

NSX-T Data Center Administration Guide

Modified on 06 MAY 2022
VMware NSX-T Data Center 2.5

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

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

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

Copyright © 2022 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

About Administering VMware NSX-T Data Center	12
1 Overview of the NSX Manager	13
2 Tier-0 Gateways	16
Add a Tier-0 Gateway	16
Create an IP Prefix List	20
Create a Community List	21
Configure a Static Route	22
Create a Route Map	23
Using Regular Expressions to Match Community Lists When Adding Route Maps	25
Configure BGP	26
Configure BFD	29
Configure IPv6 Layer 3 Forwarding	29
Create SLAAC and DAD Profiles for IPv6 Address Assignment	30
3 Tier-1 Gateway	33
Add a Tier-1 Gateway	33
4 Segments	36
Segment Profiles	36
Understanding QoS Segment Profile	37
Understanding IP Discovery Segment Profile	40
Understanding SpoofGuard Segment Profile	42
Understanding Segment Security Segment Profile	43
Understanding MAC Discovery Segment Profile	45
Add a Segment	46
5 Virtual Private Network (VPN)	49
Understanding IPSec VPN	50
Using Policy-Based IPSec VPN	50
Using Route-Based IPSec VPN	51
Understanding Layer 2 VPN	53
Adding VPN Services	54
Add an IPSec VPN Service	55
Add an L2 VPN Service	56
Adding IPSec VPN Sessions	59
Add a Policy-Based IPSec Session	59

Add an IP Address Block 136

10 Security 137

Security Configuration Overview 137

Security Terminology 138

Identity Firewall 138

Identity Firewall Workflow 139

Layer 7 Context Profile 141

Layer 7 Firewall Rule Workflow 143

Attributes 143

Distributed Firewall 147

Firewall Drafts 148

Add a Distributed Firewall 150

Distributed Firewall Packet Logs 154

Select a Default Connectivity Strategy 156

Manage a Firewall Exclusion List 157

Filtering Specific Domains (FQDN/URLs) 157

Extending Security Policies to Physical Workloads 159

Shared Address Sets 166

East-West Network Security - Chaining Third-party Services 166

Key Concepts of Network Protection East-West 167

NSX-T Data Center Requirements for East-West Traffic 168

High-Level Tasks for East-West Network Security 168

Deploy a Service for East-West Traffic Introspection 169

Add a Service Profile 170

Add a Service Chain 171

Add Redirection Rules for East-West Traffic 172

Configuring a Gateway Firewall 174

Add a Gateway Firewall Policy and Rule 174

North-South Network Security - Inserting Third-party Service 177

High-Level Tasks for North-South Network Security 177

Deploy a Service for North-South Traffic Introspection 177

Configure Traffic Redirection 179

Add Redirection Rules for North-South Traffic 180

Monitor Traffic Redirection 182

Endpoint Protection 182

Understand Endpoint Protection 182

Configure Endpoint Protection 187

Manage Endpoint Protection 202

Security Profiles 214

Create a Session Timer 214

- Flood Protection 216
- Configure DNS Security 218
- Manage Group to Profile Precedence 219

11 Inventory 221

- Add a Service 221
- Add a Group 222
- Add a Context Profile 224

12 Monitoring 226

- Add a Firewall IPFIX Profile 226
- Add a Switch IPFIX Profile 227
- Add an IPFIX Collector 228
- Add a Port Mirroring Profile 229
- Simple Network Management Protocol (SNMP) 230
- Using vRealize Log Insight for System Monitoring 230
- Using vRealize Operations Manager for System Monitoring 231
- Using vRealize Network Insight Cloud for System Monitoring 235
- Advanced Monitoring Tools 247
 - View Port Connection Information 248
 - Traceflow 248
 - Monitor Port Mirroring Sessions 251
 - Configure Filters for a Port Mirroring Session 254
 - Configure IPFIX 255
 - Monitor a Logical Switch Port Activity 425

13 Logical Switches 426

- Understanding BUM Frame Replication Modes 427
- Create a Logical Switch 429
- Connecting a VM to a Logical Switch 430
 - Attach a VM Hosted on vCenter Server to an NSX-T Data Center Logical Switch 430
 - Attach a VM Hosted on Standalone ESXi to an NSX-T Data Center Logical Switch 432
 - Attach a VM Hosted on KVM to an NSX-T Data Center Logical Switch 438
- Create a Logical Switch Port 439
- Test Layer 2 Connectivity 440
- Create a VLAN Logical Switch for the NSX Edge Uplink 443
- Switching Profiles for Logical Switches and Logical Ports 445
 - Understanding QoS Switching Profile 446
 - Understanding Port Mirroring Switching Profile 449
 - Understanding IP Discovery Switching Profile 452
 - Understanding SpoofGuard 454

- Understanding Switch Security Switching Profile 456
- Understanding MAC Management Switching Profile 458
- Associate a Custom Profile with a Logical Switch 459
- Associate a Custom Profile with a Logical Port 460
- Enhanced Networking Stack 461
 - Automatically Assign ENS Logical Cores 462
 - Configure Guest Inter-VLAN Routing 463
- Layer 2 Bridging 464
 - Create an Edge Bridge Profile 465
 - Configure Edge-Based Bridging 465
 - Create a Layer 2 Bridge-Backed Logical Switch 468

14 Logical Routers 471

- Tier-1 Logical Router 471
 - Create a Tier-1 Logical Router 473
 - Add a Downlink Port on a Tier-1 Logical Router 474
 - Add a VLAN Port on a Tier-0 or Tier-1 Logical Router 475
 - Configure Route Advertisement on a Tier-1 Logical Router 476
 - Configure a Tier-1 Logical Router Static Route 478
 - Create a Standalone Tier-1 Logical Router 479
- Tier-0 Logical Router 481
 - Create a Tier-0 Logical Router 482
 - Attach Tier-0 and Tier-1 483
 - Connect a Tier-0 Logical Router to a VLAN Logical Switch for the NSX Edge Uplink 486
 - Add a Loopback Router Port 489
 - Add a VLAN Port on a Tier-0 or Tier-1 Logical Router 490
 - Configure a Static Route 490
 - BGP Configuration Options 494
 - Configure BFD on a Tier-0 Logical Router 500
 - Enable Route Redistribution on the Tier-0 Logical Router 501
 - Understanding ECMP Routing 504
 - Create an IP Prefix List 508
 - Create a Community List 509
 - Create a Route Map 509
 - Configure Forwarding Up Timer 510

15 Advanced NAT 512

- Network Address Translation 512
 - Tier-1 NAT 514
 - Tier-0 NAT 521
 - Reflexive NAT 522

16 Advanced Grouping Objects 526

- Create an IP Set 526
- Create an IP Pool 527
- Create a MAC Set 527
- Create an NSGroup 528
- Configuring Services and Service Groups 530
 - Create an NSService 530
- Manage Tags for a VM 531

17 Advanced DHCP 532

- DHCP 532
 - Create a DHCP Server Profile 532
 - Create a DHCP Server 533
 - Attach a DHCP Server to a Logical Switch 534
 - Detach a DHCP Server from a Logical Switch 534
 - Create a DHCP Relay Profile 534
 - Create a DHCP Relay Service 535
 - Add a DHCP Relay Service to a Logical Router Port 535
 - Delete a DHCP Lease 536
- Metadata Proxies 536
 - Add a Metadata Proxy Server 536
 - Attach a Metadata Proxy Server to a Logical Switch 537
 - Detach a Metadata Proxy Server from a Logical Switch 537

18 Advanced IP Address Management 539

- Manage IP Blocks 539
- Manage Subnets for IP Blocks 540

19 Advanced Load Balancing 541

- Key Load Balancer Concepts 542
- Configuring Load Balancer Components 542
 - Create a Load Balancer 543
 - Configure an Active Health Monitor 544
 - Configure Passive Health Monitors 547
 - Add a Server Pool for Load Balancing 549
 - Configuring Virtual Server Components 553

20 Advanced Firewall 574

- Add or Delete a Firewall Rule to a Logical Router 574
- Configure Firewall for a Logical Switch Bridge Port 575
- Firewall Sections and Firewall Rules 575

- Enable and Disable Distributed Firewall 576
 - Add a Firewall Rule Section 576
 - Delete a Firewall Rule Section 577
 - Enable and Disable Section Rules 578
 - Enable and Disable Section Logs 578
 - Configure a Firewall Exclusion List 578
- About Firewall Rules 579
 - Add a Firewall Rule 580
 - Delete a Firewall Rule 583
 - Edit the Default Distributed Firewall Rule 583
 - Change the Order of a Firewall Rule 584
 - Filter Firewall Rules 585
- 21 Operations and Management 586**
 - View Monitoring Dashboards 587
 - View the Usage and Capacity of Categories of Objects 589
 - Checking the Realized State of a Configuration Change 591
 - Search for Objects 595
 - Filter by Object Attributes 596
 - Add a Compute Manager 596
 - Add an Active Directory 598
 - Add an LDAP Server 599
 - Synchronize Active Directory 600
 - Managing User Accounts and Role-Based Access Control 601
 - Manage a User's Password 601
 - Resetting the Passwords of an Appliance 603
 - Authentication Policy Settings 604
 - Obtain the Certificate Thumbprint from a vIDM Host 605
 - Configure VMware Identity Manager Integration 606
 - Validate VMware Identity Manager Functionality 608
 - Time Synchronization between NSX Manager, vIDM, and Related Components 610
 - Role-Based Access Control 611
 - Add a Role Assignment or Principal Identity 622
 - Backing Up and Restoring the NSX Manager 624
 - Configure Backups 625
 - Removing Old Backups 626
 - Listing Available Backups 627
 - Restore a Backup 627
 - Backup and Restore During Upgrade 630
 - Remove NSX-T Data Center Extension from vCenter Server 630
 - Managing the NSX Manager Cluster 631

View the Configuration and Status of the NSX Manager Cluster	631
Shut Down and Power On the NSX Manager Cluster	634
Reboot an NSX Manager	634
Change the IP Address of an NSX Manager	635
Resize an NSX Manager Node	637
Adding and Removing an ESXi Host Transport Node to and from vCenter Servers	637
Replacing an NSX Edge Transport Node in an NSX Edge Cluster	638
Replace an NSX Edge Transport Node Using the NSX Manager UI	638
Replace an NSX Edge Transport Node Using the API	639
Recovering NSX-T When vCenter Server Is Lost and Cannot Be Recovered	641
Multisite Deployment of NSX-T Data Center	642
Configuring Appliances	649
Add a License Key and Generate a License Usage Report	650
Setting Up Certificates	651
Import a Certificate	651
Create a Certificate Signing Request File	652
Import a CA Certificate	653
Create a Self-Signed Certificate	654
Replace the Certificate for an NSX Manager Node or an NSX Manager Cluster Virtual IP	655
Import a Certificate Revocation List	656
Configuring NSX Manager to Retrieve a Certificate Revocation List	657
Import a Certificate for a CSR	657
Storage of Public Certificates and Private Keys	658
Compliance-Based Configuration	658
View Compliance Status Report	658
Compliance Status Report Codes	659
Configure Global FIPS Compliance Mode for Load Balancer	662
Collect Support Bundles	665
Log Messages and Error Codes	666
Configure Remote Logging	668
Log Message IDs	675
Troubleshooting Syslog Issues	677
Configure Serial Logging on an Appliance VM	677
Customer Experience Improvement Program	678
Edit the Customer Experience Improvement Program Configuration	678
Add Tags to an Object	679
Find the SSH Fingerprint of a Remote Server	680
View Data about Applications Running on VMs	681
Configuring an External Load Balancer	682

22 Using NSX Cloud 683

- A Quick Tour of the Cloud Service Manager 683
 - Clouds 683
 - System 690
- Threat Detection using the NSX Cloud Quarantine Policy 692
 - Quarantine Policy in the NSX Enforced Mode 693
 - Quarantine Policy in the Native Cloud Enforced Mode 699
 - Whitelisting VMs 699
- NSX Enforced Mode 700
 - Currently Supported Operating Systems for Workload VMs 700
 - Onboarding VMs in the NSX Enforced Mode 701
 - Managing VMs in the NSX Enforced Mode 710
- Native Cloud Enforced Mode 711
 - Managing VMs in the Native Cloud Enforced Mode 711
- NSX-T Data Center Features Supported with NSX Cloud 715
 - Group VMs using NSX-T Data Center and Public Cloud Tags 716
 - Use Native-Cloud Services 719
 - Service Insertion for your Public Cloud 720
 - Enable NAT on NSX-managed VMs 727
 - Enable Syslog Forwarding 728
 - Set up VPN in the NSX Enforced Mode 728
- Frequently Asked Questions (FAQs) 734

23 Using NSX Intelligence 737

- Getting Started with NSX Intelligence 737
 - Tour of the NSX Intelligence Home Page 737
 - Getting Familiar with NSX Intelligence Graphic Elements 740
- Understanding NSX Intelligence Views and Flows 742
 - Working with the Groups View 742
 - Working with the VMs View 747
 - Working with Traffic Flows 749
- Working with NSX Intelligence Recommendations 751
 - Understanding NSX Intelligence Recommendations 751
 - Generate a New NSX Intelligence Recommendation 752
 - Review and Publish a Generated Recommendation 753
- Backing Up and Restoring NSX Intelligence 755
 - Configure NSX Intelligence Backups 756
 - Back Up NSX Intelligence 757
 - Restore NSX Intelligence Backups 758
- Troubleshooting NSX Intelligence Issues 759
 - Check the Status of the NSX Intelligence Appliance 759
 - Collect NSX Intelligence Support Bundles 764

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>.

Overview of the NSX Manager

1

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

The NSX Manager web interface provides two methods of configuring resources.

- The Policy interface: the **Networking, Security, Inventory, and Plan & Troubleshoot** tabs.
- The Advanced interface: the **Advanced Networking & Security** tab.

When to Use Policy or Advanced Interfaces

Be consistent about which user interface you use. There are a few reasons to use one user interface over another.

- If you are deploying a new environment with NSX-T Data Center 2.4 or later, using the new policy-based user interface to create and manage your environment is the best choice in most situations.
 - Some features are not available in the policy-based user interface. If you need these features, use the Advanced user interface for all configurations.
- If you are upgrading to NSX-T Data Center 2.4 or later, continue to make configuration changes using the **Advanced Networking & Security** user interface.

Table 1-1. When to Use Policy or Advanced Interfaces

Policy Interface	Advanced Interface
Most new deployments should use the policy-based interface.	Deployments which were created using the advanced interface, for example, upgrades from versions before the policy-based interface was present.
NSX Cloud deployments	Deployments which integrate with other plugins. For example, NSX Container Plug-in, Openstack, and other cloud management platforms.

Table 1-1. When to Use Policy or Advanced Interfaces (continued)

Policy Interface	Advanced Interface
Networking features available in the Policy interface only: <ul style="list-style-type: none"> ■ DNS Services and DNS Zones ■ VPN ■ Forwarding policies for NSX Cloud 	Networking features available in the Advanced interface only: <ul style="list-style-type: none"> ■ Forwarding up timer ■ Static routes with BFD and interface as next-hop ■ Metadata proxy ■ DHCP server attached to an isolated segment and static binding
Security features available in the Policy interface only: <ul style="list-style-type: none"> ■ Endpoint Protection ■ Network Introspection (East-West Service Insertion) ■ Context Profiles <ul style="list-style-type: none"> ■ L7 applications ■ FQDN ■ New Distributed Firewall and Gateway Firewall Layout <ul style="list-style-type: none"> ■ Categories ■ Auto service rules ■ Drafts 	Security features available in the Advanced interface only: <ul style="list-style-type: none"> ■ CPU and memory thresholds ■ Bridge Firewall ■ Distributed Firewall rules based on IPs in source and destination

Using the Policy Interface

If you decide to use the policy interface, use it to create all objects. Do not use the advanced interface to create objects.

You can use the advanced interface to modify objects that have been created in the policy interface. The settings for a policy-created object might include a link for **Advanced Configuration**. This link takes you to the advanced interface where you can fine-tune the configuration. You can also view policy-created objects in the advanced interface directly. Settings that are managed by policy but are visible in the advanced interface have this icon next to them:

. You cannot modify them from the advanced user interface.

Where to Find the Policy Interfaces and Advanced Interfaces

The policy-based and advanced interfaces appear in different parts of the NSX Manager user interface, and use different API URIs.

Table 1-2. Policy Interfaces and Advanced Interfaces

Policy Interface	Advanced Interface
<ul style="list-style-type: none"> ■ Networking tab ■ Security tab ■ Inventory tab ■ Plan & Troubleshoot tab 	Advanced Networking & Security tab
API URIs that begin with <code>/policy/api</code>	API URIs that begin with <code>/api</code>

Note The **System** tab is used for all environments. If you modify Edge nodes, Edge clusters, or transport zones, it can take up to 5 minutes for those changes to be visible on the policy-based user interface. You can synchronize immediately using `POST /policy/api/v1/infra/sites/default/enforcement-points/default?action=reload`.

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

Names for Objects Created in the Policy and Advanced Interfaces

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

Table 1-3. Object Names

Objects Created Using the Policy Interface	Objects Created Using the Advanced Interface
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
Rule	Firewall rule
Gateway firewall	Edge firewall

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 In the **Advanced Networking & Security** tab, the term tier-0 logical router is used to refer to a tier-0 gateway.

This chapter includes the following topics:

- [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 IPv6 Layer 3 Forwarding](#)
- [Create SLAAC and DAD Profiles for IPv6 Address Assignment](#)

Add a Tier-0 Gateway

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

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

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** .

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.

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.

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.

Important After you create the gateway, the HA mode cannot be changed.

- 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) Add one or more tags.
- 9 (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 169.254.0.0/28.
 - 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.
 - c Select an **ND Profile** and a **DAD Profile** for IPv6 address configuration.
These profiles are used to configure Stateless Address Autoconfiguration (SLAAC) and Duplicate Address Detection (DAD) for IPv6 addresses. The default profile is created.
- 10 Click **Save**.

- 11 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, 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**.

- 12 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) Add tags and select an ND profile.

- 13 (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.

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

- 15 Click **BGP** to configure BGP.

- 16 Click **Advanced Configuration** to go to the **Advanced Networking & Security > Routers** page to make additional configurations.
 - a To configure the layer 3 forwarding mode, click the **Global Config** tab.
 - b Click **Edit**.
 - c Select **IPv4** or **IPv4 and IPv6**.

The default is IPv4 only. IPv6 only is not supported. To enable IPv6, select **IPv4 and IPv6**.
 - d Click **Save**.

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](#).

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.
- 11 Specify the administrative distance.
- 12 Select an interface from the dropdown list.
- 13 Click the **Add** button.

What to do next

Check that the static route is configured properly. See [Verify the Static Route](#).

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](#) 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: <ul style="list-style-type: none"> ■ 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.

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 Configure **Out Filter**.
- e Configure **In Filter**.
- f Enable or disable the **Allows-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.

- g 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.
- h In the **IP Address Family** field, select **IPv4**, **IPv6**, or **Disabled**.
- i Enter a value for **Max Hop Limit**.
- j 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	<ul style="list-style-type: none"> ■ 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	<ul style="list-style-type: none"> ■ 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	<ul style="list-style-type: none"> ■ 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.

k Click **Timers & Password**.

l Enter a value for **BFD Interval**.

The unit is milliseconds. For an Edge node running in a VM, the minimum value is 1000. For a bare-metal Edge node, the minimum value is 300.

m Enter a value for **BFD Multiplier**.

n Enter a value for **Hold Down Time**.

o Enter a value for **Keep Alive Time**.

p 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 **Advanced Configuration**.

This takes you to the **Advanced Networking & Security > Routers** page. The gateway will appear as one of the logical routers. Follow the instructions in [Configure BFD on a Tier-0 Logical Router](#).

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 > Tier-0 Gateways**.
- 3 Edit a tier-0 gateway by clicking the menu icon (three dots) and select **Edit**.
- 4 Click **Advanced Configuration**.

This takes you to the **Advanced Networking & Security > Routers** page. The gateway will appear as one of the logical routers.

- 5 Click the **Global Config** tab.
- 6 In the **L3 Forwarding Mode** field, select **IPv4 and IPv6**.
IPv6 only is not supported.
- 7 Edit the gateway again by going to the **Networking** tab.
- 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 **T0-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 **Advanced Networking & Security > Routers > Global Config** 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)**

Tier-1 Gateway

3

A tier-1 gateway performs the functions of a tier-1 logical router. It has downlink connections to segments and uplink connections to tier-0 gateways.

Note In the **Advanced Networking & Security** tab, the term tier-1 logical router is used to refer to a tier-1 gateway.

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

This chapter includes the following topics:

- [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.

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 Select a failover mode.

Option	Description
Preemptive	If the preferred NSX Edgenode 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.

- 7 (Optional) Select an NSX Edge cluster if you want this tier-1 gateway to host stateful services (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.

- 8 (Optional) Select an NSX Edge node.
- 9 (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.

- 10 Click **Save**.
- 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**

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

12 (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 In the **ND Profile** field, select a profile.
- f Click **Save**.

13 (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**.

Segments

4

A segment performs the functions of a logical switch.

Note In the **Advanced Networking & Security** tab, the term logical switch is used to refer to a segment.

This chapter includes the following topics:

- [Segment Profiles](#)
- [Add a Segment](#)

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 GoS switching profile details.

Option	Description
Name	Name of the profile.
Mode	<p>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.</p> <p>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.</p> <p>Note DSCP settings work only on tunneled traffic. These settings do not apply to traffic inside the same hypervisor.</p>
Priority	<p>Set the CoS priority value.</p> <p>The priority values range from 0 to 63, where 0 has the highest priority.</p>
Class of Service	<p>Set the CoS value.</p> <p>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.</p> <p>The CoS values range from 0 to 7, where 0 is the best effort service.</p>
Ingress	<p>Set custom values for the outbound network traffic from the VM to the logical network.</p> <p>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.</p> <p>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.</p> <p>The default value 0 disables rate limiting on the ingress traffic.</p>
Ingress Broadcast	<p>Set custom values for the outbound network traffic from the VM to the logical network based on broadcast.</p> <p>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.</p> <p>The default value 0 disables rate limiting on the ingress broadcast traffic.</p>
Egress	<p>Set custom values for the inbound network traffic from the logical network to the VM.</p> <p>The default value 0 disables rate limiting on the egress traffic.</p>

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 the 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 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. By navigating to **Advanced Networking and Security > Switching > Ports** and selecting 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 switching profiles. You can also create additional ones.

Prerequisites

Familiarize yourself with the IP Discovery switching profile concepts. See [Understanding IP Discovery 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 **IP Discovery**.
- 4 Specify the IP Discovery switching 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 (the default) and the maximum is 256.
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.

Option	Description
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.

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 end 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 whitelist 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 whitelist 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.
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 management 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. This property is disabled 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

A segment connects to gateways and VMs. A segment performs the functions of a logical switch.

For information about find the VIF ID of a VM, see [Connecting a VM to a Logical Switch](#).

Note An N-VDS switch configured in the Enhanced Datapath mode supports IP Discovery, SpoofGuard, and IPFIX profiles.

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 a connected gateway.

You can select an existing Tier-0 or Tier-1 gateway, or select **None**. The default value is **None**, which means the segment is simply a logical switch. With a subnet configured, it can link to a Tier-0 or Tier-1 gateway.

- 6 If the connected gateway is a Tier-1 gateway, select a type, either **Flexible** or **Fixed**.

A flexible segment can be unlinked from gateways. A fixed segment can be deleted but not unlinked from a gateway.

- 7 To specify a subnet, click **Set Subnets**.

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

- 9 If the transport zone is of type VLAN, specify a list of VLAN IDs.

- 10 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.

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

- 12 Click **Save**.

- 13 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: **Parent**, **Child**, or **Independent**.

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 container host VM, select **Parent**. If this port is for a bare metal container or server, select **Independent**.

f Enter a context ID.

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

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.

- 14 To select segment profiles, click **Segment Profiles**.

15 Click **Save**.

Virtual Private Network (VPN)

5

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 allowing 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.

This chapter includes the following topics:

- [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.

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

IPsec VPN service in NSX-T Data Center uses the gateway-level failover functionality to support a high-availability service. Tunnels are re-established on failover and VPN configuration data is synchronized. The IPSec VPN state is not synchronized as tunnels are re-established.

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.

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

Table 5-1. 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
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

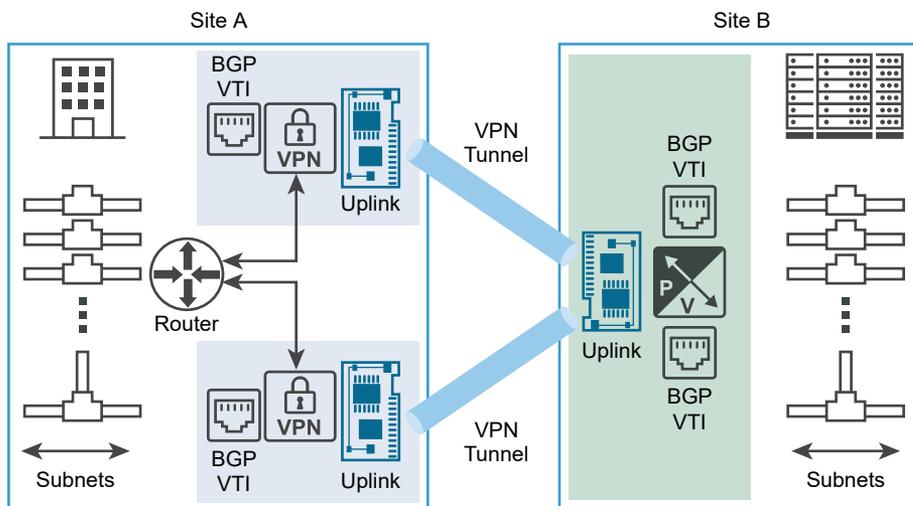
You can configure VPN tunnel redundancy with a route-based IPsec VPN session that is configured on a Tier-0 gateway. With tunnel redundancy, multiple tunnels can be set up between two sites, with one tunnel being used as the primary with failover to the other tunnels when the primary tunnel becomes unavailable. 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 5-1. 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, so VMs remain on the same subnet when they are moved between sites and their IP addresses do not change. 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-premise 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.

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.

In NSX-T Data Center, L2 VPN services are supported only on Tier-0 gateways. Segments can be connected to either Tier-0 or Tier-1 gateways and use L2 VPN services.

Starting with NSX-T Data Center 2.5 release, VLAN-based segments can be extended using L2 VPN service on an NSX Edge that is managed in an NSX-T Data Center environment. This support allows the extension of L2 networks from VLAN to VNI, VLAN to VLAN, and VNI to VNI.

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

The L2 VPN service support is provided in the following 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)

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 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 gateway and configure an L2 VPN Client service using that Tier-0 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 **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 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, even if the IP addresses are specified in the IPSec session rules. In the **Local Networks** and **Remote Networks** text boxes, enter the list of local and remote subnets between which the bypass rules are applied.
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.
- 10 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 gateway. To enable the L2 VPN service, you must first create an IPSec VPN service on the Tier-0 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 gateway configured and available for use. See [Add a Tier-0 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 the Tier-0 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 Gateway** drop-down menu, select a Tier-0 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 Gateway** drop-down menu, select the same Tier-0 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 one does not exist yet, create an IPsec VPN service for the L2 VPN client 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 Gateway** drop-down menu, select a Tier-0 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 Gateway** drop-down menu, select the same Tier-0 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 [Setting Up 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 5-2. 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 [Setting Up 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 5-3. 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

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

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 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**.

- 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.

```
MCw3ZjBjYzdzLHsic210ZU5hbWUiOiJSb3V0ZWJhc2UxIiwic3JjVGFwSXAiOiIxNjkuMjU0LjY0LjIiLCJkc3RUYXB
JcCI6IjE2OS4yNTQuNjQuMSIsImlrZU9wdG1
vbiI6ImlrZXYyIiwic2V5jYXBQcm90byI6ImdyZS9pcHNlYyIsImRoR3JvdXAiOiJkaDE0Iiwic2V5jcnlwdEFuZERpZ2
VzdCI6ImFlcylnY20vc2hhLTI1NiIsInBzayI
6IlZNd2FyZTEyMyIsInR1bm5lbHMlOiI7ImxvY2FsSWQiOiI2MC42MC42MC4xIiwicGVlcklkIjoiaWNTAuNTAuNTAuMS
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 certificate-based VPN sessions defined using 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.
 - If the X509v3 extension `x509v3 Subject Alternative Name` is found in the certificate, then one of the Subject Alternative Name is taken as the local ID value.
- 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 Specify a tag, if needed.
- 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.

If you decide not to use the default profiles provided by NSX-T Data Center, you can configure your own 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 5-4. Default IKE Profiles Used for IPSec VPN or L2 VPN Services

Default IKE Profile Name	Description
nsx-default-l2vpn-ike-profile	<ul style="list-style-type: none"> ■ Used for an L2 VPN service configuration. ■ Configured with IKE V2, AES 128 encryption algorithm, SHA2 256 algorithm, and Diffie-Hellman group14 key exchange algorithm.
nsx-default-l3vpn-ike-profile	<ul style="list-style-type: none"> ■ Used for an IPSec VPN service configuration. ■ Configured with IKE V2, AES 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 Click the **Networking > VPN > 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 5-5. 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 5-6. Algorithms Used

Type of Algorithm	Valid Values	Description
Encryption	<ul style="list-style-type: none"> ■ AES 128 (default) ■ AES 256 ■ AES GCM 128 ■ AES GCM 192 ■ AES GCM 256 	<p>The encryption algorithm used during the Internet Key Exchange (IKE) negotiation.</p> <p>The AES-GCM algorithms are supported when used with IKEv2. They are not supported when used with IKEv1.</p>
Digest	<ul style="list-style-type: none"> ■ SHA2 256 (default) ■ SHA1 ■ SHA2 384 ■ SHA2 512 	<p>The secure hashing algorithm used during the IKE negotiation.</p> <p>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 Psuedo-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.</p> <p>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.</p>
Diffie-Hellman Group	<ul style="list-style-type: none"> ■ Group 14 (default) ■ Group 2 ■ Group 5 ■ Group 15 ■ Group 16 ■ Group 19 ■ Group 20 ■ Group 21 	<p>The cryptography schemes that the peer site and the NSX Edge use to establish a shared secret over an insecure communications channel.</p>

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 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 5-7. Default IPsec Profiles Used for IPsec VPN or L2 VPN Services

Name of Default IPsec Profile	Description
nsx-default-l2vpn-tunnel-profile	<ul style="list-style-type: none"> ■ Used for L2 VPN. ■ Configured with AES GCM 128 encryption algorithm and Diffie-Hellman group 14 key exchange algorithm.
nsx-default-l3vpn-tunnel-profile	<ul style="list-style-type: none"> ■ 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 Navigate to the **Networking > VPN > 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 5-8. Algorithms Used

Type of Algorithm	Valid Values	Description
Encryption	<ul style="list-style-type: none"> ■ 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.
Digest	<ul style="list-style-type: none"> ■ SHA1 ■ SHA2 256 ■ SHA2 384 ■ SHA2 512 	The secure hashing algorithm used during the IPSec negotiation.
Diffie-Hellman Group	<ul style="list-style-type: none"> ■ 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 5-9. 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 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.

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 the **DPD Profiles** profile type, and click **Add DPD Profile**.
- 4 Enter a name for the DPD profile.

- 5 In the **DPD Probe Interval** text box, enter the number of seconds you want NSX-T Data Center to wait before sending the next DPD probe. The default is 60 seconds.

If the NSX Edge node receives a response from the remote peer site, the DPD probe interval timer is restarted. If the NSX Edge node does not hear back from the peer site within 0.5 seconds after the next DPD probe is sent, a retransmission timer is set to 0.5 seconds. The NSX Edge node retransmits the next DPD probe after the retransmission timer is reached. If the remote peer site continues not to respond, the retransmission timer is exponentially increased to the maximum limit of 6 seconds. The NSX Edge node continues to retransmit the DPD probe every time the retransmission timer expires. The NSX Edge node retransmits up to a maximum of 30 times before it declares the peer site to be dead and it tears down the security association (SA) on the dead peer's link. The total time it takes to retransmit the DPD probe 30 times is about 2 minutes and 45 seconds.

- 6 Provide a description and add a tag, as needed.
- 7 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 HA for VPN redundancy by deploying primary and secondary autonomous L2 VPN Edge clients.

Prerequisites

- Create a port group and bind it to the vSwitch on your host.
- 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 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**.

- f (Optional) To configure High Availability, enter the **HA Port** details where the exit interface maps to the appropriate HA Network.
- g (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:
- Enter a session name.
 - Enter the local IP address and the remote IP address.
 - 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:
- Select the L2 VPN session that you created.
 - Select the L2 extension port that you created.
 - 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

6

Network address translation (NAT) maps one IP address space to another. You can configure NAT on tier-0 and tier-1 gateways.

This chapter includes the following topics:

- [Configure NAT on a Gateway](#)

Configure NAT on a Gateway

You can configure source NAT (SNAT), destination NAT (DNAT), or reflexive NAT on a tier-0 or tier-1 gateway.

If a tier-0 gateway is running in active-active mode, you cannot configure SNAT or DNAT because asymmetrical paths might cause issues. You can only configure reflexive NAT (sometimes called stateless NAT). If a tier-0 gateway is running in active-standby mode, you can configure SNAT, DNAT, or reflexive NAT.

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 will process traffic from a tier-1 gateway as well as from another external interface on the 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 > NAT**.
- 3 Select a gateway.
- 4 Click **Add NAT Rule**.
- 5 Select an action.

For a tier-1 gateway, the available actions are **SNAT**, **DNAT**, **Reflexive**, **NO SNAT**, and **NO DNAT**.

For a tier-0 gateway in active-standby mode, the available actions are **SNAT**, **DNAT**, **NO SNAT**, and **NO DNAT**.

For a tier-0 gateway in active-active mode, the available action is **Reflexive**.

6 In the **Service** column, click **Set** to select services.

7 (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.

8 For **Destination 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 Enter a value for **Translated Port**.

11 Select a firewall setting from the following options:

- **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.

12 (Required) Change the logging status.

13 (Required) For **Applied To**, select objects that this rule applies to.

The available objects are **Tier-0 Gateways**, **Interfaces**, **Labels**, **Service Instance Endpoints**, and **Virtual Endpoints**.

14 Specify a priority value.

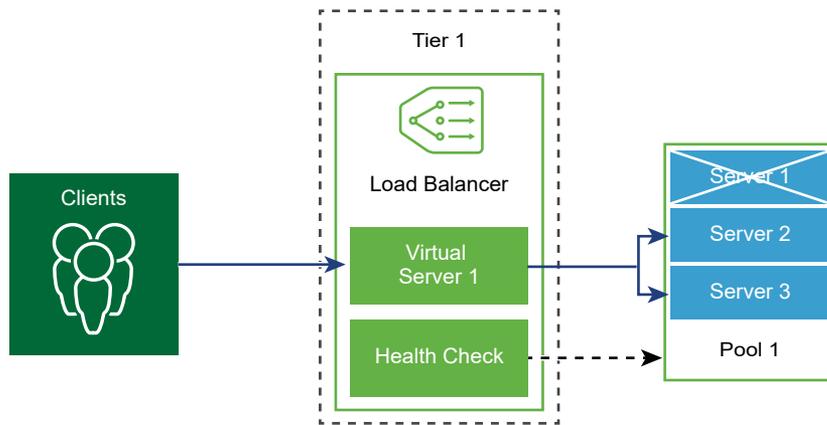
A lower value means a higher priority. The default is 100.

15 Click **Save**.

Load Balancing

7

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.

This chapter includes the following topics:

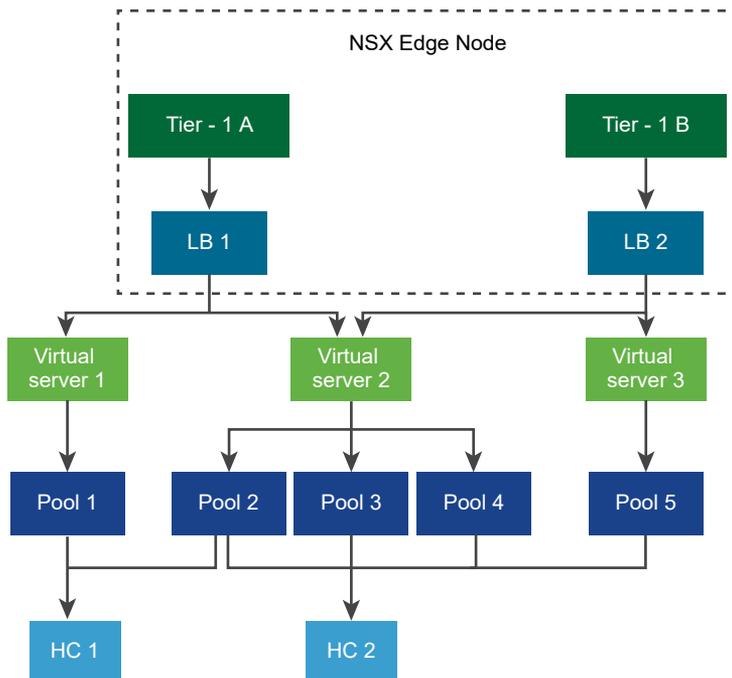
- [Key Load Balancer Concepts](#)
- [Setting Up Load Balancer Components](#)
- [Groups Created for Server Pools and Virtual Servers](#)

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, or 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, or large) determines the number of load balancers that the NSX Edge node can support. Note that in the **Advanced Networking & Security** tab, the term logical router is used to refer to a gateway.

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 the **Networking** tab to configure load balancing, run the following command:

