** from the drop-down menu.
- 5 Click **Edit**.
 - a Enter the local AS number.
For example, 64510.
 - b Click the **Status** toggle to enable or disable BGP.
 - c Click the **ECMP** toggle to enable or disable ECMP.
 - d Click the **Graceful Restart** toggle to enable or disable graceful restart.
Graceful restart is only supported if the NSX Edge cluster associated with the tier-0 router has only one edge node.
 - e If this logical router is in active-active mode, click the **Inter SR Routing** toggle to enable or disable inter-SR routing.
 - f Configure route aggregation.
 - g Click **Save**.
- 6 Click **Add** to add a BGP neighbor.
- 7 Enter the neighbor IP address.
For example, 192.168.100.254.
- 8 Specify the maximum hop limit.
The default is 1.
- 9 Enter the remote AS number.
For example, 64511 (eBGP neighbor) or 64510 (iBGP neighbor).
- 10 Configure the timers (keep alive time and hold down time) and a password.
- 11 Click the **Local Address** tab to select a local address.
 - a (Optional) Uncheck **All Uplinks** to see loopback ports as well as uplink ports.
- 12 Click the **Address Families** tab to add an address family.
- 13 Click the **BFD Configuration** tab to enable BFD.

14 Click Save.**What to do next**

Test whether BGP is working properly. See [Verify BGP Connections from a Tier-0 Service Router](#).

Verify BGP Connections from a Tier-0 Service Router

Use the CLI to verify from the tier-0 service router that a BGP connection to a neighbor is established.

Prerequisites

Verify that BGP is configured. See [Configure BGP on a Tier-0 Logical Router in Manager Mode](#).

Procedure

- 1 Log in to the NSX Manager CLI.
- 2 On the NSX Edge, run the `get logical-routers` command to find the VRF number of the tier-0 service router.

```

nsx-edge-1> get logical-routers
Logical Router
UUID          : 736a80e3-23f6-5a2d-81d6-bbefb2786666
vrf           : 0
type          : TUNNEL

Logical Router
UUID          : 421a2d0d-f423-46f1-93a1-2f9e366176c8
vrf           : 5
type          : SERVICE_ROUTER_TIER0

Logical Router
UUID          : f3ce9d7d-7123-47d6-aba6-45cf1388ca7b
vrf           : 6
type          : DISTRIBUTED_ROUTER

Logical Router
UUID          : c8e64eff-02b2-4462-94ff-89f3788f1a61
vrf           : 7
type          : SERVICE_ROUTER_TIER1

Logical Router
UUID          : fb6c3f1f-599f-4421-af8a-99692dff3dd4
vrf           : 8
type          : DISTRIBUTED_ROUTER

```

- 3 Run the `vrf <number>` command to enter the tier-0 service router context.

```
nsx-edge-1> vrf 5
nsx-edge1(tier0_sr)>
```

- 4 Verify that the BGP state is `Established`, `up`.

```
get bgp neighbor
```

```
BGP neighbor: 192.168.100.254    Remote AS: 64511
BGP state: Established, up
Hold Time: 180s    Keepalive Interval: 60s
Capabilities:
    Route Refresh: advertised and received
    Address Family: IPv4 Unicast:advertised and received
    Graceful Restart: none
    Restart Remaining Time: 0
Messages: 28 received, 31 sent
Minimum time between advertisements: 30s (default)
For Address Family IPv4 Unicast:advertised and received
    Route Refresh: 0 received, 0 sent
    Prefixes: 2 received, 2 sent, 2 advertised
1 Connections established, 2 dropped
Local host: 192.168.100.3, Local port: 179
Remote host: 192.168.100.254, Remote port: 33044
```

What to do next

Check the BGP connection from the external router. See [Verify North-South Connectivity and Route Redistribution on a Tier-0 Router](#).

Configure BFD on a Tier-0 Logical Router in Manager Mode

BFD (Bidirectional Forwarding Detection) is a protocol that can detect forwarding path failures.

Note In this release, BFD over Virtual Tunnel Interface (VTI) ports is not supported.

Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Tier-0 Logical Routers**.
- 3 Select the tier-0 logical router.

- 4 Click the **Routing** tab and select **BFD** from the drop-down menu.
- 5 Click **Edit** to configure BFD.
- 6 Click the **Status** toggle button to enable BFD.

You can optionally change the global BFD properties **Receive interval**, **Transmit interval**, and **Declare dead interval**.

- 7 (Optional) Click **Add** under BFD Peers for Static Route Next Hops to add a BFD peer.

Specify the peer IP address and set the admin status to **Enabled**. Optionally, you can override the global BFD properties **Receive interval**, **Transmit interval**, and **Declare dead interval**.

Enable Route Redistribution on the Tier-0 Logical Router in Manager Mode

When you enable route redistribution, the tier-0 logical router starts sharing specified routes with its northbound router.

Prerequisites

- Verify that the tier-0 and tier-1 logical routers are connected so that you can advertise the tier-1 logical router networks to redistribute them on the tier-0 logical router. See [Attach Tier-1 Router to a Tier-0 Router in Manager Mode](#).
- If you want to filter specific IP addresses from route redistribution, verify that route maps are configured. See [Create a Route Map in Manager Mode](#).
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Tier-0 Logical Routers**.
- 3 Select the tier-0 logical router.
- 4 Click the **Routing** tab and select **Route Redistribution** from the drop-down menu.
- 5 Click **Edit** to enable or disable route redistribution.

6 Click **Add** to add a set of route redistribution criteria.

Option	Description
Name and Description	Assign a name to the route redistribution. You can optionally provide a description. An example name, advertise-to-bgp-neighbor.
Sources	Select one or more of the following sources: ■ TO Connected ■ TO Uplink ■ TO Downlink ■ TO CSP ■ TO Loopback ■ TO Static ■ TO NAT ■ TO DNS Forwarder IP ■ TO IPSec Local IP ■ T1 Connected ■ T1 CSP ■ T1 Downlink ■ T1 Static ■ T1 LB SNAT ■ T1 NAT ■ T1 LB VIP ■ T1 DNS Forwarder IP
Route Map	(Optional) Assign a route map to filter a sequence of IP addresses from route redistribution.

Verify North-South Connectivity and Route Redistribution on a Tier-O Router

Use the CLI to verify that the BGP routes are learned. You can also check from the external router that the NSX-T Data Center-connected VMs are reachable.

Prerequisites

- Verify that BGP is configured. See [Configure BGP on a Tier-O Logical Router in Manager Mode](#).
- Verify that NSX-T Data Center static routes are set to be redistributed. See [Enable Route Redistribution on the Tier-O Logical Router in Manager Mode](#).

Procedure

- 1 Log in to the NSX Manager CLI.
- 2 View the routes learned from the external BGP neighbor.

```
nsx-edge1(tier0_sr)> get route bgp

Flags: c - connected, s - static, b - BGP, ns - nsx_static
```

```
nc - nsx_connected, rl - router_link, t0n: Tier0-NAT, t1n: Tier1-NAT
```

```
b    10.10.10.0/24      [20/0]      via 192.168.100.254
```

- 3 From the external router, check that BGP routes are learned and that the VMs are reachable through the NSX-T Data Center overlay.

- a List the BGP routes.

```
user@router# run show ip route bgp
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       I - ISIS, B - BGP, > - selected route, * - FIB route

B>* 172.16.10.0/24 [20/0] via 192.168.100.2, eth2, 00:00:48
B>* 172.16.20.0/24 [20/0] via 192.168.100.2, eth2, 00:00:48
B>* 172.16.30.0/24 [20/0] via 192.168.100.2, eth2, 00:00:48
```

- b From the external router, ping the NSX-T Data Center-connected VMs.

```
ping 172.16.10.10
```

```
PING 172.16.10.10 (172.16.10.10) 56(84) bytes of data.
64 bytes from 172.16.10.10: icmp_req=1 ttl=62 time=127 ms
64 bytes from 172.16.10.10: icmp_req=2 ttl=62 time=1.96 ms
^C
--- 172.16.10.10 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1001ms
rtt min/avg/max/mdev = 1.966/64.793/127.620/62.827 ms
```

- c Check the path through the NSX-T Data Center overlay.

```
traceroute 172.16.10.10
```

```
traceroute to 172.16.10.10 (172.16.10.10), 30 hops max, 60 byte packets
 1  192.168.100.3 (192.168.100.3)  0.640 ms  0.575 ms  0.696 ms
 2  100.91.176.1 (100.91.176.1)  0.656 ms  0.604 ms  0.578 ms
 3  172.16.10.10 (172.16.10.10)  3.397 ms  3.703 ms  3.790 ms
```

- 4 From the internal VMs, ping the external IP address.

```
ping 10.10.10.10
```

```
PING 10.10.10.10 (10.10.10.10) 56(84) bytes of data.
64 bytes from 10.10.10.10: icmp_req=1 ttl=62 time=119 ms
64 bytes from 10.10.10.10: icmp_req=2 ttl=62 time=1.93 ms
```

```
^C
--- 10.10.10.10 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1001ms
rtt min/avg/max/mdev = 1.936/60.865/119.795/58.930 ms
```

What to do next


Configure additional routing functionality, such as ECMP.

Understanding ECMP Routing

Equal cost multi-path (ECMP) routing protocol increases the north and south communication bandwidth by adding an uplink to the tier-0 logical router and configure it for each Edge node in an NSX Edge cluster. The ECMP routing paths are used to load balance traffic and provide fault tolerance for failed paths.

The tier-0 logical router must be in active-active mode for ECMP to be available. A maximum of eight ECMP paths are supported. The implementation of ECMP on NSX Edge is based on the 5-tuple of the protocol number, source address, destination address, source port, and destination port. The algorithm used to distribute the data among the ECMP paths is not round robin. Therefore, some paths might carry more traffic than others. Note that if the protocol is IPv6 and the IPv6 header has more than one extension header, ECMP will be based only on the source and destination addresses.

Figure 20-8. ECMP Routing Topology



For example, the topology above shows a single tier-0 logical router in active-active mode running on a 2-node NSX Edge cluster. Two uplink ports are configured, one on each Edge node.

Add an Uplink Port for the Second Edge Node for ECMP in Manager Mode

Before you enable ECMP, you must configure an uplink to connect the tier-0 logical router to the VLAN logical switch.

Prerequisites

- Verify that a transport zone and two transport nodes are configured. See the *NSX-T Data Center Installation Guide*.
- Verify that two Edge nodes and an Edge cluster are configured. See the *NSX-T Data Center Installation Guide*.

- Verify that a VLAN logical switch for uplink is available. See [Create a VLAN Logical Switch for the NSX Edge Uplink in Manager Mode](#).
- Verify that a tier-0 logical router is configured. See [Create a Tier-0 Logical Router in Manager Mode](#).
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Tier-0 Logical Routers**.
- 3 Select the tier-0 logical router.
- 4 Click the **Configuration** tab to add a router port.
- 5 Click **Add**.
- 6 Complete the router port details.

Option	Description
Name	Assign a name for the router port.
Description	Provide additional description that shows that the port is for ECMP configuration.
Type	Accept the default type Uplink .
MTU	If you leave this field empty, the default is 1500.
Transport Node	Assign the Edge transport node from the drop-down menu.
URPF Mode	uRPF (unicast Reverse Path Forwarding) is enabled by default on external, internal and service interfaces. From a security standpoint, it is a best practice to keep uRPF enabled on these interfaces. uRPF is also recommended in architectures that leverage ECMP. It is possible to disable uRPF in complex routing architectures where asymmetric routing exists.
Logical Switch	Assign the VLAN logical switch from the drop-down menu.
Logical Switch Port	Assign a new switch port name. You can also use an existing switch port.
IP Address/Mask	Enter an IP address that is in the same subnet as the connected port on the ToR switch.

- 7 Click **Save**.

Results

A new uplink port is added to the tier-0 router and the VLAN logical switch. The tier-0 logical router is configured on both of the edge nodes.

What to do next

Create a BGP connection for the second neighbor and enable the ECMP routing. See [Add a Second BGP Neighbor and Enable ECMP Routing in Manager Mode](#).

Add a Second BGP Neighbor and Enable ECMP Routing in Manager Mode

Before you enable ECMP routing, you must add a BGP neighbor and configure it with the newly added uplink information.

Prerequisites

- Verify that the second edge node has an uplink port configured. See [Add an Uplink Port for the Second Edge Node for ECMP in Manager Mode](#).
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Tier-0 Logical Routers**.
- 3 Select the tier-0 logical router.
- 4 Click the **Routing** tab and select **BGP** from the drop-down menu.
- 5 Click **Add** under the Neighbors section to add a BGP neighbor.
- 6 Enter the neighbor IP address.
For example, 192.168.200.254.
- 7 (Optional) Specify the maximum hop limit.
The default is 1.
- 8 Enter the remote AS number.
For example, 64511.
- 9 (Optional) Click the **Local Address** tab to select a local address.
 - a (Optional) Uncheck **All Uplinks** to see loopback ports as well as uplink ports.
- 10 (Optional) Click the **Address Families** tab to add an address family.
- 11 (Optional) Click the **BFD Configuration** tab to enable BFD.
- 12 Click **Save**.
The newly added BGP neighbor appears.
- 13 Click **Edit** next to the BGP Configuration section.

- 14 Click the **ECMP** toggle button to enable ECMP.

The Status button must be appear as Enabled.

- 15 Click **Save**.

Results

Multiple ECMP routing paths connect the VMs attached to logical switches and the two Edge nodes in the Edge cluster.

What to do next

Test whether the ECMP routing connections are working properly. See [Verify ECMP Routing Connectivity on a Tier-0 Router](#).

Verify ECMP Routing Connectivity on a Tier-0 Router

Use CLI to verify that the ECMP routing connection to neighbor is established.

Prerequisites

Verify that ECMP routing is configured. See [Add an Uplink Port for the Second Edge Node for ECMP in Manager Mode](#) and [Add a Second BGP Neighbor and Enable ECMP Routing in Manager Mode](#).

Procedure

- 1 Log in to the NSX Manager CLI.
- 2 Get the distributed router UUID information.

```
get logical-routers
```

```
Logical Router
UUID          : 736a80e3-23f6-5a2d-81d6-bbefb2786666
vrf           : 2
type          : TUNNEL

Logical Router
UUID          : d40bbfa4-3e3d-4178-8615-6f42ea335037
vrf           : 4
type          : SERVICE_ROUTER_TIER0

Logical Router
UUID          : d0289ba4-250e-41b4-8ffc-7cab4a46c3e4
vrf           : 5
type          : DISTRIBUTED_ROUTER

Logical Router
UUID          : a6ee6316-2212-4171-99cc-930c98bcad7f
vrf           : 6
type          : DISTRIBUTED_ROUTER
```

- 3 Locate the UUID information from the output.

```
Logical Router
UUID       : d0289ba4-250e-41b4-8ffc-7cab4a46c3e4
vrf        : 5
type       : DISTRIBUTED_ROUTER
```

- 4 Type the VRF for the tier-0 distributed router.

```
vrf 5
```

- 5 Verify that the tier-0 distributed router is connected to the Edge nodes.

```
get forwarding
```

For example, edge-node-1 and edge-node-2.

- 6 Enter **exit** to leave the vrf context.
- 7 Verify that the tier-0 distributed router is connected.

```
get logical-router <UUID> route
```

The route type for the UUID should appear as `NSX_CONNECTED`.

- 8 Start a SSH session on the two Edge nodes.
 - 9 Start a session to capture packets.
- ```
set capture session 0 interface fp-eth1 dir tx
set capture session 0 expression src net <IP_Address>
```
- 10 Use any tool that can generate traffic from a source VM connected to the tier-0 router to a destination VM.
  - 11 Observe the traffic on the two Edge nodes.

## Create an IP Prefix List in Manager Mode

An IP prefix list contains single or multiple IP addresses that are assigned access permissions for route advertisement. The IP addresses in this list are processed sequentially. IP prefix lists are referenced through BGP neighbor filters or route maps with in or out direction.

For example, you can add the IP address 192.168.100.3/27 to the IP prefix list and deny the route from being redistributed to the northbound router. You can also append an IP address with less-than-or-equal-to (le) and greater-than-or-equal-to (ge) modifiers to grant or limit route redistribution. For example, 192.168.100.3/27 ge 24 le 30 modifiers match subnet masks greater than or equal to 24-bits and less than or equal to 30-bits in length.

---

**Note** The default action for a route is **Deny**. When you create a prefix list to deny or permit specific routes, be sure to create an IP prefix with no specific network address (select **Any** from the dropdown list) and the **Permit** action if you want to permit all other routes.

---

### Prerequisites

- Verify that you have a tier-0 logical router configured. See [Create a Tier-0 Logical Router in Manager Mode](#).
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Tier-0 Logical Routers**.
- 3 Select the tier-0 logical router.
- 4 Click the **Routing** tab and select **IP Prefix Lists** from the drop-down menu.
- 5 Click **Add**.
- 6 Enter a name for the IP prefix list.
- 7 Click **Add** to specify a prefix.
  - a Enter an IP address in CIDR format.  
For example, 192.168.100.3/27.
  - b Select **Deny** or **Permit** from the drop-down menu.
  - c (Optional) Set a range of IP address numbers in the **le** or **ge** modifiers.  
For example, set **le** to 30 and **ge** to 24.
- 8 Repeat the previous step to specify additional prefixes.
- 9 Click **Add** at the bottom of the window.

## Create a Community List in Manager Mode

You can create BGP community lists so that you can configure route maps based on community lists.

### Prerequisites

- Verify that you have a tier-0 logical router configured. See [Create a Tier-0 Logical Router in Manager Mode](#).
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

**Procedure**

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Tier-0 Logical Routers**.
- 3 Select the tier-0 logical router.
- 4 Click the **Routing** tab and select **Community Lists** from the drop-down menu.
- 5 Click **Add**.
- 6 Enter a name for the community list.
- 7 Specify a community using the aa:nn format, for example, 300:500, and press Enter. Repeat to add additional communities.

In addition, you can click the dropdown arrow and select one or more of the following:

- NO\_EXPORT\_SUBCONFED - Do not advertise to EBGp peers.
- NO\_ADVERTISE - Do not advertise to any peer.
- NO\_EXPORT - Do not advertise outside BGP confederation

- 8 Click **Add**.

**Create a Route Map in Manager Mode**

A route map consists of a sequence of IP prefix lists, BGP path attributes, and an associated action. The router scans the sequence for an IP address match. If there is a match, the router performs the action and scans no further.

Route maps can be referenced at the BGP neighbor level and route redistribution. When IP prefix lists are referenced in route maps and the route map action of permitting or denying is applied, the action specified in the route map sequence overrides the specification within the IP prefix list.

**Prerequisites**

- Verify that an IP prefix list is configured. See [Create an IP Prefix List in Manager Mode](#).
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

**Procedure**

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Tier-0 Logical Routers**.
- 3 Select the tier-0 logical router.
- 4 Select **Routing > Route Maps**.

- 5 Click **Add**.
- 6 Enter a name and an optional description for the route map.
- 7 Click **Add** to add an entry in the route map.
- 8 Edit the column **Match IP Prefix List/Community List** to select either IP prefix lists, or community lists, but not both.
- 9 (Optional) Set BGP attributes.

| BGP Attribute   | Description                                                                                                                                                                                                                                                                                                                                                     |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| AS-path Prepend | Prepend a path with one or more AS (autonomous system) numbers to make the path longer and therefore less preferred.                                                                                                                                                                                                                                            |
| MED             | Multi-Exit Discriminator indicates to an external peer a preferred path to an AS.                                                                                                                                                                                                                                                                               |
| Weight          | Set a weight to influence path selection. The range is 0 - 65535.                                                                                                                                                                                                                                                                                               |
| Community       | Specify a community using the aa:nn format, for example, 300:500. Or use the drop-down menu to select one of the following: <ul style="list-style-type: none"> <li>■ NO_EXPORT_SUBCONFED - Do not advertise to EBGp peers.</li> <li>■ NO_ADVERTISE - Do not advertise to any peer.</li> <li>■ NO_EXPORT - Do not advertise outside BGP confederation</li> </ul> |

- 10 In the Action column, select **Permit** or **Deny**.

You can permit or deny IP addresses in the IP prefix lists from advertising their addresses.

- 11 Click **Save**.

## Configure Forwarding Up Timer in Manager Mode

You can configure forwarding up timer for a tier-0 logical router.

Forwarding up timer defines the time in seconds that the router must wait before sending the up notification after the first BGP session is established. This timer (previously known as forwarding delay) minimizes downtime in case of fail-overs for active-active or active-standby configurations of logical routers on NSX Edge that use dynamic routing (BGP). It should be set to the number of seconds an external router (TOR) takes to advertise all the routes to this router after the first BGP/BFD session. The timer value should be directly proportional to the number of northbound dynamic routes that the router must learn. This timer should be set to 0 on single edge node setups.

### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Tier-0 Logical Routers**.
- 3 Select the tier-0 logical router.
- 4 Select **Routing > Global Configuration**.
- 5 Click **Edit**.
- 6 Enter a value for the forwarding up timer.
- 7 Click **Save**.

## NAT in Manager Mode

You can configure Network Address Translation (NAT) in **Manager** mode.

---

**Note** If you use **Manager** mode to modify objects created in the **Policy** mode, some settings might not be configurable. These read-only settings have this icon next to them: . See [Chapter 1 NSX Manager](#) for more information.

---

## Network Address Translation


Network address translation (NAT) in NSX-T Data Center can be configured on tier-0 and tier-1 logical routers.

For example, the following diagram shows two tier-1 logical routers with NAT configured on Tenant2NAT. The web VM is simply configured to use 172.16.10.10 as its IP address and 172.16.10.1 as its default gateway.


NAT is enforced at the uplink of the Tenant2NAT logical router on its connection to the tier-0 logical router.

To enable NAT configuration, Tenant2NAT must have a service component on an NSX Edge cluster. Thus, Tenant2NAT is shown inside the NSX Edge. For comparison, Tenant1 can be outside of the NSX Edge because it is not using any Edge services.

Figure 20-9. NAT Topology



Note: In the following scenario, NAT hairpinning is not supported. The tier-0 logical router has DNAT and SNAT configured. Tier-1 Logical Router 2 has NO\_SNAT and SNAT configured. VM2 will not be able to access VM1 using VM1's external address 80.80.80.10.



The following sections describe how to create NAT rules using the manager UI. You can also make an API call (`POST /api/v1/logical-routers/<logical-router-id>/nat/rules?action=create_multiple`) to create multiple NAT rules at the same time. For more information, see the *NSX-T Data Center API Guide*.

## Tier-1 NAT

A tier-1 logical router supports source NAT (SNAT), destination NAT (DNAT) and reflexive NAT.

### Configure Source NAT on a Tier-1 Router in Manager Mode

Source NAT (SNAT) changes the source address in the IP header of a packet. It can also change the source port in the TCP/UDP headers. The typical usage is to change a private (rfc1918) address/port into a public address/port for packets leaving your network.

You can create a rule to either enable or disable source NAT.

In this example, as packets are received from the web VM, the Tenant2NAT tier-1 router changes the source IP address of the packets from 172.16.10.10 to 80.80.80.1. Having a public source IP address enables destinations outside of the private network to route back to the original source.

### Prerequisites

- The tier-0 router must have an uplink connected to a VLAN-based logical switch. See [Connect a Tier-0 Logical Router to a VLAN Logical Switch for the NSX Edge Uplink in Manager Mode](#).



- The tier-0 router must have routing (static or BGP) and route redistribution configured on its uplink to the physical architecture. See [Configure a Static Route in Manager Mode](#), [Configure BGP on a Tier-0 Logical Router in Manager Mode](#), and [Enable Route Redistribution on the Tier-0 Logical Router in Manager Mode](#).
- The tier-1 routers must each have an uplink to a tier-0 router configured. Tenant2NAT must be backed by an NSX Edge cluster. See [Attach Tier-1 Router to a Tier-0 Router in Manager Mode](#).
- The tier-1 routers must have downlink ports and route advertisement configured. See [Add a Downlink Port on a Tier-1 Logical Router in Manager Mode](#) and [Configure Route Advertisement on a Tier-1 Logical Router in Manager Mode](#).
- The VMs must be attached to the correct logical switches.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

#### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Tier-0 Logical Routers**.
- 3 Click a tier-1 logical router on which you want to configure NAT.
- 4 Select **Services > NAT**.
- 5 Click **ADD**.
- 6 Specify a priority value.  
A lower value means a higher precedence for this rule.
- 7 For **Action**, select **SNAT** to enable source NAT, or **NO\_SNAT** to disable source NAT.
- 8 Select the protocol type.  
By default, **Any Protocol** is selected.
- 9 (Optional) For **Source IP**, specify an IP address or an IP address range in CIDR format.  
If you leave this field blank, all sources on router's downlink ports are translated. In this example, the source IP address is 172.16.10.10.
- 10 (Optional) For **Destination IP**, specify an IP address or an IP address range in CIDR format.  
If you leave this field blank, the NAT applies to all destinations outside of the local subnet.
- 11 If **Action** is **SNAT**, for **Translated IP**, specify an IP address or an IP address range in CIDR format.  
In this example, the translated IP address is 80.80.80.1.
- 12 (Optional) For **Applied To**, select a router port.

**13** (Optional) Set the status of the rule.

The rule is enabled by default.

**14** (Optional) Change the logging status.

Logging is disabled by default.

**15** (Optional) Change the firewall bypass setting.

The setting is enabled by default.

**Results**

The new rule is listed under NAT. For example:

Tenant2NAT

Summary

Configuration

Routing

NAT

NAT

No Statistics were collected

ADD

EDIT

DELETE

COLUMNS

| ID                         | Action | Match    |              |              |                |                   | Translated |       | Stats                  |
|----------------------------|--------|----------|--------------|--------------|----------------|-------------------|------------|-------|------------------------|
|                            |        | Protocol | Source IP    | Source Ports | Destination IP | Destination Ports | IP         | Ports |                        |
| Priority: 1024             |        |          |              |              |                |                   |            |       |                        |
| <div><div></div>4100</div> | SNAT   | Any      | 172.16.10.10 | Any          | Any            | Any               | 80.80.80.1 | Any   | <div><div></div></div> |

**What to do next**

Configure the tier-1 router to advertise NAT routes.

To advertise the NAT routes upstream from the tier-0 router to the physical architecture, configure the tier-0 router to advertise tier-1 NAT routes.

**Configure Destination NAT on a Tier-1 Router in Manager Mode**

Destination NAT changes the destination address in IP header of a packet. It can also change the destination port in the TCP/UDP headers. The typical usage of this is to redirect incoming packets with a destination of a public address/port to a private IP address/port inside your network.

You can create a rule to either enable or disable destination NAT.

In this example, as packets are received from the app VM, the Tenant2NAT tier-1 router changes the destination IP address of the packets from 172.16.10.10 to 80.80.80.1. Having a public destination IP address enables a destination inside a private network to be contacted from outside of the private network.

**Prerequisites**

- The tier-0 router must have an uplink connected to a VLAN-based logical switch. See [Connect a Tier-0 Logical Router to a VLAN Logical Switch for the NSX Edge Uplink in Manager Mode](#).

- The tier-0 router must have routing (static or BGP) and route redistribution configured on its uplink to the physical architecture. See [Configure a Static Route in Manager Mode](#), [Configure BGP on a Tier-0 Logical Router in Manager Mode](#), and [Enable Route Redistribution on the Tier-0 Logical Router in Manager Mode](#).
- The tier-1 routers must each have an uplink to a tier-0 router configured. Tenant2NAT must be backed by an NSX Edge cluster. See [Attach Tier-1 Router to a Tier-0 Router in Manager Mode](#).
- The tier-1 routers must have downlink ports and route advertisement configured. See [Add a Downlink Port on a Tier-1 Logical Router in Manager Mode](#) and [Configure Route Advertisement on a Tier-1 Logical Router in Manager Mode](#).
- The VMs must be attached to the correct logical switches.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

#### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Tier-1 Logical Routers**.
- 3 Click a tier-1 logical router on which you want to configure NAT.
- 4 Select **Services > NAT**.
- 5 Click **ADD**.
- 6 Specify a priority value.  
A lower value means a higher precedence for this rule.
- 7 For **Action**, select **DNAT** to enable destination NAT, or **NO\_DNAT** to disable destination NAT.
- 8 Select the protocol type.  
By default, **Any Protocol** is selected.
- 9 (Optional) For **Source IP**, specify an IP address or an IP address range in CIDR format.  
If you leave Source IP blank, the NAT applies to all sources outside of the local subnet.
- 10 For **Destination IP**, specify an IP address or a comma-separated IP address list.  
In this example, the destination IP address is 80.80.80.1.
- 11 If **Action** is **DNAT**, for **Translated IP**, specify an IP address or an IP address range in CIDR format.  
In this example, the inside/translated IP address is 172.16.10.10.
- 12 (Optional) If **Action** is **DNAT**, for **Translated Ports**, specify the translated ports.

13 (Optional) For **Applied To**, select a router port.

14 (Optional) Set the status of the rule.

The rule is enabled by default.

15 (Optional) Change the logging status.

Logging is disabled by default.

16 (Optional) Change the firewall bypass setting.

The setting is enabled by default.

## Results

The new rule is listed under NAT. For example:

Tenant2NAT

Summary

Configuration

Routing

NAT

NAT

No Statistics were collected

+

ADDEDITDELETECOLUMNS

| ID               | Action | Match    |           |              |                |                   | Translated   |       | Stats |
|------------------|--------|----------|-----------|--------------|----------------|-------------------|--------------|-------|-------|
|                  |        | Protocol | Source IP | Source Ports | Destination IP | Destination Ports | IP           | Ports |       |
| ▲ Priority: 1024 |        |          |           |              |                |                   |              |       |       |
| 4101             | DNAT   | Any      | Any       | Any          | 80.80.80.1     | Any               | 172.16.10.10 | Any   |       |

## What to do next

Configure the tier-1 router to advertise NAT routes.

To advertise the NAT routes upstream from the tier-0 router to the physical architecture, configure the tier-0 router to advertise tier-1 NAT routes.

## Advertise Tier-1 NAT Routes to the Upstream Tier-0 Router in Manager Mode

Advertising tier-1 NAT routes enables the upstream tier-0 router to learn about these routes.

### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Tier-1 Logical Routers**.
- 3 Click a tier-1 logical router on which you have configured NAT.

- 4 From the tier-1 router, select **Routing > Route Advertisement**.
- 5 Click **Edit** to edit the route advertisement configuration.

You can toggle the following switches:

- **Status**
- **Advertise All NSX Connected Routes**
- **Advertise All NAT Routes**
- **Advertise All Static Routes**
- **Advertise All LB VIP Routes**
- **Advertise All LB SNAT IP Routes**
- **Advertise All DNS Forwarder Routes**

- 6 Click **Save**.

#### What to do next

Advertise tier-1 NAT routes from the tier-0 router to the upstream physical architecture.

#### Advertise Tier-1 NAT Routes to the Physical Architecture in Manager Mode

Advertising tier-1 NAT routes from the tier-0 router enables the upstream physical architecture to learn about these routes.

#### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

#### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Tier-0 Logical Routers**.
- 3 Click a tier-0 logical router that is connected to a tier-1 router on which you have configured NAT.
- 4 From the tier-0 router, select **Routing > Route Redistribution**.
- 5 Click **Edit** to enable or disable route redistribution.

## 6 Click **Add** to add a set of route redistribution criteria.

| Option                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Name and Description</b> | Assign a name to the route redistribution. You can optionally provide a description.<br>An example name, advertise-to-bgp-neighbor.                                                                                                                                                                                                                                                                                                                                                     |
| <b>Sources</b>              | Select one or more of the following sources: <ul style="list-style-type: none"> <li>■ TO Connected</li> <li>■ TO Uplink</li> <li>■ TO Downlink</li> <li>■ TO CSP</li> <li>■ TO Loopback</li> <li>■ TO Static</li> <li>■ TO NAT</li> <li>■ TO DNS Forwarder IP</li> <li>■ TO IPSec Local IP</li> <li>■ T1 Connected</li> <li>■ T1 CSP</li> <li>■ T1 Downlink</li> <li>■ T1 Static</li> <li>■ T1 LB SNAT</li> <li>■ T1 NAT</li> <li>■ T1 LB VIP</li> <li>■ T1 DNS Forwarder IP</li> </ul> |
| <b>Route Map</b>            | (Optional) Assign a route map to filter a sequence of IP addresses from route redistribution.                                                                                                                                                                                                                                                                                                                                                                                           |

### Verify Tier-1 NAT

Verify that SNAT and DNAT rules are working correctly.

#### Procedure

- 1 Log in the NSX Edge.
- 2 Run `get logical-routers` to determine the VRF number for the tier-0 services router.
- 3 Enter the tier-0 services router context by running the `vrf <number>` command.
- 4 Run the `get route` command and make sure that the tier-1 NAT address appears.

```
nsx-edge(tier0_sr)> get route
```

```
Flags: c - connected, s - static, b - BGP, ns - nsx_static
nc - nsx_connected, rl - router_link, t0n: Tier0-NAT, t1n: Tier1-NAT
```

```
Total number of routes: 8
t1n 80.80.80.1/32 [3/3] via 169.0.0.1
...
```

- 5 If your Web VM is set up to serve Web pages, make sure you can open a Web page at `http://80.80.80.1`.
- 6 Make sure that the tier-0 router's upstream neighbor in the physical architecture can ping 80.80.80.1.
- 7 While the ping is still running, check the stats column for the DNAT rule.  
There should be one active session.

## Tier-0 NAT

A tier-0 logical router in active-standby mode supports source NAT (SNAT), destination NAT (DNAT) and reflexive NAT. A tier-0 logical router in active-active mode supports reflexive NAT only.

### Configure Source and Destination NAT on a Tier-0 Logical Router in Manager Mode

You can configure source and destination NAT on a tier-0 logical router that is running in active-standby mode.

You can also disable SNAT or DNAT for an IP address or a range of addresses. If multiple NAT rules apply to an address, the rule with the highest priority is applied.

SNAT configured on a tier-0 logical router's uplink will process traffic from a tier-1 logical router as well as from another uplink on the tier-0 logical router.

#### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

#### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Tier-0 Logical Routers**.
- 3 Click a tier-0 logical router.
- 4 Select **Services > NAT**.
- 5 Click **ADD** to add a NAT rule.
- 6 Specify a priority value.  
A lower value means a higher priority.
- 7 For **Action**, select **SNAT**, **DNAT**, **Reflexive**, **NO\_SNAT**, or **NO\_DNAT**.

**8** Select the protocol type.

By default, **Any Protocol** is selected.

**9** (Required) For **Source IP**, specify an IP address or an IP address range in CIDR format.

If you leave this field blank, this NAT rule applies to all sources outside of the local subnet.

**10** For **Destination IP**, specify an IP address or an IP address range in CIDR format.**11** For **Translated IP**, specify an IP address or an IP address range in CIDR format.**12** (Optional) If **Action** is **DNAT**, for **Translated Ports**, specify the translated ports.**13** (Optional) For **Applied To**, select a router port.**14** (Optional) Set the status of the rule.

The rule is enabled by default.

**15** (Optional) Change the logging status.

Logging is disabled by default.

**16** (Optional) Change the firewall bypass setting.


The setting is enabled by default.

## Reflexive NAT


When a tier-0 logical router is running in active-active mode, you cannot configure stateful NAT where asymmetrical paths might cause issues. For active-active routers, you can configure reflexive NAT (sometimes called stateless NAT).

In this example, as packets are received from the web VM, the Tenant2NAT tier-1 router changes the source IP address of the packets from 172.16.10.10 to 80.80.80.1. Having a public source IP address enables destinations outside of the private network to route back to the original source.





When there are two active-active tier-0 routers involved, as shown below, reflexive NAT must be configured.



### Configure Reflexive NAT on a Tier-0 or Tier-1 Logical Router in Manager Mode

When a tier-0 or tier-1 logical router is running in active-active mode, you cannot configure stateful NAT where asymmetrical paths might cause issues. For active-active routers, you can use reflexive NAT, which is sometimes called stateless NAT.

For reflexive NAT, you can configure a single source address to be translated, or a range of addresses. If you configure a range of source addresses, you must also configure a range of translated addresses. The size of the two ranges must be the same. The address translation will be deterministic, meaning that the first address in the source address range will be translated to the first address in the translated address range, the second address in the source range will be translated to the second address in the translated range, and so on.

### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Locate the logical router you want to modify in **Networking > Tier-0 Logical Routers** or **Networking > Tier-1 Logical Routers**.
- 3 Click the tier-0 or tier-1 logical router on which you want to configure reflexive NAT.
- 4 Select **Services > NAT**.
- 5 Click **ADD**.
- 6 Specify a priority value.  
A lower value means a higher precedence for this rule.
- 7 For **Action**, select **Reflexive**.
- 8 For **Source IP**, specify an IP address or an IP address range in CIDR format.
- 9 For **Translated IP**, specify an IP address or an IP address range in CIDR format.
- 10 (Optional) Set the status of the rule.  
The rule is enabled by default.
- 11 (Optional) Change the logging status.  
Logging is disabled by default.
- 12 (Optional) Change the firewall bypass setting.  
The setting is enabled by default.

### Results

The new rule is listed under NAT. For example:

Tier0-LR-1

Overview Configuration Routing Services

NAT | REFRESH

Total Rule Statistics | Last Updated: 11/6/2018, 12:40:13 PM

0 Active sessions0 Packet count0 Bytes Data

+ ADD EDIT DELETE

| ID               | Action    | Match    |            |              |                |                   | Translated   |       | Applied To | Stats |
|------------------|-----------|----------|------------|--------------|----------------|-------------------|--------------|-------|------------|-------|
|                  |           | Protocol | Source IP  | Source Ports | Destination IP | Destination Ports | IP           | Ports |            |       |
| ▼ Priority: 1024 |           |          |            |              |                |                   |              |       |            |       |
| 1034             | Reflexive | Any      | 80.80.80.1 | Any          | Any            | Any               | 172.16.10.10 | Any   |            |       |

## Grouping Objects in Manager Mode

You can create IP sets, IP pools, MAC sets, NSGroups, and NSServices in **Manager** mode. You can also manage tags for VMs.

**Note** If you use **Manager** mode to modify objects created in the **Policy** mode, some settings might not be configurable. These read-only settings have this icon next to them: . See [Chapter 1 NSX Manager](#) for more information.

### Create an IP Set in Manager Mode

An IP set is a group of IP addresses that can be used as sources and destinations in firewall rules.

An IP set can contain a combination of individual IP addresses, IP ranges, and subnets. You can specify IPv4 or IPv6 addresses, or both. An IP set can be a member of NSGroups.

**Note** Any IP set created by this method will not be visible in Policy mode. In Policy mode, we can create a group and add members as IP addresses, ranges, network addresses, or MAC addresses by navigating to **Inventory > Groups > Set Members** and specifying IP or MAC addresses.

**Note** IPv4 addresses and IPv6 addresses are supported for source or destination ranges for firewall rules.

#### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

#### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Inventory > Groups > IP Sets > Add**.

- 3 Enter a name.
- 4 (Optional) Enter a description.
- 5 In **Members**, enter individual IP addresses, IP ranges, and subnets in a comma separated list.
- 6 Click **Save**.

## Create an IP Pool in Manager Mode

You can use an IP Pool to allocate IP addresses or subnets when you create L3 subnets.

### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > IP Management > IP Address Pools**.
- 3 Enter a name for the new IP pool.
- 4 (Optional) Enter a description.
- 5 Click **Add**.
- 6 Click the IP Ranges cell and enter IP Ranges.  
Mouse over the upper right corner of any cell and click the pencil icon to edit it.
- 7 (Optional) Enter a Gateway.
- 8 Enter a CIDR IP address with suffix.
- 9 (Optional) Enter DNS Servers.
- 10 (Optional) Enter a DNS Suffix.
- 11 Click **Save**.

## Create a MAC Set in Manager Mode

A MAC Set is a group of MAC addresses that you can use as sources and destinations in layer 2 firewall rules and as a member of an NS Group.

### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

**Procedure**

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Inventory > Groups > MAC Sets > Add**.
- 3 Enter a name.
- 4 (Optional) Enter a description.
- 5 Enter the MAC addresses in a comma-separated list.
- 6 Click **ADD**.

**Create an NSGroup in Manager Mode**

NSGroups can be configured to contain a combination of IP sets, MAC sets, logical ports, logical switches, and other NSGroups. You can specify NSGroups with Logical Switches, Logical ports and VMs as sources and destinations, and in the `Applied To` field of a firewall rule. NSGroups with IPset and MACSet will be ignored in a distributed firewall `Applied To` field.

---

**NSX Cloud Note** If using NSX Cloud, see [NSX-T Data Center Features Supported with NSX Cloud](#) for a list of auto-generated logical entities, supported features, and configurations required for NSX Cloud.

---

An NSGroup has the following characteristics:

- An NSGroup has direct members and effective members. Effective members include members that you specify using membership criteria, as well as all the direct and effective members that belong to this NSGroup's members. For example, assuming NSGroup-1 has direct member LogicalSwitch-1. You add NSGroup-2 and specify NSGroup-1 and LogicalSwitch-2 as members. Now NSGroup-2 has direct members NSGroup-1 and LogicalSwitch-2, and an effective member, LogicalSwitch-1. Next, you add NSGroup-3 and specify NSGroup-2 as a member. NSGroup-3 now has direct member NSGroup-2 and effective members LogicalSwitch-1 and LogicalSwitch-2. From the main groups table, clicking on a group and selecting **Related > NSGroups** would show NSGroup-1, NSGroup-2, and NSGroup-3 because all three have LogicalSwitch-1 as a member, either directly or indirectly.
- An NSGroup can have a maximum of 500 direct members.
- The recommended limit for the number of effective members in an NSGroup is 5000. The NSX Manager check the NSGroups regarding the limit twice a day, at 7 AM and 7 PM. Exceeding this limit does not affect any functionality but might have a negative impact on performance.
  - When the number of effective members for an NSGroup exceeds 80% of 5000, the warning message `NSGroup xyz is about to exceed the maximum member limit. Total number in NSGroup is ...` appears in the log file. When the number exceeds 5000, the warning message `NSGroup xyz has reached the maximum numbers limit. Total number in NSGroup = ...` appears.

- When the number of translated VIFs/IPs/MACs in an NSGroup exceeds 5000, the warning message `Container xyz has reached the maximum IP/MAC/VIF translations limit. Current translations count in Container - IPs:..., MACs:..., VIFs:... appears in the log file.`
- The maximum supported number of VMs is 10,000.
- You can create a maximum of 10,000 NSGroups.

For all the objects that you can add to an NSGroup as members, you can navigate to the screen for any of the objects and select **Related > NSGroups**.

### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Inventory > Groups > Add**.
- 3 Enter a name for the NSGroup.
- 4 (Optional) Enter a description.
- 5 (Optional) Click **Membership Criteria**.

For each criterion, you can specify up to five rules, which are combined with the logical AND operator. The available member criterion can apply to the following:

- **Logical Port** - can specify a tag and optional scope.
  - **Logical Switch** - can specify a tag and optional scope.
  - **Virtual Machine** - can specify a name, tag, computer OS name, or computer name that equals, contains, starts with, ends with, or doesn't equal a particular string.
  - **Transport Node** - can specify a node type that equals an edge node or a host node.
  - **IP Set** - can specify a tag and optional scope.
- 6 (Optional) Click **Members** to select members.

The available member types are:

- **AD Group** - NSGroups with ADGroups can only be used in the `extended_source` field of a distributed firewall rule, and must be the only members in the group. For example, there cannot be an NSGroup with both ADGroup and IPSet together as members.
- **IP Set** - can include both IPv4 and IPv6 addresses.
- **Logical Port** - can include both IPv4 and IPv6 addresses.

- **Logical Switch** - can include both IPv4 and IPv6 addresses.
- **MAC Set**
- **NSGroup**
- **Transport Node**
- **VIF**
- **Virtual Machine**

7 Click **ADD**.

The group is added to the table of groups. Click a group name to display an overview and edit group information including membership criteria, members, applications, and related groups. Scroll to the bottom of the **Overview** tab to add and delete tags. See [Add Tags to an Object](#) for more information. Selecting **Related> NSGroups** displays all the NSGroups that have the selected NSGroup as a member.

## Configuring Services and Service Groups

You can configure an NSService and specify parameters for matching network traffic such as a port and protocol pairing. You can also use an NSService to allow or block certain types of traffic in firewall rules.

An NSService can be of the following types:

- Ether
- IP
- IGMP
- ICMP
- ALG
- L4 Port Set

An L4 Port Set supports the identification of source ports and destination ports. You can specify individual ports or a range of ports, up to a maximum of 15 ports.

An NSService can also be a group of other NSServices. An NSService that is a group can be of the following types:

- Layer 2
- Layer 3 and above

You cannot change the type after you create an NSService. Some NSServices are predefined. You cannot modify or delete them.

### Create an NSService in Manager Mode

You can create an NSService to specify the characteristics that network matching uses, or to define the type of traffic to block or allow in firewall rules.

### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Inventory > Services > Add**.
- 3 Enter a name.
- 4 (Optional) Enter a description.
- 5 Select **Specify a protocol** to configure an individual service, or select **Group existing services** to configure a group of NSServices.
- 6 For an individual service, select a type of service and a protocol.  
The available types are **Ether**, **IP**, **IGMP**, **ICMP**, **ALG**, and **L4 Port Set**.
- 7 For a service group, select a type and members for the group.  
The available types are **Layer 2** and **Layer 3 and above**.
- 8 Click **ADD**.

## Manage Tags for a VM in Manager Mode

You can see the list of VMs in the inventory. You can also add tags to a VM to make searching easier.

### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Inventory > Virtual Machines** from the navigation panel.

The list of VMs is displayed with 4 columns: Virtual Machine, External ID, Source, and Tag. Click the filter icon in the first three columns' heading to filter the list. Enter a string of characters to do a partial match. If the string in the column contains the string that you entered, the entry is displayed. Enter a string of characters enclosed in double quotes to do an exact match. If the string in the column exactly matches the string that you entered, the entry is displayed.



- 3 Select **Inventory** > **Virtual machines** from the navigation panel.
- 4 Select a VM.
- 5 Click **MANAGE TAGS**.
- 6 Add or delete tags.

| Option       | Action                                                    |
|--------------|-----------------------------------------------------------|
| Add a tag    | Click <b>ADD</b> to specify a tag and optionally a scope. |
| Delete a tag | Select an existing tag and click <b>DELETE</b> .          |

The maximum number of tags that can be assigned from the NSX Manager to a virtual machine is 25. The maximum number of tags for all other managed objects such as logical switches or ports, is 30 .

- 7 Click **Save**.

## DHCP in Manager Mode

You can configure DHCPv4 in **Manager** mode.

You cannot configure or modify a DHCPv6 server configuration in **Manager** mode. You must use any of the following to configure or modify a DHCPv6 server:

- Policy mode
- Policy API
- Manager API

**Note** If you use **Manager** mode to modify objects created in the **Policy** mode, some settings might not be configurable. These read-only settings have this icon next to them: . See [Chapter 1 NSX Manager](#) for more information.

## DHCP

DHCP (Dynamic Host Configuration Protocol) allows clients to automatically obtain network configuration, such as IP address, subnet mask, default gateway, and DNS configuration, from a DHCP server.

You can create DHCP servers to handle DHCP requests and create DHCP relay services to relay DHCP traffic to external DHCP servers. However, you should not configure a DHCP server on a logical switch and also configure a DHCP relay service on a router port that the same logical switch is connected to. In such a scenario, DHCP requests will only go to the DHCP relay service.

If you configure DHCP servers, to improve security, configure a DFW rule to allow traffic on UDP ports 67 and 68 only for valid DHCP server IP addresses.

To block DHCP packets for ports 67 and 68 configure a DFW rule with the following:

| Source | Destination | Service | Rule  |
|--------|-------------|---------|-------|
| ANY    | ANY         | ANY     | BLOCK |

To allow DHCP packets configure a DFW rule with the following:

| Source | Destination | Service              | Rule  |
|--------|-------------|----------------------|-------|
| ANY    | ANY         | ports 67 and 68, TCP | ALLOW |

**Note** In this release, the DHCP server does not support guest VLAN tagging.

## Create a DHCP Server Profile in Manager Mode

A DHCP server profile specifies an NSX Edge cluster or members of an NSX Edge cluster. A DHCP server with this profile services DHCP requests from VMs on logical switches that are connected to the NSX Edge nodes that are specified in the profile.

### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > DHCP > Server Profiles > Add**.
- 3 Enter a name and optional description.
- 4 Select an NSX Edge cluster from the drop-down menu.
- 5 (Optional) Select members of the NSX Edge cluster.

You can specify up to 2 members.

### What to do next

Create a DHCP server. See [Create a DHCP Server in Manager Mode](#).

## Create a DHCP Server in Manager Mode

You can create DHCP servers to service DHCP requests from VMs that are connected to logical switches.

### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

**Procedure**

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > DHCP > Servers > Add**.
- 3 Enter a name and optional description.
- 4 Enter the IP address of the DHCP server and its subnet mask in CIDR format.  
For example, enter `192.168.1.2/24`.
- 5 (Required) Select a DHCP profile from the drop-down menu.
- 6 (Optional) Enter common options such as domain name, default gateway, DNS servers, and subnet mask.
- 7 (Optional) Enter classless static route options.
- 8 (Optional) Enter other options.
- 9 Click **Save**.
- 10 Select the newly created DHCP server.
- 11 Expand the IP Pools section.
- 12 Click **Add** to add IP ranges, default gateway, lease duration, warning threshold, error threshold, classless static route option, and other options.
- 13 Expand the Static Bindings section.
- 14 Click **Add** to add static bindings between MAC addresses and IP addresses, default gateway, hostname, lease duration, classless static route option, and other options.

**What to do next**

Attach a DHCP server to a logical switch. See [Attach a DHCP Server to a Logical Switch in Manager Mode](#).

**Attach a DHCP Server to a Logical Switch in Manager Mode**

You must attach a DHCP server to a logical switch before the DHCP server can process DHCP requests from VMs connected to the switch.

**Prerequisites**

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

**Procedure**

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

- 2 Select **Networking > Logical Switches > Switches**.
  - a Select a logical switch.
  - b Click **Actions > Attach to a DHCP Server**.
- 3 Alternatively, select **Networking > DHCP > Servers**.
  - a Select a DHCP server.
  - b Click **Actions > Attach to Logical Switch**.

## Detach a DHCP Server from a Logical Switch in Manager Mode

You can detach a DHCP server from a logical switch to reconfigure your environment.

### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Logical Switches**.
- 3 Click the logical switch that you intend to detach a DHCP server from.
- 4 Click **Actions > Detach from the DHCP Server**.

## Create a DHCP Relay Profile in Manager Mode

A DHCP relay profile specifies one or more external DHCP or DHCPv6 servers. When you create a DHCP/DHCPv6 relay service, you must specify a DHCP relay profile.

### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > DHCP > Relay Profiles > Add**.
- 3 Enter a name and optional description.
- 4 Enter one or more external DHCP/DHCPv6 server addresses.

**What to do next**

Create a DHCP/DHCPv6 relay service. See [Create a DHCP Relay Service in Manager Mode](#).

**Create a DHCP Relay Service in Manager Mode**

You can create a DHCP relay service to relay traffic between DHCP clients and DHCP servers that are not created in NSX-T Data Center.

**Prerequisites**

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

**Procedure**

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > DHCP > Relay Services > Add**.
- 3 Enter a name and optional description.
- 4 Select a DHCP relay profile from the drop-down menu.

**What to do next**

Add a DHCP service to a logical router port. See [Add a DHCP Relay Service to a Logical Router Port in Manager Mode](#).

**Add a DHCP Relay Service to a Logical Router Port in Manager Mode**

You can add a DHCP relay service to a logical router port. VMs on the logical switch that is attached to that port can communicate with the DHCP servers that are configured in the relay service.

**Prerequisites**

- Verify you have a configured DHCP relay service. See [Create a DHCP Relay Service in Manager Mode](#).
- Verify that the router port is of type **Downlink**.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

**Procedure**

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Tier-0 Logical Routers**.

- 3 Select the appropriate router to display more information and configuration options.
- 4 Select **Configuration > Router Ports**.
- 5 Select the router port that connects to the desired logical switch and click **Edit**.
- 6 Select a DHCP relay service from the **Relay Service** drop-down list and click **Save**.

You can also select a DHCP relay service when you add a new logical router port.

## Delete a DHCP Lease

In some situations, you might want to delete a DHCP lease. For example, if you want a DHCP client to get a different IP address, or if a client shuts down without releasing its IP address and you want the address to be available to other clients.

You can use the following API to delete a DHCP lease:

```
DELETE /api/v1/dhcp/servers/<server-id>/leases?ip=<ip>&mac=<mac>
```

To ensure that the correct lease is deleted, call the following API before and after the `DELETE` API:

```
GET /api/v1/dhcp/servers/<server-id>/leases
```

After calling the `DELETE` API, make sure that the output of the `GET` API does not show the lease that was deleted.

For more information, see the *NSX-T Data Center API Reference*.

## Metadata Proxies

With a metadata proxy server, VM instances can retrieve instance-specific metadata from an OpenStack Nova API server.

The following steps describe how a metadata proxy works:

- 1 A VM sends an HTTP GET to `http://169.254.169.254:80` to request some metadata.
- 2 The metadata proxy server that is connected to the same logical switch as the VM reads the request, makes appropriate changes to the headers, and forwards the request to the Nova API server.
- 3 The Nova API server requests and receives information about the VM from the Neutron server.
- 4 The Nova API server finds the metadata and sends it to the metadata proxy server.
- 5 The metadata proxy server forwards the metadata to the VM.

A metadata proxy server runs on an NSX Edge node. For high availability, you can configure metadata proxy to run on two or more NSX Edge nodes in an NSX Edge cluster.

## Add a Metadata Proxy Server in Manager Mode

A metadata proxy server enables VMs to retrieve metadata from an OpenStack Nova API server.

### Prerequisites

- Verify that you have created an NSX Edge cluster. For more information, see *NSX-T Data Center Installation Guide*.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > DHCP > Metadata Proxies > Add** .
- 3 Enter a name for the metadata proxy server.
- 4 (Optional) Enter a description.
- 5 Enter the URL and port for the Nova server.  
The valid port range is 3000 - 9000.
- 6 Enter a value for **Secret**.
- 7 Select an NSX Edge cluster from the drop-down list.
- 8 (Optional) Select members of the NSX Edge cluster.

### What to do next

Attach the metadata proxy server to a logical switch.

## Attach a Metadata Proxy Server to a Logical Switch in Manager Mode

To provide metadata proxy services to VMs that are connected to a logical switch, you must attach a metadata proxy server to the switch.

### Prerequisites

- Verify that you have created a logical switch. For more information, see [Create a Logical Switch in Manager Mode](#).
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

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

- 2 Select **Networking > DHCP > Metadata Proxies**.
- 3 Select a metadata proxy server.
- 4 Select the menu option **Actions > Attach to Logical Switch**
- 5 Select a logical switch from the drop-down list.

#### Results

You can also attach a metadata proxy server to a logical switch by navigating to **Networking > Logical Switches > Switches**, selecting a switch, and selecting the menu option **Actions > Add to a Metadata Proxy**.

### Detach a Metadata Proxy Server from a Logical Switch in Manager Mode

To stop providing metadata proxy services to VMs that are connected to a logical switch or use a different metadata proxy server, you can detach a metadata proxy server from a logical switch.

#### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > DHCP > Metadata Proxies**.
- 3 Select a metadata proxy server.
- 4 Select the menu option **Actions > Detach from Logical Switch**
- 5 Select a logical switch from the drop-down list.

#### Results

You can also detach a metadata proxy server from a logical switch by navigating to **Networking > Logical Switches > Switches**, selecting a switch, and selecting the menu option **Actions > Detach from the Metadata Proxy**.

## IP Address Management in Manager Mode

With IP address management (IPAM), you can create IP blocks to support NSX Container Plug-in (NCP). For more info about NCP, see the *NSX-T Container Plug-in for Kubernetes - Installation and Administration Guide*.

---

**Note** If you use **Manager** mode to modify objects created in the **Policy** mode, some settings might not be configurable. These read-only settings have this icon next to them: . See [Chapter 1 NSX Manager](#) for more information.