```
GET /policy/api/v1/infra/lb-node-usage?node_path=<node-path>
```

If you use the **Advanced Networking & Security** tab 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*.

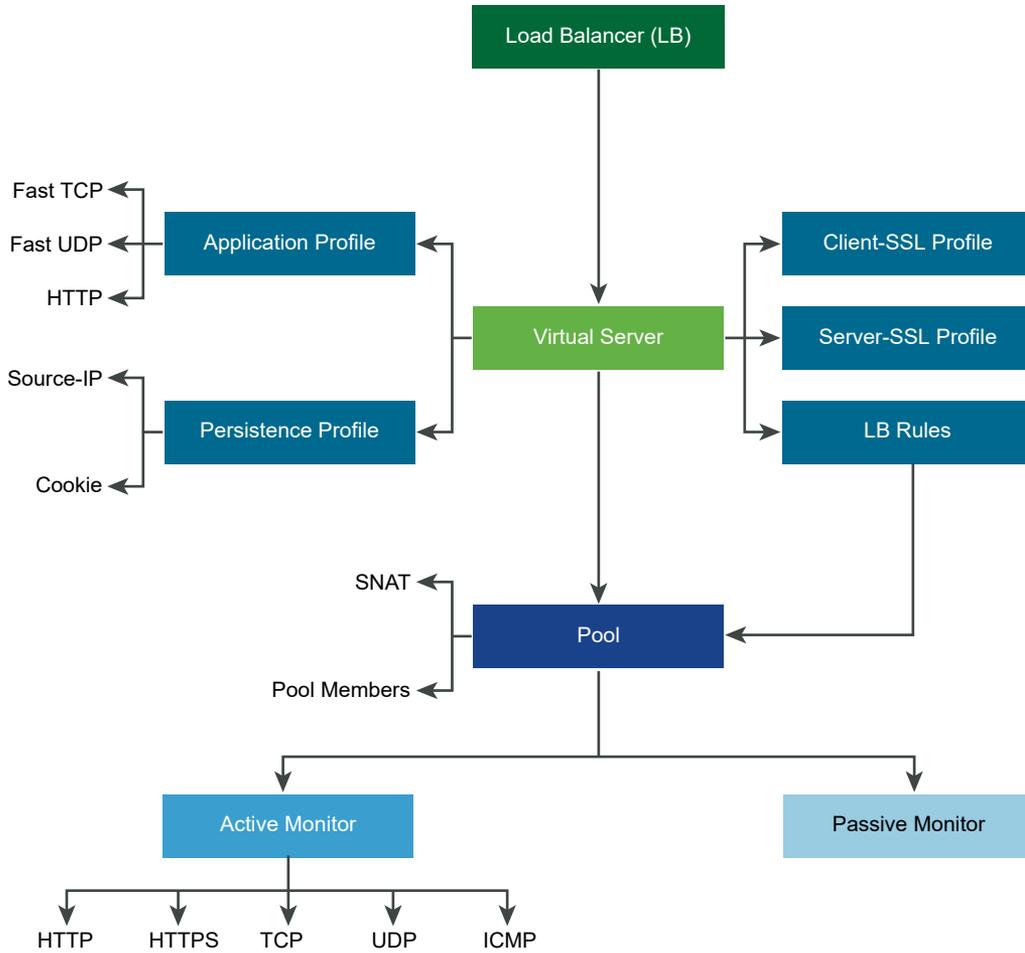
Supported Load Balancer Features

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.

To detect an inactive client or server communication, the load balancer uses the HTTP application profile response timeout feature set to 60 seconds. If the server does not send traffic during the 60 seconds interval, NSX-T Data Center ends the connection on the client and server side.

Note: SSL -Terminate-mode and proxy-mode is not supported in NSX-T Data Center limited export release.

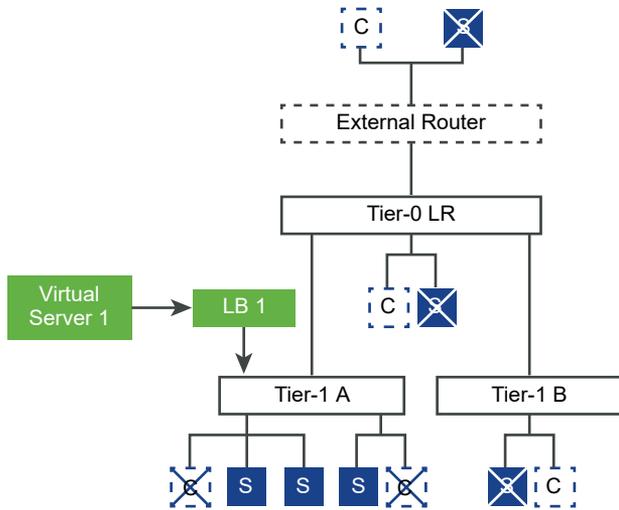


Load Balancer Topologies

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

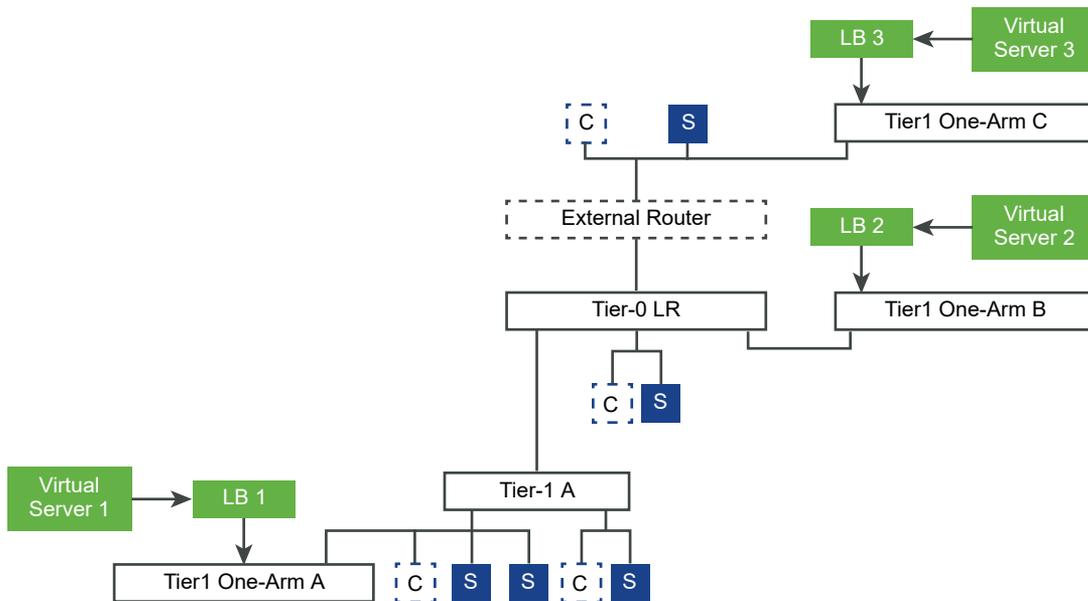
In the inline mode, the load balancer is in the traffic path between the client and the server. Clients and servers should not be connected to overlay segments on the same tier-1 logical router if SNAT on the load balancer is not desired. If clients and servers are connected to overlay segments on the 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.



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:

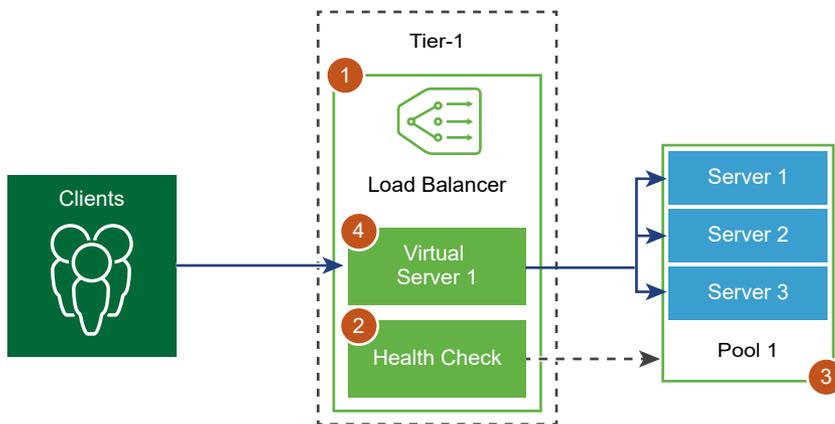
- Ingress
 - DNAT - Firewall - Load Balancer
 - Note: If DNAT is configured with Firewall Bypass, Firewall is skipped but not Load Balancer.
- Egress
 - Load Balancer - Firewall - SNAT

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.

Note In the **Advanced & Security** tab, the term tier-1 logical router is used to refer 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.



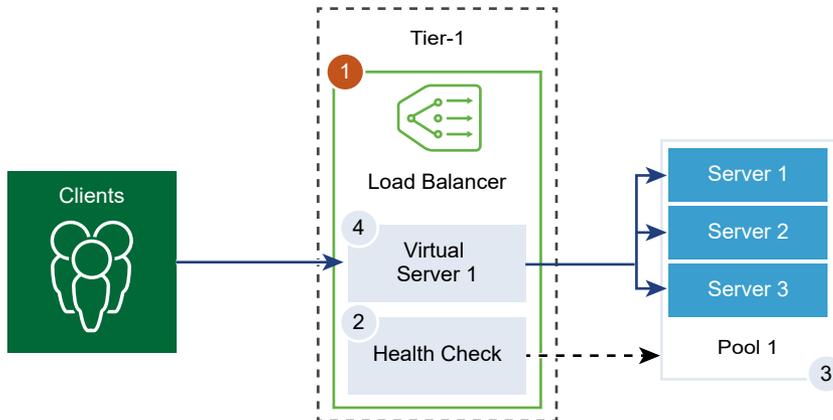
Add Load Balancers

Load balancer is created and attached to the tier-1 gateway.

Note In the **Advanced & Security** tab, the term tier-1 logical router is used to refer to a 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 virtual server size and number of pool members based on your 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

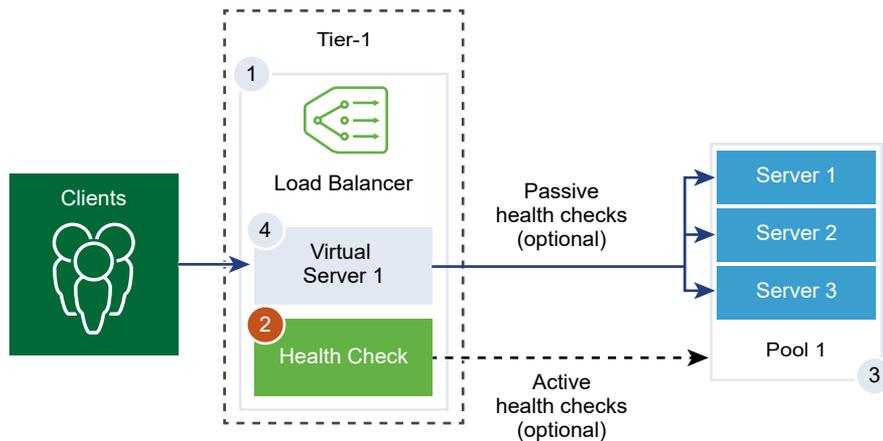
The active health monitor is used to test whether a server is available. The active health monitor uses several types of tests such as sending a basic ping to servers or advanced HTTP requests to monitor an application health.

Note In the **Advanced & Security** tab, the term tier-1 logical router is used to refer to a tier-1 gateway.

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 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.
Monitoring Interval	Set the time in seconds that the monitor sends another connection request to the server.
Timeout Period	Set the number of times the server is tested before it is considered as DOWN.
Fall Count	Set a value when the consecutive failures reach this value, the server is considered temporarily unavailable.
Rise Count	Set a number after this timeout period, the server is tried again for a new connection to see if it is available.
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, 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.

- 6 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.
HTTP Request Version	Select the supported request version from the drop-down menu. You can also accept the default version, HTTP_VERSION_1.
HTTP Response Header	Click Add and enter the HTTP response header name and corresponding value. The default header value is 4000. The maximum header value is 64,000.
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 Select the **HTTPS** protocol.
- 9 Complete step 5.
- 10 Click **Configure**.
- 11 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.
HTTP Request Version	Select the supported request version from the drop-down menu. You can also accept the default version, HTTP_VERSION_1.
HTTP Response Header	Click Add and enter the HTTP response header name and corresponding value. The default header value is 4000. The maximum header value is 64,000.
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.
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.

Option	Description
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.

12 Select the **ICMP** protocol.

13 Complete step 5 and assign the data size in byte of the ICMP health check packet.

14 Select the **TCP** protocol.

15 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.

16 Select the **UDP** protocol.

17 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
Fall 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

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.

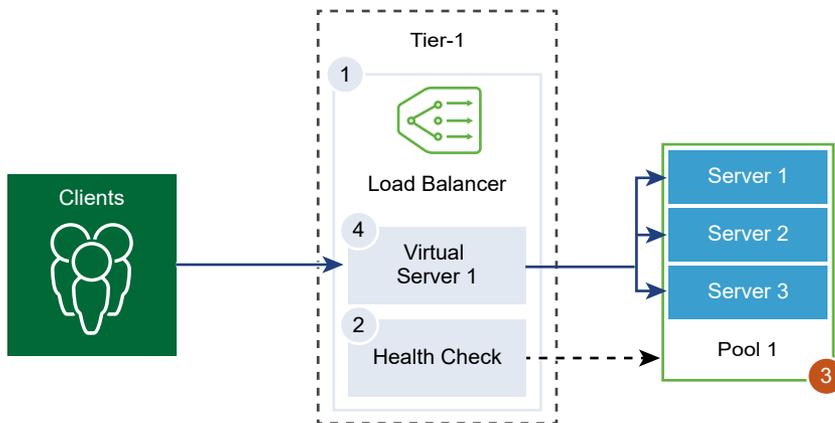
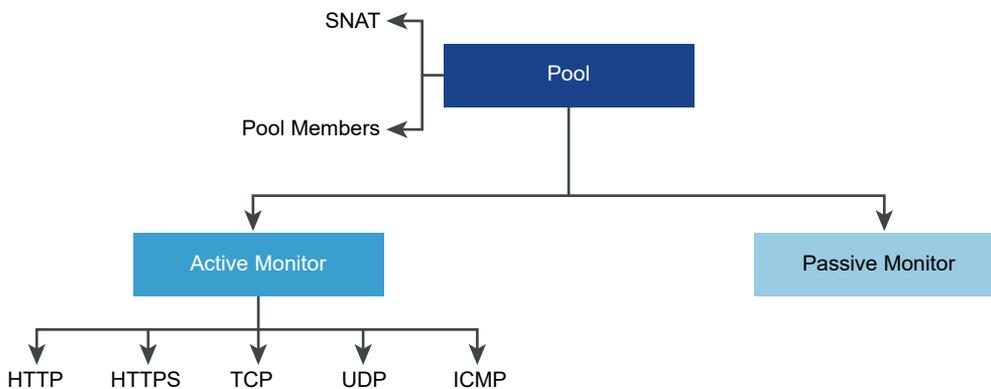


Figure 7-1. Server Pool Parameter Configuration



Prerequisites

- If you use dynamic pool members, a NSGroup must be configured. See [Create an NSGroup](#).
- Verify that a passive health monitors 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 Select the server pool members.

Server pool consists of single or multiple pool members.

Option	Description
Enter individual members	<p>Enter a pool member name, IP address, and a port.</p> <p>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.</p> <p>You can set the server pool admin state. By default, the option is enable when a server pool member is added.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>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.</p>
Select a group	<p>Select a pre-configured group of server pool members.</p> <p>Enter a group name and an optional description.</p> <p>Set the compute member from existing list or create one. You can specify membership criteria, select members of the group, add IP and MAC addresses as group members, and add Active Directory groups. The identity members intersect with the compute member to define membership of the group.</p> <p>Enter tags to make searching easier. You can specify a tag to set a scope of the tag.</p> <p>You can optionally define the maximum group IP address list.</p>

6 Select the active health check monitor for the server pool from the drop-down menu.

The load balancer periodically sends an ICMP ping to the servers to verify health independent of data traffic. You can configure only 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	<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>
Disable	Disable the SNAT translation mode.
IP Pool	<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 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 pool 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 Enter tags to make searching easier.

You can specify a tag to set a scope of the tag.

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.

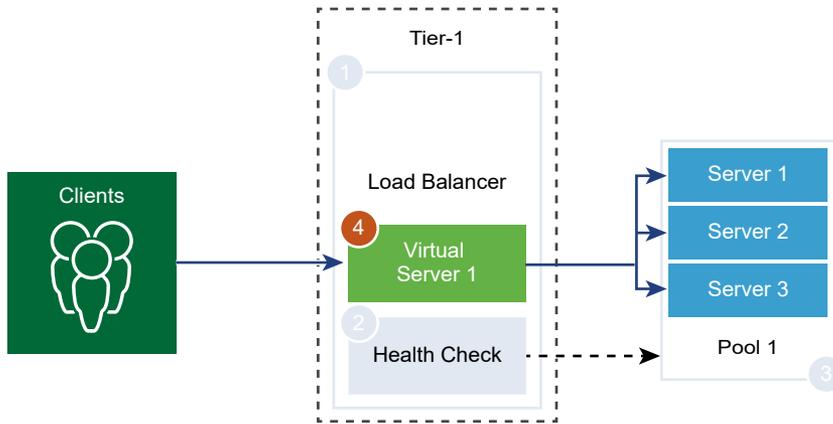
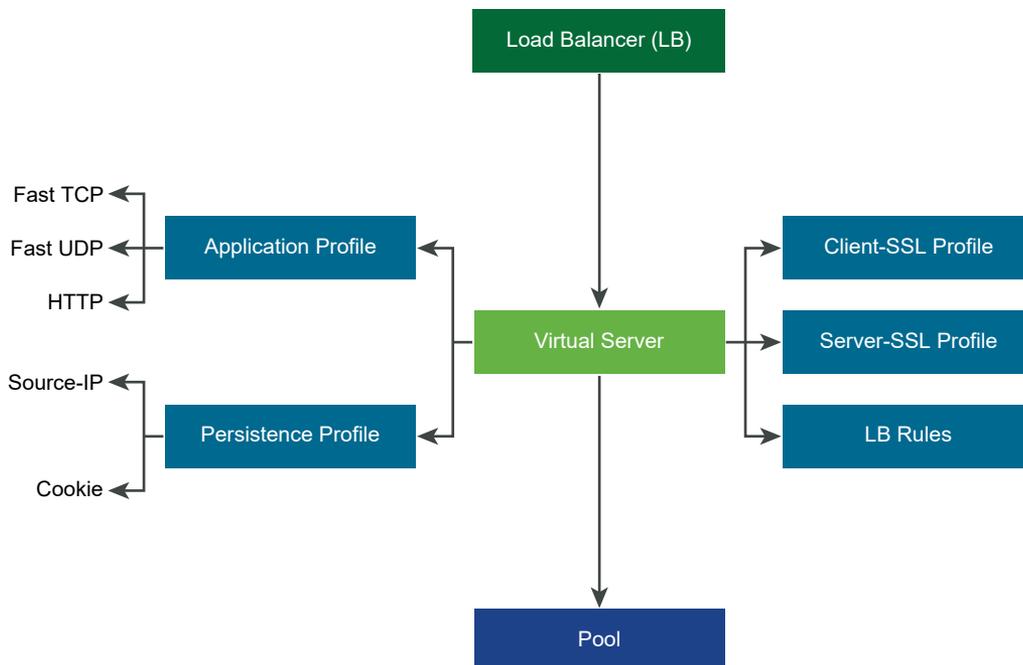


Figure 7-2. Virtual Server Components



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 7-3. Layer 4 TCP and UDP Application Profile

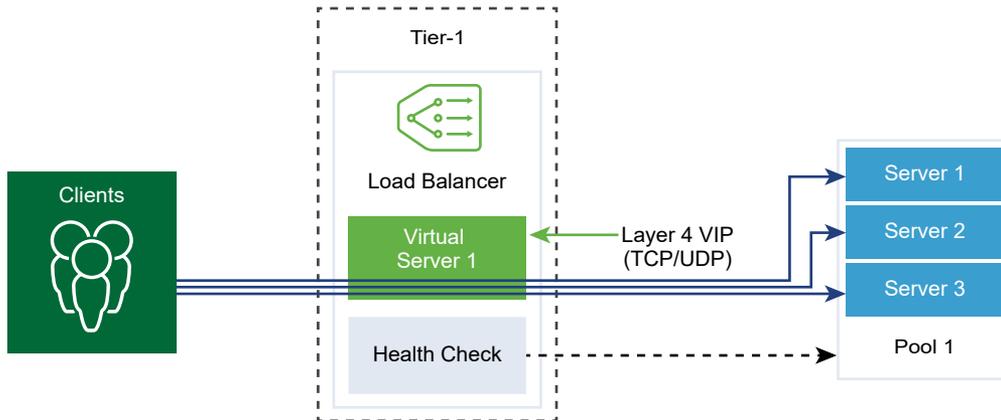
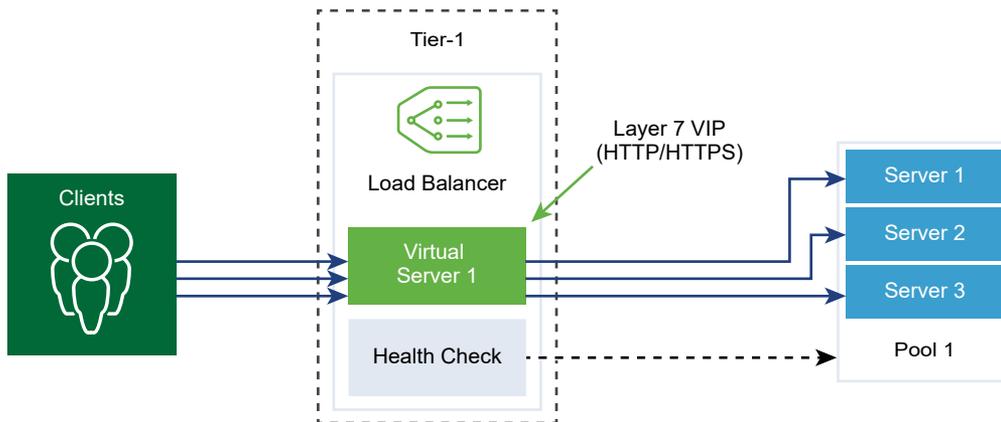


Figure 7-4. Layer 7 HTTPS Application Profile



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.
X-Forwarded-For (XFF)	<ul style="list-style-type: none"> ■ 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. <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>
Request Body Size	<p>Enter value for the maximum size of the buffer used to store the HTTP request body.</p> <p>If the size is not specified, then the request body size is unlimited.</p>
Redirection	<ul style="list-style-type: none"> ■ 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. <ul style="list-style-type: none"> 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. <ul style="list-style-type: none"> 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.

Option	Description
NTLM Authentication	<p>Toggle the button for the load balancer to turn off TCP multiplexing and enable HTTP keep-alive.</p> <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>
Tags	<p>Enter tags to make searching easier.</p> <p>You can specify a tag to set a scope of the tag.</p>

Add a Persistence Profile

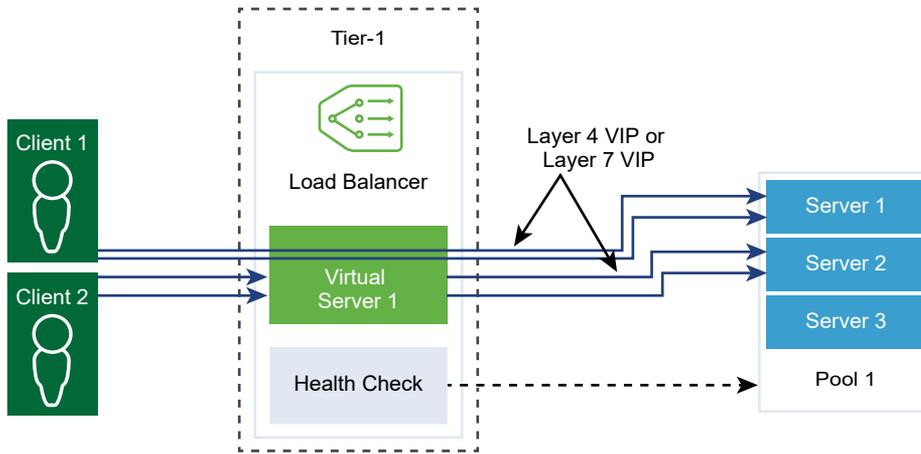
To ensure stability of stateful applications, load balancers implement persistence which directs all related connections to the same server. Different types of persistence are supported to address different types of application needs.

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	<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 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.</p>
Persistence Entry Timeout	<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> <p>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>.</p> <ul style="list-style-type: none"> ■ 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.

Option	Description
Purge Entries When Full	<p>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.</p> <p>When this option is disabled, if the source IP persistence table is full, new client connections are rejected.</p>
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	<p>Enter tags to make searching easier.</p> <p>You can specify a tag to set a scope of the tag.</p>

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	<p>Toggle the button to share persistence across multiple virtual servers that are associated to the same pool members.</p> <p>The Cookie persistence profile inserts a cookie with the format, <i><name>.<profile-id>.<pool-id></i>.</p> <p>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, <i><name>.<virtual_server_id>.<pool_id></i>.</p>
Cookie Mode	<p>Select a mode from the drop-down menu.</p> <ul style="list-style-type: none"> ■ 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	<p>Enter the domain name.</p> <p>HTTP cookie domain can be configured only in the INSERT mode.</p>
Cookie Fallback	<p>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.</p> <p>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.</p>
Cookie Path	<p>Enter the cookie URL path.</p> <p>HTTP cookie path can be set only in the INSERT mode.</p>
Cookie Garbling	<p>Toggle the button to disable encryption.</p> <p>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.</p>
Cookie Type	<p>Select a cookie type from the drop-down menu.</p> <p>Session Cookie - Not stored. Will be lost when the browser is closed.</p> <p>Persistence Cookie - Stored by the browser. Not lost when the browser is closed.</p>

Option	Description
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.
Persistence Entry Timeout	<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> <p>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>.</p> <ul style="list-style-type: none"> ■ 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 7-5. SSL Offloading

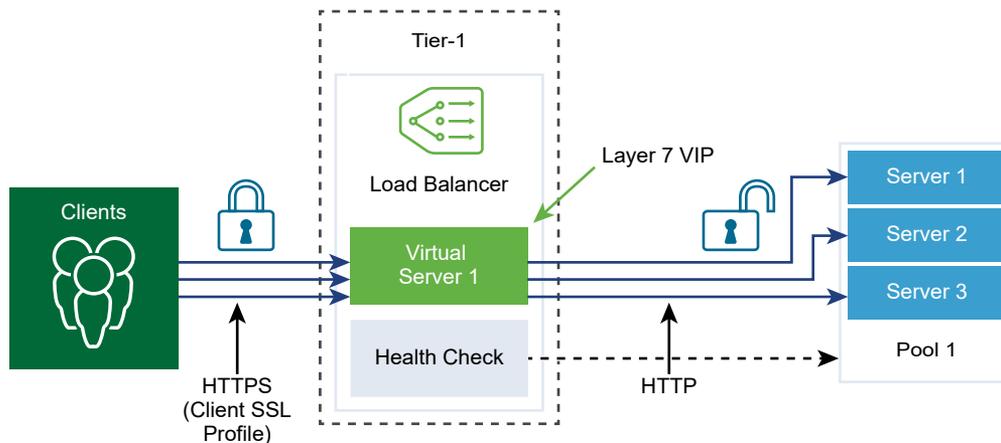
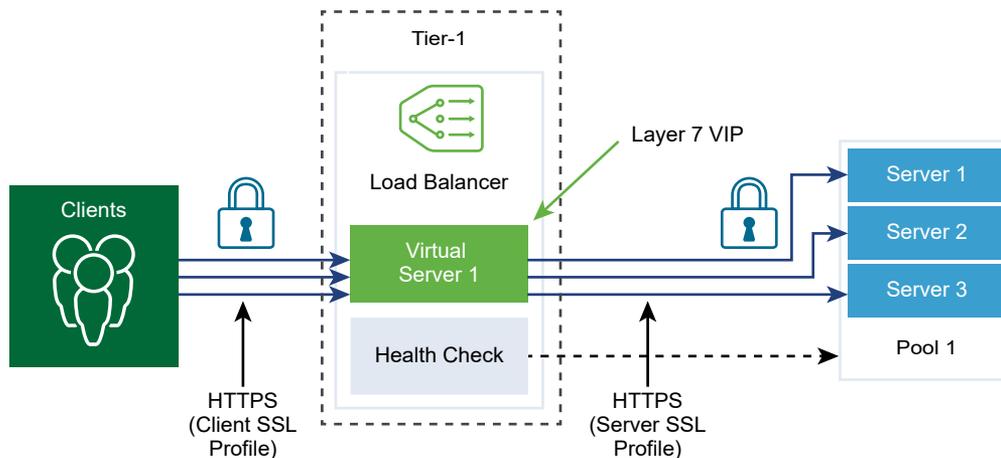


Figure 7-6. End-to-End SSL



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.
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.

Option	Description
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** 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.

Option	Description
Name and Description	Enter a name and a description for the Layer 4 virtual server.
IP Address	Enter the virtual server IP address.
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.
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 the 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 4 virtual server.
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.

4 Select a **L4 UDP** protocol and enter the protocol details.

Option	Description
Name and Description	Enter a name and a description for the Layer 4 virtual server.
IP Address	Enter the virtual server IP address.

Option	Description
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.
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 the 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 4 virtual server.
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.

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.

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.

Note Layer 7 SSL passthrough is supported in NSX-T Data Center 3.0 and later.

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.

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 a **L7 HTTP** protocol 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.
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 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 the 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.
Tags	Enter tags to make searching easier. 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 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:sensitive)"` instead use `"Case ((?i:sensitive))"`.
- 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"
- (NSX-T Data Center 2.5.0 only) `$_ssl_client_cert` - returns the client certificate in the PEM format for an established SSL connection, with each line except the first prepended with the tab character.
- (NSX-T Data Center 2.5.1 and later) `$_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

Prerequisites

Verify 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, click **Set > Add Rule** 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. 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
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	Match an IP source or destination address. 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.

Actions	Description
HTTP Request URI Rewrite	Modify an URI. http_request.uri - URI (without query arguments) to write http_request.uri_args - URI query arguments to write
HTTP Request Header Rewrite	Modify value of an HTTP header. http_request.header_name - header name http_request.header_value - value to write
HTTP Request Header Delete	Delete HTTP header. http_request.header_delete - header name http_request.header_delete - value to write

- 3 Click **Request Forwarding > Add Rule** 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. http_request.method - value to match
HTTP Request URI	Match an HTTP request URI. http_request.uri - 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
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	Match an IP source or destination address. 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.
<hr/>	
Action	Description
HTTP Reject	Reject a request, for example, by setting status to 5xx. http_forward.reply_status - HTTP status code used to reject http_forward.reply_message - HTTP rejection message
HTTP Redirect	Redirect a request. Status code must be set to 3xx. http_forward.redirect_status - HTTP status code for redirect 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. 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 very long, hashing the variable will ensure 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 very 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.
Reply Status	Shows the status of the reply.
Reply Message	Server responds with a reply message that contains confirmed addresses and configuration.

- 4 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	Match any HTTP response header. 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

Supported Match Condition	Description
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	Match an IP source or destination address. 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.

Action	Description
HTTP Response Header Rewrite	Modify the value of an HTTP response header. http_response.header_name - header name http_response.header_value - value to write
HTTP Response Header Delete	Delete HTTP header. 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 very long, hashing the variable will ensure 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 very 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.

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 no LB-SNAT Automap: T1-Uplink IP 100.64.x.y and T1-ServiceInterface IP
- Pool configured with no 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*.

Forwarding Policies



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 22 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.

This chapter includes the following topics:

- [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)

9

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*.

This chapter includes the following topics:

- [Add a DNS Zone](#)
- [Add a DNS Forwarder Service](#)
- [Add a DHCP Server](#)
- [Configure a DHCP Relay Server for a Tier-0 or Tier-1 Gateway](#)
- [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 Address 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 Address 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 Server

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.

Note The DHCP server that is created using this procedure is not supported on a VLAN-backed segment. You must use the DHCP feature under **Advanced Networking & Security** to create a DHCP server that is supported on a VLAN-backed 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 > IP Address Management > DHCP**.
- 3 Click **Add Server**.
- 4 Select **DHCP Server** as the server type.
- 5 Enter a name for the server.
- 6 Enter the server's IP address in CIDR format.

This step will create a two logical ports (one for a logical interface and one for the DHCP server itself) and connect the DHCP server to a specific DHCP logical switch. This interface will appear on the tier-0 or tier-1 gateway as a connected interface, so make sure you choose a non-overlapping subnet for the tier-1 or tier-0 gateway that you want to assign the DHCP server to. You can specify `<IP address>/30` for this purpose. The subnet range used here does not get advertised to the connected tier-0 gateway, but does appear in the tier-1 gateway's forwarding table.

- 7 Enter a lease time.
- 8 Select an NSX Edge cluster.
- 9 Click **Save**.
- 10 To assign a DHCP server to a tier-0 or tier-1 gateway:
 - a Navigate to **Networking > Tier-0 Gateways** or **Networking > Tier-1 Gateways**.
 - b Edit an existing gateway.
 - c In the **IP Address Management** field, click **No IP Allocation**.
 - d Select **DHCP Local Server** from the Type dropdown list.
 - e Select a DHCP server.

- f Click **Save**.
 - g Click **Save**.
- 11 To assign a DHCP server to a segment:
- a Navigate to **Networking > Segments**.
 - b Add or edit a segment.
The segment must be associated with a tier-0 or tier-1 gateway.
 - c Click **Set Subnets** if you are adding a new segment, or click the number under **Subnets** to add or modify a subnet.
 - d Enter the appropriate DHCP ranges.
 - e Click **Apply**.
 - f Click **Save**.

Configure a DHCP Relay Server for a Tier-0 or Tier-1 Gateway

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 a DHCP relay server to relay DHCP traffic to external DHCP servers.

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 Management > DHCP**.
- 3 Click **Add Server**.
- 4 Select **DHCP Relay** as the server type.
- 5 Enter a name for the relay server.
- 6 Enter one or more IP addresses for the server.
- 7 Click **Save**.
- 8 Go to **Networking > Tier-0 Gateways** or **Networking > Tier-1 Gateways** to configure a DHCP relay server for a gateway.
- 9 Edit the appropriate gateway.
- 10 In the **IP Address Management** field, click **No IP Allocation** for a tier-0 gateway or **No IP Allocation Set** for a tier-1 gateway.
- 11 In the **Type** field, select **DHCP Relay**.
- 12 In the **DHCP Relay** field, select the DHCP relay server you created earlier.

- 13 Click **Save**.
- 14 For each segment connected to the gateway that will use this DHCP relay service, you must specify DHCP ranges for the relay to function.
 - a Go to **Networking > Segments**.
 - b Add or edit a segment.
 - c Click **Set Subnets** if you are adding a new segment, or click the number under **Subnets** to modify a subnet.
 - d Specify one or more DHCP ranges.

This is required for the relay to function.
 - e Click **Apply**.
 - f Click **Save**.

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 Address 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 Sunet > 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**.

Add an IP Address Block

You can configure IP address blocks for use by other components.

Note You can also add an IP address block by navigating to **Advanced Networking & Security > Networking > IPAM**.

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 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**.

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.

This chapter includes the following topics:

- [Security Configuration Overview](#)
- [Security Terminology](#)
- [Identity Firewall](#)
- [Layer 7 Context Profile](#)
- [Distributed Firewall](#)
- [East-West Network Security - Chaining Third-party Services](#)
- [Configuring a Gateway Firewall](#)
- [North-South Network Security - Inserting Third-party Service](#)
- [Endpoint Protection](#)
- [Security Profiles](#)

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 specific rules.

Security Terminology

The following terms are used throughout distributed firewall.

Table 10-1. 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 on to in to 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 Wall 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 left 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.

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 - Version 11

- VMCI Driver
- NSX File Introspection Driver
- NSX Network Introspection Driver

Layer 7 Context Profile

Layer 7 Ap 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 whitelisting or blacklisting. Currently a predefined list of domains is supported. FQDN 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.

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

- 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.
-

When a context-profile has been used in a rule, any traffic coming in from a virtual machine is matched against the rule-table based on 5-tuple. If the rule matches the flow also includes a Layer 7 context profile, that packet is redirected to a user-space component called the vDPI engine. A few subsequent packets are punted to that vDPI engine for each flow, and after it has determined the App Id, this 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 on 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 process in the kernel, and matched against the connection table. For fully matched DROP rule a reject packet is generated. Logs generated by the firewall will include the Layer 7 App Id and applicable URL, if that flow was punted to DPI.

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 whitelisting and blacklisting. 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. 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.

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 above 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.

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. 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

Attribute (App Id)	Description	Type
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
CTRXCOTO	Hosting Citrix GoToMeeting, or similar sessions based on the GoToMeeting platform. Includes voice, video, and limited crowd management functions	Collaboration
CTRICA	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
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

Attribute (App Id)	Description	Type
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
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

Attribute (App Id)	Description	Type
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
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. Rules are evaluated top down, and left to right.

Table 10-2. Distributed Firewall Rule Categories

Category	Description
Ethernet	Used for Layer 2 based rules
Emergency	Used for quarantine and allow rules
Infrastructure	Define access to shared services. Global rules - AD, DNS, NTP, DHCP, Backup, Management Servers
Environment	Rules between zones - production vs development, inter business unit rules
Application	Rules between applications, application tiers, or the rules between 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 aurodraft 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 (DFW) monitors all the East-West traffic on your virtual machines.

Prerequisites

Guest VMs to be DFW-protected must have their vNIC connected to an N-VDS logical switch associated with a transport zone.

If you are creating rules for Identity Firewall, first create a group with Active Directory members. IDFW only supports TCP-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.

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 above a Layer 7 any/any rule will not be executed.

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 Enable Distributed Firewall by selecting **Actions > General Settings**, and toggling the Distributed Firewall Status. Click **Save**.
- 4 Ensure that you are in the correct pre-defined category, and click **Add Policy**. For more about categories, see [Distributed Firewall](#) .
- 5 Enter a **Name** for the new policy section.

- 6 (Optional) To configure the following policy settings, click the gear icon:

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

- 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. Groups with Active Directory members can be used for the source field 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](#).

- 11 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.
- 12 In the **Services** column, click the edit icon and select services. The service matches **Any** if not defined.
- 13 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, and TFTP), one context profile is supported per rule.

- 14 Click **Apply** to apply the context profile to the rule.
- 15 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.

- 16 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.

- 17 Click the status toggle button to enable or disable the rule.

18 (Optional) 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> 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.

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

20 On each rule, click the **Info** icon to view the rule ID number, and where it is enforced.

This icon is grayed out until you publish the rule. You can also specify a rule ID when you click the filter icon to display only policies and rules that satisfy the filter criteria.

21 The realization status API has been enhanced at a security policy level to provide additional realization status information. This can be achieved by specifying the query parameter `include_enforced_status=true` along with the `intent_path`. Make the following API call.