---

### Manage IP Blocks in Manager Mode

Setting up NSX Container Plug-in requires that you create IP blocks for the containers.



### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > IP Address Pools > IP Blocks**.
- 3 To add an IP block, click **Add**.
  - a Enter a name and optionally a description.
  - b Enter an IP block in CIDR format. For example, 10.10.10.0/24.
- 4 To edit an IP block, click the name of an IP block.
  - a In the **Overview** tab, click **Edit**.

You can change the name, description, or the IP block value.
- 5 To manage the tags of an IP block, click the name of an IP block.
  - a In the **Overview** tab, click **Manage**.

You can add or delete tags.
- 6 To delete one or more IP blocks, select the blocks.
  - a Click **Delete**.

You cannot delete an IP block that has its subnet allocated.

## Manage Subnets for IP Blocks in Manager Mode

You can add or delete subnets for IP blocks.

### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > IP Address Pools > IP Blocks**.
- 3 Click the name of an IP block.
- 4 Click the **Subnets** tab.

- 5 To add a subnet, click **Add**.
  - a Enter a name and optionally a description.
  - b Enter the size of the subnet.
- 6 To delete one or more subnets, select the subnets.
  - a Click **Delete**.


## Load Balancing in Manager Mode

This information covers the NSX-T Data Center load balancing configuration in **Manager** mode.

For information about NSX Advanced Load Balancer (Avi Networks) see <https://www.vmware.com/products/nsx-advanced-load-balancer.html>.

**Note** If you use **Manager** mode to modify objects created in the **Policy** mode, some settings might not be configurable. These read-only settings have this icon next to them: . See [Chapter 1 NSX Manager](#) for more information.

The NSX-T Data Center logical load balancer offers high-availability service for applications and distributes the network traffic load among multiple servers.



The load balancer distributes incoming service requests evenly among multiple servers in such a way that the load distribution is transparent to users. Load balancing helps in achieving optimal resource utilization, maximizing throughput, minimizing response time, and avoiding overload.

You can map a virtual IP address to a set of pool servers for load balancing. The load balancer accepts TCP, UDP, HTTP, or HTTPS requests on the virtual IP address and decides which pool server to use.

Depending on your environment needs, you can scale the load balancer performance by increasing the existing virtual servers and pool members to handle heavy network traffic load.


**Note** Logical load balancer is supported only on the Tier-1 logical router. One load balancer can be attached only to a Tier-1 logical router.

## Key Load Balancer Concepts

Load balancer includes virtual servers, server pools, and health checks monitors.

A load balancer is connected to a Tier-1 logical router. The load balancer hosts single or multiple virtual servers. A virtual server is an abstract of an application service, represented by a unique combination of IP, port, and protocol. The virtual server is associated to single to multiple server pools. A server pool consists of a group of servers. The server pools include individual server pool members.


To test whether each server is correctly running the application, you can add health check monitors that check the health status of a server.



## Configuring Load Balancer Components

To use logical load balancers, you must start by configuring a load balancer and attaching it to a Tier-1 logical router.

Next, you can set up health check monitoring for your servers. You must then configure server pools for your load balancer. Finally, you must create a layer 4 or layer 7 virtual server for your load balancer.




## Create a Load Balancer in Manager Mode

Load balancer is created and attached to the Tier-1 logical router.

**Important** The information in this topic is specific to administering your environment in manager mode. For more information about manager mode and policy mode, see [Chapter 1 NSX Manager](#). For information about load balancers in policy mode, see [Chapter 8 Load Balancing](#).

You can configure the level of error messages you want the load balancer to add to the error log. Avoid setting the log level to DEBUG on load balancers with significant traffic due to the number of messages printed to the log that affect performance.



### Prerequisites

- Verify that a Tier-1 logical router is configured. See [Create a Tier-1 Logical Router in Manager Mode](#).
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

**Procedure**

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Load Balancer > Add**.
- 3 Enter a name and a description for the load balancer.
- 4 Select the load balancer virtual server size and number of pool members based on your available resources.
- 5 Define the severity level of the error log from the drop-down menu.  
Load balancer collects information about encountered issues of different severity levels to the error log.
- 6 Click **OK**.
- 7 Associate the newly created load balancer to a virtual server.
  - a Select the load balancer and click **Actions > Attach to a Virtual Server**.
  - b Select an existing virtual server from the drop-down menu.
  - c Click **OK**.
- 8 Attach the newly created load balancer to a Tier-1 logical router.
  - a Select the load balancer and click **Actions > Attach to a Logical Router**.
  - b Select an existing Tier-1 logical router from the drop-down menu.  
The Tier-1 router must be in the Active-Standby mode.
  - c Click **OK**.
- 9 (Optional) Delete the load balancer.  
If you no longer want to use this load balancer, you must first detach the load balancer from the virtual server and Tier-1 logical router.

**Configure an Active Health Monitor in Manager Mode**


The active health monitor is used to test whether a server is available. The active health monitor uses several types of tests such as sending a basic ping to servers or advanced HTTP requests to monitor the application health.

Servers that fail to respond within a certain time period or respond with errors are excluded from future connection handling until a subsequent periodic health check finds these servers to be healthy.

Active health checks are performed on server pool members after the pool member is attached to a virtual server and that virtual server is attached to a Tier-1 gateway (previously called a Tier-1 logical router).

If the Tier-1 gateway is connected to a Tier-0 gateway, a router link port is created and its IP address (typically in the 100.64.x.x format) is used to perform the health check for the load balancer service. If the Tier-1 gateway is standalone (has only one centralized service port and is not connected to a Tier-0 gateway), the centralized service port IP address is used to perform the health check for the load balancer service. See [Create a Standalone Tier-1 Logical Router in Manager Mode](#) for information about standalone Tier-1 gateways.

**Note** More than one active health monitor can be configured per server pool.



### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Load Balancing > Monitors > Active Health Monitors > Add**.
- 3 Enter a name and description for the active health monitor.
- 4 Select a health check protocol for the server from the drop-down menu.  
You can also use predefined protocols in NSX Manager; `http-monitor`, `https-monitor`, `Icmp-monitor`, `Tcp-monitor`, and `Udp-monitor`.
- 5 Set the value of the monitoring port.

## 6 Configure the values to monitor a service pool.

You can also accept the default active health monitor values.

| Option                     | Description                                                                                                       |
|----------------------------|-------------------------------------------------------------------------------------------------------------------|
| <b>Monitoring Interval</b> | Set the time in seconds that the monitor sends another connection request to the server.                          |
| <b>Fall Count</b>          | Set a value when the consecutive failures reach this value, the server is considered temporarily unavailable.     |
| <b>Rise Count</b>          | Set a number after this timeout period, the server is tried again for a new connection to see if it is available. |
| <b>Timeout Period</b>      | Set the number of times the server is tested before it is considered as DOWN.                                     |

For example, if the monitoring interval is set as 5 seconds and the timeout as 15 seconds, the load balancer send requests to the server every 5 seconds. In each probe, if the expected response is received from the server within 15 seconds, then the health check result is OK. If not, then the result is CRITICAL. If the recent three health check results are all UP, the server is considered as UP.

## 7 If you select HTTP as the health check protocol, complete the following details.

| Option                      | Description                                                                                                                                                                                                                                                                                                                                                |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>HTTP Method</b>          | Select the method for detecting the server status from the drop-down menu, GET, OPTIONS, POST, HEAD, and PUT.                                                                                                                                                                                                                                              |
| <b>HTTP Request URL</b>     | Enter the request URI for the method.<br>ASCII control characters (backspace, vertical tab, horizontal tab, line feed, etc), unsafe characters such as a space, \, <, >, {, }, and any character outside the ASCII character set are not allowed in the request URL and should be encoded. For example, replace a space with a plus (+) sign, or with %20. |
| <b>HTTP Request Version</b> | Select the supported request version from the drop-down menu.<br>You can also accept the default version, HTTP_VERSION_1_1.                                                                                                                                                                                                                                |
| <b>HTTP Request Body</b>    | Enter the request body.<br>Valid for the POST and PUT methods.                                                                                                                                                                                                                                                                                             |
| <b>HTTP Response Code</b>   | Enter the string that the monitor expects to match in the status line of HTTP response body.<br>The response code is a comma-separated list.<br>For example, 200,301,302,401.                                                                                                                                                                              |
| <b>HTTP Response Body</b>   | If the HTTP response body string and the HTTP health check response body match, then the server is considered as healthy.                                                                                                                                                                                                                                  |

8 If you select HTTPs as the health check protocol, complete the following details.

- a Select the SSL protocol list.

TLS versions TLS1.1 and TLS1.2 versions are supported and enabled by default. TLS1.0 is supported, but disabled by default.

- b Click the arrow and move the protocols into the selected section.

- c Assign a default SSL cipher or create a custom SSL cipher.

- d Complete the following details for HTTP as the health check protocol.

| Option                      | Description                                                                                                                                                                                                                                                                                                                                                |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>HTTP Method</b>          | Select the method for detecting the server status from the drop-down menu: GET, OPTIONS, POST, HEAD, and PUT.                                                                                                                                                                                                                                              |
| <b>HTTP Request URL</b>     | Enter the request URI for the method.<br>ASCII control characters (backspace, vertical tab, horizontal tab, line feed, etc), unsafe characters such as a space, \, <, >, {, }, and any character outside the ASCII character set are not allowed in the request URL and should be encoded. For example, replace a space with a plus (+) sign, or with %20. |
| <b>HTTP Request Version</b> | Select the supported request version from the drop-down menu.<br>You can also accept the default version, HTTP_VERSION_1_1.                                                                                                                                                                                                                                |
| <b>HTTP Request Body</b>    | Enter the request body.<br>Valid for the POST and PUT methods.                                                                                                                                                                                                                                                                                             |
| <b>HTTP Response Code</b>   | Enter the string that the monitor expects to match in the status line of HTTP response body.<br>The response code is a comma-separated list.<br>For example, 200,301,302,401.                                                                                                                                                                              |
| <b>HTTP Response Body</b>   | If the HTTP response body string and the HTTP health check response body match, then the server is considered as healthy.                                                                                                                                                                                                                                  |

9 If you select ICMP as the health check protocol, assign the data size in byte of the ICMP health check packet.

10 If you select TCP as the health check protocol, you can leave the parameters empty.

If both the sent and expected are not listed, then a three-way handshake TCP connection is established to validate the server health. No data is sent. Expected data if listed has to be a string and can be anywhere in the response. Regular expressions are not supported.

11 If you select UDP as the health check protocol, complete the following required details.

| Required Option          | Description                                                                                                                                        |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>UDP Data Sent</b>     | Enter the string to be sent to a server after a connection is established.                                                                         |
| <b>UDP Data Expected</b> | Enter the string expected to receive from the server.<br>Only when the received string matches this definition, is the server is considered as UP. |



## 12 Click **Finish**.

### What to do next

Associate the active health monitor with a server pool. See [Add a Server Pool for Load Balancing in Manager Mode](#).

## Configure Passive Health Monitors in Manager Mode

Load balancers perform passive health checks to monitor failures during client connections and mark servers causing consistent failures as DOWN.

Passive health check monitors client traffic going through the load balancer for failures. For example, if a pool member sends a TCP Reset (RST) in response to a client connection, the load balancer detects that failure. If there are multiple consecutive failures, then the load balancer considers that server pool member to be temporarily unavailable and stops sending connection requests to that pool member for some time. After some time, the load balancer sends a connection request to check if the pool member has recovered. If that connection is successful, then the pool member is considered healthy. Otherwise, the load balancer waits for some time and tries again.

Passive health check considers the following scenarios to be failures in client traffic.

- For server pools associated with Layer 7 virtual servers, if the connection to the pool member fails. For example, if the pool member sends a TCP RST when the load balancer tries to connect or perform a SSL handshake between load balancer and the pool member fails.
- For server pools associated with Layer 4 TCP virtual servers, if the pool member sends a TCP RST in response to client TCP SYN or does not respond at all.
- For server pools associated with Layer 4 UDP virtual servers, if a port is unreachable or a destination unreachable ICMP error message is received in response to a client UDP packet.

Server pools associated to Layer 7 virtual servers, the failed connection count is incremented when any TCP connection errors, for example, TCP RST failure to send data or SSL handshake failures occur.

Server pools associated to Layer 4 virtual servers, if no response is received to a TCP SYN sent to the server pool member or if a TCP RST is received in response to a TCP SYN, then the server pool member is considered as DOWN. The failed count is incremented.

For Layer 4 UDP virtual servers, if an ICMP error such as, port or destination unreachable message is received in response to client traffic, then it is considered as DOWN.

---

**Note** One passive health monitor can be configured per server pool.

---

### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

**Procedure**

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Networking > Load Balancing > Monitors > Passive Health Monitors > Add**.
- 3 Enter a name and description for the passive health monitor.
- 4 Configure the values to monitor a service pool.

You can also accept the default active health monitor values.

| Option                | Description                                                                                                   |
|-----------------------|---------------------------------------------------------------------------------------------------------------|
| <b>Fall Count</b>     | Set a value when the consecutive failures reach this value, the server is considered temporarily unavailable. |
| <b>Timeout Period</b> | Set the number of times the server is tested before it is considered as DOWN.                                 |

For example, when the consecutive failures reach the configured value 5, that member is considered temporarily unavailable for 5 seconds. After this period, that member is tried again for a new connection to see if it is available. If that connection is successful, then the member is considered available and the failed count is set to zero. However, if that connection fails, then it is not used for another timeout interval of 5 seconds.

- 5 Click **OK**.

**What to do next**

Associate the passive health monitor with a server pool. See [Add a Server Pool for Load Balancing in Manager Mode](#).

**Add a Server Pool for Load Balancing in Manager Mode**

Server pool consists of one or more servers that are configured and running the same application. A single pool can be associated to both Layer 4 and Layer 7 virtual servers.





Figure 20-10. Server Pool Parameter Configuration



### Prerequisites

- If you use dynamic pool members, a NSGroup must be configured. See [Create an NSGroup in Manager Mode](#).
- Depending on the monitoring you use, verify that active or passive health monitors are configured. See [Configure an Active Health Monitor in Manager Mode](#) or [Configure Passive Health Monitors in Manager Mode](#).
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Load Balancing > Server Pools > Add**.
- 3 Enter a name and description for the load balancer pool.

You can optionally describe the connections managed by the server pool.

- 4 Select the algorithm balancing method for the server pool.

Load balancing algorithm controls how the incoming connections are distributed among the members. The algorithm can be used on a server pool or a server directly.

All load balancing algorithms skip servers that meet any of the following conditions:

- Admin state is set to DISABLED.
- Admin state is set to GRACEFUL\_DISABLED and no matching persistence entry.
- Active or passive health check state is DOWN.

- Connection limit for the maximum server pool concurrent connections is reached.

| Option                           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>ROUND_ROBIN</b>               | Incoming client requests are cycled through a list of available servers capable of handling the request.<br>Ignores the server pool member weights even if they are configured.                                                                                                                                                                                                                                                                                       |
| <b>WEIGHTED_ROUND_ROBIN</b>      | Each server is assigned a weight value that signifies how that server performs relative to other servers in the pool. The value determines how many client requests are sent to a server compared to other servers in the pool.<br>This load balancing algorithm focuses on fairly distributing the load among the available server resources.                                                                                                                        |
| <b>LEAST_CONNECTION</b>          | Distributes client requests to multiple servers based on the number of connections already on the server.<br>New connections are sent to the server with the fewest connections. Ignores the server pool member weights even if they are configured.                                                                                                                                                                                                                  |
| <b>WEIGHTED_LEAST_CONNECTION</b> | Each server is assigned a weight value that signifies how that server performs relative to other servers in the pool. The value determines how many client requests are sent to a server compared to other servers in the pool.<br>This load balancing algorithm focuses on using the weight value to distribute the load among the available server resources fairly.<br>By default, the weight value is 1 if the value is not configured and slow start is enabled. |
| <b>IP-HASH</b>                   | Selects a server based on a hash of the source IP address and the total weight of all the running servers.                                                                                                                                                                                                                                                                                                                                                            |

- 5 Toggle the TCP Multiplexing button to enable this menu item.

With TCP multiplexing, you can use the same TCP connection between a load balancer and the server for sending multiple client requests from different client TCP connections.

- 6 Set the maximum number of TCP multiplexing connections per pool that are kept alive to send future client requests.

## 7 Select the Source NAT (SNAT) mode.

Depending on the topology, SNAT might be required so that the load balancer receives the traffic from the server destined to the client. SNAT can be enabled per server pool.

| Mode                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Transparent Mode</b> | <p>Load balancer uses the client IP address and port spoofing while establishing connections to the servers.</p> <p>SNAT is not required.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Auto Map Mode</b>    | <p>Load Balancer uses the interface IP address and ephemeral port to continue the communication with a client initially connected to one of the server's established listening ports.</p> <p>SNAT is required.</p> <p>Enable port overloading to allow the same SNAT IP and port to be used for multiple connections if the tuple (source IP, source port, destination IP, destination port, and IP protocol) is unique after the SNAT process is performed.</p> <p>You can also set the port overload factor to allow the maximum number of times a port can be used simultaneously for multiple connections.</p>                                                                                                                                                                                                                                                               |
| <b>IP List Mode</b>     | <p>Specify a single IP address range, for example, 1.1.1.1-1.1.1.10 to be used for SNAT while connecting to any of the servers in the pool.</p> <p>By default, from 4000 through 64000 port range is used for all configured SNAT IP addresses. Port ranges from 1000 through 4000 are reserved for purposes such as, health checks and connections initiated from Linux applications. If multiple IP addresses are present, then they are selected in a Round Robin manner.</p> <p>Enable port overloading to allow the same SNAT IP and port to be used for multiple connections if the tuple (source IP, source port, destination IP, destination port, and IP protocol) is unique after the SNAT process is performed.</p> <p>You can also set the port overload factor to allow the maximum number of times a port can be used simultaneously for multiple connections.</p> |

## 8 Select the server pool members.

Server pool consists of single or multiple pool members. Each pool member has an IP address and a port.

Each server pool member can be configured with a weight for use in the load balancing algorithm. The weight indicates how much more or less load a given pool member can handle relative to other members in the same pool.

Designating a pool member as a backup member works with the health monitor to provide an active/standby state. If active members fail a health check, traffic failover occurs for backup members.

| Option         | Description                                                                                                                                                               |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Static</b>  | Click <b>Add</b> to include a static pool member.<br>You can also clone an existing static pool member.                                                                   |
| <b>Dynamic</b> | Select the NSGroup from the drop-down menu.<br>The server pool membership criteria is defined in the group. You can optionally, define the maximum group IP address list. |

- 9 Enter the minimum number of active members the server pool must always maintain.
- 10 Select an active and passive health monitor for the server pool from the drop-down menu.

Setting an active and passive health monitor for the server pool is optional. When you select an active health monitor and if the Tier-1 gateway is connected to a Tier-0 gateway, a router link port is created. The router link port's IP address (typically in the 100.64.x.x format) is used to perform the health check for the load balancer service. If the Tier-1 gateway is standalone (has only one centralized service port and is not connected to a Tier-0 gateway), the centralized service port IP address is used to perform the health check for the load balancer service. See [Create a Standalone Tier-1 Logical Router in Manager Mode](#) for information about standalone Tier-1 gateways.

Add a firewall rule to allow the IP address to perform the health check for the load balancer service.

- 11 Click **Finish**.

## Configuring Virtual Server Components

With the virtual server there are several components that you can configure such as, application profiles, persistent profiles, and load balancer rules.





Figure 20-11. Virtual Server Components



### Configure Application Profiles in Manager Mode

Application profiles are associated with virtual servers to enhance load balancing network traffic and simplify traffic-management tasks.

Application profiles define the behavior of a particular type of network traffic. The associated virtual server processes network traffic according to the values specified in the application profile. Fast TCP, Fast UDP, and HTTP application profiles are the supported types of profiles.

TCP application profile is used by default when no application profile is associated to a virtual server. TCP and UDP application profiles are used when an application is running on a TCP or UDP protocol and does not require any application level load balancing such as, HTTP URL load balancing. These profiles are also used when you only want Layer 4 load balancing, which has faster performance and supports connection mirroring.

HTTP application profile is used for both HTTP and HTTPS applications when the load balancer needs to take actions based on Layer 7 such as, load balancing all images requests to a specific server pool member or terminating HTTPS to offload SSL from pool members. Unlike the TCP application profile, the HTTP application profile terminates the client TCP connection before selecting the server pool member.

Figure 20-12. Layer 4 TCP and UDP Application Profile





Figure 20-13. Layer 7 HTTPS Application Profile



### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Load Balancing > Profiles > Application Profiles**.
- 3 Create a Fast TCP application profile.
  - a Select **Add > Fast TCP Profile** from the drop-down menu.
  - b Enter a name and a description for the Fast TCP application profile.



- c Complete the application profile details.

You can also accept the default FAST TCP profile settings.

| Option                          | Description                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Connection Idle Timeout</b>  | Enter the time in seconds on how long the server can remain idle after a TCP connection is established.<br><br>Set the idle time to the actual application idle time and add a few more seconds so that the load balancer does not close its connections before the application does. |
| <b>Connection Close Timeout</b> | Enter the time in seconds that the TCP connection both FINs or RST must be kept for an application before closing the connection.<br><br>A short closing timeout might be required to support fast connection rates.                                                                  |
| <b>HA Flow Mirroring</b>        | Toggle the button to make all the flows to the associated virtual server mirrored to the HA standby node.                                                                                                                                                                             |

- d Click **OK**.

#### 4 Create a Fast UDP application profile.

You can also accept the default UDP profile settings.

- a Select **Add > Fast UDP Profile** from the drop-down menu.
- b Enter a name and a description for the Fast UDP application profile.
- c Complete the application profile details.

| Option                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Idle Timeout</b>      | Enter the time in seconds on how long the server can remain idle after a UDP connection is established.<br><br>UDP is a connectionless protocol. For load balancing purposes, all the UDP packets with the same flow signature such as, source and destination IP address or ports and IP protocol received within the idle timeout period are considered to belong to the same connection and sent to the same server.<br><br>If no packets are received during the idle timeout period, the connection which is an association between the flow signature and the selected server is closed. |
| <b>HA Flow Mirroring</b> | Toggle the button to make all the flows to the associated virtual server mirrored to the HA standby node.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

- d Click **OK**.

#### 5 Create an HTTP application profile.

You can also accept the default HTTP profile settings.

HTTP application profile is used for both HTTP and HTTPS applications.

- a Select **Add > Fast HTTP Profile** from the drop-down menu.
- b Enter a name and a description for the HTTP application profile.

## c Complete the application profile details.

| Option                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Redirection             | <ul style="list-style-type: none"> <li>■ None - If a website is temporarily down, user receives a page not found error message.</li> <li>■ HTTP Redirect - If a website is temporarily down or has moved, incoming requests for that virtual server can be temporarily redirected to a URL specified here. Only a static redirection is supported.<br/><br/>For example, if HTTP Redirect is set to <code>http://sitedown.abc.com/sorry.html</code>, then irrespective of the actual request, for example, <code>http://original_app.site.com/home.html</code> or <code>http://original_app.site.com/somepage.html</code>, incoming requests are redirected to the specified URL when the original website is down.</li> <li>■ HTTP to HTTPS Redirect - Certain secure applications might want to force communication over SSL, but instead of rejecting non-SSL connections, they can redirect the client request to use SSL. With HTTP to HTTPS Redirect, you can preserve both the host and URI paths and redirect the client request to use SSL.<br/><br/>For HTTP to HTTPS redirect, the HTTPS virtual server must have port 443 and the same virtual server IP address must be configured on the same load balancer.<br/><br/>For example, a client request for <code>http://app.com/path/page.html</code> is redirected to <code>https://app.com/path/page.html</code>. If either the host name or the URI must be modified while redirecting, for example, redirect to <code>https://secure.app.com/path/page.html</code>, then load balancing rules must be used.</li> </ul> |
| X-Forwarded-For (XFF)   | <ul style="list-style-type: none"> <li>■ <b>Insert</b> - If the XFF HTTP header is not present in the incoming request, the load balancer inserts a new XFF header with the client IP address. If the XFF HTTP header is present in the incoming request, the load balancer appends the XFF header with the client IP address.</li> <li>■ <b>Replace</b> - If the XFF HTTP header is present in the incoming request, the load balancer replaces the header.</li> </ul> <p>Web servers log each request they handle with the requesting client IP address. These logs are used for debugging and analytics purposes. If the deployment topology requires SNAT on the load balancer, then server uses the client SNAT IP address which defeats the purpose of logging.</p> <p>As a workaround, the load balancer can be configured to insert XFF HTTP header with the original client IP address. Servers can be configured to log the IP address in the XFF header instead of the source IP address of the connection.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Connection Idle Timeout | Enter the time in seconds on how long an HTTP application can remain idle, instead of the TCP socket setting which must be configured in the TCP application profile.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Request Header Size     | Specify the maximum buffer size in bytes used to store HTTP request headers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| NTLM Authentication     | Toggle the button for the load balancer to turn off TCP multiplexing and enable HTTP keep-alive.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

| Option | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|        | <p>NTLM is an authentication protocol that can be used over HTTP. For load balancing with NTLM authentication, TCP multiplexing must be disabled for the server pools hosting NTLM-based applications. Otherwise, a server-side connection established with one client's credentials can potentially be used for serving another client's requests.</p> <p>If NTLM is enabled in the profile and associated to a virtual server, and TCP multiplexing is enabled at the server pool, then NTLM takes precedence. TCP multiplexing is not performed for that virtual server. However, if the same pool is associated to another non-NTLM virtual server, then TCP multiplexing is available for connections to that virtual server.</p> <p>If the client uses HTTP/1.0, the load balancer upgrades to HTTP/1.1 protocol and the HTTP keep-alive is set. All HTTP requests received on the same client-side TCP connection are sent to the same server over a single TCP connection to ensure that reauthorization is not required.</p> |

- d Click **OK**.


### Configure Persistent Profiles in Manager Mode

To ensure stability of stateful applications, load balancers implement persistence which directs all related connections to the same server. Different types of persistence are supported to address different types of application needs.

Some applications maintain the server state such as, shopping carts. Such state might be per client and identified by the client IP address or per HTTP session. Applications might access or modify this state while processing subsequent related connections from the same client or HTTP session.

Source IP persistence profile tracks sessions based on the source IP address. When a client requests a connection to a virtual server that enables the source address persistence, the load balancer checks if that client was previously connected, if so, returns the client to the same server. If not, you can select a server pool member based on the pool load balancing algorithm. Source IP persistence profile is used by Layer 4 and Layer 7 virtual servers.

Cookie persistence profile inserts a unique cookie to identify the session the first time a client accesses the site. The HTTP cookie is forwarded by the client in subsequent requests and the load balancer uses that information to provide the cookie persistence. Cookie persistence profile can only be used by Layer 7 virtual servers. Note that blank space in a cookie name is **not** supported.



### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Load Balancing > Profiles > Persistence Profiles**.
- 3 Create a Source IP persistence profile.
  - a Select **Add > Source IP Persistence** from the drop-down menu.
  - b Enter a name and a description for the Source IP persistence profile.

- c Complete the persistence profile details.

You can also accept the default Source IP profile settings.

| Option                           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Share Persistence</b>         | <p>Toggle the button to share the persistence so that all virtual servers this profile is associated with can share the persistence table.</p> <p>If persistence sharing is not enabled in the Source IP persistence profile associated to a virtual server, each virtual server that the profile is associated to maintain a private persistence table.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Persistence Entry Timeout</b> | <p>Enter the persistence expiration time in seconds.</p> <p>The load balancer persistence table maintains entries to record that client requests are directed to the same server.</p> <ul style="list-style-type: none"> <li>■ If no new connection requests are received from the same client within the timeout period, the persistence entry expires and is deleted.</li> <li>■ If a new connection request from the same client is received within the timeout period, the timer is reset, and the client request is sent to a sticky pool member.</li> </ul> <p>After the timeout period has expired, new connection requests are sent to a server allocated by the load balancing algorithm. For the L7 load balancing TCP source IP persistence scenario, the persistence entry times out if no new TCP connections are made for some time, even if the existing connections are still alive.</p> |
| <b>HA Persistence Mirroring</b>  | <p>Toggle the button to synchronize persistence entries to the HA peer.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Purge Entries When Full</b>   | <p>Purge entries when the persistence table is full.</p> <p>A large timeout value might lead to the persistence table quickly filling up when the traffic is heavy. When the persistence table fills up, the oldest entry is deleted to accept the newest entry.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

- d Click **OK**.

#### 4 Create a Cookie persistence profile.

- Select **Add > Cookie Persistence** from the drop-down menu.
- Enter a name and a description for the Cookie persistence profile.
- Toggle the **Share Persistence** button to share persistence across multiple virtual servers that are associated to the same pool members.

The Cookie persistence profile inserts a cookie with the format, *<name>.<profile-id>.<pool-id>*.

If the persistence shared is not enabled in the Cookie persistence profile associated with a virtual server, the private Cookie persistence for each virtual server is used and is qualified by the pool member. The load balancer inserts a cookie with the format, *<name>.<virtual\_server\_id>.<pool\_id>*.

- d Click **Next**.

- e Complete the persistence profile details.

| Option          | Description                                                                                                                                                                                                                                                                                |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cookie Mode     | Select a mode from the drop-down menu. <ul style="list-style-type: none"> <li>■ INSERT - Adds a unique cookie to identify the session.</li> <li>■ PREFIX - Appends to the existing HTTP cookie information.</li> <li>■ REWRITE - Rewrites the existing HTTP cookie information.</li> </ul> |
| Cookie Name     | Enter the cookie name. A blank space in a cookie name is <b>not</b> supported.                                                                                                                                                                                                             |
| Cookie Domain   | Enter the domain name.<br>HTTP cookie domain can be configured only in the INSERT mode.                                                                                                                                                                                                    |
| Cookie Path     | Enter the cookie URL path.<br>HTTP cookie path can be set only in the INSERT mode.                                                                                                                                                                                                         |
| Cookie Garbling | Encrypt the cookie server IP address and port information.<br>Toggle the button to disable encryption. When garbling is disabled, the cookie server IP address and port information is in a plain text.                                                                                    |
| Cookie Fallback | Select a new server to handle a client request if the cookie points to a server that is in a DISABLED or is in a DOWN state.<br>Toggle the button so that the client request is rejected if cookie points to a server that is in a DISABLED or is in a DOWN state.                         |

- f Complete the Cookie expiry details.

| Option             | Description                                                                                                                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cookie Time Type   | Select a cookie time type from the drop-down menu.<br><b>Session Cookie</b> is not stored and will be lost when the browser is closed.<br><b>Persistence Cookie</b> is stored by the browser and is not lost when the browser is closed. |
| Maximum Idle Time  | Enter the time in seconds that a cookie can be idle before it expires.                                                                                                                                                                   |
| Maximum Cookie Age | For <b>Session Cookie</b> only. Enter the maximum age in seconds that a cookie can be active.                                                                                                                                            |

- g Click **Finish**.

### Configure SSL Profile in Manager Mode

SSL profiles configure application-independent SSL properties such as, cipher lists and reuse these lists across multiple applications. SSL properties are different when the load balancer is acting as a client and as a server, as a result separate SSL profiles for client-side and server-side are supported.

**Note** SSL profile is not supported in the NSX-T Data Center limited export release.

Client-side SSL profile refers to the load balancer acting as an SSL server and terminating the client SSL connection. Server-side SSL profile refers to the load balancer acting as a client and establishing a connection to the server.

You can specify a cipher list on both the client-side and server-side SSL profiles.

SSL session caching allows the SSL client and server to reuse previously negotiated security parameters avoiding the expensive public key operation during the SSL handshake. SSL session caching is disabled by default on both the client-side and server-side.

SSL session tickets are an alternate mechanism that allow the SSL client and server to reuse previously negotiated session parameters. In SSL session tickets, the client and server negotiate whether they support SSL session tickets during the handshake exchange. If supported by both, server can send an SSL ticket, which includes encrypted SSL session parameters to the client. The client can use that ticket in subsequent connections to reuse the session. SSL session tickets are enabled on the client-side and disabled on the server-side.

Figure 20-14. SSL Offloading





Figure 20-15. End-to-End SSL



### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

**Procedure**

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

- 2 Select **Networking > Load Balancing > Profiles > SSL Profiles**.

- 3 Create a Client SSL profile.

- a Select **Add > Client Side SSL** from the drop-down menu.
- b Enter a name and a description for the Client SSL profile.
- c Assign the SSL Ciphers to be included in the Client SSL profile.

You can also create custom SSL Ciphers.

- d Click the arrow to move the ciphers to the Selected section.
- e Click the **Protocols and Sessions** tab.
- f Select the SSL protocols to be included in the Client SSL profile.

SSL protocol versions TLS1.1 and TLS1.2 are enabled by default. TLS1.0 is also supported, but disabled by default.

- g Click the arrow to move the protocol to the Selected section.
- h Complete the SSL protocol details.

You can also accept the default SSL profile settings.

| Option                             | Description                                                                                                                                                                                                   |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Session Caching</b>             | SSL session caching allows the SSL client and server to reuse previously negotiated security parameters avoiding the expensive public key operation during an SSL handshake.                                  |
| <b>Session Cache Entry Timeout</b> | Enter the cache timeout in seconds to specify how long the SSL session parameters must be kept and can be reused.                                                                                             |
| <b>Prefer Server Cipher</b>        | Toggle the button so that the server can select the first supported cipher from the list it can support.<br><br>During an SSL handshake, the client sends an ordered list of supported ciphers to the server. |

- i Click **OK**.

- 4 Create a Server SSL profile.

- a Select **Add > Server Side SSL** from the drop-down menu.
- b Enter a name and a description for the Server SSL profile.
- c Select the SSL Ciphers to be included in the Server SSL profile.

You can also create custom SSL Ciphers.

- d Click the arrow to move the ciphers to the Selected section.



- e Click the **Protocols and Sessions** tab.
- f Select the SSL protocols to be included in the Server SSL profile.  
 SSL protocol versions TLS1.1 and TLS1.2 are enabled by default. TLS1.0 is also supported, but disabled by default.
- g Click the arrow to move the protocol to the Selected section.
- h Accept the default session caching setting.  
 SSL session caching allows the SSL client and server to reuse previously negotiated security parameters avoiding the expensive public key operation during an SSL handshake.
- i Click **OK**.

### Configure Layer 4 Virtual Servers in Manager Mode

Virtual servers receive all the client connections and distribute them among the servers. A virtual server has an IP address, a port, and a protocol. For Layer 4 virtual servers, lists of ports ranges can be specified instead of a single TCP or UDP port to support complex protocols with dynamic ports.

A Layer 4 virtual server must be associated to a primary server pool, also called a default pool.

If a virtual server status is disabled, any new connection attempts to the virtual server are rejected by sending either a TCP RST for the TCP connection or ICMP error message for UDP. New connections are rejected even if there are matching persistence entries for them. Active connections continue to be processed. If a virtual server is deleted or disassociated from a load balancer, then active connections to that virtual server fail.

#### Prerequisites

- Verify that application profiles are available. See [Configure Application Profiles in Manager Mode](#).
- Verify that persistent profiles are available. See [Configure Persistent Profiles in Manager Mode](#).
- Verify that SSL profiles for the client and server are available. See [Configure SSL Profile in Manager Mode](#).
- Verify that server pools are available. See [Add a Server Pool for Load Balancing in Manager Mode](#).
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

#### Procedure

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

- 2 Select **Networking > Load Balancing > Virtual Servers > Add**.

- 3 Enter a name and a description for the Layer 4 virtual server.

- 4 Select a Layer 4 protocol from the drop-down menu.

Layer 4 virtual servers support either the Fast TCP or Fast UDP protocol, but not both. For Fast TCP or Fast UDP protocol support on the same IP address and port, for example DNS, a virtual server must be created for each protocol.

Based on the protocol type, the existing application profile is automatically populated.

- 5 Toggle the Access Log button to enable logging for the Layer 4 virtual server.

- 6 Click **Next**.

- 7 Enter the virtual server IP address and port number.

You can enter the virtual server port number or port range.

- 8 Complete the advanced properties details.

| Option                               | Description                                                                                                                                                                                                                                                                                                                                                             |
|--------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Maximum Concurrent Connection</b> | Set the maximum concurrent connection allowed to a virtual server so that the virtual server does not deplete resources of other applications hosted on the same load balancer.                                                                                                                                                                                         |
| <b>Maximum New Connection Rate</b>   | Set the maximum new connection to a server pool member so that a virtual server does not deplete resources.                                                                                                                                                                                                                                                             |
| <b>Default Pool Member Port</b>      | Enter a default pool member port if the pool member port for a virtual server is not defined.<br><br>For example, if a virtual server is defined with port range 2000-2999 and the default pool member port range is set as 8000-8999, then an incoming client connection to the virtual server port 2500 is sent to a pool member with a destination port set to 8500. |

- 9 Select an existing server pool from the drop-down menu.

The server pool consists of one or more servers, also called pool members that are similarly configured and running the same application.

- 10 Select an existing sorry server pool from the drop-down menu.

The sorry server pool serves the request when a load balancer cannot select a backend server to the serve the request from the default pool.

- 11 Click **Next**.

- 12 Select the existing persistence profile from the drop-down menu.

Persistence profile can be enabled on a virtual server to allow related client connections to be sent to the same server.

- 13 Click **Finish**.

## Configure Layer 7 Virtual Servers in Manager Mode

Virtual servers receive all the client connections and distribute them among the servers. A virtual server has an IP address, a port, and a protocol TCP.

Load balancer rules are supported for only Layer 7 virtual servers with an HTTP application profile. Different load balancer services can use load balancer rules.

Each load balancer rule consists of single or multiple match conditions and single or multiple actions. If the match conditions are not specified, then the load balancer rule always matches and is used to define default rules. If more than one match condition is specified, then the matching strategy determines if all conditions must match or any one condition must match for the load balancer rule to be considered a match.

Each load balancer rule is implemented at a specific phase of the load balancing processing; HTTP Request Rewrite, HTTP Request Forwarding, and HTTP Response Rewrite. Not all the match conditions and actions are applicable to each phase.

If a virtual server status is disabled, any new connection attempts to the virtual server are rejected by sending either a TCP RST for the TCP connection or ICMP error message for UDP. New connections are rejected even if there are matching persistence entries for them. Active connections continue to be processed. If a virtual server is deleted or disassociated from a load balancer, then active connections to that virtual server fail.

### Prerequisites

- Verify that application profiles are available. See [Configure Application Profiles in Manager Mode](#).
- Verify that persistent profiles are available. See [Configure Persistent Profiles in Manager Mode](#).
- Verify that SSL profiles for the client and server are available. See [Configure SSL Profile in Manager Mode](#).
- Verify that server pools are available. See [Add a Server Pool for Load Balancing in Manager Mode](#).
- Verify that CA and client certificate are available. See [Create a Certificate Signing Request File](#).
- Verify that a certification revocation list (CRL) is available. See [Import a Certificate Revocation List](#).
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).
- [Configure Layer 7 Virtual Server Pool and Rules](#)

With Layer 7 virtual servers, you can optionally configure load balancer rules and customize load balancing behavior using match or action rules.

## ■ Configure Layer 7 Virtual Server Load Balancing Profiles

With Layer 7 virtual servers, you can optionally configure load balancer persistence, client-side SSL, and server-side SSL profiles.

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Load Balancing > Virtual Servers > Add**.
- 3 Enter a name and a description for the Layer 7 virtual server.
- 4 Select the Layer 7 menu item.  
  
Layer 7 virtual servers support the HTTP and HTTPS protocols.  
  
The existing HTTP application profile is automatically populated.
- 5 (Optional) Click **Next** to configure server pool and load balancing profiles.
- 6 Click **Finish**.

### Configure Layer 7 Virtual Server Pool and Rules

With Layer 7 virtual servers, you can optionally configure load balancer rules and customize load balancing behavior using match or action rules.

Load Balancer rules support REGEX for match types. PCRE style REGEX patterns is supported with a few limitations on advanced use cases. When REGEX is used in match conditions, named capturing groups are supported.

REGEX restrictions include:

- Character unions and intersections are not supported. For example, do not use `[a-z[0-9]]` and `[a-z&&[aeiou]]` instead use `[a-z0-9]` and `[aeiou]` respectively.
- Only 9 back references are supported and `\1` through `\9` can be used to refer to them.
- Use `\Odd` format to match octal characters, not the `\ddd` format.
- Embedded flags are not supported at the top level, they are only supported within groups. For example, do not use `"Case (?i:s)ensitive"` instead use `"Case ((?i:s)ensitive)"`.
- Preprocessing operations `\l`, `\u`, `\L`, `\U` are not supported. Where `\l` - lowercase next char `\u` - uppercase next char `\L` - lower case until `\E` `\U` - upper case to `\E`.
- `(?(condition)X)`, `(?{code})`, `(?{Code})` and `(?#comment)` are not supported.
- Predefined Unicode character class `\X` is not supported
- Using named character construct for Unicode characters is not supported. For example, do not use `\N{name}` instead use `\u2018`.

When REGEX is used in match conditions, named capturing groups are supported. For example, REGEX match pattern `/news/(?<year>\d+)-(?<month>\d+)-(?<day>\d+)/(?<article>.* )` can be used to match a URI like `/news/2018-06-15/news1234.html`.

Then variables are set as follows, \$year = "2018" \$month = "06" \$day = "15" \$article = "news1234.html". After the variables are set, these variables can be used in load balancer rule actions. For example, URI can be rewritten using the matched variables like, /news.py?year=\$year&month=\$month&day=\$day&article=\$article. Then the URI gets rewritten as /news.py?year=2018&month=06&day=15&article=news1234.html.

Rewrite actions can use a combination of named capturing groups and built-in variables. For example, URI can be written as /news.py?year=\$year&month=\$month&day=\$day&article=\$article&user\_ip=\$\_remote\_addr. Then the example URI gets rewritten as /news.py?year=2018&month=06&day=15&article=news1234.html&user\_ip=1.1.1.1.

---

**Note** For named capturing groups, the name cannot start with an \_ character.

---

In addition to named capturing groups, the following built-in variables can be used in rewrite actions. All the built-in variable names start with \_.

- \$\_args - arguments from the request
- \$\_arg\_<name> - argument <name> in the request line
- \$\_cookie\_<name> - value of <name> cookie
- \$\_upstream\_cookie\_<name> - cookie with the specified name sent by the upstream server in the "Set-Cookie" response header field
- \$\_upstream\_http\_<name> - arbitrary response header field and <name> is the field name converted to lower case with dashes replaced by underscores
- \$\_host - in the order of precedence - host name from the request line, or host name from the "Host" request header field, or the server name matching a request
- \$\_http\_<name> - arbitrary request header field and <name> is the field name converted to lower case with dashes replaced by underscores
- \$\_https - "on" if connection operates in SSL mode, or "" otherwise
- \$\_is\_args - "?" if a request line has arguments, or "" otherwise
- \$\_query\_string - same as \$\_args
- \$\_remote\_addr - client address
- \$\_remote\_port - client port
- \$\_request\_uri - full original request URI (with arguments)
- \$\_scheme - request scheme, "http" or "https"
- \$\_server\_addr - address of the server which accepted a request
- \$\_server\_name - name of the server which accepted a request
- \$\_server\_port - port of the server which accepted a request
- \$\_server\_protocol - request protocol, usually "HTTP/1.0" or "HTTP/1.1"

- `$_ssl_client_cert` - returns the client certificate in the PEM format for an established SSL connection, with each line except the first prepended with the tab character
- `$_ssl_server_name` - returns the server name requested through SNI
- `$_uri` - URI path in request
- `$_ssl_ciphers`: returns the client SSL ciphers
- `$_ssl_client_i_dn`: returns the "issuer DN" string of the client certificate for an established SSL connection according to RFC 2253
- `$_ssl_client_s_dn`: returns the "subject DN" string of the client certificate for an established SSL connection according to RFC 2253
- `$_ssl_protocol`: returns the protocol of an established SSL connection
- `$_ssl_session_reused`: returns "r" if an SSL session was reused, or "." otherwise

### Prerequisites

Verify a Layer 7 virtual server is available. See [Configure Layer 7 Virtual Servers in Manager Mode](#).

### Procedure

- 1 Open the Layer 7 virtual server.
- 2 Skip to the Virtual Server Identifiers page.
- 3 Enter the virtual server IP address and port number.  
You can enter the virtual server port number or port range.
- 4 Complete the advanced properties details.

| Option                               | Description                                                                                                                                                                                                                                                                                                                                                             |
|--------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Maximum Concurrent Connection</b> | Set the maximum concurrent connection allowed to a virtual server so that the virtual server does not deplete resources of other applications hosted on the same load balancer.                                                                                                                                                                                         |
| <b>Maximum New Connection Rate</b>   | Set the maximum new connection to a server pool member so that a virtual server does not deplete resources.                                                                                                                                                                                                                                                             |
| <b>Default Pool Member Port</b>      | Enter a default pool member port if the pool member port for a virtual server is not defined.<br><br>For example, if a virtual server is defined with port range 2000–2999 and the default pool member port range is set as 8000–8999, then an incoming client connection to the virtual server port 2500 is sent to a pool member with a destination port set to 8500. |

- 5 (Optional) Select an existing default server pool from the drop-down menu.

The server pool consists of one or more servers, called pool members that are similarly configured and running the same application.

- 6 Click **Add** to configure the load balancer rules for the HTTP Request Rewrite phase.

Supported match types are, REGEX, STARTS\_WITH, ENDS\_WITH, etc and inverse option.

| Supported Match Condition  | Description                                                                                                                                                      |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| HTTP Request Method        | Match an HTTP request method.<br>http_request.method - value to match                                                                                            |
| HTTP Request URI           | Match an HTTP request URI without query arguments.<br>http_request.uri - value to match                                                                          |
| HTTP Request URI arguments | Match an HTTP request URI query argument.<br>http_request.uri_arguments - value to match                                                                         |
| HTTP Request Version       | Match an HTTP request version.<br>http_request.version - value to match                                                                                          |
| HTTP Request Header        | Match any HTTP request header.<br>http_request.header_name - header name to match<br>http_request.header_value - value to match                                  |
| HTTP Request Payload       | Match an HTTP request body content.<br>http_request.body_value - value to match                                                                                  |
| TCP Header Fields          | Match a TCP source or the destination port.<br>tcp_header.source_port - source port to match<br>tcp_header.destination_port - destination port to match          |
| IP Header Fields           | Match an IP source or destination address.<br>ip_header.source_address - source address to match<br>ip_header.destination_address - destination address to match |

| Action                      | Description                                                                                                                         |
|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------------|
| HTTP Request URI Rewrite    | Modify an URI.<br>http_request.uri - URI (without query arguments) to write<br>http_request.uri_args - URI query arguments to write |
| HTTP Request Header Rewrite | Modify value of an HTTP header.<br>http_request.header_name - header name<br>http_request.header_value - value to write             |

- 7 Click **Add** to configure the load balancer rules for the HTTP Request Forwarding.

All match values accept regular expressions.

| Supported Match Condition | Description                                                                         |
|---------------------------|-------------------------------------------------------------------------------------|
| HTTP Request Method       | Match an HTTP request method.<br>http_request.method - value to match               |
| HTTP Request URI          | Match an HTTP request URI.<br>http_request.uri - value to match                     |
| HTTP Request URI args     | Match an HTTP request URI query argument.<br>http_request.uri_args - value to match |

| Supported Match Condition | Description                                                                                                                                             |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| HTTP Request Version      | Match an HTTP request version.<br>http_request.version - value to match                                                                                 |
| HTTP Request Header       | Match any HTTP request header.<br>http_request.header_name - header name to match<br>http_request.header_value - value to match                         |
| HTTP Request Payload      | Match an HTTP request body content.<br>http_request.body_value - value to match                                                                         |
| TCP Header Fields         | Match a TCP source or the destination port.<br>tcp_header.source_port - source port to match<br>tcp_header.destination_port - destination port to match |
| IP Header Fields          | Match an IP source address.<br>ip_header.source_address - source address to match                                                                       |

| Action      | Description                                                                                                                                                                                             |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Reject      | Reject a request, for example, by setting status to 5xx.<br>http_forward.reply_status - HTTP status code used to reject<br>http_forward.reply_message - HTTP rejection message                          |
| Redirect    | Redirect a request. Status code must be set to 3xx.<br>http_forward.redirect_status - HTTP status code for redirect<br>http_forward.redirect_url - HTTP redirect URL                                    |
| Select Pool | Force the request to a specific server pool. Specified pool member's configured algorithm (predictor) is used to select a server within the server pool.<br>http_forward.select_pool - server pool UUID |

- 8 Click **Add** to configure the load balancer rules for the HTTP Response Rewrite.

All match values accept regular expressions.

| Supported Match Condition | Description                                                                                                                        |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------|
| HTTP Response Header      | Match any HTTP response header.<br>http_response.header_name - header name to match<br>http_response.header_value - value to match |

| Action                       | Description                                                                                                                            |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------|
| HTTP Response Header Rewrite | Modify the value of an HTTP response header.<br>http_response.header_name - header name<br>http_response.header_value - value to write |

- 9 (Optional) Click **Next** to configure load balancing profiles.

- 10 Click **Finish**.



## Configure Layer 7 Virtual Server Load Balancing Profiles

With Layer 7 virtual servers, you can optionally configure load balancer persistence, client-side SSL, and server-side SSL profiles.

---

**Note** SSL profile is not supported in the NSX-T Data Center limited export release.

---

If a client-side SSL profile binding is configured on a virtual server but not a server-side SSL profile binding, then the virtual server operates in an SSL-terminate mode, which has an encrypted connection to the client and plain text connection to the server. If both the client-side and server-side SSL profile bindings are configured, then the virtual server operates in SSL-proxy mode, which has an encrypted connection both to the client and the server.

Associating server-side SSL profile binding without associating a client-side SSL profile binding is currently not supported. If a client-side and a server-side SSL profile binding is not associated with a virtual server and the application is SSL-based, then the virtual server operates in an SSL-unaware mode. In this case, the virtual server must be configured for Layer 4. For example, the virtual server can be associated to a fast TCP profile.

### Prerequisites

Verify a Layer 7 virtual server is available. See [Configure Layer 7 Virtual Servers in Manager Mode](#).

### Procedure

- 1 Open the Layer 7 virtual server.
- 2 Skip to the Load Balancing Profiles page.
- 3 Toggle the Persistence button to enable the profile.

Persistence profile allows related client connections to be sent to the same server.

- 4 Select either the Source IP Persistence or Cookie Persistence profile.
- 5 Select the existing persistence profile from the drop-down menu.
- 6 Click **Next**.

- 7 Toggle the Client Side SSL button to enable the profile.

Client-side SSL profile binding allows multiple certificates, for different host names to be associated to the same virtual server.

The associated Client-side SSL profile is automatically populated.

- 8 Select a default certificate from the drop-down menu.

This certificate is used if the server does not host multiple host names on the same IP address or if the client does not support Server Name Indication (SNI) extension.

- 9 Select the available SNI certificate and click the arrow to move the certificate to the Selected section.
- 10 (Optional) Toggle the Mandatory Client Authentication to enable this menu item.

- 11 Select the available CA certificate and click the arrow to move the certificate to the Selected section.
- 12 Set the certificate chain depth to verify the depth in the server certificates chain.
- 13 Select the available CRL and click the arrow to move the certificate to the Selected section.  
A CRL can be configured to disallow compromised server certificates.
- 14 Click **Next**.
- 15 Toggle the Server Side SSL button to enable the profile.  
The associated Server-side SSL profile is automatically populated.
- 16 Select a client certificate from the drop-down menu.  
The client certificate is used if the server does not host multiple host names on the same IP address or if the client does not support Server Name Indication (SNI) extension.
- 17 Select the available SNI certificate and click the arrow to move the certificate to the Selected section.
- 18 (Optional) Toggle the Server Authentication to enable this menu item.  
Server-side SSL profile binding specifies whether the server certificate presented to the load balancer during the SSL handshake must be validated or not. When validation is enabled, the server certificate must be signed by one of the trusted CAs whose self-signed certificates are specified in the same server-side SSL profile binding.
- 19 Select the available CA certificate and click the arrow to move the certificate to the Selected section.
- 20 Set the certificate chain depth to verify the depth in the server certificates chain.
- 21 Select the available CRL and click the arrow to move the certificate to the Selected section.  
A CRL can be configured to disallow compromised server certificates. OCSP and OCSP stapling are not supported on the server-side.
- 22 Click **Finish**.

## Firewall in Manager Mode

You can configure Distributed Firewall and logical router Firewall in **Manager Mode**.

---

**Note** If you use **Manager** mode to modify objects created in the **Policy** mode, some settings might not be configurable. These read-only settings have this icon next to them: . See [Chapter 1 NSX Manager](#) for more information.

---

## Add or Delete a Firewall Rule to a Logical Router in Manager Mode

You can add firewall rules to a tier-0 or tier-1 logical router to control communication into the router.

Edge fire-walling is implemented on uplink router ports, meaning that firewall rules will be applicable only if traffic hits uplink router ports on edge. To apply firewall rules to particular IP destination, you must configure groups with /32 network. If you provide a subnet other than /32, firewall rules will be applied to the complete subnet.

### Prerequisites

- Familiarize yourself with the parameters of a firewall rule. See [Add a Firewall Rule in Manager Mode](#).
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Locate the router in **Networking > Tier-0 Logical Routers** or **Networking > Tier-1 Logical Routers**.
- 3 Click the name of the logical router.
- 4 Select **Services > Edge Firewall**.
- 5 Click an existing section or rule.
- 6 To add a rule, click **Add Rule** on the menu bar and select **Add Rule Above** or **Add Rule Below**, or click the menu icon in the first column of a rule and select **Add Rule Above** or **Add Rule Below**, and specify the rule parameters.

The Applied To field is not shown because this rule applies only to the logical router.

- 7 To delete a rule, select the rule, click **Delete** on the menu bar or click the menu icon in the first column and select **Delete**.

### Results

---

**Note** If you add a firewall rule to a tier-0 logical router and the NSX Edge cluster backing the router is running in active-active mode, the firewall can only run in stateless mode. If you configure the firewall rule with stateful services such as HTTP, SSL, TCP, and so on, the firewall rule will not work as expected. To avoid this issue, configure the NSX Edge cluster to run in active-standby mode.

---

## Configure Firewall for a Logical Switch Bridge Port in Manager Mode

You can configure firewall sections and firewall rules for the bridge port of a layer 2 bridge-backed logical switch. The bridge must be created using NSX Edge nodes.

## Prerequisites

- Verify that the switch is attached to a bridge profile. See [Create a Layer 2 Bridge-Backed Logical Switch in Manager Mode](#).
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

## Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Security > Bridge Firewall**.
- 3 Select a logical switch.  
The switch must be attached to a bridge profile.
- 4 Follow the same steps in previous sections for configuring layer 2 or layer 3 firewall.

## Firewall Sections and Firewall Rules

Firewall sections are used to group a set of firewall rules.

A firewall section is made up from one or more individual firewall rules. Each individual firewall rule contains instructions that determine whether a packet should be allowed or blocked; which protocols it is allowed to use; which ports it is allowed to use and so forth. Sections are used for multi-tenancy, such as specific rules for sales and engineering departments in separate sections.

A section can be defined as enforcing stateful or stateless rules. Stateless rules are treated as traditional stateless ACLs. Reflexive ACLs are not supported for stateless sections. A mix of stateless and stateful rules on a single logical switch port is not recommended and may cause undefined behavior.

Rules can be moved up and down within a section. For any traffic attempting to pass through the firewall, the packet information is subjected to the rules in the order shown in the section, beginning at the top and proceeding to the default rule at the bottom. The first rule that matches the packet has its configured action applied, and any processing specified in the rule's configured options is performed and all subsequent rules are ignored (even if a later rule is a better match). Thus, you should place specific rules above more general rules to ensure those rules are not ignored. The default rule, located at the bottom of the rule table, is a "catchall" rule; packets not matching any other rules will be enforced by the default rule.