```
GET https://<nsx>/policy/api/v1/infra/realized-state/status?intent_path=/
infra/domains/default/security-policies/<security-policy-
id>&include_enforced_status=true
```

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->ff02:0:0:0:0:0:1:2/547
```

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 #)
- source IP address/source port>destination IP address/destination port
- 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 #)
- source IP address/source port>destination IP address/destination port
- APP_XXX is the discovered application

Select a Default Connectivity Strategy

You can select a default connectivity strategy to enforce your security model.

The default connectivity strategy creates either an allow-all (blacklist) or deny-all (whitelist) firewall policy on top of the other firewall rules you create, instead of having to modifying individual rules. To set a default connectivity strategy, go to **Distributed Firewall**. At the top of the page, click the connectivity status to select another option.

Firewall policy and rules must have already been created to change the default selected connectivity strategy, and have it to go into effect immediately. If no policy or rules are created, the default connectivity strategy remains until a policy and rules are created.

The following options are available:

- **Blacklist (with or without logging):** This is the default option and creates an allow-all rule on the DFW.
- **Whitelist (with or without logging):** Creates a deny-all traffic firewall rule. Only communication from sites or applications that have been defined in firewall rules is allowed, and all other communication is denied access, including DHCP traffic.
- **None:** Select this option to disable both blacklisting or whitelisting of firewall rules. This is useful if you have a set of rules already configured using previous versions of NSX-T Data Center.

Manage a Firewall Exclusion List

Firewall exclusion lists are made of groups that can be excluded from a firewall rule based on group membership.

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 AD 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 Manager and 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`.

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.

Option	Description
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.
Applied To	Select DFW or a group as required.
Action	Select Allow , Drop , or Reject .

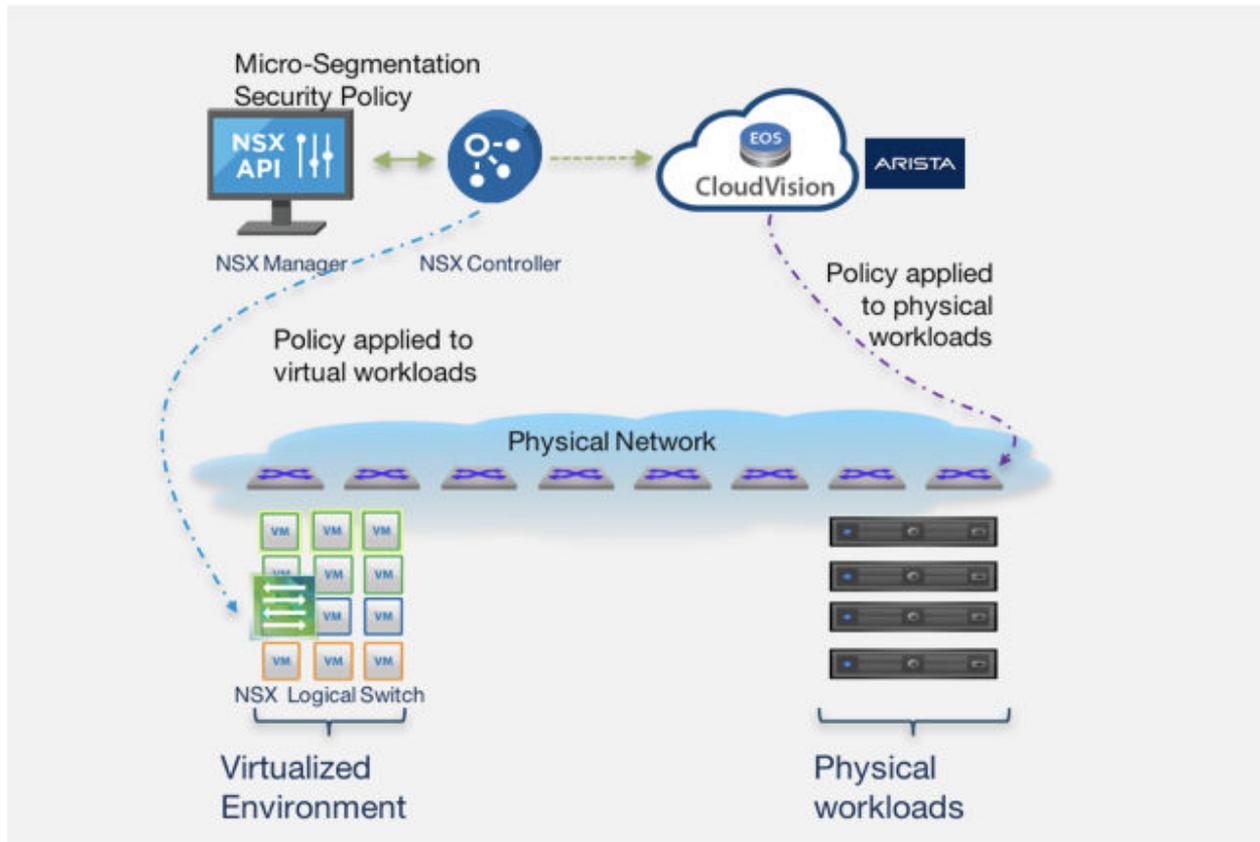
- 9 Click **Publish**.

Extending Security Policies to Physical Workloads

NSX-T Data Center can act as a single point of administration for both virtual and physical workloads.

Starting in NSX-T Data Center 2.5.1, integration with Arista CloudVision eXchange (CVX) is supported. 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.

Starting in NSX-T Data Center 2.5.1, integration with Arista EOS 4.22.1FX-PCS and later is supported.



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 dvpkg'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 Arista CVX to interact with NSX-T Manager

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:

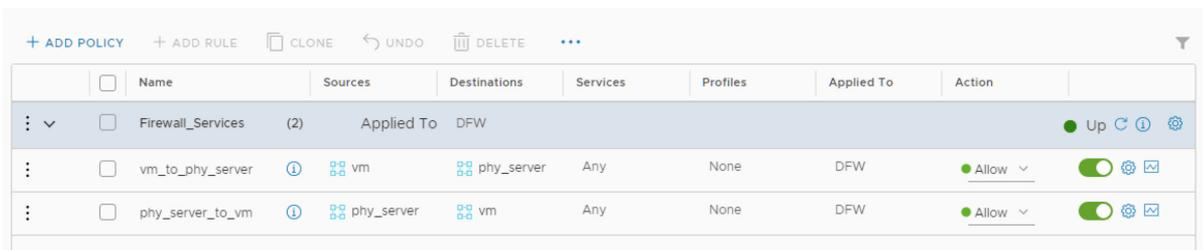
```
show running-config section tag
```

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 10 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](#)

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.

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.

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.
- **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 10-3. 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.
- 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 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 > East West Security > Network Introspection > 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 Select **Security > East West Security > Network Introspection > Service Chain > Add Chain**.
- 3 Enter the service chain name.
- 4 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.
- 5 To set the forward path, click the **Set Forward Path field** and click **Add Profile in Sequence**.
- 6 Add the first profile in the service chain and click **Add**.
- 7 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.
- 8 Click **Save** to finish adding a forward path for the service chain.
- 9 In the Reverse Path column, select **Inverse Forward Path** for the service plane to use the service profile you set for the forward path.
- 10 To set a new service profile for the reverse path, click **Set Reverse Path** and add a service profile.
- 11 Click **Save** to finish adding a reverse path for the service chain.
- 12 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.
- 13 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](#).

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 does not forward it to the regular distributed firewall, rather it redirects that traffic along 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 **Security > East West Security > Network Introspection > Rules > Add Policy.**
A policy section is similar to a firewall section where you define rules that determine how traffics flows.
- 3 Select a service chain.
- 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.
 - a Define membership criteria using one of these entities:
 - Virtual Machine
 - Logical Switch
 - Logical Port
 - IP Set
 - b Define static member list using one of these entities:
 - Group
 - Segment
 - Segment Port
 - Virtual Network Interface
 - 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:
 - Select **DFW** to apply the rule to all virtual NICs attached to the logical switch.
 - Select **VM groups** to apply the rule on virtual NICs of member VMs of the group. Members can be selected from a static list or based on dynamic criteria. The supported NSX-T Data Center objects are: Virtual Machine, Logical Switch, Logical Port, IP Set and so on.
- 10 In the Action field, select **Redirect** to redirect traffic along the service chain or **Do Not Redirect** not to apply network introspection on the traffic.
- 11 Click **Publish**.
- 12 To revert a published rule, select a rule and click **Revert**.
- 13 To add a policy, click **+ Add Policy**.
- 14 To clone a policy or a rule, select the policy or rule and click **Clone**.
- 15 To enable a rule, enable the Enable/Disable icon or select the rule and from the menu click **Enable > Enable Rule**.
- 16 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.

Configuring a 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 10-4. 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 [Configuring a 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, tag, and notes.

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 .
Tag	Tag 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.

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
```

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 10-5. 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 Manager UI
Configure Traffic Redirection	Administrator	API and NSX Manager UI

Deploy a Service for North-South Traffic Introspection

After you register a service, you must deploy an instance of the service 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 **Advanced Networking & Security > Partner Services > Service Instances > Catalog**.
- 3 The Catalog tab displays the registered services.
- 4 Select the service displayed in OVF form factor and click **Deploy** to begin deployment of the service instance.
- 5 In the Partner Service Insertion window, click **Proceed**.
- 6 In the Partner Service window, enter the details.

Table 10-6. Partner Service Details

Field	Description
Instance Name	Enter a name to identify the service instance.
Description	Description about the service instance.
Partner Service	Select the partner service registered with NSX-T Data Center.
Deployment Specification	Select the form factor to deploy.
Logical Router	Select the tier-0 logical router where the service instance must be deployed.

- 7 Click **Next**.
- 8 In the Instance Configuration window, enter the details.

Table 10-7. Service Instance Details

Field	Description
Deployment Mode	Select Standalone to deploy a single service instance at the tier-0 logical router. Select High Availability to deploy a couple of service instances in active-standby mode at the tier-0 logical router.
Failure Policy	Select Allow or Block .
Service Instance IP Address	Enter the IP address to be used by the service instance.
Gateway	Enter the gateway address.
Subnet Mask	Enter the subnet mask.

Table 10-7. Service Instance Details (continued)

Field	Description
Network ID	Enter the network ID of the logical switch where you want to connect the management network.
Compute Manager	Select the registered vCenter Server.
Resource Pool	Select the resource pool that provides resources to deploy the service instance.
Datastore	Select the repository to store service instance data.

9 Click **Next**.

10 In the Advanced Configuration window, enter the details.

Table 10-8.

Field	Description
Deployment Template	Select the template to be used during deployment of the service instance.
License	Enter the license of the template.

11 Click **Finish**.

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 logical router.

Alternatively, go to the vCenter Server and verify the deployment status.

What to do next

Configure rules to redirect traffic to the service instance deployed at the tier-0 router. See [Configure Traffic Redirection](#)

Configure Traffic Redirection

After you deploy a service instance, configure the type of traffic that the router redirects to the service. Configuring traffic redirection is similar to configuring a firewall.

For information about configuring a firewall, see [Firewall Sections and Firewall Rules](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Partner Services > Service Instances** .

- 3 Click the service instance.
- 4 Click the **Traffic Redirection** tab.
- 5 To add a section, select an existing section and click **Add Section**.
 - ◆ From the menu, select **Add Section Above** or **Add Section Below**.

A new section is created. The traffic type to be redirected is set to **L3 Redirect**, service is of the type **Stateless**, the **Applied To** field is associated to a Tier-0 logical router that is configured on the host. After you define rules, the **Rules** field is auto-populated.
- 6 Click **Publish** to persist configuration details of the section.
- 7 To add a rule within that section, select the section and click **Add Rule**.
- 8 In the rule row, enter the following details:
 - a Enter rule name.
 - b Enter the source and destination of L3 traffic. The partner service VM introspects traffic flowing in from the source before redirecting it to the destination VM.
 - c In the **Applied To** field, select the uplink of Tier-0 router.
 - d In the **Action** field, select **Redirect** if traffic needs to be introspected by the service VMs or select **Don't Redirect** if traffic does not need to be introspected for north-south introspection.
- 9 Each rule can be enabled individually. After you enable a rule, it is applied to the traffic that matches the rule.
- 10 Click Advanced Settings to configure the traffic direction and to enable logging.
- 11 At the end of a section containing rules, click **Publish** to persist the rules in the section or click **Revert** to cancel the operation.

Results

The traffic is sent to network introspection rules where policy rules are applied to the traffic.

What to do next

See [Add Redirection Rules for North-South Traffic](#).

Add Redirection Rules for North-South Traffic

Use the **Advanced Networking and Security** UI to set up north-south redirection rules. Traffic redirection happens only for services inserted at the Tier-0 router.

Follow instructions at [Configure Traffic Redirection](#).

Prerequisites

- Register and deploy third-party services on NSX-T.
- Configure Tier-0 router.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 **Security > North South Firewall > 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** to the service instance that is registered with NSX-T 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:
 - Select **DFW** to apply the rule to all virtual NICs attached to the logical switch.
 - Select **VM groups** to apply the rule on virtual NICs of member VMs of the group. Members can be selected from a static list or based on dynamic criteria. The supported NSX-T Data Center objects are: Virtual Machine, Logical Switch, Logical Port, IP Set and so on.
- 10 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.
- 11 Click **Publish**.
- 12 To revert a published rule, select a rule and click **Revert**.
- 13 To add a policy, click **+ Add Policy**.
- 14 To clone a policy or a rule, select the policy or rule and click **Clone**.
- 15 To enable a rule, enable the Enable/Disable icon or select the rule and from the menu click **Enable > Enable Rule**.
- 16 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.

Monitor Traffic Redirection

After you deploy a service instance and configure traffic redirection, you can monitor the amount of traffic that goes into and out of the service instance.

Procedure

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

2 Select **Advanced Networking & Security > Partner Services > Service Instances**.

3 Click the name of a service instance.

The **Overview** tab shows the configuration and status of the service instance.

4 Click the **Statistics** tab.

Information about the number of packets and the amount of data that go into and out of the service instance is displayed.

5 Click **Refresh** to update the statistics..

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.

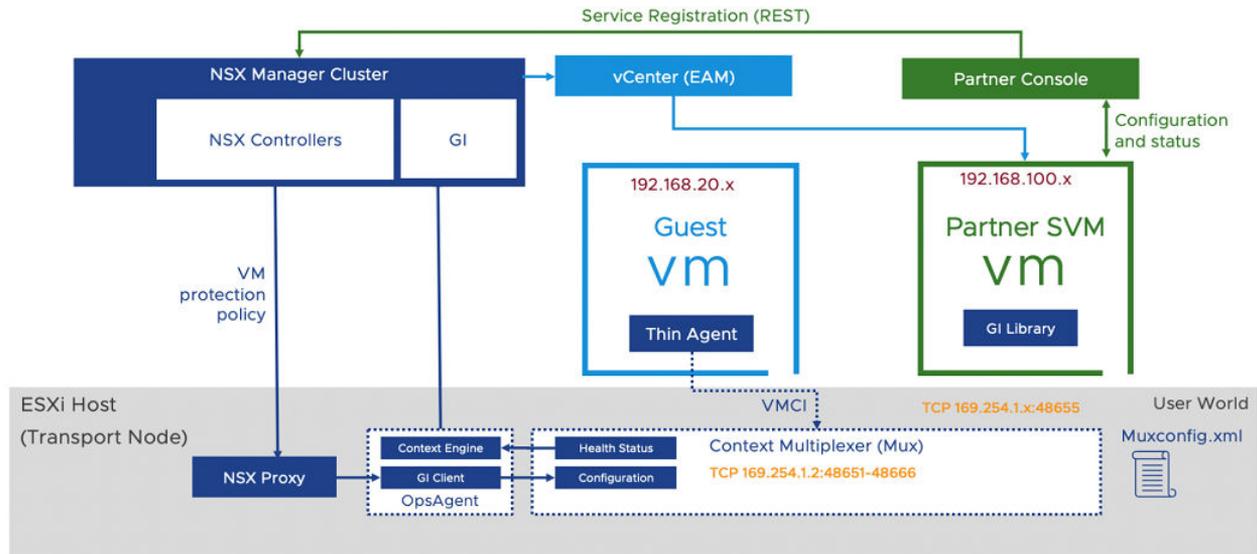
The benefits of using the guest introspection platform to protect guest VM endpoints:

- **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.

Guest Introspection Architecture

Understand the architecture of service insertion and guest introspection components in NSX-T Data Center.

Figure 10-1. Guest Introspection 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 10-2. 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 Linux 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/rhel7/x86_64
enabled = 1
gpgcheck = 1
metadata_expire = 86400
ui_repoid_vars = basearch
```

3 Install the package.

```
yum install vmware-nsx-gi-file
```

4 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
```

5 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/centos7/x86_64
enabled = 1
gpgcheck = 1
metadata_expire = 86400
ui_repoid_vars = basearch
```

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/
```

- 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**.

- From the Partner Service drop-down menu, select the partner service to view details related to service instances.

Table 10-9.

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	<p>When the service instance is deployed, the health status is <i>Ready</i>. To bring the health status from <i>Ready</i> to <i>Up</i>, make the required configuration changes. See Bring up Service Instance.</p> <p>After the following parameters are successfully realized by NSX-T Data Center, the health status changes from <i>Ready</i> to <i>Up</i>.</p> <ul style="list-style-type: none"> ■ 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

- From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- Select **System > Service Deployments > Service Instances**.
- From the Partner Service drop-down menu, select the partner service to view details related to service instances.
- 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.

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%

- a **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.
- b **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.

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 10-10. 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 10-10. 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 10-10. 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 10-10. 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 on.	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 10-10. ESXi Agency Manager Issues (continued)

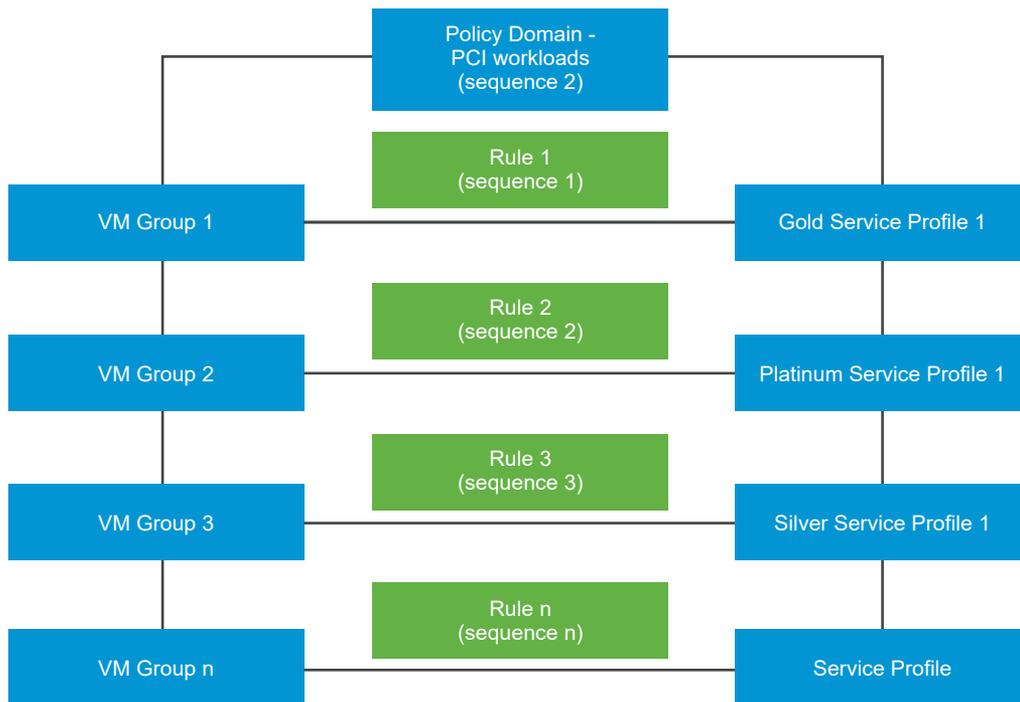
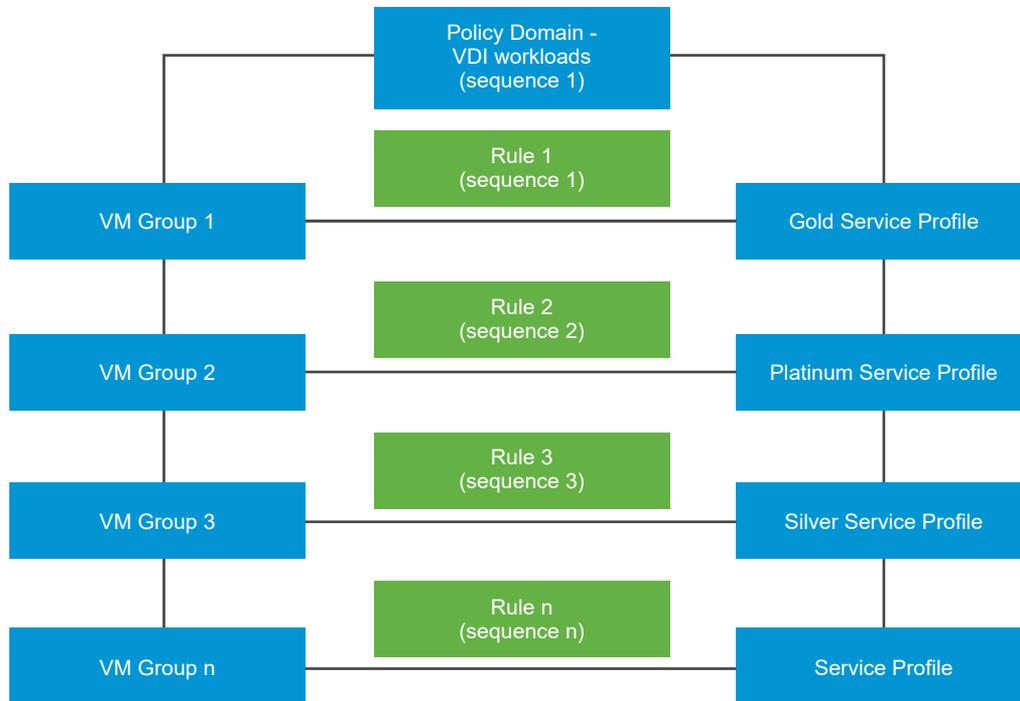
VM Wrong Resource Pool	VM Issue	An agent virtual machine is expected to be located in a designated agent virtual machine resource pool, but is found in a different resource pool.	Click Resolve . ESXi Agency Manager (EAM) service attempts to place the agent virtual machine to a designated resource pool.
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 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 10-11. 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> <p>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.</p>
Create a service configuration.	<pre>POST /api/v1/service-configs</pre> <p>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.</p>
Delete a service configuration.	<pre>DELETE /api/v1/service-configs/ <config-set-id></pre> <p>The service configuration API deletes the configuration applied to the VM group.</p>
Get details of a specific configuration.	<pre>GET /api/v1/service-configs/ <config-set-id></pre> <p>Get details of a specific configuration</p>
Update a service configuration.	<pre>PUT /api/v1/service-configs/ <config-set-id></pre> <p>Update a service configuration.</p>
Get effective profiles.	<pre>GET /api/v1/service-configs/ effective-profiles?resource_id=<resource-id> &resource_type=<resource-type></pre> <p>The service configuration API returns only that profile which is applied to a particular VM group.</p>

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, GI 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 10-12. Resolve policy conflicts

Scenario	Expected Endpoint Protection Flow	Resolution
<p>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.</p>	<p>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.</p> <p>Consider VM 1 is part of Group 1 and Group 2.</p> <ul style="list-style-type: none"> ■ 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 <p>Endpoint protection policy runs the Gold service profile on VM 1 but does not run Platinum service profile on VM1.</p>	<p>Change the Sequence Number of Rule 2 such that it runs before Rule 1.</p> <ul style="list-style-type: none"> ■ On the NSX-T Policy Manager UI, drag the Rule 2 before Rule 1 on the rule list. ■ Using NSX-T Policy Manager API, manually add a higher sequence number for Rule 2.
<p>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.</p>	<p>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.</p> <p>Consider VM 1 is part of Group 1 and Group 2.</p> <p>Rule 1: Group 1 (by OS name) is applied Gold (service profile)</p> <p>Rule 2: Group 2 (by tag) is applied Gold (service profile)</p>	<ul style="list-style-type: none"> ■ 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 10-13. 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 10-14. Troubleshoot Health Status

Health Status Attribute	Resolution
Solution Status is <code>Down</code> Or Not available.	<ol style="list-style-type: none"> 1 Verify that service deployment status is <code>Up</code> (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 <code>Down</code> .	<ol style="list-style-type: none"> 1 Verify that service deployment status is <code>Up</code> (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 <code>Unavailable</code> .	<ol style="list-style-type: none"> 1 Verify that service deployment status is <code>Up</code> (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 <code>Unavailable</code> .	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.

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.
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.

TCP Variables	Description
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.
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.

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. 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 profile timer seen when using API and UI. GFW default session timers are optimized for North-South traffic, and are lower by default.
- FW 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 & 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

Timer Property	Edge Default (secs)	Minimum (secs)	Maximum (secs)
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 10-15. 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.
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.

Table 10-15. Parameters for Firewall and Edge Gateway Profiles (continued)

Parameter	Minimum and maximum values	Default	
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	<p>This field captures the Time to live for the DNS cache entry in seconds. You have the following options:</p> <p>TTL 0 - cached entry never expires.</p> <p>TTL 1 to 3599 - invalid</p> <p>TTL 3600 to 864000 – valid</p> <p>TTL left empty – automatic TTL, set from the DNS response packet.</p> <p>Note DNS Security Profile has a default DNS cache timeout of 24 hours.</p>
Applied To	<p>You can select a group based on any criteria to apply the DNS security profile to.</p> <p>Note Only one DNS server profile is applied to a VM.</p>
Tags	<p>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.</p>

- 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.

Inventory

11

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

This chapter includes the following topics:

- [Add a Service](#)
- [Add a Group](#)
- [Add a Context Profile](#)

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 New Service**.
- 4 Enter a name.
- 5 Click **Set Service Entries**. Click **Add New Service Entry**.
- 6 For a new service, select a type of service, and specify additional properties.
The available types are **IP**, **IGMP**, **ICMPv4**, **ICMPv6**, **ALG**, **TCP**, **UDP** and **Ether**.
- 7 Click **Save**.
- 8 (Optional) Add one or more tags.
- 9 (Optional) Enter a description.
- 10 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. Groups based on dynamic or logical objects cannot be used in the **Applied to** field of distributed firewall rules.

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'.

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.

NSX Cloud Note If 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.

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.

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.

Procedure

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

2 Click **Add Group**.

3 Enter a group name.

4 (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** - can specify a tag and optional scope.
- **Segment** - 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 does not equal a particular string.
- **IP Set** - can specify a tag, and optional scope.

5 (Optional) Click **Members** to select members.

The available member types are:

- **Group**
- **Segment**
- **Segment Port**
- **Virtual Network Interface**
- **Virtual Machine**

6 (Optional) Click **IP/MAC Addresses** to add IP and MAC addresses as group members.

IPv4, IPv6, and multicast addresses are supported.

7 (Optional) Click **AD Groups** to add Active Directory Groups. Groups with Active Directory members can be used in the source field of a distributed firewall rule for Identity Firewall. Groups can contain both AD and compute members.

8 (Optional) Enter a description and tag.

9 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`.

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. Gateway firewall rules do not support the use of FQDN attributes, or other sub attributes.
-

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** 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).

There are multiple ways to monitor the NSX-T environment as well as network traffic.

This chapter includes the following topics:

- Add a Firewall IPFIX Profile
- Add a Switch IPFIX Profile
- Add an IPFIX Collector
- Add a Port Mirroring Profile
- Simple Network Management Protocol (SNMP)
- Using vRealize Log Insight for System Monitoring
- Using vRealize Operations Manager for System Monitoring
- Using vRealize Network Insight Cloud for System Monitoring
- Advanced Monitoring Tools

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**.

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 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.

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	<ul style="list-style-type: none"> ■ 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	<ul style="list-style-type: none"> ■ 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**.

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.

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*.

Example:

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.

Alert Name	Description
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.
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 12-1. Alerts in the Management Pack for NSX-T

Alert	Description	Recommendation
NSX-T Management service has failed	Triggered when the management service on the NSX-T Data Center host is not running.	Please log in to the NSX-T 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.	Please log in to NSX-T 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.	Please 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: <ul style="list-style-type: none"> ■ 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, please 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. Please enable the BFD service if required.
NAT rule not configured	Triggered when the NAT rule on the logical router is not configured.	Please log in to the NSX-T 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.	Please log in to the NSX-T Manager and add the static routes for the Logical Router if required.
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, please log in to NSX-T Manager and enable the service.

Table 12-1. Alerts in the Management Pack for NSX-T (continued)

Alert	Description	Recommendation
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, please log in to NSX-T 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, please log in to NSX-T Manager and enable the service.
Controller Node Connectivity is broken	Triggered when the controller node connection status is down in NSX-T Data Center	Please log in to NSX-T 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.	Please check the status of controller cluster.
Management Status is not stable	Triggered when the status of any node on the management cluster is down.	Please 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, please 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, please 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, please check and clean the File System to make more space.
Edge Cluster Status is down	Triggered when edge cluster status is down.	Please check the edge cluster status and if required follow standard troubleshooting steps recommended by NSX-T documentation and VMware documentation.

Table 12-1. Alerts in the Management Pack for NSX-T (continued)

Alert	Description	Recommendation
Logical Switch State has failed	Triggered when the state of logical switch has failed.	Please check the logical switch state and if necessary follow standard troubleshooting steps recommended by NSX-T documentation and VMware documentation.
Load Balancer Service operational status down	Triggered when the operational status of load balancer service is down.	Please check the operational status of load balancer service and if necessary follow standard troubleshooting steps recommended by NSX-T documentation and VMware documentation.
Load balancer service operational status error	Triggered when the operational status of load balancer service contains error.	Please check the operational status of load balancer service and if necessary follow standard troubleshooting steps recommended by NSX-T documentation and VMware documentation.
Load Balancer virtual server operational state down	Triggered when the operational state of load balancer virtual server is down.	Please check the operational state of load balancer virtual server and if necessary follow standard troubleshooting steps recommended by NSX-T documentation and VMware documentation.
Load Balancer virtual server operational state detached	Triggered when the operational state of load balancer virtual server is detached.	Please check the operational state of load balancer virtual server and if necessary follow standard troubleshooting steps recommended by NSX-T documentation and VMware documentation.
Edge node configuration state has failed	Triggered when the configuration state of edge node has failed.	Please check the configuration state of the edge node and if necessary follow standard troubleshooting steps recommended by NSX-T documentation and VMware documentation.
Management service monitor runtime state has failed	Triggered when the monitor runtime state of the management service stops running.	Please log in to the NSX-T Manager VA and restart the failed management service.

Table 12-1. Alerts in the Management Pack for NSX-T (continued)

Alert	Description	Recommendation
Management cluster's management status is not stable	Triggered when the management status of a management cluster is not stable.	Please check the status of management cluster.
Less than 3 manager nodes are deployed	Triggered when the NSX-T 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.	Please log in to NSX-T Manager and check the manager connectivity of manager node and follow standard troubleshooting steps recommended by NSX-T 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, please 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, please 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 75 percent.	File system usage is more than 70, please 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 12-2. 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.
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.

Table 12-2. vRealize Network Insight Computed NSX-T Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80212	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.
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.

Table 12-2. vRealize Network Insight Computed NSX-T Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80249	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'.
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'.

Table 12-2. vRealize Network Insight Computed NSX-T Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80241	NSXTEdgeNodeTunnelStatusDegradedEvent	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	NSXTHostNodeStatusDegradedEvent	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	NSXTEdgeNodeStatusDegradedEvent	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'.
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.

Table 12-2. vRealize Network Insight Computed NSX-T Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80254	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
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

Table 12-2. vRealize Network Insight Computed NSX-T Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80262	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.
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.

Table 12-2. 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	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	NSXTMPNodeService ControllerStatusEvent	Warning	Controller Service has stopped running.	Controller Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80404	NSXTMPNodeService DataStoreStatusEvent	Warning	DataStore Service has stopped running.	DataStore Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80405	NSXTMPNodeService HttpStatusEvent	Warning	HTTP Service has stopped running.	HTTP Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80406	NSXTMPNodeService InstallUpgradeEvent	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	NSXTMPNodeService LiagentStatusEvent	Warning	Liagent service has stopped running.	Liagent Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80408	NSXTMPNodeService ManagerStatusEvent	Warning	Manager Service has stopped running.	Manager Service status has turned to stopped.
1.3.6.1.4.1.6876.100.1.0.80409	NSXTMPNodeService MgmtPlaneBusStatusEvent	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	NSXTMPNodeService MigrationCoordinatorStatusEvent	Warning	Migration Coordinator Service has stopped running.	Migration Coordinator Service status has turned to stopped.
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.

Table 12-2. vRealize Network Insight Computed NSX-T Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80416	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.
1.3.6.1.4.1.6876.100.1.0.80402	NSXTMPNodeService CmInventoryStatusEvent	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.

Table 12-2. vRealize Network Insight Computed NSX-T Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80405	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.
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 NodeMgmtStatusEvent	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 NodeStatsStatusEvent	Critical	Node Statistics Service has stopped	One of the Services of the NSX-T Management Node, namely Node Statistics has stopped running.

Table 12-2. vRealize Network Insight Computed NSX-T Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80413	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.
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.

Table 12-2. vRealize Network Insight Computed NSX-T Events (continued)

OID	Event Name	Default Severity	UI Name	Description
1.3.6.1.4.1.6876.100.1.0.80421	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.

NSX-T System Events

Here is the list of NSX-T 2.2 to 2.5 events supported in vRealize Network Insight. The Object ID (OID) for all these NSX-T system events is 1.3.6.1.4.1.6876.100.1.0.80203.

Table 12-3. NSX-T System Events

Event Name	Description
vmwNSXPlatformSysCpuUsage	CPU Usage on both manager and edge appliances (NSX-T 2.2).
vmwNSXPlatformSysDiskUsage	Disk Space Usage on both manager and edge appliance for /var/log partition (NSX-T 2.2).
vmwNSXPlatformSysMemUsage	Memory Usage on both manager and edge appliance (NSX-T 2.2).
vmwNSXPlatformSysConfigDiskUsage	Disk Usage for Manager and Edge Appliances for /config partition (NSX-T 2.4).
vmwNSXPlatformSysVarDumpDiskUsage	Disk Usage for Manager and Edge Appliances for /var/dump partition (NSX-T 2.5).
vmwNSXPlatformSysRepositoryDiskUsage	Disk Usage for Manager and Edge Appliances for /repository partition (NSX-T 2.5).

Table 12-3. NSX-T System Events (continued)

Event Name	Description
vmwNSXPlatformSysRootDiskUsage	Disk usage for Manager and Edge appliances for root partition (NSX-T 2.5).
vmwNSXPlatformSysTmpDiskUsage	Disk usage for Manager and Edge appliances for tmp partition (NSX-T 2.5).
vmwNSXPlatformSysImageDiskUsage	Disk Usage for Manager and Edge appliances for /image partition (NSX-T 2.5).
vmwNSXDhcpPoolUsageOverloadedEvent	DHCP pool overloaded/normal (NSX-T 2.5).
vmwNSXDhcpPoolLeaseAllocationFailedEvent	DHCP pool lease allocation failed/succeeded (NSX-T 2.5).
vmwNSXPlatformPasswordExpiryStatus	Password expiration for manager (NSX-T 2.4).
vmwNSXPlatformCertificateExpiryStatus	Certificate expiration for manager (NSX-T 2.4).
vmwNSXRoutingBgpNeighborStatus	BGP neighbor status (NSX-T 2.2).
vmwNSXVpnTunnelState	VPN Tunnel up/down (NSX-T 2.2).
vmwNSXVpnL2TunnelStatus	L2 VPN Session up/down (NSX-T 2.2).
vmwNSXVpnIkeSessionStatus	IKE Session up/down (NSX-T 2.2).
vmwNSXDnsForwarderStatus	DNS Forwarder Status (NSX-T 2.4).
vmwNSXClusterNodeStatus	Cluster Node status (NSX-T 2.4).
vmwNSXFabricCryptoStatus	Edge crypto mux driver failed/passed Known_Answer_Tests(KAT) (NSX-T 2.4).
Manager Disk Utilization is not OK	
BGP Neighbor down	Need an alert when the BGP neighbor is down.
BGP Neighbor Up	Clear Alarm when a neighbor comes up.
Storage usage over X	Alarm for Storage over X - Event is raised for all appliance VM (MP, CCP) or transport nodes (edge, host).
Memory usage over X	Alarm for Memory over X - Event is raised for all appliance VM (MP, CCP) or transport nodes (edge, host).
CPU usage over X	Alarm for CPU over X - Event is raised for all appliance VM (MP, CCP) or transport nodes (edge, host).

Advanced Monitoring Tools

NSX-T support advanced monitoring methods, including viewing port connections, traceflow, port mirroring, activity monitoring, and so on.

View Port Connection Information

You can use the port connection tool to quickly visualize and troubleshoot the connection between two VMs.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Tools > 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.

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.

Trace flow can be used on transport nodes and supports both IPV4 and IPV6 protocols including: ICMP, TCP, UDP, DHCP, DNS and ARP/NDP.

You can construct packets with custom header fields and packet sizes. The source or destination for the trace flow can be a logical switch port, logical router uplink port, CSP or DHCP port. The destination endpoint can be any device in the NSX 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 distributed logical routers.

If NSX 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.

Traceflow observations may include observations of broadcasted traceflow packets. The ESXi host broadcasts a traceflow packet if it does not know the destination host's MAC address. For broadcast traffic, the source is a VM vNIC. The Layer 2 destination MAC address for broadcast traffic 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 to calculate an IP network address for the packet.

Trace the Path of a Packet with 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.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > 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 Specify the source and destination information according to the traffic type.

Traffic Type	Source	Destination
Unicast	<p>Select a VM or a logical port. For a VM:</p> <ul style="list-style-type: none"> ■ 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. <p>For a logical port:</p> <ul style="list-style-type: none"> ■ Select an attachment type: VIF, DHCP, Edge Uplink, or Edge Centralized Service. ■ Select a port. 	<p>Select a VM, a logical port, or IP-MAC. For a VM:</p> <ul style="list-style-type: none"> ■ 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. <p>For a logical port:</p> <ul style="list-style-type: none"> ■ Select an attachment type: VIF, DHCP, Edge Uplink, or Edge Centralized Service. ■ Select a port. <p>For IP-MAC:</p> <ul style="list-style-type: none"> ■ 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

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.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 3 Select **Advanced Networking & Security > Tools > Port Mirroring Session**.

4 Click **Add** and select a session type.

The available types are **Local SPAN**, **Remote SPAN**, **Remote L3 SPAN**, and **Logical SPAN**.

5 Enter a session name and optionally a description.

6 Provide additional parameters.

Session Type	Parameters
Local SPAN	<ul style="list-style-type: none"> ■ Transport Node - Select a transport node. ■ Direction - Select Bidirectional, Ingress, or Egress. ■ Packet Truncation - Select a packet truncation value.
Remote SPAN	<ul style="list-style-type: none"> ■ 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	<ul style="list-style-type: none"> ■ 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	<ul style="list-style-type: none"> ■ Logical Switch - Select a logical switch. ■ Direction - Select Bidirectional, Ingress, or Egress. ■ Packet Truncation - Select a packet truncation value.

7 Click **Next**.

8 Provide source information.

Session Type	Parameters
Local SPAN	<ul style="list-style-type: none"> ■ Select an N-VDS. ■ Select physical interfaces. ■ Enable or disable encapsulated packet. ■ Select virtual machines. ■ Select virtual interfaces.
Remote SPAN	<ul style="list-style-type: none"> ■ Select virtual machines. ■ Select virtual interfaces.
Remote L3 SPAN	<ul style="list-style-type: none"> ■ Select virtual machines. ■ Select virtual interfaces. ■ Select a logical switch.
Logical SPAN	<ul style="list-style-type: none"> ■ Select logical ports.

9 Click **Next**.

10 Provide destination information.

Session Type	Parameters
Local SPAN	<ul style="list-style-type: none"> ■ Select virtual machines. ■ Select virtual interfaces.
Remote SPAN	<ul style="list-style-type: none"> ■ Select an N-VDS. ■ Select physical interfaces.
Remote L3 SPAN	<ul style="list-style-type: none"> ■ Specify an IPv4 address.
Logical SPAN	<ul style="list-style-type: none"> ■ Select logical ports.

11 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

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

You can configure IPFIX collectors for switches.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > Tools > 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

You can configure IPFIX profiles for switches.

Procedure

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

- 2 Select **Advanced Networking & Security > Tools > 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.
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

You can configure IPFIX collectors for firewalls.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > Tools > IPFIX**
- 3 Click the **Firewall IPFIX Collectors** tab.
- 4 Click **Add** to add a collector.
- 5 Enter a name and optionally a description.

- Click **Add** and enter the IP address and port of a collector.

You can add up to 4 collectors.

- Click **Add**.

Configure Firewall IPFIX Profiles

You can configure IPFIX profiles for firewalls.

Procedure

- From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- Select **Advanced Networking & Security > Tools > IPFIX**
- Click the **Firewall IPFIX Profiles** tab.
- 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 G1oba1 . 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.

- 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

Element ID	Parameter Name	Data Type	Unit
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
889	vxlانExportRole	unsigned8	1 byte
890	ingressInterfaceAttr	unsigned16	2 bytes
898	virtualObsID	string	variable length

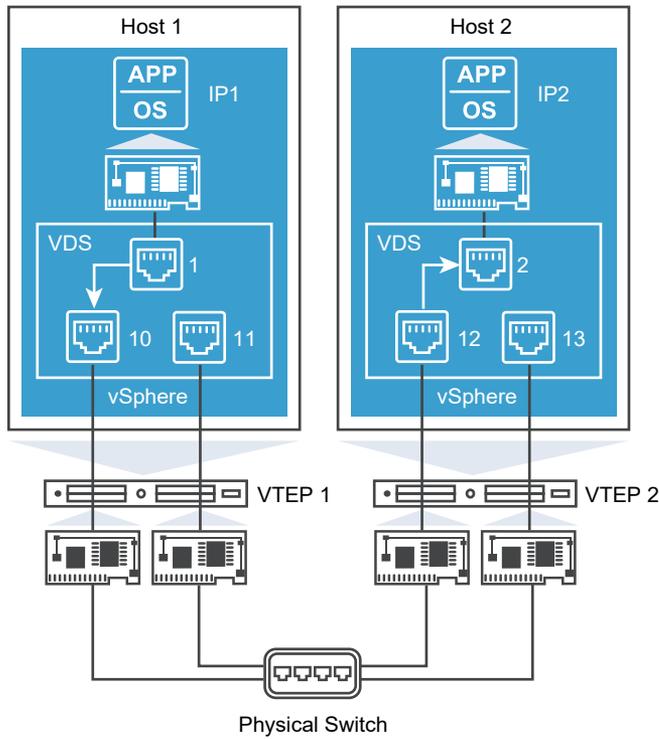
The following table lists VMware-specific elements in distributed firewall IPFIX packets.

Element ID	Parameter Name	Data Type	Unit
950	ruleId	unsigned32	4 bytes
951	vmUuid	string	16 bytes
952	vnicIndex	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)