---

**Note** A logical switch has a property called N-VDS mode. This property comes from the transport zone that the switch belongs to. If the N-VDS mode is **ENS** (also known as **Enhanced Datapath**), then you cannot create a firewall rule or section with the switch or its ports in the **Source**, **Destination**, or **Applied To** fields.

---

## Activate and Deactivate Distributed Firewall in Manager Mode

You can activate and deactivate the distributed firewall feature.

If it is deactivated, no firewall rules are enforced at the dataplane level. Upon reactivation rules are enforced.

### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 Navigate to **Security > Distributed Firewall**.
- 2 Click the **Settings** tab.
- 3 Click Distributed Firewall **Edit**.
- 4 In the dialog box, toggle the firewall status to green (activated) or gray (deactivated).
- 5 Click **Save**.

## Add a Firewall Rule Section in Manager Mode

A firewall rule section is edited and saved independently and is used to apply separate firewall configuration to tenants.

### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 Select **Security > Distributed Firewall**.
- 2 Click the **General** tab for layer 3 (L3) rules or the **Ethernet** tab for layer 2 (L2) rules.
- 3 Click an existing section or rule.
- 4 Click the section icon on the menu bar and select **Add Section Above** or **Add Section Below**.

---

**Note** For any traffic attempting to pass through the firewall, the packet information is subjected to the rules in the order shown in the Rules table, beginning at the top and proceeding to the default rules at the bottom. In some cases, the order of precedence of two or more rules might be important in determining the disposition of a packet.

---

- 5 Enter the section name.

- 6 To make the firewall stateless, select the **Enable Stateless Firewall**. This option is applicable for L3 only.

Stateless firewalls watch network traffic, and restrict or block packets based on source and destination addresses or other static values. For TCP and UDP flows, after the first packet, a cache is created and maintained for the traffic tuple in either direction, if the firewall result is ALLOW. This means that the traffic no longer needs to check with the firewall rules, resulting in lower latency. Stateless firewalls are thus typically faster and perform better under heavier traffic loads.

Stateful firewalls can watch traffic streams from end to end. The firewall is always consulted for every packet, to validate the state and sequence numbers. Stateful firewalls are better at identifying unauthorized and forged communications.

There is no toggling between stateful and stateless once it is defined.

- 7 Select one or more objects to apply the section.

The types of object are logical ports, logical switches, and NSGroups. If you select an NSGroup, it must contain one or more logical switches or logical ports. If the NSGroup contains only IP sets or MAC sets, it will be ignored.

---

**Note** The **Applied To** in a section it will override any **Applied To** settings in the rules in that section.

---

- 8 Click **OK**.

#### What to do next

Add Firewall rules to the section.

### Delete a Firewall Rule Section in Manager Mode

A firewall rule section can be deleted when it is no longer used.

When you delete a firewall rule section, all rules in that section are deleted. You cannot delete a section and add it again at a different place in the firewall table. To do so, you must delete the section and publish the configuration. Then add the deleted section to the firewall table and re-publish the configuration.

#### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

#### Procedure

- 1 Select **Security > Distributed Firewall**.
- 2 Click the **General** tab for L3 rules or the **Ethernet** tab for L2 rules.

- 3 Click the menu icon in the first column of the section and select **Delete Section**.

You can also select the section and click the delete icon on the menu bar.

## Enable and Disable Section Rules in Manager Mode

You can enable or disable all rules in a firewall rule section.

### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 Select **Security > Distributed Firewall**.
- 2 Click the **General** tab for L3 rules or the **Ethernet** tab for L2 rules.
- 3 Click the menu icon in the first column of the section and select **Enable All Rules** or **Disable All Rules**.
- 4 Click **Publish**.

## Enable and Disable Section Logs in Manager Mode

Enabling logs for section rules records information on packets for all of the rules in a section. Depending on the number of rules in a section, a typical firewall section will generate large amounts of log information and can affect performance.

Logs are stored in the /var/log/dfwpktlogs.log file on ESXi and KVM hosts.

### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 Select **Security > Distributed Firewall**.
- 2 Click the **General** tab for L3 rules or the **Ethernet** tab for L2 rules.
- 3 Click the menu icon in the first column of the section and select **Enable Logs** or **Disable Logs**.
- 4 Click **Publish**.

## Configure a Firewall Exclusion List in Manager Mode

A logical port, logical switch, or NSGroup can be excluded from a firewall rule.

After you've created a section with firewall rules you may want to exclude an NSX-T Data Center appliance port from the firewall rules.

---

**Note** NSX-T Data Center automatically adds NSX Edge node virtual machines to the firewall exclusion list.

---

### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 Select **Security > Distributed Firewall > Exclusion List > Add**.
- 2 Select a type and an object.  
The available types are **Logical Port**, **Logical Switch**, and **NSGroup**.
- 3 Click **OK**.
- 4 To remove an object from the exclusion list, select the object and click **Delete** on the menu bar.

## About Firewall Rules

NSX-T Data Center uses firewall rules to specify traffic handling in and out of the network.

Firewall offers multiple sets of configurable rules: Layer 3 rules (General tab) and Layer 2 rules (Ethernet tab). Layer 2 firewall rules are processed before Layer 3 rules and if allowed in the Layer 2 rules will then be processed by the Layer 3 rules. You can configure an exclusion list that contains logical switches, logical ports, or groups that are to be excluded from firewall enforcement.

Firewall Rules are enforced as follows:

- Rules are processed in top-to-bottom ordering.
- Each packet is checked against the top rule in the rule table before moving down the subsequent rules in the table.
- The first rule in the table that matches the traffic parameters is enforced.

No subsequent rules can be enforced as the search is then terminated for that packet. Because of this behavior, it is always recommended to put the most granular policies at the top of the rule table. This will ensure they will be enforced before more specific rules.



The default rule, located at the bottom of the rule table, is a catchall rule; packets not matching any other rules will be enforced by the default rule. After the host preparation operation, the default rule is set to allow action. This ensures that VM-to-VM communication is not broken during staging or migration phases. It is a best practice to then change this default rule to block action and enforce access control through a positive control model (i.e., only traffic defined in the firewall rule is allowed onto the network).

**Note** TCP strict can be enabled on a per section basis to turn off mid-session pick-up and enforce the requirement for a three-way handshake. When enabling TCP strict mode for a particular Distributed Firewall Section, and using a default ANY-ANY Block rule, packets that do not complete the three-way handshake connection requirements, and that match a TCP-based rule in this section are dropped. Strict is only applied to stateful TCP rules and is enabled at the distributed firewall section level. TCP strict is not enforced for packets that match a default ANY-ANY Allow which as no TCP service specified.

**Table 20-1. Properties of a Firewall Rule**

| Property    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name        | Name of the firewall rule.                                                                                                                                                                                                                                                                                                                                                                                                                              |
| ID          | Unique system generated ID for each rule.                                                                                                                                                                                                                                                                                                                                                                                                               |
| Source      | The source of the rule can be either an IP or MAC address or an object other than an IP address. The source will match any if not defined. Both IPv4 and IPv6 are supported for source or destination range.                                                                                                                                                                                                                                            |
| Destination | The destination IP or MAC address/netmask of the connection that is affected by the rule. The destination will match any if not defined. Both IPv4 and IPv6 are supported for source or destination range.                                                                                                                                                                                                                                              |
| Service     | The service can be a predefined port protocol combination for L3. For L2 it can be ether-type. For both L2 and L3 you can manually define a new service or service group. The service will match any, if it is not specified.                                                                                                                                                                                                                           |
| Applied To  | Defines the scope at which this rule is applicable. If not defined the scope will be all logical ports. If you have added "applied to" in a section it will overwrite the rule.                                                                                                                                                                                                                                                                         |
| Log         | Logging can be turned off or on. Logs are stored at /var/log/dfwpklogs.log file on ESX and KVM hosts.                                                                                                                                                                                                                                                                                                                                                   |
| Action      | The action applied by the rule can be <b>Allow</b> , <b>Drop</b> , or <b>Reject</b> . The default is <b>Allow</b> .                                                                                                                                                                                                                                                                                                                                     |
| IP Protocol | The options are <b>IPv4</b> , <b>IPv6</b> , and <b>IPv4_IPv6</b> . The default is <b>IPv4_IPv6</b> . To access this property, click the <b>Advanced Settings</b> icon.                                                                                                                                                                                                                                                                                  |
| Direction   | The options are <b>In</b> , <b>Out</b> , and <b>In/Out</b> . The default is <b>In/Out</b> . This field refers to the direction of traffic from the point of view of the destination object. <b>In</b> means that only traffic to the object is checked, <b>Out</b> means that only traffic from the object is checked, and <b>In/Out</b> means traffic in both directions is checked. To access this property, click the <b>Advanced Settings</b> icon. |

Table 20-1. Properties of a Firewall Rule (continued)

| Property        | Description                                                                                                        |
|-----------------|--------------------------------------------------------------------------------------------------------------------|
| Rule Tags       | Tags that have been added to the rule. To access this property, click the <b>Advanced Settings</b> icon.           |
| Flow Statistics | Read-only field that displays the byte, packet count, and sessions. To access this property, click the graph icon. |

**Note** If SpoofGuard is not enabled, automatically discovered address bindings cannot be guaranteed to be trustworthy because a malicious virtual machine can claim the address of another virtual machine. SpoofGuard, if enabled, verifies each discovered binding so that only approved bindings are presented.

## Add a Firewall Rule in Manager Mode

A firewall is a network security system that monitors and controls the incoming and outgoing network traffic based on predetermined firewall rules.

Firewall rules are added at the NSX Manager scope. Using the Applied To field, you can then narrow down the scope at which you want to apply the rule. You can add multiple objects at the source and destination levels for each rule, which helps reduce the total number of firewall rules to be added.

**Note** By default, a rule matches on the default of any source, destination, and service rule elements, matching all interfaces and traffic directions. If you want to restrict the effect of the rule to particular interfaces or traffic directions, you must specify the restriction in the rule.

### Prerequisites

- To use a group of addresses, first manually associate the IP and MAC address of each VM with their logical switch.
- Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 Select **Security > Distributed Firewall**.
- 2 Click the **General** tab for L3 rules or the **Ethernet** tab for L2 rules.
- 3 Click an existing section or rule.

- 4 Click the menu icon in the first column of a rule and select **Add Rule Above** or **Add Rule Below**.

A new row appears to define a firewall rule.

**Note** For any traffic attempting to pass through the firewall, the packet information is subjected to the rules in the order shown in the Rules table, beginning at the top and proceeding to the default rules at the bottom. In some cases, the order of precedence of two or more rules might be important in determining the disposition of a packet.

- 5 In the **Name** column, enter the rule name.
- 6 In the **Source** column, click the edit icon and select the source of the rule. The source will match any if not defined.

| Option            | Description                                                                                                                                        |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| IP Address es     | Enter multiple IP or MAC addresses in a comma-separated list. The list can contain up to 255 characters. Both IPv4 and IPv6 formats are supported. |
| Container Objects | The available objects are IP Set, Logical Port, Logical Switch, and NS Group. Select the objects and click <b>OK</b> .                             |

- 7 In the **Destination** column, click the edit icon and select the destination. The destination will match any if not defined.

| Option            | Description                                                                                                                                                |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IP Address es     | You can enter multiple IP or MAC addresses in a comma-separated list. The list can contain up to 255 characters. Both IPv4 and IPv6 formats are supported. |
| Container Objects | The available objects are IP Set, Logical Port, Logical Switch, and NS Group. Select the objects and click <b>OK</b> .                                     |

- 8 In the **Service** column, click the edit icon and select services. The service will match any if not defined.
- 9 To select a predefined service, select one of more available services.
- 10 To define a new service, click the **Raw Port-Protocol** tab and click **Add..**

| Option          | Description                                                                                                                  |
|-----------------|------------------------------------------------------------------------------------------------------------------------------|
| Type of Service | <ul style="list-style-type: none"> <li>■ ALG</li> <li>■ ICMP</li> <li>■ IGMP</li> <li>■ IP</li> <li>■ L4 Port Set</li> </ul> |
| Protocol        | Select one of the available protocols.                                                                                       |

| Option            | Description                  |
|-------------------|------------------------------|
| Source Ports      | Enter the source port.       |
| Destination Ports | Select the destination port. |

11 In the **Applied To** column, click the edit icon and select objects.

12 In the **Log** column, set the logging option.

Logs are in the `/var/log/dfwpktlogs.log` file on ESXi and KVM hosts. Enabling logging can affect performance.

13 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.                                                                                                                                                                                                                                 |
| 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. |

14 Click the **Advanced Settings** icon to specify IP protocol, direction, rule tags, and comments.

15 Click **Publish**.

## Delete a Firewall Rule in Manager Mode

A firewall is a network security system that monitors and controls the incoming and outgoing network traffic based on predetermined firewall rules. Custom defined rules can be added and deleted.

### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 Select **Security > Distributed Firewall**.
- 2 Click the **General** tab for L3 rules or the **Ethernet** tab for L2 rules.
- 3 Click the menu icon in the first column of the rule and select **Delete Rule**.

#### 4 Click **Publish**.

### Edit the Default Distributed Firewall Rule in Manager Mode

You can edit the default firewall settings that apply to traffic that does not match any of the user-defined firewall rules.

The default firewall rules apply to traffic that does not match any of the user-defined firewall rules. The default Layer 3 rule is under the **General** tab and the default Layer 2 rule is under the **Ethernet** tab.

The default firewall rules allow all L3 and L2 traffic to pass through all prepared clusters in your infrastructure. The default rule is always at the bottom of the rules table and cannot be deleted. However, you can change the **Action** element of the rule from **Allow** to **Drop** or **Reject**, and indicate whether traffic for that rule should be logged.

The default Layer 3 firewall rule applies to all traffic, including DHCP. If you change the **Action** to **Drop** or **Reject**, DHCP traffic will be blocked. You will need to create a rule to allow DHCP traffic.

#### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

#### Procedure

- 1 Select **Security > Distributed Firewall**.
- 2 Click the **General** tab for L3 rules or the **Ethernet** tab for L2 rules.
- 3 In the **Name** column, enter a new name.
- 4 In the **Action** column, select one of the options.
  - **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.
  - **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.

- 5 In the **Log**, enable or disable logging.

Enabling logging can affect performance.

- 6 Click **Publish**.

## Change the Order of a Firewall Rule in Manager Mode

Rules are processed in top-to-bottom ordering. You can change the order of the rules in the list.

For any traffic attempting to pass through the firewall, the packet information is subjected to the rules in the order shown in the rules table, beginning at the top and proceeding to the default rules at the bottom. In some cases, the order of precedence of two or more rules might be important in determining the traffic flow.

You can move a custom rule up or down in the table. The default rule is always at the bottom of the table and cannot be moved.

### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 Select **Security > Distributed Firewall**.
- 2 Click the **General** tab for L3 rules or the **Ethernet** tab for L2 rules.
- 3 Select the rule and click the **Move Up** or **Move Down** icon on the menu bar.
- 4 Click **Publish**.

## Filter Firewall Rules in Manager Mode

When you navigate to the firewall section, initially all the rules are displayed. You can apply a filter to control what is displayed so that you see only a subset of the rules. This can make it easier to manage the rules.

### Prerequisites

Verify that **Manager** mode is selected in the NSX Manager user interface. See [Chapter 1 NSX Manager](#). If you do not see the **Policy** and **Manager** mode buttons, see [Configure User Interface Settings](#).

### Procedure

- 1 Select **Security > Distributed Firewall**.
- 2 Click the **General** tab for L3 rules or the **Ethernet** tab for L2 rules.

- 3 In the search text field on the right side of the menu bar, select an object or enter the beginning characters of an object's name to narrow down the list of objects to select.

After you select an object, the filter is applied and the list of rules is updated, showing only rules that contain the object in any of the following columns:

- Sources
- Destinations
- Applied To
- Services

- 4 To remove the filter, delete the object name from the text field.

# Backing Up and Restoring NSX Manager

# 21

If an NSX Manager or a Global Manager appliance becomes inoperable, or if you want to restore your environment to a previous state, you can restore from a backup. While the appliance is inoperable, the data plane is not affected, but you cannot make configuration changes.

You can restore an NSX-T Data Center configuration back to the state that is captured in any of the backups.

When restoring a backup, you must restore to new appliances running the same version of NSX-T Data Center as the appliances that were backed up.

There are two backup methods:

- **Recurring:** Recurring backups run based on a customizable schedule and ensure that you have up-to-date backups. You can also trigger a recurring backup for configuration changes.

When you set up recurring backups, the system backs up the inventory every five minutes. Inventory backups provide the latest inventory updates, such as the addition or removal of a Transport Node, to the restore.

Inventory backups do not get collected for Global Manager.

- **Manual:** You manually run the backup at any time.

Manual backups do not include inventory backups.

---

**Important** Do not use snapshots to back up NSX-T Data Center appliances. See [Disable Snapshots on NSX-T Data Center Appliances](#) in the *NSX-T Data Center Installation Guide* for more information and instructions.

---

Read the following topics next:

- [Configure Backups](#)
- [Remove Old Backups](#)
- [Restore a Backup](#)
- [Certificate Management after Restore](#)



## Configure Backups

Before backups can occur, you must configure a backup file server. After a backup file server is configured, you can start a backup at any time, or schedule recurring backups.

### Prerequisites

- Verify that the SFTP server is running supported OS and SFTP software. The following table displays the supported and tested software for backup, although other software versions might work.

| Currently supported OS | Specifically tested version  | SFTP software version     |
|------------------------|------------------------------|---------------------------|
| CentOS                 | 7.7                          | OpenSSH_7.4p1             |
| RHEL                   | 7.7                          | OpenSSH_7.4p1             |
| Ubuntu                 | 18.04                        | OpenSSH_7.6p1             |
| Windows                | Windows Server 2019 Standard | OpenSSH_for_Windows_7.7p1 |

- Verify that the SFTP server is ready for use and is running SSH and SFTP, using the following commands:

- `$ ssh backup_user@sftp_server`

- `$ sftp backup_user@sftp_server`

- Ensure that the directory path exists where you want to store your backups. You cannot use the root directory (/).
- If you have multiple NSX-T Data Center deployments, you must use a different directory for storing the backup of each deployment.
- You can take backups using either the IP address or the FQDN of the NSX Manager or Global Manager appliance:
  - If you are using the IP address for backup and restore, do not publish the appliance's FQDN.
  - If you are using FQDN for backup and restore, you must configure and publish the FQDN before starting the backup. Backup and restore only support lowercase FQDN.

Use this API to publish the NSX Manager or Global Manager FQDN.

Example request:

```
PUT https://<nsx-mgr OR global-mgr>/api/v1/configs/management

{
 "publish_fqdns": true,
 "_revision": 0
}
```

See the *NSX-T Data Center API Guide* for API details.

**Procedure**

- 1 From a browser, log in with admin privileges to the NSX Manager or Global Manager at `https://<manager-ip-address>`.
- 2 Select **System > Backup & Restore**.
- 3 Click **Edit** under the **SFTP Server** label to configure your SFTP server.
- 4 Enter the IP address or FQDN of the backup file server.
- 5 Change the default port if necessary. The default port is 22.
- 6 The protocol text box is already filled in.

SFTP is the only supported protocol.

- 7 In the **Directory Path** text box, enter the absolute directory path where the backups will be stored.

The directory must already exist and cannot be the root directory (/). Avoid using path drive letters or spaces in directory names; they are not supported. If the backup file server is a Windows machine, you must use the forward slash when you specify the destination directory. For example, if the backup directory on the Windows machine is `c:\SFTP_Root\backup`, specify `/SFTP_Root/backup` as the destination directory.

The path to the backup directory can contain only the following characters: alphanumeric (a-z, A-Z, 0-9), underscore ( \_ ), plus and minus sign ( + - ), tilde and percent sign ( ~ % ), forward slash ( / ), and period ( . ).

The backup process generates a name for the backup file that can be quite long. On a Windows server, the length of the full path name of the backup file can exceed the limit set by Windows and cause backups to fail. To avoid this issue, see the KB article <https://kb.vmware.com/s/article/76528>.

- 8 Enter the user name and password required to log in to the backup file server.  
The first time you configure a file server, you must provide a password. Subsequently, if you reconfigure the file server, and the server IP or FQDN, port, and user name are the same, you do not need to enter the password again.
- 9 You can leave the SSH Fingerprint blank and accept or reject the fingerprint provided by the server after you click **Save** in a later step. If necessary, you can retrieve the SSH fingerprint by using this API: `POST /api/v1/cluster/backups?action=retrieve_ssh_fingerprint`. Note that only SHA256 hashed ECDSA (256 bit) host key is accepted as a fingerprint.
- 10 Enter a passphrase.

---

**Important** You will need this passphrase to restore a backup. If you forget the passphrase, you cannot restore any backups.

---

## 11 Click **Edit** under the **Schedule** label.

You can schedule recurring backups. You can also trigger backups for configuration changes. You can select both options for recurring backups. When you set up recurring backups, the system automatically backs up the inventory if there is a change in inventory, such as the addition or removal of a Transport Node. This feature is not available for manual backups.

Inventory backups do not get collected for Global Manager.

To enable recurring backups:

- a Click the **Recurring Backup** toggle.
- b Click **Weekly** and set the days and time of the backup, or click **Interval** and set the interval between backups.
- c Enabling the **Detect NSX configuration change** option will trigger an unscheduled full configuration backup when it detects any runtime or non-configuration related changes, or any change in user configuration. For Global Manager, this setting triggers backup if any changes in the database are detected, such as the addition or removal of a Local Manager or Tier-0 gateway or DFW policy.

You can specify a time interval for detecting database configuration changes. The valid range is 5 minutes to 1,440 minutes (24 hours). This option can potentially generate a large number of backups. Use it with caution.

## 12 Click **Save**.

### Results

After you configure a backup file server, you can click **Backup Now** to manually start a backup at any time. Automatic backups run as scheduled.

You see a progress bar of your in-progress backup.

When the manual or scheduled backup completes, it is listed in the Backup History section of the page. The **Last Backup Status** label indicates whether the backup was successful and lists the timestamp, node, and cluster details of the appliance backed up. If the backup fails, you can see an error message.

If you need to see a list of available backups but do not have access to an NSX Manager or Global Manager appliance see [Listing Available Backups](#) for details.

## Remove Old Backups

Backups can accumulate on the backup file server and consume a large amount of storage. You can run a script that comes with NSX-T Data Center to automatically delete old backups.

You can find the Python script `nsx_backup_cleaner.py` in the directory `/var/vmware/nsx/file-store` on NSX Manager. You must log in as root to access this file. Typically, you schedule a job on the backup file server to run this script periodically to clean up old backups. The following usage information describes how to run the script:

```
nsx_backup_cleaner.py -d backup_dir [-k 1] [-l 5] [-h]
Or
nsx_backup_cleaner.py --dir backup_dir [--retention-period 1] [--min-count 5] [--help]
```

Required parameters:

- d/--dir: Backup root directory
- k/--retention-period: Number of days need to retain a backup file

Optional parameters:

- l/--min-count: Minimum number of backup files to be kept, default value is 100
- h/--help: Display help message

The age of a backup is calculated as the difference between the backup's timestamp and the time the script is run. If this value is larger than the retention period, the backup is deleted if there are more backups on the disk than the minimum number of backups.

For more information about setting up the script to run periodically on a Linux or Windows server, see the comments at the beginning of the script.

## Restore a Backup

Restoring a backup restores the state of the network at the time of the backup. In addition, the configurations maintained by NSX Manager or Global Manager appliances are also restored. For NSX Manager, any changes, such as adding or deleting nodes, that were made to the fabric since the backup was taken, are reconciled.

---

**Note** DNS entries (name servers and search domains) are not retained when you restore from a backup.

---

You must restore the backup to a new NSX Manager or Global Manager appliance.

If you had a cluster of the NSX Manager appliance when the backup was taken, the restore process restores one node first and then prompts you to add the other nodes. You can add the other nodes during the restore process or after the first node is restored.

If you had a cluster of Global Manager appliances, you can only restore one node using the restore process. You must create the cluster after the restore of the first node completes.

---

**Important** If any nodes in the appliance cluster are still available, you must power them off before you start the restore.

---

### Prerequisites

- Verify that you have the login credentials for the backup file server.

- Verify that you have the SSH fingerprint of the backup file server. Only SHA256 hashed ECDSA (256 bit) host key is accepted as a fingerprint.
- Verify that you have the passphrase of the backup file.
- Identify which backup you want to restore by following the procedure in [Listing Available Backups](#). Take note of the IP or FQDN of the NSX-T Data Center appliance that took the backup.
- If you had an NSX Intelligence appliance installed on the NSX Manager when the backup was taken, you must upload the OVA file of that same version of the NSX Intelligence appliance on the new NSX Manager before you restore the backup. See *Installing and Upgrading VMware NSX Intelligence* for the installation information.

### Procedure

- 1 If any nodes in the appliance cluster that you are restoring are online, power them off.
- 2 Install one new appliance node on which to restore the backup.
  - If the backup listing for the backup you are restoring contains an IP address, you must deploy the new NSX Manager or Global Manager node with the same IP address. Do not configure the node to publish its FQDN.
  - If the backup listing for the backup you are restoring contains an FQDN, you must configure the new appliance node with this FQDN and publish the FQDN. Only lowercase FQDN is supported for backup and restore.

---

**Note** Until the FQDN is configured and published, the Restore button for the backup is disabled in the newly deployed NSX Manager or Global Manager UI.

---

Use this API to publish the NSX Manager or Global Manager FQDN.

Example request:

```
PUT https://<nsx-mgr OR global-mgr>/api/v1/configs/management

{
 "publish_fqdns": true,
 "_revision": 0
}
```

See the *NSX-T Data Center API Guide* for API details.

In addition, if the new manager node has a different IP address than the original one, you must update the DNS server's forward and reverse lookup entries for the manager node with the new IP address.

After the new manager node is running and online, you can proceed with the restore.

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

**4** Select **System > Backup & Restore**.**5** To configure the backup file server, click **Edit**.

Do not configure automatic backup if you are going to perform a restore.

**6** Enter the IP address or FQDN.**7** Change the port number, if necessary.

The default is 22.

**8** To log in to the server, enter the user name and password.**9** In the **Destination Directory** text box, enter the absolute directory path where the backups are stored.

The path to the backup directory can contain only the following characters: alphanumeric (a-z, A-Z, 0-9), underscore ( \_ ), plus and minus sign ( + - ), tilde and percent sign ( ~ % ), forward slash ( / ), and period ( . ).

Avoid using path drive letters or spaces in directory names; they are not supported. If the backup file server is a Windows machine, you must use the forward slash when you specify the destination directory. For example, if the backup directory on the Windows machine is c:\SFTP\_Root\backup, specify /SFTP\_Root/backup as the destination directory.

**10** Enter the passphrase that was used to encrypt the backup data.**11** You can leave the SSH Fingerprint blank and accept or reject the fingerprint provided by the server after you click **Save** in a later step. If necessary, you can retrieve the SSH fingerprint by using this API: `POST /api/v1/cluster/backups?action=retrieve_ssh_fingerprint`.**12** Click **Save**.**13** Select a backup.**14** Click **Restore**.**15** The restore process prompts you to take action, if necessary, as it progresses.

---

**Note** If you are restoring a Global Manager appliance, the following steps do not appear. After restoring the first Global Manager node, you must manually join the other nodes to form the cluster.

---

- a Confirm CM/VC Connectivity: If you want to restore existing compute managers, ensure that they are registered with the new NSX Manager node and available during the restore process.
- b If you have deleted or added fabric nodes or transport nodes, you are prompted to take certain actions, for example, log in to a node and run a script. If you have created a logical switch or segment since the backup, the logical switch or segment will not appear after the restore.

- c If the backup has information about a manager cluster, you are prompted to add other nodes. If you decide not to add nodes, you can still proceed with the restore and manually add other nodes to form the cluster after the restore of this node completes.
- d If there are fabric nodes that did not discover the new manager node, you are provided a list of them.

A progress bar displays the status of the restore operation noting the step the restore process is on. During the restore process, services on the manager appliance get restarted and the control plane becomes unavailable until restore completes.

After the restore operation is finished, the **Restore Complete** screen shows the result of the restore, the timestamp of the backup file, and the start and end time of the restore operation. Any segments created after the backup was taken are not restored.

If the restore fails, the screen displays the step where the failure occurred, for example, `Current Step: Restoring Cluster (DB)` OR `Current Step: Restoring Node`. If either cluster restore or node restore fails, the error might be transient. In that case, there is no need to click **Retry**. You can restart or reboot the manager and the restore continues.

You can also determine if there was a cluster restore or node restore failure by selecting the log files. Run `get log-file syslog` to view the system log file and search for the strings `Cluster restore failed` and `Node restore failed`.

To restart the manager, run the `restart service manager` command.

To reboot the manager, run the `reboot` command.

---

**Note** If you added a compute manager after the backup, and you try to add the compute manager again after the restore, you get an error message indicating that registration failed. Click the **Resolve** button to resolve the error and successfully add the compute manager. For more information, see [Add a Compute Manager](#), step 4. If you want to remove information about NSX-T Data Center that is stored in a vCenter Server, follow the steps in [Remove NSX-T Data Center Extension from vCenter Server](#)

If the vCenter Server was registered with custom ports in the backup, you must manually open all the custom ports on the restored manager appliances.

---

- 16 If you have only one node deployed, after the restored manager node is up and functional, you can deploy additional nodes to form a cluster.

See the *NSX-T Data Center Installation Guide* for instructions.

- 17 If you had other manager cluster VMs that you powered down in Step 1, delete them after the new manager cluster is deployed.

## Listing Available Backups

The backup file server stores backups from all the NSX Manager or Global Manager nodes. To get the list of backups so that you can find the one you want to restore, you must run the `get_backup_timestamps.sh` script.

The script can be found on each NSX Manager or Global Manager appliance at `/var/vmware/nsx/file-store/get_backup_timestamps.sh`. You can run this script on any Linux machine or NSX-T Data Center appliance. As a best practice, copy this script after installing NSX-T Data Center to a machine that is not an NSX Manager or Global Manager so that you can run this script even if all the NSX Manager or Global Manager nodes become inaccessible. If you need to restore a backup but have no access to this script, you can install a new NSX Manager or Global Manager node and run the script there.

You can copy the script to another machine or to the backup file server by logging in to the NSX Manager or Global Manager as `admin` and running a CLI command. For example:

```
nsxmgr-1> copy file get_backup_timestamps.sh url scp://admin@server1/tmp/
admin@server's password:
nsxmgr-1>
```

The script is interactive and prompts you for the information that you specified when you configured the backup file server. You can specify the number of backups to display. Each backup is listed with a timestamp, the NSX Manager or Global Manager node's IP address or FQDN if the NSX Manager or Global Manager node is set up to publish its FQDN, and the node ID. For example,

```
admin@host1:/home/admin# ./get_backup_timestamps.sh
Enter file server ip:
10.10.10.20
Enter port:
22
Enter directory path:
/home/nsx/backups
Enter number of latest backup or press Enter to list all backups:

root@10.10.10.20's password:
Latest backups:
[Backup timestamp; IP address/FQDN; Node id]
2019-01-22;09:16:43 nsxmgr.example.com 41893642-597b-915f-5117-7da576df4ff2
2019-01-22;09:14:42 nsxmgr.example.com 41893642-597b-915f-5117-7da576df4ff2
2019-01-22;09:13:30 nsxmgr.example.com 41893642-597b-915f-5117-7da576df4ff2
2019-01-22;09:01:52 10.10.10.77 35163642-6623-8f6d-7af0-52e03f16faed
2019-01-22;09:00:33 10.10.10.77 35163642-6623-8f6d-7af0-52e03f16faed
```

## Certificate Management after Restore

After restoring your NSX Manager appliances, certificates in the system get into an inconsistent state and you must update all self-signed or CA-signed certificates.

See [Chapter 19 Certificates](#) for more information on the type of certificates used in NSX-T Data Center and for instructions on updating them.

If you are using NSX-T Data Center version 3.0.1 or later, after you restore the first NSX Manager node, certificates are applied on this restored node, however, these certificates are not applied to the other nodes that are installed to form the restored NSX Manager cluster.



If you are using NSX-T Data Center version 3.0.0, none of the nodes have the original certificates applied and you must restore certificates manually for each node.

Follow these steps to update certificates after you complete the restore process:

- 1 If you are using NSX-T Data Center version 3.0.1 or later, update tomcat certificates on the two nodes installed and joined with the restored NSX Manager node to form a three-node cluster.

If you are using NSX-T Data Center version 3.0.0, update tomcat certificates for all of the NSX Manager nodes, including the one that was restored.

Use the following POST request to bring the nodes back to the same state as the backed-up cluster.

```
POST https://<nsx-mgr>/api/v1/node/services/http?
action=apply_certificate&certificate_id=<cert-id>
```

The certificate ID corresponds to the ID of the tomcat certificate that was in use on the original setup.

- 2 Verify the certificates by running the following GET request and confirm cluster stability.

```
GET https://<nsx-mgr>/api/v1/trust-management/certificates
```

You may need to change the configuration of the appliances you've installed, for example, adding licenses, certificates, and changing passwords. There are also routine maintenance tasks that you should perform, including running backups. Additionally, there are tools to help you find information about the appliances that are part of the NSX-T Data Center infrastructure and the logical networks created by NSX-T Data Center, including remote system logging, traceflow, and port connections.

Read the following topics next:

- [View the Usage and Capacity of Categories of Objects](#)
- [Configure User Interface Settings](#)
- [Configure a Node Profile](#)
- [Checking the Realized State of a Configuration Change](#)
- [View Network Topology](#)
- [Search for Objects](#)
- [Filter by Object Attributes](#)
- [Add a Compute Manager](#)
- [Add an Active Directory](#)
- [Add an LDAP Server](#)
- [Synchronize Active Directory](#)
- [Remove NSX-T Data Center Extension from vCenter Server](#)
- [Managing the NSX Manager Cluster](#)
- [Replacing an NSX Edge Transport Node in an NSX Edge Cluster](#)
- [Managing Resource Reservations for an Edge VM Appliance](#)
- [Adding and Removing an ESXi Host Transport Node to and from vCenter Servers](#)
- [Changing the Distributed Router Interfaces' MAC Address](#)
- [Configuring Appliances](#)
- [Add a License Key and Generate a License Usage Report](#)

- [Compliance-Based Configuration](#)
- [Collect Support Bundles](#)
- [Log Messages and Error Codes](#)
- [Customer Experience Improvement Program](#)
- [Find the SSH Fingerprint of a Remote Server](#)
- [Configuring an External Load Balancer](#)
- [Configure Proxy Settings](#)
- [View Container-Related Information](#)

## View the Usage and Capacity of Categories of Objects

You can view the usage and capacity of various categories of manager objects. You can also set alerts to let you easily see when certain thresholds in usage are reached.

This feature is available in manager mode only. To see the usage and capacity of different categories of objects, click one of the following tabs:

- **Networking > Network Overview > Capacity**
- **Security > Security Overview > Capacity**
- **Inventory > Inventory Overview > Capacity**
- **System > System Overview > Capacity**

You can also navigate to **Plan & Troubleshoot > Consolidated Capacity** to see all the object categories on one page.

On each capacity page, for each category of objects, the following information is displayed:

- **Maximum capacity** - This value is based on the capacity of a large appliance.
- **Current inventory** - The number of objects that have been successfully created or configured. A color-coded bar is displayed to indicate the usage percentage. If usage is below the minimum capacity threshold, the color is green. If usage is at or above the minimum capacity threshold but below the maximum capacity threshold, the color is orange. If usage is at or above the maximum capacity threshold, the color is red.
- **Minimum Capacity Threshold** - This is the usage level at which the usage bar mentioned above will show an orange color. You can change this value. The default is 70%.
- **Maximum Capacity Threshold** - This is the usage level at which the usage bar mentioned above will show a red color. You can change this value. The default is 100%.

When you change the warning alert or critical alert value, you can click **Revert** to go back to the last saved value. You can click **Reset Values** to restore the default values for all the object categories.

The networking capacity page shows the following object categories:

- Tier-0 logical routers
- Tier-1 logical routers
- Prefix lists
- System-wide NAT rules
- DHCP server instances
- System-wide DHCP ranges and pools
- Tier-1 logical routers with NAT enabled
- Logical switches
- System-wide logical switch ports

The security capacity page shows the following object categories:

- System-wide endpoint protection-enabled hosts
- System-wide endpoint protection-enabled virtual machines
- Active Directory groups
- Active Directory domains
- Distributed firewall rules
- System-wide firewall rules
- System-wide firewall sections
- Distributed firewall sections

The inventory capacity page shows the following object categories:

- Groups
- IP sets
- Groups based on IP sets
- vSphere clusters
- Hypervisor hosts


The system capacity page shows the following object categories:

- Edge clusters
- System-wide edge nodes

## Configure User Interface Settings

There are two possible modes in the NSX Manager web interface: Policy and Manager. You can control which mode is default, and if users can switch between them using the user interface mode buttons.

If present, you can use the **Policy** and **Manager** buttons to switch between the Policy and Manager modes. Switching modes controls which menus items are available to you.



- By default, if your environment contains only objects created through Policy mode, your user interface is in Policy mode and you do not see the **Policy** and **Manager** buttons.
- By default, if your environment contains any objects created through Manager mode, you see the **Policy** and **Manager** buttons in the top-right corner.

You can use the User Interface Settings to modify these defaults.

See [Chapter 1 NSX Manager](#) for more information about the modes.

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Navigate to **System > User Interface Settings** and click **Edit**.
- 3 Modify the user interface settings: **Toggle Visibility** and **Default Mode**.

| Toggle Visibility                               | Description                                                                                                       |
|-------------------------------------------------|-------------------------------------------------------------------------------------------------------------------|
| Visible to All Users                            | If Manager mode objects are present, the mode buttons are visible to all users.                                   |
| Visible to Users with the Enterprise Admin Role | If Manager mode objects are present, the mode buttons are visible to users with the <b>Enterprise Admin</b> role. |
| Hidden from All Users                           | Even if Manager mode objects are present, the mode buttons are hidden from all users.                             |

**Default Mode** can be set to Policy or Manager.

## Configure a Node Profile

You can configure settings such as time zone, NTP servers, SNMP, and syslog servers to apply to all NSX Manager and Edge nodes. In addition to NSX Manager and Edge nodes, the SNMP configuration is applied to the VMware SNMP agent on all KVM hypervisors.

In this release, only one node profile is supported. This profile represents a collection of time zone, NTP servers, SNMP configuration and syslog servers. By default, the node profile is applied to all nodes, unless the node is configured to not accept such configuration from the NSX Manager. To prevent a node from accepting the node profile, use the CLI command `set node central-config disabled` on that node.

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **System > Fabric > Profiles**.
- 3 Click the **Node Profiles** tab.
- 4 Click **All NSX Nodes** in the **Name** column.
- 5 Click **Edit** to configure the time zone and NTP servers.
- 6 In the **Syslog Servers** section, click **Add** to add a Syslog server.
  - a Enter the FQDN or IP address of the Syslog server.
  - b Specify a port number.
  - c Select a protocol.

The available protocols are **TCP**, **UDP**, and **LI** (Log Insight).
  - d Select a log level.

The available levels are **Emergency**, **Alert**, **Critical**, **Error**, **Warning**, **Notice**, **Information**, and **Debug**.
- 7 In the **SNMP Polling** section, under **v2c**, click **Add** to add an SNMPv2c community.
  - a Enter a name for the community.
  - b Enter a **Community String** value.

This value is used for authentication.
- 8 In the **SNMP Polling** section, under **v3**, click **Add** to add an SNMPv3 user.
  - a Enter a user name.
  - b Enter an authentication password.

You can click the icon on the right to show or hide the password.
  - c Enter a private password.

You can click the icon on the right to show or hide the password.
- 9 In the **SNMP Traps** section, under **v2c**, click **Add** to add an SNMPv2c trap configuration.
  - a Enter a FQDN or IP address.
  - b Specify a port number.

- c Enter a name for the community.
- d Enter a **Community String** value.

This value is used for authentication.

**10** In the **SNMP Traps** section, under **v3**, click **Add** to add an SNMPv3 trap configuration.

- a Enter a FQDN or IP address.
- b Specify a port number.
- c Enter a user name.

## Checking the Realized State of a Configuration Change

When you make a configuration change, NSX Manager typically sends a request to another component to implement the change. For some layer 3 entities, if you make the configuration change using the API, you can track the status of the request to see if the change is successfully implemented.

The configuration change that you initiate is called the desired state. The result of implementing the change is called the realized state. If NSX Manager implements the change successfully, the realized state will be the same as the desired state. If there is an error, the realized state will not be the same as the desired state.

For some layer 3 entities, when you call an API to make a configuration change, the response will include the parameter `request_id`. You can use the parameters `request_id` and the `entity_id` to make an API call to find out the status of the request.

This feature supports the following entities and APIs:

```
EdgeCluster
 POST /edge-clusters
 PUT /edge-clusters/<edge-cluster-id>
 DELETE /edge-clusters/<edge-cluster-id>
 POST /edge-clusters/<edge-cluster-id>?action=replace_transport_node

LogicalRouter
 POST /logical-routers
 PUT /logical-routers/<logical-router-id>
 DELETE /logical-routers/<logical-router-id>
 POST /logical-routers/<logical-router-id>?action=reprocess
 POST /logical-routers/<logical-router-id>?action=reallocate

LogicalRouterPort
 POST /logical-router-ports
 PUT /logical-router-ports/<logical-router-port-id>
 DELETE /logical-router-ports/<logical-router-port-id>

StaticRoute
 POST /logical-routers/<logical-router-id>/routing/static-routes
 PUT /logical-routers/<logical-router-id>/routing/static-routes/<static-route-id>
 DELETE /logical-routers/<logical-router-id>/routing/static-routes/<static-route-id>
```

```

BGPConfig
 PUT /logical-routers/<logical-router-id>/routing/bgp

BgpNeighbor
 POST /logical-routers/<logical-router-id>/routing/bgp/neighbors
 PUT /logical-routers/<logical-router-id>/routing/bgp/neighbors/<bgp-neighbor-id>
 DELETE /logical-routers/<logical-router-id>/routing/bgp/neighbors/<bgp-neighbor-id>
 POST /logical-routers/<logical-router-id>/routing/bgp/neighbors/<bgp-neighbor-id>

BGPCommunityList
 POST /logical-routers/<logical-router-id>/routing/bgp/community-lists
 PUT /logical-routers/<logical-router-id>/routing/bgp/community-lists/<community-list-id>
 DELETE /logical-routers/<logical-router-id>/routing/bgp/community-lists/<community-list-id>

AdvertisementConfig
 PUT /logical-routers/<logical-router-id>/routing/advertisement

AdvertiseRouteList
 PUT /logical-routers/<logical-router-id>/routing/advertisement/rules

NatRule
 POST /logical-routers/<logical-router-id>/nat/rules
 PUT /logical-routers/<logical-router-id>/nat/rules/<rule-id>
 DELETE /logical-routers/<logical-router-id>/nat/rules/<rule-id>

DhcpRelayService
 POST /dhcp/relays
 PUT /dhcp/relays/<relay-id>
 DELETE /dhcp/relays/<relay-id>

DhcpRelayProfile
 POST /dhcp/relay-profiles
 PUT /dhcp/relay-profiles/<relay-profile-id>
 DELETE /dhcp/relay-profiles/<relay-profile-id>

StaticHopBfdPeer
 POST /logical-routers/<logical-router-id>/routing/static-routes/bfd-peers
 PUT /logical-routers/<logical-router-id>/routing/static-routes/bfd-peers/<bfd-peers-id>
 DELETE /logical-routers/<logical-router-id>/routing/static-routes/bfd-peers/<bfd-peers-id>

IPPrefixList
 POST /logical-routers/<logical-router-id>/routing/ip-prefix-lists
 PUT /logical-routers/<logical-router-id>/routing/ip-prefix-lists/<ip-prefix-list-id>
 DELETE /logical-routers/<logical-router-id>/routing/ip-prefix-lists/<ip-prefix-list-id>

RouteMap
 POST /logical-routers/<logical-router-id>/routing/route-maps
 PUT /logical-routers/<logical-router-id>/routing/route-maps/<route-map-id>
 DELETE /logical-routers/<logical-router-id>/routing/route-maps/<route-map-id>

RedistributionConfig
 PUT /logical-routers/<logical-router-id>/routing/redistribution

RedistributionRuleList

```



```

 PUT /logical-routers/<logical-router-id>/routing/redistribution/rules

BfdConfig
 PUT /logical-routers/<logical-router-id>/routing/bfd-config

MplsConfig
 PUT /logical-routers/<logical-router-id>/routing/mps

RoutingGlobalConfig
 PUT /logical-routers/<logical-router-id>/routing

IPSecVPNIKEProfile
 POST /vpn/ipsec/ike-profiles
 PUT /vpn/ipsec/ike-profiles/<ike-profile-id>
 DELETE /vpn/ipsec/ike-profiles/<ike-profile-id>

IPSecVPNDPDProfile
 POST /vpn/ipsec/dpd-profiles
 PUT /vpn/ipsec/dpd-profiles/<dpd-profile-id>
 DELETE /vpn/ipsec/dpd-profiles/<dpd-profile-id>

IPSecVPNTunnelProfile
 POST /vpn/ipsec/tunnel-profiles
 PUT /vpn/ipsec/tunnel-profiles/<tunnel-profile-id>
 DELETE /vpn/ipsec/tunnel-profiles/<tunnel-profile-id>

IPSecVPNLocalEndpoint
 POST /vpn/ipsec/local-endpoints
 PUT /vpn/ipsec/local-endpoints/<local-endpoint-id>
 DELETE /vpn/ipsec/local-endpoints/<local-endpoint-id>

IPSecVPNPeerEndpoint
 POST /vpn/ipsec/peer-endpoints
 PUT /vpn/ipsec/peer-endpoints/<peer-endpoint-id>
 DELETE /vpn/ipsec/peer-endpoints/<peer-endpoint-id>

IPSecVPNService
 POST /vpn/ipsec/services
 PUT /vpn/ipsec/services/<service-id>
 DELETE /vpn/ipsec/services/<service-id>

IPSecVPNSession
 POST /vpn/ipsec/sessions
 PUT /vpn/ipsec/sessions/<session-id>
 DELETE /vpn/ipsec/sessions/<session-id>

DhcpServer
 POST /dhcp/servers
 PUT /dhcp/servers/<server-id>
 DELETE /dhcp/servers/<server-id>

DhcpStaticBinding
 POST /dhcp/servers/static-bindings
 PUT /dhcp/servers/<server-id>/static-bindings/<binding-id>
 DELETE /dhcp/servers/<server-id>/static-bindings/<binding-id>

```

```
DhcpIpPool
 POST /dhcp/servers/ip-pools
 PUT /dhcp/servers/<server-id>/ip-pools/<pool-id>
 DELETE /dhcp/servers/<server-id>/ip-pools/<pool-id>

DnsForwarder
 POST /dns/forwarders
 PUT /dns/forwarders/<forwarder-id>
 DELETE /dns/forwarders/<forwarder-id>
```

You can call the following APIs to get the realized states:

```
EdgeCluster
Request - GET /edge-clusters/<edge-cluster-id>/state?request_id=<request-id>
Response - An instance of EdgeClusterStateDto which will inherit ConfigurationState. If the
edge cluster is deleted then the state will be unknown and it will return the common entity
not found error.

LogicalRouter / All L3 Entities - All L3 entities can use this API to get realization state
Request - GET /logical-routers/<logical-router-id>/state?request_id=<request-id>
Response - An instance of LogicalRouterStateDto which will inherit ConfigurationState. Delete
operation of any entity other than logical router can be covered by getting the state of
logical router but if the logical router itself is deleted then the state will be unknown and
it will return the common entity not found error.

LogicalServiceRouterCluster - All L3 entities which are the part of services can use this API
to get the realization state
Request - GET /logical-routers/<logical-router-id>/service-cluster/state?request_id=<request-
id>
Response - An instance of LogicalServiceRouterClusterState which will inherit
ConfigurationState.

LogicalRouterPort / DhcpRelayService / DhcpRelayProfile
Request - GET /logical-router-ports/<logical-router-port-id>/state?request_id=<request-id>
Response - An instance of LogicalRouterPortStateDto which will inherit ConfigurationState.

IPSecVPNIKEProfile / IPSecVPNDPDProfile / IPSecVPNTunnelProfile / IPSecVPNLocalEndpoint /
IPSecVPNPeerEndpoint / IPSecVPNService / IPSecVPNSession
Request - GET /vpn/ipsec/sessions/<session-id>/state?request_id=<request-id>
Response - An instance of IPSecVPNSessionStateDto which will inherit ConfigurationState. If
the session is deleted then the state will be unknown and it will return the common entity
not found error. When IPSecVPNService is disabled, IKE itself is down and it does not
respond. It will return unknown state in such a case.

DhcpServer
Request - GET /dhcp/servers/<server-id>/state?request_id=<request-id>
Response - An instance of ConfigurationState.

DhcpStaticBinding
Request - GET /dhcp/servers/<server-id>/static-bindings/<binding-id>/state?
request_id=<request-id>
Response - An instance of ConfigurationState.

DhcpIpPool
```

```
Request - GET /dhcp/servers/<server-id>/ip-pools/<pool-id>/state?request_id=<request-id>
Response - An instance of ConfigurationState.
```

```
DnsForwarder
```

```
Request - GET /dns/forwarders/<forwarder-id>/state?request_id=<request-id>
Response - An instance of ConfigurationState.
```

For more information about the APIs, see the *NSX-T Data Center API Reference*.

## View Network Topology

View the network topology of your NSX-T Data Center environment to get an overview of the logical entities in your network. The graphical representation of the network topology is helpful when you are verifying your network configuration or troubleshooting errors.

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Networking > Network Topology**.
- 3 Navigate through the network topology to see more information:
  - Zoom in to view more details.
  - Point to an object to view its logical path in the network.
  - Click an object to display the details panel for that object.
  - Click **Export** on the tool bar to save the topology to a PDF file.
  - Apply filters to focus on specific objects. See [Filter by Object Attributes](#) for more details about filters.

## Search for Objects

You can search for objects using various criteria throughout the NSX-T Data Center inventory.

The search results are sorted by relevance and you can filter these results based on your search query.

---

**Note** If you have special characters in your search query that also function as operators, then you must add a leading backslash. The characters that function as operators are: +, -, =, &&, ||, <, >, !, (, ), {, }, [, ], ^, ", ~, ?, :, /, \.

---

### Procedure

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

- 2 On the homepage, enter a search pattern for an object or object type.

As you enter your search pattern, the search feature provides assistance by showing the applicable keywords.

| Search                                       | Search Query         |
|----------------------------------------------|----------------------|
| Objects with Logical as the name or property | Logical              |
| Exact logical switch name                    | display_name:LSP-301 |
| Names with special characters such as, !     | Logical\!            |

All the related search results are listed and grouped by resource type in different tabs.

You can click the tabs for specific search results for a resource type.

- 3 (Optional) In the search bar, click the save icon to save your refined search criteria.
- 4 In the search bar, click the  icon to open the advanced search column where you can refine your search.
- 5 Specify one or more criteria to refine your search.
  - Name
  - Resource Type
  - Description
  - ID
  - Created by
  - Modified by
  - Tags
  - Creation Date
  - Modified Date

You can also view your recent search results and saved search criteria.

- 6 (Optional) Click **Clear All** to reset your advanced search criteria.

## Filter by Object Attributes

When viewing objects in NSX Manager, you can filter the objects by one or more of their attributes. For example, when viewing details of Tier 0 gateways you can choose to filter by **Status** and view only those gateways that are **Down**.

The following types of filters are available:

- Predefined filters – A list of commonly used filters that you can apply to your objects.

- Text-based filter – A filter based on the attribute value that you enter. This filter is applicable only to the **Name**, **Tag**, **Path**, and **Description** attributes of the objects.
- Attribute-value pairs – An attribute drop-down menu that you can use to specify attribute-value pairs for filtering.

You can either use multiple attributes of an object or multiple values of a single attribute to filter objects. The AND operator is applied when you select multiple attributes whereas the OR operator is used when you specify multiple values of a single attribute.

#### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Navigate to the tab that displays the objects you want to view.
- 3 Specify the attributes that you want to use to filter the objects.

- Click  and select from a list of predefined filters.
- Enter a value for the **Name**, **Tag**, **Path**, or **Description** attributes.
- Select an attribute from the drop-down menu and specify its value. For example, **Status: Down**

Objects satisfying your filter criteria are displayed.

- 4 (Optional) Click **Clear** to reset your filters.

## Add a Compute Manager

A compute manager, for example, vCenter Server, is an application that manages resources such as hosts and VMs.

NSX-T Data Center polls compute managers to collect cluster information from vCenter Server.

When you add a vCenter Server compute manager, you must provide a vCenter Server user's credentials. You can provide the vCenter Server administrator's credentials, or create a role and a user specifically for NSX-T Data Center and provide this user's credentials. This role must have the following vCenter Server privileges:

---

Extension.Register extension

---

Extension.Unregister extension

---

Extension.Update extension

---

Sessions.Message

---

Sessions.Validate session

---

Sessions.View and stop sessions

---

Host.Configuration.Maintenance

---

|                                                   |
|---------------------------------------------------|
| Host.Configuration.NetworkConfiguration           |
| Host.Local Operations.Create virtual machine      |
| Host.Local Operations.Delete virtual machine      |
| Host.Local Operations.Reconfigure virtual machine |
| Tasks                                             |
| Scheduled task                                    |
| Global.Cancel task                                |
| Permissions.Reassign role permissions             |
| Resource.Assign vApp to resource pool             |
| Resource.Assign virtual machine to resource pool  |
| Virtual Machine.Configuration                     |
| Virtual Machine.Guest Operations                  |
| Virtual Machine.Provisioning                      |
| Virtual Machine.Inventory                         |
| Network.Assign network                            |
| vApp                                              |

For more information about vCenter Server roles and privileges, see the *vSphere Security* document.

### Prerequisites

- Verify that you use the supported vSphere version. See [Supported vSphere version](#).
- IPv6 and IPv4 communication with vCenter Server.
- Verify that you use the recommended number of compute managers. See <https://configmax.vmware.com/home>.

---

**Note** NSX-T Data Center does not support the same vCenter Server to be registered with more than one NSX Manager.

---

- Use of custom ports such as HTTP or HTTPS on vCenter Server is not supported.

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **System > Fabric > Compute Managers > Add**.

### 3 Complete the compute manager details.

| Option                             | Description                                                                                                                                                                               |
|------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Name and Description</b>        | Type the name to identify the vCenter Server.<br>You can optionally describe any special details such as, the number of clusters in the vCenter Server.                                   |
| <b>FQDN or IP Address</b>          | Type the FQDN or IP address of the vCenter Server.                                                                                                                                        |
| <b>Type</b>                        | The default compute manager type is set to vCenter Server.                                                                                                                                |
| <b>HTTPS Port of Reverse Proxy</b> | The default port is 443. If you use another port, verify that the port is open on all the NSX Manager appliances.<br>Set the reverse proxy port to register the compute manager in NSX-T. |
| <b>Username and Password</b>       | Type the vCenter Server login credentials.                                                                                                                                                |
| <b>SHA-256 Thumbprint</b>          | Type the vCenter Server SHA-256 thumbprint algorithm value.                                                                                                                               |
| <b>Enable Trust</b>                | Supported only on vCenter Server 7.0 and later versions.<br>Enable this field to trust compute manager for authentication.                                                                |

If you left the thumbprint value blank, you are prompted to accept the server provided thumbprint.

After you accept the thumbprint, it takes a few seconds for NSX-T Data Center to discover and register the vCenter Server resources.

**Note** If the FQDN, IP, or thumbprint of the compute manager changes after registration, edit the computer manager and enter the new values.

### 4 If the progress icon changes from **In progress** to **Not registered**, perform the following steps to resolve the error.

- a Select the error message and click **Resolve**. One possible error message is the following:

```
Extension already registered at CM <vCenter Server name> with id <extension ID>
```

- b Enter the vCenter Server credentials and click **Resolve**.

If an existing registration exists, it will be replaced.

### Results

It takes some time to register the compute manager with vCenter Server and for the connection status to appear as **UP**.