```

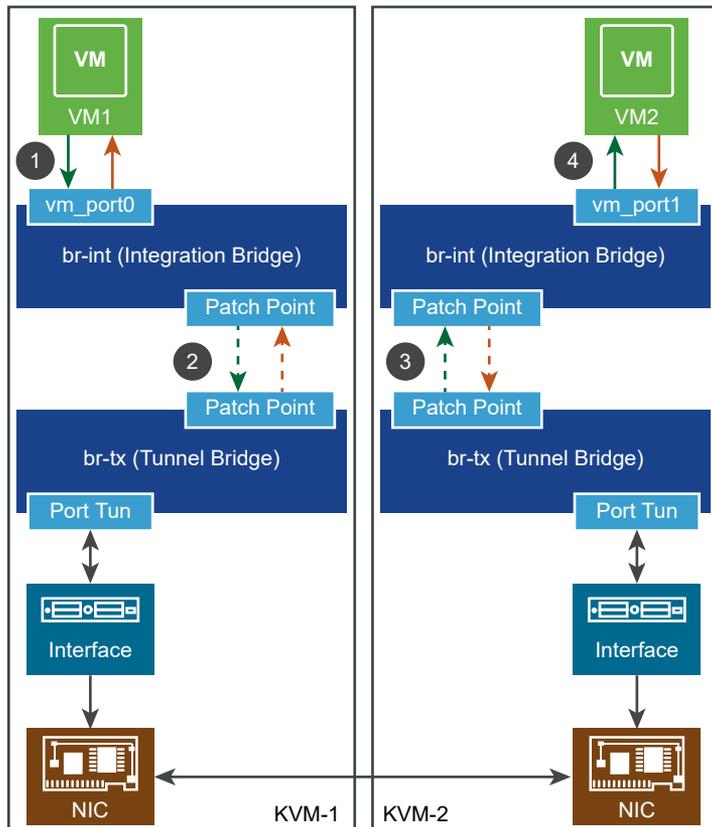
KVM IPFIX Templates

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

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](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Networking > Switching > 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](#).

Logical Switches

13

You can configure logical switches and related objects from the **Advanced Networking & Security** tab. A logical switch reproduces switching functionality, broadcast, unknown unicast, multicast (BUM) traffic, in a virtual environment decoupled from the underlying hardware.

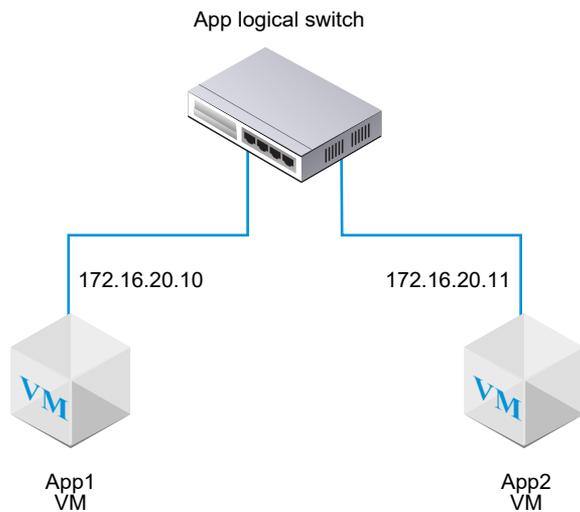
Note If you use the **Advanced Networking & Security** user interface to modify objects created in the policy interface, some settings might not be configurable. These read-only settings have this icon next to them: . See [Chapter 1 Overview of the NSX Manager](#) for more information.

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 13-1. 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.

This chapter includes the following topics:

- [Understanding BUM Frame Replication Modes](#)
- [Create a Logical Switch](#)
- [Connecting a VM to a Logical Switch](#)
- [Create a Logical Switch Port](#)
- [Test Layer 2 Connectivity](#)
- [Create a VLAN Logical Switch for the NSX Edge Uplink](#)
- [Switching Profiles for Logical Switches and Logical Ports](#)
- [Enhanced Networking Stack](#)
- [Layer 2 Bridging](#)

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

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 [Chapter 13 Logical Switches](#) and [Understanding BUM Frame Replication Modes](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Networking > Switching > 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](#).

Connecting a VM to a Logical Switch

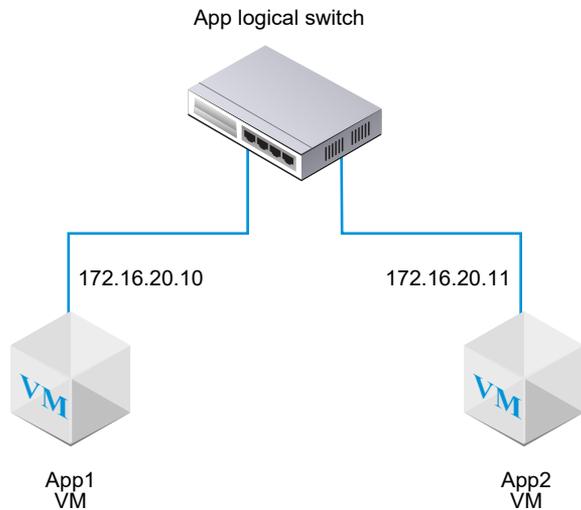
Depending on your host, the configuration for connecting a VM to a logical switch can vary.

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 an NSX-T Data Center Logical Switch

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 an NSX-T Data Center Logical Switch](#).

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)	<input checked="" type="checkbox"/> Connect...
CD/DVD drive 1	Client Device	<input type="checkbox"/> Connect...
Floppy drive 1	Client Device	<input type="checkbox"/> 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 in **Advanced Networking & Security > Networking > Switching > 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*.

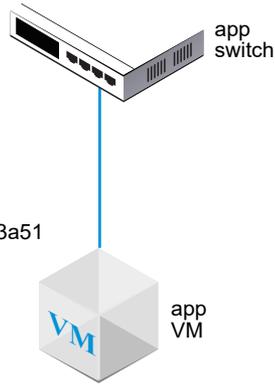
Attach a VM Hosted on Standalone ESXi to an NSX-T Data Center Logical Switch

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



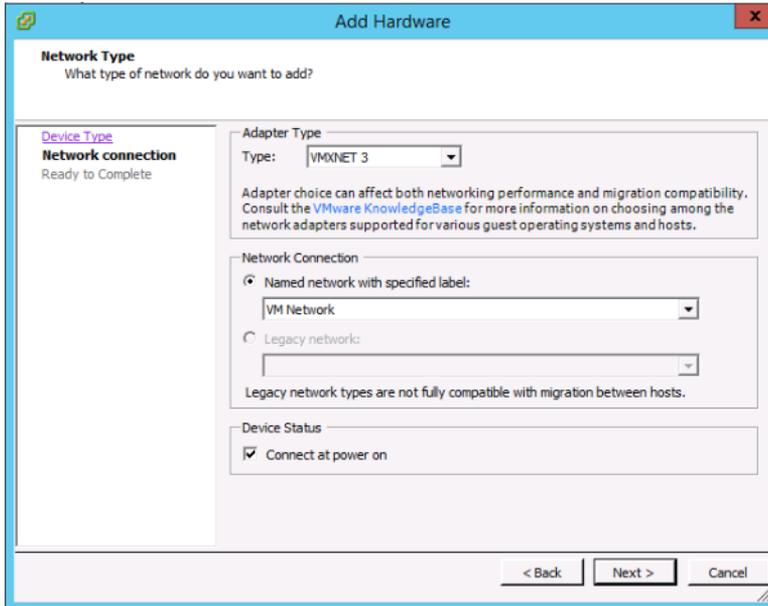
Prerequisites

- 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 under **Advanced Networking & Security > Networking > Switching > 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

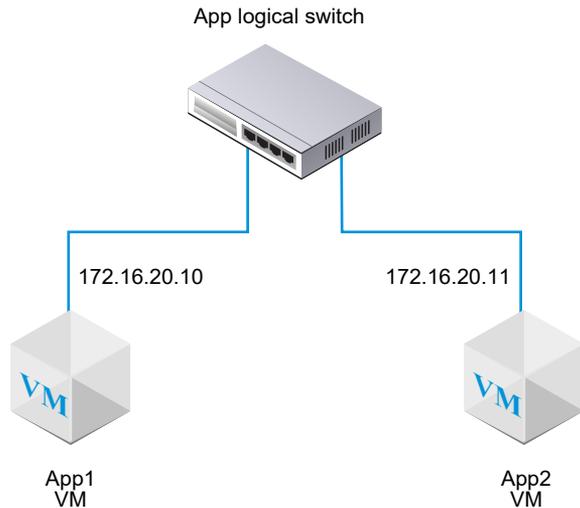
Add a logical router.

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 an NSX-T Data Center Logical Switch

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 under **Advanced Networking & Security > Networking > Switching > 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

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](#).

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](#).

Prerequisites

Verify that a logical switch is created. See [Chapter 13 Logical Switches](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Networking > Switching > 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 .
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.

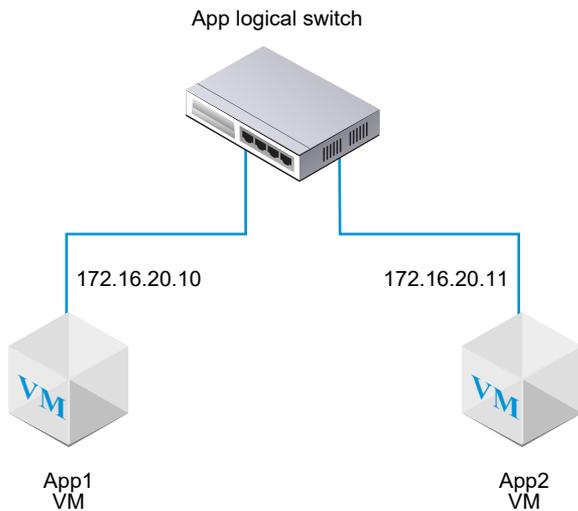
- 5 Click **Save**.

Test Layer 2 Connectivity

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 13-2. Logical Switch Topology



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 **Advanced Networking & Security > Networking > Switching > Switches**.

- e Click the logical switch and note the UUID and VNI information.
- f Run the following commands to troubleshoot the problem.

Command	Description
<code>get logical-switch <vni-or-uuid> arp-table</code>	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>
<code>get logical-switch <vni-or-uuid> connection-table</code>	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>
<code>get logical-switch <vni-or-uuid> mac-table</code>	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>
<code>get logical-switch <vni-or-uuid> stats</code>	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>
<code>get logical-switch <vni-or-uuid> stats-sample</code>	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>
<code>get logical-switch <vni-or-uuid> vtep</code>	<p>Displays all virtual tunnel end points related to the specified logical switch.</p> <p>Sample output.</p> <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>

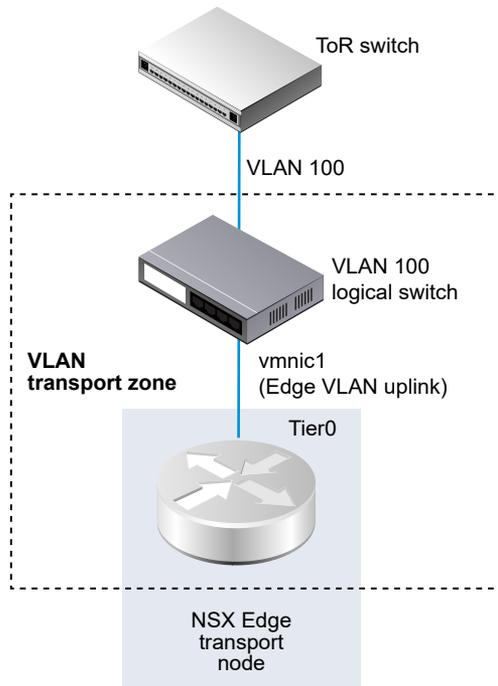
Results

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

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-edg1> 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*.

Procedure

- 1 From a browser, log in to an NSX Manager at `https://<nsx-mgr>`.

- 2 Select **Advanced Networking & Security > Networking > Switching > 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

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 **Advanced Networking & Security > Networking > Switching > 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	<p>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.</p> <p>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.</p> <p>Note DSCP settings work only on tunneled traffic. These settings do not apply to traffic inside the same hypervisor.</p>
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	<p>Set custom values for the outbound network traffic from the VM to the logical network based on broadcast.</p> <p>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.</p> <p>The default value 0 disables rate limiting on the ingress broadcast traffic.</p>
Egress	<p>Set custom values for the inbound network traffic from the logical network to the VM.</p> <p>The default value 0 disables rate limiting on the egress traffic.</p>

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](#) or [Associate a Custom Profile with a Logical Port](#).

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

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.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Networking > Switching > 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	<p>Enter a random 32-bit value to identify mirrored packets from the logical port.</p> <p>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.</p> <p>The default 32-bit value is made of the following values.</p> <ul style="list-style-type: none"> ■ 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	<p>Enter the destination ID of the collector for the mirroring session.</p> <p>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.</p>

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](#) or [Associate a Custom Profile with a Logical Port](#).

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](#).
- Verify that the customized port mirroring switching profile is attached to a logical switch. See [Associate a Custom Profile with a Logical Switch](#).

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](#).

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. By navigating to **Advanced Networking and Security > Switching > Ports** and selecting 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

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](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Networking > Switching > 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.
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.

Option	Description
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](#) or [Associate a Custom Profile with a Logical Port](#).

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

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.

Procedure

- 1 In NSX Manager, select to **Advanced Networking & Security > Networking > Switching > 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](#).

Configure a SpoofGuard Switching 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/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.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Networking > Switching > 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](#) or [Associate a Custom Profile with a Logical Port](#).

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

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](#).

Procedure

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

- 2 Select **Advanced Networking & Security > Networking > Switching**.
- 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.
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](#) or [Associate a Custom Profile with a Logical Port](#).

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. This property is disabled by default, except when the guest VM is deployed using VMware Integrated OpenStack, in which case the 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

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](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Networking > Switching > 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](#) or [Associate a Custom Profile with a Logical Port](#).

Associate a Custom Profile with a Logical Switch

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](#).
- Verify that a custom switching profile is configured. See [Switching Profiles for Logical Switches and Logical Ports](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Networking > Switching > 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](#).

Associate a Custom Profile with a Logical Port

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](#).
- Verify that a custom switching profile is configured. See [Switching Profiles for Logical Switches and Logical Ports](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > Networking > Switching > 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 the *NSX-T Data Center Administration Guide*.

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-ic-mode>
```

Where, *<ens-ic-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.

Layer 2 Bridging

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

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.

Procedure

- 1 Select **System > Fabric > Profiles > Edge Bridge Profiles > Add**.
- 2 Enter a name for the Edge bridge profile and optionally a description.
- 3 Select an NSX Edge cluster.
- 4 Select a primary node.
- 5 Select a backup node.
- 6 Select a failover mode.

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

- 7 Click the **Add** button.

What to do next

You can now associate a logical switch with the bridge profile.

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="vmodl.Binary">AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
```

```

AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA=
```

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="vmodl.Binary">AAABAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA=
```

Use this input to disable a SINK port:

```

<opaqueDataSpec>
  <operation>edit</operation>
  <opaqueData>
    <key>com.vmware.etherswitch.port.extraEthFRP</key>
    <opaqueData
xsi:type="vmodl.Binary">AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA=
```

- e Set `isRuntime` to `false`.
- f Click **Invoke Method**.

Create a Layer 2 Bridge-Backed Logical Switch

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.
- 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.

Procedure

- 1 From a browser, log in to an NSX Manager at `https://<nsx-mgr>`.
- 2 Select **Advanced Networking & Security > Networking > Switching**.
- 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"  
}
```

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 the **Advanced Networking & Security** user interface to modify objects created in the policy interface, some settings might not be configurable. These read-only settings have this icon next to them: . See [Chapter 1 Overview of the NSX Manager](#) for more information.

This chapter includes the following topics:

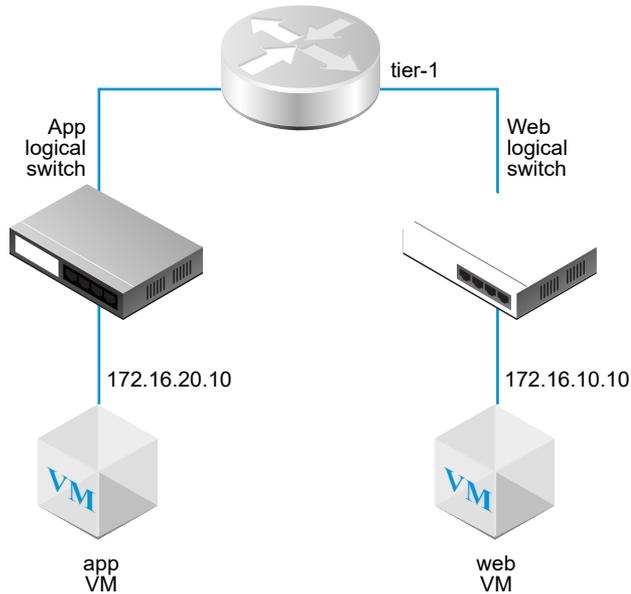
- [Tier-1 Logical Router](#)
- [Tier-0 Logical Router](#)

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 14-1. 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

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](#).
- 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](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Routers > Routers > Add**.
- 3 Select **Tier-1 Router** and 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](#).

Add a Downlink Port on a Tier-1 Logical Router

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](#) .

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Networking > Routers**.
- 3 Click the name of a 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 the router port IP address in CIDR notation.
For example, the IP address can be 172.16.10.1/24.
- 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](#).

Add a VLAN Port on a Tier-0 or Tier-1 Logical Router

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.

Procedure

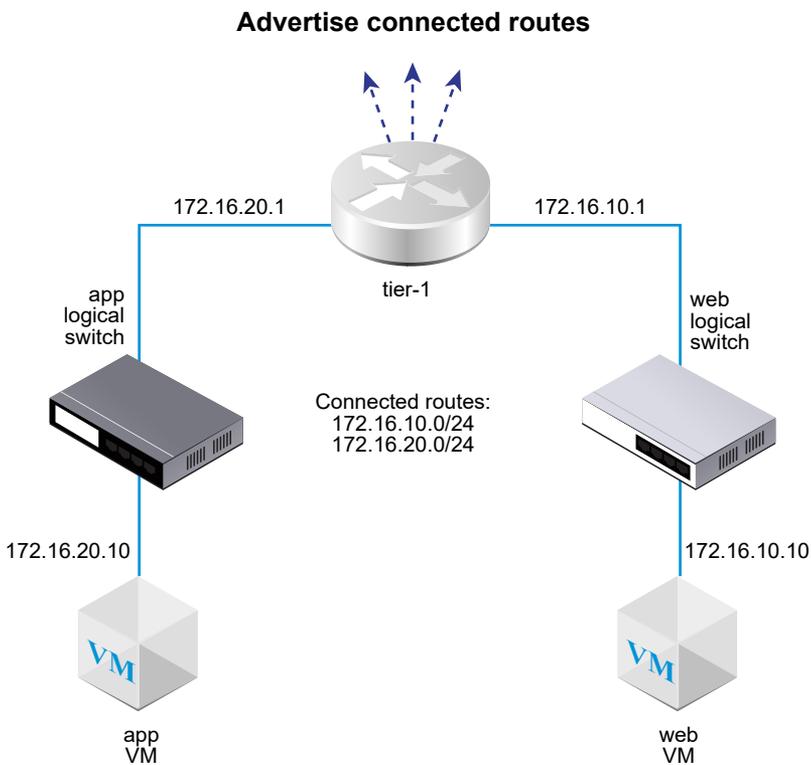
- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Networking > Routers**.
- 3 Click the name of a router.
- 4 Click the **Configuration** tab and select **Router Ports**.
- 5 Click **Add**.
- 6 Enter a name for the router port and optionally a description.
- 7 In the **Type** field, select **Centralized**.
- 8 For **URPF Mode**, select **Strict** or **None**.
URPF (unicast Reverse Path Forwarding) is a security feature.
- 9 (Required) 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 the router port IP address in CIDR notation.

12 Click **Add**.

Configure Route Advertisement on a Tier-1 Logical Router

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 [Chapter 13 Logical Switches](#).
- Verify that downlink ports for the tier-1 logical router are configured. See [Add a Downlink Port on a Tier-1 Logical Router](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > Networking > 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: <ul style="list-style-type: none"> ■ Any ■ NSX Connected ■ Tier-1 LB VIP ■ Static ■ Tier-1 NAT ■ Tier-1 LB SNAT
Prefix operator	Select GE or EQ .

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](#).

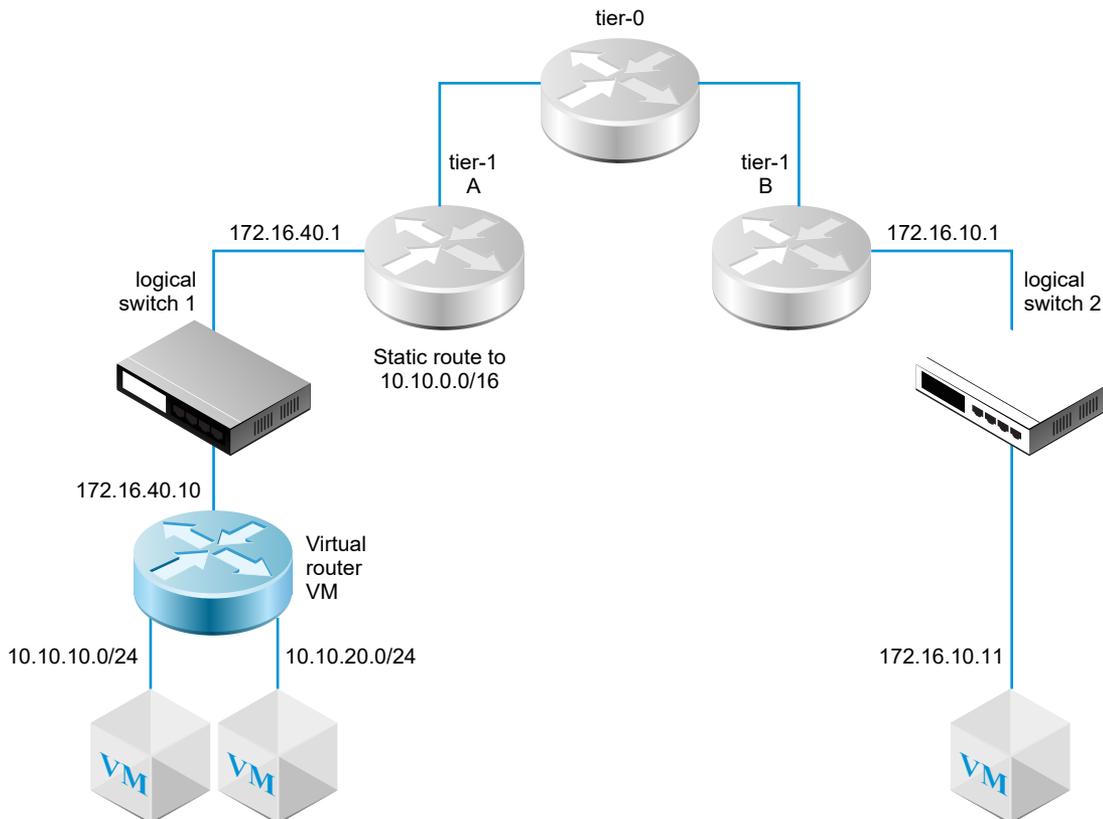
If you already have a tier-0 logical router connected to the tier-1 logical router, you can verify that the tier-0 router is learning the tier-1 router connected routes. See [Verify that a Tier-0 Router Has Learned Routes from a Tier-1 Router](#).