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

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

---

**Note** After a vCenter Server (VC) compute manager is successfully added, it cannot be removed if you successfully performed any of the following actions:

- Transport nodes are prepared using VDS that is dependent on the VC.
- Service VMs deployed on a host or a cluster in the VC using NSX service insertion.
- You use the NSX Manager UI to deploy Edge VMs, NSX Intelligence VM, or NSX Manager nodes on a host or a cluster in the VC.

If you try to perform any of these actions and you encounter an error (for example, installation failed), you can remove the VC if you have not successfully performed any of the actions listed above.

If you have successfully prepared any transport node using VDS that is dependent on the VC or deployed any VM, you can remove the VC after you have done the following:

- Unprepare all transport nodes. If uninstalling a transport node fails, you must force delete the transport node.
- Undeploy all service VMs, any NSX Intelligence VM, all NSX Edge VMs and all NSX Manager nodes. The undeployment must be successful or in a failed state.
- If an NSX Manager cluster consists of nodes deployed from the VC (manual method) and nodes deployed from the NSX Manager UI, and you had to undeploy the manually deployed nodes, then you cannot remove the VC. To successfully remove the VC, ensure that you re-deploy an NSX Manager node from the VC.

This restriction applies to a fresh installation of NSX-T Data Center 3.0 as well as an upgrade.

---

## Add an Active Directory

Active Directory is used in creating user-based Identity Firewall rules.

Windows 2008 is not supported as an Active Directory server or RDSH Server OS.

You can register one or more Windows domains with an NSX Manager. NSX Manager gets group and user information, and the relationship between them from each domain that it is registered. NSX Manager also retrieves Active Directory (AD) credentials.

You can register an entire AD (Active Directory) domain to be used by IDFW (Identity Firewall), or you can synchronize a subset of a large domain. Once a domain is registered, NSX synchronizes all AD data required by IDFW.



Once the Active Directory is synced to the NSX Manager, you can create security groups based on user identity, and create identity-based firewall rules.

---

**Note** For Identity Firewall rule enforcement, Windows Time service should be **on** for all VMs using Active Directory. This ensures that the date and time is synchronized between Active Directory and VMs. AD group membership changes, including enabling and deleting users, do not immediately take effect for logged in users. For changes to take effect, users must log out and then log back in. AD administrator's should force a logout when group membership is modified. This behavior is a limitation of Active Directory.

---

#### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Navigate to **System > Identity Firewall AD > Active Directory**.
- 3 Click **Add Active Directory**.
- 4 Enter the name of the active directory.
- 5 Enter the **NetBios Name** and **Base Distinguished Name**.

To retrieve the netBIOS name for your domain, enter `nbtstat -n` in a command window on a Windows Workstation that is part of a domain, or on a domain controller. In the NetBIOS Local Name Table, the entry with a <00> prefix and type Group is the NetBIOS name.

A base distinguished name (Base DN) is needed to add an Active Directory domain. A Base DN is the starting point that an LDAP server uses when searching for users authentication within an Active Directory domain. For example, if your domain name is corp.local the DN for the Base DN for Active Directory would be "DC=corp,DC=local".

- 6 Set the **Delta Synchronization Interval** if necessary. A delta synchronization updates local AD objects that have changed since the last synchronization event.

Any changes made in Active Directory are NOT seen on NSX Manager until a delta or full synchronization has been performed.

- 7 Click **Save**.

## Add an LDAP Server

LDAP (Lightweight Directory Access Protocol) server configuration and functionality is only for use with Identity Firewall. LDAP provides a central place for authentication, meaning that when you configure a connection to your LDAP server, the user records are stored in your external LDAP server.

#### Prerequisites

The domain account must have AD read permission for all objects in the domain tree.

When there is a cluster of NSX Managers, all nodes need to be able to reach the LDAP server.

**Procedure**

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Navigate to **System > Identity Firewall AD > Active Directory**.
- 3 Select the **LDAP Server** tab.
- 4 Click **Add LDAP Server**.
- 5 Enter the **Host** name of the LDAP server.
- 6 Select the active directory the LDAP server is connected to from the **Connected to (Directory)** drop-down menu.
- 7 (Optional) Select the **protocol**: LDAP (unsecured) or LDAPS (secured).
- 8 If LDAPS was selected, select the SHA-256 Thumbprint suggested by NSX Manager, or enter a SHA-256 Thumbprint.
- 9 Enter the **port** number of the LDAP server.  
  
For local domain controllers, the default LDAP port 389 and LDAPS port 636 are used for the Active Directory sync, and should not be edited from the default values.
- 10 Enter the **username** and **password** of an Active Directory account with a minimum of read-only access to the Active Directory domain.
- 11 Click **Save**.
- 12 To verify that you can connect to the LDAP server, click **Test Connection**.

## Synchronize Active Directory

Active Directory objects can be used to create security groups based on user identity, and identity-based firewall rules.

---

**Note** Do not enable Distributed Intrusion Detection Service (IDS) in an environment that is using Distributed Load Balancer. NSX-T Data Center does not support using IDS with a Distributed Load Balancer.

---

To enable selective sync, use domain create/update API with selective sync enabled, and a list of selected Organization Units (OUs). When selective sync is enabled, NSX-T only synchronizes the AD data inside the selected OUs. During a selective delta sync, only the Active Directory data which is inside the selected OUs and has been created or changed since last sync are updated. If any directory-groups are removed from the selected OUs, they will not be updated during a selective delta sync. They will be updated during a full sync when all directory groups are updated. For more information, see the *NSX-T Data Center API Guide*.

If you use the API to manually end a full sync after it has begun, the sync stats will not be updated correctly.

**Note** IDFW relies on the security and integrity of the guest operating system. There are multiple methods for a malicious local administrator to spoof their identity to bypass firewall rules. User identity information is provided by the Guest Introspection Agent inside guest VMs. Security administrators must ensure that NSX Guest Introspection Agent is installed and running in each guest VM. Logged-in users should not have the privilege to remove or stop the agent.

#### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Navigate to **System > Identity Firewall AD > Active Directory**.
- 3 Click the three button menu icon next to the Active Directory that you want to synchronize, and select one of the following:

| Menu Item  | Description                                                                                                           |
|------------|-----------------------------------------------------------------------------------------------------------------------|
| Sync Delta | Perform a delta synchronization, where local AD objects that have changed since the last synchronization are updated. |
| Sync All   | Perform a full synchronization, where the local state of all AD objects is updated.                                   |

- 4 Click **View Sync Status** to see the current state of the Active Directory, the previous synchronization state, the synchronization status, and the last synchronization time.

## Remove NSX-T Data Center Extension from vCenter Server

When you add a compute manager, NSX Manager adds its identity as an extension in vCenter Server. If you remove the compute manager, the extension in vCenter Server will be removed automatically. If the extension is not removed for some reason, you can manually remove the extension with the following procedure.

#### Prerequisites

Enable access to the vCenter Server Managed Object Browser (MOB) by following the procedure in <https://kb.vmware.com/s/article/2042554>.

#### Procedure

- 1 Login to the MOB at `https://<vCenter Server hostname or IP address>/mob`.
- 2 Click the **content** link, which is the value for the **content** property in the Properties table.
- 3 Click the **ExtensionManager** link, which is the value for **extensionManager** property in the Properties table.

- 4 Click the **UnregisterExtension** link in the Methods table.
- 5 Enter `com.vmware.nsx.management.nsx` in the **value** text field.
- 6 Click the **Invoke Method** link on the right hand side of the page below the Parameters table.  
The method result says `void` but the extension will be removed.
- 7 To make sure the extension is removed, click the **FindExtension** method on the previous page and invoke it by entering the same value for the extension.  
The result should be `void`.

## Managing the NSX Manager Cluster

You can reboot an NSX Manager if it becomes inoperable. You can also change the IP address of an NSX Manager.

In a production environment, it is highly recommended that the NSX Manager cluster has three members to provide high availability. If you delete an NSX Manager and deploy a new one, the new NSX Manager can have the same or a different IP address.

---

**Note** The primary NSX Manager node is the node that you create first, before you create a manager cluster. This node cannot be deleted. After you deploy two more manager nodes from the primary manager node's UI to form a cluster, only the second and the third manager nodes have the option (from the gear icon) to be deleted. For information about removing and adding a manager node, see [Change the IP Address of an NSX Manager](#).

---

## View the Configuration and Status of the NSX Manager Cluster

You can view the configuration and status of the NSX Manager cluster from the NSX Manager UI. You can get additional information using the CLI.

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://nsx-manager-ip-address`.
- 2 Select **System > Overview**  
The status of the NSX Manager cluster is displayed.
- 3 To see additional information about the configuration, run the following CLI command:

```
manager1> get cluster config
Cluster Id: 18807edd-56d1-4107-b7b7-508d766a08e3
Cluster Configuration Version: 3
Number of nodes in the cluster: 3

Node UUID: 43cd0642-275c-af1d-fe46-1f5200f9e5f9
Node Status: JOINED
```

| ENTITY  |      |      | UUID | IP |
|---------|------|------|------|----|
| ADDRESS | PORT | FQDN |      |    |
|         |      |      |      |    |

|                                                 |      |                                   |                                      |    |
|-------------------------------------------------|------|-----------------------------------|--------------------------------------|----|
| HTTPS                                           |      |                                   | 5c8d01f1-f3ee-4f94-b517-a093d8fbfad3 |    |
| 10.160.71.225                                   | 443  | ychin-nsxmanager-ob-12065118-1-F5 |                                      |    |
| CONTROLLER                                      |      |                                   | 06fd0574-69c0-432e-a8af-53d140dbef8f |    |
| 10.160.71.225                                   | -    | ychin-nsxmanager-ob-12065118-1-F5 |                                      |    |
| CLUSTER_BOOT_MANAGER                            |      |                                   | da8d535e-7a0c-4dd8-8919-d88bdde006b8 |    |
| 10.160.71.225                                   | -    | ychin-nsxmanager-ob-12065118-1-F5 |                                      |    |
| DATASTORE                                       |      |                                   | 3c9c4ec1-afef-47bd-aadb-1ed6a5536bc4 |    |
| 10.160.71.225                                   | 9000 | ychin-nsxmanager-ob-12065118-1-F5 |                                      |    |
| MANAGER                                         |      |                                   | eb5e8922-23bd-4c3a-ae22-d13d9195a6bc |    |
| 10.160.71.225                                   | -    | ychin-nsxmanager-ob-12065118-1-F5 |                                      |    |
| POLICY                                          |      |                                   | f9da1039-08ad-4a20-bacc-5b91c5d67730 |    |
| 10.160.71.225                                   | -    | ychin-nsxmanager-ob-12065118-1-F5 |                                      |    |
| Node UUID: 8ebb0642-201e-6a5f-dd47-a1e38542e672 |      |                                   |                                      |    |
| Node Status: JOINED                             |      |                                   |                                      |    |
| ENTITY                                          |      |                                   | UUID                                 | IP |
| ADDRESS                                         | PORT | FQDN                              |                                      |    |
| HTTPS                                           |      |                                   | 3757f155-8a5d-4b53-828f-d67041d5a210 |    |
| 10.160.93.240                                   | 443  | ychin-nsxmanager-ob-12065118-2-F5 |                                      |    |
| CONTROLLER                                      |      |                                   | 7b1c9952-8738-4900-b68b-ca862aa4f6a9 |    |
| 10.160.93.240                                   | -    | ychin-nsxmanager-ob-12065118-2-F5 |                                      |    |
| CLUSTER_BOOT_MANAGER                            |      |                                   | b5e12db1-5e0d-4e33-a571-6ba258dceb2e |    |
| 10.160.93.240                                   | -    | ychin-nsxmanager-ob-12065118-2-F5 |                                      |    |
| DATASTORE                                       |      |                                   | bee1f629-4e23-4ab8-8083-9e0f0bb83178 |    |
| 10.160.93.240                                   | 9000 | ychin-nsxmanager-ob-12065118-2-F5 |                                      |    |
| MANAGER                                         |      |                                   | 45ccd6e3-1497-4334-944c-e6bbcd5c723e |    |
| 10.160.93.240                                   | -    | ychin-nsxmanager-ob-12065118-2-F5 |                                      |    |
| POLICY                                          |      |                                   | d5ba5803-b059-4fbc-897c-3aace8cf1219 |    |
| 10.160.93.240                                   | -    | ychin-nsxmanager-ob-12065118-2-F5 |                                      |    |
| Node UUID: 2e7e0642-df4a-b2ec-b9e8-633d1469f1ea |      |                                   |                                      |    |
| Node Status: JOINED                             |      |                                   |                                      |    |
| ENTITY                                          |      |                                   | UUID                                 | IP |
| ADDRESS                                         | PORT | FQDN                              |                                      |    |
| HTTPS                                           |      |                                   | bce3cc4c-7d60-45e2-aa7b-cdc75e445a14 |    |
| 10.160.76.33                                    | 443  | ychin-nsxmanager-ob-12065118-3-F5 |                                      |    |
| CONTROLLER                                      |      |                                   | ced46f5c-9e52-4b31-a1cb-b3dead991c71 |    |
| 10.160.76.33                                    | -    | ychin-nsxmanager-ob-12065118-3-F5 |                                      |    |
| CLUSTER_BOOT_MANAGER                            |      |                                   | 88b70d31-3428-4ccc-ab57-55859f45030c |    |
| 10.160.76.33                                    | -    | ychin-nsxmanager-ob-12065118-3-F5 |                                      |    |
| DATASTORE                                       |      |                                   | fb4aec3c-cae3-4386-b5b9-c0b99b7d9048 |    |
| 10.160.76.33                                    | 9000 | ychin-nsxmanager-ob-12065118-3-F5 |                                      |    |
| MANAGER                                         |      |                                   | 82b07440-3ff6-4f67-a1c9-e9327d1686ad |    |
| 10.160.76.33                                    | -    | ychin-nsxmanager-ob-12065118-3-F5 |                                      |    |
| POLICY                                          |      |                                   | 61f21a78-a56c-4af1-867b-3f24132d53c7 |    |
| 10.160.76.33                                    | -    | ychin-nsxmanager-ob-12065118-3-F5 |                                      |    |

#### 4 To see additional information about the status, run the following CLI command:

```

manager1> get cluster status
Cluster Id: 18807edd-56d1-4107-b7b7-508d766a08e3
Group Type: DATASTORE
Group Status: STABLE

Members:

```

| UUID                                 |        | FQDN                              |
|--------------------------------------|--------|-----------------------------------|
| IP                                   | STATUS |                                   |
| 43cd0642-275c-af1d-fe46-1f5200f9e5f9 |        | ychin-nsxmanager-ob-12065118-1-F5 |
| 10.160.71.225                        | UP     |                                   |
| 8ebb0642-201e-6a5f-dd47-a1e38542e672 |        | ychin-nsxmanager-ob-12065118-2-F5 |
| 10.160.93.240                        | UP     |                                   |
| 2e7e0642-df4a-b2ec-b9e8-633d1469f1ea |        | ychin-nsxmanager-ob-12065118-3-F5 |
| 10.160.76.33                         | UP     |                                   |
| Group Type: CLUSTER_BOOT_MANAGER     |        |                                   |
| Group Status: STABLE                 |        |                                   |
| Members:                             |        |                                   |
| UUID                                 |        | FQDN                              |
| IP                                   | STATUS |                                   |
| 43cd0642-275c-af1d-fe46-1f5200f9e5f9 |        | ychin-nsxmanager-ob-12065118-1-F5 |
| 10.160.71.225                        | UP     |                                   |
| 8ebb0642-201e-6a5f-dd47-a1e38542e672 |        | ychin-nsxmanager-ob-12065118-2-F5 |
| 10.160.93.240                        | UP     |                                   |
| 2e7e0642-df4a-b2ec-b9e8-633d1469f1ea |        | ychin-nsxmanager-ob-12065118-3-F5 |
| 10.160.76.33                         | UP     |                                   |
| Group Type: CONTROLLER               |        |                                   |
| Group Status: STABLE                 |        |                                   |
| Members:                             |        |                                   |
| UUID                                 |        | FQDN                              |
| IP                                   | STATUS |                                   |
| 7b1c9952-8738-4900-b68b-ca862aa4f6a9 |        | ychin-nsxmanager-ob-12065118-2-F5 |
| 10.160.93.240                        | UP     |                                   |
| ced46f5c-9e52-4b31-a1cb-b3dead991c71 |        | ychin-nsxmanager-ob-12065118-3-F5 |
| 10.160.76.33                         | UP     |                                   |
| 06fd0574-69c0-432e-a8af-53d140dbef8f |        | ychin-nsxmanager-ob-12065118-1-F5 |
| 10.160.71.225                        | UP     |                                   |
| Group Type: MANAGER                  |        |                                   |
| Group Status: STABLE                 |        |                                   |
| Members:                             |        |                                   |
| UUID                                 |        | FQDN                              |
| IP                                   | STATUS |                                   |
| 43cd0642-275c-af1d-fe46-1f5200f9e5f9 |        | ychin-nsxmanager-ob-12065118-1-F5 |
| 10.160.71.225                        | UP     |                                   |
| 8ebb0642-201e-6a5f-dd47-a1e38542e672 |        | ychin-nsxmanager-ob-12065118-2-F5 |
| 10.160.93.240                        | UP     |                                   |
| 2e7e0642-df4a-b2ec-b9e8-633d1469f1ea |        | ychin-nsxmanager-ob-12065118-3-F5 |
| 10.160.76.33                         | UP     |                                   |
| Group Type: POLICY                   |        |                                   |
| Group Status: STABLE                 |        |                                   |
| Members:                             |        |                                   |
| UUID                                 |        | FQDN                              |
| IP                                   | STATUS |                                   |
| 43cd0642-275c-af1d-fe46-1f5200f9e5f9 |        | ychin-nsxmanager-ob-12065118-1-F5 |

```

10.160.71.225 UP
 8ebb0642-201e-6a5f-dd47-a1e38542e672 ychin-nsxmanager-ob-12065118-2-F5
10.160.93.240 UP
 2e7e0642-df4a-b2ec-b9e8-633d1469f1ea ychin-nsxmanager-ob-12065118-3-F5
10.160.76.33 UP

Group Type: HTTPS
Group Status: STABLE

Members:
 UUID FQDN
IP STATUS
 43cd0642-275c-af1d-fe46-1f5200f9e5f9 ychin-nsxmanager-ob-12065118-1-F5
10.160.71.225 UP
 8ebb0642-201e-6a5f-dd47-a1e38542e672 ychin-nsxmanager-ob-12065118-2-F5
10.160.93.240 UP
 2e7e0642-df4a-b2ec-b9e8-633d1469f1ea ychin-nsxmanager-ob-12065118-3-F5
10.160.76.33 UP

```

## Update API Service Configuration of the NSX Manager Cluster

You can modify the API service properties of the NSX Manager cluster, such as TLS protocol version, cipher suites, and so on.

The following procedure explains the workflow of running the NSX API service calls to disable the TLS 1.1 protocol, and to enable or disable the cipher suites in the API service configuration.

For a detailed information about the API schema, example request, example response, and error messages of the NSX API service, you must read the *NSX-T Data Center API Guide*.

### Procedure

- 1 Run the following GET API to read the configuration of the NSX API service:

```
GET https://<NSX-Manager-IP>/api/v1/cluster/api-service
```

The API response contains the list of cipher suites and TLS protocols.

- 2 Disable the TLS 1.1 protocol.

- a Set TLSv1.1 to `enabled = false`.

- b Run the following PUT API to send the changes to the NSX API server:

```
PUT https://<NSX-Manager-IP>/api/v1/cluster/api-service
```

- 3 Enable or disable the cipher suites.

- a Set one or more cipher names to `enabled = false` or `enabled = true` depending on your requirement.

- b Run the following PUT API to send the changes to the NSX API server:

```
PUT https://<NSX-Manager-IP>/api/v1/cluster/api-service
```

## Results

The API service on each NSX Manager node restarts after it is updated using the API. There might be a delay of up to a minute between the time the API call completes and when the new configuration comes into effect. The changes in the API service configuration are applied to all the nodes in the NSX Manager cluster.

## Shut Down and Power On the NSX Manager Cluster

If you need to shut down the NSX Manager cluster, use the following procedure.

### Procedure

- 1 To shut down an NSX Manager cluster, shut down one manager node at a time. You can log in to the command-line interface (CLI) of a manager node as `admin` and run the command `shutdown`, or shut down the manager node VM from vCenter Server.

Make sure that the VM is powered off in vCenter Server before proceeding to the next one.

- 2 To power on an NSX Manager cluster, power on one manager node VM at a time in vCenter Server.

Make sure that the node is up and running before proceeding to the next one.

## Reboot an NSX Manager

You can reboot an NSX Manager with a CLI command to recover from critical errors.

If you need to reboot multiple NSX Managers, you must reboot them one at a time. Wait for the rebooted NSX Manager to be online before rebooting another.

### Procedure

- 1 Log in to the CLI of the NSX Manager.
- 2 Run the following command.

```
nsx-manager> reboot
Are you sure you want to reboot (yes/no): y
```

## Change the IP Address of an NSX Manager

You can change the IP address of an NSX Manager in an NSX Manager cluster. This section describes several approaches.

For example, if you have a cluster consisting of Manager A, Manager B, and Manager C, you can change the IP address of one or more of the managers in the following ways:

- Scenario A:
  - Manager A has IP address 172.16.1.11.
  - Manager B has IP address 172.16.1.12.



- Manager C has IP address 172.16.1.13.
- Add Manager D with a new IP address, for example, 192.168.55.11.
- Remove Manager A.
- Add Manager E with a new IP address, for example, 192.168.55.12.
- Remove Manager B.
- Add Manager F with a new IP address, for example, 192.168.55.13.
- Remove Manager C.
- Scenario B:
  - Manager A has IP address 172.16.1.11.
  - Manager B has IP address 172.16.1.12.
  - Manager C has IP address 172.16.1.13.
  - Add Manager D with a new IP address, for example, 192.168.55.11.
  - Add Manager E with a new IP address, for example, 192.168.55.12.
  - Add Manager F with a new IP address, for example, 192.168.55.13.
  - Remove Manager A, Manager B, and Manager C.
- Scenario C:
  - Manager A has IP address 172.16.1.11.
  - Manager B has IP address 172.16.1.12.
  - Manager C has IP address 172.16.1.13.
  - Remove Manager A.
  - Add Manager D with a new IP address, for example, 192.168.55.11.
  - Remove Manager B.
  - Add Manager E with a new IP address, for example, 192.168.55.12.
  - Remove Manager C.
  - Add Manager F with a new IP address, for example, 192.168.55.13.

The first two scenarios require additional virtual RAM, CPU and disk for the additional NSX Managers during this IP address change.

Scenario C is not recommended because it temporarily reduces the number of NSX Managers and a loss of one of the two active managers during the IP address change will have an impact on the operations of NSX-T Data Center. This scenario is for a situation where additional virtual RAM, CPU and disk are not available and an IP address change is required.

---

**Note** If you are using the cluster VIP feature, you must either use the same subnet for the new IP addresses or disable the cluster VIP during the IP address changes because the cluster VIP requires all NSX Managers to be in the same subnet.

---

### Prerequisites

Familiarize yourself with how to deploy an NSX Manager into a cluster. For more information, see the *NSX-T Data Center Installation Guide*.

### Procedure

- 1 If the NSX Manager you want to remove was deployed manually, perform the following steps.
  - a Run the following CLI command to detach the NSX Manager from the cluster.
 

```
detach node <node-id>
```
  - b Delete the NSX Manager VM.
- 2 If the NSX Manager you want to delete was deployed automatically through the NSX Manager UI, perform the following steps.
  - a From your browser, log in with administrator privileges to an NSX Manager at `https://nsx-manager-ip-address`.
 

This NSX Manager must not be the one that you want to delete.
  - b From the **Systems** tab, click **NSX Management Nodes**.
 

The status of the NSX Manager cluster is displayed.
  - c For the NSX Manager that you want to delete, click the gear icon and select **Delete**.
- 3 Deploy a new NSX Manager.

## Resize an NSX Manager Node

You can change the number of CPU cores or memory of an NSX Manager node at any time.

Note that in normal operating conditions all three manager nodes must have the same number of CPU cores and memory. A mismatch of CPU or memory between NSX Managers in an NSX management cluster should only be done when transitioning from one size of NSX Manager to another size of NSX Manager.

If you have configured resource allocation reservation for the NSX Manager VMs in vCenter Server, you might need to adjust the reservation. For more information, see the vSphere documentation.

**Prerequisites**

- Verify that the new size satisfies the system requirements for a manager node. For more information, see "NSX Manager VM System Requirements" in the *NSX-T Data Center Installation Guide*.
- Familiarize yourself with how to deploy an NSX Manager into a cluster. For more information, see the *NSX-T Data Center Installation Guide*.
- For information about how to remove a manager node from a cluster, see [Change the IP Address of an NSX Manager](#).

**Procedure**

- 1 Deploy a new manager node with the new size.
- 2 Add the new manager node to the cluster.
- 3 Remove an old manager node.
- 4 Repeat steps 1 to 3 to replace the other two old manager nodes.

## Replacing an NSX Edge Transport Node in an NSX Edge Cluster

You can replace an NSX Edge transport node in an NSX Edge cluster using the NSX Manager UI or the API.

### Replace an NSX Edge Transport Node Using the NSX Manager UI

The following procedure describes replacing an NSX Edge transport node in an NSX Edge cluster using the NSX Manager UI. You can replace the Edge transport node regardless of whether it is running or not.

If the Edge node to be replaced is not running, the new Edge node can have the same management IP address and TEP IP address. If the Edge node to be replaced is running, the new Edge node must have a different management IP address and TEP IP address.

**Prerequisites**

Familiarize yourself with the procedure to install an NSX Edge node, join the Edge node with the management plane, and create an NSX Edge transport node. For more information, see the *NSX-T Data Center Installation Guide*.

**Procedure**

- 1 If you want the new Edge transport node to have the same configurations as the Edge transport node to be replaced, make the following API call to find the configurations:

```
GET https://<nsx-manager-IP>/api/v1/transport-nodes/<tn-id>
```

- 2 Follow the procedures in the *NSX-T Data Center Installation* guide to install and configure an Edge transport node.

If you want this Edge transport node to have the same configurations as the Edge transport node to be replaced, use the configurations obtained in step 1.

- 3 In NSX Manager, select **System > Fabric > Nodes > Edge Clusters**.
- 4 Select an Edge cluster by clicking the checkbox in the first column.
- 5 Click **Actions > Replace Edge Cluster Member**.

It is recommended that you place the transport node being replaced in maintenance mode. If the transport node is not running, you can safely ignore this recommendation.

- 6 Select the node to be replaced from the dropdown list.
- 7 Select the replacement node from the dropdown list.
- 8 Click **Save**.

### Results

If you are running an NSX-T version earlier than 3.1.3, after replacing the NSX Edge transport node, you might see the alarm "All BGP/BFD sessions are down." To resolve the issue, follow the workaround instructions in the KB article <https://kb.vmware.com/s/article/83983>.

## Replace an NSX Edge Transport Node Using the API

The following procedure describes replacing an NSX Edge transport node in an NSX Edge cluster using the NSX-T API. You can replace the Edge transport node regardless of whether it is running or not.

If the Edge node to be replaced is not running, the new Edge node can have the same management IP address and TEP IP address. If the Edge node to be replaced is running, the new Edge node must have a different management IP address and TEP IP address.

### Prerequisites

Familiarize yourself with the procedure to install an NSX Edge node, join the Edge node with the management plane, and create an NSX Edge transport node. For more information, see the *NSX-T Data Center Installation Guide*.

### Procedure

- 1 If you want the new Edge transport node to have the same configurations as the Edge transport node to be replaced, make the following API call to find the configurations:

```
GET https://<nsx-manager-IP>/api/v1/transport-nodes/<tn-id>
```

- 2 Follow the procedures in the NSX-T Data Center Installation guide to install and configure an Edge transport node.

If you want this Edge transport node to have the same configurations as the Edge transport node to be replaced, use the configurations obtained in step 1.

- 3 Make an API call to get the ID of the new transport node and the transport node to be replaced. The `id` field contains the transport node ID. For example,

```
GET https://<nsx-manager-IP>/api/v1/transport-nodes
...
{
 "resource_type": "TransportNode",
 "description": "",
 "id": "73cb00c9-70d0-4808-abfe-a12a43251133",
 "display_name": "TN-edgenode-01a",
 ...
 {
 "resource_type": "TransportNode",
 "description": "",
 "id": "890f0e3c-aa81-46aa-843b-8ac25fe30bd3",
 "display_name": "TN-edgenode-03a",
```

- 4 Make an API call to get the ID of the NSX Edge cluster. The `id` field contains the NSX Edge cluster ID. Get the members of the NSX Edge cluster from the `members` array. For example,

```
GET https://<nsx-manager-IP>/api/v1/edge-clusters
....
{
 "resource_type": "EdgeCluster",
 "description": "",
 "id": "9a302df7-0833-4237-af1f-4d826c25ad78",
 "display_name": "Edge-Cluster-1",
 ...
 "members": [
 {
 "member_index": 0,
 "transport_node_id": "73cb00c9-70d0-4808-abfe-a12a43251133"
 },
 {
 "member_index": 1,
 "transport_node_id": "e5d17b14-cdeb-4e63-b798-b23a0757463b"
 }
],
}
```

- 5 Make an API to replace a transport node in an NSX Edge cluster. The `member_index` must match the index of the transport node to be replaced.

For example, the transport node `TN-edgenode-01a` (73cb00c9-70d0-4808-abfe-a12a43251133) has failed and is replaced by transport node `TN-edgenode-03a` (890f0e3c-aa81-46aa-843b-8ac25fe30bd3) in NSX Edge cluster `Edge-Cluster-1` (9a302df7-0833-4237-af1f-4d826c25ad78).

```
POST http://<nsx-manager-IP>/api/v1/edge-clusters/9a302df7-0833-4237-af1f-4d826c25ad78?
action=replace_transport_node
{
 "member_index": 0,
 "transport_node_id" : "890f0e3c-aa81-46aa-843b-8ac25fe30bd3"
}
```

### Results

If you are running an NSX-T version earlier than 3.1.3, after replacing the NSX Edge transport node, you might see the alarm "All BGP/BFD sessions are down." To resolve the issue, follow the workaround instructions in the KB article <https://kb.vmware.com/s/article/83983>.

## Managing Resource Reservations for an Edge VM Appliance

NSX-T Data Center uses vSphere resource allocation to reserve resources for NSX Edge appliances. You can tune the CPU and memory resources reserved for NSX Edge to ensure optimal use of resources on an NSX Edge.

For maximum performance NSX Edge VM appliance must be assigned 100% of the available resources. If you customize resources allocated to the NSX Edge VM, turn back the allocation later to 100% to get maximum performance.

For auto-deployed NSX Edge appliances, you can change the resource allocation from NSX Manager. However, if an NSX Edge appliance is deployed from vSphere, you can only manage resource reservations for that NSX Edge VM from vSphere.

As per the resource requirements of the Edge VM deployed in your environment, there are two ways to manage reservations:

- Default values assigned to give 100% resource reservations.
- Custom values assigned to give 0–100% resource reservations.

### Default Reservations

Assumes the NSX Edge set to the **High** priority. The level of priority importance defines the number of vCPU shares and memory assigned to the NSX Edge. To assign custom values, you can change the relative priority assigned to the NSX Edge.

Resource constraints for different form factors set with Normal priority:

| Form Factor | Number of vCPUs | vCPU Shares | RAM (GB) |
|-------------|-----------------|-------------|----------|
| Small       | 2               | 2000        | 4        |
| Medium      | 4               | 4000        | 8        |
| Large       | 8               | 8000        | 32       |
| XLarge      | 16              | 16000       | 64       |

You can tune reservations of an NSX Edge appliance by considering two parameters:

- Relative priority assigned to a VM
- Pre-assigned resource constraints for a VM form factor

## Custom Reservations

Assign relative priority for an NSX Edge appliance. You can change the relative importance of an NSX Edge appliance to assign the following resource requirements:

| Relative Importance | CPU Shares (shares per vCPU) | Memory (shares per MB configured virtual machine memory) |
|---------------------|------------------------------|----------------------------------------------------------|
| Extra High          | 4000                         | 40                                                       |
| High                | 2000                         | 20                                                       |
| Normal              | 1000                         | 10                                                       |
| Low                 | 500                          | 5                                                        |

For example, a High relative importance to an NSX Edge appliance deployed in a medium form factor assigns the following vCPU and memory shares:

- $4 \text{ (vCPUs)} \times 8000 \text{ (vCPU share value)} = 32000 \text{ shares of vCPU}$
- $20 \text{ (GB RAM)} \times 1000 = 20000 \text{ shares of memory}$

**Note** Before assigning a CPU value in MHz to guarantee the allocated CPU cycles for an NSX EdgeVM, ensure that the relative importance is set to Low. If the relative importance is set to Normal or High with a custom CPU value in MHz, the VM deployment might face issues due to resource constraints.

## Tune Resource Reservations for an NSX Edge Appliance

You can tune resource reservations on an NSX Edge VM appliance. By default, 100% resources are allocated to an NSX Edge VM. Flexibility to change resource reservations avoids the need to add additional capacity to the vCenter Server and the need to reduce current reservations on other non-Edge VMs.

### Prerequisites

- Verify that the cluster has sufficient capacity to avoid failures.

## Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **System > Fabric > Nodes > Edge Transport Nodes**.
- 3 Select the NSX Edge transport node.
- 4 Click **Actions > Change Edge VM Resource Reservations**.
- 5 In the **Change Edge VM Resource Reservations** window, you can customize the existing resource allocation applied to the Edge transport node.

| Action                   | Description                                                                                                                                                                                                                                                                                                                                           |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CPU Reservation Priority | Low - 2000 shares<br>Normal - 4000 shares<br>High - 8000 shares<br>Extra High - 10000 shares                                                                                                                                                                                                                                                          |
| Memory Reservation (%)   | Reservation percentage is relative to the pre-defined value in the form factor.<br>100 indicates 100% of memory is reserved for the NSX Edge VM.<br>If you enter 50, it indicates that 50% of memory is assigned to the Edge transport node.                                                                                                          |
| CPU Reservation (MHz)    | Enter CPU reservation in MHz.<br>The maximum amount of MHz is equal to the number of vCPUs multiplied by the normal CPU operation rate of the physical CPU core.<br><br><b>Note</b> 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. |

- 6 Click **Save**.

If changes made to the resource reservations do not take effect, you might need to reboot the NSX Edge VM from vCenter Server.

The NSX Edge VM appliance autostarts on ESXi host reboot provided the NSX Edge cluster has vSphere HA turned off. For more details on vSphere HA, see the vSphere documentation.

## Adding and Removing an ESXi Host Transport Node to and from vCenter Servers

You can move an ESXi host transport node from one vCenter Server (VC) to another, and also from one NSX Manager cluster to another.



## Scenario 1: VC1 connected to NSX Manager cluster 1, and VC2 connected to NSX Manager cluster 2

Assuming ESX1, an ESXi host transport node, is in VC1, you can move it to VC2 by performing the following steps:

- 1 Uninstall NSX from ESX1.
- 2 Move ESX1 to VC2.
- 3 Apply a transport node profile to ESX1.

## Scenario 2: Both VC1 and VC2 connected to NSX Manager cluster

Assuming ESX1, an ESXi host transport node, is in VC1, you can move it to VC2 by performing the following steps:

- 1 Uninstall NSX from ESX1.
- 2 Move ESX1 to VC2.
- 3 Apply a transport node profile to ESX1.

## Scenario 3: VC1 connected to NSX Manager cluster 1

Assuming ESX1, an ESXi host transport node, is in VC1, you can move it to NSX Manager cluster 2 as a standalone host by performing the following steps:

- 1 Uninstall NSX from ESX1.
- 2 Add ESX1 to NSX Manager cluster 2.

## Changing the Distributed Router Interfaces' MAC Address

All logical router interfaces in NSX-T and NSX for vSphere setups have the same MAC address (02:50:56:56:44:52). Starting with NSX-T Data Center 3.0.2, you can change this address in NSX-T to avoid issues when migrating VMs from an NSX for vSphere setup to an NSX-T setup.

Changing the MAC address involves making two API calls.

If you have not created any transport node, make the following GET API call. For example:

```
GET https://10.40.79.126/api/v1/global-configs/RoutingGlobalConfig
```

Response:

```
{
 "l3_forwarding_mode" : "IPV4_ONLY",
 "logical_uplink_mtu" : 1500,
 "vdr_mac" : "02:50:56:56:44:77",
 "vdr_mac_nested" : "02:50:56:56:44:52",
 "allow_changing_vdr_mac_in_use" : true,
 "resource_type" : "RoutingGlobalConfig",
 "id" : "49b261fe-f4e4-46ad-958c-da9cb4271e32",
 "display_name" : "49b261fe-f4e4-46ad-958c-da9cb4271e32",
```

```

 "_create_user" : "system",
 "_create_time" : 1595313890595,
 "_last_modified_user" : "admin",
 "_last_modified_time" : 1595465694142,
 "_system_owned" : false,
 "_protection" : "NOT_PROTECTED",
 "_revision" : 14
 }

```

Take the response of the call, change the `vdr_mac` value, and use it to make the following PUT API call. For example:

```

PUT https://10.40.79.126/api/v1/global-configs/RoutingGlobalConfig
{
 "l3_forwarding_mode" : "IPv4_ONLY",
 "logical_uplink_mtu" : 1500,
 "vdr_mac" : "02:50:56:56:44:99",
 "vdr_mac_nested" : "02:50:56:56:44:53",
 "allow_changing_vdr_mac_in_use" : true,
 "resource_type" : "RoutingGlobalConfig",
 "id" : "49b261fe-f4e4-46ad-958c-da9cb4271e32",
 "display_name" : "49b261fe-f4e4-46ad-958c-da9cb4271e32",
 "_create_user" : "system",
 "_create_time" : 1595313890595,
 "_last_modified_user" : "admin",
 "_last_modified_time" : 1595465694142,
 "_system_owned" : false,
 "_protection" : "NOT_PROTECTED",
 "_revision" : 14
}

```

Response:

```

{
 "l3_forwarding_mode" : "IPv4_ONLY",
 "logical_uplink_mtu" : 1500,
 "vdr_mac" : "02:50:56:56:44:99",
 "vdr_mac_nested" : "02:50:56:56:44:53",
 "allow_changing_vdr_mac_in_use" : true,
 "resource_type" : "RoutingGlobalConfig",
 "id" : "49b261fe-f4e4-46ad-958c-da9cb4271e32",
 "display_name" : "49b261fe-f4e4-46ad-958c-da9cb4271e32",
 "_create_user" : "system",
 "_create_time" : 1595313890595,
 "_last_modified_user" : "admin",
 "_last_modified_time" : 1595466163148,
 "_system_owned" : false,
 "_protection" : "NOT_PROTECTED",
 "_revision" : 15
}

```

If you have already created transport nodes, make the same GET and PUT API calls, except that for the PUT call, set the parameter `allow_changing_vdr_mac_in_use` to `true`.

## Configuring Appliances

Some system configuration tasks must be done using the command line or API.

For complete command line interface information, see the *NSX-T Data Center Command-Line Interface Reference*. For complete API information, see the *NSX-T Data Center API Guide*.

**Table 22-1. System configuration commands and API requests.**

| Task                  | Command Line<br>(NSX Manager and NSX Edge)       | API Request<br>(NSX Manager only)                                           |
|-----------------------|--------------------------------------------------|-----------------------------------------------------------------------------|
| Set system timezone   | <code>set timezone &lt;timezone&gt;</code>       | PUT <code>https://&lt;nsx-mgr&gt;/api/v1/node</code>                        |
| Set NTP Server        | <code>set ntp-server &lt;ntp-server&gt;</code>   | PUT <code>https://&lt;nsx-mgr&gt;/api/v1/node/services/ntp</code>           |
| Set a DNS server      | <code>set name-servers &lt;dns-server&gt;</code> | PUT <code>https://&lt;nsx-mgr&gt;/api/v1/node/network/name-servers</code>   |
| Set DNS Search Domain | <code>set search-domains &lt;domain&gt;</code>   | PUT <code>https://&lt;nsx-mgr&gt;/api/v1/node/network/search-domains</code> |

**Note** The recommended method to configure an NTP server for all appliances is to configure a node profile. See [Configure a Node Profile](#). If you configure an NTP server individually on an appliance, be sure to configure the same NTP server on all the appliances.

## Add a License Key and Generate a License Usage Report

You can add license keys and generate a license usage report. The usage report is a file in CSV format.

The following regular license types are available:

- NSX Data Center Standard
- NSX Data Center Professional
- NSX Data Center Advanced
- NSX Data Center Enterprise Plus
- NSX Data Center Remote Office Branch Office (ROBO)
- NSX Data Center Evaluation
- NSX for vSphere - Standard
- NSX for vSphere - Advanced
- NSX for vSphere - Enterprise
- NSX for vShield Endpoint

The following limited export license types are available:

- VMware NSX Enterprise per Processor (Limited Export)

- NSX Data Center Advanced per Processor (for Limited Export)
- NSX Data Center Evaluation
- NSX for vShield Endpoint

Add-on license type for both regular and limited export licenses:

- NSX Distributed Threat Prevention (Distributed IDS included)

You can add this license only if the NSX Data Center Advanced or NSX Data Center Enterprise Plus license exists. And you cannot delete the NSX Data Center Advanced or NSX Data Center Enterprise Plus license until this add-on license is deleted.

Note: For Limited Export Release version, you can add this license only if the VMware NSX Enterprise per Processor (Limited Export) or NSX Data Center Advanced per Processor (for Limited Export) license exists. And you cannot delete the VMware NSX Enterprise per Processor (Limited Export) or NSX Data Center Advanced per Processor (for Limited Export) license until this add-on license is deleted.

When you install NSX Manager, the default license is NSX for vShield Endpoint. This license never expires but has certain restrictions. You cannot create or update the following objects:

- Tier-0 and tier-1 logical router
- Tier-0 and tier-1 gateway
- Logical switch
- Layer 2 segment (Note: You can create and update service segments.)
- Distributed firewall
- VPN
- NAT
- Load balancer
- Service Insertion
- NSX Intelligence

If you upgrade from a previous release, both the default vShield Endpoint license and the previous default NSX Data Center Evaluation will be available. Note the following:

- You cannot delete the default vShield Endpoint license key.
- You can delete the previous default NSX Data Center Evaluation license key.
- If you add a new vShield Endpoint license, the default vShield Endpoint license will be hidden. If you remove the new vShield Endpoint license, the default vShield Endpoint license will be available again.

- If you add a new NSX Data Center Evaluation license, the default NSX Data Center Evaluation license, if it exists because of an upgrade, will be permanently deleted. If you remove the new NSX Data Center Evaluation license, you will not have any evaluation license.

---

**Note** About the evaluation license:

If you install the NSX Data Center Evaluation license, it will be valid for 60 days.

Note the following:

- You have the NSX for vShield Endpoint license and the NSX Data Center Evaluation license only.
    - If the NSX Data Center Evaluation license is valid, it will be used.
    - If the NSX Data Center Evaluation license has expired, the NSX for vShield Endpoint license will be used. (Enforcement will take effect.)
  - You have the NSX for vShield Endpoint license, the NSX Data Center Evaluation license, and other licenses.
    - If the NSX Data Center Evaluation license is valid, it and the other licenses will be used.
    - If the NSX Data Center Evaluation license has expired, the other licenses will be used.
- 

If a license has expired or will expire within 60 days, an alarm will be generated each time you log in. You can view alarms by going to **Home > Alarms**.

If you have only the NSX for vShield Endpoint license, after you log in, an informational banner message will let you know that the license has certain restrictions (see above). If you have any expired or expiring license, after you log in, a warning banner message will let you know.

If you have multiple keys of the same license type and want to combine the keys, you must go to <https://my.vmware.com> and use the Combine Keys functionality. The NSX Manager UI does not provide this functionality.

#### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **System > Licenses > Add**.
- 3 Enter a license key.
- 4 To generate a license usage report, select **Export > License Usage Report**.

The CSV report lists the VM, CPU, unique concurrent user, vCPU and core usage numbers of the following features:

- Switching and Routing
- NSX Edge load balancer
- VPN

- Distributed Firewall
- Context Aware Micro-Segmentation - Application identification
- Context Aware Micro-Segmentation - Identity firewall for remote desktop session host
- Service Insertion
- Identity Firewall
- Enhanced Guest Introspection
- Micro-Segmentation Planning (L4) (CPU, core, VM, and concurrent user information only)
- Federation (CPU, core, VM, and concurrent user information only)
- Distributed IDS (CPU information only)

---

**Note** The following features are disabled for the Limited Export Release version:

- IPsec VPN
  - HTTPS-based Load Balancer
- 

## Compliance-Based Configuration

NSX-T Data Center can be configured to use FIPS 140-2 validated cryptographic modules to run in FIPS-compliant mode. The modules are validated to FIPS 140-2 standards by the NIST Cryptographic Module Validation Program (CMVP).

All exceptions to FIPS compliance can be retrieved using the compliance report. See [View Compliance Status Report](#) for more information.

The following validated modules are used:

- VMware OpenSSL FIPS Object Module version 2.0.9: [Certificate #2839](#)
- VMware's OpenSSL FIPS Object Module version 2.0.20-vmw: [Certificate #3550](#)
- BC-FJA (Bouncy Castle FIPS Java API) version 1.0.1: [Certificate #3152](#)
- VMware's IKE Crypto Module version 1.1.0: [Certificate #3435](#)
- VMware's VPN Crypto Module version 1.0: [Certificate #3542](#)

You can find more information about the cryptographic modules that VMware has validated against the FIPS 140-2 standard here: <https://www.vmware.com/security/certifications/fips.html>.

By default, load balancer uses modules that have FIPS mode turned off. You can turn on FIPS mode for the modules used by load balancer. See [Configure Global FIPS Compliance Mode for Load Balancer](#) for more information.

Details about southbound and northbound connections to the NSX Controller:

- For southbound connections between the controller component of the NSX Manager appliance and other nodes, X509 certificate-based authentication is used with FIPS 140-2 validated OpenSSL algorithm. The connections support TLS 1.2-based cipher suites with AES 128-bit, 256-bit, or 384-bit encryption keys.
- The controller function and the management function of the NSX Manager appliance run on the same node. Hence, there is no north-bound cross-node communication between the controller and manager components of the NSX Manager appliance.

## View Compliance Status Report

You can view a compliance report for NSX-T Data Center features. You can use the report to configure your NSX-T Data Center environment to adhere to your IT policies and industry standards.

The compliance report includes information about each non-compliant configuration.

**Table 22-2. Compliance Report Information**

| Compliance Report Column   | Description                                                                                                                       | Example                       |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------|-------------------------------|
| <b>Non Compliance Code</b> | Code to identify the type of non-compliance.                                                                                      | 72301                         |
| <b>Description</b>         | Description of the type of non-compliance.                                                                                        | Certificate is not CA signed. |
| <b>Resource Name</b>       | Name or ID of the affected resource.                                                                                              | nsx-manager-1                 |
| <b>Resource Type</b>       | Type of resource affected.                                                                                                        | CertificateComplianceReporter |
| <b>Affected Resources</b>  | Number of affected resources. The number can be 0 if there are non-compliant configurations present, but the feature is not used. | 1                             |

You can also retrieve the report using the API: `GET /policy/api/v1/compliance/status`.

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 From the **Home** page, click **Monitoring Dashboards > Compliance Report**.

## Compliance Status Report Codes

You can find more information about the meaning of the compliance status report.

Table 22-3. Compliance Report Codes

| Code  | Description                                                                                                                                                      | Compliance Status Source                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Remediation                                                                                                                                                                         |
|-------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 72001 | Encryption is disabled.                                                                                                                                          | <p>This status is reported if a VPN IPSec Profile configuration contains NO_ENCRYPTION, NO_ENCRYPTION_AUTH_AES_GMAC_128, NO_ENCRYPTION_AUTH_AES_GMAC_192, or NO_ENCRYPTION_AUTH_AES_GMAC_256 encryption_algorithms.</p> <p>This status affects IPSec VPN session configurations which use the reported non-compliant configurations.</p>                                                                                                                                       | To remediate this status, add a VPN IPSec Profile that uses compliant encryption algorithms and use the profile in all VPN configurations. See <a href="#">Add IPSec Profiles</a> . |
| 72011 | BGP messages with neighbor bypass integrity check. No message authentication defined.                                                                            | <p>This status is reported if no password is configured for BGP neighbors.</p> <p>This status affects the BGP neighbor configuration.</p>                                                                                                                                                                                                                                                                                                                                      | To remediate this status, configure a password on the BGP neighbor and update the tier-0 gateway configuration to use the password. See <a href="#">Configure BGP</a> .             |
| 72012 | Communication with BGP neighbor uses weak integrity check. MD5 is used for message authentication.                                                               | <p>This status is reported if MD5 authentication is used for the BGP neighbor password.</p> <p>This status affects the BGP neighbor configuration.</p>                                                                                                                                                                                                                                                                                                                         | No remediation available as NSX-T Data Center supports only MD5 authentication for BGP.                                                                                             |
| 72021 | SSL version 3 used for establishing secure socket connection. It is recommended to run TLSv 1.1 or higher and fully disable SSLv3 that have protocol weaknesses. | <p>This status is reported if SSL version 3 is configured in the load balancer client SSL profile, load balancer server SSL profile, or load balancer HTTPS monitor.</p> <p>This status affects the following configurations:</p> <ul style="list-style-type: none"> <li>■ Load balancer pools that are associated with HTTPS monitors.</li> <li>■ Load balancer virtual servers that are associated with load balancer client SSL profiles or server SSL profiles.</li> </ul> | To remediate this status, configure an SSL profile to use TLS 1.1 or later and use this profile in all load balancer configurations. See <a href="#">Add an SSL Profile</a> .       |



Table 22-3. Compliance Report Codes (continued)

| Code  | Description                                                                                                                                                          | Compliance Status Source                                                                                                                                                                                                                                                                                                                                                                                                                                             | Remediation                                                                                                                                                                   |
|-------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 72022 | TLS version 1.0 used for establishing secure socket connection. It is recommended to run TLSv 1.1 or higher and fully disable TLSv1.0 that have protocol weaknesses. | <p>This status is reported if TLSv1.0 is configured in load balancer client SSL profile, load balancer server SSL profile, or load balancer HTTPS monitor.</p> <p>This status affects the following configurations:</p> <ul style="list-style-type: none"> <li>■ Load balancer pools that are associated with HTTPS monitors.</li> <li>■ Load balancer virtual servers that are associated with load balancer client SSL profiles or server SSL profiles.</li> </ul> | To remediate this status, configure an SSL profile to use TLS 1.1 or later and use this profile in all load balancer configurations. See <a href="#">Add an SSL Profile</a> . |
| 72023 | Weak Diffie-Hellman group is used.                                                                                                                                   | <p>This error is reported if a VPN IPSec Profile or VPN IKE Profile configuration includes the following Diffie-Hellman groups: 2, 5, 14, 15 or 16. Groups 2 and 5 are weak Diffie-Hellman groups. Groups 14, 15, and 16 are not weak groups, but are not FIPS-compliant.</p> <p>This status affects IPSec VPN session configurations which use the reported non-compliant configurations.</p>                                                                       | To remediate this status, configure the VPN Profiles to use Diffie-Hellman group 19, 20, or 21. See <a href="#">Adding Profiles</a> .                                         |
| 72024 | Load balancer FIPS global setting is disabled.                                                                                                                       | <p>This error is reported if the load balancer FIPS global setting is disabled.</p> <p>This status affects all load balancer services.</p>                                                                                                                                                                                                                                                                                                                           | To remediate this status, enable FIPS for load balancer. See <a href="#">Configure Global FIPS Compliance Mode for Load Balancer</a> .                                        |

Table 22-3. Compliance Report Codes (continued)

| Code  | Description                          | Compliance Status Source                                                                                                                                                                                                                                                                                                                     | Remediation                                                                                                                                                                                                                                         |
|-------|--------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 72200 | Insufficient true entropy available. | <p>This status is reported when a pseudo random number generator is used to generate entropy rather than relying on hardware-generated entropy.</p> <p>Hardware-generated entropy is not used because the NSX Manager node does not have the required hardware acceleration support to create sufficient true entropy.</p>                   | <p>To remediate this status, you might need to use newer hardware to run the NSX Manager node. Most recent hardware supports this feature.</p> <hr/> <p><b>Note</b> If the underlying infrastructure is virtual, you will not get true entropy.</p> |
| 72201 | Entropy source unknown.              | This status is reported when no entropy status is available for the indicated node.                                                                                                                                                                                                                                                          | To remediate this status, verify that the indicated node is functioning properly.                                                                                                                                                                   |
| 72301 | Certificate is not CA signed.        | <p>This status is reported when one of the NSX Manager certificates is not CA signed. NSX Manager uses the following certificates:</p> <ul style="list-style-type: none"> <li>■ Syslog certificate.</li> <li>■ API certificates for the individual NSX Manager nodes.</li> <li>■ Cluster certificate used for the NSX ManagerVIP.</li> </ul> | To remediate this status, install CA-signed certificates. See <a href="#">Chapter 19 Certificates</a> .                                                                                                                                             |

## Configure Global FIPS Compliance Mode for Load Balancer

There is a global setting for FIPS compliance for load balancers. By default, the setting is turned off to improve performance.

Changing the global configuration for FIPS compliance for load balancers affects new load balancer instances, but does not affect any existing load balancer instances.

If the global setting for FIPS for load balancer (`lb_fips_enabled`) is set to *true*, new load balancer instances use modules that comply with FIPS 140-2. Existing load balancer instances might be using non-compliant modules.

To make the change take effect on existing load balancers, you must detach and reattach the load balancer from the tier-1 gateway.

You can check the global FIPS compliance status for load balancer using `GET /policy/api/v1/compliance/status`.

```
...
{
 "non_compliance_code": 72024,
 "description": "Load balancer FIPS global setting is disabled.",
 "reported_by": {
 "target_id": "971ca477-df1a-4108-8187-7918c2f8c3ba",
 "target_display_name": "971ca477-df1a-4108-8187-7918c2f8c3ba",
 "target_type": "FipsGlobalConfig",
 "is_valid": true
 },
 "affected_resources": [
 {
 "path": "/infra/lb-services/LB_Service",
 "target_id": "/infra/lb-services/LB_Service",
 "target_display_name": "LB_1",
 "target_type": "LBService",
 "is_valid": true
 }
]
},
...
```

**Note** The compliance report displays the global setting for FIPS compliance for load balancer. Any given load balancer instance can have a FIPS compliance status that is different from the global setting.

## Procedure

- 1 Retrieve the global FIPS setting for load balancer.

`GET https://nsx-mgr1/policy/api/v1/infra/global-config`

Example response body:

```
{
 "fips": {
 "lb_fips_enabled": false
 },
 "resource_type": "GlobalConfig",
 "id": "global-config",
 "display_name": "global-config",
 "path": "/infra/global-config",
 "relative_path": "global-config",
 "marked_for_delete": false,
 "_create_user": "system",
 "_create_time": 1561225479619,
 "_last_modified_user": "admin",
 "_last_modified_time": 1561937915337,
}
```

```

 "_system_owned": true,
 "_protection": "NOT_PROTECTED",
 "_revision": 2
 }

```

## 2 Change the global FIPS setting for load balancer.

The global setting is used when you create new load balancer instances. Changing the setting does not affect existing load balancer instances.

PUT <https://nsx-mgr1/policy/api/v1/infra/global-config>

Example request body:

```

{
 "fips": {
 "lb_fips_enabled": true
 },
 "resource_type": "GlobalConfig",
 "_revision": 2
}

```

Example response body:

```

{
 "fips": {
 "lb_fips_enabled": true
 },
 "resource_type": "GlobalConfig",
 "id": "global-config",
 "display_name": "global-config",
 "path": "/infra/global-config",
 "relative_path": "global-config",
 "marked_for_delete": false,
 "_create_user": "system",
 "_create_time": 1561225479619,
 "_last_modified_user": "admin",
 "_last_modified_time": 1561937960950,
 "_system_owned": true,
 "_protection": "NOT_PROTECTED",
 "_revision": 3
}

```

## 3 If you want any existing load balancer instances to use this global setting, you must detach and reattach the load balancer from the tier-1 gateway.

---

**Caution** Detaching a load balancer from the tier-1 gateway results in a traffic interruption for the load balancer instance.

---

- a Navigate to **Networking > Load Balancing**.
- b On the load balancer you want to detach, click the three dots menu (⋮), then click **Edit**.

- c Click , then click **Save** to detach the load balancer from the tier-1 gateway.

| Name   | Size    | Tier-1 Gateway                                                                            |
|--------|---------|-------------------------------------------------------------------------------------------|
| LB_1 * | Small ▾ | TLR1_LR  |

- d Click the three dots menu (⋮), then click **Edit**.
- e Select the correct gateway from the **Tier-1 Gateway** drop-down menu, then click **Save** to reattach the load balancer to the tier-1 gateway.

## Collect Support Bundles

You can collect support bundles on registered cluster and fabric nodes and download the bundles to your machine or upload them to a file server.

If you choose to download the bundles to your machine, you get a single archive file consisting of a manifest file and support bundles for each node. If you choose to upload the bundles to a file server, the manifest file and the individual bundles are uploaded to the file server separately.

**NSX Cloud Note** If you want to collect the support bundle for CSM, log in to CSM, go to **System > Utilities > Support Bundle** and click on **Download**. The support bundle for PCG is available from NSX Manager using the following instructions. The support bundle for PCG also contains logs for all the workload VMs.

### Procedure

- 1 From your browser, log in as a local admin user to an NSX Manager at <https://nsx-manager-ip-address/login.jsp?local=true>.
- 2 Select **System > Support Bundle**
- 3 Select the target nodes.

The available types of nodes are **Management Nodes**, **Edges**, **Hosts**, and **Public Cloud Gateways**.

- 4 (Optional) Specify log age in days to exclude logs that are older than the specified number of days.
- 5 (Optional) Toggle the switch that indicates whether to include or exclude core files and audit logs.

**Note** Core files and audit logs might contain sensitive information such as passwords or encryption keys.

- 6 (Optional) Select the check box to upload the bundles to a remote file server.

- 7 Click **Start Bundle Collection** to start collecting support bundles.

Depending on how many log files exist, each node might take several minutes.

- 8 Monitor the status of the collection process.

The status tab shows the progress of collecting support bundles.

- 9 Click **Download** to download the bundle if the option to send the bundle to a file remote server was not set.

The bundle collection may fail for a manager node if there is not enough disk space. If you encounter an error, check whether older support bundles are present on the failed node. Log in to the NSX Manager UI of the failed manager node using its IP address and initiate the bundle collection from that node. When prompted by the NSX Manager, either download the older bundle or delete it.

## Log Messages and Error Codes

NSX-T Data Center components write to log files in the directory `/var/log`. On NSX-T appliances and KVM hosts, NSX syslog messages conform with RFC 5424. On ESXi hosts, syslog messages conform with RFC 3164.

### Viewing Logs

On NSX-T appliances syslog messages are in `/var/log/syslog`. On KVM hosts, syslog messages are in `/var/log/vmware/nsx-syslog`.

On NSX-T appliances, you can run the following NSX-T CLI command to view the logs:

```
get log-file <auth.log | controller | controller-error | http.log | kern.log | manager.log |
node-mgmt.log | policy.log | syslog> [follow]
```

The log files are:

| Name                | Description          |
|---------------------|----------------------|
| auth.log            | Authorization log    |
| controller          | Controller log       |
| controller-error    | Controller error log |
| http.log            | HTTP service log     |
| kern.log            | Kernel log           |
| manager.log         | Manager service log  |
| node-mgmt.log       | Node management log  |
| nsx-audit-write.log | NSX audit write log  |
| nsx-audit.log       | NSX audit log        |
| policy.log          | Policy service log   |
| syslog              | System log           |

On hypervisors, you can use Linux commands such as `tail`, `grep`, and `more` to view the logs.

Each syslog message has the component (`comp`) and sub-component (`subcomp`) information to help identify the source of the message.

NSX-T Data Center produces logs with facility `local6`, which has a numerical value of 22.

The audit log is part of syslog. An audit log message can be identified by the string `audit="true"` in the `structured-data` field. You can configure an external log server to receive log messages. You can also access audit logs using the API `/api/v1/administration/audit-logs`. The file `nsx-audit.log` contains syslog messages with `audit="true"` in the `structured-data` field. The file `nsx-audit-write.log` contains syslog messages with both `audit="true"` and `update="true"` in the `structured-data` field.

Each syslog and audit log message contains a timestamp generated by an NTP server, if configured, or by the system clock. An example of an audit log message:

```
<182>1 2020-05-05T00:29:02.900Z nsx-manager1 NSX 14389 - [nsx@6876 audit="true"
comp="nsx-manager" level="INFO" reqId="fe75651d-c3e7-4680-8753-9ae9d92d7f0c" subcomp="policy"
username="admin"] UserName="admin", ModuleName="AAA", Operation="GetCurrentUserInfo",
Operation status="success"
```

An API call can come from NSX Manager, a policy API client, or an NSX node. All API calls are subject to authentication and authorization, and will generate audit logs. This logging is enabled by default and cannot be disabled. An audit log that is associated with an API call has the following information:

- An entity ID parameter `entId` to identify the object of the API.
- A request ID parameter `req-id` to identify a specific API call.
- An external request ID parameter `ereqId` if the API call contains the header `X-NSX-EREQID:<string>`.
- An external user parameter `euser` if the API call contains the header `X-NSX-EUSER:<string>`.

An audit log message from a policy or manager API call will have the following additional fields. Note that node API (NAPI) and CLI audit logs will not have these fields.

- An `update` flag that shows whether the API operation is a read (GET) or write (PUT/POST/DELETE/...) operation.
- An `operation name` field that shows the name of the API operation.
- An `operation status` field that shows whether the API operation succeeded or failed.
- A `new value` field that shows all parameter values of the API request.

NSX-T does not have the concept of a privileged mode. API calls from all sources and users are audited.

An example of login and logout syslog messages showing a successful login, a failed login, and logins from 2 different devices (note the different IP addresses):

```
2020-07-07T16:33:20.339Z svc.nsxmanager NSX 1513 SYSTEM [nsx@6876 audit="true"
comp="nsx-manager" level="INFO" subcomp="http"] UserName="admin@10.166.61.56",
ModuleName="ACCESS_CONTROL", Operation="LOGIN", Operation status="success"

2020-07-07T16:33:58.779Z svc.nsxmanager NSX 1513 SYSTEM [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" subcomp="http"] UserName="admin", ModuleName="ACCESS_CONTROL",
Operation="LOGOUT", Operation status="success"

2020-07-07T16:50:21.301Z svc.nsxmanager NSX 1513 SYSTEM [nsx@6876 audit="true"
comp="nsx-manager" level="INFO" subcomp="http"] UserName="admin@10.166.61.80",
ModuleName="ACCESS_CONTROL", Operation="LOGIN", Operation status="success"

2020-07-07T16:43:20.339Z svc.nsxmanager NSX 1513 SYSTEM [nsx@6876 audit="true"
comp="nsx-manager" level="INFO" subcomp="http"] UserName="admin@10.166.61.56",
ModuleName="ACCESS_CONTROL", Operation="LOGIN", Operation status="failure"
```

An example of a syslog message of a policy API call:

```
<182>1 2020-07-06T18:09:14.210Z svc.nsxmanager NSX 2326 FABRIC [nsx@6876 audit="true"
comp="nsx-manager" entId="68d5a9d0-4691-4c9c-94ed-64fd1c96150f" level="INFO" reqId="4c2335aa-
c973-4f74-983f-331a4f7041ca" subcomp="manager" update="true" username="admin"]
UserName="admin", ModuleName="TransportZone", Operation="CreateTransportZone", Operation
status="success", New
value=[{"transport_type":"OVERLAY","host_switch_name":"nsxvswitch","host_switch_mode":"STANDAR
D","nested_nsx":false,"is_default":false,"display_name":"1-
transportzone-1307","_protection":"UNKNOWN"}]
```

An example of syslog messages of CLI access:

```
2020-07-07T16:36:41.783Z svc.nsxmanager NSX 21018 - [nsx@6876 comp="nsx-manager"
subcomp="cli" username="admin" level="INFO"] NSX CLI started (Manager, Policy, Controller)
for user: admin
2020-07-07T16:36:53.469Z svc.nsxmanager NSX 21018 - [nsx@6876 comp="nsx-manager"
subcomp="cli" username="admin" level="INFO"] NSX CLI stopped for user: admin
```

An example of a syslog message when a user runs a CLI command (in this example, set user admin password-expiration 100):

```
<182>1 2020-07-22T20:51:49.017Z manager2 NSX 1864 - [nsx@6876 comp="nsx-manager"
subcomp="cli" username="admin" level="INFO" audit="true"] CMD: set user admin password-
expiration 100 (duration: 2.185s), Operation status: CMD_EXECUTED
```

An example of a syslog message of an NAPI call:

```
<182>1 2020-07-21T21:01:38.803Z manager2 NSX 4690 - [nsx@6876 comp="nsx-manager"
subcomp="node-mgmt" username="admin" level="INFO" audit="true"] admin 'GET /api/v1/node/
services/syslog/exporters' 200 731 "" "PostmanRuntime/7.26.1" 0.004588
```



An example of a syslog message of a CLI command:

```
<182>1 2020-07-21T20:54:40.018Z manager2 NSX 16915 - [nsx@6876 comp="nsx-manager"
subcomp="cli" username="admin" level="INFO" audit="true"] CMD: set logging-server 1.1.1.1
proto udp level info (duration: 4.356s), Operation status: CMD_EXECUTED
```

RFC 5424 and RFC 3164 define the following severity levels:

| Severity Level | Description                              |
|----------------|------------------------------------------|
| 0              | Emergency: system is unusable            |
| 1              | Alert: action must be taken immediately  |
| 2              | Critical: critical conditions            |
| 3              | Error: error conditions                  |
| 4              | Warning: warning conditions              |
| 5              | Notice: normal but significant condition |
| 6              | Informational: informational messages    |
| 7              | Debug: debug-level messages              |

All logs with a severity of emergency, alert, critical, or error contain a unique error code in the structured data portion of the log message. The error code consists of a string and a decimal number. The string represents a specific module.

## Log Message Formats

For more information about RFC 5424, see <https://tools.ietf.org/html/rfc5424>. For more information about RFC 3164, see <https://tools.ietf.org/html/rfc3164>.

RFC 5424 defines the following format for log messages:

```
<facility * 8 + severity> version UTC-TZ hostname APP-NAME procid MSGID [structured-data] msg
```

A sample log message:

```
<187>1 2016-03-15T22:53:00.114Z nsx-manager NSX - SYSTEM [nsx@6876 comp="nsx-manager"
errorCode="MP4039" subcomp="manager"] Connection verification failed for broker
'10.160.108.196'. Marking broker unhealthy.
```

## Error Codes

For a list of error codes, see the knowledge base article 71077 [NSX-T Data Center 2.x Error Codes](#).

## Configure Remote Logging

You can configure NSX-T Data Center appliances and hypervisors to send log messages to a remote log server.

Remote logging is supported on NSX Manager, NSX Edge, and hypervisors. You must configure remote logging on each node individually.

On an KVM host, the NSX-T Data Center installation package automatically configures the rsyslog daemon by putting configuration files in the `/etc/rsyslog.d` directory.

### Prerequisites

- Familiarize yourself with the CLI command `set logging-server`. For more information, see the *NSX-T CLI Reference*.
- If you are using protocols TLS or LI-TLS in NSX CLI to configure a secure connection to a log server, the server and client certificates must be stored in `/image/vmware/nsx/file-store` on each NSX-T Data Center NSX-T appliance. Note that certificates in file store is needed only if the exporter is configured using NSX CLI. If you use API, then there is no need for using the file store. Once you complete the syslog exporter configuration, you must delete all certificates and keys from this location to avoid potential security vulnerabilities.
- To configure a secure connection to a log server, verify that the server is configured with a CA-signed certificates. For example, if you have a Log Insight server `vrli.prome.local` as the log server, you can run the following command from a client to see the certificate chain on the server:

```
root@caserver:~# echo -n | openssl s_client -connect vrli.prome.local:443 | sed -ne '/
^Certificate chain/,/^---/p'
depth=2 C = US, L = California, O = GS, CN = Orange Root Certification Authority
verify error:num=19:self signed certificate in certificate chain
Certificate chain
 0 s:/C=US/ST=California/L=HTG/O=GSS/CN=vrli.prome.local
 i:/C=US/L=California/O=GS/CN=Green Intermediate Certification Authority
 1 s:/C=US/L=California/O=GS/CN=Green Intermediate Certification Authority
 i:/C=US/L=California/O=GS/CN=Orange Root Certification Authority
 2 s:/C=US/L=California/O=GS/CN=Orange Root Certification Authority
 i:/C=US/L=California/O=GS/CN=Orange Root Certification Authority

DONE
```

### Procedure

- 1 To configure remote logging on an NSX-T Data Center appliance, run the following command to configure a log server and the types of messages to send to the log server. Multiple facilities or message IDs can be specified as a comma delimited list, without spaces.

```
set logging-server <hostname-or-ip-address[:port]> proto <proto> level <level> [facility
<facility>] [messageid <messageid>] [serverca <filename>] [clientca <filename>]
[certificate <filename>] [key <filename>] [structured-data <structured-data>]
```

You can run the command multiple times to add multiple configurations. For example:

```
nsx> set logging-server 192.168.110.60 proto udp level info facility local6 messageid
SYSTEM,FABRIC
nsx> set logging-server 192.168.110.60 proto udp level info facility auth,user
```

To forward only audit logs to the remote server, specify `audit="true"` in the `structured-data` parameter. For example:

```
set logging-server <server-ip> proto udp level info structured-data audit="true"
```

- 2 To configure secure remote logging using the protocol LI-TLS, specify the parameter `proto li-tls`. For example:

```
set logging-server vrli.prome.local proto li-tls level info messageid
SWITCHING,ROUTING,FABRIC,SYSTEM,POLICY,HEALTHCHECK,SHA,MONITORING serverca intermed-ca-
full-chain.crt
```

If the configuration is successful, you will get a prompt without any text. To see the content of the server certificate chain (intermediate followed by root), log in as `root` and run the following command:

```
root@nsx1:~# keytool -printcert -file /image/vmware/nsx/file-store/intermed-ca-full-
chain.crt
Certificate[1]:
Owner: CN=Green Intermediate Certification Authority, O=GS, L=California, C=US
Issuer: CN=Orange Root Certification Authority, O=GS, L=California, C=US
Serial number: 3e726e7fbb3b0a7a6b4edd767f867fd2
Valid from: Sun Mar 15 00:00:00 UTC 2020 until: Mon Mar 17 00:00:00 UTC 2025
Certificate fingerprints:
 MD5: 94:C8:9F:92:56:60:EB:DB:ED:4B:11:17:33:27:C0:C9
 SHA1: 42:9C:3C:51:E8:8E:AC:2E:5E:62:95:82:D7:22:E0:FB:08:B8:64:29
 SHA256:
58:B8:63:3D:0C:34:35:39:FC:3D:1E:BA:AA:E3:CE:A9:C0:F3:58:53:1F:AD:89:A5:01:0D:D3:89:9E:7B:C
5:69
Signature algorithm name: SHA256WITHRSA
Subject Public Key Algorithm: 4096-bit RSA key
Version: 3
Certificate[2]:
Owner: CN=Orange Root Certification Authority, O=GS, L=California, C=US
Issuer: CN=Orange Root Certification Authority, O=GS, L=California, C=US
Serial number: 3e726e7fbb3b0a7a6b4edd767f867fd1
Valid from: Mon Mar 16 07:16:07 UTC 2020 until: Fri Mar 10 07:16:07 UTC 2045
Certificate fingerprints:
 MD5: ED:AC:F1:7F:88:05:83:2A:83:C0:09:03:D5:00:CA:7B
 SHA1: DC:B5:3F:37:DF:BD:E0:5C:A4:B7:F4:4C:96:12:75:7A:16:C7:61:37
 SHA256:
F2:5B:DE:8A:F2:31:9D:E6:EF:35:F1:30:6F:DA:05:FF:92:B4:15:96:AA:82:67:E3:3C:C1:69:A3:E5:27:B
9:A5
Signature algorithm name: SHA256WITHRSA
Subject Public Key Algorithm: 4096-bit RSA key
Version: 3
```

The logs for both successful and failure conditions are in `/var/log/loginsight-agent/liagent_2020-MM-DD-<file-num>.log`. If the configuration is successful, you can view the Log Insight configuration with the following command:

```
root@nsx1:/image/vmware/nsx/file-store# cat /var/lib/loginsight-agent/liagent-effective.ini
; Dynamic file representing the effective configuration of VMware Log Insight Agent
(merged server-side and client-side configuration)
; DO NOT EDIT THIS FILE BY HAND -- YOUR CHANGES WILL BE OVERWRITTEN
; Creation time: 2020-03-22T19:41:21.648800

[server]
hostname=vrli.prome.local
proto=cfapi
ssl=yes
ssl_ca_path=/config/vmware/nsx-node-api/syslog/bb466082-996f-4d77-b6e3-1fa93f4a20d4_ca.pem
ssl_accept_any_trusted=yes
port=9543
filter={filelog; nsx-syslog; pri_severity <= 6 and (msgid == "SWITCHING" or msgid ==
"ROUTING" or msgid == "FABRIC" or msgid == "SYSTEM" or msgid == "POLICY" or msgid ==
"HEALTHCHECK" or msgid == "SHA" or msgid == "MONITORING") }

[filelog|nsx-syslog]
directory=/var/log
include=syslog;syslog.*
parser=nsx-syslog_parser

[parser|nsx-syslog_parser]
base_parser=syslog
extract_sd=yes

[update]
auto_update=no
```

- 3 To configure secure remote logging using the protocol TLS, specify the parameter `proto` `tls`. For example:

```
set logging-server vrli.prome.local proto tls level info serverca Orange-CA.crt.pem
clientca Orange-CA.crt.pem certificate gc-nsxt-mgr-full.crt.pem key gc-nsxt-mgr.key.pem
```

Note the following:

- For the `serverCA` parameter, only the root certificate is required, not the full chain.
- If `clientCA` is different from `serverCA`, only the root certificate is required.
- The certificate should hold the full chain of the NSX Manager (they should be NDcPP compliant - ECU, BASIC and CDP (CDP - this check can be ignored))

You can inspect the content of each certificate. For example::

```
root@gc3:~# keytool -printcert -file /image/vmware/nsx/file-store/Orange-CA.crt.pem
Owner: CN=Orange Root Certification Authority, O=GS, L=California, C=US
Issuer: CN=Orange Root Certification Authority, O=GS, L=California, C=US
Serial number: 3e726e7fbb3b0a7a6b4edd767f867fd1
```

```

Valid from: Mon Mar 16 07:16:07 UTC 2020 until: Fri Mar 10 07:16:07 UTC 2045
Certificate fingerprints:
 MD5: ED:AC:F1:7F:88:05:83:2A:83:C0:09:03:D5:00:CA:7B
 SHA1: DC:B5:3F:37:DF:BD:E0:5C:A4:B7:F4:4C:96:12:75:7A:16:C7:61:37
 SHA256:
F2:5B:DE:8A:F2:31:9D:E6:EF:35:F1:30:6F:DA:05:FF:92:B4:15:96:AA:82:67:E3:3C:C1:69:A3:E5:27:B
9:A5
Signature algorithm name: SHA256WITHRSA
Subject Public Key Algorithm: 4096-bit RSA key
Version: 3
root@gc3:~#

root@gc3:/image/vmware/nsx/file-store# keytool -printcert -file gc-nsxt-mgr-full.crt.pem
Certificate[1]:
Owner: CN=gc.prome.local, O=GS, L=HTG, ST=California, C=US
Issuer: CN=Green Intermediate Certification Authority, O=GS, L=California, C=US
Serial number: bdf43ab31340b87f323b438a2895a075
Valid from: Mon Mar 16 07:26:51 UTC 2020 until: Wed Mar 16 07:26:51 UTC 2022
Certificate fingerprints:
 MD5: 36:3C:1F:57:96:07:84:C0:6D:B7:33:9A:8D:25:4D:27
 SHA1: D1:4E:F9:45:2D:0D:34:79:D2:B4:FA:65:28:E0:5C:DC:74:50:CA:3B
 SHA256:
3C:FF:A9:5D:AA:68:44:44:DD:07:2F:DD:E2:BE:9C:32:19:7A:03:D5:26:8D:5F:AD:56:CA:D2:6C:91:96:2
7:6F
Signature algorithm name: SHA256WITHRSA
Subject Public Key Algorithm: 4096-bit RSA key
Version: 3
Certificate[2]:
Owner: CN=Green Intermediate Certification Authority, O=GS, L=California, C=US
Issuer: CN=Orange Root Certification Authority, O=GS, L=California, C=US
Serial number: 3e726e7fbb3b0a7a6b4edd767f867fd2
Valid from: Sun Mar 15 00:00:00 UTC 2020 until: Mon Mar 17 00:00:00 UTC 2025
Certificate fingerprints:
 MD5: 94:C8:9F:92:56:60:EB:DB:ED:4B:11:17:33:27:C0:C9
 SHA1: 42:9C:3C:51:E8:8E:AC:2E:5E:62:95:82:D7:22:E0:FB:08:B8:64:29
 SHA256:
58:B8:63:3D:0C:34:35:39:FC:3D:1E:BA:AA:E3:CE:A9:C0:F3:58:53:1F:AD:89:A5:01:0D:D3:89:9E:7B:C
5:69
Signature algorithm name: SHA256WITHRSA
Subject Public Key Algorithm: 4096-bit RSA key
Version: 3
Certificate[3]:
Owner: CN=Orange Root Certification Authority, O=GS, L=California, C=US
Issuer: CN=Orange Root Certification Authority, O=GS, L=California, C=US
Serial number: 3e726e7fbb3b0a7a6b4edd767f867fd1
Valid from: Mon Mar 16 07:16:07 UTC 2020 until: Fri Mar 10 07:16:07 UTC 2045
Certificate fingerprints:
 MD5: ED:AC:F1:7F:88:05:83:2A:83:C0:09:03:D5:00:CA:7B
 SHA1: DC:B5:3F:37:DF:BD:E0:5C:A4:B7:F4:4C:96:12:75:7A:16:C7:61:37
 SHA256:
F2:5B:DE:8A:F2:31:9D:E6:EF:35:F1:30:6F:DA:05:FF:92:B4:15:96:AA:82:67:E3:3C:C1:69:A3:E5:27:B
9:A5
Signature algorithm name: SHA256WITHRSA
Subject Public Key Algorithm: 4096-bit RSA key
Version: 3

```

## Examples of successful logging in /var/log/syslog::

```
<182>1 2020-03-22T21:54:34.501Z gc3.prome.local NSX 5187 - [nsx@6876 comp="nsx-manager"
subcomp="node-mgmt" username="admin" level="INFO"] Successfully created CA PEM file /
config/vmwarensx-node-api/syslog/92a78d8a-acfd-4515-b05a-2927b70ae920_ca.pem for logging
server vrli.prome.local:6514
<182>1 2020-03-22T21:54:36.269Z gc3.prome.local NSX 5187 - [nsx@6876 comp="nsx-manager"
subcomp="node-mgmt" username="admin" level="INFO"] Successfully created client CA PEM
file /config/vmwarensx-node-api/syslog/92a78d8a-acfd-4515-b05a-2927b70ae920_client_ca.pem
for logging server vrli.prome.local:6514
<182>1 2020-03-22T21:54:36.495Z gc3.prome.local NSX 5187 - [nsx@6876 comp="nsx-manager"
subcomp="node-mgmt" username="root" level="INFO"] cert issuer = /C=US/L=California/O=GS/
CN=Green IntermediateCertification Authority
<182>1 2020-03-22T21:54:36.514Z gc3.prome.local NSX 5187 - [nsx@6876 comp="nsx-manager"
subcomp="node-mgmt" username="root" level="INFO"] cert subject = /C=US/ST=California/L=HTG/
O=GS/CN=gc.promelocal
<182>1 2020-03-22T21:54:36.539Z gc3.prome.local NSX 5187 - [nsx@6876 comp="nsx-manager"
subcomp="node-mgmt" username="root" level="INFO"] certificate trust check succeeded.
status: 200, result: {'status': 'OK'}
<182>1 2020-03-22T21:54:36.612Z gc3.prome.local NSX 5187 - [nsx@6876 comp="nsx-manager"
subcomp="node-mgmt" username="root" level="INFO"] Certificate already exists, skip import
<182>1 2020-03-22T21:54:37.322Z gc3.prome.local NSX 5187 - [nsx@6876 comp="nsx-manager"
subcomp="node-mgmt" username="admin" level="INFO"] Successfully created certificate PEM
file /config/vmwarensx-node-api/syslog/92a78d8a-acfd-4515-b05a-2927b70ae920_cert.pem for
logging server vrli.prome.local:6514
<182>1 2020-03-22T21:54:38.020Z gc3.prome.local NSX 5187 - [nsx@6876 comp="nsx-manager"
subcomp="node-mgmt" username="admin" level="INFO"] Successfully created key PEM file /
config/vmwarensx-node-api/syslog/92a78d8a-acfd-4515-b05a-2927b70ae920_key.pem for logging
server vrli.prome.local:6514
```

## Examples of logging failure in /var/log/syslog::

```
<182>1 2020-03-22T21:33:30.424Z gc3.prome.local NSX 5187 - [nsx@6876 comp="nsx-manager"
subcomp="node-mgmt" username="admin" level="INFO"] Successfully created client CA PEM
file /config/vmwarensx-node-api/syslog/76332782-1ec6-483a-95d4-2adeaf2ef112_client_ca.pem
for logging server vrli.prome.local:6514
<182>1 2020-03-22T21:33:30.779Z gc3.prome.local NSX 5187 - [nsx@6876 comp="nsx-manager"
subcomp="node-mgmt" username="root" level="INFO"] cert issuer = /C=US/L=California/O=GS/
CN=Green IntermediateCertification Authority
<182>1 2020-03-22T21:33:30.803Z gc3.prome.local NSX 5187 - [nsx@6876 comp="nsx-manager"
subcomp="node-mgmt" username="root" level="INFO"] cert subject = /C=US/ST=California/L=HTG/
O=GS/CN=gc.promelocal
<179>1 2020-03-22T21:33:30.823Z gc3.prome.local NSX 5187 - [nsx@6876 comp="nsx-
manager" subcomp="node-mgmt" username="root" level="ERROR" errorCode="NODE10"]
Certificate trust check failed. status:200, result: {'error_message': 'Certificate
CN=gc.prome.local,O=GS,L=HTG,ST=California,C=US was not verifiably signed by
CN=gc.prome.local,O=GS,L=HTG,ST=California,C=US: certificate does not verifywith supplied
key', 'status': 'ERROR'}
<179>1 2020-03-22T21:33:30.824Z gc3.prome.local NSX 5187 - [nsx@6876 comp="nsx-
manager" subcomp="node-mgmt" username="admin" level="ERROR" errorCode="NODE10"]
Failed to create certificate PEM file config/vmware/nsx-node-api/syslog/
76332782-1ec6-483a-95d4-2adeaf2ef112_cert.pem for logging server vrli.prome.local:6514
<182>1 2020-03-22T21:33:31.578Z gc3.prome.local NSX 5187 - [nsx@6876 comp="nsx-manager"
subcomp="node-mgmt" username="admin" level="INFO"] Successfully deleted CA PEM file /
```

```
config/vmwarensx-node-api/syslog/76332782-1ec6-483a-95d4-2adeaf2ef112_ca.pem
<182>1 2020-03-22T21:33:32.342Z gc3.prome.local NSX 5187 - [nsx@6876 comp="nsx-manager"
subcomp="node-mgmt" username="admin" level="INFO"] Successfully deleted client CA PEM
file /config/vmwarensx-node-api/syslog/76332782-1ec6-483a-95d4-2adeaf2ef112_ca.pem
<182>1 2020-03-22T21:33:32.346Z gc3.prome.local NSX 16698 - [nsx@6876 comp="nsx-cli"
subcomp="node-mgmt" username="admin" level="INFO" audit="true"] CMD: set logging-server
vrli.prome.local prototls level info serverca Orange-CA.crt.pem clientca Orange-CA.crt.pem
certifi
cate gc-nsxt-mgr.crt.pem key gc-nsxt-mgr.key.pem (duration: 6.365s), Operation status:
CMD_EXECUTED
```

You can check if the certificate and private key match with the following command. If they match, the output will be writing RSA key. Any other output means they do not match. For example:

```
root@caserter:~/server-certs# diff <(openssl x509 -in certs/gc-nsxt-mgr.crt.pem -pubkey
-noout) <(openssl rsa -in private/gc-nsxt-mgr.key.pem -pubout)
writing RSA key
```

Example of a corrupt private key:

```
root@caserter:~/server-certs# diff <(openssl x509 -in certs/gc-nsxt-mgr.crt.pem -pubkey
-noout) <(openssl rsa -in private/gc-nsxt-mgr-corrupt.key.pem -pubout)
unable to load Private Key
140404188370584:error:0D07209B:asn1 encoding routines:ASN1_get_object:too
long:asn1_lib.c:147:
140404188370584:error:0D068066:asn1 encoding routines:ASN1_CHECK_TLEN:bad object
header:tasn_dec.c:1205:
140404188370584:error:0D07803A:asn1 encoding routines:ASN1_ITEM_EX_D2I:nested asn1
error:tasn_dec.c:386:Type=RSA
140404188370584:error:04093004:rsa routines:OLD_RSA_PRIV_DECODE:RSA lib:rsa_ameth.c:119:
140404188370584:error:0D07209B:asn1 encoding routines:ASN1_get_object:too
long:asn1_lib.c:147:
140404188370584:error:0D068066:asn1 encoding routines:ASN1_CHECK_TLEN:bad object
header:tasn_dec.c:1205:
140404188370584:error:0D07803A:asn1 encoding routines:ASN1_ITEM_EX_D2I:nested asn1
error:tasn_dec.c:386:Type=PKCS8_PRIV_KEY_INFO
140404188370584:error:0907B00D:PEM routines:PEM_READ_BIO_PRIVATEKEY:ASN1
lib:pem_pkey.c:141:
1,14d0
< -----BEGIN PUBLIC KEY-----
< MIIICjANBgkqhkiG9w0BAQEFAAOCAg8AMIICGKCAgEAv3yH7pZidfkLrEP3zVa9
< EcOKXlFFjkThZRZMfguenlm8s6QHYVvuUX8IRB48Li3/DufOj0bzaPWktpv+Q2P0
< N/j4LoX2RzjV/DPxYfLP6GMNMc21L3s9ruBeWUtthtUP8khCwd2d2rZ09cUZVl0P9
< kIYBb5RMFC7Z1Outh3bKdepEf+sXz3DaKZ/WySzYq9x86QDaA3ABO3Q0i7txBscI
< FvXuMDOMQaC3pPp9FWO6IPRAWB57wahLJv6K5qGIfwubSBFg53grT4snf1lDZAHz
< 9hz5JgGr80GVyWyb7rgigpl9iUWAZx8U9De9XoxmvBN5iEGTIuKGaEgICL176crb
< RMkhjncQnHI+z6sQvpyJ7U0zZc72eBIWoHUKcWWk3eU6Oy4OiyW6jYuXG7hZYlly
< nSkme3mZUWJKvcoX05+3zeCP623/HzE7X2sNyWFjzeF3XEvauZrIbsJh/xp2ShDa
< uKKEY0gUGhLtCa3TpV9l8d6tFWVy8XjVjdjoVt4s7MfUo/airVmRykfsWrKyNUOQ
```

```
< qRZvSbqjt8pm+3bSvKdXX4ul7ptPG2GF20ETWHPwjK2JwQpGhR9zK8fsKzvm6hXi
< kq76zi4FefuVps3e1r39+0F+p6d6i2oUoo24sC1iSePTDhU74efVp6iv8HmnDgYX
< Ylm6Kusr0JT5TJFDfASmrj8CAwEAAQ==
< -----END PUBLIC KEY-----
```

Example of a valid private key and certificate but they are not made for each other:

```
root@caserver:~/server-certs# diff <(openssl x509 -in certs/gc-nsxt-mgr.crt.pem -pubkey
-noout) <(openssl rsa -in private/vrli.key.pem -pubout)
writing RSA key
2,13c2,13
< MIICijANBgqhkiG9w0BAQEFAAOCAg8AMIICCgKCAgEAv3yH7pZidfkLrEP3zVa9
< EcOKXlFFjkThZRMfguenlm8s6QHYVvuUX8IRB48Li3/DUfOj0bzaPWktpv+Q2P0
< N/j4LoX2RzjV/DPxYfLP6GMNMc21L3s9ruBeWUthtUP8khCWd2d2rZ09cUZV10P9
< kIYBb5RMFC7Z10Uth3bKdepEf+sXz3DaKZ/WySzYq9x86QDa3ABO3Q0i7txBscI
< FvXuMDOMQaC3pPp9FWO6IPRAWB57wahLJv6K5qGIfwubSBFg53grT4snf11DZAhZ
< 9hz5JgGr80GVyWyb7rgigpl9iUWAZx8U9De9XoxmvBN5iEGTIuKGaEgICL176crb
< RMkhjnCqNHI+z6sQvpYJ7U0zZc72eBIWoHukcWWk3eU6Oy40iyW6jYuXG7hZy1ly
< nSkme3mZUWJKvcoX05+3zeCP623/HzE7X2sNyWFjzeF3XEvauZrIbsJh/xp2ShDa
< uKKEY0gUGhLtCa3TpV918d6tFWVy8XjVjdjoVt4s7MfUo/airVmRykfsWrKyNUOQ
< qRZvSbqjt8pm+3bSvKdXX4ul7ptPG2GF20ETWHPwjK2JwQpGhR9zK8fsKzvm6hXi
< kq76zi4FefuVps3e1r39+0F+p6d6i2oUoo24sC1iSePTDhU74efVp6iv8HmnDgYX
< Ylm6Kusr0JT5TJFDfASmrj8CAwEAAQ==

> MIICijANBgqhkiG9w0BAQEFAAOCAg8AMIICCgKCAgEAqvsjay7+o7gCW7szT3ho
> bc34XX216u5J14/X/pUDI/YHmIf06bsZ1r/14bTL4Q7BM6+9MI6UYEE7DxUoINGO
> o4FEEQE32KWVFe3gw3homHU39q4pQjsJsxTcTE3oDMLIY0nWJ0PRUst3DffYUH1L
> W0NUN9yDN+faL2Uf021iuDqVy9V8AH3ON6fu+QCA8nt71ZkzeTxSA0ldpl2NA17F
> rD8rm05wxnV7WtuV7V8PstISiClzhHgZRM1+B0r300itnyAzEGLaRT3//PKfe0Oe
> HCdxGMLrUtMqxIItJahEsqvMufyqNYecVscYXLHPelizKCsQfy8c08LnznG8VAdc
> YILSn3uYGZap6aF1SgVxsvZicwv1YnssmgE13Af0nScmfM96k9h5joHVEkWK608v
> oT5DGG1kVL2Qly97x0b6EnzUorzivv5zJMKvFcOektR8HdMHQit5uvmMRY3S5zow
> FtvfSDFWxxKyTy6GBRpp+8F+Jq91yGy/qa9lhKBzT2lg+rJp7T8k7/Nm9Tjyx7jL
> EggEKZEL4chxpo8ucF98hbvXWRuaFHC2iDzGuUmuS1FfjVvHTuIbEMQfjapLZrHx
> 8jHfOP/PL+6kPbvNZZ2rTpczuEoGTQFFW9vX48GzIEyMeR6QWpPR0F7r4xak68P5
> 2PJmMveinDhU35IqWEXHAwcCAwEAAQ==
```

- 4 To view the logging configuration, run the `get logging-server` command. For example,

```
nsx> get logging-servers
192.168.110.60 proto udp level info facility local6 messageid SYSTEM,FABRIC
192.168.110.60 proto udp level info facility auth,user
```

- 5 To clear the remote logging configuration, run the following command:

```
nsx> clear logging-servers
```



**6** To configure remote logging on an ESXi host:

- a Run the following commands to configure syslog and send a test message:

```
esxcli network firewall ruleset set -r syslog -e true
esxcli system syslog config set --loghost=udp://<log server IP>:<port> --log-level=info
esxcli system syslog reload
esxcli system syslog mark -s "This is a test message"
```

- b You can run the following command to display the configuration:

```
esxcli system syslog config get
```

**7** To configure remote logging on a KVM host:

- a Edit the file `/etc/rsyslog.d/10-vmware-remote-logging.conf` for your environment.
- b Add the following line to the file:

```
. @<ip>:514;RFC5424fmt
```

- c Run the following command:

```
service rsyslog restart
```

## Log Message IDs

In a log message, the message ID field identifies the type of message. You can use the `messageid` parameter in the `set logging-server` command to filter which log messages are sent to a logging server.

**Table 22-4. Log Message IDs**

| Message ID | Examples                 |
|------------|--------------------------|
| FABRIC     | Host node                |
|            | Host preparation         |
|            | Edge node                |
|            | Transport zone           |
|            | Transport node           |
|            | Uplink profiles          |
|            | Cluster profiles         |
|            | Edge cluster             |
| SWITCHING  | Logical switch           |
|            | Logical switch ports     |
|            | Switching profiles       |
|            | switch security features |

Table 22-4. Log Message IDs (continued)

| Message ID      | Examples                                                                                                                                                                                                                                    |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ROUTING         | Logical router<br>Logical router ports<br>Static routing<br>Dynamic routing<br>NAT                                                                                                                                                          |
| FIREWALL        | Firewall rules<br>Firewall rule sections                                                                                                                                                                                                    |
| FIREWALL-PKTLOG | Firewall connection logs<br>Firewall packet logs                                                                                                                                                                                            |
| GROUPING        | IP sets<br>Mac sets<br>NSGroups<br>NSServices<br>NSService groups<br>VNI Pool<br>IP Pool                                                                                                                                                    |
| DHCP            | DHCP relay                                                                                                                                                                                                                                  |
| SYSTEM          | Appliance management (remote syslog, ntp, etc)<br>Cluster management<br>Trust management<br>Licensing<br>User and roles<br>Task management<br>Install<br>Upgrade (NSX Manager, NSX Edge and host-packages upgrades )<br>Realization<br>Tags |
| MONITORING      | SNMP<br>Port connection<br>Traceflow                                                                                                                                                                                                        |
| -               | All other log messages.                                                                                                                                                                                                                     |

## Troubleshooting Syslog Issues

If logs are not received by the remote log server, perform the following steps.

- Verify the remote log server's IP address.
- Verify that the `level` parameter is configured correctly.
- Verify that the `facility` parameter is configured correctly.
- If the protocol is TLS, set the protocol to UDP to see if there is a certificate mismatch.

- If the protocol is TLS, verify that port 6514 is open on both ends.
- Remove the message ID filter and see if logs are received by the server.
- Restart the rsyslog service with the command `restart service rsyslogd`.

## Configure Serial Logging on an Appliance VM

You can configure serial logging on an appliance VM to capture log messages when the VM crashes.

### Procedure

- 1 Log in to the VM as `root`.
- 2 Edit `/etc/default/grub`.
- 3 Find the parameter `GRUB_CMDLINE_LINUX_DEFAULT` and append `console=ttyS0 console=tty0`.
- 4 Run the command `update-grub2`.
- 5 Verify that the `/boot/grub/grub.cfg` file has the change made in step 3.
- 6 Power off the VM.
- 7 Edit the VM's configuration (`.vmx`) file and add the following lines:

```
serial0.present = "TRUE"
serial0.fileType = "file"
serial0.fileName = "serial.out"
serial0.yieldOnMsrRead = "TRUE"
answer.msg.serial.file.open = "Append"
```

- 8 Power on the VM.

### Results

If a kernel panic occurs in the VM, you can find the file `serial.out` containing log messages at the same location as that of the `.vmx` file.

## Firewall Audit Log Messages

Firewall configuration changes are audited. Below are examples of audit log messages related to these changes.

### Distributed firewall changes in Policy mode

Adding a firewall section (`SecurityPolicy-1`) with a rule (`Rule1_1`):

```
<182>1 2020-08-11T21:58:50.319Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="2aff6b4f-3d4f-4d62-a639-61291f7e879e" splitId="a5mxlu78"
splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin",
ModuleName="DfwSecurityPolicy", Operation="PatchSecurityPolicyForDomain", Operation
status="success", Old
```

```

value=[{"precedence":10,"category":"Application","resource_type":"CommunicationMap","id":"SecurityPolicy-1","display_name":"SecurityPolicy-1","path":"/infra/domains/default/security-policies/SecurityPolicy-1","relative_path":"SecurityPolicy-1","parent_path":"/infra/domains/default","unique_id":"895eeac5-641b-4306-be7f-a43fdd969ee5","marked_for_delete":false,"overridden":false,"_create_user":"admin","_create_time":1597183130247,"_last_modified_user":"admin","_last_modified_time":1597183130247,"_system_owned":false,"_protection":"NOT_PROTECTED","_revision":0}], New value=["default"
"SecurityPolicy-1"
{"resource_type":"SecurityPolicy","id":"SecurityPolicy-1","display_name":"SecurityPolicy-1","path":"/infra/domains/default/security-policies/SecurityPolicy-1","children":[{"Rule":
{"action":"ALLOW","resource_type":"Rule","id":"Rule1_1","display_name":"Rule1_1","path":

<182>1 2020-08-11T21:58:50.320Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="2aff6b4f-3d4f-4d62-a639-61291f7e879e" splitId="a5mxlu78"
splitIndex="2 of 2" subcomp="policy" update="true"] "/infra/domains/default/security-policies/
SecurityPolicy-1/rules/
Rule1_1","marked_for_delete":false,"overridden":false,"sequence_number":10,"sources_excluded":
false,"destinations_excluded":false,"source_groups":["ANY"],"destination_groups":
["ANY"],"services":["ANY"],"profiles":["ANY"],"logged":false,"scope":
["ANY"],"disabled":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","_protection":"UNKNOWN"
},"resource_type":"ChildRule","marked_for_delete":false,"mark_for_override":false,"_protection":
"UNKNOWN"}],"marked_for_delete":false,"overridden":false,"sequence_number":10,"category":"A
pplication","stateful":true,"locked":false,"scope":["ANY"],"_protection":"UNKNOWN"}]

<182>1 2020-08-11T21:58:50.404Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" entId="Rule1_1" level="INFO" reqId="2aff6b4f-3d4f-4d62-a639-61291f7e879e"
splitId="E993J2LF" splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin",
ModuleName="DfwSecurityPolicy", Operation="UpdateSecurityRule", Operation status="success",
Old value=[{"sequence_number":10,"source_groups":["ANY"],"destination_groups":
["ANY"],"services":["ANY"],"action":"ALLOW","logged":false,"scope":
["ANY"],"disabled":false,"direction":"IN_OUT","resource_type":"CommunicationEntry","id":"Rule1
_1","display_name":"Rule1_1","path":"/infra/domains/default/security-policies/
SecurityPolicy-1/rules/Rule1_1","relative_path":"Rule1_1","parent_path":"/infra/domains/
default/security-policies/
SecurityPolicy-1","unique_id":"2024","marked_for_delete":false,"overridden":false,"_create_use
r":"admin","_create_time":1597183130364,"_last_modified_user":"admin","_last_modified_time":15
97183130364,"_system_owned":false,"_protection":"NOT_PROTECTED","_revision":0}], New
value=["default" "SecurityPolicy-1" "Rule1_1"
{"action":"ALLOW","resource_type":"Rule","id":"Rule1_1","display_name":"Rule1_1","path":

<182>1 2020-08-11T21:58:50.404Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" entId="Rule1_1" level="INFO" reqId="2aff6b4f-3d4f-4d62-a639-61291f7e879e"
splitId="E993J2LF" splitIndex="2 of 2" subcomp="policy" update="true"] "/infra/domains/
default/security-policies/SecurityPolicy-1/rules/
Rule1_1","marked_for_delete":false,"overridden":false,"sequence_number":10,"sources_excluded":
false,"destinations_excluded":false,"source_groups":["ANY"],"destination_groups":
["ANY"],"services":["ANY"],"profiles":["ANY"],"logged":false,"scope":
["ANY"],"disabled":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","_protection":"UNKNOWN"
"}]

<182>1 2020-08-11T21:58:50.466Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="2aff6b4f-3d4f-4d62-a639-61291f7e879e" splitId="iMHW1shi"

```

```
splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin", ModuleName="Policy",
Operation="PatchInfra", Operation status="success", New
value=[{"enforce_revision_check":true} {"resource_type":"Infra","children":[{"children":
[{"SecurityPolicy":
{"resource_type":"SecurityPolicy","id":"SecurityPolicy-1","display_name":"SecurityPolicy-1","p
ath":"/infra/domains/default/security-policies/SecurityPolicy-1","children":[{"Rule":
{"action":"ALLOW","resource_type":"Rule","id":"Rule1_1","display_name":"Rule1_1","path":"/
infra/domains/default/security-policies/SecurityPolicy-1/rules/
Rule1_1","marked_for_delete":false,"overridden":false,"sequence_number":10,"sources_excluded":
false,"destinations_excluded":false,"source_groups":["ANY"],"destination_groups":
["ANY"],"services":["ANY"],"profiles":["ANY"],"logged":false,"scope":
["ANY"],"disabled":false,"direction":"IN_OUT","ip_protocol":"IPv4_IPv6","_protection":"UNKNOWN
"},"resource_type":"ChildRule","marked_for_delete":false,"mark_for_override":false,"_protectio
n":"UNKNOWN"}],"marked_for_delete"

<182>1 2020-08-11T21:58:50.466Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="2aff6b4f-3d4f-4d62-a639-61291f7e879e" splitId="iMHW1shi"
splitIndex="2 of 2" subcomp="policy"
update="true"] :false,"overridden":false,"sequence_number":10,"category":"Application","stateful":true,"locked":false,"scope":
["ANY"],"_protection":"UNKNOWN"},"resource_type":"ChildSecurityPolicy","marked_for_delete":false,"mark_for_override":false,"_protection":"UNKNOWN"}],"target_type":"Domain","resource_type":
"ChildResourceReference","id":"default","marked_for_delete":false,"mark_for_override":false,"_protection":"UNKNOWN"}],"marked_for_delete":false,"overridden":false,"_protection":"UNKNOWN",
_revision":-1}]
```

### Updating a rule (from Rule1\_1 to Rule1\_1\_updated) in a section (SecurityPolicy-1):

```
<182>1 2020-08-11T22:22:06.303Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="6aadd8de-d157-4479-b84c-8410dd48c2aa" splitId="mJ7hQGhg"
splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin",
ModuleName="DfwSecurityPolicy", Operation="PatchSecurityPolicyForDomain", Operation
status="success", New value=["default" "SecurityPolicy-1"
{"resource_type":"SecurityPolicy","id":"SecurityPolicy-1","display_name":"SecurityPolicy-1","p
ath":"/infra/domains/default/security-policies/
SecurityPolicy-1","unique_id":"895eeac5-641b-4306-be7f-a43fdd969ee5","children":[{"Rule":
{"action":"ALLOW","resource_type":"Rule","id":"Rule1_1","display_name":"Rule1_1_updated","path
":"/infra/domains/default/security-policies/SecurityPolicy-1/rules/
Rule1_1","unique_id":"2024","marked_for_delete":false,"overridden":false,"rule_id":2024,"seque
nce_number":10,"sources_excluded":false,"destinations_excluded":false,"source_groups":
["ANY"],"destination_groups":["ANY"],"services":["ANY"],"profiles":
["ANY"],"logged":false,"scope":
["ANY"],"disabled":false,"direction":"IN_OUT","ip_protocol":"IPv4_IPv6","is_default":false,"_p
rotection":"UNKNOWN",_revision":0},"resource_type":
<182>1 2020-08-11T22:22:06.303Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="6aadd8de-d157-4479-b84c-8410dd48c2aa" splitId="mJ7hQGhg"
splitIndex="2 of 2" subcomp="policy" update="true"]
"ChildRule","marked_for_delete":false,"mark_for_override":false,"_protection":"UNKNOWN"}],"mar
ked_for_delete":false,"overridden":false,"sequence_number":10,"internal_sequence_number":13000
010,"category":"Application","stateful":true,"tcp_strict":true,"locked":false,"lock_modified_t
ime":0,"scope":["ANY"],"is_default":false,"_protection":"UNKNOWN",_revision":0}]

<182>1 2020-08-11T22:22:06.324Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" entId="Rule1_1" level="INFO" reqId="6aadd8de-d157-4479-b84c-8410dd48c2aa"
splitId="JKVilI6n" splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin",
```

```

ModuleName="DfwSecurityPolicy", Operation="UpdateSecurityRule", Operation status="success",
Old value=[{"sequence_number":10,"source_groups":["ANY"],"destination_groups":
["ANY"],"services":["ANY"],"action":"ALLOW","logged":false,"scope":
["ANY"],"disabled":false,"direction":"IN_OUT","resource_type":"CommunicationEntry","id":"Rule1
_1","display_name":"Rule1_1","path":"/infra/domains/default/security-policies/
SecurityPolicy-1/rules/Rule1_1","relative_path":"Rule1_1","parent_path":"/infra/domains/
default/security-policies/
SecurityPolicy-1","unique_id":"2024","marked_for_delete":false,"overridden":false,"_create_use
r":"admin","_create_time":1597183130364,"_last_modified_user":"admin","_last_modified_time":15
97183130369,"_system_owned":false,"_protection":"NOT_PROTECTED","_revision":0}], New
value=["default" "SecurityPolicy-1" "Rule1_1"
{"action":"ALLOW","resource_type":"Rule","id":"Rule1_1","display_name":"Rule1_1_updated","path
":

```

```

<182>1 2020-08-11T22:22:06.324Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" entId="Rule1_1" level="INFO" reqId="6aadd8de-d157-4479-b84c-8410dd48c2aa"
splitId="JKVilI6n" splitIndex="2 of 2" subcomp="policy" update="true"] "/infra/domains/
default/security-policies/SecurityPolicy-1/rules/
Rule1_1","unique_id":"2024","marked_for_delete":false,"overridden":false,"rule_id":2024,"seque
nce_number":10,"sources_excluded":false,"destinations_excluded":false,"source_groups":
["ANY"],"destination_groups":["ANY"],"services":["ANY"],"profiles":
["ANY"],"logged":false,"scope":
["ANY"],"disabled":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","is_default":false,"_p
rotection":"UNKNOWN","_revision":0}]

```

```

<182>1 2020-08-11T22:22:06.363Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="6aadd8de-d157-4479-b84c-8410dd48c2aa" splitId="9MtbEpd8"
splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin", ModuleName="Policy",
Operation="PatchInfra", Operation status="success", New
value=[{"enforce_revision_check":true} {"resource_type":"Infra","children":[{"children":
[{"SecurityPolicy":
{"resource_type":"SecurityPolicy","id":"SecurityPolicy-1","display_name":"SecurityPolicy-1","p
ath":"/infra/domains/default/security-policies/
SecurityPolicy-1","unique_id":"895eeac5-641b-4306-be7f-a43fdd969ee5","children":[{"Rule":
{"action":"ALLOW","resource_type":"Rule","id":"Rule1_1","display_name":"Rule1_1_updated","path
":"/infra/domains/default/security-policies/SecurityPolicy-1/rules/
Rule1_1","unique_id":"2024","marked_for_delete":false,"overridden":false,"rule_id":2024,"seque
nce_number":10,"sources_excluded":false,"destinations_excluded":false,"source_groups":
["ANY"],"destination_groups":["ANY"],"services":["ANY"],"profiles":
["ANY"],"logged":false,"scope":
["ANY"],"disabled":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","is_default":false,"_p
rotection":"UNKNOWN","_revision":

```

```

<182>1 2020-08-11T22:22:06.363Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="6aadd8de-d157-4479-b84c-8410dd48c2aa" splitId="9MtbEpd8"
splitIndex="2 of 2" subcomp="policy" update="true"]
0},"resource_type":"ChildRule","marked_for_delete":false,"mark_for_override":false,"_protectio
n":"UNKNOWN"}],"marked_for_delete":false,"overridden":false,"sequence_number":10,"internal_seq
uence_number":13000010,"category":"Application","stateful":true,"tcp_strict":true,"locked":fal
se,"lock_modified_time":0,"scope":
["ANY"],"is_default":false,"_protection":"UNKNOWN","_revision":0},"resource_type":"ChildSecuri
tyPolicy","marked_for_delete":false,"mark_for_override":false,"_protection":"UNKNOWN"}],"targe
t_type":"Domain","resource_type":"ChildResourceReference","id":"default","marked_for_delete":f
alse,"mark_for_override":false,"_protection":"UNKNOWN"}],"marked_for_delete":false,"overridden
":false,"_protection":"UNKNOWN","_revision":-1}]

```

## Deleting a rule (Rule1\_2) from a section (SecurityPolicy-1):

```
<182>1 2020-08-11T22:12:24.444Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" entId="Rule1_2" level="INFO" reqId="1a58e753-460c-443f-8a28-0d40d8af9b76"
subcomp="policy" update="true" username="admin"] UserName="admin",
ModuleName="DfwSecurityPolicy", Operation="DeleteSecurityRule", Operation status="success",
Old value=[{"sequence_number":20,"source_groups":["ANY"],"destination_groups":
["ANY"],"services":["ANY"],"action":"ALLOW","logged":false,"scope":
["ANY"],"disabled":false,"direction":"IN_OUT","resource_type":"CommunicationEntry","id":"Rule1
_2","display_name":"Rule1_2","path":"/infra/domains/default/security-policies/
SecurityPolicy-1/rules/Rule1_2","relative_path":"Rule1_2","parent_path":"/infra/domains/
default/security-policies/
SecurityPolicy-1","unique_id":"2026","marked_for_delete":false,"overridden":false,"_create_use
r":"admin","_create_time":1597183904580,"_last_modified_user":"admin","_last_modified_time":15
97183904582,"_system_owned":false,"_protection":"NOT_PROTECTED","_revision":0}], New
value=["default" "SecurityPolicy-1" "Rule1_2"]
```

```
<182>1 2020-08-11T22:12:24.463Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="1a58e753-460c-443f-8a28-0d40d8af9b76" splitId="hoDI5YJQ"
splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin",
ModuleName="DfwSecurityPolicy", Operation="PatchSecurityPolicyForDomain", Operation
status="success", New value=["default" "SecurityPolicy-1"
{"resource_type":"SecurityPolicy","id":"SecurityPolicy-1","display_name":"SecurityPolicy-1","p
ath":"/infra/domains/default/security-policies/
SecurityPolicy-1","unique_id":"895eeac5-641b-4306-be7f-a43fdd969ee5","children":[{"Rule":
{"resource_type":"Rule","id":"Rule1_2","path":"/infra/domains/default/security-policies/
SecurityPolicy-1/rules/
Rule1_2","marked_for_delete":true,"overridden":false,"sources_excluded":false,"destinations_ex
cluded":false,"logged":false,"disabled":false,"direction":"IN_OUT","_protection":"UNKNOWN"},"r
esource_type":"ChildRule","marked_for_delete":true,"mark_for_override":false,"_protection":"UN
KNOWN"}]], "marked_for_delete":false,"overridden":false,"sequence_number":10,"internal_sequence_
number":13000010,"category":"Application","stateful":true,"tcp_strict":true,"locked":false,"lo
ck_modified_time":0,"scope":
```

```
<182>1 2020-08-11T22:12:24.463Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="1a58e753-460c-443f-8a28-0d40d8af9b76" splitId="hoDI5YJQ"
splitIndex="2 of 2" subcomp="policy" update="true"]
["ANY"],"is_default":false,"_protection":"UNKNOWN","_revision":0}]
```

```
<182>1 2020-08-11T22:12:24.497Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="1a58e753-460c-443f-8a28-0d40d8af9b76" splitId="mxpzQHfF"
splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin", ModuleName="Policy",
Operation="PatchInfra", Operation status="success", New
value=[{"enforce_revision_check":true} {"resource_type":"Infra","children":[{"children":
{"SecurityPolicy":
{"resource_type":"SecurityPolicy","id":"SecurityPolicy-1","display_name":"SecurityPolicy-1","p
ath":"/infra/domains/default/security-policies/
SecurityPolicy-1","unique_id":"895eeac5-641b-4306-be7f-a43fdd969ee5","children":[{"Rule":
{"resource_type":"Rule","id":"Rule1_2","path":"/infra/domains/default/security-policies/
SecurityPolicy-1/rules/
Rule1_2","marked_for_delete":true,"overridden":false,"sources_excluded":false,"destinations_ex
cluded":false,"logged":false,"disabled":false,"direction":"IN_OUT","_protection":"UNKNOWN"},"r
esource_type":"ChildRule","marked_for_delete":true,"mark_for_override":false,"_protection":"UN
KNOWN"}]], "marked_for_delete":false,"overridden":false,"sequence_number":10,"internal_sequence_
```

```
number":13000010,"category":"Application","stateful":true,"tcp_strict":
<182>1 2020-08-11T22:12:24.497Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="1a58e753-460c-443f-8a28-0d40d8af9b76" splitId="mxpzQHfF"
splitIndex="2 of 2" subcomp="policy" update="true"]
true,"locked":false,"lock_modified_time":0,"scope":
["ANY"],"is_default":false,"_protection":"UNKNOWN","_revision":0},"resource_type":"ChildSecuri
tyPolicy","marked_for_delete":false,"mark_for_override":false,"_protection":"UNKNOWN"}],"targe
t_type":"Domain","resource_type":"ChildResourceReference","id":"default","marked_for_delete":f
alse,"mark_for_override":false,"_protection":"UNKNOWN"}],"marked_for_delete":false,"overridden
":false,"_protection":"UNKNOWN","_revision":-1}]
```

Deleting a section (SecurityPolicy-1) that contains a rule (Rule1\_1):

```
<182>1 2020-08-11T22:24:24.898Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" entId="SecurityPolicy-1" level="INFO" reqId="724b5494-10cd-4124-a431-56ba7d922bbf"
splitId="4WIXz9qL" splitIndex="1 of 2" subcomp="policy" update="true" username="admin"]
UserName="admin", ModuleName="DfwSecurityPolicy", Operation="DeleteSecurityPolicyForDomain",
Operation status="success", Old
value=[{"precedence":10,"category":"Application","resource_type":"CommunicationMap","id":"Secu
rityPolicy-1","display_name":"SecurityPolicy-1","path":"/infra/domains/default/security-
policies/SecurityPolicy-1","relative_path":"SecurityPolicy-1","parent_path":"/infra/domains/
default","unique_id":"895eeac5-641b-4306-be7f-
a43fdd969ee5","marked_for_delete":false,"overridden":false,"_create_user":"admin","_create_tim
e":1597183130247,"_last_modified_user":"admin","_last_modified_time":1597183130251,"_system_ow
ned":false,"_protection":"NOT_PROTECTED","_revision":0}{"sequence_number":10,"source_groups":
["ANY"],"destination_groups":["ANY"],"services":
["ANY"],"action":"ALLOW","logged":false,"scope":
["ANY"],"disabled":false,"direction":"IN_OUT","resource_type":"CommunicationEntry","id":"Rule1
_1","display_name":"Rule1_1_updated","path":
```

```
<182>1 2020-08-11T22:24:24.898Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" entId="SecurityPolicy-1" level="INFO" reqId="724b5494-10cd-4124-a431-56ba7d922bbf"
splitId="4WIXz9qL" splitIndex="2 of 2" subcomp="policy" update="true" username="admin"] "/"
infra/domains/default/security-policies/SecurityPolicy-1/rules/
Rule1_1","relative_path":"Rule1_1","parent_path":"/infra/domains/default/security-policies/
SecurityPolicy-1","unique_id":"2024","marked_for_delete":false,"overridden":false,"_create_use
r":"admin","_create_time":1597183130364,"_last_modified_user":"admin","_last_modified_time":15
97184526313,"_system_owned":false,"_protection":"NOT_PROTECTED","_revision":1}], New
value=["default" "SecurityPolicy-1"]
```

```
<182>1 2020-08-11T22:24:24.938Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="724b5494-10cd-4124-a431-56ba7d922bbf" subcomp="policy"
update="true"] UserName="admin", ModuleName="Policy", Operation="PatchInfra", Operation
status="success", New value=[{"enforce_revision_check":true}
{"resource_type":"Infra","children":[{"children":[{"SecurityPolicy":
{"resource_type":"SecurityPolicy","id":"SecurityPolicy-1","path":"/infra/domains/default/
security-policies/
SecurityPolicy-1","marked_for_delete":true,"overridden":false,"locked":false,"_protection":"UN
KNOWN"},"resource_type":"ChildSecurityPolicy","marked_for_delete":true,"mark_for_override":fal
se,"_protection":"UNKNOWN"}],"target_type":"Domain","resource_type":"ChildResourceReference","
id":"default","marked_for_delete":false,"mark_for_override":false,"_protection":"UNKNOWN"}],"m
arked_for_delete":false,"overridden":false,"_protection":"UNKNOWN","_revision":-1}]
```



## Gateway firewall changes in Policy mode

Note that log messages for a tier-0 gateway and a tier-1 gateway are similar.