Configure a Tier-1 Logical Router Static Route

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 14-2. 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](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Networking > 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

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.
- Must not have a downlink.
- 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, DNAT, 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 - DNAT - 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.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > Routers > Routers > Add**.
- 3 Select **Tier-1 Router** and 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

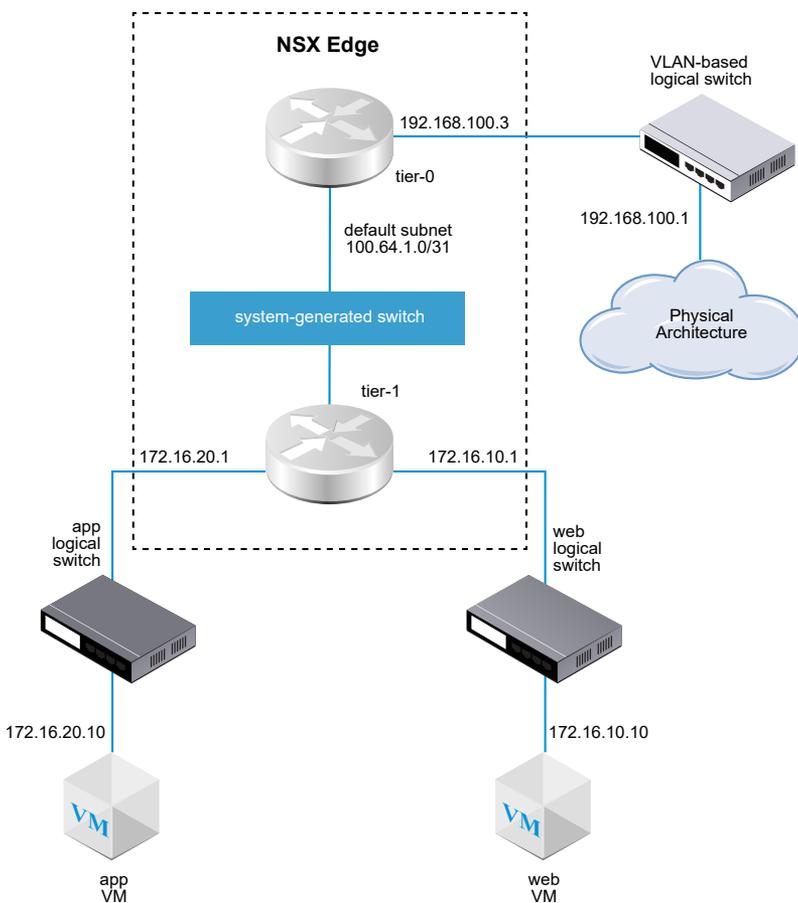
A tier-0 logical router provides a gateway service between the logical and physical network.

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 14-3. 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

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](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Routers > Routers > Add**.
- 3 Select **Tier-0 Router** from the drop-down menu.

- 4 Assign a name for the tier-0 logical router.
- 5 Select an existing NSX Edge cluster from the drop-down menu to back this tier-0 logical router.

- 6 (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.

- 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-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.

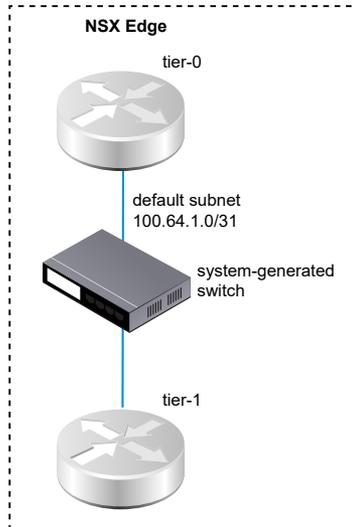
Configure the tier-0 logical router to connect it to a VLAN logical switch to create an uplink to an external network. See [Connect a Tier-0 Logical Router to a VLAN Logical Switch for the NSX Edge Uplink](#).

Attach Tier-0 and Tier-1

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.



Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > Networking > Routers**.
- 3 Select the tier-1 logical router.
- 4 From the **Summary** tab, click **Edit**.
- 5 Select the 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)>

```

- 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

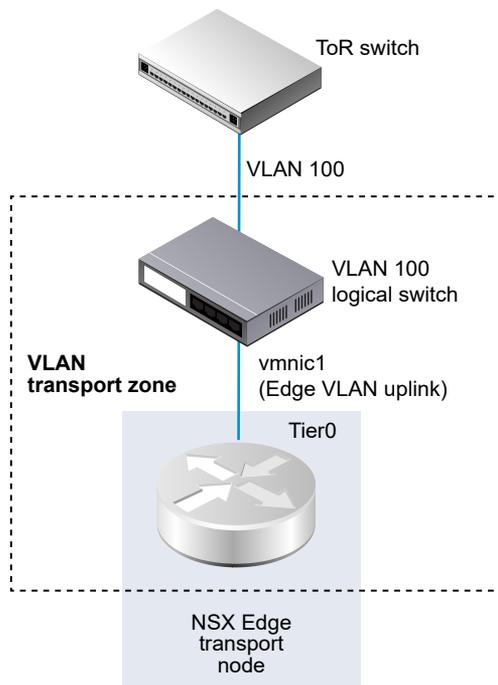
Total number of routes: 7

b   10.10.10.0/24      [20/0]      via 192.168.100.254
rl  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

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](#).

Create a tier-0 router.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Networking > 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 Type an IP address in CIDR format 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](#).

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-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: 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
rl  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

```

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](#) or [Configure BGP on a Tier-0 Logical Router](#).

Add a Loopback Router Port

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

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > Networking > 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

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.

Procedure

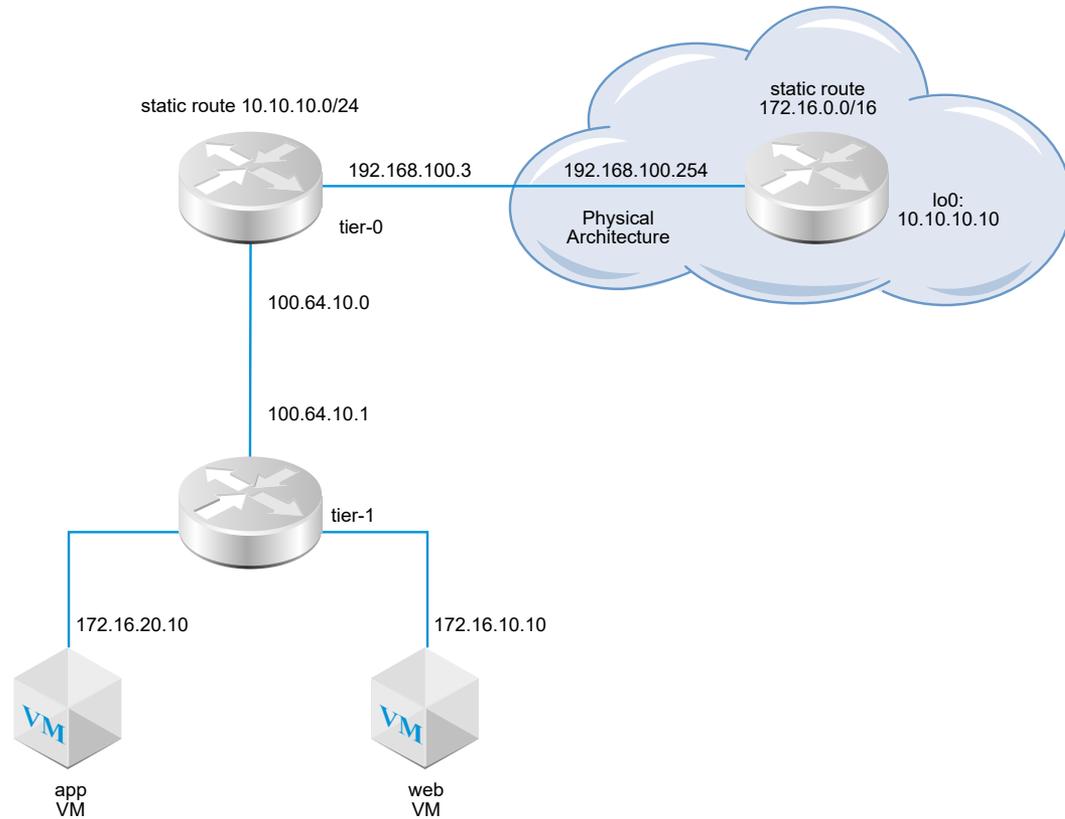
- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Networking > Routers**.
- 3 Click the name of a router.
- 4 Click the **Configuration** tab and select **Router Ports**.
- 5 Click **Add**.
- 6 Enter a name for the router port and optionally a description.
- 7 In the **Type** field, select **Centralized**.
- 8 For **URPF Mode**, select **Strict** or **None**.
URPF (unicast Reverse Path Forwarding) is a security feature.
- 9 (Required) 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 the router port IP address in CIDR notation.
- 12 Click **Add**.

Configure a Static Route

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 14-4. 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](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Networking > 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](#).

Verify the Static Route

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](#).

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.

```
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
 2 100.64.1.1 (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.

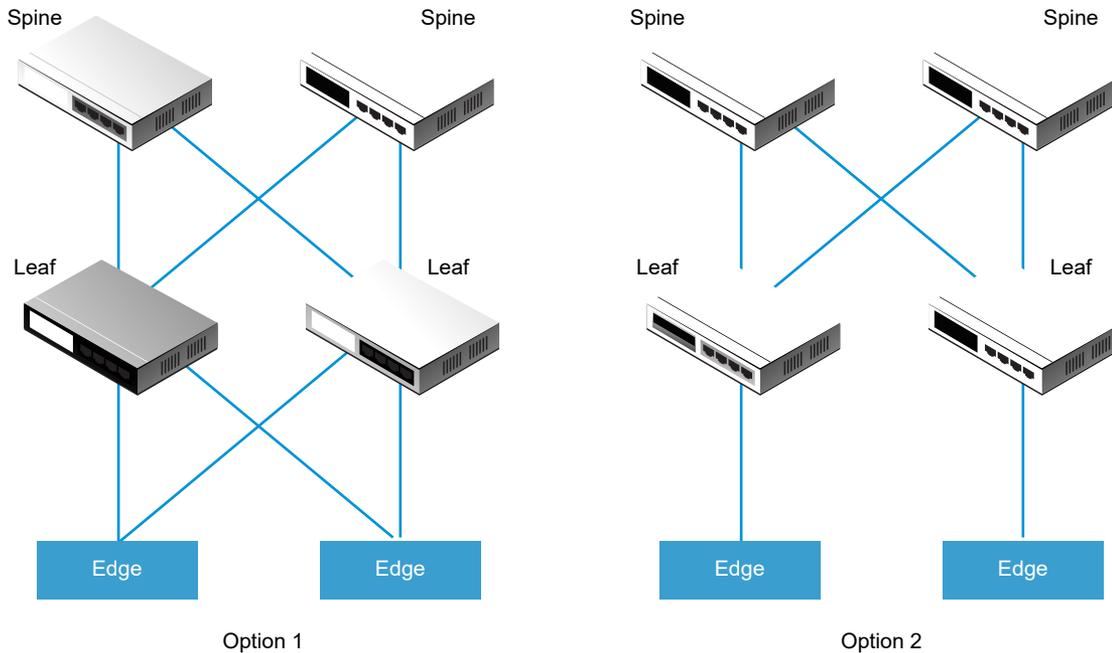
```
ping 10.10.10.10
```

```
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

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.

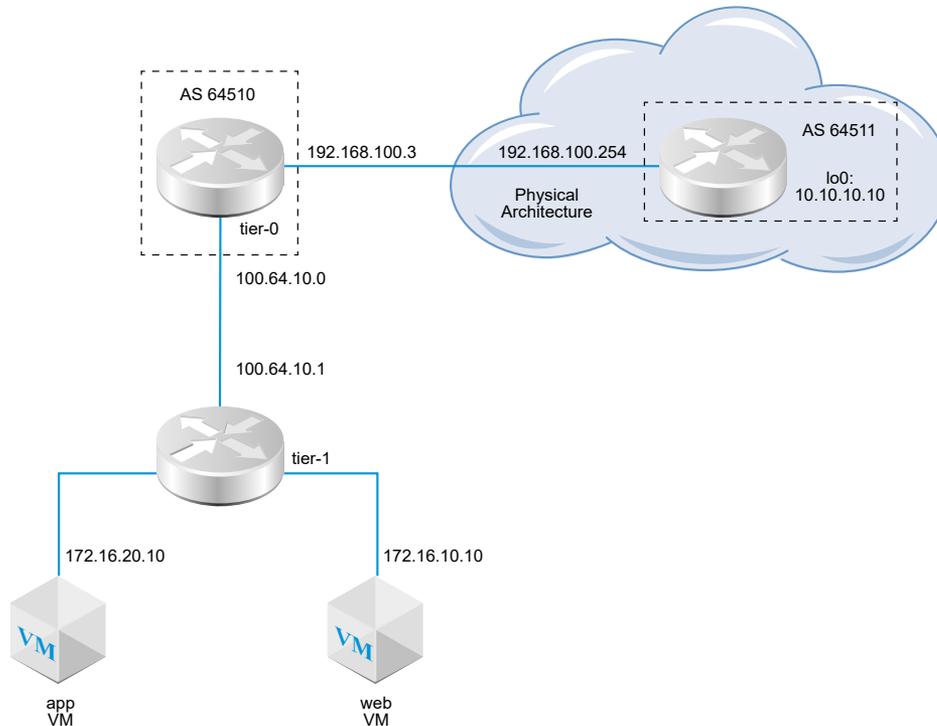
A tier-0 logical router in active-active mode supports inter-SR (service router) routing. If router #1 is unable to communicate with a northbound physical router, traffic is re-routed to router #2 in the active-active cluster. If router #2 is able to communicate with the physical router, traffic between 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 14-5. 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](#). 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](#).
- 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](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > Networking > 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](#).

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-bbfeb2786666
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-edg1(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](#).

Configure BFD on a Tier-0 Logical Router

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.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Networking > 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

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-0 and Tier-1](#).
- If you want to filter specific IP addresses from route redistribution, verify that route maps are configured. See [Create a Route Map](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Networking > 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: <ul style="list-style-type: none"> ■ 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

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-0 Logical Router](#).
- Verify that NSX-T Data Center static routes are set to be redistributed. See [Enable Route Redistribution on the Tier-0 Logical Router](#).