Adding a section (T1-Policies) with a rule (myT1\_Rule1) for a tier-1 gateway (myT1):

```
<182>1 2020-08-11T22:31:26.800Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="c2790fbd-db29-46d3-9a0e-1003455ee9ea" splitId="Ta8faYzQ"
splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin",
ModuleName="PolicyEdgeFirewall", Operation="PatchGatewayPolicyForDomain", Operation
status="success", Old
value=[{"precedence":10,"category":"LocalGatewayRules","resource_type":"CommunicationMap","id":
"T1-Policies","display_name":"T1-Policies","path":"/infra/domains/default/gateway-
policies/T1-Policies","relative_path":"T1-Policies","parent_path":"/infra/domains/
default","unique_id":"a73c1345-6b4e-43e0-
b4ee-9a91c7ba9df6","marked_for_delete":false,"overridden":false,"_create_user":"admin","_creat
e_time":1597185086789,"_last_modified_user":"admin","_last_modified_time":1597185086789,"_syst
em_owed":false,"_protection":"NOT_PROTECTED","_revision":0}], New value=[{"default" "T1-
Policies" {"resource_type":"GatewayPolicy","id":"T1-Policies","display_name":"T1-
Policies","path":"/infra/domains/default/gateway-policies/T1-Policies","children":[{"Rule":
{"action":"ALLOW","resource_type":"Rule","id":"myT1_Rule1","display_name":"myT1_Rule1","path":

<182>1 2020-08-11T22:31:26.801Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="c2790fbd-db29-46d3-9a0e-1003455ee9ea" splitId="Ta8faYzQ"
splitIndex="2 of 2" subcomp="policy" update="true"] "/infra/domains/default/gateway-
policies/T1-Policies/rules/
myT1_Rule1","marked_for_delete":false,"overridden":false,"sequence_number":10,"sources_exclude
d":false,"destinations_excluded":false,"source_groups":["ANY"],"destination_groups":
["ANY"],"services":["ANY"],"profiles":["ANY"],"logged":false,"scope":["/infra/tier-1s/
myT1"],"disabled":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","_protection":"UNKNOWN"
},"resource_type":"ChildRule","marked_for_delete":false,"mark_for_override":false,"_protection
":"UNKNOWN"}],"marked_for_delete":false,"overridden":false,"sequence_number":10,"category":"Lo
calGatewayRules","stateful":true,"locked":false,"_protection":"UNKNOWN"}]

<182>1 2020-08-11T22:31:26.878Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="c2790fbd-db29-46d3-9a0e-1003455ee9ea" splitId="aZfgiFKt"
splitIndex="1 of 2" subcomp="policy" update="true" username="admin"] UserName="admin",
ModuleName="PolicyEdgeFirewall", Operation="PatchGatewayRule", Operation status="success",
Old value=[{"sequence_number":10,"source_groups":["ANY"],"destination_groups":
["ANY"],"services":["ANY"],"action":"ALLOW","logged":false,"scope":["/infra/tier-1s/
myT1"],"disabled":false,"direction":"IN_OUT","resource_type":"CommunicationEntry","id":"myT1_R
ule1","display_name":"myT1_Rule1","path":"/infra/domains/default/gateway-policies/T1-Policies/
rules/myT1_Rule1","relative_path":"myT1_Rule1","parent_path":"/infra/domains/default/gateway-
policies/T1-
Policies","unique_id":"2028","marked_for_delete":false,"overridden":false,"_create_user":"admi
n","_create_time":1597185086809,"_last_modified_user":"admin","_last_modified_time":1597185086
809,"_system_owed":false,"_protection":"NOT_PROTECTED","_revision":0}], New value=[{"default"
"T1-Policies" "myT1_Rule1"
{"action":"ALLOW","resource_type":"Rule","id":"myT1_Rule1","display_name":"myT1_Rule1","path":

<182>1 2020-08-11T22:31:26.878Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="c2790fbd-db29-46d3-9a0e-1003455ee9ea" splitId="aZfgiFKt"
splitIndex="2 of 2" subcomp="policy" update="true" username="admin"] "/infra/domains/default/
gateway-policies/T1-Policies/rules/
myT1_Rule1","marked_for_delete":false,"overridden":false,"sequence_number":10,"sources_exclude
```

```
d":false,"destinations_excluded":false,"source_groups":["ANY"],"destination_groups":
["ANY"],"services":["ANY"],"profiles":["ANY"],"logged":false,"scope":["/infra/tier-1s/
myT1"],"disabled":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","_protection":"UNKNOWN"
}}

<182>1 2020-08-11T22:31:26.890Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="c2790fbd-db29-46d3-9a0e-1003455ee9ea" splitId="0s7tdCjN"
splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin", ModuleName="Policy",
Operation="PatchInfra", Operation status="success", New
value=[{"enforce_revision_check":true} {"resource_type":"Infra","children":[{"children":
[{"GatewayPolicy":{"resource_type":"GatewayPolicy","id":"T1-Policies","display_name":"T1-
Policies","path":"/infra/domains/default/gateway-policies/T1-Policies","children":[{"Rule":
{"action":"ALLOW","resource_type":"Rule","id":"myT1_Rule1","display_name":"myT1_Rule1","path":
"/infra/domains/default/gateway-policies/T1-Policies/rules/
myT1_Rule1","marked_for_delete":false,"overridden":false,"sequence_number":10,"sources_exclude
d":false,"destinations_excluded":false,"source_groups":["ANY"],"destination_groups":
["ANY"],"services":["ANY"],"profiles":["ANY"],"logged":false,"scope":["/infra/tier-1s/
myT1"],"disabled":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","_protection":"UNKNOWN"
},"resource_type":"ChildRule","marked_for_delete":false,"mark_for_override":false,"_protection
":

<182>1 2020-08-11T22:31:26.890Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="c2790fbd-db29-46d3-9a0e-1003455ee9ea" splitId="0s7tdCjN"
splitIndex="2 of 2" subcomp="policy" update="true"]
"UNKNOWN"}], "marked_for_delete":false, "overridden":false, "sequence_number":10, "category":"Loca
lGatewayRules", "stateful":true, "locked":false, "_protection":"UNKNOWN"}, {"resource_type":"ChildG
atewayPolicy", "marked_for_delete":false, "mark_for_override":false, "_protection":"UNKNOWN"}], "t
arget_type":"Domain", "resource_type":"ChildResourceReference", "id":"default", "marked_for_delet
e":false, "mark_for_override":false, "_protection":"UNKNOWN"}], "marked_for_delete":false, "overri
dden":false, "_protection":"UNKNOWN", "_revision":-1}]
```

### Updating a rule (from myT1\_Rule1 to myT1\_Rule1\_Updated) in a section (T1-Policies):

```
<182>1 2020-08-11T22:36:19.410Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="a17fcbdc-1aed-4526-93e9-40a3730eeb7f" splitId="BiHDjsY8"
splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin",
ModuleName="PolicyEdgeFirewall", Operation="PatchGatewayPolicyForDomain", Operation
status="success", New value=["default" "T1-Policies"
{"resource_type":"GatewayPolicy","id":"T1-Policies","display_name":"T1-Policies","path":"/
infra/domains/default/gateway-policies/T1-Policies","unique_id":"a73c1345-6b4e-43e0-
b4ee-9a91c7ba9df6","children":[{"Rule":
{"action":"ALLOW","resource_type":"Rule","id":"myT1_Rule1","display_name":"myT1_Rule1_Updated"
,"path":"/infra/domains/default/gateway-policies/T1-Policies/rules/
myT1_Rule1","unique_id":"2028","marked_for_delete":false,"overridden":false,"rule_id":2028,"se
quence_number":10,"sources_excluded":false,"destinations_excluded":false,"source_groups":
["ANY"],"destination_groups":["ANY"],"services":["ANY"],"profiles":
["ANY"],"logged":false,"scope":["/infra/tier-1s/
myT1"],"disabled":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","is_default":false,"_pr
otection":"UNKNOWN", "_revision":0}, {"resource_type":

<182>1 2020-08-11T22:36:19.410Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-
manager" level="INFO" reqId="a17fcbdc-1aed-4526-93e9-40a3730eeb7f" splitId="BiHDjsY8"
splitIndex="2 of 2" subcomp="policy" update="true"]
"ChildRule", "marked_for_delete":false, "mark_for_override":false, "_protection":"UNKNOWN"}], "mar
ked_for_delete":false, "overridden":false, "sequence_number":10, "internal_sequence_number":13000
```

```
010,"category":"LocalGatewayRules","stateful":true,"tcp_strict":true,"locked":false,"lock_modified_time":0,"is_default":false,"_protection":"UNKNOWN","_revision":0}]
```

```
<182>1 2020-08-11T22:36:19.430Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-manager" level="INFO" reqId="a17fcbdc-1aed-4526-93e9-40a3730eeb7f" splitId="HqtdDMqz" splitIndex="1 of 2" subcomp="policy" update="true" username="admin"] UserName="admin", ModuleName="PolicyEdgeFirewall", Operation="PatchGatewayRule", Operation status="success", Old value=[{"sequence_number":10,"source_groups":["ANY"],"destination_groups":["ANY"],"services":["ANY"],"action":"ALLOW","logged":false,"scope":["/infra/tier-1s/myT1"],"disabled":false,"direction":"IN_OUT","resource_type":"CommunicationEntry","id":"myT1_Rule1","display_name":"myT1_Rule1","path":"/infra/domains/default/gateway-policies/T1-Policies/rules/myT1_Rule1","relative_path":"myT1_Rule1","parent_path":"/infra/domains/default/gateway-policies/T1-Policies","unique_id":"2028","marked_for_delete":false,"overridden":false,"_create_user":"admin","_create_time":1597185086809,"_last_modified_user":"admin","_last_modified_time":1597185086841,"_system_owned":false,"_protection":"NOT_PROTECTED","_revision":0}], New value=["default" "T1-Policies" "myT1_Rule1" {"action":"ALLOW","resource_type":"Rule","id":"myT1_Rule1","display_name":"myT1_Rule1_Updated","path":
```

```
<182>1 2020-08-11T22:36:19.430Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-manager" level="INFO" reqId="a17fcbdc-1aed-4526-93e9-40a3730eeb7f" splitId="HqtdDMqz" splitIndex="2 of 2" subcomp="policy" update="true" username="admin"] "/infra/domains/default/gateway-policies/T1-Policies/rules/myT1_Rule1","unique_id":"2028","marked_for_delete":false,"overridden":false,"rule_id":2028,"sequence_number":10,"sources_excluded":false,"destinations_excluded":false,"source_groups":["ANY"],"destination_groups":["ANY"],"services":["ANY"],"profiles":["ANY"],"logged":false,"scope":["/infra/tier-1s/myT1"],"disabled":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","is_default":false,"_protection":"UNKNOWN","_revision":0}]
```

```
<182>1 2020-08-11T22:36:19.443Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-manager" level="INFO" reqId="a17fcbdc-1aed-4526-93e9-40a3730eeb7f" splitId="fMYsYjV5" splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin", ModuleName="Policy", Operation="PatchInfra", Operation status="success", New value=[{"enforce_revision_check":true} {"resource_type":"Infra","children":[{"children":[{"GatewayPolicy":{"resource_type":"GatewayPolicy","id":"T1-Policies","display_name":"T1-Policies","path":"/infra/domains/default/gateway-policies/T1-Policies","unique_id":"a73c1345-6b4e-43e0-b4ee-9a91c7ba9df6","children":[{"Rule":{"action":"ALLOW","resource_type":"Rule","id":"myT1_Rule1","display_name":"myT1_Rule1_Updated","path":"/infra/domains/default/gateway-policies/T1-Policies/rules/myT1_Rule1","unique_id":"2028","marked_for_delete":false,"overridden":false,"rule_id":2028,"sequence_number":10,"sources_excluded":false,"destinations_excluded":false,"source_groups":["ANY"],"destination_groups":["ANY"],"services":["ANY"],"profiles":["ANY"],"logged":false,"scope":["/infra/tier-1s/myT1"],"disabled":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","is_default":false,"_protection":"UNKNOWN","_revision":
```

```
<182>1 2020-08-11T22:36:19.443Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-manager" level="INFO" reqId="a17fcbdc-1aed-4526-93e9-40a3730eeb7f" splitId="fMYsYjV5" splitIndex="2 of 2" subcomp="policy" update="true"] 0},"resource_type":"ChildRule","marked_for_delete":false,"mark_for_override":false,"_protection":"UNKNOWN"}],{"marked_for_delete":false,"overridden":false,"sequence_number":10,"internal_sequence_number":13000010,"category":"LocalGatewayRules","stateful":true,"tcp_strict":true,"locked":false,"lock_modified_time":0,"is_default":false,"_protection":"UNKNOWN","_revision":0},"res
```

```
source_type":"ChildGatewayPolicy","marked_for_delete":false,"mark_for_override":false,"_protection":"UNKNOWN"}], "target_type":"Domain", "resource_type":"ChildResourceReference", "id":"default", "marked_for_delete":false, "mark_for_override":false, "_protection":"UNKNOWN"}], "marked_for_delete":false, "overridden":false, "_protection":"UNKNOWN", "_revision":-1}]
```

Delete a rule (myT1\_Rule2) from a section (T1-Policies):

```
<182>1 2020-08-11T22:38:03.262Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-manager" entId="myT1_Rule2" level="INFO" reqId="ccb8d0bb-0fe2-415a-9979-ala3a80a7038" subcomp="policy" update="true" username="admin"] UserName="admin", ModuleName="PolicyEdgeFirewall", Operation="DeleteGatewayRule", Operation status="success", Old value=[{"sequence_number":20,"source_groups":["ANY"],"destination_groups":["ANY"],"services":["ANY"],"action":"ALLOW","logged":false,"scope":["/infra/tier-1s/myT1"],"disabled":false,"direction":"IN_OUT","resource_type":"CommunicationEntry","id":"myT1_Rule2","display_name":"myT1_Rule2","path":"/infra/domains/default/gateway-policies/T1-Policies/rules/myT1_Rule2","relative_path":"myT1_Rule2","parent_path":"/infra/domains/default/gateway-policies/T1-Policies","unique_id":"2029","marked_for_delete":false,"overridden":false,"_create_user":"admin","_create_time":1597185467310,"_last_modified_user":"admin","_last_modified_time":1597185467314,"_system_owned":false,"_protection":"NOT_PROTECTED","_revision":0}], New value=["default" "T1-Policies" "myT1_Rule2"]

<182>1 2020-08-11T22:38:03.280Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-manager" level="INFO" reqId="ccb8d0bb-0fe2-415a-9979-ala3a80a7038" splitId="G1UhKvqu" splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin", ModuleName="PolicyEdgeFirewall", Operation="PatchGatewayPolicyForDomain", Operation status="success", New value=["default" "T1-Policies" {"resource_type":"GatewayPolicy","id":"T1-Policies","display_name":"T1-Policies","path":"/infra/domains/default/gateway-policies/T1-Policies","unique_id":"a73c1345-6b4e-43e0-b4ee-9a91c7ba9df6","children":[{"Rule":{"resource_type":"Rule","id":"myT1_Rule2","path":"/infra/domains/default/gateway-policies/T1-Policies/rules/myT1_Rule2","marked_for_delete":true,"overridden":false,"sources_excluded":false,"destinations_excluded":false,"logged":false,"disabled":false,"direction":"IN_OUT","_protection":"UNKNOWN"},"resource_type":"ChildRule","marked_for_delete":true,"mark_for_override":false,"_protection":"UNKNOWN"}], "marked_for_delete":false, "overridden":false, "sequence_number":10, "internal_sequence_number":13000010, "category":"LocalGatewayRules", "stateful":true, "tcp_strict":true, "locked":false, "lock_modified_time":0, "is_default":

<182>1 2020-08-11T22:38:03.280Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-manager" level="INFO" reqId="ccb8d0bb-0fe2-415a-9979-ala3a80a7038" splitId="G1UhKvqu" splitIndex="2 of 2" subcomp="policy" update="true"] false, "_protection":"UNKNOWN", "_revision":0}]

<182>1 2020-08-11T22:38:03.295Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-manager" level="INFO" reqId="ccb8d0bb-0fe2-415a-9979-ala3a80a7038" splitId="xnO9T8NE" splitIndex="1 of 2" subcomp="policy" update="true"] UserName="admin", ModuleName="Policy", Operation="PatchInfra", Operation status="success", New value=[{"enforce_revision_check":true} {"resource_type":"Infra","children":[{"children":[{"GatewayPolicy":{"resource_type":"GatewayPolicy","id":"T1-Policies","display_name":"T1-Policies","path":"/infra/domains/default/gateway-policies/T1-Policies","unique_id":"a73c1345-6b4e-43e0-b4ee-9a91c7ba9df6","children":[{"Rule":{"resource_type":"Rule","id":"myT1_Rule2","path":"/infra/domains/default/gateway-policies/T1-Policies/rules/myT1_Rule2","marked_for_delete":true,"overridden":false,"sources_excluded":false,"destinations_excluded":false,"logged":false,"disabled":false,"direction":"IN_OUT","_protection":"UNKNOWN"}]
```

```
, "resource_type": "ChildRule", "marked_for_delete": true, "mark_for_override": false, "_protection": "UNKNOWN"}], "marked_for_delete": false, "overridden": false, "sequence_number": 10, "internal_sequence_number": 13000010, "category": "LocalGatewayRules", "stateful": true, "tcp_strict": true, "locked": false, "lock_modified_time": 0, "is_default": false, "_protection": "UNKNOWN", "_revision": 0}, {"resource_type": "ChildGatewayPolicy", "marked_for_delete": false, "mark_for_override": false, "_protection": "UNKNOWN"}], "target_type": "Domain", "resource_type": "ChildResourceReference", "id": "default", "marked_for_delete": false, "mark_for_override": false, "_protection": "UNKNOWN"}], "marked_for_delete": false, "overridden": false, "_protection": "UNKNOWN", "_revision": -1}]
```

Deleting a section (T1-Policies) that contains a rule (myT1\_Rule1\_Updated):

```
<182>1 2020-08-11T22:41:30.726Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-manager" entId="T1-Policies" level="INFO" reqId="d751343c-32ab-46ee-b176-752b8aelec0d" splitId="WZc3oxDG" splitIndex="1 of 2" subcomp="policy" update="true" username="admin"]
UserName="admin", ModuleName="PolicyEdgeFirewall", Operation="DeleteGatewayPolicy", Operation status="success", Old value=[{"precedence":10,"category":"LocalGatewayRules","resource_type":"CommunicationMap","id":"T1-Policies","display_name":"T1-Policies","path":"/infra/domains/default/gateway-policies/T1-Policies","relative_path":"T1-Policies","parent_path":"/infra/domains/default","unique_id":"a73c1345-6b4e-43e0-b4ee-9a91c7ba9df6","marked_for_delete":false,"overridden":false,"_create_user":"admin","_create_time":1597185086789,"_last_modified_user":"admin","_last_modified_time":1597185086790,"_system_owned":false,"_protection":"NOT_PROTECTED","_revision":0} {"sequence_number":10,"source_groups":["ANY"],"destination_groups":["ANY"],"services":["ANY"],"action":"ALLOW","logged":false,"scope":["/infra/tier-1s/myT1"],"disabled":false,"direction":"IN_OUT","resource_type":"CommunicationEntry","id":"myT1_Rule1","display_name":"myT1_Rule1_Updated","path":

<182>1 2020-08-11T22:41:30.726Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-manager" entId="T1-Policies" level="INFO" reqId="d751343c-32ab-46ee-b176-752b8aelec0d" splitId="WZc3oxDG" splitIndex="2 of 2" subcomp="policy" update="true" username="admin"] "/infra/domains/default/gateway-policies/T1-Policies/rules/myT1_Rule1","relative_path":"myT1_Rule1","parent_path":"/infra/domains/default/gateway-policies/T1-Policies","unique_id":"2028","marked_for_delete":false,"overridden":false,"_create_user":"admin","_create_time":1597185086809,"_last_modified_user":"admin","_last_modified_time":1597185379419,"_system_owned":false,"_protection":"NOT_PROTECTED","_revision":1}], New value=["default" "T1-Policies"]

<182>1 2020-08-11T22:41:30.733Z manager1 NSX 22164 - [nsx@6876 audit="true" comp="nsx-manager" level="INFO" reqId="d751343c-32ab-46ee-b176-752b8aelec0d" subcomp="policy" update="true"] UserName="admin", ModuleName="Policy", Operation="PatchInfra", Operation status="success", New value=[{"enforce_revision_check":true} {"resource_type":"Infra","children":[{"children":[{"GatewayPolicy":{"resource_type":"GatewayPolicy","id":"T1-Policies","path":"/infra/domains/default/gateway-policies/T1-Policies","marked_for_delete":true,"overridden":false,"locked":false,"_protection":"UNKNOWN"}},
```

```
"resource_type":"ChildGatewayPolicy","marked_for_delete":true,"mark_for_override":false,"_protection":"UNKNOWN"}]], "target_type":"Domain", "resource_type":"ChildResourceReference", "id":"default", "marked_for_delete":false, "mark_for_override":false, "_protection":"UNKNOWN"}]], "marked_for_delete":false, "overridden":false, "_protection":"UNKNOWN", "_revision":-1}]
```

## Distributed firewall changes in Manager mode

Adding a firewall section (FirewallSection-2):

```
<182>1 2020-08-12T00:25:53.300Z manager1 NSX 1503 - [nsx@6876 audit="true" comp="nsx-manager" level="INFO" reqId="244e8a97-93d4-4047-b817-81b59b94ce13" subcomp="manager" username="admin"] UserName="admin", ModuleName="NSX-Firewall", Operation="CREATE", Operation status="success", New value=[FirewallSectionLock [Id=0ffb0688-9f4e-4096-a19f-2d98ce8cfbeb, sectionId=f5226cab-525b-4e33-a26d-e5053fbba0a1, sectionRevision=0, locked=false, comments=Default section unlock comment, created_by=admin, create_time=1597191953299, last_modified_by=admin, last_modified_time=1597191953299]]

<182>1 2020-08-12T00:25:53.313Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-manager" entId="f5226cab-525b-4e33-a26d-e5053fbba0a1" level="INFO" reqId="244e8a97-93d4-4047-b817-81b59b94ce13" subcomp="manager" update="true" username="admin"] UserName="admin", ModuleName="Firewall", Operation="AddSection", Operation status="success", New value=[{"operation":"insert_before","id":"ffffffff-8a04-4924-a5b4-54d30e81befe"} {"locked":false,"autoplumbed":false,"tcp_strict":false,"display_name":"FirewallSection-2","section_type":"LAYER3","stateful":true,"_protection":"UNKNOWN"}]
```

Adding a rule (mp\_Rule1) to a section (FirewallSection-2):

```
<182>1 2020-08-12T00:27:21.252Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-manager" entId="f5226cab-525b-4e33-a26d-e5053fbba0a1" level="INFO" reqId="3562bb3e-bf18-4aa0-aldd-abde13e8559c" splitId="ScK9FB8V" splitIndex="1 of 2" subcomp="manager" update="true" username="admin"] UserName="admin", ModuleName="Firewall", Operation="UpdateSectionWithRules", Operation status="success", Old value=[{"locked":false,"comments":"Default section unlock comment","lock_modified_by":"admin","lock_modified_time":1597191953299,"autoplumbed":false,"enforced_on":"VIF","tcp_strict":false,"category":"Default","resource_type":"FirewallSection","id":"f5226cab-525b-4e33-a26d-e5053fbba0a1","display_name":"FirewallSection-2","section_type":"LAYER3","stateful":true,"rule_count":0,"is_default":false,"_create_user":"admin","_create_time":1597191953297,"_last_modified_user":"admin","_last_modified_time":1597191953297,"_system_owned":false,"_protection":"NOT_PROTECTED","_revision":0}], New value=[{"f5226cab-525b-4e33-a26d-e5053fbba0a1" {"rules": [{"display_name":"mp_Rule1","sources_excluded":false,"destinations_excluded":false,"action":"ALLOW","disabled":false,"logged":false,"direction":"IN_OUT","ip_protocol":"IPV4_IPV6","is_default":false}], "resource_type":"FirewallSection","id":

<182>1 2020-08-12T00:27:21.252Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-manager" entId="f5226cab-525b-4e33-a26d-e5053fbba0a1" level="INFO" reqId="3562bb3e-bf18-4aa0-aldd-abde13e8559c" splitId="ScK9FB8V" splitIndex="2 of 2" subcomp="manager" update="true" username="admin"] "f5226cab-525b-4e33-a26d-e5053fbba0a1","display_name":"FirewallSection-2","section_type":"LAYER3","stateful":true,"rule_count":0,"is_default":false,"locked":false,"comments":"Default section unlock comment","lock_modified_by":"admin","lock_modified_time":1597191953299,"autoplumbed":false,"enforced_on":"VIF","tcp_strict":false,"category":"Default","_protection":"UNKNOWN","_revision":0}]
```

## Updating a rule (from mp\_Rule1 to mp\_Rule1\_updated) in a section (FirewallSection-2):

```
<182>1 2020-08-12T00:28:54.226Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="f5226cab-525b-4e33-a26d-e5053fbba0a1" level="INFO"
reqId="37954994-8d59-448e-923d-940813087640" splitId="KcUAlRY1" splitIndex="1 of 2"
subcomp="manager" update="true" username="admin"] UserName="admin", ModuleName="Firewall",
Operation="UpdateSectionWithRules", Operation status="success", Old
value=[{"section_id":"f5226cab-525b-4e33-a26d-
e5053fbba0a1","resource_type":"FirewallRule","id":"536870917","display_name":"mp_Rule1","sourc
es_excluded":false,"destinations_excluded":false,"action":"ALLOW","disabled":false,"logged":fa
lse,"direction":"IN_OUT","ip_protocol":"IPv4_IPV6","is_default":false}
{"locked":false,"comments":"Default section unlock
comment","lock_modified_by":"admin","lock_modified_time":1597191953299,"autoplumbed":false,"en
forced_on":"VIF","tcp_strict":false,"category":"Default","resource_type":"FirewallSection","id
":"f5226cab-525b-4e33-a26d-
e5053fbba0a1","display_name":"FirewallSection-2","section_type":"LAYER3","stateful":true,"rule
_count":1,"is_default":false,"_create_user":"admin","_create_time":1597191953297,"_last_modifi
ed_user":"admin","_last_modified_time":1597192041235,"_system_owned":false,"_protection":"NOT_
PROTECTED","_revision":
<182>1 2020-08-12T00:28:54.226Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="f5226cab-525b-4e33-a26d-e5053fbba0a1" level="INFO"
reqId="37954994-8d59-448e-923d-940813087640" splitId="KcUAlRY1" splitIndex="2 of 2"
subcomp="manager" update="true" username="admin"] 1}], New value=[{"f5226cab-525b-4e33-a26d-
e5053fbba0a1" {"rules":[{"section_id":"f5226cab-525b-4e33-a26d-
e5053fbba0a1","resource_type":"FirewallRule","id":"536870917","display_name":"mp_Rule1_updated
","sources_excluded":false,"destinations_excluded":false,"action":"ALLOW","disabled":false,"lo
gged":false,"direction":"IN_OUT","ip_protocol":"IPv4_IPV6","is_default":false,"_revision":1}],
"resource_type":"FirewallSectionRuleList","id":"f5226cab-525b-4e33-a26d-
e5053fbba0a1","display_name":"FirewallSection-2","section_type":"LAYER3","stateful":true,"rule
_count":1,"is_default":false,"locked":false,"comments":"Default section unlock
comment","lock_modified_by":"admin","lock_modified_time":1597191953299,"autoplumbed":false,"en
forced_on":"VIF","tcp_strict":false,"category":"Default","_protection":"UNKNOWN","_revision":1
}]}
```

## Deleting a rule (mp\_Rule2) from a section (FirewallSection-2):

```
<182>1 2020-08-12T00:33:58.355Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="f5226cab-525b-4e33-a26d-e5053fbba0a1" level="INFO"
reqId="2db867e0-0407-44a2-8a6c-96895ff14a2f" splitId="m9SdpPw2" splitIndex="1 of 3"
subcomp="manager" update="true" username="admin"] UserName="admin", ModuleName="Firewall",
Operation="UpdateSectionWithRules", Operation status="success", Old
value=[{"section_id":"f5226cab-525b-4e33-a26d-
e5053fbba0a1","resource_type":"FirewallRule","id":"536870918","display_name":"mp_Rule2","sourc
es_excluded":false,"destinations_excluded":false,"action":"ALLOW","disabled":false,"logged":fa
lse,"direction":"IN_OUT","ip_protocol":"IPv4_IPV6","is_default":false}
{"section_id":"f5226cab-525b-4e33-a26d-
e5053fbba0a1","resource_type":"FirewallRule","id":"536870917","display_name":"mp_Rule1_updated
","sources_excluded":false,"destinations_excluded":false,"action":"ALLOW","disabled":false,"lo
gged":false,"direction":"IN_OUT","ip_protocol":"IPv4_IPV6","is_default":false}
{"locked":false,"comments":"Default section unlock
comment","lock_modified_by":"admin","lock_modified_time":1597191953299,"autoplumbed":false,"en
forced_on":"VIF","tcp_strict":false,"category":"Default","resource_type":"FirewallSection","id
":
<182>1 2020-08-12T00:33:58.355Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
```

```

manager" entId="f5226cab-525b-4e33-a26d-e5053fbba0a1" level="INFO"
reqId="2db867e0-0407-44a2-8a6c-96895ff14a2f" splitId="m9SdpPw2" splitIndex="2 of 3"
subcomp="manager" update="true" username="admin"] "f5226cab-525b-4e33-a26d-
e5053fbba0a1", "display_name": "FirewallSection-2", "section_type": "LAYER3", "stateful": true, "rule
_count": 2, "is_default": false, "create_user": "admin", "create_time": 1597191953297, "last_modifi
ed_user": "admin", "last_modified_time": 1597192378372, "system_owned": false, "protection": "NOT_
PROTECTED", "revision": 3}], New value=[{"f5226cab-525b-4e33-a26d-e5053fbba0a1" {"rules":
[{"section_id": "f5226cab-525b-4e33-a26d-
e5053fbba0a1", "resource_type": "FirewallRule", "id": "536870917", "display_name": "mp_Rule1_updated
", "sources_excluded": false, "destinations_excluded": false, "action": "ALLOW", "disabled": false, "lo
gged": false, "direction": "IN_OUT", "ip_protocol": "IPv4_IPV6", "is_default": false, "revision": 3}],
"resource_type": "FirewallSectionRuleList", "id": "f5226cab-525b-4e33-a26d-
e5053fbba0a1", "display_name": "FirewallSection-2", "section_type": "LAYER3", "stateful": true, "rule
_count": 2, "is_default": false, "locked": false, "comments": "Default section unlock
comment", "lock_modified_by":
<182>1 2020-08-12T00:33:58.355Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="f5226cab-525b-4e33-a26d-e5053fbba0a1" level="INFO"
reqId="2db867e0-0407-44a2-8a6c-96895ff14a2f" splitId="m9SdpPw2" splitIndex="3 of 3"
subcomp="manager" update="true" username="admin"]
"admin", "lock_modified_time": 1597191953299, "autoplumbed": false, "enforced_on": "VIF", "tcp_strict
": false, "category": "Default", "protection": "UNKNOWN", "revision": 3}]

```

Deleting a section (FirewallSection-2) that contains a rule (mp\_Rule1):

```

<182>1 2020-08-12T00:35:01.304Z manager1 NSX 1503 - [nsx@6876 audit="true" comp="nsx-manager"
level="INFO" reqId="f23e091f-aa6e-47a6-945a-98291cc3f0ba" subcomp="manager" username="admin"]
UserName="admin", ModuleName="NSX-Firewall", Operation="DELETE", Operation status="success",
Old value=[FirewallSectionLock [Id=0ffb0688-9f4e-4096-a19f-2d98ce8cfbeb,
sectionId=f5226cab-525b-4e33-a26d-e5053fbba0a1, sectionRevision=0, locked=false,
comments=Default section unlock comment, created_by=admin, create_time=1597191953299,
last_modified_by=admin, last_modified_time=1597191953299]]
<182>1 2020-08-12T00:35:01.324Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="f5226cab-525b-4e33-a26d-e5053fbba0a1" level="INFO" reqId="f23e091f-
aa6e-47a6-945a-98291cc3f0ba" subcomp="manager" update="true" username="admin"]
UserName="admin", ModuleName="Firewall", Operation="DeleteSection", Operation
status="success", Old value=[null{"section_id": "f5226cab-525b-4e33-a26d-
e5053fbba0a1", "resource_type": "FirewallRule", "id": "536870917", "display_name": "mp_Rule1_updated
", "sources_excluded": false, "destinations_excluded": false, "action": "ALLOW", "disabled": false, "lo
gged": false, "direction": "IN_OUT", "ip_protocol": "IPv4_IPV6", "is_default": false}
{"locked": false, "autoplumbed": false, "enforced_on": "VIF", "tcp_strict": false, "category": "Default
", "resource_type": "FirewallSection", "id": "f5226cab-525b-4e33-a26d-
e5053fbba0a1", "display_name": "FirewallSection-2", "section_type": "LAYER3", "stateful": true, "rule
_count": 1, "is_default": false, "create_user": "admin", "create_time": 1597191953297, "last_modifi
ed_user": "admin", "last_modified_time": 1597192438335, "system_owned": false, "protection": "NOT_
PROTECTED", "revision": 4}], New value=[{"f5226cab-525b-4e33-a26d-e5053fbba0a1"
{"cascade": true}]

```

## Edge firewall changes in Manager mode

Note that log messages for a tier-0 logical router and a tier-1 logical router are similar.



Adding a firewall section (FirewallSection-1) for a tier-1 logical router (myT1\_mp):

```
<182>1 2020-08-12T00:09:55.661Z manager1 NSX 1503 - [nsx@6876 audit="true" comp="nsx-manager"
level="INFO" reqId="14af9252-ddc3-4949-8e01-b2c5676ac258" subcomp="manager" username="admin"]
UserName="admin", ModuleName="NSX-Firewall", Operation="CREATE", Operation status="success",
New value=[FirewallSectionLock [I
d=15b61818-2a65-48cf-a98e-7c2f3fccc845, sectionId=9808d1ec-de08-48b3-8173-12f26fb0ae9c,
sectionRevision=0, locked=false, comments=Default section unlock comment, created_by=admin,
create_time=1597190995659, last_modified_by=admin, last_modified_time=1597190995659]]

<182>1 2020-08-12T00:09:55.687Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="9808d1ec-de08-48b3-8173-12f26fb0ae9c" level="INFO" reqId="14af9252-
ddc3-4949-8e01-b2c5676ac258" subcomp="manager" update="true" username="admin"]
UserName="admin", ModuleName="Firewall", Operation="AddSection", Operation status="success",
New value=[{"operation":"insert_before","id":"095b443a-115d-4bf7-b4f7-192305321e95"}
{"locked":false,"autoplumbed":false,"tcp_strict":false,"display_name":"FirewallSection-1","app
lied_tos":
[{"target_id":"6562738e-73b9-4f21-9461-460ead581daf","target_display_name":"myT1_mp","target_t
ype":"LogicalRouter","is_valid":true}],{"section_type":"LAYER3","stateful":true,"is_default":fa
lse,"_system_owned":false,"_protection":"UNKNOWN","_revision":0}]
```

Adding a rule (myT1\_mp\_Rule1) to a section (FirewallSection-1):

```
<182>1 2020-08-12T00:13:44.092Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="9808d1ec-de08-48b3-8173-12f26fb0ae9c" level="INFO"
reqId="d4e7bdef-0cc6-45e9-8884-061b0f688fec" splitId="snErcGKF" splitIndex="1 of 2"
subcomp="manager" update="true" username="admin"] UserName="admin", ModuleName="Firewall",
Operation="UpdateSectionWithRules", Operation status="success", Old
value=[{"locked":false,"comments":"Default section unlock
comment","lock_modified_by":"admin","lock_modified_time":1597190995659,"autoplumbed":false,"en
forced_on":"LOGICALROUTER","tcp_strict":false,"category":"Default","resource_type":"FirewallSe
ction","id":"9808d1ec-
de08-48b3-8173-12f26fb0ae9c","display_name":"FirewallSection-1","applied_tos":
[{"target_id":"6562738e-73b9-4f21-9461-460ead581daf","target_display_name":"myT1_mp","target_t
ype":"LogicalRouter","is_valid":true}],{"section_type":"LAYER3","stateful":true,"rule_count":0,
"is_default":false,"_create_user":"admin","_create_time":1597190995657,"_last_modified_user":
"admin","_last_modified_time":1597190995657,"_system_owned":false,"_protection":"NOT_PROTECTED"
,"_revision":0}], New value=[{"9808d1ec-de08-48b3-8173-12f26fb0ae9c" {"rules":
[{"display_name":"myT1_mp_Rule1","sources_excluded":false,"destinations_excluded":
```

Updating a rule (from myT1\_mp\_Rule1 to myT1\_mp\_Rule1\_updated) in a section (FirewallSection-1):

```
<182>1 2020-08-12T00:15:31.078Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="9808d1ec-de08-48b3-8173-12f26fb0ae9c" level="INFO" reqId="eb880eee-5798-42fc-
a8aa-58b70e4aa152" splitId="WviId4ja" splitIndex="1 of 3" subcomp="manager" update="true"
username="admin"] UserName="admin", ModuleName="Firewall",
Operation="UpdateSectionWithRules", Operation status="success", Old
value=[{"locked":false,"comments":"Default section unlock
comment","lock_modified_by":"admin","lock_modified_time":1597190995659,"autoplumbed":false,"en
forced_on":"LOGICALROUTER","tcp_strict":false,"category":"Default","resource_type":"FirewallSe
ction","id":"9808d1ec-
de08-48b3-8173-12f26fb0ae9c","display_name":"FirewallSection-1","applied_tos":
[{"target_id":"6562738e-73b9-4f21-9461-460ead581daf","target_display_name":"myT1_mp","target_t
```

```

ype":"LogicalRouter","is_valid":true}], "section_type":"LAYER3", "stateful":true, "rule_count":1,
"is_default":false, "create_user":"admin", "create_time":1597190995657, "last_modified_user":"
admin", "last_modified_time":1597191224058, "system_owned":false, "protection":"NOT_PROTECTED"
, "revision":1} {"section_id":"9808d1ec-
de08-48b3-8173-12f26fb0ae9c", "resource_type":"FirewallRule", "id":"536870914", "display_name":"m
yTl_mp_Rule1", "sources_excluded":
<182>1 2020-08-12T00:15:31.078Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="9808d1ec-de08-48b3-8173-12f26fb0ae9c" level="INFO" reqId="eb880eee-5798-42fc-
a8aa-58b70e4aa152" splitId="WviLd4ja" splitIndex="2 of 3" subcomp="manager" update="true"
username="admin"]
false, "destinations_excluded":false, "action":"ALLOW", "disabled":false, "logged":false, "directio
n":"IN_OUT", "ip_protocol":"IPV4_IPV6", "is_default":false}], New value=["9808d1ec-
de08-48b3-8173-12f26fb0ae9c" {"rules":[{"section_id":"9808d1ec-
de08-48b3-8173-12f26fb0ae9c", "resource_type":"FirewallRule", "id":"536870914", "display_name":"m
yTl_mp_Rule1_updated", "sources_excluded":false, "destinations_excluded":false, "action":"ALLOW",
"disabled":false, "logged":false, "direction":"IN_OUT", "ip_protocol":"IPV4_IPV6", "revision":1}]
, "resource_type":"FirewallSectionRuleList", "id":"9808d1ec-
de08-48b3-8173-12f26fb0ae9c", "display_name":"FirewallSection-1", "applied_tos":
[{"target_id":"6562738e-73b9-4f21-9461-460ead581daf", "target_display_name":"myTl_mp", "target_t
ype":"LogicalRouter", "is_valid":true}], "section_type":"LAYER3", "stateful":true, "is_default":fa
lse, "locked":false, "comments":"Default section unlock
comment", "lock_modified_by":"admin", "lock_modified_time":1597190995659, "autoplumbed":
<182>1 2020-08-12T00:15:31.078Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="9808d1ec-de08-48b3-8173-12f26fb0ae9c" level="INFO" reqId="eb880eee-5798-42fc-
a8aa-58b70e4aa152" splitId="WviLd4ja" splitIndex="3 of 3" subcomp="manager" update="true"
username="admin"]
false, "enforced_on":"LOGICALROUTER", "tcp_strict":false, "category":"Default", "system_owned":fa
lse, "protection":"UNKNOWN", "revision":1}]

```

### Deleting a rule (myTl\_mp\_Rule2) from a section (FirewallSection-1):

```

<182>1 2020-08-12T00:18:05.341Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="9808d1ec-de08-48b3-8173-12f26fb0ae9c" level="INFO" reqId="bc95016c-5ec2-4b25-
ab17-0b10b6c5a4f0" splitId="damZHQkr" splitIndex="1 of 3" subcomp="manager" update="true"
username="admin"] UserName="admin", ModuleName="Firewall",
Operation="UpdateSectionWithRules", Operation status="success", Old
value=[{"locked":false, "comments":"Default section unlock
comment", "lock_modified_by":"admin", "lock_modified_time":1597190995659, "autoplumbed":false, "en
forced_on":"LOGICALROUTER", "tcp_strict":false, "category":"Default", "resource_type":"FirewallSe
ction", "id":"9808d1ec-
de08-48b3-8173-12f26fb0ae9c", "display_name":"FirewallSection-1", "applied_tos":
[{"target_id":"6562738e-73b9-4f21-9461-460ead581daf", "target_display_name":"myTl_mp", "target_t
ype":"LogicalRouter", "is_valid":true}], "section_type":"LAYER3", "stateful":true, "rule_count":2,
"is_default":false, "create_user":"admin", "create_time":1597190995657, "last_modified_user":"
admin", "last_modified_time":1597191475552, "system_owned":false, "protection":"NOT_PROTECTED"
, "revision":3} {"section_id":"9808d1ec-
de08-48b3-8173-12f26fb0ae9c", "resource_type":"FirewallRule", "id":"536870914", "display_name":
<182>1 2020-08-12T00:18:05.341Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-
manager" entId="9808d1ec-de08-48b3-8173-12f26fb0ae9c" level="INFO" reqId="bc95016c-5ec2-4b25-
ab17-0b10b6c5a4f0" splitId="damZHQkr" splitIndex="2 of 3" subcomp="manager" update="true"
username="admin"]
"myTl_mp_Rule1_updated", "sources_excluded":false, "destinations_excluded":false, "action":"ALLOW
", "disabled":false, "logged":false, "direction":"IN_OUT", "ip_protocol":"IPV4_IPV6", "is_default":
false} {"section_id":"9808d1ec-
de08-48b3-8173-12f26fb0ae9c", "resource_type":"FirewallRule", "id":"536870915", "display_name":"m

```

```

yTl_mp_Rule2", "sources_excluded": false, "destinations_excluded": false, "action": "ALLOW", "disabled": false, "logged": false, "direction": "IN_OUT", "ip_protocol": "IPV4_IPV6", "is_default": false}],
New value=["9808dlec-de08-48b3-8173-12f26fb0ae9c" {"rules":[{"section_id": "9808dlec-de08-48b3-8173-12f26fb0ae9c", "resource_type": "FirewallRule", "id": "536870914", "display_name": "myTl_mp_Rule1_updated", "sources_excluded": false, "destinations_excluded": false, "action": "ALLOW", "disabled": false, "logged": false, "direction": "IN_OUT", "ip_protocol": "IPV4_IPV6", "_revision": 3}], "resource_type": "FirewallSectionRuleList", "id": "9808dlec-de08-48b3-8173-12f26fb0ae9c", "display_name": "FirewallSection-1", "applied_tos":
<182>1 2020-08-12T00:18:05.341Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-manager" entId="9808dlec-de08-48b3-8173-12f26fb0ae9c" level="INFO" reqId="bc95016c-5ec2-4b25-ab17-0b10b6c5a4f0" splitId="damZHQkr" splitIndex="3 of 3" subcomp="manager" update="true" username="admin"] :
[{"target_id": "6562738e-73b9-4f21-9461-460ead581daf", "target_display_name": "myTl_mp", "target_type": "LogicalRouter", "is_valid": true}], "section_type": "LAYER3", "stateful": true, "is_default": false, "locked": false, "comments": "Default section unlock comment", "lock_modified_by": "admin", "lock_modified_time": 1597190995659, "autoplumbed": false, "enforced_on": "LOGICALROUTER", "tcp_strict": false, "category": "Default", "_system_owned": false, "_protection": "UNKNOWN", "_revision": 3}]

```

Deleting a section (FirewallSection-1) that contains a rule (myTl\_mp\_Rule2):

```

<182>1 2020-08-12T00:21:27.646Z manager1 NSX 1503 - [nsx@6876 audit="true" comp="nsx-manager" level="INFO" reqId="781f43d5-0b4c-494e-89a1-cbc2998fc232" subcomp="manager" username="admin"]
UserName="admin", ModuleName="NSX-Firewall", Operation="DELETE", Operation status="success",
Old value=[FirewallSectionLock [Id=15b61818-2a65-48cf-a98e-7c2f3fccc845, sectionId=9808dlec-de08-48b3-8173-12f26fb0ae9c, sectionRevision=0, locked=false, comments=Default section unlock comment, created_by=admin, create_time=1597190995659, last_modified_by=admin, last_modified_time=1597190995659]]
<182>1 2020-08-12T00:21:27.669Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-manager" entId="9808dlec-de08-48b3-8173-12f26fb0ae9c" level="INFO"
reqId="781f43d5-0b4c-494e-89a1-cbc2998fc232" splitId="u3AofFMr" splitIndex="1 of 2" subcomp="manager" update="true" username="admin"]
UserName="admin", ModuleName="Firewall", Operation="DeleteSection", Operation status="success", Old value=[{"locked": false, "autoplumbed": false, "enforced_on": "LOGICALROUTER", "tcp_strict": false, "category": "Default", "resource_type": "FirewallSection", "id": "9808dlec-de08-48b3-8173-12f26fb0ae9c", "display_name": "FirewallSection-1", "applied_tos": [{"target_id": "6562738e-73b9-4f21-9461-460ead581daf", "target_display_name": "myTl_mp", "target_type": "LogicalRouter", "is_valid": true}], "section_type": "LAYER3", "stateful": true, "rule_count": 1, "is_default": false, "_create_user": "admin", "_create_time": 1597190995657, "_last_modified_user": "admin", "_last_modified_time": 1597191671601, "_system_owned": false, "_protection": "NOT_PROTECTED", "_revision": 6} {"section_id": "9808dlec-de08-48b3-8173-12f26fb0ae9c", "resource_type": "FirewallRule", "id": "536870916", "display_name": "myTl_mp_Rule1", "sources_excluded": false, "destinations_excluded": false, "action": "ALLOW", "disabled": false, "logged": false, "direction":
<182>1 2020-08-12T00:21:27.669Z manager1 NSX 1503 FIREWALL [nsx@6876 audit="true" comp="nsx-manager" entId="9808dlec-de08-48b3-8173-12f26fb0ae9c" level="INFO"
reqId="781f43d5-0b4c-494e-89a1-cbc2998fc232" splitId="u3AofFMr" splitIndex="2 of 2" subcomp="manager" update="true" username="admin"]
"IN_OUT", "ip_protocol": "IPV4_IPV6", "is_default": false} null], New value=["9808dlec-de08-48b3-8173-12f26fb0ae9c" {"cascade": true}]

```

## Customer Experience Improvement Program

NSX-T Data Center participates in VMware's Customer Experience Improvement Program (CEIP).

Details regarding the data collected through CEIP and the purposes for which it is used by VMware are set forth at the Trust & Assurance Center at <https://www.vmware.com/solutions/trustvmware/ceip-products.html>.

To join or leave the CEIP for NSX-T Data Center, or edit program settings, see [Edit the Customer Experience Improvement Program Configuration](#).

## Edit the Customer Experience Improvement Program Configuration

When you install or upgrade NSX Manager, you can decide to join the CEIP and configure data collection settings.

You can also edit the existing CEIP configuration to join or leave the CEIP program, define the frequency and the days the information is collected, and proxy server configuration.

### Prerequisites

- Verify that the NSX Manager is connected and can synchronize with your hypervisor.
- Verify that NSX-T Data Center is connected to a public network for uploading data.

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **System > Customer Program**.
- 3 Click **Edit** in the Customer Experience Improvement Program section.
- 4 In the Edit Customer Experience Program dialog box, select the **Join the VMware Customer Experience Improvement Program** check box.
- 5 Toggle the **Schedule** switch to disable or enable the data collection.  
The schedule is enabled by default.
- 6 (Optional) Configure the data collection and upload recurrence settings.
- 7 Click **Save**.

## Find the SSH Fingerprint of a Remote Server

Some tasks that involve communication with a remote server require that you provide the SSH fingerprint for the remote server. The SSH fingerprint is derived from a host key on the remote server.

To connect using SSH, the NSX Manager and the remote server must have a host key type in common. NSX Manager supports the ECDSA (256 bit) key. The default location of this key is `/etc/ssh/ssh_host_ecdsa_key.pub`.

Having the fingerprint for a remote server helps you confirm you are connecting to the correct server, protecting you from man-in-the-middle attacks. You can ask the administrator of the remote server to provide the SSH fingerprint of the server. Or you can connect to the remote server to find the fingerprint. Connecting to the server over console is more secure than over the network.

#### Procedure

- 1 Log in to the remote server as root.

Logging in using a console is more secure than over the network.

- 2 Locate the ECDSA (256 bit) key. The default location of the key is `/etc/ssh/ssh_host_ecdsa_key.pub`.

```
$ ls -al /etc/ssh/*pub
-rw-r--r-- 1 root root 93 Apr 8 18:10 ssh_host_ecdsa_key.pub
-rw-r--r-- 1 root root 393 Apr 8 18:10 ssh_host_rsa_key.pub
```

- 3 Get the fingerprint of the key.

```
ssh-keygen -lf /etc/ssh/ssh_host_ecdsa_key.pub | awk '{print $2}'
```

## Configuring an External Load Balancer

You can configure an external load balancer to distribute traffic to the NSX Managers in a manager cluster.

An NSX Manager cluster does not require an external load balancer. The NSX Manager virtual IP (VIP) provides resiliency in the event of a Manager node failure but has the following limitations:

- VIP does not perform load balancing across the NSX Managers.
- VIP requires all the NSX Managers to be in the same subnet.
- VIP recovery takes about 1 - 3 minutes in the event of a Manager node failure.

An external load balancer can provide the following benefits:

- Load balance across the NSX Managers.
- The NSX Managers can be in different subnets.
- Fast recovery time in the event of a Manager node failure.

An external load balancer will not work with the NSX Manager VIP. Do not configure an NSX Manager VIP if you use an external load balancer.

## Authentication Methods When Accessing NSX Manager

The following authentication methods are supported by NSX Manager. For more information about the authentication methods, see the *NSX-T Data Center API Guide*.

- HTTP Basic Authentication

- Session-Based Authentication
- Authentication using an X.509 certificate and a Principal Identity
- Authentication in VMware Cloud on AWS

The session-based authentication method (used when you access NSX Manager from a browser) requires source-IP persistence (all requests from the client must go to the same NSX Manager). The other methods do not require source-IP persistence (requests from the client can go to different NSX Managers).


## Recommendations

- Create a single VIP on the load balancer with source-IP persistence configured to handle all the authentication methods.
- If you have applications or scripts that might generate a lot of requests to NSX Manager, create a second VIP without source-IP persistence for these applications or scripts. Use the first VIP for browser access to NSX Manager only.

The VIP must have the following configurations:

- Type: Layer4-TCP
- Port: 443
- Pool: NSX Manager Pool
- Persistence: Source-IP persistence for the first VIP. None for the second VIP (if present).

Example of an external load balancer configuration:



## NSX Manager's Certificate

The clients access NSX Manager using FQDN name (for example, nsx.mycompany.com). This FQDN is resolved to the load balancer's VIP. To avoid any certificate mismatch, each NSX Manager must have a certificate that is valid for the VIP's FQDN name. Therefore, you must configure each NSX Manager with a SAN certificate that is valid for its own name (for example, nsxmgr1.mycompany.com) and the VIP's FQDN.

## Monitoring the Health of NSX Managers

The load balancer can check that each NSX Manager is running with the following API:

```
GET /api/v1/reverse-proxy/node/health
```

The request headers are:

- Header1
  - Name: Authorization
  - Value: Basic <Base64 Value>

Note: <Base64 Value> is Username:Password encoded in Base64. You can use <https://www.base64encode.net> to do the encoding. For example, Header1 might be `Authorization: Basic YWRtaW46Vk13YXJlMSFWTXdhcmUxIQ==` for admin:VMware1!VMware1!.

- Header2
  - Name: Content-Type
  - Value: application/json
- Header3
  - Name: Accept
  - Value: application/json

A response indicating that the NSX Manager is running will be:

```
"healthy" : true
```

Note that the format of the response is `"healthy"<space>:<space>true`.

If you change the password of the user that you specify in Header1, you must update Header1 accordingly.

## Configure Proxy Settings

You can configure proxy settings for your NSX-T Data Center environment.

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **System > Proxy Settings**.
- 3 In the **Scheme** column, select **HTTP** or **HTTPS**.
- 4 In the **Host** column, enter an IP address.
- 5 In the **Port** column, enter a port number.
- 6 In the **Username** column, enter a user name.
- 7 In the **Password** column, enter a password.
- 8 Click **Save**.

## View Container-Related Information

You can view container-related information such as hyperbus status or the status of NCP (NSX Container Plug-in) clusters.

### Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **System > Fabric > Nodes** from the navigation panel.
- 3 To see the hyperbus and NSX node agent status, click the **Host Transport Nodes** tab.
- 4 Click the node where NSX node agent is running.
- 5 Click the **Monitor** tab.

The **Hyperbus Status** field is near the bottom of the page. The hyperbus status and the status of every node agent VIFID are displayed.

- 6 To see information about NCP Clusters, click the **NCP Clusters** tab.

For each cluster, the cluster name, ID, status and type are displayed.



# Using NSX Cloud

# 23

NSX Cloud enables you to manage and secure your public cloud inventory using NSX-T Data Center.

See "Installing NSX Cloud Components" in the *NSX-T Data Center Installation Guide* for the NSX Cloud deployment workflow.

See also: [public cloud](#).

Read the following topics next:

- [Cloud Service Manager: UI Walkthrough](#)
- [Threat Detection using the NSX Cloud Quarantine Policy](#)
- [NSX Enforced Mode](#)
- [Native Cloud Enforced Mode](#)
- [NSX-T Data Center Features Supported with NSX Cloud](#)
- [NSX Cloud FAQs and Troubleshooting](#)

## Cloud Service Manager: UI Walkthrough

The Cloud Service Manager (CSM) provides a single pane of glass management endpoint for your public cloud inventory.

The CSM interface is divided into the following categories:

- **Search:** You can use the search text box to find public cloud accounts or related constructs.
- **Clouds:** Your public cloud inventory is managed through the sections under this category.
- **System:** You can access **Settings**, **Utilities**, and **Users** for Cloud Service Manager from this category.

You can perform all public cloud operations by going to the **Clouds** subsection of CSM.

To perform system-based operations, such as, backup, restore, upgrade, and user management, go to the **System** subsection.

## Clouds

These are the sections under **Clouds**:

## Overview

Access your public cloud account by clicking **Clouds > Overview**.

Each tile on this page represents your public cloud account with the number of accounts, regions, VPCs or VNets, and instances (workload VMs) it contains.

You can perform the following tasks:

|                                            |                                                                                                                                                                                                                                                                                                                                      |
|--------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Add a public cloud account or subscription | <p>You can add one or more public cloud accounts or subscriptions. This enables you to view your public cloud inventory in CSM. It also shows the number of VMs that are managed by NSX-T Data Center and their state.</p> <p>See "Add Your Public Account" in the <i>NSX-T Data Center Installation Guide</i> for instructions.</p> |
| Deploy/Undeploy NSX Public Cloud Gateway   | <p>You can deploy or undeploy one or two (for High Availability) PCG(s). You can also undeploy PCG from CSM.</p> <p>See "Deploy or Link PCGs" in the <i>NSX-T Data Center Installation Guide</i> for instructions.</p>                                                                                                               |
| Enable or Disable Quarantine Policy        | <p>You can enable or disable Quarantine Policy. See <a href="#">Threat Detection using the NSX Cloud Quarantine Policy</a> for details.</p>                                                                                                                                                                                          |
| Switch between Grid and Card view          | <p>The cards display an overview of your inventory. The grid displays more details. Click the icons to switch between the view types.</p>                                                                                                                                                                                            |

CSM provides a single pane of glass view of all your public cloud accounts that you have connected with NSX Cloud by presenting your public cloud inventory in different ways:

- You can view the number of regions you are operating in.
- You can view the number of VPCs/VNets per region.
- You can view the number of workload VMs per VPC/VNet.

There are four tabs under **Clouds**.

## Accounts

You can add your public cloud account by navigating to **Clouds > <your public cloud> > Accounts** section of CSM. You can also view information on the public cloud accounts you have already added.

Each card represents a public cloud account of the cloud provider you selected.

You can perform the following actions from this section:

- Add Account
- Edit Account
- Delete Account
- Resync Account

## Regions

Navigate to **Clouds > <your public cloud> > Regions** to see your inventory for a selected region.

You can filter the regions by your public cloud account. Each region has VPCs/VNets and instances. If you have deployed any PCGs, you can see them here as **Gateways** with an indicator for the PCG's health.

If you do not have any VPCs/VNets in a public cloud region, that region is not displayed in CSM.

## VPCs or VNets

Navigate to **Clouds > <your public cloud> > VPCs or VNets** to view the VPCs or VNets in your public cloud account or subscription.

You can filter the inventory by account and region.

- Each card represents one VPC/VNet.
- You can have one or two (for HA) PCGs deployed on Transit VPCs/VNets.
- You can link Compute VPCs/VNets to Transit VPCs/VNets.
- You can view more details for each VPC or VNet by switching to the grid view.

In the grid view you can see three tabs: **Overview**, **Instances**, and **Segments**.

- **Overview** lists the options under Actions as described in the next step.
- **Instances** displays a list of instances in the VPC/VNet.
- **Segments** displays overlay segments in NSX-T Data Center.

---

**Note** This feature is not supported in the current release for NSX Cloud. Do not tag your workload VMs in AWS or Microsoft Azure with tags shown on this screen.

---

- Click **Actions** to access the following:
  - **Edit Configuration** (only available for Transit VPCs/VNets):
    - Enable or disable Quarantine Policy if in the NSX Enforced Mode.
    - Change your proxy server selection.
  - **Link to Transit VPC/VNet**: This option is only available to VPCs/VNets that do not have any PCG deployed on them. Click to select a Transit VPC/VNet to link to.
  - **Deploy NSX Cloud Gateway**: This option is only available to VPCs/VNets that do not have a PCG deployed on them. Click this option to get started with deploying PCG on this VPC/VNet and make it a Transit or self-managed VPC/VNet. See **Deploy or Link NSX Public Cloud Gateways** in the *NSX-T Data Center Installation Guide* for detailed instructions.

## Instances

The **Clouds > <your public cloud> > Instances** section displays details of the NSX-managed instances in your VPC or VNet.

- You can filter the instance inventory by account, region, and VPC or VNet.
- Each card represents an instance (workload VM) and displays a summary.

- For details on the instance, click on the card or switch to the grid view. Among other details, you can see the following:
  - **Rules Realization:** For workload VMs managed in the Native Cloud Enforced Mode, you can see the status of DFW rules created in NSX Manager. The realization status can be *Successful* or *Failed*. You can click on *failed* status to view the error message. See [Set up Micro-segmentation for Workload VMs in the Native Cloud Enforced Mode](#) for details.



You can add instances to or remove instances from the CSM User Managed list. See [User Managed List for VMs](#) for details.

## CSM Icons





CSM displays the state and health of your public cloud constructs using descriptive icons.





**Note** In the Native Cloud Enforced Mode: Quarantine Policy is always enabled and all VMs are always NSX-managed. Only the states where Quarantine Policy is enabled for NSX-managed VMs apply in this mode.

In the NSX Enforced Mode: Quarantine Policy can be disabled and it is possible to have unmanaged VMs in the VPC/VNet. All relevant states apply to this mode.

| CSM Section and Icon                                                                | Description      |
|-------------------------------------------------------------------------------------|------------------|
| VPCs/VNets                                                                          |                  |
|  | Transit VPC/VNet |
|  | Compute VPC/VNet |

| CSM Section and Icon | Description                                              |
|----------------------|----------------------------------------------------------|
|                      | Self-Managed VPC/VNet                                    |
|                      | VPC/VNet showing healthy PCGs                            |
|                      | VPC/VNet showing PCGs in error state                     |
|                      | VPC/VNet showing one PCG in error state and one healthy. |

| CSM Section and Icon                                                                | Description                       |
|-------------------------------------------------------------------------------------|-----------------------------------|
|    | VPC/VNet showing NSX-managed VMs. |
|    | VPC/VNet showing unmanaged VMs.   |
|   | VPC/VNet showing VMs with errors. |
|  | VPC/VNet showing powered-off VMs. |
| Instances                                                                           |                                   |

| CSM Section and Icon                                                                | Description                                                 |
|-------------------------------------------------------------------------------------|-------------------------------------------------------------|
|    | NSX-managed VMs with no errors.                             |
|    | NSX-managed VMs with errors and Quarantine Policy disabled. |
|  | NSX-managed VMs with errors and Quarantine Policy enabled.  |
|  | Unmanaged VMs quarantined.                                  |

## System

These are the sections under **System**:

## System > Settings

These settings are first configured when you install CSM. You can edit them thereafter.

### Join CSM with NSX Manager

You must connect the CSM appliance with NSX Manager to allow these components to communicate with each other.

#### Prerequisites

- NSX Manager must be installed and you must have the username and password for the admin account to log in to NSX Manager.
- CSM must be installed and you must have the Enterprise Administrator role assigned in CSM.
- You must have the NSX Data Center Enterprise Plus license.

#### Procedure

- 1 From a browser, log in to CSM.
- 2 When prompted in the setup wizard, click **Begin Setup**.
- 3 Enter the following details in the NSX Manager Credentials screen:

| Option                       | Description                                                                                                                                                   |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NSX Manager Host Name</b> | Enter the fully qualified domain name (FQDN) of the NSX Manager, if available. You may also enter the IP address of NSX Manager.                              |
| <b>Admin Credentials</b>     | Enter an Enterprise Administrator username and password for NSX Manager.                                                                                      |
| <b>Manager Thumbprint</b>    | Optionally, enter the NSX Manager's thumbprint value. If you leave this field blank, the system identifies the thumbprint and displays it in the next screen. |

- 4 (Optional) If you did not provide a thumbprint value for NSX Manager, or if the value was incorrect, the **Verify Thumbprint** screen appears. Select the checkbox to accept the thumbprint discovered by the system.
- 5 Click **Connect**.

**Note** If you missed this setting in the setup wizard or if you want to change the associated NSX Manager, log in to CSM, click **System > Settings**, and click **Configure** on the panel titled **Associated NSX Node**.

CSM verifies the NSX Manager thumbprint and establishes connection.

- 6 (Optional) Set up the Proxy server. See instructions in [\(Optional\) Configure Proxy Servers](#).

### (Optional) Configure Proxy Servers

If you want to route and monitor all internet-bound HTTP/HTTPS traffic through a reliable HTTP Proxy, you can configure up to five proxy servers in CSM.



All public cloud communication from PCG and CSM is routed through the selected proxy server.

Proxy settings for PCG are independent of proxy settings for CSM. You can choose to have none or a different proxy server for PCG.

You can choose the following levels of authentication:

- Credentials-based authentication.
- Certificate-based authentication for HTTPS interception.
- No authentication.

#### Procedure

- 1 Click **System > Settings**. Then click **Configure** on the panel titled **Proxy Servers**.

**Note** You can also provide these details when using the CSM Setup Wizard that is available when you first install CSM.

- 2 In the Configure Proxy Servers screen, enter the following details:

| Option                | Description                                                                                                                                                              |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>        | Use this radio button to indicate the default proxy server.                                                                                                              |
| <b>Profile Name</b>   | Provide a proxy server profile name. This is mandatory.                                                                                                                  |
| <b>Proxy Server</b>   | Enter the proxy server's IP address. This is mandatory.                                                                                                                  |
| <b>Port</b>           | Enter the proxy server's port. This is mandatory.                                                                                                                        |
| <b>Authentication</b> | Optional. If you want to set up additional authentication, select this checkbox and provide valid username and password.                                                 |
| <b>Username</b>       | This is required if you select the Authentication checkbox.                                                                                                              |
| <b>Password</b>       | This is required if you select the Authentication checkbox.                                                                                                              |
| <b>Certificate</b>    | Optional. If you want to provide an authentication certificate for HTTPS interception, select this checkbox and copy-paste the certificate in the text box that appears. |
| <b>No Proxy</b>       | Select this option if you do not want to use any of the proxy servers configured.                                                                                        |

## System > Utilities

The following utilities are available.

### Backup and Restore

Follow the same instructions for backing up and restoring CSM, as you do for NSX Manager. See [Backup and Restore of the CSM Appliance](#).

### Support Bundle

Click **Download** to retrieve the support bundle for CSM. This is used for troubleshooting. See the *NSX-T Data Center Troubleshooting Guide* for more information.

## Backup and Restore of the CSM Appliance

You can back up a CSM appliance and restore from a backup.

Currently only IP address backups and restores are supported for CSM.

### Backing up CSM

You can backup the CSM appliance manually or set up a recurring backup after you configure a backup server.

You can only restore the CSM appliance from an IP address backup. Do not configure FQDN for the CSM appliance.

### Prerequisites

- Verify that you have the SSH fingerprint of the backup file server. Only SHA256 hashed ECDSA key is accepted as a fingerprint. See [Find the SSH Fingerprint of a Remote Server](#).
- Ensure that the directory path already exists where you want to store your backups. You cannot use the root directory (/).

### Procedure

- 1 From your browser, log in with admin privileges to CSM at `https://<csm-ip-address>`.
- 2 Select **System > Utilities > Tools**.
- 3 On the **Backup** tab, click **Configure**.
- 4 Enter the IP address or host name of the backup file server.
- 5 Change the default port if required.
- 6 The protocol field is already filled in. Do not change the value.  
SFTP is the only supported protocol.
- 7 Enter the username and password required to log in to the backup file server.  
  
The first time you configure a file server, you must provide a password. Subsequently, if you reconfigure the file server, and the server IP (or hostname), port, and user name are the same, you do not need to enter the password again.
- 8 In the **Destination Directory** field, enter the absolute directory path where the backups will be stored.

The directory must already exist and cannot be /. If the backup file server is a Windows machine, you still use the forward slash when you specify the destination directory. For example, if the backup directory on the Windows machine is `c:\SFTP_Root\backup`, specify `/SFTP_Root/backup` as the destination directory.

---

**Note** The backup process will generate a name for the backup file that can be quite long. On a Windows server, the length of the full path name of the backup file can exceed the limit set by Windows and cause backups to fail. To avoid this issue, see the KB article <https://kb.vmware.com/s/article/76528>.

---

**9** To encrypt the backups, enter an **Encryption Passphrase**.

You will need this passphrase to restore a backup. If you forget the passphrase, you cannot restore any backups.

**10** Enter the SSH fingerprint of the server that stores the backups.

You can leave this blank and accept or reject the fingerprint provided by the server.

**11** Click the **Schedule** tab.**12** To enable automatic backups, click the **Automatic Backup** toggle.**13** Click **Weekly** and set the days and time of the backup, or click **Interval** and set the interval between backups.**14** Enabling the **Detect NSX configuration change** option will trigger an unscheduled full configuration backup when it detects any runtime or non-configuration related changes, or any change in user configuration.

You can specify a time interval for detecting database configuration changes. The valid range is 5 minutes to 1,440 minutes (24 hours).

---

**Note** This option can potentially generate a large number of backups. Use it with caution.

---

**15** Click **Save**.**Results**

After you configure a backup file server, you can click **Backup Now** to start a backup at any time.

If your backup server is getting full, see instructions for removing backups: [Remove Old Backups](#).

**Restoring CSM from a backup**

You can restore a CSM appliance if you have a backup.

You must restore a backup on a new installation of CSM. If the old CSM node is still available, you must power it off, before you start the restore process.

---

**Note** You can only restore CSM from an IP address backup. FQDN backups are not supported for CSM.

---

**Prerequisites**

- Verify that you have the login credential for the backup file server.
- Verify that you have the SSH fingerprint of the backup file server. Only an SHA256 hashed ECDSA key is accepted as a fingerprint. See [Find the SSH Fingerprint of a Remote Server](#).
- Verify that you have the passphrase of the backup file.

**Procedure**

- 1** If the old CSM node is still available, power it off.
- 2** Deploy a new CSM node with the same IP address of the original CSM node.

- 3 From your browser, log in with admin privileges to a new CSM appliance.
- 4 Select **System > Utilities > Tools**.
- 5 Click the **Restore** tab.
- 6 Click **Restore Now**. The Restore wizard opens.
- 7 Select the check box on the **Prerequisites** screen.
- 8 Provide details of the remote backup server:
  - a Enter the IP address or host name.
  - b Change the port number, if necessary.  
The default is 22.
  - c To log in to the server, enter the user name and password.
  - d In the **Backup Directory** text box, enter the absolute directory path where the backups are stored.
  - e Enter the passphrase that was used to encrypt the backup data.
  - f Enter the SSH fingerprint of the server that stores the backups.
- 9 Click **Next**.
- 10 Select a backup. You can also get a list of available backups by logging in to the backup file server. See [Listing Available Backups](#). Replace NSX Manager with CSM in these instructions, for example, when you are asked to log in to the NSX Manager to run a CLI command, log in to CSM instead.
- 11 Click **Restore**.  
  
You lose connectivity until the restore completes. The status of the restore operation is displayed. After the restore operation is completed, the Restore Complete screen is displayed, showing the result of the restore, the timestamp of the backup file, and the start and end time of the restore operation. If the restore failed, the screen displays the step where the failure occurred.  
  
You can also determine the reason for a restore failure by selecting the log files. Run `get log-file syslog` to view the system log file.  
  
To restart CSM, run the `service nsx-cloud-service-manager restart` command.  
  
To reboot the CSM node, run the `reboot` command.
- 12 After the new CSM node is deployed, delete the original CSM VM that you powered down in Step 1.

## System > Users

Users are managed using role-based access control (RBAC).

See [Chapter 18 Authentication and Authorization](#) for details.

## Threat Detection using the NSX Cloud Quarantine Policy

The Quarantine Policy feature in NSX Cloud provides a threat detection mechanism for your NSX-managed workload VMs.

Quarantine Policy is implemented differently in the two VM-management modes.

**Table 23-1. Quarantine Policy Implementation in the NSX Enforced Mode and the Native Cloud Enforced Mode**

| Configurations related to Quarantine Policy                     | In the NSX Enforced Mode                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | In the Native Cloud Enforced Mode                                                                                                                                            |
|-----------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Default state                                                   | Disabled when deploying PCG using NSX Tools. You can enable it from the PCG-deployment screen or later. See <a href="#">How to Enable or Disable Quarantine Policy</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Always enabled. Cannot be disabled.                                                                                                                                          |
| Auto-created security groups unique to each mode                | All healthy NSX-managed VMs are assigned the <code>vm-underlay-sg</code> security group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <code>nsx-&lt;NSX GUID&gt;</code> security groups are created for and applied to NSX-managed workload VMs that are matched with a Distributed Firewall Policy in NSX Manager |
| Auto-created Public Cloud Security Groups common to both modes: | <p>The <b>gw</b> security groups are applied to the respective PCG interfaces in AWS and Microsoft Azure.</p> <ul style="list-style-type: none"> <li>■ <code>gw-mgmt-sg</code></li> <li>■ <code>gw-uplink-sg</code></li> <li>■ <code>gw-vtep-sg</code></li> </ul> <p>The <b>vm</b> security groups are applied to NSX-managed VMs depending on their current state and whether Quarantine Policy is enabled or disabled:</p> <ul style="list-style-type: none"> <li>■ <code>default-vnet-&lt;vnet-id&gt;-sg</code> in Microsoft Azure and <code>default</code> in AWS.</li> </ul> <p><b>Note</b> In AWS, the <code>default</code> security group already exists. It is not created by NSX Cloud.</p> |                                                                                                                                                                              |

### General Recommendation for NSX Enforced Mode :

Start with *disabled* for **Brownfield** deployments: Quarantine Policy is disabled by default. When you already have VMs set up in your public cloud environment, use the disabled mode for Quarantine Policy until you onboard your workload VMs. This ensures that your existing VMs are not automatically quarantined.

Start with *enabled* for **Greenfield** deployments: For greenfield deployments, it is recommended that you enable Quarantine Policy to allow threat detection for your VMs to be managed by NSX Cloud.

### Quarantine Policy in the NSX Enforced Mode

Enabling Quarantine Policy is optional in the NSX Enforced Mode.

## How to Enable or Disable Quarantine Policy

In the NSX Enforced Mode, you can elect to enable Quarantine Policy in two ways.

The first possibility to enable Quarantine Policy is when you deploy PCG on a Transit VPC/VNet or link a Compute VPC/VNet to a Transit. Move the slider for **Quarantine Policy on the Associated VPC/VNet** to **Enabled** from the default **Disabled** state. See **Deploy PCG** in the *NSX-T Data Center Installation Guide*.

You can also enable Quarantine Policy later following the steps here.

### Prerequisites

If enabling Quarantine Policy after deploying or linking to a PCG, you must have one or more Transit or Compute VPCs/VNets onboarded in the NSX Enforced Mode, that is you elected to use NSX Tools for managing your workload VMs.

### Procedure

- 1 Log in to CSM and go to your public cloud:
  - a If using AWS, go to **Clouds > AWS > VPCs**. Click on the Transit or Compute VPC.
  - b If using Microsoft Azure, go to **Clouds > Azure > VNets**. Click on the Transit or Compute VNet.

- 2 Enable the option using any one of the following:

- In the tile view, click on **ACTIONS > Edit Configuration**.



- If you are in the grid view, select the checkbox next to the VPC or VNet and click



- ◆ If you are in the VPC or VNet's page, click the ACTIONS icon to go to **Edit Configurations**.



- 3 Turn **Default Quarantine** on or off.

- 4 Click **SAVE**.

## Quarantine Policy Impact when Disabled

NSX Cloud does not manage the public cloud security groups of untagged VMs when Quarantine Policy is disabled.

However, for VMs tagged with `nsx.network=default` in the public cloud, NSX Cloud assigns appropriate security groups depending on the VM's state. This behavior is similar to when the Quarantine Policy is enabled, but the rules in the quarantine security groups: `default-vnet-<vnet-id>-sg` in Microsoft Azure and `default` in AWS are configured similar to default public cloud security groups, allowing everything within the VPC/VNet and denying all other inbound traffic. Any manual changes to the security groups of tagged VMs are reverted to the NSX Cloud-assigned security group within two minutes.

**Note** If you do not want NSX Cloud to assign security groups to your NSX-managed (tagged) VMs, add them to the User Managed list in CSM. See [User Managed List for VMs](#).

The following table shows how NSX Cloud manages the public cloud security groups of workload VMs when Quarantine Policy is disabled.

**Table 23-2. NSX Cloud assignment of public cloud security groups when Quarantine Policy is disabled**

| Is VM tagged with <i>nsx.network=default</i> in the public cloud? | Is VM added to the User Managed List? | VM's Public cloud security group when Quarantine Policy is disabled and explanation                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------------------------------------------|---------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VM could be tagged or not tagged                                  | Added to the User Managed list.       | Retains existing public cloud security group because NSX Cloud doesn't take any action on VMs in the User Managed list.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Not tagged                                                        | Not added to the User Managed List    | Retains existing public cloud security group because NSX Cloud doesn't take action on untagged VMs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Tagged                                                            | Not added to the User Managed List    | <ul style="list-style-type: none"> <li>■ If VM has no threats: <code>vm-underlay-sg</code></li> <li>■ If VM has potential threats (see note): <code>default-vnet-&lt;vnet-id&gt;-sg</code> in Microsoft Azure; <code>default</code> in AWS</li> </ul> <p><b>Note</b> The assignment of public cloud security groups is triggered within 90 seconds of applying the <code>nsx.network=default</code> tag to your workload VMs. You still need to install NSX Tools for the VMs to be NSX-managed. Until NSX Tools are installed, your tagged workload VMs remain in the default security group.</p> |

The following table shows how NSX Cloud manages the public cloud security groups of VMs if Quarantine policy was enabled before and is now disabled:

**Table 23-3. NSX Cloud assignment of public cloud security groups when Quarantine Policy is disabled from being enabled at first**

| Is VM tagged with <i>nsx.network=default</i> in the public cloud? | Is VM in the User Managed list?     | VM's existing public cloud security group when Quarantine Policy is enabled                                  | VM's public cloud security group after Quarantine Policy is disabled                                                                                                                                                                                                                                                                                                            |
|-------------------------------------------------------------------|-------------------------------------|--------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VM could be tagged or not tagged                                  | Yes, VM is in the User Managed list | Any existing public cloud security group                                                                     | Retains existing public cloud security group because NSX Cloud doesn't take any action on VMs in the User Managed list.<br><br><b>Note</b> If you have a VM in the User Managed list in any NSX Cloud-assigned security groups, you must manually move it to <i>default</i> security group in AWS and <i>default-vnet-&lt;vnet-id&gt;-sg</i> security group in Microsoft Azure. |
| Not tagged                                                        | Not added to the User Managed List  | <i>default-vnet-&lt;vnet-id&gt;-sg</i> (Microsoft Azure)<br>Or <i>default</i> (AWS)                          | Remains in the existing security groups when disabling the Quarantine Policy because it is untagged and not considered NSX-managed. You can manually assign any other security group to this VM as required.                                                                                                                                                                    |
| Tagged                                                            | Not added to the User Managed List  | <i>vm-underlay-sg</i><br>Or <i>default-vnet-&lt;vnet-id&gt;-sg</i> (Microsoft Azure) Or <i>default</i> (AWS) | Retains the NSX Cloud-assigned security group because that is consistent for tagged VMs in the Quarantine enabled or disabled modes.                                                                                                                                                                                                                                            |

## Quarantine Policy Impact when Enabled

NSX Cloud manages the public cloud security group of all workload VMs in this VPC/VNet when Quarantine Policy is enabled.

Any manual changes to the security groups are reverted to the NSX Cloud-assigned security group within two minutes. If you do not want NSX Cloud to assign security groups to your VMs, add them to the User Managed list in CSM. See [User Managed List for VMs](#).

**Note** Removing the VM from the User Managed list causes the VM to revert to the NSX Cloud-assigned security group.



**Table 23-4. NSX Cloud assignment of public cloud security groups when Quarantine Policy is enabled**

| Is VM tagged with <i>nsx.network=default</i> in the public cloud)? | Is VM in the User Managed list?     | VM's public cloud security group when Quarantine Policy is enabled and explanation                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------------------------------------------------------|-------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Tagged                                                             | Not added to the User Managed List  | <ul style="list-style-type: none"> <li>■ If VM has no threats: <code>vm-underlay-sg</code></li> <li>■ If VM has potential threats (see note): <code>default-vnet-&lt;vnet-ID&gt;-sg</code> in Microsoft Azure; <code>default</code> in AWS</li> </ul> <p><b>Note</b> The assignment of public cloud security groups is triggered within 90 seconds of applying the <code>nsx.network=default</code> tag to your workload VMs. You still need to install NSX Tools for the VMs to be NSX-managed. Until NSX Tools are installed your tagged workload VMs are quarantined.</p> |
| Not Tagged                                                         | Not added to the User Managed List  | <code>default-vnet-&lt;vnet-ID&gt;-sg</code> in Microsoft Azure; <code>default</code> in AWS. Untagged VMs are considered unmanaged and therefore quarantined by NSX Cloud.                                                                                                                                                                                                                                                                                                                                                                                                  |
| Tagged                                                             | Yes, VM is in the User Managed List | Retains existing public cloud security group because NSX Cloud doesn't take action on VMs in the User Managed list.                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Not Tagged                                                         |                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

The following table captures the impact on security group assignments if the Quarantine Policy was disabled at first and then you enable it:

**Table 23-5. NSX Cloud assignment of public cloud security groups when Quarantine Policy is enabled from being disabled at first**

| Is VM tagged with <i>nsx.network=default</i> in the public cloud? | Is VM in the User Managed list?    | VM's existing public cloud security group when Quarantine Policy is disabled                                                      | VM's public cloud security group after Quarantine Policy is enabled                                                          |
|-------------------------------------------------------------------|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------|
| Not Tagged                                                        | Not added to the User Managed List | Any existing public cloud security group                                                                                          | <code>default-vnet-&lt;vnet-ID&gt;-sg</code> (Microsoft Azure) Or <code>default</code> (AWS)                                 |
| Tagged                                                            | Not added to the User Managed List | <code>vm-underlay-sg</code><br>Or <code>default-vnet-&lt;vnet-ID&gt;-sg</code> (Microsoft Azure)<br>Or <code>default</code> (AWS) | Retains the NSX Cloud-assigned security group that is consistent for tagged VMs in the Quarantine enabled or disabled modes. |

**Table 23-5. NSX Cloud assignment of public cloud security groups when Quarantine Policy is enabled from being disabled at first (continued)**

| Is VM tagged with <i>nsx.network=default</i> in the public cloud? | Is VM in the User Managed list?     | VM's existing public cloud security group when Quarantine Policy is disabled | VM's public cloud security group after Quarantine Policy is enabled                                                     |
|-------------------------------------------------------------------|-------------------------------------|------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| Tagged                                                            | Yes, VM is in the User Managed List | Any existing public cloud security group.                                    | Retains existing public cloud security group because NSX Cloud doesn't take any action on VMs in the User Managed list. |
| Not Tagged                                                        |                                     |                                                                              |                                                                                                                         |

## Quarantine Policy in the Native Cloud Enforced Mode

Quarantine Policy is always enabled in the Native Cloud Enforced Mode.

**Table 23-6. Assignment of Public Cloud Security Groups in the Native Cloud Enforced Mode**

| Is VM part of a valid NSX-T Security policy?          | Is VM added to the User Managed List? | VM's public cloud security group and explanation                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------------------------------------------|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Yes, VM is matched with a valid NSX-T Security Policy | Not added to User Managed List        | NSX Cloud-created public cloud security group named like <i>nsx-{NSX-GUID}</i> which is the corresponding public cloud security group for the NSX-T Security Policy.                                                                                                                                                                                                                                                     |
| No, VM does not have a valid NSX-T firewall policy    | Not added to User Managed List        | <p><i>default-vnet-&lt;vnet-ID&gt;-sg</i> in Microsoft Azure or <i>default</i> in AWS because this is the threat detection behavior of NSX Cloud. In the Native Cloud Enforced Mode, the NSX Cloud-created security groups <i>default-vnet-&lt;vnet-ID&gt;-sg</i> in Microsoft Azure or <i>default</i> in AWS mimic the default public cloud security policy.</p> <p><b>Note</b> In CSM the VM shows an Error state.</p> |
| Yes, VM has valid NSX-T Security policy               | Added to User Managed list            | Retains existing public cloud security group because NSX Cloud doesn't take any action on VMs added to the User Managed list.                                                                                                                                                                                                                                                                                            |
| No, VM does not have a valid NSX-T Security policy    |                                       |                                                                                                                                                                                                                                                                                                                                                                                                                          |

## User Managed List for VMs

Adding VMs to the **User Managed** list is an option available from CSM for all workload VMs in your public cloud inventory.

You can add VMs to the **User Managed** in both the VM-management modes: NSX Enforced Mode and the Native Cloud Enforced Mode.

### Why to add VMs to the User Managed list?

- In the NSX Enforced Mode: If you have the Quarantine Policy enabled and you need to verify any specific DFW policies with existing applications on the VM, add such a VM to the **User Managed** list before onboarding it with NSX Cloud.

- In either the NSX Enforced Mode or the Native Cloud Enforced Mode:
  - If you have VMs with errors and want to access them to resolve the errors, add such VMs to the **User Managed** list so you can move them out of the quarantine state and use debugging tools as required.
  - Add VMs to the **User Managed** list in your public cloud inventory that you don't want NSX-T to manage, e.g. DNS Forwarder, Proxy server etc.

## How to use the User Managed List

Follow these instructions to add VMs to the **User Managed** list or remove them.

### Prerequisites

You must have one or more public cloud accounts added to CSM.

### Procedure

- 1 Log in to CSM using an Enterprise Admin account and go to your public cloud account.
  - a If using AWS, go to **Clouds > AWS > VPCs > Instances**.
  - b If using Microsoft Azure, go to **Clouds > Azure > VNets > Instances**.
- 2 If in Tiles mode, switch to Grid mode by clicking the mode selector in the right corner of the instances view.
- 3 Select the VMs (instances) that you want to add to or remove from the **User Managed** list.
- 4 Click **Actions** and select either **Add to User Managed List** or **Remove from User Managed List**.
- 5 Go back to the Accounts tab, select the account tile and click **Actions > Resync Account**.

### Results

Each VM added to the **User Managed** list remains in the security group it was assigned before adding to the **User Managed** list. You can now apply any other security group to the VM as required. NSX Cloud ignores VMs in the **User Managed** list regardless of the status of Quarantine Policy.

If you remove a VM from the **User Managed** list in the Native Cloud Enforced Mode or remove an NSX-managed VM from the **User Managed** list in the NSX Enforced Mode, NSX Cloud starts assigning security groups to that VM depending on its state.

## NSX Enforced Mode

In the NSX Enforced Mode, that is, by using NSX Tools, you must first onboard VMs by tagging them in the public cloud and installing NSX Tools on them, before starting to manage these VMs using NSX-T Data Center.

## Currently Supported Operating Systems for Workload VMs

This is the list of operating systems currently supported by NSX Cloud for your workload VMs in the NSX Enforced Mode.

Currently, the following operating systems are supported:

---

**Note** See the NSX Cloud Known Issues section in the *NSX-T Data Center Release Notes* for exceptions. For supported operating systems it is assumed that you are using the standard Linux kernel versions. Public cloud marketplace images with custom kernels, for example, upstream Linux kernel with modified sources, are not supported.

---

- SUSE Linux Enterprise Server (SLES) 12 SP3
- Red Hat Enterprise Linux (RHEL) 7.2, 7.3, 7.4, 7.5, 7.6
- CentOS 7.2, 7.3, 7.4, 7.5, 7.6

---

**Note** RHEL Extended Update Support (EUS) kernel in RHEL and CentOS are not supported.

---

**Note** Only the CentOS marketplace images whose distribution versions match their expected minor kernel versions are supported for NSX Cloud. For example, the distribution versions and their corresponding kernel versions are expected to be as follows:

| RHEL version | Kernel version |
|--------------|----------------|
| RHEL 7.6     | 3.10.0-957     |
| RHEL 7.5     | 3.10.0-862     |
| RHEL 7.4     | 3.10.0-693     |
| RHEL 7.3     | 3.10.0-514     |
| RHEL 7.2     | 3.10.0-327     |

---

- Ubuntu 14.04, 16.04, 18.04
- Microsoft Windows Server 2016 - Service based release, Desktop experience(1709, 1803, 1809)
- Microsoft Windows Server 2019 Datacenter
- Microsoft Windows Server 2012 R2
- Microsoft Windows 10 versions 1809, 1803, 1709 (only supported in Microsoft Azure in the current NSX Cloud release)

## Onboarding VMs in the NSX Enforced Mode

Refer to this workflow for an overview of the steps involved in onboarding and managing workload VMs from your public cloud in the NSX Enforced Mode.

Table 23-7. Day-N Workflow for onboarding your workload VMs into NSX Cloud

| Task                                                                                                                                                                                                                                         | Instructions                                                                     |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------|
|  Tag workload VMs with the key-value <code>nsx.network=default</code> .                                                                                     | Follow instructions in your public cloud documentation for tagging workload VMs. |
|  Install NSX Tools on your Windows and Linux workload VMs.                                                                                                  | See <a href="#">Install NSX Tools</a>                                            |
| <b>Note</b> If <b>Auto- Install NSX Tools</b> is enabled in CSM for Microsoft Azure VNets, NSX Tools are automatically installed.                                                                                                            |                                                                                  |
|  (Optional) If you have added VMs to the User Managed list in CSM, remove those VMs from the User Managed list that you want to bring under NSX management. | See <a href="#">How to use the User Managed List</a> .                           |

## Tag VMs in the Public Cloud

Apply the `nsx.network=default` tag to VMs that you want to manage using NSX-T Data Center.

### Procedure

- 1 Log in to your public cloud account and go to your VPC or VNet where you want your workload VMs to be managed by NSX-T Data Center.
- 2 Select the VMs that you want to manage using NSX-T Data Center.
- 3 Add the following tag details for the VMs and save your changes.

```
Key: nsx.network
Value: default
```

**Note** Apply this tag at the VM level.

### Results

You may have already onboarded the VPCs/VNets where you applied the `nsx.network=default` tags to workload VMs. You can also onboard these VPCs/VNets after applying the tag. Successful onboarding of the VPC/VNet results in the workload VMs to be considered NSX-managed.

### What to do next

Install NSX Tools on these VMs. See [Install NSX Tools](#).

If using Microsoft Azure, you have the option to auto-install NSX Tools on tagged VMs. See [Install NSX Tools Automatically](#) for details.

## Install NSX Tools

Install NSX Tools on your workload VMs

There are several options available to install NSX Tools:

- Download and install NSX Tools in individual workload VMs. Linux and Windows VMs have some variations.
- Use replicable images with NSX Tools installed on them using your public cloud's supported method, for example, create an AMI in AWS or a Managed Image in Microsoft Azure.
- AWS-only: When launching VMs, provide the NSX Tools download location and installation command in **User Data**.
- Microsoft Azure-only: Enable auto-installation of NSX Tools when deploying PCG in a Microsoft Azure VNet or while linking to a Transit VNet, or by editing a Transit/Compute VNet's Configuration.

---

**Note** If you have whitelisted workload VMs on which you want to install NSX Tools, ensure the following ports are open in the security groups you have assigned to such VMs:

- Inbound UDP 6081 : For overlay data packets. This should be allowed for (Active/Standby) PCG's VTEP IP address (eth1 interface).
  - Outbound TCP 5555 : For control packets. This should be allowed for (Active/Standby) PCG's management IP address (eth0 interface).
  - TCP 8080 : For install/upgrade on the PCG's management IP address.
  - TCP 80: For downloading any third party dependencies while installing NSX Tools.
  - UDP 67,68: For DHCP packets.
  - UDP 53: For DNS resolution.
- 

### Install NSX Tools on Linux VMs

To install NSX Tools on your Linux workload VMs, follow these instructions.

See [Currently Supported Operating Systems for Workload VMs](#) for a list of Linux distributions currently supported.

---

**Note** To verify the checksum of this script, go to **VMware Downloads > Drivers & Tools > NSX Cloud Scripts**.

---

### Prerequisites

You need the following commands to run the NSX Tools installation script:

- **wget**
- **nslookup**
- **dmidecode**

## Procedure

- 1 Log in to CSM and go to your public cloud:
  - a If using AWS, go to **Clouds > AWS > VPCs**. Click a Transit or Compute VPC.
  - b If using Microsoft Azure, go to **Clouds > Azure > VNets**. Click the VNet on which one or a pair of PCGs is deployed and running.

**Note:** Transit VPC/VNet is where one or a pair of PCGs is deployed and running. Compute VPC/VNet is the one linked to a Transit and can use the PCG instances deployed there.

- 2 From the **NSX Tools Download & Installation** section of the screen, make a note of the **Download Location** and the **Installation Command** under **Linux**.

---

**Note** For VNets, the DNS Suffix in the Installation Command is dynamically generated to match the DNS settings you select when deploying PCG. For Transit VNets, the `-dnsServer <dns-server-ip>` parameter is optional. For Compute VNets, you must provide the DNS Forwarder IP address to complete this command.

---

- 3 Log in to the Linux workload VM with superuser privileges.
- 4 Use `wget` or equivalent to download the installation script on your Linux VM from the **Download Location** you noted from CSM. The installation script is downloaded in the directory where you run the `wget` command.

---

**Note** To verify the checksum of this script, go to **VMware Downloads > Drivers & Tools > NSX Cloud Scripts**.

---

- 5 Change permissions on the installation script to make it executable if necessary, and run it:

```
$ chmod +x install_nsx_vm_agent.sh && sudo ./install_nsx_vm_agent.sh
```

**Note:** On Red Hat Enterprise Linux and its derivatives, SELinux is not supported. To install NSX Tools, disable SELinux.

- 6 You lose connectivity with your Linux VM after installation of NSX Tools begins. Messages such as the following appear on your screen: `Installation completed!!! Starting NSX Agent service. SSH connection will now be lost..` To complete the onboarding process, log in to your VM again.

## Results

NSX Tools are installed on your workload VM.

---

**Note**

- After NSX Tools are successfully installed, port 8888 shows as open on the workload VM but it is blocked for VMs in the underlay mode and must be used only when required for advanced troubleshooting. You can access workload VMs over port 8888 using a jumphost if the jumphost is also in the same VPC as the workload VMs that you want to access.
  - The script uses `eth0` as the default interface.
- 

What to do next

### Managing VMs in the NSX Enforced Mode

#### Install NSX Tools on Windows VMs

Follow these instructions to install NSX Tools on your Windows workload VM.

See [Currently Supported Operating Systems for Workload VMs](#) for a list of Microsoft Windows versions currently supported.

---

**Note** To verify the checksum of this script, go to **VMware Downloads > Drivers & Tools > NSX Cloud Scripts**.

---

#### Procedure

- 1 Log in to CSM and go to your public cloud:
  - a If using AWS, go to **Clouds > AWS > VPCs**. Click on a Transit or Compute VPC.
  - b If using Microsoft Azure, go to **Clouds > Azure > VNets**. Click on the VNet on which one or a pair of PCGs is deployed and running.

**Note:** Transit VPC/VNet is where one or a pair of PCGs is deployed and running. Compute VPC/VNet is the one linked to a Transit and can use the PCGs deployed there.

- 2 From the **NSX Tools Download & Installation** section of the screen, make a note of the **Download Location** and the **Installation Command** under **Windows**.

---

**Note** For VNets, the DNS Suffix in the Installation Command is dynamically generated to match the DNS settings you choose when deploying PCG. For Transit VNets, the `-dnsServer <dns-server-ip>` parameter is optional. For Compute VNets, you must provide the DNS Forwarder IP address to complete this command.

---

- 3 Connect to your Windows workload VM as Administrator.



- 4 Download the installation script on your Windows VM from the **Download Location** you noted from CSM. You can use any browser, for example, Internet Explorer, to download the script. It is downloaded in your browser's default downloads directory, for example, *C:\Downloads*.

---

**Note** To verify the checksum of this script, go to **VMware Downloads > Drivers & Tools > NSX Cloud Scripts**

---

**Note:**

- 5 Open a PowerShell prompt and go to the directory containing the downloaded script.
- 6 Use the **Installation command** you noted from CSM to run the downloaded script.

For example:

```
c:\> powershell -file 'nsx_install.ps1' -operation install -dnsSuffix <>
```

---

**Note** The file argument needs the full path unless you are in the same directory or if the PowerShell script is already in the path. For example, if you download the script to *C:\Downloads*, and you are currently not in that directory, then the script must contain the location: *powershell -file 'C:\Downloads\nsx\_install.ps1' ...*

---

- 7 The script runs and when completed, displays a message indicating whether NSX Tools was installed successfully.

---

**Note** The script considers the primary network interface as the default.

---

**What to do next**

[Managing VMs in the NSX Enforced Mode](#)

### Generate Replicable Images

You can generate an AMI in AWS or a Managed Image in Microsoft Azure of a VM with the NSX agent installed on it.

With this feature, you can launch multiple VMs with the agent configured and running.

There are two ways in which you can generate an AMI/Managed Image (image in the rest of this topic) of a VM with the NSX agent installed on it:

- **Generate image with an unconfigured NSX agent:** You can generate an image from a VM that has the NSX agent installed on it but not configured by using the `-noStart` option. This option allows the NSX agent package to be fetched and installed but the NSX services are not started. Also, no NSX configurations such as certificate generation, are made.
- **Generate image after removing existing NSX agent configurations:** You can remove configurations from an existing NSX-managed VM and use it for generating an image.

## Generating AMI with an unconfigured NSX agent

You can generate an AMI of a VM with the NSX agent installed on it and not configured.

To generate an image from a VM that has the NSX agent installed on it using the **-noStart** option, do the following:

### Procedure

- 1 Copy paste the NSX agent Installation Command from CSM. See instructions at [Install NSX Tools](#)

- a Edit the command for Windows as follows:

```
c:\> powershell -file 'nsx_install.ps1' -operation install -dnsSuffix <> -noStart true
```

- b Edit the command for Linux as follows:

```
$ chmod +x install_nsx_vm_agent.sh && sudo ./install_nsx_vm_agent.sh --no-start
```

- 2 Go to this VM in your public cloud and create an image.

## Generating an Image After Removing Existing NSX Agent Configurations

You can generate an image of a VM that has a configured NSX agent.

To remove configurations from an existing NSX-managed VM and use it for generating images, do the following:

### Procedure

- 1 Removing NSX agent configurations from a Windows or Linux VM:

- a Log in to the workload VM using preferably using a jumphost.
  - b Open the NSX-T CLI:

```
sudo nsxcli
```

- c Enter the following commands:

```
hostname> set debug
hostname> clear nsx-vm-agent state
```

- 2 Locate this VM in your public cloud and create an image.

## Install NSX Tools Automatically

Currently only supported for Microsoft Azure.

In Microsoft Azure, if the following criteria are met, NSX Tools are installed automatically:

- Azure VM Extensions are installed on the VMs in the VNet added into NSX Cloud. See [Microsoft Azure documentation on VM Extensions](#) for more details.

- The security group applied to VMs in Microsoft Azure must allow access to install NSX Tools. If Quarantine Policy is enabled, you can add your VMs to the User Managed list in CSM before installation and remove them from the User Managed list after installation.
- VMs tagged using the key `nsx.network` and value `default`.

To enable this feature:

- 1 Go to **Clouds > Azure > VNets**.
- 2 Select the VNet on whose VMs you want to auto-install CSM.
- 3 Enable the option using any one of the following:
  - In the tile view, click on **ACTIONS > Edit Configuration**.



- If you are in the grid view, select the checkbox next to the VNet and click **ACTIONS > Edit**



- If you are in the VNet tab, click the ACTIONS icon to go to **Edit Configurations**.



- 4 Move the slider next to **Auto-Install NSX Tools** to the ON position.

---

**Note** If NSX Tools installation fails, do the following:

- 1 Log in to the Microsoft Azure portal and navigate to the VM where NSX Tools installation failed.
- 2 Go to the VM's Extensions and uninstall the extension named `VMwareNsxAgentInstallCustomScriptExtension`.
- 3 Remove the `nsx.network=default` tag from this VM.
- 4 Add the `nsx.network=default` tag on this VM again.

Within about three minutes, NSX Tools are installed on this VM.

---

### Install NSX Tools with User Data in AWS

When launching a new workload VM in an AWS VPC, you can install NSX Tools by providing the NSX Tools download and installation instructions in the User Data field.

When you launch an AWS EC2 instance, you have the option of passing `user data` to the instance that can be used to perform common automated configuration tasks, including running scripts after the instance starts. You can pass two types of user data to AWS EC2: shell scripts and cloud-init directives.

Copy the download and installation instructions for NSX Tools from CSM and paste into User Data when launching a new workload VM.

## Prerequisites

Before installing NSX Tools using User Data, ensure that the Transit and Compute VPCs are peered. This is required so that FQDN given in the download command, such as, `nsx-gw.vmware.local` can be resolved from the launching instance.

## Procedure

1 Log in to AWS console and start the process of launching a new workload VM.

2 In another browser window, log in to CSM.

a Go to **Clouds > AWS > VPCs**

---

**Note** Transit VPC/VNet is where one or a pair of PCGs is deployed and running. Compute VPC/VNet is the one linked to a Transit and can use the PCGs deployed there.

---

b Click on a Transit or Compute VPC.

c From the **NSX Tools Download & Installation** section of the screen, copy the **Download Location** and the **Installation Command** under **Linux** or **Windows** depending on what OS you are using for your workload VM. You can also copy-paste the following shell script:

```
#!/bin/bash
sudo wget http://nsx-gw.vmware.local:8080/factory_default/linux/install_nsx_vm_agent.sh
sudo chmod +x install_nsx_vm_agent.sh && sudo ./install_nsx_vm_agent.sh
```

3 In AWS, in the steps for launching a new workload VM instance, paste the download location and the installation command as **Text** in User Data in the Advanced Details section.

## Results

The workload VM is launched and NSX Tools are installed on it automatically.

## Uninstalling NSX Tools

Use these OS-specific commands to uninstall NSX Tools.

### Uninstalling NSX Tools from a Windows VM

---

**Note** To see other options available for the installation script, use `-help`.

---

1 Remote log in to the VM using RDP.

2 Run the installation script with the uninstall option:

```
\nsx_install.ps1 -operation uninstall
```

### Uninstalling NSX Tools from a Linux VM

---

**Note** To see other options available for the installation script, use `--help`.

---

1 Remote log in to the VM using SSH.

2 Run the installation script with the uninstall option:

```
sudo ./install_nsx_vm_agent.sh --uninstall
```

## Security Groups after Onboarding in the NSX Enforced Mode

The following security group configurations take place automatically:

If Quarantine Policy is enabled:

- Healthy NSX-managed VMs are moved to the `vm-underlay-sg` in the public cloud.
- Unmanaged VMs or NSX-managed VMs with errors are moved to the `default` Security Group in AWS and `default-vnet-<vnet-ID>-sg` Network Security Group in Microsoft Azure.
- VMs in the User Managed list are not affected.

If Quarantine Policy is disabled:

- Healthy NSX-managed VMs are moved to the `vm-underlay-sg` in the public cloud.
- NSX-managed VMs with errors are moved to the `default` Security Group in AWS and `default-vnet-<vnet-ID>-sg` Network Security Group in Microsoft Azure.
- Unmanaged VMs and VMs in the User Managed list are not affected.

## Managing VMs in the NSX Enforced Mode

Follow these steps to start managing successfully onboarded VMs in the NSX Enforced Mode.

**Table 23-8. Micro-segmentation workflow for your NSX-managed workload VMs in the NSX Enforced Mode**

| Task                                                                                                                                                                                         | Instructions                                                                                                                                                                   |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  To allow inbound access to workload VMs, create distributed firewall (DFW) rules as required.            | See <a href="#">Default Connectivity Strategy for NSX-Managed Workload VMs in the NSX Enforced Mode</a> .                                                                      |
|  Group your workload VMs using public cloud tags or NSX-T Data Center tags and set up micro-segmentation. | See <a href="#">Set up Micro-segmentation for Workload VMs in the NSX Enforced Mode</a> .<br>See also: <a href="#">Group VMs using NSX-T Data Center and Public Cloud Tags</a> |

## Default Connectivity Strategy for NSX-Managed Workload VMs in the NSX Enforced Mode

When you deploy the PCG on your Transit VPC/VNet or when you link a Compute VPC/VNet to a Transit, NSX Cloud creates default Security Policies and DFW rules therein for NSX-managed workload VMs.

The two stateless rules are for DHCP access and they do not affect access to your workload VMs.

The two stateful rules are as follows:

| DFW Rules created by NSX Cloud under Policy: <code>cloud-stateful-cloud-&lt;VPC/VNet ID&gt;</code> | Properties                                                           |
|----------------------------------------------------------------------------------------------------|----------------------------------------------------------------------|
| <code>cloud-&lt;VPC/VNet ID&gt;-managed</code>                                                     | Allows access to the VMs within the same VPC/VNet.                   |
| <code>cloud-&lt;VPC/VNet ID&gt;-inbound</code>                                                     | Blocks access to NSX-managed VMs from anywhere outside the VPC/VNet. |

**Note** Do not edit any of the default rules.

You can create a copy of the existing inbound rule, adjust the sources and destinations, and set to **Allow**. Place the **Allow** rule above the default **Reject** rule. You can also add new policies and rules. See [Add a Distributed Firewall](#) for instructions.

## Set up Micro-segmentation for Workload VMs in the NSX Enforced Mode

You can set up micro-segmentation for managed workload VMs.

**Note** DFW rules depend on the tags assigned to VMs. Since these tags can be modified by anyone with the appropriate public cloud permissions, NSX-T Data Center assumes that such users are trustworthy and the responsibility of ensuring and auditing that VMs are correctly tagged at all times lies with the public cloud network administrator.

Do the following to apply distributed firewall rules to NSX-managed workload VMs:

- 1 Create groups using VM names or tags or other membership criteria, for example, for **web**, **app**, **DB** tiers. For instructions, see [Add a Group](#).

You can use any of the following tags for membership criteria. See [Group VMs using NSX-T Data Center and Public Cloud Tags](#) for details.

- system-defined tags
- tags from your VPC or VNet that are discovered by NSX Cloud
- or your own custom tags

- 2 Create an East-West distributed firewall policy and rule and apply to the group you created. See [Add a Distributed Firewall](#). You can also use Context Profiles to create rules specific to App IDs and FQDN/URLs. A predefined list of public cloud FQDN/URLs is available when you create an FQDN/URL context profile. See [Layer 7 Context Profile](#) for details.

This micro-segmentation takes effect when the inventory is either manually re-synchronized from CSM, or within about three minutes when the changes are pulled into CSM from your public cloud.

## Native Cloud Enforced Mode

In the Native Cloud Enforced Mode, all your workload VMs are automatically NSX-managed. Follow the workflow outlined here to start managing these VMs using NSX-T Data Center.

**Note** All operating systems are supported for your workload VMs in the Native Cloud Enforced Mode.

### Managing VMs in the Native Cloud Enforced Mode

In the Native Cloud Enforced Mode, NSX Cloud utilizes NSX-T Data Center Groups and Distributed Firewall rules to create corresponding Application Security Groups and Network Security Groups in Microsoft Azure and Security Groups in AWS.

All workload VMs in your VPCs/VNets onboarded in the Native Cloud Enforced Mode are NSX-managed.

Follow this workflow:

**Table 23-9. Micro-segmentation workflow for your workload VMs in the Native Cloud Enforced Mode**

| Task                                                                                                                                                    | Instructions                                                                                                                                                                              |
|---------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <input type="checkbox"/> Create one or more Groups in NSX Manager to include workload VMs from your public cloud.                                       | See <a href="#">Set up Micro-segmentation for Workload VMs in the Native Cloud Enforced Mode</a><br><br>See also: <a href="#">Group VMs using NSX-T Data Center and Public Cloud Tags</a> |
| <input type="checkbox"/> Create one or more Security Policies in NSX Manager that apply to the Group(s) you created for your public cloud workload VMs. |                                                                                                                                                                                           |
| <input type="checkbox"/> Remove workload VMs from the User Managed list in CSM if you want them managed by NSX-T Security Policies.                     |                                                                                                                                                                                           |
| <input type="checkbox"/> Resync your public cloud account in CSM.                                                                                       |                                                                                                                                                                                           |
| <input type="checkbox"/> From your VPC/VNet, switch to the details view in CSM for troubleshooting Security policies if there are any errors.           | See <a href="#">Current Limitations and Common Errors</a>                                                                                                                                 |

### Set up Micro-segmentation for Workload VMs in the Native Cloud Enforced Mode

You can configure Security Policy in NSX Manager for workload VMs in the Native Cloud Enforced Mode.