Procedure

- 1 Log in to the NSX Manager CLI.
- 2 View the routes learned from the external BGP neighbor.

```
nsx-edgel1(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

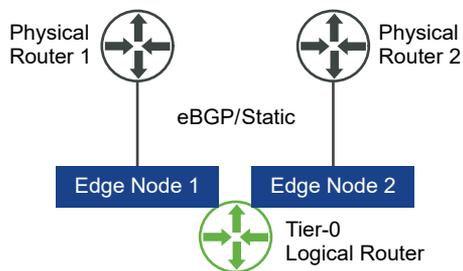
Configure additional routing functionality, such as ECMP.

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 14-6. 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

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](#).
- Verify that a tier-0 logical router is configured. See [Create a Tier-0 Logical Router](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > Networking > 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	Unicast Reverse Path Forwarding is a security feature. Setting it to None is recommended if you have multiple active-active Edge nodes in ECMP mode. The default is Strict .
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](#).

Add a Second BGP Neighbor and Enable ECMP Routing

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](#) .

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Networking > 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](#).

Verify ECMP Routing Connectivity

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](#) and [Add a Second BGP Neighbor and Enable ECMP Routing](#).

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

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](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Networking > 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

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](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Networking > 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

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](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > Networking > 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"> ■ 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

- 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

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.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > Networking > 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**.

You can configure NAT from the **Advanced Networking & Security** tab.

Note If you use the **Advanced Networking & Security** user interface to modify objects created in the policy interface, some settings might not be configurable. These read-only settings have this icon next to them: . See [Chapter 1 Overview of the NSX Manager](#) for more information.

This chapter includes the following topics:

- [Network Address Translation](#)

Network Address Translation

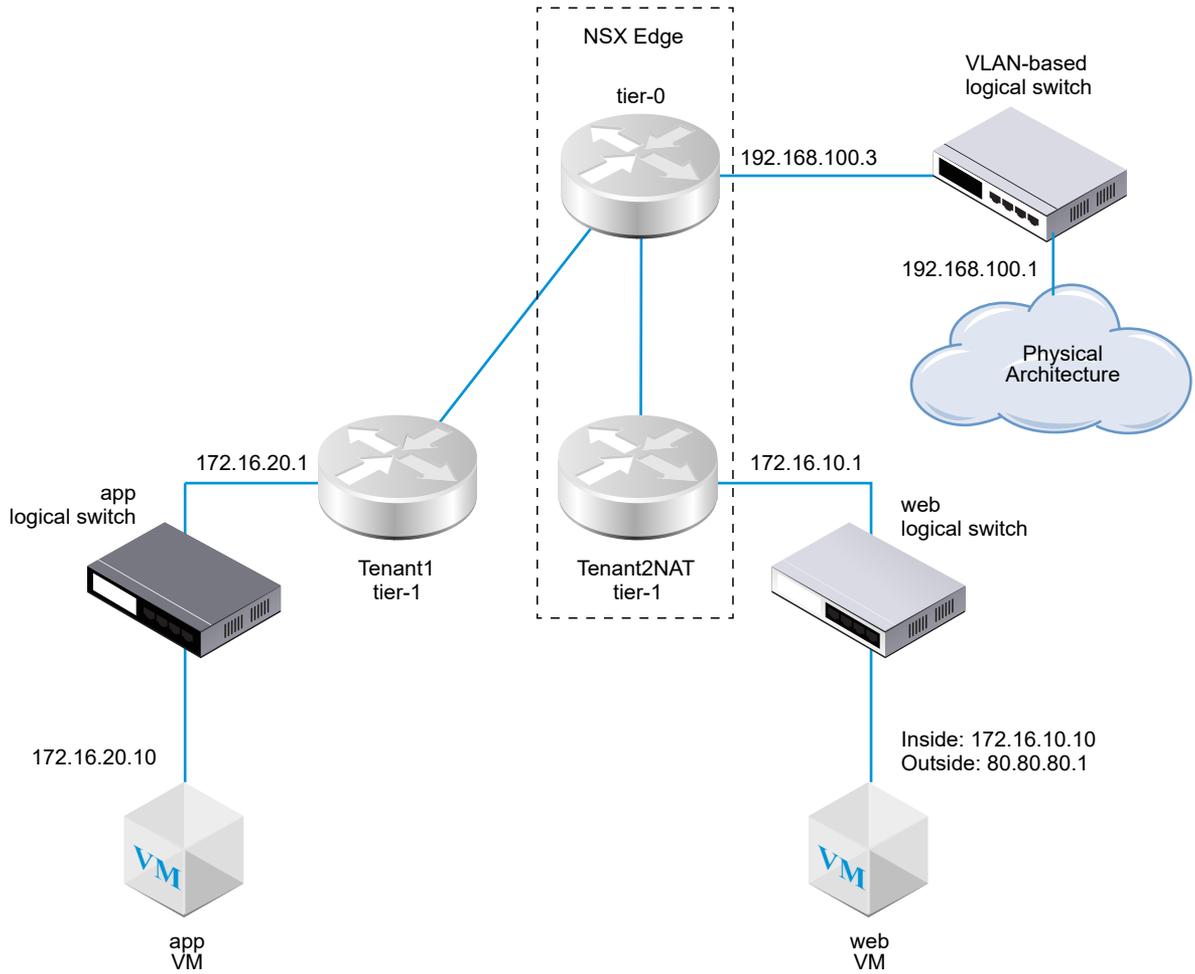
Network address translation (NAT) in NSX-T Data Center can be configured on tier-0 and tier-1 logical routers.

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.

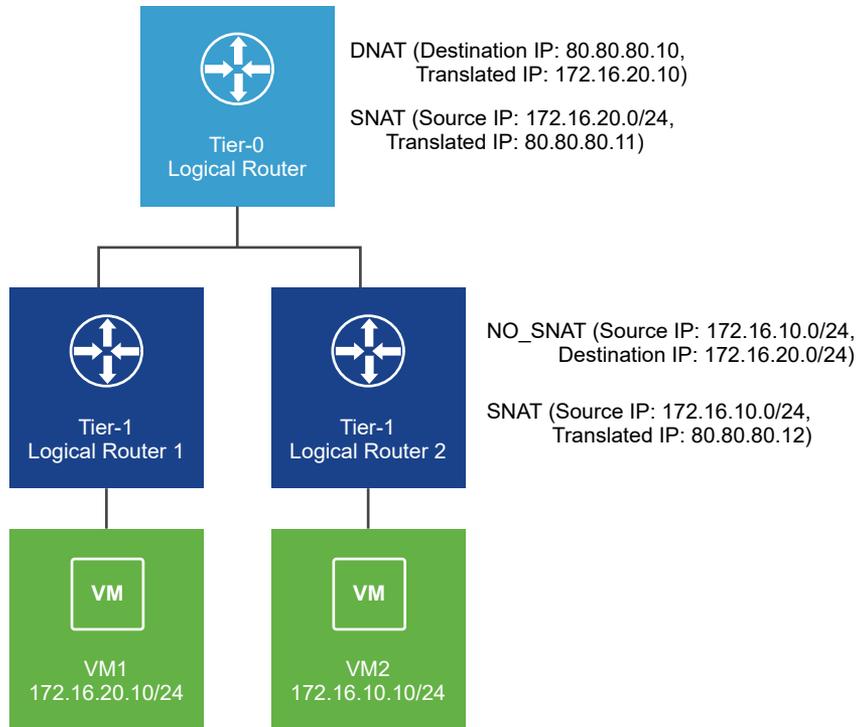
NAT is enforced at the uplink of the Tenant2NAT logical router on its connection to the tier-0 logical router.

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 15-1. 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

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](#).

- 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](#), [Configure BGP on a Tier-0 Logical Router](#), and [Enable Route Redistribution on the Tier-0 Logical Router](#).
- 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-0 and Tier-1](#).
- The tier-1 routers must have downlink ports and route advertisement configured. See [Add a Downlink Port on a Tier-1 Logical Router](#) and [Configure Route Advertisement on a Tier-1 Logical Router](#).
- The VMs must be attached to the correct logical switches.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Networking > 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	
Priority: 1024								
4100	SNAT	Any	172.16.10.10	Any	Any	Any	80.80.80.1	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.

Configure Destination NAT on a Tier-1 Router

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](#).
- 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](#), [Configure BGP on a Tier-0 Logical Router](#), and [Enable Route Redistribution on the Tier-0 Logical Router](#).
- 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-0 and Tier-1](#).

- The tier-1 routers must have downlink ports and route advertisement configured. See [Add a Downlink Port on a Tier-1 Logical Router](#) and [Configure Route Advertisement on a Tier-1 Logical Router](#).
- The VMs must be attached to the correct logical switches.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Networking > 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 an IP address range in CIDR format.
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

 ADD  EDIT  DELETE  COLUMNS

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

Advertising tier-1 NAT routes enables the upstream tier-0 router to learn about these routes.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > Networking > 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

Advertising tier-1 NAT routes from the tier-0 router enables the upstream physical architecture to learn about these routes.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Routing**.
- 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
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: <ul style="list-style-type: none"> ■ 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 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

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.

Procedure

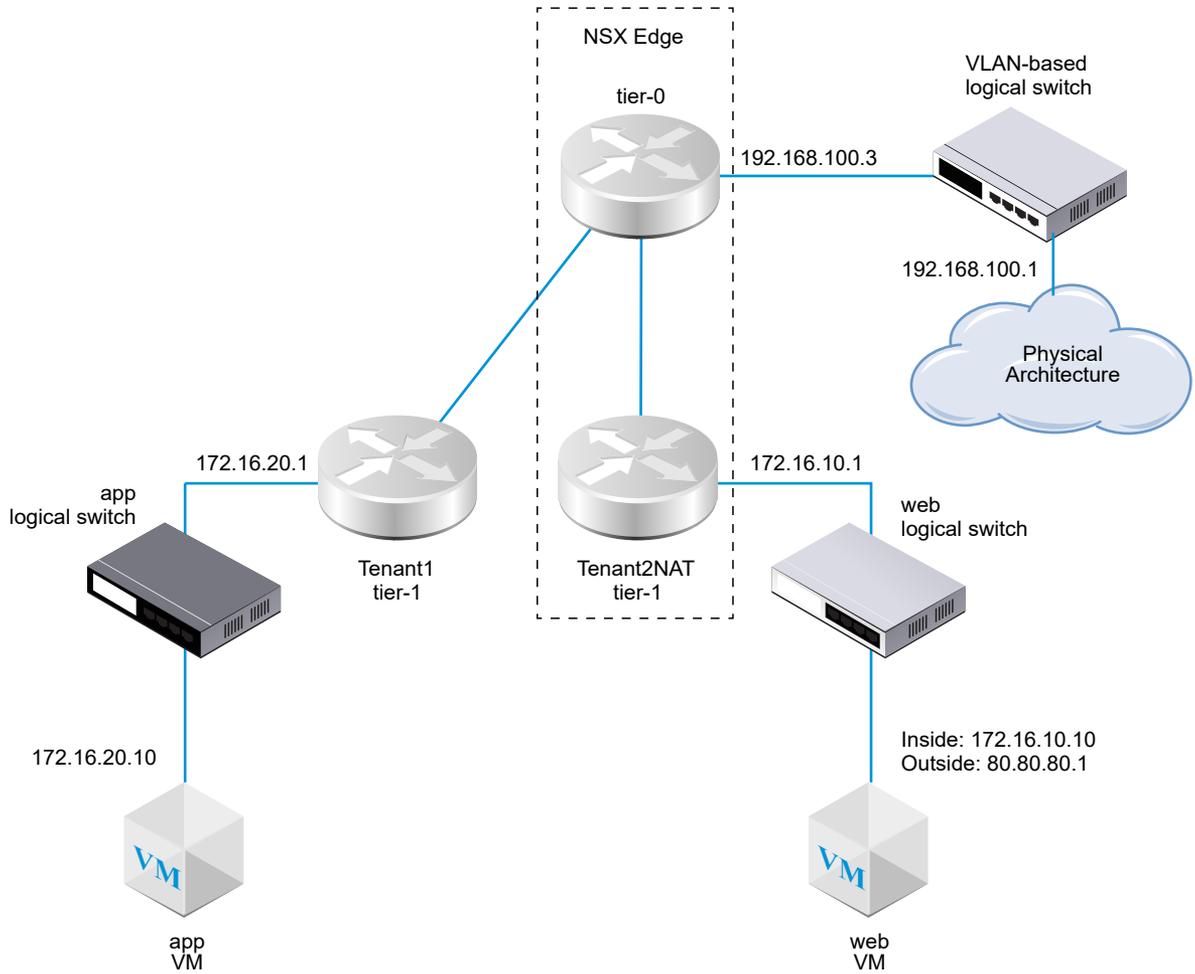
- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Networking > Routers**.
- 3 Click a tier-0 logical router.
- 4 Select **Services > NAT**.
- 5 Click **ADD** to add a NAT rule.
- 6 Specify a priority 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.

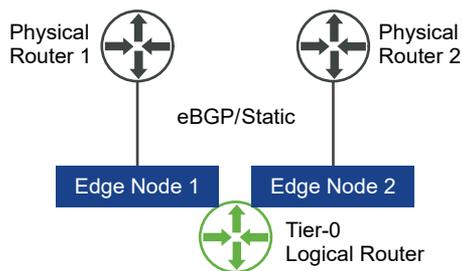
Reflexive NAT

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

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.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > Networking > Routers**.
- 3 Click a 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:

TierO-LR-1

Overview Configuration Routing **Services**

NAT | REFRESH

Total Rule Statistics | Last Updated: 11/6/2018, 12:40:13 PM

0 Active sessions 0 Packet count 0 Bytes Data

+ ADD EDIT DELETE

ID	Action	Match				Translated		Applied To	Stats
		Protocol	Source IP	Source Ports	Destination IP	Destination Ports	IP		
▼ Priority: 1024									
1034	Reflexive	Any	80.80.80.1	Any	Any	Any	172.16.10.10	Any	

Advanced Grouping Objects

16

You can create IP sets, IP pools, MAC sets, NSGroups, and NSServices. You can also manage tags for VMs.

Note If you use the **Advanced Networking & Security** user interface to modify objects created in the policy interface, some settings might not be configurable. These read-only settings have this icon next to them: . See [Chapter 1 Overview of the NSX Manager](#) for more information.

This chapter includes the following topics:

- [Create an IP Set](#)
- [Create an IP Pool](#)
- [Create a MAC Set](#)
- [Create an NSGroup](#)
- [Configuring Services and Service Groups](#)
- [Manage Tags for a VM](#)

Create an IP Set

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. 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.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > 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

You can use an IP Pool to allocate IP addresses or subnets when you create L3 subnets.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > Inventory > Groups > IP Pools > Add**.
- 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

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.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > 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

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**.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > 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

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.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > 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

You can see the list of VMs in the inventory. You can also add tags to a VM to make searching easier.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > 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 ADD to specify a tag and optionally a scope.
Delete a tag	Select an existing tag and click DELETE .

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**.

You can configure DHCP from the **Advanced Networking & Security** tab.

Note If you use the **Advanced Networking & Security** user interface to modify objects created in the policy interface, some settings might not be configurable. These read-only settings have this icon next to them: . See [Chapter 1 Overview of the NSX Manager](#) for more information.

This chapter includes the following topics:

- [DHCP](#)
- [Metadata Proxies](#)

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.

Note A DFW rule that has `Logical Switch/Logical Port/NSGroup` as the source, `Any` as the destination, and is configured to drop DHCP packets for ports 67 and 68, will fail to block DHCP traffic. To block DHCP traffic, configure `Any` as the source as well as the destination.

In this release, the DHCP server does not support guest VLAN tagging.

Create a DHCP Server Profile

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.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > 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](#).

Create a DHCP Server

You can create DHCP servers to service DHCP requests from VMs that are connected to logical switches.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > 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](#).

Attach a DHCP Server to a Logical Switch

You must attach a DHCP server to a logical switch before the DHCP server can process DHCP requests from VMs connected to the switch.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > Networking > Switching**.
 - a Click the checkbox of a logical switch.
 - b Click **Actions > Attach DHCP Server**.
- 3 Alternatively, select **Advanced Networking & Security > DHCP**.
 - a Click the **Servers** tab.
 - b Click the checkbox of a DHCP server.
 - c Click **Actions > Attach to Logical Switch**.

Detach a DHCP Server from a Logical Switch

You can detach a DHCP server from a logical switch to reconfigure your environment.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > Networking > Switching**.
- 3 Click the logical switch that you intend to detach a DHCP server from.
- 4 Click **Actions > Detach DHCP Server**.

Create a DHCP Relay Profile

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.

Procedure

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

- 2 Select **Advanced Networking & Security > 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](#).

Create a DHCP Relay Service

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.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > 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](#).

Add a DHCP Relay Service to a Logical Router Port

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](#).
- Verify that the router port is of type **Downlink**.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > Networking > 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

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*.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > 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

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](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > 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 **Switching > Switches**, selecting a switch, and selecting the menu option **Actions > Attach Metadata Proxy**.

Detach a Metadata Proxy Server from a Logical Switch

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 **Advanced Networking & Security > 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 **Switching > Switches**, selecting a switch, and selecting the menu option **Actions > Detach Metadata Proxy**.

Advanced IP Address Management

18

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 the **Advanced Networking & Security** user interface to modify objects created in the policy interface, some settings might not be configurable. These read-only settings have this icon next to them: . See [Chapter 1 Overview of the NSX Manager](#) for more information.

This chapter includes the following topics:

- [Manage IP Blocks](#)
- [Manage Subnets for IP Blocks](#)

Manage IP Blocks

Setting up NSX Container Plug-in requires that you create IP blocks for the containers.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Networking > IPAM**.
- 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

You can add or delete subnets for IP blocks.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > Networking > IPAM**.
- 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**.

Advanced Load Balancing

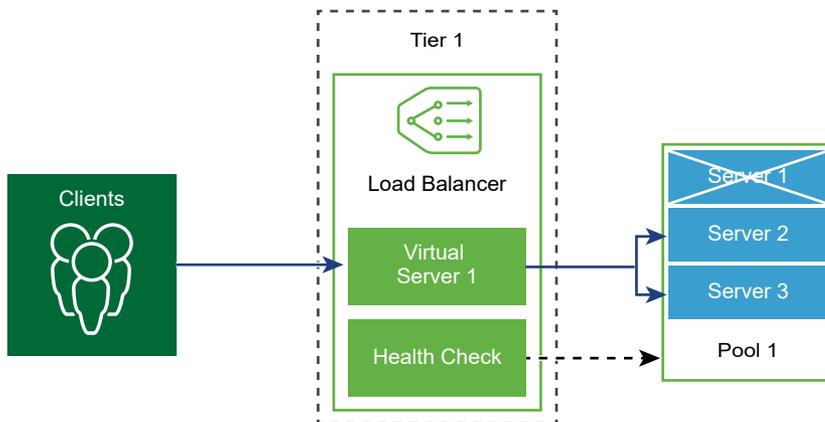
19

This information covers the NSX-T Data Center load balancing configuration found under the **Advanced Networking & Security** tab.

For information about NSX Advanced Load Balancer (Avi Networks) see <https://www.vmware.com/products/nsx-advanced-load-balancer.html>.

Note If you use the **Advanced Networking & Security** user interface to modify objects created in the policy interface, some settings might not be configurable. These read-only settings have this icon next to them: . See [Chapter 1 Overview of the 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.

This chapter includes the following topics:

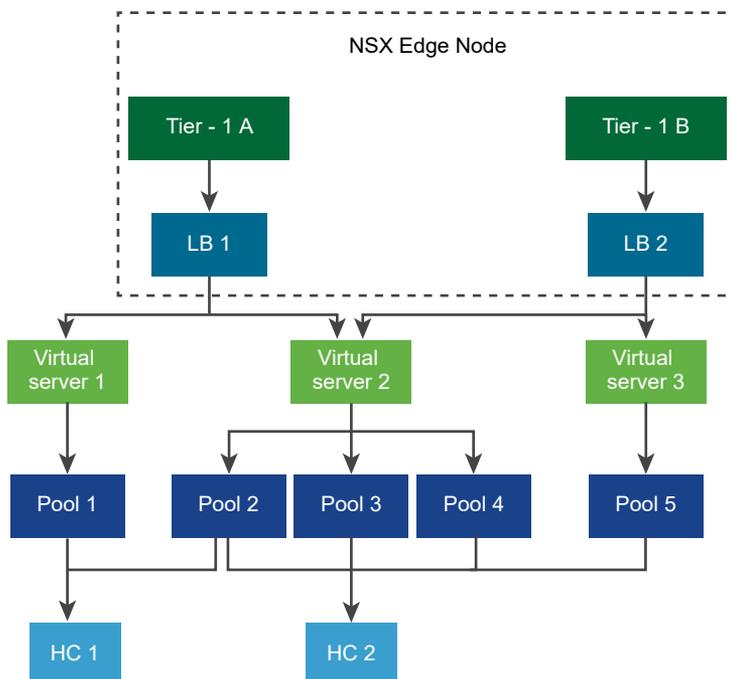
- [Key Load Balancer Concepts](#)
- [Configuring Load Balancer Components](#)

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.

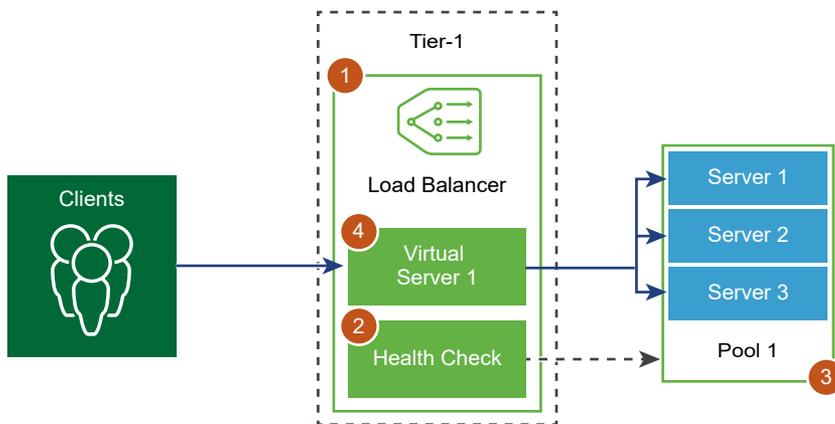
To test whether each server is correctly running the application, you can add health check monitors that check the health status of a server.



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.

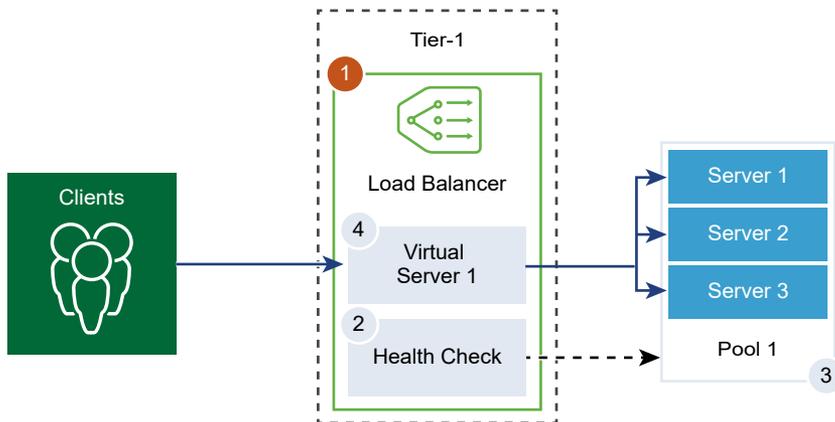


Create a Load Balancer

Load balancer is created and attached to the Tier-1 logical router.

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 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](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > 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

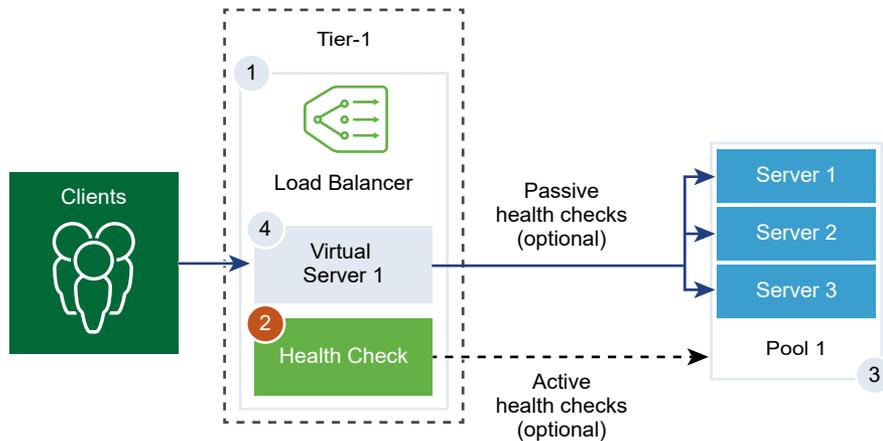
The active health monitor is used to test whether a server is available. The active health monitor uses several types of tests such as sending a basic ping to servers or advanced HTTP requests to monitor the application health.

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](#) for information about standalone Tier-1 gateways.

Note 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 **Advanced Networking & Security > Networking > Load Balancer > 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
Monitoring Interval	Set the time in seconds that the monitor sends another connection request to the server.
Fall Count	Set a value when the consecutive failures reach this value, the server is considered temporarily unavailable.

Option	Description
Rise Count	Set a number after this timeout period, the server is tried again for a new connection to see if it is available.
Timeout Period	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
HTTP Method	Select the method for detecting the server status from the drop-down menu, GET, OPTIONS, POST, HEAD, and PUT.
HTTP Request URL	Enter the request URI for the method.
HTTP Request Version	Select the supported request version from the drop-down menu. You can also accept the default version, HTTP_VERSION_1_1.
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.

- 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
HTTP Method	Select the method for detecting the server status from the drop-down menu: GET, OPTIONS, POST, HEAD, and PUT.
HTTP Request URL	Enter the request URI for the method.
HTTP Request Version	Select the supported request version from the drop-down menu. You can also accept the default version, HTTP_VERSION_1_1.
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.

- 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
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.

- 12 Click **Finish**.

What to do next

Associate the active health monitor with a server pool. See [Add a Server Pool for Load Balancing](#).

Configure Passive Health Monitors

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.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > Networking > Load Balancer > 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
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.

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](#).

Add a Server Pool for Load Balancing

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.

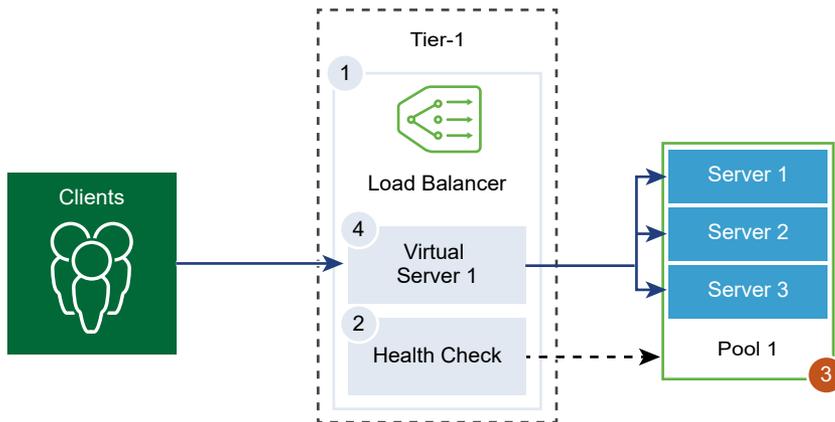
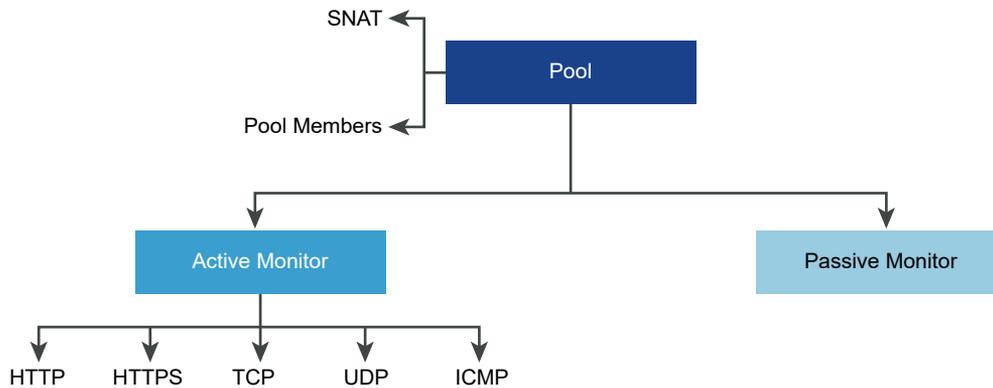


Figure 19-1. Server Pool Parameter Configuration



Prerequisites

- If you use dynamic pool members, a NSGroup must be configured. See [Create an NSGroup](#).
- Depending on the monitoring you use, verify that active or passive health monitors are configured. See [Configure an Active Health Monitor](#) or [Configure Passive Health Monitors](#).

Procedure

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

- 2 Select **Advanced Networking & Security > Networking > Load Balancer > 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
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 fairly. 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 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
Transparent Mode	<p>Load balancer uses the client IP address and port spoofing while establishing connections to the servers.</p> <p>SNAT is not required.</p>
Auto Map Mode	<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>
IP List Mode	<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
Static	Click Add to include a static pool member. You can also clone an existing static pool member.
Dynamic	Select the NSGroup from the drop-down menu. 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](#) 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.

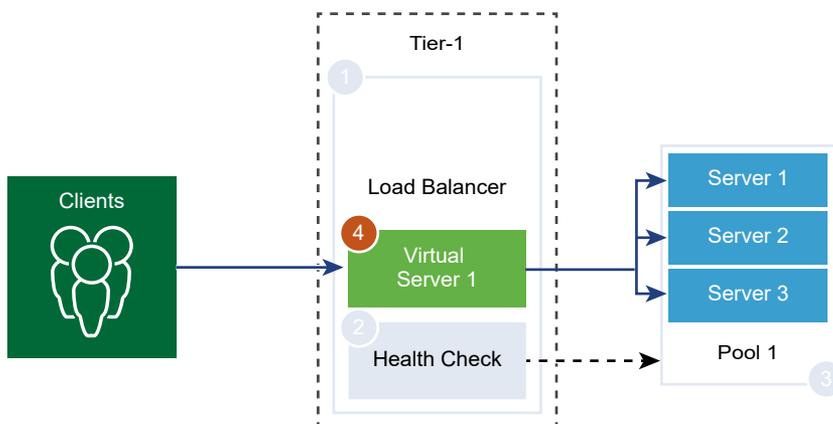
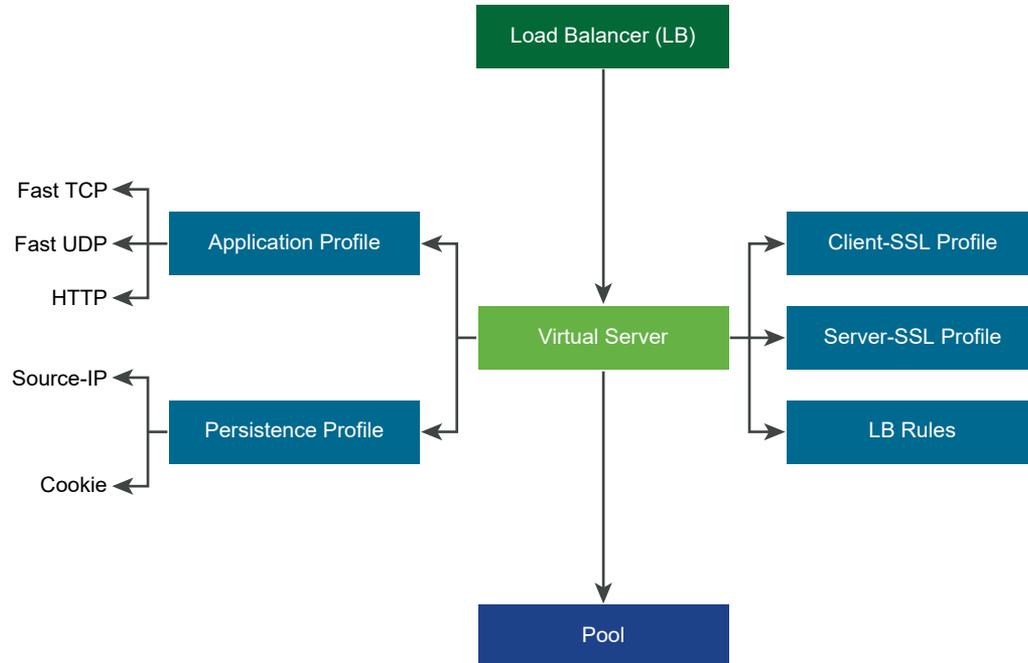


Figure 19-2. Virtual Server Components



Configure Application Profiles

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 19-3. Layer 4 TCP and UDP Application Profile

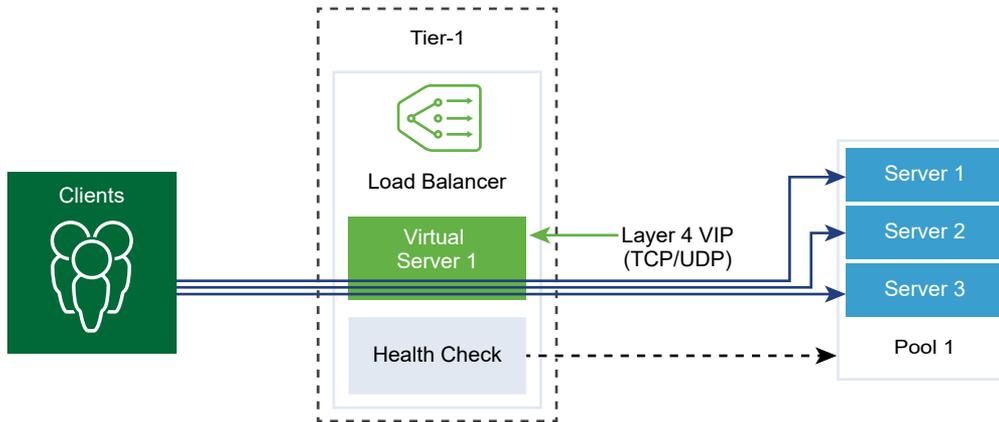
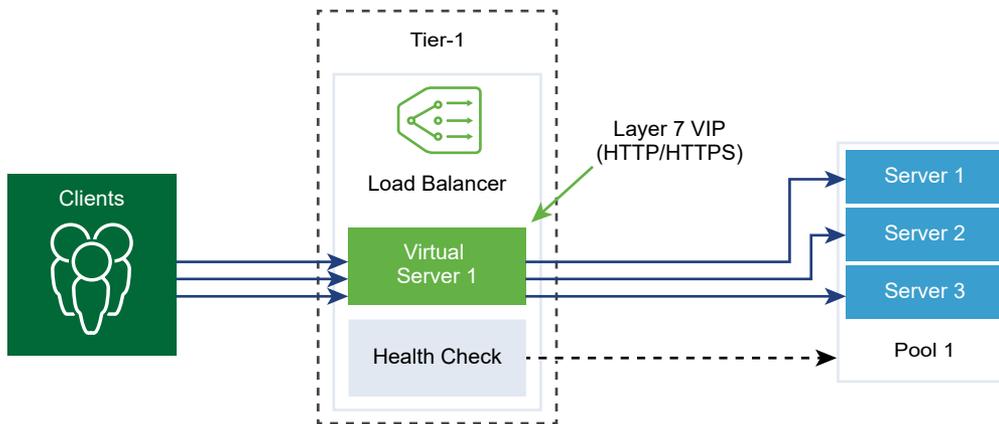


Figure 19-4. Layer 7 HTTPS Application Profile



Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > Networking > Load Balancer > 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
Connection 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.
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.
HA Flow Mirroring	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
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.

- 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"> ■ 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.
X-Forwarded-For (XFF)	<ul style="list-style-type: none"> ■ 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. <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. 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**.

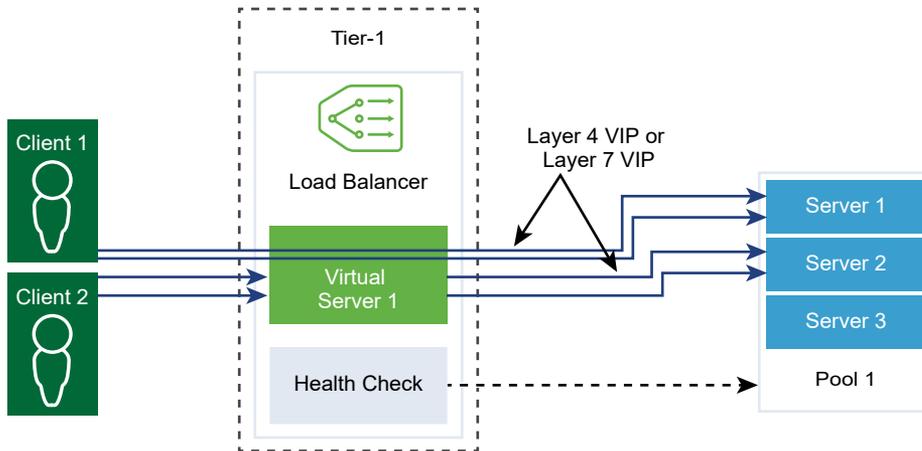
Configure Persistent Profiles

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 a blank space in a cookie name 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 **Advanced Networking & Security > Networking > Load Balancer > 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
Share Persistence	<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>
Persistence Entry Timeout	<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"> ■ If no new connection requests are received from the same client within the timeout period, the persistence entry expires and is deleted. ■ 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. <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>
HA Persistence Mirroring	<p>Toggle the button to synchronize persistence entries to the HA peer.</p>
Purge Entries When Full	<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.

- a Select **Add > Cookie Persistence** from the drop-down menu.
- b Enter a name and a description for the Cookie persistence profile.
- c 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"> ■ 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. Note that 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 Path	Enter the cookie URL path. HTTP cookie path can be set only in the INSERT mode.
Cookie Garbling	Encrypt the cookie server IP address and port information. 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. 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. Session Cookie is not stored and will be lost when the browser is closed. Persistence Cookie 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 Session Cookie only. Enter the maximum age in seconds that a cookie can be active.

- g Click **Finish**.

Configure 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 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 19-5. SSL Offloading

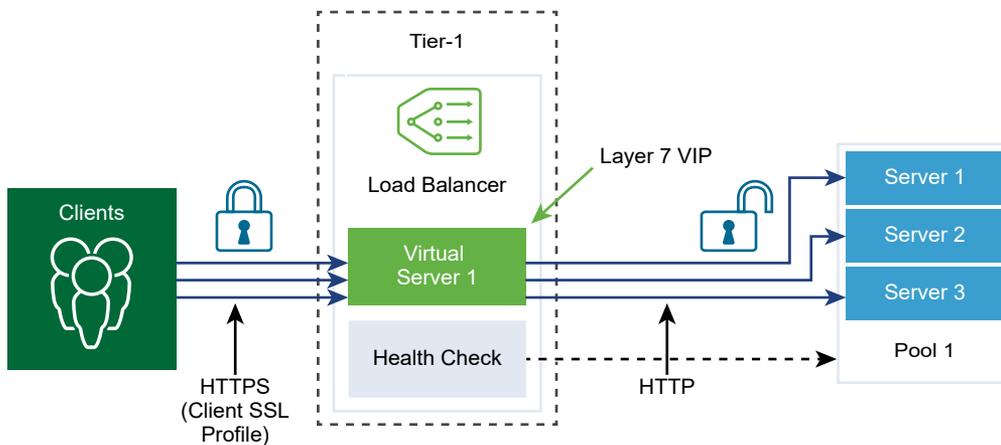
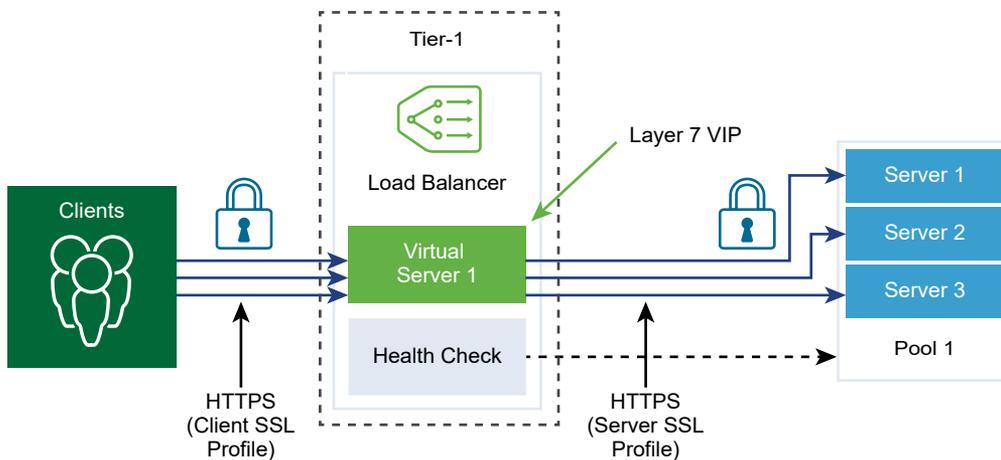


Figure 19-6. End-to-End SSL



Procedure

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

2 Select **Advanced Networking & Security > Networking > Load Balancer > 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
Session Caching	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.
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.

- 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

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](#).
- Verify that persistent profiles are available. See [Configure Persistent Profiles](#).
- Verify that SSL profiles for the client and server are available. See [Configure SSL Profile](#).
- Verify that server pools are available. See [Add a Server Pool for Load Balancing](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Networking > Load Balancer > 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
Maximum 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.
Maximum New Connection Rate	Set the maximum new connection to a server pool member so that a virtual server does not deplete resources.
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.

- 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

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](#).
- Verify that persistent profiles are available. See [Configure Persistent Profiles](#).
- Verify that SSL profiles for the client and server are available. See [Configure SSL Profile](#).
- Verify that server pools are available. See [Add a Server Pool for Load Balancing](#).
- 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](#).
- [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 **Advanced Networking & Security > Networking > Load Balancer > 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:sensitive)"` instead use `"Case ((?i:sensitive)"`.
- 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](#).

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
Maximum 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.
Maximum New Connection Rate	Set the maximum new connection to a server pool member so that a virtual server does not deplete resources.
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.

- 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. 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

Supported Match Condition	Description
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 Payload	Match an HTTP request body content. http_request.body_value - value to match
TCP Header Fields	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 Fields	Match an IP source or destination address. ip_header.source_address - source address to match ip_header.destination_address - destination address to match

Action	Description
HTTP Request URI Rewrite	Modify an URI. http_request.uri - URI (without query arguments) to write http_request.uri_args - URI query arguments to write
HTTP Request Header Rewrite	Modify value of an HTTP header. http_request.header_name - header name 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. http_request.method - value to match
HTTP Request URI	Match an HTTP request URI. http_request.uri - value to match
HTTP Request URI args	Match an HTTP request URI query argument. http_request.uri_args - 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 Payload	Match an HTTP request body content. http_request.body_value - value to match

Supported Match Condition	Description
TCP Header Fields	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 Fields	Match an IP source address. ip_header.source_address - source address to match
Action	Description
Reject	Reject a request, for example, by setting status to 5xx. http_forward.reply_status - HTTP status code used to reject http_forward.reply_message - HTTP rejection message
Redirect	Redirect a request. Status code must be set to 3xx. http_forward.redirect_status - HTTP status code for redirect 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. 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. http_response.header_name - header name to match http_response.header_value - value to match
Action	Description
HTTP Response Header Rewrite	Modify the value of an HTTP response header. http_response.header_name - header name 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](#).

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**.

Note If you use the **Advanced Networking & Security** user interface to modify objects created in the policy interface, some settings might not be configurable. These read-only settings have this icon next to them: . See [Chapter 1 Overview of the NSX Manager](#) for more information.

This chapter includes the following topics:

- [Add or Delete a Firewall Rule to a Logical Router](#)
- [Configure Firewall for a Logical Switch Bridge Port](#)
- [Firewall Sections and Firewall Rules](#)
- [About Firewall Rules](#)

Add or Delete a Firewall Rule to a Logical Router

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](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Select **Advanced Networking & Security > Networking > Routers**.
- 3 Click the **Routers** tab if it is not already selected.
- 4 Click the name of a logical router.
- 5 Select **Services > Edge Firewall**.

- 6 Click an existing section or rule.
- 7 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.

- 8 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

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](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > 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.

Enable and Disable Distributed Firewall

You can enable or disable the distributed firewall feature.

If it is disabled, no firewall rules are enforced at the dataplane level. Upon re-enablement rules are re-enforced.

Procedure

- 1 Navigate to **Advanced Networking & Security > Security > Distributed Firewall**.
- 2 Click the **Settings** tab.
- 3 Click Distributed Firewall **Edit**.
- 4 In the dialog box, toggle the firewall status to green (enabled) or gray (disabled).
- 5 Click **Save**.

Add a Firewall Rule Section

A firewall rule section is edited and saved independently and is used to apply separate firewall configuration to tenants.

Procedure

- 1 Select **Advanced Networking & Security > 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.

- 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.

- Enter the section name.
- 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.

- 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.

- Click **OK**.

What to do next

Add Firewall rules to the section.

Delete a Firewall Rule Section

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.

Procedure

- Select **Advanced Networking & Security > 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

You can enable or disable all rules in a firewall rule section.

Procedure

- 1 Select **Advanced Networking & Security > 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

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.

Procedure

- 1 Select **Advanced Networking & Security > 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

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 Manager and NSX Edge node virtual machines to the firewall exclusion list.

Procedure

- 1 Select **Advanced Networking & Security > 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. 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/dfwpktlogs.log file on ESX and KVM hosts.
Action	The action applied by the rule can be Allow , Drop , or Reject . The default is Allow .
IP Protocol	The options are IPv4 , IPv6 , and IPv4_IPv6 . The default is IPv4_IPv6 . To access this property, click the Advanced Settings icon.
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 traffic in both directions is checked. To access this property, click the Advanced Settings icon.
Rule Tags	Tags that have been added to the rule. To access this property, click the Advanced Settings 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

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.

Procedure

- 1 Select **Advanced Networking & Security > 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.
Contain er Objects	The available objects are IP Set, Logical Port, Logical Switch, and NS Group. Select the objects and click OK .

- 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.
Contain er Objects	The available objects are IP Set, Logical Port, Logical Switch, and NS Group. Select the objects and click OK .

- 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"> ■ ALG ■ ICMP ■ IGMP ■ IP ■ L4 Port Set
Protocol	Select one of the available protocols.
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

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.

Procedure

- 1 Select **Advanced Networking & Security > 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

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** (not recommended), 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.

Procedure

- 1 Select **Advanced Networking & Security > 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.

Note Selecting **Reject** as the action for the default rule is not recommended.

- 5 In the **Log**, enable or disable logging.
Enabling logging can affect performance.
- 6 Click **Publish**.

Change the Order of a Firewall Rule

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.

Procedure

- 1 Select **Advanced Networking & Security > 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

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.

Procedure

- 1 Select **Advanced Networking & Security > 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.

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.

This chapter includes the following topics:

- [View Monitoring Dashboards](#)
- [View the Usage and Capacity of Categories of Objects](#)
- [Checking the Realized State of a Configuration Change](#)
- [Search for Objects](#)
- [Filter by Object Attributes](#)
- [Add a Compute Manager](#)
- [Add an Active Directory](#)
- [Add an LDAP Server](#)
- [Synchronize Active Directory](#)
- [Managing User Accounts and Role-Based Access Control](#)
- [Backing Up and Restoring the NSX Manager](#)
- [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](#)
- [Recovering NSX-T When vCenter Server Is Lost and Cannot Be Recovered](#)
- [Multisite Deployment of NSX-T Data Center](#)
- [Configuring Appliances](#)
- [Add a License Key and Generate a License Usage Report](#)
- [Setting Up Certificates](#)

- [Compliance-Based Configuration](#)
- [Collect Support Bundles](#)
- [Log Messages and Error Codes](#)
- [Customer Experience Improvement Program](#)
- [Add Tags to an Object](#)
- [Find the SSH Fingerprint of a Remote Server](#)
- [View Data about Applications Running on VMs](#)
- [Configuring an External Load Balancer](#)

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 21-1. 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 21-1. 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 21-2. Networking & Security Dashboards

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 21-3. Advanced Networking & Security Dashboards

Dashboard	Sources	Description
Load Balancers	Advanced Networking & Security > Load Balancers	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	Advanced Networking & Security > Security > Distributed Firewall Advanced Networking & Security > Security > Bridge Firewall Advanced Networking & Security > Networking > 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	Advanced Networking & Security > Switching	Shows the status of logical switches and logical ports, including both VM and container ports.

Table 21-4. 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.

View the Usage and Capacity of Categories of Objects

You can view the usage and capacity of various categories of objects in your NSX-T Data Center environment. You can also set alerts to let you easily see when certain thresholds in usage are reached.

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 (realized) - The number of objects that have been successfully created or configured. This number reflects the NSX Manager objects that are shown in the **Advanced Networking & Security** tab. These objects can include some that you create in the **Networking, Security, Inventory, or System** tabs. A color-coded bar is displayed to indicate the usage percentage. If usage is below the warning alert level, the color is green. If usage is at or above the warning alert level but below the critical alert level, the color is orange. If usage is at or above the critical alert level, the color is red.
- Warning alert - This is the usage level at which the usage bar mentioned above will show an orange color. You can change this value.
- Critical alert - This is the usage level at which the usage bar mentioned above will show a red color. You can change this value.

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:

- Networking and security groups
- IP sets
- Groups based on IP sets
- vCenter clusters
- Hypervisor hosts

The system capacity page shows the following object categories:

- System-wide virtual interfaces
- Edge clusters
- System-wide edge nodes

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*.

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.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.

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
Name and Description	Type the name to identify the vCenter Server. You can optionally describe any special details such as, the number of clusters in the vCenter Server.
Domain Name/IP Address	Type the IP address of the vCenter Server.
Type	Keep the default option.
Username and Password	Type the vCenter Server login credentials.
Thumbprint	Type the vCenter Server SHA-256 thumbprint algorithm value.

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.

- 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.

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.

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 > 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. The event log reader account must have read permissions for security event logs.

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 > 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.

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 > 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.

Managing User Accounts and Role-Based Access Control

NSX-T Data Center appliances have two built-in users: admin and audit. You can integrate NSX-T Data Center with VMware Identity Manager (vIDM) and configure role-based access control (RBAC) for users that vIDM manages.

For users managed by vIDM, the authentication policy that applies is the one configured by the vIDM administrator, and not NSX-T Data Center's authentication policy, which applies to users admin and audit only.

Manage a User's Password

Each NSX Manager and NSX Edge appliance has three local accounts, admin, audit, and root. You can manage the password for these users but you cannot add or delete users.

The audit user is not active by default. To activate it, log in as admin and run the `set user audit` 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 a local user on the NSX Manager will expire within 30 days, the NSX Manager web interface displays a password expiration notification. If you set a local user's password expiration to 30 days or fewer the notification is always present.

Starting in NSX-T Data Center 2.5.1, the notification includes a "Change Password" link. Click the link to change the local user's password from the web interface.

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.
- 2 To change the password, run the `set user` command. For example:

```
nsx> set user admin
Current password:
New password:
Confirm new password:
nsx>
```

- 3 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>
```

- 4 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>
```

- 5 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, and Cloud Service Manager appliances.

Note If you have an NSX Manager cluster, resetting the password for the `root`, `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 are running NSX-T Data Center 2.5.2 or later and 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 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. If you are running NSX-T Data Center 2.5.0 or 2.5.1 and want to reset the password for `admin` and `audit`, use the following procedure as well. If you are running NSX-T Data Center 2.5.2 or later you can use the above procedure to reset the password for `admin` or `audit` after you reset the password for `root`.

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`.

- 8 If you are running NSX-T Data Center 2.5.0 or 2.5.1, perform the following steps:

- a Remove all options after `root=UUID=<ID number>` and add `rw single init=/bin/bash` after the UUID.

- b Press **Ctrl-X** to boot.

- c When the log messages stop, press Enter.

You will see the prompt `root@(none) :/#`.

- d If you are resetting the password for `root`, run the command `passwd`.

If you are resetting the password for `admin` or `audit`, run the command `passwd <admin or audit user ID>`.

You can run the `passwd` command multiple times.

- e Enter a new password and enter it again to confirm.

- f If you are resetting the password on an NSX Manager, run the command `touch /var/vmware/nsx/reset_cluster_credentials`.

- g Run the command `sync`.

- h Run the command `reboot -f`.

- 9 If you are running NSX-T Data Center 2.5.2 or later, perform the following steps:

- a Add `systemd.wants=PasswordRecovery.service` to the end of the line.

- b Press **Ctrl-X** to boot.

- c Enter a new password for `root` and enter it again to confirm.

After the boot process completes, 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
```

Obtain the Certificate Thumbprint from a vIDM Host

Before you configure the integration of vIDM with NSX-T, you must get the certificate thumbprint from the vIDM host.

You must use OpenSSL version 1.x or higher for the thumbprint. On the vIDM host, the command `openssl` runs an older OpenSSL version and therefore you must use the command `openssl1` on the vIDM host. This command is only available from the vIDM host.

On a server that is not the vIDM host, you can use the `openssl` command that is running OpenSSL version 1.x or higher.

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 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 get an error running the command, you might need to run `openssl` with the `sudo` command, that is, `sudo openssl`

- 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
```

Configure VMware Identity Manager 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.

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**.
- 3 Click the **Configuration** 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.
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.

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 recommended 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..

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
- Execute

- Read
- None

Full access gives the user all permissions. The execute permission includes the read permission.

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

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 21-5. Roles and Permissions](#) and [Table 21-6. Roles and Permissions for Advanced Networking and Security](#) 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 21-5. 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	FA	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
Networking > Segments	FA	R	FA	FA	R	R	FA	R	R	R	R	R	R
Networking > Segments > Segment Profiles	FA	R	FA	FA	R	R	FA	R	R	R	R	R	R

Table 21-5. 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
Networking > IP Address Pools	FA	R	FA	FA	R	R	FA	R	R	R	None	None	None
Networking Forwarding Policies	FA	R	FA	R	FA	R	FA	R	None	None	None	None	None
Networking > DNS	FA	R	FA	FA	R	R	FA	R	R	R	None	None	None
Networking > Load Balancing	FA	R	None	None	R	None	FA	R	FA	R	None	None	None
Networking > NAT	FA	R	FA	R	FA	R	FA	R	R	R	None	None	None
Networking > VPN	FA	R	FA	R	FA	R	FA	R	None	None	FA	None	None
Networking > IPv6 Profiles													
Security > Distributed Firewall	FA	R	R	R	FA	R	FA	R	R	R	R	R	R
Security > Gateway Firewall	FA	R	R	R	FA	R	FA	R	None	None	None	None	FA
Security > Network Introspection	FA	R	R	R	R	R	FA	R	None	None	None	None	FA

Table 21-5. 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
Security > Endpoint Protection Rules	FA	R	R	R	R	R	FA	R	None	None	None	FA	None
Inventories > Context Profiles	FA	R	FA	R	FA	R	FA	R	R	R	R	R	R
Inventories > Virtual Machines	R	R	R	R	R	R	R	R	R	R	R	R	R
Plan & Troubleshoot > Port Mirroring	FA	R	FA	R	R	R	FA	R	None	None	None	None	None
Plan & Troubleshoot > Port Mirroring Binding	FA	R	FA	FA	R	R	FA	R	R	R	R	R	R
Plan & Troubleshoot > Monitoring Profile Binding	FA	R	FA	FA	R	R	FA	R	R	R	R	R	R
Plan & Troubleshoot > Firewall IPFIX Profiles	FA	R	FA	R	FA	R	FA	R	R	R	R	R	R

Table 21-5. 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
Plan & Troubleshoot > 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	None	None	None	None	None
System > Fabric > Nodes > Nodes	FA	R	FA	R	FA	R	R	R	R	R	None	None	None
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

Table 21-5. 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 > 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
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

Table 21-5. 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 > 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
System > Utilities > Restore	FA	R	None	None	None	None	None	None	None	None	None	None	None
System > Utilities > Upgrade	FA	R	R	R	R	R	None	None	None	None	None	None	None
System > Users > Role Assignments	FA	R	None	None	None	None	None	None	None	None	None	None	None
System > Active Directory	FA	R	FA	R	FA	FA	R	R	R	R	R	R	R
System > Users > Configuration	FA	R	None	None	None	None	None	None	None	None	None	None	None

Table 21-5. 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 > License s	FA	R	R	R	R	R	None	Non e	Non e	No ne	None	None	None
System > System Adminis tration	FA	R	R	R	R	R	R	R	Non e	No ne	None	None	None
Custom Dashbo ard Configu ration	FA	R	R	R	R	R	FA	R	R	R	R	R	R
System > Lifecycl e Manage ment > Migrate	FA	Non e	None	None	None	None	None	Non e	Non e	No ne	None	None	None

Table 21-6. Roles and Permissions for Advanced Networking and Security

Operati on	EA	A	NE	NO	SE	SO	CS Adm	CS Aud	LB Ad m	LB Au d	VPN Adm	GI Adm	NI Adm
Tools > Port Connect ion	E	R	E	E	E	E	E	R	E	E	None	None	None
Tools > Traceflo w	E	R	E	E	E	E	E	R	E	E	None	None	None
Tools > Port Mirrorin g	FA	R	FA	R	R	R	FA	R	Non e	No ne	None	None	None
Tools > IPFIX	FA	R	FA	R	FA	R	FA	R	R	R	R	R	R

Table 21-6. Roles and Permissions for Advanced Networking and Security (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
Firewall > Distribu ted Firewall > General	FA	R	R	R	FA	R	FA	R	Non e	No ne	None	None	R
Firewall > Distribu ted Firewall > Configu ration	FA	R	R	R	FA	R	FA	R	Non e	No ne	None	None	None
Firewall > Edge Firewall	FA	R	R	R	FA	R	FA	R	Non e	No ne	None	None	FA
Routing > Routers	FA	R	FA	FA	R	R	FA	R	R	R	R	None	R
Routing > NAT	FA	R	FA	R	FA	R	FA	R	R	R	None	None	None
DHCP > Server Profiles	FA	R	FA	R	None	None	FA	R	Non e	No ne	None	None	None
DHCP > Servers	FA	R	FA	R	None	None	FA	R	Non e	No ne	None	None	None
DHCP > Relay Profiles	FA	R	FA	R	None	None	FA	R	Non e	No ne	None	None	None
DHCP > Relay Services	FA	R	FA	R	None	None	FA	R	Non e	No ne	None	None	None
DHCP > Metadat a Proxies	FA	R	FA	R	None	None	None	Non e	Non e	No ne	None	None	None
IPAM	FA	R	FA	FA	R	R	None	Non e	R	R	None	None	None
Switchin g > Switche s	FA	R	FA	FA	R	R	FA	R	R	R	R	None	R

Table 21-6. Roles and Permissions for Advanced Networking and Security (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
Switchin g > Ports	FA	R	FA	FA	R	R	FA	R	R	R	R	None	R
Switchin g > Switchin g Profiles	FA	R	FA	FA	R	R	FA	R	R	R	None	None	None
Networ king > Load Balance rs	FA	R	None	None	R	None	FA	R	FA	R	None	None	None
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 > IP Sets	FA	R	FA	R	FA	R	FA	R	R	R	R	R	R
Inventor y > IP Pools	FA	R	FA	R	None	None	None	None	R	R	R	R	R
Inventor y > 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

Table 21-6. Roles and Permissions for Advanced Networking and Security (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 > Virtual Machine s	R	R	R	R	R	R	R	R	R	R	R	R	R
Inventor y > Virtual Machine s > Configu re Tags	FA	Non e	None	None	None	None	None	Non e	Non e	No ne	None	None	None

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. 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>
```

Prerequisites

- If you want to assign roles to users, verify that a vIDM host is associated with NSX-T. For more information, see [Configure VMware Identity Manager Integration](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **System > Users**.
- 3 To assign roles to users, select **Add > Role Assignment**.
 - 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 (Optional) If using NSX Cloud, log in to the CSM appliance instead of NSX Manager and repeat steps 1 through 4.