Starting in NSX-T Data Center 3.0, you can create security policies and rules in VPCs/VNets from different accounts or subscriptions.

---

**Note** DFW rules depend on the tags assigned to VMs. Since these tags can be modified by anyone with the appropriate public cloud permissions, NSX-T Data Center assumes that such users are trustworthy and the responsibility of ensuring and auditing that VMs are correctly tagged at all times lies with the public cloud network administrator.

---

### Prerequisites

Verify that you have a Transit or Compute VPC/VNet in the Native Cloud Enforced Mode.

### Procedure

- 1 In NSX Manager, edit or create Groups for workload VMs, for example, VM names starting with web, app, db, could be three separate Groups. See [Add a Group](#) for instructions. Also see [Group VMs using NSX-T Data Center and Public Cloud Tags](#) for information on using public cloud tags to create Groups for your workload VMs.

Workload VMs that match the criteria are added to the Group. VMs that do not match any grouping criteria are placed in the `default` Security Group in AWS and the `default-vnet-<vnet-ID>-sg` Network Security Group in Microsoft Azure.

---

**Note** You cannot use the Groups that are auto-created by NSX Cloud.

---

- 2 In NSX Manager, create Distributed Firewall (DFW) rules with these Groups in the **Source**, **Destination** or **Applied To** fields. See [Add a Distributed Firewall](#) for instructions.

---

**Note** Only Stateful policies are supported for public cloud workload VMs. Stateless policies can be created in NSX Manager but they will not be matched with any Groups that contain your public cloud workload VMs.

L7 Context Profiles are not supported for DFW rules for workload VMs in the Native Cloud Enforced Mode.

---

- 3 In CSM, remove those VMs from the User Managed list that you want to bring under NSX management. See [How to use the User Managed List](#) for instructions.

---

**Note** Adding VMs to the User Managed list is a manual step that is strongly recommended in the day-0 workflow as soon as you add your public cloud inventory in CSM. If you have not added any VMs to the User Managed list, you do not need to remove them from it.

---



- 4 For Groups and DFW rules that find a match in the public cloud, the following takes place automatically:
  - a In AWS, NSX Cloud creates a new Security Group named like `nsx-<NSX_GUID>`.
  - b In Microsoft Azure, NSX Cloud creates an Application Security Group (ASG) corresponding with the Group created in NSX Manager and a Network Security Group (NSG) corresponding to the DFW rules that are matched with grouped workload VMs.

NSX Cloud synchronizes NSX Manager and public cloud groups and DFW rules every 30 seconds.
- 5 Resynchronize your public cloud account in CSM:
  - a Log in to CSM and go to your public cloud account.
  - b From the public cloud account, click **Actions > Resync Account**. Wait for the resynch to complete.
  - c Go to the VPC/VNet and click on the red-colored **Errors** indicator. This takes you to the instances view.
  - d Switch the view to Details if viewing in Grid and click on **Failed** in the Rules Realization column to view errors, if any.

#### What to do next

See [Current Limitations and Common Errors](#).

### Current Limitations and Common Errors

Refer to these known limitations and common errors to troubleshoot managing your public cloud workload VMs in the Native Cloud Enforced Mode.

---

**Note** The following limits are set by your public cloud:

- The number of security groups that can be applied to a workload VM.
- The number of rules that can be realized for a workload VM.
- The number of rules that can be realized per security group.
- The scope of the security group assignment, for example, the scope of the Network Security Group (NSG) in Microsoft Azure is limited to that region, whereas the scope of the Security Group (SG) in AWS is limited to that VPC.

Refer to the public cloud documentation for more information on these limits.

---

#### Current Limitations

The current release has the following limitations for DFW rules for workload VMs:

- Nested Groups are not supported.
- Groups without VM and/or IP address as member are not supported, for example, Segment or Logical Port based criteria are not supported.

- Both Source and Destination as IP address or CIDR based Group is not supported.
- Both Source and Destination as "ANY" is not supported.
- **Applied\_To** Group can be only Source or Destination or Source + Destination Groups. Other options are not supported.
- Only TCP, UDP, and ICMP are supported.

---

**Note** Only in AWS:

Deny rules created for workload VMs in your AWS VPCs are not realized on AWS because in AWS, everything is in the denied list by default. This leads to the following results in NSX-T Data Center:

- If there is a Deny rule between VM1 and VM2 then traffic is not allowed between VM1 and VM2 because of the default AWS behavior, not because of the Deny rule. The Deny rule is not realized in AWS.
- Assuming the following two rules are created in NSX Manager for the same VMs with rule 1 having a higher priority than rule 2:

- a VM1 to VM2 DENY SSH
- b VM1 to VM2 Allow SSH

the Deny rule is ignored because it is not realized in AWS and therefore the Allow SSH rule is realized. This is contrary to expectation but a limitation because of the default AWS behavior.

---

## Common Errors and their Resolution

### Error: No NSX policy applied to VM.

If you see this error, none of the DFW rules were applied to the particular VM. Edit the rule or the Group in NSX Manager to include this VM.

### Error: Stateless NSX rule is not supported.

If you see this error, it means that you have added DFW rules for public cloud workload VMs in a Stateless Security Policy. This is not supported. Create a new or use an existing Security Policy in the Stateful mode.

## NSX-T Data Center Features Supported with NSX Cloud

NSX Cloud creates a network topology for your public cloud VPC or VNet by generating logical networking entities in NSX-T Data Center.

Use this list as a reference for what is auto-generated and how you should use NSX-T Data Center features as they apply to the public cloud.

## NSX Manager Configurations

See **Auto-created NSX-T Logical Entities** in the *NSX-T Data Center Installation Guide* for details on the logical entities created after a PCG is successfully deployed.

**Important** Do not edit or delete any of these auto-created entities.

**Note** If you are not able to access some features on Windows workload VMs ensure that the Windows firewall settings are correctly configured.

Table 23-10.

| NSX-T Data Center Feature    | Details                                                                                      | NSX Cloud Note                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------------|----------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Segments or Logical Switches | See <a href="#">Chapter 4 Segments</a>                                                       | A segment is created for every public cloud subnet to which a managed VM is attached. This is a hybrid segment.                                                                                                                                                                                                                                                                                 |
| Gateways or Logical Routers  | See <a href="#">Chapter 2 Tier-0 Gateways</a> and <a href="#">Chapter 3 Tier-1 Gateway</a> . | When PCG is deployed on a Transit VPC or VNet, a tier-0 logical router is auto-created by NSX Cloud. A tier-1 router is created for each Compute VPC/VNet when it's linked to a Transit VPC/VNet                                                                                                                                                                                                |
| IPFIX                        | See <a href="#">Configure IPFIX in Manager Mode</a> .                                        | <ul style="list-style-type: none"> <li>■ IPFIX is supported in NSX Cloud only on UDP port 4739.</li> <li>■ <b>Switch and DFW IPFIX:</b> If the collector is in the same subnet as the Windows VM on which IPFIX profile has been applied, a static ARP entry for the collector on the Windows VM is needed because Windows silently discards UDP packets when no ARP entry is found.</li> </ul> |
| Port Mirroring               | See <a href="#">Monitor Port Mirroring Sessions in Manager Mode</a> .                        | <p>Port Mirroring is supported only in AWS in the current release.</p> <ul style="list-style-type: none"> <li>■ For NSX Cloud, configure Port Mirroring from <b>Tools &gt; Port Mirroring Session</b>.</li> <li>■ Only L3SPAN Port Mirroring is supported.</li> <li>■ The collector must be in the same VPC as the source workload VM.</li> </ul>                                               |
| Gateway Firewall             | See <a href="#">Gateway Firewall</a> .                                                       | Only supported on tier-0 gateways.                                                                                                                                                                                                                                                                                                                                                              |

## Group VMs using NSX-T Data Center and Public Cloud Tags

NSX Cloud allows you to use the public cloud tags assigned to your workload VMs.

NSX Manager uses tags to group VMs, as do public clouds. Therefore, to facilitate grouping VMs, NSX Cloud pulls in the public cloud tags applied to your workload VMs provided they meet predefined size and reserved-words criteria, into NSX Manager.

---

**Note** DFW rules depend on the tags assigned to VMs. Since these tags can be modified by anyone with the appropriate public cloud permissions, NSX-T Data Center assumes that such users are trustworthy and the responsibility of ensuring and auditing that VMs are correctly tagged at all times lies with the public cloud network administrator.

---

## Tags terminology

A **tag** in NSX Manager refers to what is known as **value** in a public cloud context. The **key** of a public cloud tag, is referred to as **scope** in NSX Manager.

| Components of tags |                                                   |
|--------------------|---------------------------------------------------|
| in NSX Manager     | Equivalent components of tags in the public cloud |
| Scope              | Key                                               |
| Tag                | Value                                             |

## Tag Types and Limitations

NSX Cloud allows three types of tags for NSX-managed public cloud VMs.

- **System Tags:** These tags are system-defined and you cannot add, edit, or delete them. NSX Cloud uses the following system tags:
  - azure:subscription\_id
  - azure:region
  - azure:vm\_rg
  - azure:vnet\_name
  - azure:vnet\_rg
  - azure:transit\_vnet\_name
  - azure:transit\_vnet\_rg
  - aws:account
  - aws:availabilityzone
  - aws:region
  - aws:vpc
  - aws:subnet
  - aws:transit\_vpc

- **Discovered Tags:** Tags that you have added to your VMs in the public cloud are automatically discovered by NSX Cloud and displayed for your workload VMs in NSX Manager inventory. These tags are not editable from within NSX Manager. There is no limit to the number of discovered tags. These tags are prefixed with `dis:azure:` to denote they are discovered from Microsoft Azure and `dis:aws` from AWS.

When you make any changes to the tags in the public cloud, the changes are reflected in NSX Manager within three minutes.

By default this feature is enabled. You can enable or disable the discovery of Microsoft Azure or AWS tags at the time of adding the Microsoft Azure subscription or AWS account.

- **User Tags:** You can create up to 25 user tags. You have add, edit, delete privileges for user tags. For information on managing user tags, see [Manage Tags for a VM in Manager Mode](#).

Table 23-11. Summary of Tag Types and Limitations

| Tag type       | Tag scope or predetermined prefix                                                                                                                                                                                                                            | Limitations                                                                                                                                                                                                                                                                                                           | Enterprise Administrator Privileges | Auditor Privileges |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------|--------------------|
| System-defined | Complete system tags:<br><ul style="list-style-type: none"> <li>■ azure:subscription_id</li> <li>■ azure:region</li> <li>■ azure:vm_rg</li> <li>■ azure:vnet_name</li> <li>■ azure:vnet_rg</li> <li>■ aws:vpc</li> <li>■ aws:availability zone</li> </ul>    | Scope (key): 20 characters<br>Tag (value): 65 characters<br>Maximum possible: 5                                                                                                                                                                                                                                       | Read only                           | Read only          |
| Discovered     | Prefix for Microsoft Azure tags that are imported from your VNet:<br><b>dis:azure:</b><br>Prefix for AWS tags that are imported from your VPC:<br><b>dis:aws:</b>                                                                                            | Scope (key): 20 characters<br>Tag (value): 65 characters<br>Maximum allowed: unlimited<br><hr/> <b>Note</b> The limits on characters excludes the prefix <b>dis:&lt;public cloud name&gt;</b> . Tags that exceed these limits are not reflected in NSX Manager.<br><hr/> Tags with the prefix <b>nsx</b> are ignored. | Read only                           | Read only          |
| User           | User tags can have any scope (key) and value within the allowed number of characters, except: <ul style="list-style-type: none"> <li>■ the scope (key) prefix <b>dis:azure:</b> or <b>dis:aws:</b></li> <li>■ the same scope (key) as system tags</li> </ul> | Scope (key): 30 characters<br>Tag (value): 65 characters<br>Maximum allowed: 25                                                                                                                                                                                                                                       | Add/Edit/Delete                     | Read only          |

## Examples of Discovered Tags

**Note** Tags are in the format **key=value** for the public cloud and **scope=tag** in NSX Manager.

Table 23-12.

| Public Cloud tag for the workload VM                                           | Discovered by NSX Cloud?            | Equivalent NSX Manager tag for the workload VM  |
|--------------------------------------------------------------------------------|-------------------------------------|-------------------------------------------------|
| Name=Developer                                                                 | Yes                                 | dis:azure:Name=Developer                        |
| ValidDisTagKeyLength=ValidDisTagValue                                          | Yes                                 | dis:azure:ValidDisTagKeyLength=ValidDisTagValue |
| Abcdefghijklmnopqrstuvwxyz=value2                                              | No (key exceeds 20 chars)           | none                                            |
| tag3=AbcdefghijklmnopqrstuvwxyzAb2369Ohgjjuytreswqacvbcdefghijklmnopqrstuvwxyz | No (value exceeds 65 characters)    | none                                            |
| nsx.name=Tester                                                                | No (key has the prefix <b>nsx</b> ) | none                                            |

## How to use Tags in NSX Manager

- See [Manage Tags for a VM in Manager Mode](#).
- See [Search for Objects](#).
- See [Add a Group](#).
- See [Set up Micro-segmentation for Workload VMs in the NSX Enforced Mode](#).

## Use Native-Cloud Services

The following native-cloud services are supported for use with your public cloud workload VMs from within NSX Manager.

When you deploy PCG, a Group is created in NSX Manager for each supported native-cloud service.

The following Groups are created for the currently supported public cloud services:

- aws-dynamo-db-service-endpoint
- aws-elb-service-endpoint
- aws-rds-service-endpoint
- aws-s3-service-endpoint
- azure-cosmos-db-service-endpoint
- azure-load-balancer-service-endpoint
- azure-sql-service-endpoint
- azure-storage-service-endpoint

To use these native-cloud services, create DFW policies that contain the native-cloud service Group in the Source or Destination fields of the rule as required.

DFW rules are enforced on VMs not on the native-cloud services.

**Note** In the NSX Enforced Mode, that is, managing your workloads with NSX Tools, currently there is no support for Microsoft Azure's native-cloud services.

## Current Limitations

| ENDPOINT        |               |           | DFW Rule with service as DESTINATION |                      | DFW Rule with service as SOURCE |                 |
|-----------------|---------------|-----------|--------------------------------------|----------------------|---------------------------------|-----------------|
| Public Cloud    | Service       | Scope     | Enforced on VM?                      | Enforced on Service? | Enforced on Service?            | Enforced on VM? |
| Microsoft Azure | BLOB Storage  | Global    | Yes                                  | No                   | No                              | Yes             |
|                 | Cosmos DB     |           |                                      |                      |                                 |                 |
|                 | SQL           |           |                                      |                      |                                 |                 |
|                 | Load Balancer |           |                                      |                      |                                 |                 |
| AWS             | S3            | VPC Local | Yes                                  | No                   | No                              | Yes             |
|                 | Dynamo DB     |           |                                      |                      |                                 |                 |
|                 | RDS           |           |                                      |                      |                                 |                 |
|                 | ELB           |           |                                      |                      |                                 |                 |

## Service Insertion for your Workload VMs in the NSX Enforced Mode

NSX Cloud supports the use of third-party services in your public cloud for NSX-managed workload VMs in the NSX Enforced Mode.

NSX Cloud supports Service Insertion for the following:

- North-south traffic from workload VMs via a service appliance hosted in a Transit VPC/VNet..
- VPN traffic from the PCG to an on-prem edge or gateway. This traffic can be routed via a service appliance in a Transit VPC/VNet as well.

Here is an overview of the configurations to allow service insertion for your NSX-managed workload VMs.



**Table 23-13. Overview of configurations required for service insertion for NSX-managed workload VMs in the NSX Enforced Mode.**

| Frequency                                                                                                        | Task                                                                                                                                                                                                                                                                                                 | Instructions                                                                           |
|------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------|
| Follow these instructions for the initial setup if you want to set up service insertion for north-south traffic. | Set up the service appliance in your public cloud preferably in a Transit VPC or VNet (where you have deployed the PCG).                                                                                                                                                                             | See instructions specific to the third-party service appliance and the public cloud.   |
|                                                                                                                  | Register the third-party service in NSX-T Data Center.                                                                                                                                                                                                                                               | See <a href="#">Create the Service Definition and a Corresponding Virtual Endpoint</a> |
|                                                                                                                  | Create a virtual instance endpoint of the service using a /32 Virtual Service IP address (VSIP) to be used only for service insertion by the service appliance. The VSIP should not conflict with the CIDR range of VPCs or VNets. This VSIP is advertised over BGP to the PCG.                      | See <a href="#">Create the Service Definition and a Corresponding Virtual Endpoint</a> |
|                                                                                                                  | Create an IPSec VPN tunnel between the service appliance and the PCG.                                                                                                                                                                                                                                | See <a href="#">Set up an IPSec VPN Session</a>                                        |
|                                                                                                                  | Configure BGP between the PCG and the service appliance and advertise the VSIP from the service appliance and the default route (0.0.0.0/0) from the PCG.                                                                                                                                            | See <a href="#">Configure BGP and Route Redistribution</a>                             |
| Follow these instructions for the initial setup for VPN traffic from the public cloud to on-prem.                | Create a VPN tunnel between the PCG and the on-prem edge or gateway.                                                                                                                                                                                                                                 | See <a href="#">Set up VPN in the NSX Enforced Mode.</a>                               |
| Follow these instructions for both types of service insertion as part of the initial setup.                      | Create a lowest priority default catch-all rule with the action set to <b>Do Not Redirect</b> . This ensures that no packets are redirected on the VTI interface of the PCG and the Service Appliance.                                                                                               | See <a href="#">Set up Redirection Rules.</a>                                          |
| Follow these instructions as and when necessary for each type of service insertion use case.                     | After the one-time configurations are complete, set up redirection rules to reroute selective traffic from NSX-managed workload VMs to the VSIP. These rules are applied to the uplink port of the PCG for north-south service insertion and to the VTI interface of the PCG for traffic to on-prem. | See <a href="#">Set up Redirection Rules.</a>                                          |

## What to read next

### Procedure

#### 1 Create the Service Definition and a Corresponding Virtual Endpoint

You must use NSX Manager APIs to create a service definition and virtual endpoint for the service appliance in your public cloud.

#### 2 Set up an IPSec VPN Session

Set up an IPSec VPN session between the PCG and your service appliance.

#### 3 Configure BGP and Route Redistribution

Configure BGP between the PCG and the service appliance over the IPSec VPN tunnel.

#### 4 Set up Redirection Rules

You must set up a default redirection rule as part of the initial setup for service insertion.

## Create the Service Definition and a Corresponding Virtual Endpoint

You must use NSX Manager APIs to create a service definition and virtual endpoint for the service appliance in your public cloud.

### Prerequisites

Pick out a /32 reserved IP address to serve as the Virtual Endpoint for the service appliance in your public cloud, for example, 100.100.100.100/32. This is referred to as the Virtual Service IP (VSIP).

---

**Note** If you deployed your service appliance in a High Availability pair, do not create another service definition but use the same VSIP when advertising it to the PCG during BGP configuration.

---

### Procedure

- 1 To create a Service Definition for the service appliance, run the following API call using NSX Manager credentials for authorization:

```
POST https://{NSX Manager-IP}/policy/api/v1/enforcement-points/default/service-definitions
```

Example request:

```
{
 "resource_type": "ServiceDefinition",
 "description": "NS-Service",
 "display_name": "Service_Appliance1",
 "attachment_point": [
 "TIER0_LR"
],
 "transports": [
 "L3_ROUTED"
],
 "functionalities": [
```

```

 "NG_FW", "BYOD"
],
 "on_failure_policy": "ALLOW",
 "implementations": [
 "NORTH_SOUTH"
],
 "vendor_id" : "Vendor1"
}

```

Example response:

```

{
 "resource_type": "ServiceDefinition",
 "description": "NS-Service",
 "id": "33890153-6eea-4c9d-8e34-7b6532b9d65c",
 "display_name": "Service_Appliance1",
 "attachment_point": [
 "TIER0_LR"
],
 "transports": [
 "L3_ROUTED"
],
 "functionalities": [
 "NG_FW", "BYOD"
],
 "vendor_id": "Vendor1",
 "on_failure_policy": "ALLOW",
 "implementations": [
 "NORTH_SOUTH"
],
 "_create_time": 1540424262137,
 "_last_modified_user": "nsx_policy",
 "_system_owned": false,
 "_protection": "REQUIRE_OVERRIDE",
 "_last_modified_time": 1540424262137,
 "_create_user": "nsx_policy",
 "_revision": 0
}

```

- 2 To create a Virtual Endpoint for the service appliance, run the following API call using NSX Manager credentials for authorization:

```

PATCH https://{NSX Manager-IP}/policy/api/v1/infra/tier-0s/<tier-0 router ID>/locale-services/cloud/endpoints/virtual-endpoints/Service_Appliance1_Endpoint

```

Example request:

```

{
 "resource_type": "VirtualEndpoint",
 "display_name": "Service_Appliance1_Endpoint",
 "target_ips": [
 {
 "ip_addresses": [
 "100.100.100.100"
]
 }
]
}

```

```

],
 "prefix_length": 32
 }
],
"service_names": [
 "Service_Appliance1"
]
}

```

Example response:

```
200 OK
```

**Note** The `display_name` in step 1 must match the `service_names` in step 2.

What to do next

[Set up an IPSec VPN Session](#)

## Set up an IPSec VPN Session

Set up an IPSec VPN session between the PCG and your service appliance.

### Prerequisites

- One or an HA pair of PCGs must be deployed in a Transit VPC/VNet.
- The service appliance must be set up in your public cloud, preferably in the Transit VPC/VNet.

### Procedure

- 1 Navigate to **Networking > VPN**
- 2 Add a **VPN service** of type IPSec and note the following configuration options specific to NSX Cloud. See [Add an IPSec VPN Service](#) for other details.

| Option                | Description                                                                                                                                             |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Name</b>           | The name of this VPN service is used to set up the local endpoint and the IPSec VPN sessions. Make a note of it.                                        |
| <b>Service Type</b>   | Confirm that this value is set to IPSec.                                                                                                                |
| <b>Tier-0 Gateway</b> | Select the tier-0 gateway auto-created for your Transit VPC/VNet. Its name contains your VPC/VNet ID, for example, <code>cloud-t0-vpc-6bcd2c13</code> . |

- 3 Add a **Local Endpoint** for your PCG. The IP address of the local endpoint is the value of the tag `nsx:local_endpoint_ip` for the PCG deployed in your Transit VPC/VNet. Log in to your Transit VPC/VNet for this value. Note the following configurations specific to NSX Cloud and see [Add Local Endpoints](#) for other details.

| Option      | Description                                                                                                                                                                               |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name        | The local endpoint name is used to set up the IPSec VPN sessions. Make a note of it.                                                                                                      |
| VPN Service | Select the VPN Service you added in step 2.                                                                                                                                               |
| IP Address  | Find this value by logging in to the AWS console or the Microsoft Azure portal. It is the value of the tag <code>nsx:local_endpoint_ip</code> applied to the uplink interface of the PCG. |

- 4 Create a **Route-Based IPSec session** between the PCG and the service appliance in your public cloud (preferably hosted in the Transit VPC/VNet).

| Option           | Description                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Type             | Confirm that this value is set to <b>Route Based</b> .                                                                                                                                                                                                                                                                                                                                                                 |
| VPN Service      | Select the VPN Service you added in step 2.                                                                                                                                                                                                                                                                                                                                                                            |
| Local Endpoint   | Select the local endpoint you created in step 3.                                                                                                                                                                                                                                                                                                                                                                       |
| Remote IP        | Enter the private IP address of the service appliance.<br><br><b>Note</b> If your service appliance is accessible using a public IP address, assign a public IP address to the local endpoint IP (also known as secondary IP) to the PCG's uplink interface.                                                                                                                                                           |
| Tunnel Interface | This subnet must match with the service appliance subnet for the VPN tunnel. Enter the subnet value you set up in the service appliance for the VPN tunnel or note the value you enter here and make sure the same subnet is used when setting up the VPN tunnel in the service appliance.<br><br><b>Note</b> You configure BGP on this tunnel interface. See <a href="#">Configure BGP and Route Redistribution</a> . |
| Remote ID        | Enter the private IP address of your service appliance in the public cloud.                                                                                                                                                                                                                                                                                                                                            |
| IKE Profile      | The IPSec VPN session must be associated with an IKE profile. If you created a profile, select it from the drop-down menu. You can also use the default profile.                                                                                                                                                                                                                                                       |

#### What to do next

#### [Configure BGP and Route Redistribution](#)

### Configure BGP and Route Redistribution

Configure BGP between the PCG and the service appliance over the IPSec VPN tunnel.

You set up BGP neighbors on the IPSec VPN tunnel interface that you established between PCG and the service appliance. See [Configure BGP](#) for more details.

You need to configure BGP similarly on your service appliance. See documentation for your specific service in the public cloud for details.

Next, set up route redistribution as follows:

- The PCG advertises its default route (0.0.0.0/0) to the service appliance.
- The service appliance advertises the VSIP to the PCG. This is the same IP address which is used when registering the service. See [Create the Service Definition and a Corresponding Virtual Endpoint](#).

---

**Note** If your service appliance is deployed in a High Availability pair, advertise the same VSIP from both service appliances.

---

#### Procedure

- 1 Navigate to **Networking > Tier-0 Gateways**.
- 2 Select the auto-created tier-0 gateway for your Transit VPC/VNet named like `cloud-t0-vpc-6bcd2c13` and click **Edit**.
- 3 Click the number or icon next to **BGP Neighbors** under the **BGP** section.
- 4 Note these configurations:

| Option           | Description                                                                                                                    |
|------------------|--------------------------------------------------------------------------------------------------------------------------------|
| IP Address       | Use the IP address configured on the service appliance tunnel interface for the VPN between the PCG and the service appliance. |
| Remote AS Number | This number must match the AS number of the service appliance in your public cloud.                                            |
| Route Filter     | Set an Out Filter to advertise the default route (0.0.0.0/0) from the PCG to service appliance.                                |

- 5 From the **Route Redistribution** section, enable static routes on tier-0 gateway.

## Set Route Re-distribution

Tier-0 Gateways cloud-t0-415... #Route Re-distribution 3

ADD ROUTE RE-DISTRIBUTION Search

| Name | Route Re-distribution | Route Map |
|------|-----------------------|-----------|
|      | Set *                 |           |

### Set Route Re-distribution

Tier-0 Gateways cloud-t0-415... #Selected Sources 1

Select sources below

**Tier-0 Subnets**

☒ Static Routes
 ☐ NAT IP

☐ IPsec Local IP
 ☐ DNS Forwarder IP

☐ EVPN TEP IP

☐ Connected Interfaces & Segments

☐ Service Interface Subnet
 ☐ External Interface Subnet

☐ Loopback Interface Subnet
 ☐ Connected Segment

What to do next

[Set up Redirection Rules](#)

## Set up Redirection Rules

You must set up a default redirection rule as part of the initial setup for service insertion.

After the initial setup is completed, you can create and edit redirection rules as necessary for rerouting different types of traffic for your NSX-managed workload VMs through the service appliance.

These are the two types of redirection rules:

- 1 As part of the initial service insertion setup, you must create a catch-all rule to prevent redirection for traffic for the VTI interface of the VPN tunnel between the PCG and the service appliance. This rule must have the lowest possible priority and must be created for both use cases of service insertion.
- 2 The second rule sets up specific redirection for traffic for the service appliance. You can adjust this rule and add others as necessary.

## Procedure

1 To add the default catch-all rule to complete the one-time setup, follow these steps:

- a Navigate to **Security > North South Firewall > Network Introspection (N-S)**
- b Click **Add Policy**.

| Option              | Description                                                                                                  |
|---------------------|--------------------------------------------------------------------------------------------------------------|
| <b>Name</b>         | Provide a descriptive name, for example, <b>Default_No-Redirect-Policy</b> .                                 |
| <b>Redirect To:</b> | Select the name of the Virtual Endpoint you created for this service appliance when registering the service. |
| <b>Apply To:</b>    | Select the PCG's tier-0 gateway.                                                                             |

- c Select the new policy and click **Add Rule**. Note the following values specific to service insertion:

| Option              | Description                                                         |
|---------------------|---------------------------------------------------------------------|
| <b>Sources</b>      | Any                                                                 |
| <b>Destinations</b> | Any                                                                 |
| <b>Applied To</b>   | Select the VTI interface between the PCG and the service appliance. |
| <b>Action</b>       | Select <b>Do Not Redirect</b> .                                     |

**Important** This rule must have the lowest possible priority.



2 For the second rule, follow these steps:

- a Navigate to **Security > North South Firewall > Network Introspection (N-S)**
- b Click **Add Policy**.

| Option              | Description                                                                                                                                                  |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Name:</b>        | Provide a descriptive name for the policy, for example, <b>On-Prem Service Insertion for AWS VMs</b> or <b>North-south Service Insertion for Azure VMs</b> . |
| <b>Redirect To:</b> | Select the name of the Virtual Endpoint you created for this service appliance when registering the service.                                                 |
| <b>Apply To:</b>    | Select the PCG's tier-0 gateway.                                                                                                                             |

- c Select the new policy and click **Add Rule**. Note the following values specific to service insertion:

| Option              | Description                                                                                                                                                                                                                                                                                                                                           |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Sources</b>      | Select a group of subnets whose traffic must be redirected, for example, a group of your NSX-managed workload VMs.                                                                                                                                                                                                                                    |
| <b>Destinations</b> | Select a list of destination IP addresses or services, such as <b>YouTube</b> , that you want to route through the service appliance.                                                                                                                                                                                                                 |
| <b>Applied To</b>   | <ul style="list-style-type: none"> <li>■ If you are using north-south service insertion with the service appliance in the public cloud: select the uplink port of the active and standby PCG.</li> <li>■ If you are using VPN traffic to on-prem: select the VTI interface of the active and standby PCG to the on-prem service appliance.</li> </ul> |
| <b>Action</b>       | Select <b>Redirect</b> .                                                                                                                                                                                                                                                                                                                              |

## Enable NAT on NSX-managed VMs

NSX Cloud supports enabling NAT on NSX-managed VMs.

You can enable North-South traffic on VMs in NSX-managed VMs using public cloud tags.

On the NSX-managed VM for which you want to enable NAT, apply the following tag:

**Table 23-14.**

| Key                 | Value                                                                   |
|---------------------|-------------------------------------------------------------------------|
| <b>nsx.publicip</b> | <b>public IP address from your public cloud</b> , for example, 50.1.2.3 |

**Note** The public IP address you provide here must be free to use and must not be assigned to any VM, even the workload VM you want to enable NAT for. If you assign a public IP address that was previously associated with any other instance or private IP address, NAT does not work. In that case, unassign the public IP address.

After this tag is applied, the workload VM can access internet traffic.

## Enable Syslog Forwarding

NSX Cloud supports syslog forwarding.

You can enable syslog forwarding for Distributed Firewall (DFW) packets on managed VMs. See **Configure Remote Logging** in the *NSX-T Data Center Troubleshooting Guide* for further details.

Do the following:

### Procedure

- 1 Log in to PCG using the jump host.
- 2 Type **nsxcli** to open the NSX-T Data Center CLI.
- 3 Type the following commands to enable DFW log forwarding:

```
nsx-public-cloud-gateway> set gw-controller vm-log-forwarding enabled
nsx-public-cloud-gateway> set logging-server <server-IP-address> proto udp level info
messageid FIREWALL-PKTLOG
```

After this is set, NSX agent DFW packet logs are available under `/var/log/syslog` on PCG.

- 4 To enable log forwarding per VM, enter the following command:

```
nsx-public-cloud-gateway> set gw-controller vm-log-forwarding enabled <vm-id>
```

## Set up VPN in the Native Cloud Enforced Mode

You can create a VPN tunnel between the PCG and a remote endpoint by following this workflow. These instructions are specific to workload VMs managed in the Native Cloud Enforced Mode.

### Prerequisites

- In AWS: Verify that you have deployed a VPC in the Native Cloud Enforced Mode. This must be a Transit or Self-managed VPC. VPN is not supported for Compute VPCs in AWS.
- In Microsoft Azure: Verify that you have deployed a VNet in the Native Cloud Enforced Mode. You can use both Transit and Compute VNets.
- Verify that the remote endpoint is peered with the PCG and has route-based IPsec VPN and BGP capabilities.

### Procedure

- 1 In your public cloud, find the NSX-assigned local endpoint for the PCG and assign a public IP address to if necessary:
  - a Go to your PCG instance in the public cloud and navigate to Tags.
  - b Note the IP address in the value field of the tag `nsx.local_endpoint_ip`.

- c (Optional) If your VPN tunnel requires a public IP, for example, if you want to set up a VPN to another public cloud or to the on-prem NSX-T Data Center deployment:
  - 1 Navigate to the uplink interface of the PCG instance.
  - 2 Attach a public IP address to the `nsx.local_endpoint_ip` IP address that you noted in step **b**.
- d (Optional) If you have an HA pair of PCG instances, repeat steps **a** and **b** and attach a public IP address if necessary, as described in step **c**.

- 2 In NSX Manager, enable IPsec VPN for the PCG that appears as a tier-0 gateway named like `cloud-t0-vpc/vnet-<vpc/vnet-id>` and create route-base IPsec sessions between this tier-0 gateway's endpoint and the remote IP address of the desired VPN peer. See [Add an IPsec VPN Service](#) for other details.

- a Go to **Networking > VPN > VPN Services > Add Service > IPsec**. Provide the following details:

| Option              | Description                                                                                                                               |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| Name                | Enter a descriptive name for the VPN service, for example <code>&lt;VPC-ID&gt;-AWS_VPN</code> or <code>&lt;VNet-ID&gt;-AZURE_VPN</code> . |
| Tier0/Tier1 Gateway | Select the tier-0 gateway for the PCG in your public cloud.                                                                               |

- b Go to **Networking > VPN > Local Endpoints > Add Local Endpoint**. Provide the following information and see [Add Local Endpoints](#) for other details. :

**Note** If you have an HA pair of PCG instances, create a local endpoint for each instance using the corresponding local endpoint IP address attached to it in the public cloud.

| Option      | Description                                                                                                                                                |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name        | Enter a descriptive name for the local endpoint, for example <code>&lt;VPC-ID&gt;-PCG-preferred-LE</code> or <code>&lt;VNET-ID&gt;-PCG-preferred-LE</code> |
| VPN Service | Select the VPN service for the PCG's tier-0 gateway that you created in step 2a.                                                                           |
| IP Address  | Enter the value of the PCG's local endpoint IP address that you noted in step 1b.                                                                          |

- c Go to **Networking > VPN > IPsec Sessions > Add IPsec Session > Route Based**. Provide the following information and see [Add a Route-Based IPsec Session](#) for other details:

**Note** If you are creating a VPN tunnel between PCGs deployed in a VPC and PCGs deployed in a VNet, you must create a tunnel for each PCG's local endpoint in the VPC and the remote IP address of the PCG in the VNet, and conversely from the PCGs in the VNet to the remote IP address of PCGs in the VPC. You must create a separate tunnel for the active and standby PCGs. This results in a full mesh of IPsec Sessions between the two public clouds.

| Option         | Description                                                                                                   |
|----------------|---------------------------------------------------------------------------------------------------------------|
| Name           | Enter a descriptive name for the IPsec session, for example, <code>&lt;VPC--ID&gt;-PCG1-to-remote_edge</code> |
| VPN Service    | Select the VPN service you created in step 2a.                                                                |
| Local Endpoint | Select the local endpoint you created in step 2b.                                                             |
| Remote IP      | Enter the public IP address of the remote peer with which you are creating the VPN tunnel.                    |

| Option                  | Description                                                                                                                                          |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
|                         | <b>Note</b> Remote IP can be a private IP address if you are able to reach the private IP address, for example, using DirectConnect or ExpressRoute. |
| <b>Tunnel Interface</b> | Enter the tunnel interface in a CIDR format. The same subnet must be used for the remote peer to establish the IPSec session.                        |

**Step 2a.**

**Step 2b.**

**Step 2c.**

| Name                         | Type        | VPN Service      | Local Endpoint            | Remote IP    |
|------------------------------|-------------|------------------|---------------------------|--------------|
| <VPC-ID>-PCG1-to-remote_edge | Route Based | <VPC-ID>-AWS_VPN | <VPC-ID>-PCG-preferred-LE | 3.213.92.220 |

IP address of VPN Peer

| Name                         | Type        | VPN Service      | Local Endpoint            | Remote IP    |
|------------------------------|-------------|------------------|---------------------------|--------------|
| <VPC-ID>-PCG1-to-remote_edge | Route Based | <VPC-ID>-AWS_VPN | <VPC-ID>-PCG-preferred-LE | 3.213.92.220 |

IP address of VPN Peer

| Name                         | Type        | VPN Service      | Local Endpoint            | Remote IP    |
|------------------------------|-------------|------------------|---------------------------|--------------|
| <VPC-ID>-PCG1-to-remote_edge | Route Based | <VPC-ID>-AWS_VPN | <VPC-ID>-PCG-preferred-LE | 3.213.92.220 |

IP address of VPN Peer

- 3 Set up BGP neighbors on the IPsec VPN tunnel interface that you established in step 2. See [Configure BGP](#) for more details.
  - a Navigate to **Networking > Tier-0 Gateways**
  - b Select the auto-created tier-0 gateway for which you created the IPsec session and click **Edit**.
  - c Click the number or icon next to **BGP Neighbors** under the **BGP** section and provide the following details:

| Option           | Description                                                                                                    |
|------------------|----------------------------------------------------------------------------------------------------------------|
| IP Address       | Use the IP address of the remote VTI configured on the tunnel interface in the IPsec session for the VPN peer. |
| Remote AS Number | This number must match the AS number of the remote peer.                                                       |

**Tier-0 Gateways**

ADD GATEWAY EXPAND ALL Filter by Name

| Tier-0 Gateway Name    | HA Mode     | Linked Tier-1 Gateways       | Linked Segments |
|------------------------|-------------|------------------------------|-----------------|
| Local AS               | 1000        |                              |                 |
| BGP                    | On          | ECMP                         | On              |
| Graceful Restart       | Helper Only | Multipath Relax              | On              |
| Graceful Restart Timer | 180 seconds | Graceful Restart Stale Timer | 600 seconds     |
| Route Aggregation      | 0           | BGP Neighbors                |                 |

**Step 3.**

**BGP Neighbors**

Tier-0 Gateways cloud-t0-vpc... #Neighbors

| IP Address    | BFD      | Remote AS number |
|---------------|----------|------------------|
| 192.168.50.11 | Disabled | 1000             |

Source Addresses Not Set

Max Hop Limit 1

- 4 **Important** This step is only for NSX-T Data Center 3.0.0. Skip it if your are using NSX-T Data Center 3.0.1.

If you are using Microsoft Azure, after you have configured VPN and BGP in NSX Manager, **Enable IP Forwarding** on the uplink interface of the PCG instance. If you have an active and a standby PCG instance, for HA, then enable IP Forwarding on both PCG instances.



- 5 Advertise the prefixes you want to use for the VPN using the Redistribution Profile. Do the following:

- a

---

**Important** This step is only for NSX-T Data Center 3.0.0. Skip it if you are using NSX-T Data Center 3.0.1.

---

Add a static route for the CIDR of the VPC/VNet onboarded with the Native Cloud Enforced Mode to point to the uplink IP address of the tier-0 gateway, that is, PCG. See [Configure a Static Route](#) for instructions. If you have a PCG pair for HA, set up next hops to each PCG's uplink IP address.

## Static Routes

**Step 5a.**

Tier-0 Gateways cloud-t0-vpc... #Static Routes 4

Search

| Name                                                  | Network        | Next Hops | Status  |
|-------------------------------------------------------|----------------|-----------|---------|
| <VPC/VPN-ID>-<br>NativeCloudMode_CIDR_<br>StaticRoute | 14.1.14.128/28 | 2         | Success |

## Next Hops

Tier-0 Gateways cloud-t0-vpc... | Static Route cloud-native... #Next Hops 2

Search

| IP Address  | Admin Distance | Scope                                       |
|-------------|----------------|---------------------------------------------|
| 13.1.13.129 | 30             | preferred-uplink-port-vpc-0942e6190176e19c1 |
| 13.1.13.145 | 30             | uplink-port-vpc-0942e6190176e19c1           |

- b Add a prefix list for the VPC/VNet CIDR onboarded in the Native Cloud Enforced Mode and add it as an Out Filter in BGP neighbor configuration. See [Create an IP Prefix List](#) for instructions.
- c Set up a route redistribution profile, enabling static route and selecting the route filter you created for VPC/VNet CIDRs in step b.

## Set Route Re-distribution

**Step 5c.**

Tier-0 Gateways cloud-t0-415... #Route Re-distribution 3

ADD ROUTE RE-DISTRIBUTION

Search

| Name                                  | Route Re-distribution | Route Map                            |
|---------------------------------------|-----------------------|--------------------------------------|
| <VPC/VPN-ID>-<br>NativeCloudMode_CIDR | Set *                 | Select Route Map                     |
| ADD                                   | CANCEL                | Select Route Map created in Step 5b. |

## Set Route Re-distribution

Tier-0 Gateways cloud-t0-415... #Selected Sources 1

Select sources below

Tier-0 Subnets

- ☒ Static Routes
- ☐ IPSec Local IP
- ☐ EVPN TEP IP
- ☐ Connected Interfaces & Segments
- ☐ Service Interface Subnet
- ☐ Loopback Interface Subnet
- ☐ NAT IP
- ☐ DNS Forwarder IP
- ☐ External Interface Subnet
- ☐ Connected Segment

## 6 In your public cloud:

- a Go to the routing table of the subnet where you have your workload VMs.

---

**Note** Do not use the routing table of the PCG's uplink or management subnets.

---

- b Add the tag `nsx.managed = true` to the routing table.

## 7

---

**Important** This step is only for NSX-T Data Center 3.0.0. Skip it if you are using NSX-T Data Center 3.0.1.

---

NSX Cloud creates a **default-snat** rule for the tier-0 gateway (PCG) with source `0.0.0.0/0` and destination `Any`. Because of this rule, all traffic from your VMs in the Native Cloud Enforced Mode has the PCG's uplink IP address. If you want to see the true source of your traffic, do the following:

- a Go to **Networking > NAT** and disable the **default-snat** rule for the tier-0 gateway (PCG).
- b Create a new SNAT rule with the following values if you have VMs in the NSX Enforced Mode to continue providing SNAT for such VMs:

| Option      | Description                                                                       |
|-------------|-----------------------------------------------------------------------------------|
| Source      | CIDR of the VPC/VNet in the NSX Enforced Mode .                                   |
| Destination | Any                                                                               |
| Translated  | The same IP address in <b>Translated</b> that is in the <b>default-snat</b> rule. |
| Apply To    | Select the PCG's uplink interface.                                                |

Do not edit the **default-snat** rule. It is reverted in case of a failover.

## Results

Verify that routes are created in the managed routing table for all IP prefixes advertised by the remote endpoint with next hop set to the PCG's uplink IP address.

## Set up VPN in the NSX Enforced Mode

You can set up VPN using PCGs that appear as auto-created tier-0 gateways in the on-prem NSX-T Data Center deployment. These instructions are specific to workload VMs managed in the NSX Enforced Mode.

Use PCGs in the same way as you use tier-0 gateways in NSX Manager to set up VPN by following the additional steps outlined here. You can create VPN tunnels between PCGs deployed in the same public cloud, or different public clouds, or with an on-prem gateway or router. See [Chapter 6 Virtual Private Network \(VPN\)](#) for details on VPN support in NSX-T Data Center.

## Prerequisites

- Verify that you have one or an HA pair of PCGs deployed in a VPC/VNet.

- Verify that the remote peer supports route-based VPN and BGP.

#### Procedure

- 1 In your public cloud, find the NSX-assigned local endpoint for the PCG and assign a public IP address to if necessary:
  - a Go to your PCG instance in the public cloud and navigate to Tags.
  - b Note the IP address in the value field of the tag `nsx.local_endpoint_ip`.
  - c (Optional) If your VPN tunnel requires a public IP, for example, if you want to set up a VPN to another public cloud or to the on-prem NSX-T Data Center deployment:
    - 1 Navigate to the uplink interface of the PCG instance.
    - 2 Attach a public IP address to the `nsx.local_endpoint_ip` IP address that you noted in step **b**.
  - d (Optional) If you have an HA pair of PCG instances, repeat steps **a** and **b** and attach a public IP address if necessary, as described in step **c**.

- 2 In NSX Manager, enable IPsec VPN for the PCG that appears as a tier-0 gateway named like `cloud-t0-vpc/vnet-<vpc/vnet-id>` and create route-base IPsec sessions between this tier-0 gateway's endpoint and the remote IP address of the desired VPN peer. See [Add an IPsec VPN Service](#) for other details.

- a Go to **Networking > VPN > VPN Services > Add Service > IPsec**. Provide the following details:

| Option              | Description                                                                                                                               |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| Name                | Enter a descriptive name for the VPN service, for example <code>&lt;VPC-ID&gt;-AWS_VPN</code> or <code>&lt;VNet-ID&gt;-AZURE_VPN</code> . |
| Tier0/Tier1 Gateway | Select the tier-0 gateway for the PCG in your public cloud.                                                                               |

- b Go to **Networking > VPN > Local Endpoints > Add Local Endpoint**. Provide the following information and see [Add Local Endpoints](#) for other details. :

**Note** If you have an HA pair of PCG instances, create a local endpoint for each instance using the corresponding local endpoint IP address attached to it in the public cloud.

| Option      | Description                                                                                                                                                |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name        | Enter a descriptive name for the local endpoint, for example <code>&lt;VPC-ID&gt;-PCG-preferred-LE</code> or <code>&lt;VNET-ID&gt;-PCG-preferred-LE</code> |
| VPN Service | Select the VPN service for the PCG's tier-0 gateway that you created in step 2a.                                                                           |
| IP Address  | Enter the value of the PCG's local endpoint IP address that you noted in step 1b.                                                                          |

- c Go to **Networking > VPN > IPsec Sessions > Add IPsec Session > Route Based**. Provide the following information and see [Add a Route-Based IPsec Session](#) for other details:

**Note** If you are creating a VPN tunnel between PCGs deployed in a VPC and PCGs deployed in a VNet, you must create a tunnel for each PCG's local endpoint in the VPC and the remote IP address of the PCG in the VNet, and conversely from the PCGs in the VNet to the remote IP address of PCGs in the VPC. You must create a separate tunnel for the active and standby PCGs. This results in a full mesh of IPsec Sessions between the two public clouds.

| Option         | Description                                                                                                   |
|----------------|---------------------------------------------------------------------------------------------------------------|
| Name           | Enter a descriptive name for the IPsec session, for example, <code>&lt;VPC--ID&gt;-PCG1-to-remote_edge</code> |
| VPN Service    | Select the VPN service you created in step 2a.                                                                |
| Local Endpoint | Select the local endpoint you created in step 2b.                                                             |
| Remote IP      | Enter the public IP address of the remote peer with which you are creating the VPN tunnel.                    |

| Option                  | Description                                                                                                                                          |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
|                         | <b>Note</b> Remote IP can be a private IP address if you are able to reach the private IP address, for example, using DirectConnect or ExpressRoute. |
| <b>Tunnel Interface</b> | Enter the tunnel interface in a CIDR format. The same subnet must be used for the remote peer to establish the IPSec session.                        |

**Step 2a.**

VPN SERVICES

| Name             | Service Type | Tier0/Tier1 Gateway            | Sessions |
|------------------|--------------|--------------------------------|----------|
| <VPC-ID>-AWS_VPN | IPSec        | cloud-t0-vpc-073617880a9622d93 | 6        |

Description: VPN Service on AWS Transit VPC ID vpc-073617880a9622d93

IKE Log Level: Info

Session sync: Enabled

Admin Status: Enabled

**Step 2b.**

LOCAL ENDPOINTS

| Name                      | VPN Service      | IP Address | Site Certificate | Se |
|---------------------------|------------------|------------|------------------|----|
| <VPC-ID>-PCG-preferred-LE | <VPC-ID>-AWS_VPN | 10.99.3.35 | Not Set          | 3  |

Description: Not Set

Local ID: 10.99.3.35

Trusted CA Certificates: Not Set

Certificate Revocation List: Not Set

Tags: 0

**Step 2c.**

IPSEC SESSIONS

| Name                         | Type        | VPN Service      | Local Endpoint            | Remote IP    |
|------------------------------|-------------|------------------|---------------------------|--------------|
| <VPC-ID>-PCG1-to-remote_edge | Route Based | <VPC-ID>-AWS_VPN | <VPC-ID>-PCG-preferred-LE | 3.213.92.220 |

Description: Not Set

Admin Status: Enabled

Compliance suite: None

Tunnel Interface: 192.168.50.10/24

Authentication Mode: PSK


Remote ID: 172.0.3.145

Pre-shared Key: \*\*\*\*\*

IP address of VPN Peer

- 3 Set up BGP neighbors on the IPsec VPN tunnel interface that you established in step 2. See [Configure BGP](#) for more details.
  - a Navigate to **Networking > Tier-0 Gateways**
  - b Select the auto-created tier-0 gateway for which you created the IPsec session and click **Edit**.
  - c Click the number or icon next to **BGP Neighbors** under the **BGP** section and provide the following details:

| Option           | Description                                                                                                    |
|------------------|----------------------------------------------------------------------------------------------------------------|
| IP Address       | Use the IP address of the remote VTI configured on the tunnel interface in the IPsec session for the VPN peer. |
| Remote AS Number | This number must match the AS number of the remote peer.                                                       |




**Tier-0 Gateways**

ADD GATEWAY EXPAND ALL Filter by Name

| Tier-0 Gateway Name    | HA Mode     | Linked Tier-1 Gateways       | Linked Segments |
|------------------------|-------------|------------------------------|-----------------|
| Local AS               | 1000        |                              |                 |
| BGP                    | On          | ECMP                         | On              |
| Graceful Restart       | Helper Only | Multipath Relax              | On              |
| Graceful Restart Timer | 180 seconds | Graceful Restart Stale Timer | 600 seconds     |
| Route Aggregation      | 0           | BGP Neighbors                |                 |

**Step 3.**



**BGP Neighbors**

Tier-0 Gateways cloud-t0-vpc... #Neighbors

| IP Address       | BFD      | Remote AS number |
|------------------|----------|------------------|
| 192.168.50.11    | Disabled | 1000             |
| Source Addresses | Not Set  |                  |
| Max Hop Limit    | 1        |                  |



- 4 Advertise the prefixes you want to use for the VPN using the Redistribution Profile. In NSX Enforced Mode, connect tier-1 enabled routes in the redistribution profile.

**Tier-0 Gateways**

ADD GATEWAY ▾ EXPAND ALL Filter by Name, Path and more

| Tier-0 Gateway Name                                                                                           | HA Mode | Linked Tier-1 Gateways | Linked Segments | Status ⓘ |
|---------------------------------------------------------------------------------------------------------------|---------|------------------------|-----------------|----------|
| <b>ROUTE RE-DISTRIBUTION</b><br>Route Re-distribution: <b>Step 4.</b> Route Re-distribution Status: <b>On</b> |         |                        |                 |          |

**Route Re-distribution**

Tier-0 Gateways: cloud-t0-415... #Selected Sources 1

**Tier-0 Subnets**

**Advertised Tier-1 Subnets**

- Connected Interfaces & Segments
- Service Interface Subnet
- Connected Segment

## NSX Cloud FAQs and Troubleshooting

This topic covers some frequently asked questions and troubleshooting information.

### How can I verify that my NSX Cloud components are installed and running?

- 1 To verify that NSX Tools on your workload VM are connected to PCG, do the following:
  - a Type the `nsxcli` command to open NSX CLI.
  - b Type the following command to get the gateway connection status, for example:

```
get gateway connection status
Public Cloud Gateway : nsx-gw.vmware.com:5555
Connection Status : ESTABLISHED
```

- 2 The workload VMs must have the correct tags to connect to PCG:
  - a Log in to the AWS console or the Microsoft Azure portal.
  - b Verify the VM's eth0 or interface tag.

The `nsx.network` key must have the value `default`.

## My VMs launched using cloud-init are quarantined and do not allow installation of third-party tools. What should I do?

With the Quarantine Policy enabled, when launching VMs using cloud-init scripts with the following specifications, your VMs are quarantined upon launching and you are not able to install custom applications or tools on them:

- tagged with `nsx.network=default`
- custom services auto-installed or bootstrapped when the VM is powered on

### Solution:

Update the `default` (AWS) or `default-vnet-<vnet-ID>-sg` (Microsoft Azure) security group to add inbound/outbound ports as required for the installation of custom or third-party applications.

## I tagged my VM correctly and installed NSX Tools, but my VM is quarantined. What should I do?

If you encounter this problem, try the following:

- Check whether the NSX Cloud tag: `nsx.network` and its value: `default` are correctly typed in. This is case-sensitive.
- Resync the AWS or Microsoft Azure account from CSM:
  - Log in to CSM.
  - Go to **Clouds > AWS/Azure > Accounts**.
  - Click on **Actions** from the public cloud account tile and click **Resync Account**.

## What should I do if I cannot access my workload VM?

From your Public Cloud (AWS or Microsoft Azure):

- 1 Ensure that all ports on the VM, including those managed by NSX Cloud, the OS firewall (Microsoft Windows or IPTables), and NSX-T Data Center are properly configured in order to allow traffic,

For example, to allow `ping` to a VM, the following needs to be properly configured:

- Security Group on AWS or Microsoft Azure. See [Threat Detection using the NSX Cloud Quarantine Policy](#) for more information.

- NSX-T Data Center DFW rules. See [Default Connectivity Strategy for NSX-Managed Workload VMs in the NSX Enforced Mode](#) for details.
  - Windows Firewall or IPTables on Linux.
- 2 Attempt resolving the issue by logging in to the VM using SSH or other methods, for example, the Serial Console in Microsoft Azure.
  - 3 You can reboot the locked out VM.
  - 4 If you still cannot access the VM, then attach a secondary NIC to the workload VM from which to access that workload VM.

## Do I need a PCG even in the Native Cloud Enforced Mode?

Yes.

## Can I change the IAM role for the PCG after I have onboarded my public cloud account in CSM?

Yes. You can rerun the NSX Cloud script applicable to your public cloud to regenerate the PCG role. Edit your public cloud account in CSM with the new rolename after you regenerate the PCG role . Any new PCG instances deployed in your public cloud account will use the new role.

Note that existing PCG instances continue to use the old PCG role. If you want to update the IAM role for an existing PCG instance, go to your public cloud and manually change the role for that PCG instance.

## Can I use the NSX-T Data Center on-prem licenses for NSX Cloud?

Yes, you can if your ELA has a clause for it.

## I am using the URL from CSM to deploy PCG but I get an error because the gateway name is unresolvable.

When the URL from the CSM UI for installing PCG fails because of gateway name being unresolvable, do the following in the respective public cloud for the OS of your workload VM:

- On Microsoft Windows workload VMs in Microsoft Azure, run the following command and download the install script again using the URL from CSM:

```
Add-DnsClientNrptRule -Namespace "nsx-gw.vmware.local" -NameServers "168.63.129.16"
-DnsSecEnable
```

- On Microsoft Windows workload VMs in AWS, run the following command and download the install script again using the URL from CSM::

```
Add-DnsClientNrptRule -Namespace "nsx-gw.vmware.local" -NameServers "169.254.169.253"
-DnsSecEnable
```

- On Linux workload VMs in Microsoft Azure run the following command to get PCG's IP addresses and download the install script using these IP addresses with the URL from CSM.

```
nslookup nsx-gw.vmware.local 168.63.129.16 | awk '/^Address: / { print $2 }'
```

- On Linux workload VMs in AWS run the following command to get PCG's IP addresses and download the install script using these IP addresses with the URL from CSM.:

```
nslookup nsx-gw.vmware.local 169.254.169.253 | awk '/^Address: / { print $2 }'
```