- 6 If the certificate for the principal identity expires, perform the following steps:
 - a Import a new certificate and note the certificate's ID. See [Import a Certificate](#).
 - b Call the following API to get the ID of the principal identity.

```
GET https://<nsx-mgr>/api/v1/trust-management/principal-identities
```

- c 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"
}
```

Backing Up and Restoring the NSX Manager

If the NSX Manager cluster becomes inoperable, or if you want to restore your environment to a previous state, you can restore from a backup. While the NSX Manager is inoperable, the data plane is not affected, but you cannot make configuration changes.

There are two types of backups:

Cluster backup

This backup includes the desired state of the virtual network.

Node backup

This is a backup of the NSX Manager nodes.

There are two backup methods:

Manual

You manually run the backup at any time.

Automated

Automated backups run based on a schedule that you set. Automated backups are highly recommended to ensure that you have up-to-date backups.

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 NSX Manager appliances running the same version of NSX Manager as the appliances that were backed up.

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 configure a schedule for automatic backups.

Prerequisites

Verify that you have the SSH fingerprint of the backup file server. Only an SHA256 hashed ECDSA (256 bit) key is accepted as a fingerprint. See [Find the SSH Fingerprint of a Remote Server](#).

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **System > Backup & Restore**.
- 3 Click **Edit** in the upper right of the page to configure backups.
- 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 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 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 you have multiple NSX-T Data Center deployments, you must use a different directory for each deployment. 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, click the **Change Encryption Passphrase** toggle and enter the 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** 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 set the interval between the backups triggered by configuration changes. The default is 5 minutes.

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.

Removing 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.

Listing Available Backups

The backup file server stores backups from all the NSX Managers. 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 is located on an NSX Manager. The full path name is `/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, you should copy this script after installing NSX-T Data Center to a machine that is not an NSX Manager so that you can run this script even if all the NSX Managers become inaccessible. If you need to restore a backup but have no access to this script, you can install a new NSX Manager 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 as admin and running a CLI command. For example:

```
nsxmgr-1> copy file get_backup_timestamps.sh url scp://admin@10.127.1.20/tmp/
admin@10.127.1.20's password:
nsxmgr-1>
```

The script is interactive and will prompt 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 node's IP address or FQDN if the NSX 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:
10.108.115.108
Enter port:
22
Enter directory path:
/home/nsx/backups
Enter number of latest backup or press Enter to list all backups:

root@10.108.115.108's password:
Latest backups:
[Backup timestamp; IP address/FQDN; Node id]
2019-01-22;09:00:33 10.196.196.77 35163642-6623-8f6d-7af0-52e03f16faed
2019-01-22;09:01:52 10.196.196.77 35163642-6623-8f6d-7af0-52e03f16faed
2019-01-22;09:13:30 wdc75.platformqe.com 41893642-597b-915f-5117-7da576df4ff2
2019-01-22;09:14:42 wdc75.platformqe.com 41893642-597b-915f-5117-7da576df4ff2
2019-01-22;09:16:43 wdc75.platformqe.com 41893642-597b-915f-5117-7da576df4ff2
```

Restore a Backup

Restoring a backup results in restoring the state of the network at the time of the backup. In addition, the configurations maintained by the NSX Manager are also restored and any changes,

such as adding or deleting nodes, that were made to the fabric since the backup was taken are reconciled.

You must restore the backup to a new NSX Manager appliance.

If you had an NSX Manager cluster when the backup was taken, you should also restore to an NSX Manager cluster. The restore process restores one NSX Manager node first and then prompts you to add the other NSX Manager nodes.

Important If any nodes in the NSX Manager cluster are still available, you must power them off before you start the restore.

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 (256 bit) 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.
- 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 Manager node that took the backup.
- If you configure the NSX Manager nodes to publish their FQDN, you must configure the forward and reverse lookup entries for the NSX Manager nodes on the DNS server.

Procedure

- 1 Power off all nodes in the NSX Manager cluster that you are restoring.
- 2 Install one new NSX Manager 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 node with the same IP address. Do not configure the NSX Manager node to publish its FQDN.

```
2019-01-22;09:01:52 10.196.196.77 35163642-6623-8f6d-7af0-52e03f16faed
```

- If the backup listing for the backup you are restoring contains an FQDN, you must configure the new NSX Manager node with this FQDN (see the section "Publishing the FQDNs of the NSX Managers" in the topic "NSX Manager Installation" in the *NSX-T Data Center Installation Guide* for more information). In addition, if the new NSX 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 NSX Manager node with the new IP address.

```
2019-01-22;09:16:43 nsxmgr.example.com 41893642-597b-915f-5117-7da576df4ff2
```

After the new NSX Manager node is running and online, you can proceed with the restore.

- 3 From your browser, log in with admin privileges to the new NSX Manager.

- 4 Select **System > Backup & Restore**.
- 5 Click the **Restore** tab.
- 6 To configure the backup file server, click **Edit**.
- 7 Enter the IP address or host name.
- 8 Change the port number, if necessary.
The default is 22.
- 9 To log in to the server, enter the user name and password.
- 10 In the **Destination Directory** text box, enter the absolute directory path where the backups are stored.
- 11 Enter the passphrase that was used to encrypt the backup data.
- 12 Enter the SSH fingerprint of the server that stores the backups.
- 13 Click **Save**.
- 14 Select a backup.
- 15 Click **Restore**.

The status of the restore operation is displayed. If you have deleted or added fabric nodes or transport nodes since the backup, you are prompted to take certain actions, for example, log in to a node and run a script.

If the backup has information about an NSX Manager cluster, you are prompted to add NSX Manager nodes. If you decide not to add NSX Manager nodes, you can still proceed with the restore.

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.

If the restore failed, 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 failed, 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 checking 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.

- 16 If you have only one node deployed, after the restored NSX Manager node is up and functional, you can deploy additional nodes to form a NSX Manager cluster.

See the *NSX-T Data Center Installation Guide* for instructions.

- 17 After the new NSX Manager cluster is deployed, delete the original NSX Manager cluster VMs that you powered down in Step 1.

You must also replace the certificates on the second and third node of the cluster.

Results

If you added a compute manager after the backup, and you try to add the compute manager again after the restore, you will get an error message indicating that registration failed. You can 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 the 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](#).

Backup and Restore During Upgrade

The Management Plane stops responding during the upgrade process and you need to restore a backup that was taken while the upgrade was in progress.

Problem

The Upgrade Coordinator has been upgraded and the Management Plane stops responding. You have a backup that was created while the upgrade was in progress.

Solution

- 1 Deploy your Management Plane node with the same IP address that the backup was created from.
- 2 Upload the upgrade bundle that you used at the beginning of the upgrade process.
- 3 Upgrade the Upgrade Coordinator.
- 4 Restore the backup taking during the upgrade process.
- 5 Upload a new upgrade bundle if necessary.
- 6 Continue with the upgrade process.

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.nuxt` 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	ADDRESS	PORT	FQDN	UUID	IP
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-afeb-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	ADDRESS	PORT	FQDN	UUID	IP
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	ADDRESS	PORT	FQDN	UUID	IP
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

```

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. 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.

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.

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**.

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"
}

```

Recovering NSX-T When vCenter Server Is Lost and Cannot Be Recovered

If vCenter Server (VC) is lost and cannot be recovered (perhaps because there is no backup or the backup is corrupted), use the following procedure to recover the NSX-T environment after you re-deploy VC.

The new VC must have the same FQDN and IP address as the original VC. Also, it must have the same clusters containing the same hosts. Be careful with hosts that have power-on VMs when you add them to VC. Make sure that they are added to the correct clusters and not to the VC datacenter.

Compute Manager

In NSX Manager, delete the old computer manager. Then add the new VC as a computer manager.

Host Transport Nodes

In NSX Manager, the hosts will appear in the correct VC clusters. No action is needed.

Edge Nodes

You must replace Edge nodes that were deployed from the NSX Manager UI.

- 1 Follow the procedure in [Replace an NSX Edge Transport Node Using the NSX Manager UI](#) to replace an Edge node.
- 2 Verify that the gateways (or logical routers) and tunnels are configured on the new Edge VM.
- 3 Delete the old edge node by going to **System > Fabric > Edge Transport Nodes**. Select the Edge node and click **Actions > Delete**. Errors such as "Power off failed" can be ignored.
- 4 In VC, power off the old Edge VM and delete it.
- 5 Repeat the steps above for each of the Edge nodes.

NSX Manager

You must replace NSX Managers that were deployed from the NSX Manager UI. Typically, the second and third NSX Managers are deployed this way.

- 1 Log in to the first NSX Manager's UI.
- 2 Go to **System > Appliances** and select the third NSX Manager. Click **Actions > Delete**. This will fail because the Manager VM cannot be powered off. The force delete option will now be available. Select **Actions > Force delete**.
- 3 If the force delete does not work, do the following:
 - a Log in to the first NSX Manager's CLI.

- b Run the command `get cluster status` to get the UID of the third NSX Manager.
- c Run the command `detach node <node-uid>` to detach the third NSX Manager from the cluster.
- d Make the following API call to force delete the third NSX Manager:

```
POST : https://<nsx-manager-1>/api/v1/cluster/nodes/deployments/<node-uid>?
action=delete&force_delete=true
```

- 4 In VC, power off and delete the third NSX Manager.
- 5 Deploy a new NSX Manager with the same configuration as the third NSX Manager.
- 6 Repeat the steps above to delete the second NSX Manager.
- 7 Deploy two new NSX Managers.

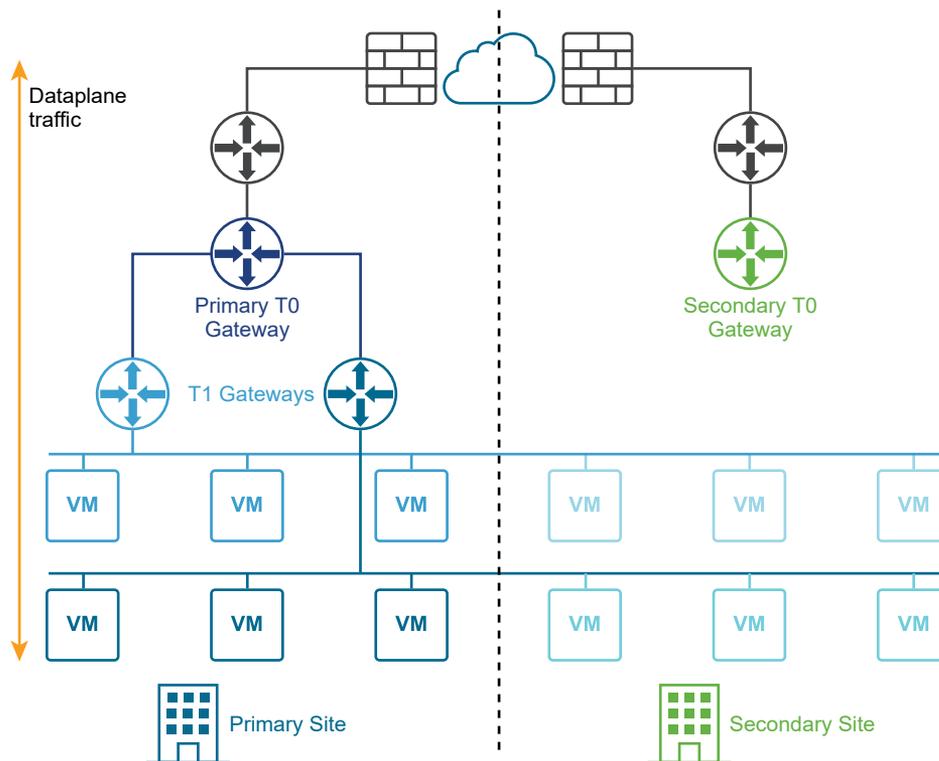
Multisite Deployment of NSX-T Data Center

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:

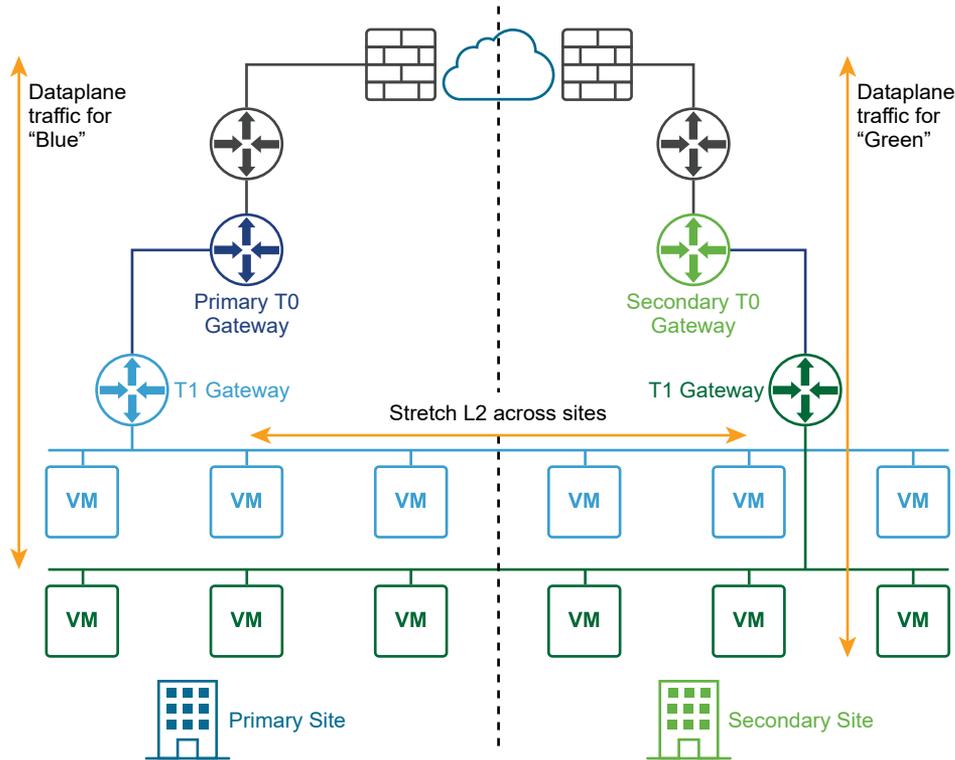
- Disaster recovery
- Active-active

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.

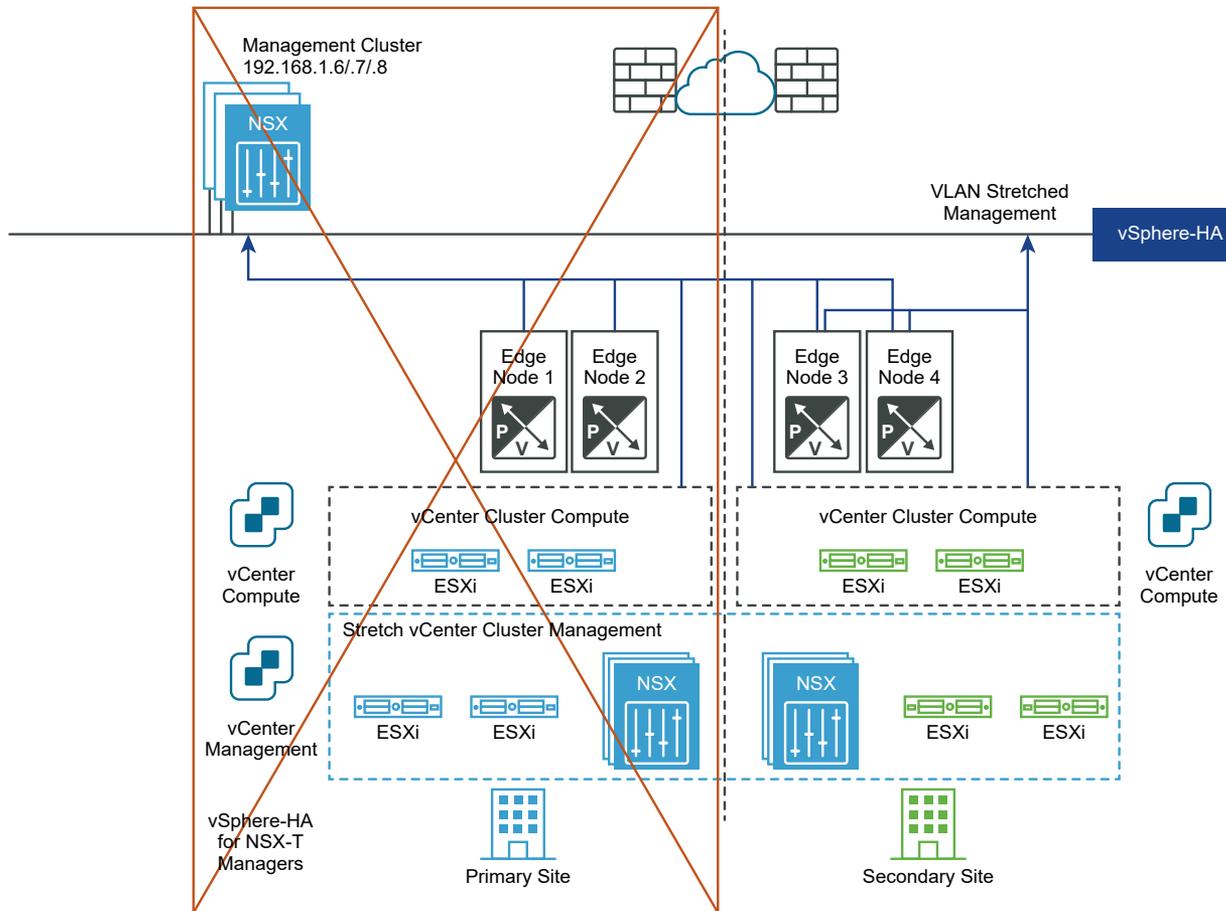
Automatic Recovery of the Management Plane

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 diagram illustrates automatic recovery of the management plane.



Automatic Recovery of the Data Plane

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.

Note: The failover mode of the tier-1 gateway can be preemptive or non-preemptive.

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>
```

```

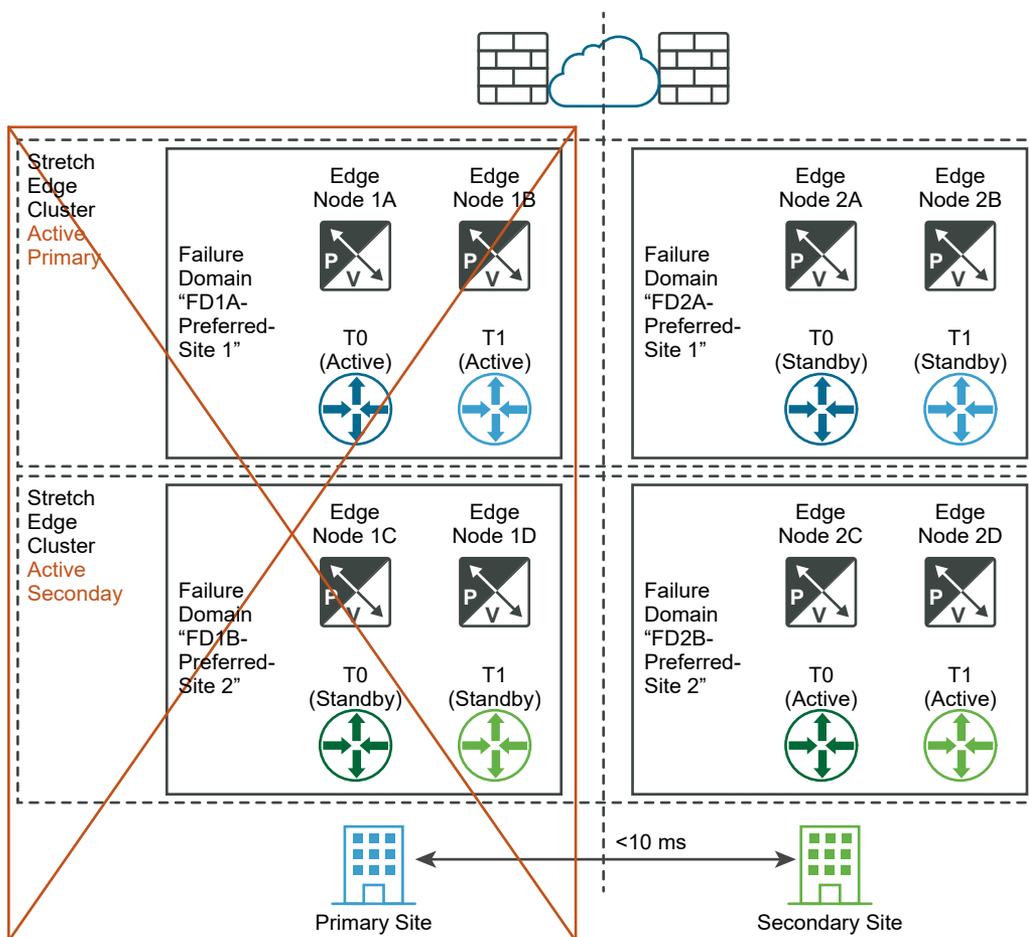
{
  "_revision": 0,
  "id": "bf8d4daf-93f6-4c23-af38-63f6d372e14e",
  "resource_type": "EdgeCluster",
  ...
  "allocation_rules": [
    {
      "action": {
        "enabled": true,
        "action_type": "AllocationBasedOnFailureDomain"
      }
    }
  ],
}

```

- Create tier-0 and tier-1 gateways using the API or NSX Manager UI.

When an Edge node in the primary site fails, the tier-0 and tier-1 gateways hosted on that node will be migrated to an Edge node in the secondary site.

The following diagram illustrates automatic recovery of the data plane.



Manual/Scripted Recovery of the Management Plane

Requirements:

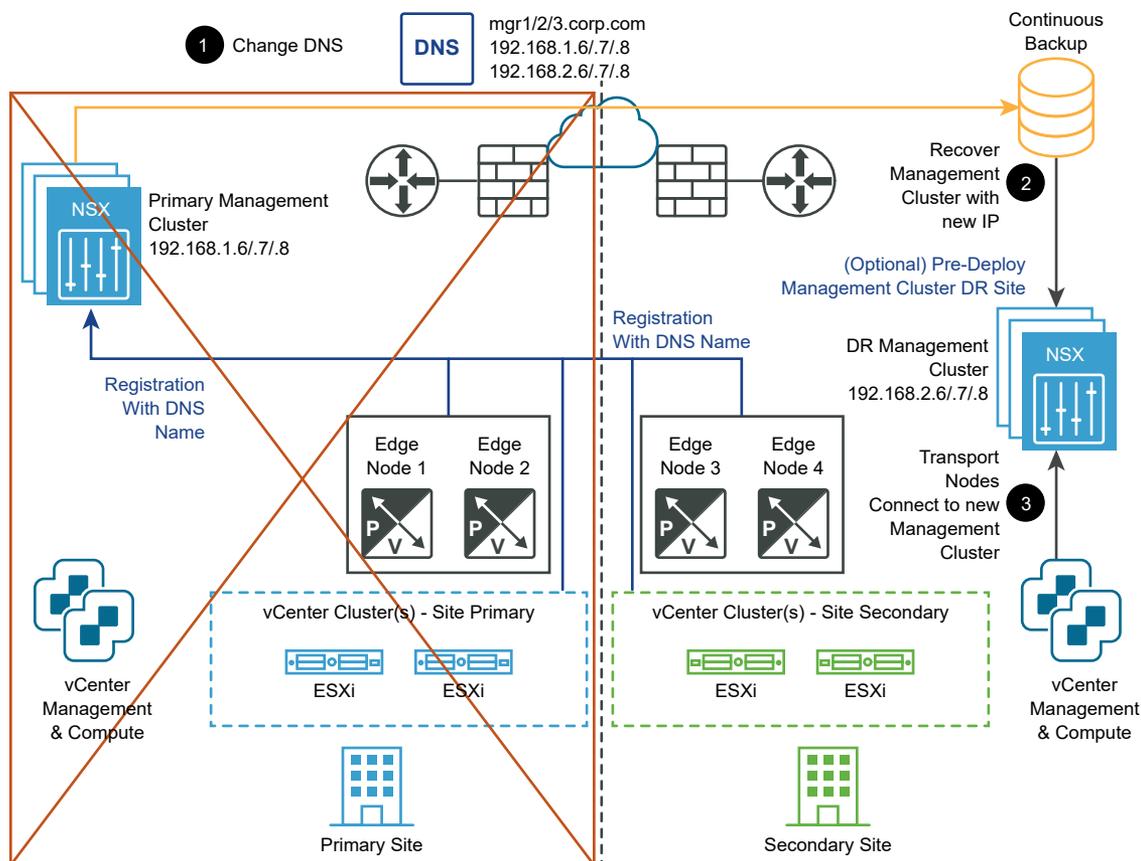
- DNS for NSX Managers with a short TTL (for example, 5 minutes).
- Continuous backup.

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.



- MTU must be at least 1600. 9000 is recommended.

NSX Manager Configuration

- Automatic backup when NSX-T Data Center configuration changes must be enabled.
- NSX Manager must be set up to use FQDN.

Data Plane Recovery

- The same internet provider must be used if public IP addresses are exposed through services such as NAT or load balancer.
- The HA mode for the tier-0 gateway must be active-standby, and the failover mode must be preemptive.

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.

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 21-7. System configuration commands and API requests.

Task	Command Line (NSX Manager and NSX Edge)	API Request (NSX Manager only)
Set system timezone	<code>set timezone <timezone></code>	<code>PUT https://<nsx-mgr>/api/v1/node</code>
Set NTP Server	<code>set ntp-server <ntp-server></code>	<code>PUT https://<nsx-mgr>/api/v1/node/ services/ntp</code>
Set a DNS server	<code>set name-servers <dns-server></code>	<code>PUT https://<nsx-mgr>/api/v1/node/ network/name-servers</code>
Set DNS Search Domain	<code>set search-domains <domain></code>	<code>PUT https://<nsx-mgr>/api/v1/node/ network/search-domains</code>

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 non-evaluation NSX-T Data Center 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 Advanced (available from NSX-T Data Center 2.5.1)
- NSX Enterprise (available from NSX-T Data Center 2.5.1)

When you install NSX Manager, a pre-installed evaluation license becomes active and is valid for 60 days. The evaluation license provides all the features of an enterprise license. You cannot install or unassign an evaluation license. You can assign a new evaluation license when the default evaluation license is present. The new evaluation license will override the default evaluation license. You can also unassign the non-default evaluation license. In that case, the default evaluation license will be restored.

You can install one or more of the non-evaluation licenses, but for each type, you can only install one key. When you install a standard, advanced, or enterprise license, the evaluation license is no longer available. You can also unassign non-evaluation licenses. If you unassign all non-evaluation licenses, the evaluation license is restored.

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.

If your license will expire within 60 days or if it has expired, after you log in to NSX Manager, a notification window will appear to inform you of the situation. You can also click the notification icon in the upper right corner of the window to see the notification.

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
- DFW
- Context Aware Micro-Segmentation - Application identification
- Context Aware Micro-Segmentation - Identity firewall for remote desktop session host
- Service Insertion
- Identity Firewall
- Enhanced Guest Introspection

Note The following features are disabled for the Limited Export Release version:

- IPSec VPN
 - HTTPS-based Load Balancer
-

Setting Up Certificates

You can import certificates, create a certificate signing request (CSR), generate self-signed certificates, and import a certificate revocation list (CRL).

After you install NSX-T Data Center, the manager nodes and cluster have self-signed certificates. To improve security, it is highly recommended that you replace the self-signed certificates with CA-signed certificates.

Import a Certificate

You can import a certificate with a private key to replace the default self-signed certificate after activation.

Note that only RSA-based certificates are supported.

Prerequisites

Verify that a 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 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.
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**.

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.
- Verify that the public and private key pairs are 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 **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.

Option	Description
Organization Unit	Enter the department in your organization that is handling this certificate For example, IT department.
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. DSA encryption - is used for digital signatures. Therefore, it is faster than RSA when creating an encrypted token but slower to analyze and validate this token. This encryption is faster to decrypt and slower 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. 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**.

8 Select **Download CSR PEM** from the drop-down menu.

You can save the CSR PEM file for your records and CA submission.

9 Use the contents of the CSR file to submit a certificate request to the CA in accordance with the CA enrollment process.

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.

Import a CA Certificate

You can import a signed CA certificate. After the import and activation, other certificates signed by that CA will be trusted by NSX-T Data Center.

Note that only RSA-based certificates are supported.

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. Set to No if this certificate is for the NSX Manager nodes.

- 4 Click **Import**.

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 10 years.

7 Click **Add**.

Results

The self-signed certificate appears in the **Certificates** tab.

Replace the Certificate for an NSX Manager Node or an NSX Manager Cluster Virtual IP

You can replace the certificate 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. To improve security, it is highly recommended that you replace the self-signed certificates with CA-signed certificates and that you use a different certificate for each node.

Prerequisites

Verify that a certificate is available in the NSX Manager. See [Import a 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 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.

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	<p>Copy all of the items in the CRL and paste them in this section.</p> <p>A sample CRL.</p> <pre>-----BEGIN X509 CRL----- MIIBODCB4zANBgkqhkiG9w0BAQQFADBGMQswCQYDVQQGEwJBVTEMMMAoGA1 UECBMD UUxEMRkwFwYDVQQKEwBNaW5jb20gUHR5LiBMdGQuMQswCQYDVQQLEwJDUz EbMBkG A1UEAxMSU1NMZWF5IGR1bW8gc2VydMvYFw0wMTAxMTUxNjI2NTdaFw0wMT AyMTQx NjI2NTdaMFkwEgYIBARcNOTUxMDA5MjMzMjA1WjASAgEDFw05NTEyMDEwMT AwMdBa MBMCAhI0Fw0wMTAxMTUxNjE5NDdaMBMCAhI1Fw0wMTAxMTUxNjIzNDZaMA OGCSqG SIB3DQEBAUAA0EAHPjQ3M93QOj8Ufi+jZM7Y78TfAzG4jJn/ E6MYBPFVQFYo/Gp UZexfjSVo5CIyyS0tYscz8o07avwBxTimpDEQg== -----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*.

Import a Certificate for a CSR

You can import a signed certificate for a CSR.

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 > 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.

Storage of Public Certificates and Private Keys

Public certificates and private keys are stored on the NSX Managers. When a load balancer or a 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.

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 in NSX-T Data Center 2.5:

- 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.

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 21-8. Compliance Report Information

Compliance Report Column	Description	Example
Non Compliance Code	Code to identify the type of non-compliance.	72301
Description	Description of the type of non-compliance.	Certificate is not CA signed.
Resource Name	Name or ID of the affected resource.	nsx-manager-1
Resource Type	Type of resource affected.	CertificateComplianceReporter
Affected Resources	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 21-9. Compliance Report Codes

Code	Description	Compliance Status Source	Remediation
72001	Encryption is disabled.	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. This status affects IPSec VPN session configurations which use the reported non-compliant configurations.	To remediate this status, add a VPN IPSec Profile that uses compliant encryption algorithms and use the profile in all VPN configurations. See Add IPSec Profiles .
72011	BGP messages with neighbor bypass integrity check. No message authentication defined.	This status is reported if no password is configured for BGP neighbors. This status affects the BGP neighbor configuration.	To remediate this status, configure a password on the BGP neighbor and update the tier-0 gateway configuration to use the password. See Configure BGP .
72012	Communication with BGP neighbor uses weak integrity check. MD5 is used for message authentication.	This status is reported if MD5 authentication is used for the BGP neighbor password. This status affects the BGP neighbor configuration.	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.	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. This status affects the following configurations: <ul style="list-style-type: none"> ■ Load balancer pools that are associated with HTTPS monitors. ■ Load balancer virtual servers that are associated with load balancer client SSL profiles or server SSL profiles. 	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 Add an SSL Profile .

Table 21-9. 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.	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. This status affects the following configurations: <ul style="list-style-type: none"> ■ Load balancer pools that are associated with HTTPS monitors. ■ Load balancer virtual servers that are associated with load balancer client SSL profiles or server SSL profiles. 	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 Add an SSL Profile .
72023	Weak Diffie-Hellman group is used.	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. This status affects IPSec VPN session configurations which use the reported non-compliant configurations.	To remediate this status, configure the VPN Profiles to use Diffie-Hellman group 19, 20, or 21. See Adding Profiles .
72024	Load balancer FIPS global setting is disabled.	This error is reported if the load balancer FIPS global setting is disabled. This status affects all load balancer services.	To remediate this status, enable FIPS for load balancer. See Configure Global FIPS Compliance Mode for Load Balancer .

Table 21-9. Compliance Report Codes (continued)

Code	Description	Compliance Status Source	Remediation
72200	Insufficient true entropy available.	This status is reported when a pseudo random number generator is used to generate entropy rather than relying on hardware-generated entropy. Hardware-generated entropy is not used because the NSX Manager node does not have the required hardware acceleration support to create sufficient true entropy.	To remediate this status, you might need to use newer hardware to run the NSX Manager node. Most recent hardware supports this feature. Note If the underlying infrastructure is virtual, you will not get true entropy.
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.	This status is reported when one of the NSX Manager certificates is not CA signed. NSX Manager uses the following certificates: <ul style="list-style-type: none"> ■ Syslog certificate. ■ API certificates for the individual NSX Manager nodes. ■ Cluster certificate used for the NSX ManagerVIP. 	To remediate this status, install CA-signed certificates. See Setting Up Certificates .

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 with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 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
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. For example:

```
<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"
```

Each API call produces one audit log message. 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>`.

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 syslog 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 vrl1.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 vrl1.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 - EKU, 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@caserver:~/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@caserver:~/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-----
< MIIICjANBgkqhkiG9w0BAQEFAAOCAg8AMIICCgKCAgEAv3yH7pZidfkLrEP3zVa9
< EcOKXlFFjkThZRZMfguenlm8s6QHYYVvuUX8IRB48Li3/DUfOj0bzaPWktpv+Q2P0
< N/j4LoX2RzjV/DPxYfLP6GMNMc21L3s9ruBeWUtthUP8khCwd2d2rZ09cUZV10P9
< kIYBb5RMFC7Z10Uth3bKdepEf+sXz3DaKZ/WySzYq9x86QDaA3ABO3Q0i7txBscI
< FvXuMDOMQaC3pPp9FWO6IPRAWB57wahLJv6K5qGIfwubSBFg53grT4snf11DZAhZ
< 9hz5JgGr80GVyWyb7rgigpl9iUWAZx8U9De9XoxmVBN5iEGTIuKGAegICL176crb
< RMkhjnCqNHI+z6sQvpYJ7U0zZc72eBIWoHUkcWWk3eU6Oy4OiyW6jYuXG7hZY1ly
< nSkme3mZUWJKvcoX05+3zeCP623/Hze7X2sNyWFjzeF3XEvauZrIbsJh/xp2ShDa
< uKKEY0gUGhLtCa3TpV9l8d6tFWVy8XjVjdjoVt4s7MfUo/airVmRykfsWrKyNUOQ
< qRZvSbqjt8pm+3bSvKdXX4ul7ptPG2GF20ETWHPwj2JwQpGhR9zK8fsKzvm6hXi
< 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
< MIIICjANBgkqhkiG9w0BAQEFAAOCAg8AMIICCgKCAgEAv3yH7pZidfkLrEP3zVa9
< EcOKXlFFjkThZRZMfguenlm8s6QHYYVvuUX8IRB48Li3/DUfOj0bzaPWktpv+Q2P0
< N/j4LoX2RzjV/DPxYfLP6GMNMc21L3s9ruBeWUtthUP8khCwd2d2rZ09cUZV10P9
< kIYBb5RMFC7Z10Uth3bKdepEf+sXz3DaKZ/WySzYq9x86QDaA3ABO3Q0i7txBscI
< FvXuMDOMQaC3pPp9FWO6IPRAWB57wahLJv6K5qGIfwubSBFg53grT4snf11DZAhZ
< 9hz5JgGr80GVyWyb7rgigpl9iUWAZx8U9De9XoxmVBN5iEGTIuKGAegICL176crb
< RMkhjnCqNHI+z6sQvpYJ7U0zZc72eBIWoHUkcWWk3eU6Oy4OiyW6jYuXG7hZY1ly
< nSkme3mZUWJKvcoX05+3zeCP623/Hze7X2sNyWFjzeF3XEvauZrIbsJh/xp2ShDa
< uKKEY0gUGhLtCa3TpV9l8d6tFWVy8XjVjdjoVt4s7MfUo/airVmRykfsWrKyNUOQ
< qRZvSbqjt8pm+3bSvKdXX4ul7ptPG2GF20ETWHPwj2JwQpGhR9zK8fsKzvm6hXi
< kq76zI4FefuVps3e1r39+0F+p6d6i2oUoo24sC1iSePTDhU74efVp6iv8HmnDgYX
< Ylm6Kusr0JT5TJFDfASmrj8CAwEAAQ==
---
> MIIICjANBgkqhkiG9w0BAQEFAAOCAg8AMIICCgKCAgEAqvsjay7+o7gCW7szT3ho
> bc34XX2l6u5Jl4/X/pUDI/YHmIf06bsZ1r/14bTL4Q7BM6+9MI6UYEE7DxUoINGO
> o4FEEQE32KwVFe3gw3homHU39q4pQjsJsxTcTE3oDMLIY0nWJ0PRUst3DdfyUH1L
> W0NUN9ydn+fa12Uf021iuDqVy9V8AH3ON6fu+QCA8nt71ZkzeTxSA0ldp12NA17F
> rD8rm05wxnV7WtuV7V8PstISiClzhHgZRM1+B0r300itnyAzEGLaRT3//PKfe00e
> HCdxGmlrUtMqxIITJahEsqvMufyqNYecVscyXLHPelizKCsQfy8c08LnznG8VAdc
> YILSn3uYGZap6aF1SgVxsvZicwv1YnssmgE13Af0nScmfM96k9h5j0HVEkWK608v
> oT5DGG1kVL2Qly97x0b6EnzUorzivv5zJMKvFcOektR8HdMHQit5uvmMRY3S5zow

```

```
> FtvfSDfWxxKyTy6GBrpP+8F+Jq91yGy/qa9lhKBzT2lg+rJp7T8k7/Nm9Tjyx7jL
> EqgEKZEL4chxpo8ucF98hbxXWRuaFHC2iDzGuUmuS1FfjVvHTuIbEMQfjapLZrHx
> 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 syslog 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>
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 21-10. 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
ROUTING	Logical router Logical router ports Static routing Dynamic routing NAT
FIREWALL	Firewall rules Firewall rule sections
FIREWALL-PKTLOG	Firewall connection logs Firewall packet logs
GROUPING	IP sets Mac sets NSGroups NSServices NSService groups VNI Pool IP Pool
DHCP	DHCP relay
SYSTEM	Appliance management (remote syslog, ntp, etc) Cluster management Trust management Licensing User and roles Task management Install Upgrade (NSX Manager, NSX Edge and host-packages upgrades) Realization Tags

Table 21-10. Log Message IDs (continued)

Message ID	Examples
MONITORING	SNMP Port connection 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 `.vmtx` file.

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.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**.

Add Tags to an Object

You can add tags to objects to make searching easier. When you specify a tag, you can also specify a scope.

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.

Most objects can have a maximum of 30 tags. For the following objects, the maximum is lower because of tags that are created and used internally.

Table 21-11. Maximum number of tags for objects created using the Advanced Networking & Security tab

Object	Maximum Number of Tags
virtual machine	25
Logical Port	29

Table 21-12. Maximum number of tags for objects created using the Networking, Security, or Inventory tabs

Object	Maximum Number of Tags
Group	29
Segment	27
Segment Port	29
Logical Router Port	30 - number of labels
NAT Rule	27
IPSec VPN Session	29

Table 21-13. Maximum number of tags for Cloud Service Manager objects

Object	Maximum Number of Tags
BFD Health Monitoring Profile, Transport Zone, Uplink Host Switch Profile, Transport Node, Edge Cluster	23

Table 21-14. Maximum number of tags for Public Cloud Manager objects

Object	Maximum Number of Tags
BFD Health Monitoring Profile, Transport Zone, Logical Switch, Node, Transport Node, Edge Cluster, Logical Router, Logical Router Uplink Port, Static Route, DHCP Profile, NSGroup, Firewall Section Rule List	23
NAT Rule	20
IP Set, NSGroup	22

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, go to the **Segments** tab and edit a segment.
- 3 Go the **Tags** field and add tags.
Each tag has a tag value, which is required, and a scope value, which is optional. The maximum length of a tag is 256 characters. The maximum length of a scope is 128 characters.
- 4 Click **Save**.

Find the SSH Fingerprint of a Remote Server

Some API requests that involve copying files to or from a remote server require that you provide the SSH fingerprint for the remote server in the request body. The SSH fingerprint is derived from a host key on the remote server.

To connect via SSH, the NSX Manager and the remote server must have a host key type in common. If there are multiple host keys types in common, whichever one is preferred according to the HostKeyAlgorithm configuration on the NSX Manager is used.

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 if they can 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.

The following table lists what NSX Manager supports in order from more preferred to less preferred.

Table 21-15. NSX Manager Host Keys in Preferred Order

Host key types supported by NSX Manager	Default Location of the Key
ECDSA (256 bit)	/etc/ssh/ssh_host_ecdsa_key.pub
ED25519	/etc/ssh/ssh_host_ed25519_key.pub

Procedure

- 1 Log in to the remote server as root.
Logging in using a console is more secure than over the network.
- 2 List the public key files in the `/etc/ssh` directory.

```
$ ls -al /etc/ssh/*pub
-rw-r--r-- 1 root root 601 Apr  8 18:10 ssh_host_dsa_key.pub
-rw-r--r-- 1 root root  93 Apr  8 18:10 ssh_host_ed25519_key.pub
-rw-r--r-- 1 root root 393 Apr  8 18:10 ssh_host_rsa_key.pub
```

- 3 Compare the available keys to what NSX Manager supports.
In this example, ED25519 is the only acceptable key.
- 4 Get the fingerprint of the key.

```
# awk '{print $2}' /etc/ssh/ssh_host_ed25519_key.pub | base64 -d | sha256sum -b | sed
's/ .*$/' | xxd -r -p | base64 | sed 's/./44g' | awk '{print "SHA256:"$1}'
SHA256:KemgftCfsd/hn7EEflhJ4m1698rRhMmNN2IW8y9iq2A
```

View Data about Applications Running on VMs

You can view information about applications running on VMs that are members of an NSGroup. This is a technical preview feature.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Select **Advanced Networking & Security > Inventory > Groups**.
- 3 Click the name of an NSGroup.
- 4 Click the **Applications** tab.
- 5 Click **COLLECT APPLICATION DATA**.

This process can take a few minutes. When the process is completed, the following information is displayed:

- The total number of processes.

- Circles representing various tiers, for example, web tier, database tier, and application tier. Also displayed is the number of processes in each tier.

6 Click a circle to see more information about the processes in that tier.

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.

Note that 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.

When accessing NSX Manager from a browser through an external load balancer, session persistence must be enabled on the load balancer.

When accessing NSX Manager from an API client through an external load balancer, four authentication methods are available (see the *NSX-T Data Center API Guide* for more information):

- HTTP Basic Authentication - Load balancer session persistence is not required.
- Client Certificate Authentication - Load balancer session persistence is not required.
- Authenticating to vIDM - Load balancer session persistence is not required.
- Session-Based Authentication - Load balancer session persistence is required.

Recommendation:

- Configure a single IP on the external load balancer for both browser and API access. The load balancer must have session persistence enabled.

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](#).

This chapter includes the following topics:

- [A Quick Tour of the Cloud Service Manager](#)
- [Threat Detection using the NSX Cloud Quarantine Policy](#)
- [NSX Enforced Mode](#)
- [Native Cloud Enforced Mode](#)
- [NSX-T Data Center Features Supported with NSX Cloud](#)
- [Frequently Asked Questions \(FAQs\)](#)

A Quick Tour of the Cloud Service Manager

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**:

Clouds > Overview

Access your public cloud account by clicking **Clouds**.

Overview: Each tile on this screen 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 and indicates the number of VMs that are managed by NSX-T Data Center and their state.</p> <p>See Add your Public Cloud Account in the <i>NSX-T Data Center Installation Guide</i> for detailed 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 PCG or Undeploy PCG in the <i>NSX-T Data Center Installation Guide</i> for detailed instructions.</p>
Enable or Disable Quarantine Policy	<p>You can enable or disable Quarantine Policy. See Threat Detection using the NSX Cloud Quarantine Policy 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 holistic 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**.

Clouds > {Your Public Cloud} > Accounts

The Accounts section of CSM provides information on the public cloud accounts you have already added.

Each card represents a public cloud account of the cloud provider you selected from Clouds.

You can perform the following actions from this section:

- Add Account
- Edit Account
- Delete Account
- Resync Account

Clouds > {Your Public Cloud} > Regions

The Regions section displays 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.

Clouds > {Your Public Cloud} > VPCs or VNets

The VPCs or VNets section displays your public cloud inventory.

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.

Note 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. 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 on **Actions** to access the following:
 - **Edit Configuration** (only available for Transit VPCs/VNets):
 - Enable or disable Quarantine Policy if in the NSX Enforced Mode.
 - Provide a Fallback Security Group that is required when the VPC/VNet is off-boarded from NSX Cloud when using the NSX Enforced Mode. See [Quarantine Policy Impact when Disabled](#).
 - 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.

Clouds > {Your Public Cloud} > Instances

The Instances section displays details of the 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.

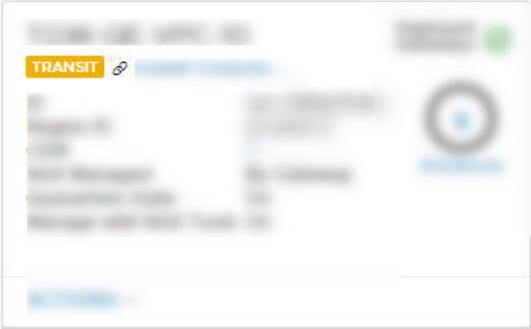
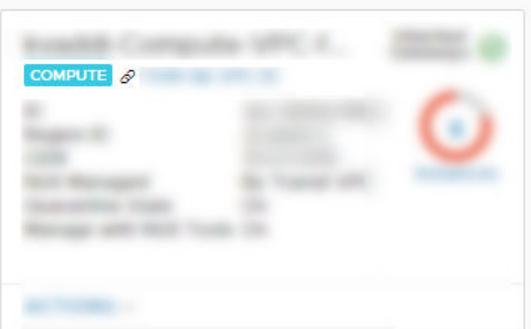
You can add instances to or remove instances from the CSM whitelist. See [Whitelisting VMs](#) for details.

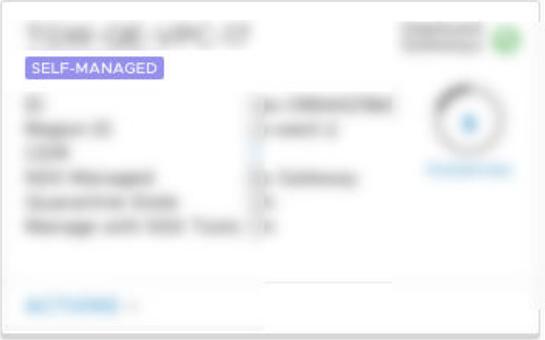
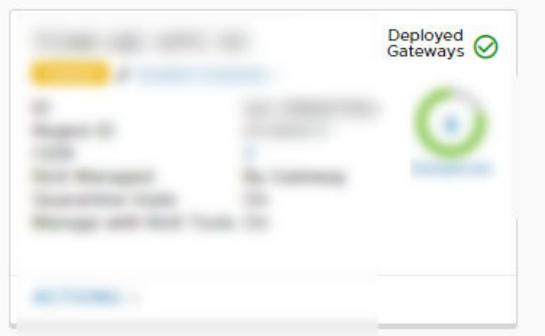
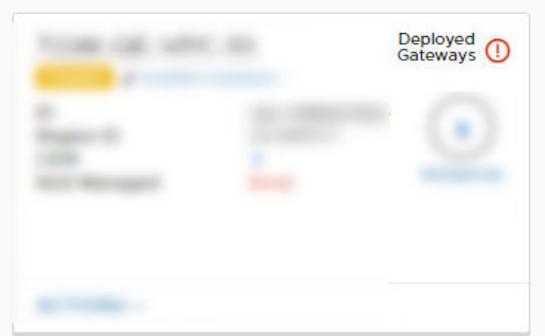
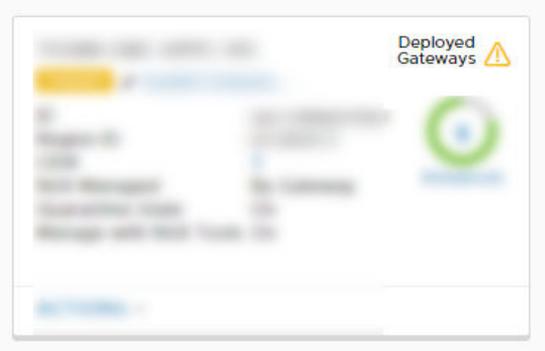
CSM Icons

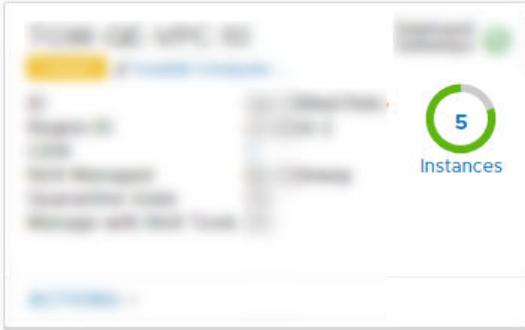
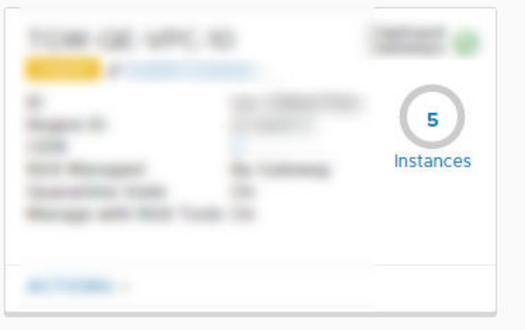
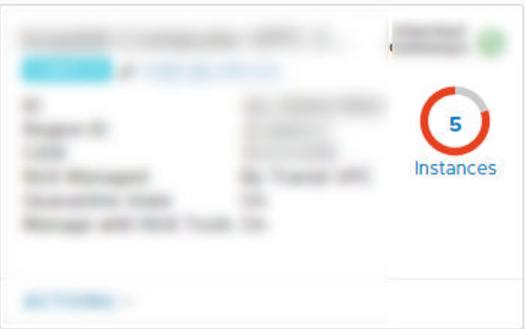
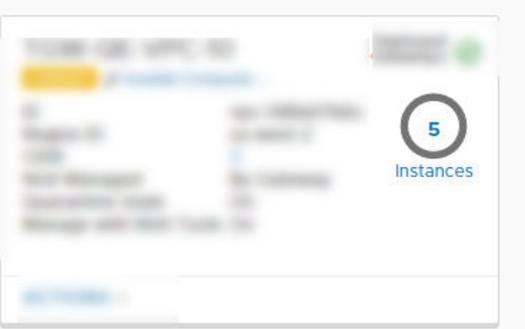
CSM displays the state and health of your public cloud constructs using descriptive icons.

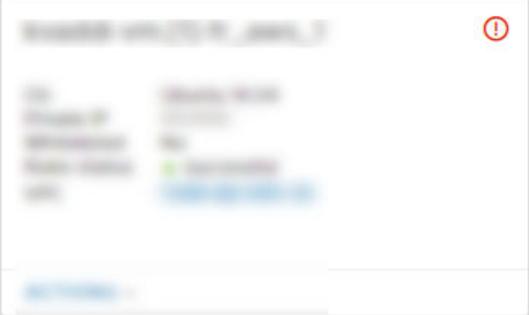
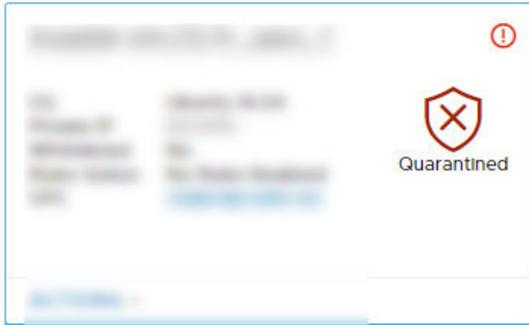
Note In the Native Cloud Enforced Mode: Quarantine Policy is always enabled and all VMs are always NSX-managed. Only the states where Quarantine Policy is enabled for NSX-managed VMs apply in this mode.

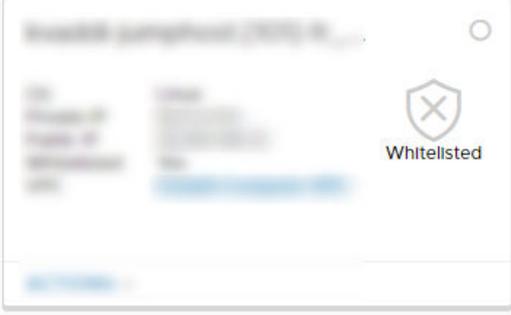
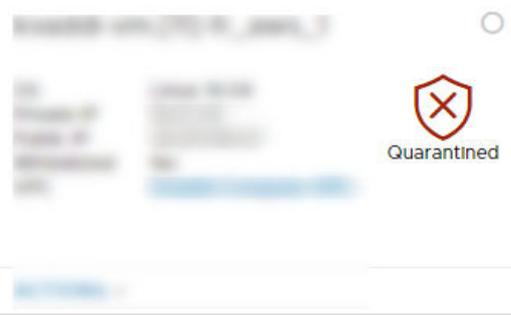
In the NSX Enforced Mode: Quarantine Policy can be disabled and it is possible to have unmanaged VMs in the VPC/VNet. All relevant states apply to this mode.

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
	<p>VPC/VNet showing NSX-managed VMs.</p>
	<p>VPC/VNet showing unmanaged VMs.</p>
	<p>VPC/VNet showing VMs with errors.</p>
	<p>VPC/VNet showing powered-off VMs.</p>
<p>Instances</p>	

CSM Section and Icon	Description
	<p>NSX-managed VMs with no errors.</p>
	<p>NSX-managed VMs with errors and Quarantine Policy disabled.</p>
	<p>NSX-managed VMs with errors and Quarantine Policy enabled.</p>

CSM Section and Icon	Description
	<p>Unmanaged VMs whitelisted.</p>
	<p>Unmanaged VMs quarantined.</p>

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.

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
NSX Manager Host Name	Enter the fully qualified domain name (FQDN) of the NSX Manager, if available. You may also enter the IP address of NSX Manager.
Admin Credentials	Enter an Enterprise Administrator username and password for NSX Manager.
Manager Thumbprint	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
Default	Use this radio button to indicate the default proxy server.
Profile Name	Provide a proxy server profile name. This is mandatory.
Proxy Server	Enter the proxy server's IP address. This is mandatory.
Port	Enter the proxy server's port. This is mandatory.
Authentication	Optional. If you want to set up additional authentication, select this check box and provide valid username and password.
Username	This is required if you select the Authentication checkbox.
Password	This is required if you select the Authentication checkbox.
Certificate	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.
No Proxy	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 [Backing Up and Restoring the NSX Manager](#) for details.

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.

System > Users

Users are managed using role-based access control (RBAC).

See [Managing User Accounts and Role-Based Access Control](#) 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 22-1. 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 How to Enable or Disable Quarantine Policy .	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-<NSX GUID></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 gw security groups are applied to the respective PCG interfaces in AWS and Microsoft Azure.</p> <ul style="list-style-type: none"> ■ gw-mgmt-sg ■ gw-uplink-sg ■ gw-vtep-sg <p>The vm 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"> ■ <code>vm-quarantine-sg</code> in Microsoft Azure and <code>default</code> in AWS. <p>Note 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**.

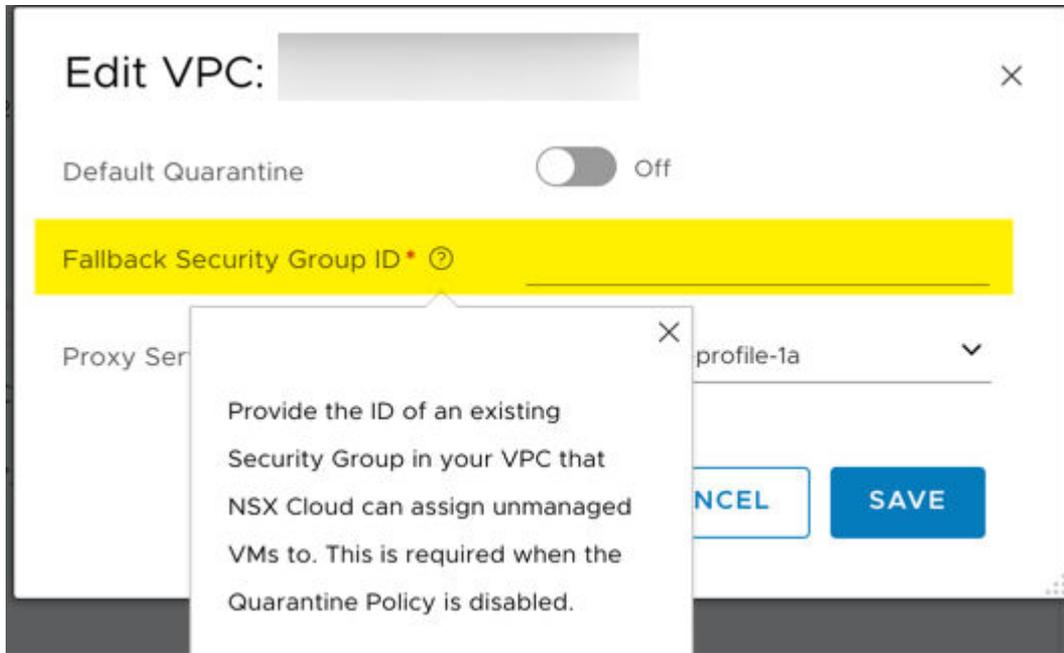


- If you are in the grid view, select the checkbox next to the VPC or VNet and click **ACTIONS > Edit Configuration**.
- ◆ If you are in the VPC or VNet's page, click the ACTIONS icon to go to **Edit Configurations**.



- 3 Turn **Default Quarantine** on or off to enable or disable it.
- 4 If you are disabling Quarantine Policy, you must provide a fallback security group.

Note The fallback security group must be an existing user-defined security group in your public cloud. You cannot use any of the NSX Cloud security groups as a fallback security group.



- All unmanaged VMs in this VPC or VNet will get the fallback security group assigned to them upon disabling Quarantine Policy.
- All managed VMs retain the security group assigned by NSX Cloud. The first time such VMs are untagged and become unmanaged after disabling Quarantine Policy, they also get the fallback security group assigned to them.

5 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: `vm-quarantine-sg` in Microsoft Azure and `default` in AWS are less restrictive. 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, whitelist them in CSM. See [Whitelisting VMs](#).

The following table shows how NSX Cloud manages the public cloud security groups of workload VMs when Quarantine Policy is disabled.

Table 22-2. 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 whitelisted?	VM's Public cloud security group when Quarantine Policy is disabled and explanation
Tagged	Not whitelisted	<ul style="list-style-type: none"> ■ If VM has no threats: <i>vm-underlay-sg</i> ■ If VM has potential threats (see note): <i>vm-quarantine-sg</i> in Microsoft Azure; <i>default</i> in AWS <p>Note The assignment of public cloud security groups is triggered within 90 seconds of applying the <i>nsx.network=default</i> 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 whitelisted	Retains existing public cloud security group because NSX Cloud doesn't take action on untagged VMs.
Tagged	Whitelisted	Retains existing public cloud security group because NSX Cloud doesn't take any action on whitelisted VMs.
Not Tagged		

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 with a Fallback Security Group configured for handling security group assignments for this VPC/VNet.

Table 22-3. 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 whitelisted?	VM's existing public cloud security group when Quarantine Policy is enabled	VM's public cloud security group after Quarantine Policy is disabled and a fallback security group provided
Not Tagged	Not Whitelisted	<i>vm-quarantine-sg</i> (Microsoft Azure) Or <i>default</i> (AWS)	This VM is assigned the fallback security group you provided when disabling the Quarantine Policy because it is untagged and not considered NSX-managed, therefore NSX Cloud reverts the security group it assigned this VM when you disable Quarantine Policy.
Tagged	Not Whitelisted	<i>vm-underlay-sg</i> Or <i>vm-quarantine-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.

Table 22-3. NSX Cloud assignment of public cloud security groups when Quarantine Policy is disabled from being enabled at first (continued)

Is VM tagged with <i>nsx.network=default</i> in the public cloud?	Is VM whitelisted?	VM's existing public cloud security group when Quarantine Policy is enabled	VM's public cloud security group after Quarantine Policy is disabled and a fallback security group provided
Tagged	Whitelisted	Any existing public cloud security group	Retains existing public cloud security group because NSX Cloud doesn't take any action on whitelisted VMs. Note If you have a whitelisted VM in any NSX Cloud-assigned security groups, you must move it to the designated fallback security group manually.
Not Tagged			

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, whitelist them in CSM. See [Whitelisting VMs](#).

Note Removing the VM from the whitelist causes the VM to revert to the NSX Cloud-assigned security group.

Table 22-4. 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 whitelisted?	VM's public cloud security group when Quarantine Policy is enabled and explanation
Tagged	Not whitelisted	<ul style="list-style-type: none"> ■ If VM has no threats: <i>vm-underlay-sg</i> ■ If VM has potential threats (see note): <i>vm-quarantine-sg</i> in Microsoft Azure; <i>default</i> in AWS <p>Note The assignment of public cloud security groups is triggered within 90 seconds of applying the <i>nsx.network=default</i> 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 whitelisted	<i>vm-quarantine-sg</i> in Microsoft Azure; <i>default</i> in AWS. Untagged VMs are considered unmanaged and therefore quarantined by NSX Cloud.
Tagged	Whitelisted	Retains existing public cloud security group because NSX Cloud doesn't take action on whitelisted VMs.
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 22-5. 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 whitelisted?	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 Whitelisted	Any existing public cloud security group	<i>vm-quarantine-sg</i> (Microsoft Azure) Or <i>default</i> (AWS)
Tagged	Not Whitelisted	<i>vm-underlay-sg</i> Or <i>vm-quarantine-sg</i> (Microsoft Azure) Or <i>default</i> (AWS)	Retains the NSX Cloud-assigned security group that is consistent for tagged VMs in the Quarantine enabled or disabled modes.
Tagged	Whitelisted	Any existing public cloud security group.	Retains existing public cloud security group because NSX Cloud doesn't take any action on whitelisted VMs.
Not Tagged			

Quarantine Policy in the Native Cloud Enforced Mode

Quarantine Policy is always enabled in the Native Cloud Enforced Mode.

Table 22-6. Assignment of Public Cloud Security Groups in the Native Cloud Enforced Mode

Is VM part of a valid NSX-T Security policy?	Is VM whitelisted?	VM's public cloud security group and explanation
Yes, VM is matched with a valid NSX-T Security Policy	Not whitelisted	NSX Cloud-created public cloud security group named like <code>nsx-{NSX-GUID}</code> 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 whitelisted	<p><code>vm-quarantine-sg</code> in Microsoft Azure or <code>default</code> in AWS because this is the threat detection behavior of NSX Cloud. In the Native Cloud Enforced Mode, the NSX Cloud-created security groups <code>vm-quarantine-sg</code> in Microsoft Azure or <code>default</code> in AWS mimic the default public cloud security policy.</p> <hr/> <p>Note In CSM the VM shows an Error state.</p>
Yes, VM has valid NSX-T Security policy	Whitelisted	Retains existing public cloud security group because NSX Cloud doesn't take any action on whitelisted VMs.
No, VM does not have a valid NSX-T Security policy		

Whitelisting VMs

Whitelisting is an option available from CSM for all workload VMs in your public cloud inventory.

Whitelisting works in both the VM-management modes: NSX Enforced Mode and the Native Cloud Enforced Mode.

Why to Whitelist VMs?

- 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, whitelist such a VM 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, whitelist such VMs so you can move them out of the quarantine state and use debugging tools as required.
 - Whitelist VMs in your public cloud inventory that you don't want NSX-T to manage, e.g. DNS Forwarder, Proxy server etc.

How to Whitelist VMs or Remove from Whitelist

Follow these instructions to add VMs to the whitelist 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 whitelist or remove from whitelist.
- 4 Click **Actions** and select either **Add to Whitelist** or **Remove from Whitelist**.
- 5 Go back to the Accounts tab, select the account tile and click **Actions > Resync Account**.

Results

Each VM added to the whitelist remains in the security group it was assigned before whitelisting. You can now apply any other security group to the VM as required. NSX Cloud ignores whitelisted VMs regardless of the status of Quarantine Policy.

If you remove a VM from whitelist in the Native Cloud Enforced Mode or remove an NSX-managed VM from whitelist 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.

- 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 22-7. Day-N Workflow for onboarding your workload VMs into NSX Cloud

Task	Instructions
<input type="checkbox"/> Tag workload VMs with the key-value <code>nsx.network=default</code> .	Follow instructions in your public cloud documentation for tagging workload VMs.
<input type="checkbox"/> Install NSX Tools on your Windows and Linux workload VMs.	See Install NSX Tools
<p>Note If Auto- Install NSX Tools is enabled in CSM for Microsoft Azure VNets, NSX Tools are automatically installed.</p>	
<input type="checkbox"/> (Optional) In CSM, remove from whitelist all VMs that you want to bring under NSX management.	See How to Whitelist VMs or Remove from Whitelist .
<p>Note Whitelisting is a manual step that is recommended in the day-0 workflow as soon as you add your public cloud inventory in CSM. You do not need to remove VMs from whitelist if you did not add any to the whitelist.</p>	

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 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 whitelist the VMs in CSM before installation and remove them from whitelist 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 Configuration**.



- 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.

Copy the download and installation instructions for NSX Tools from CSM and paste into User Data when launching a new workload VM.

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.

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 `vm-quarantine-sg` Network Security Group in Microsoft Azure.

- Whitelisted VMs 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 `vm-quarantine-sg` Network Security Group in Microsoft Azure.
- Unmanaged VMs and whitelisted VMs 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 22-8. 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 Default Connectivity Strategy for NSX-Managed Workload VMs in the NSX Enforced Mode .
 Group your workload VMs using public cloud tags or NSX-T Data Center tags and set up micro-segmentation.	See Set up Micro-segmentation for Workload VMs in the NSX Enforced Mode . See also: Group VMs using NSX-T Data Center and Public Cloud Tags

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>ccloud-stateful-ccloud-<VPC/VNet ID></code>	Properties
<code>ccloud-<VPC/VNet ID>-managed</code>	Allows access to the VMs within the same VPC/VNet.
<code>ccloud-<VPC/VNet ID>-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.

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](#).

Note 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
-

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.

- 2 Create an East-West distributed firewall policy and rule and apply to the group you created. See [Add a Distributed Firewall](#).

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 22-9. 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 Set up Micro-segmentation for Workload VMs in the Native Cloud Enforced Mode See also: Group VMs using NSX-T Data Center and Public Cloud Tags
<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 Whitelist 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 Current Limitations and Common Errors

Set up Micro-segmentation for Workload VMs in the Native Cloud Enforced Mode

Refer to this workflow for configuring Security Policy in NSX Manager for workload VMs in the Native Cloud Enforced Mode, that is by not installing NSX Tools on the workload VMs.

Prerequisites

You must 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 be added to the Group. VMs that do not match any grouping criteria are placed in the `default` Security Group in AWS and the `vm-quarantine-sg` Network Security Group in Microsoft Azure.

Note You cannot use the Groups auto-created by NSX Cloud.

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.

- 2 In NSX Manager create Distributed Firewall (DFW) rules with the 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.

- 3 In CSM, remove those VMs from whitelist that you want to bring under NSX management. See [How to Whitelist VMs or Remove from Whitelist](#) for instructions.

Note Whitelisting 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 whitelisted any VMs, you do not need to remove them from the whitelist.

- 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.

Note NSX Cloud synchronizes NSX Manager and public cloud groups and DFW rules every 30 seconds.

- 5 Resync 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 resync 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 local VPC/VNet rule enforcement is supported. You can create Groups in NSX Manager that span across VPC/VNets. However, the realization of such rules only works within the VPC/VNet. Cross-VPC/VNet DFW rules are not realized.

- Only TCP and UDP 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 blacklisted 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 22-10.

NSX-T Data Center Feature	Details	NSX Cloud Note
Segments or Logical Switches	See Chapter 4 Segments	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 Chapter 2 Tier-0 Gateways and Chapter 3 Tier-1 Gateway .	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 Configure IPFIX .	<ul style="list-style-type: none"> ■ IPFIX is supported in NSX Cloud only on UDP port 4739. ■ Switch and DFW IPFIX: 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.
Port Mirroring	See Monitor Port Mirroring Sessions .	<p>Port Mirroring is supported only in AWS in the current release.</p> <ul style="list-style-type: none"> ■ For NSX Cloud, configure Port Mirroring from Tools > Port Mirroring Session. ■ Only L3SPAN Port Mirroring is supported. ■ The collector must be in the same VPC as the source workload VM.
Gateway Firewall	See Configuring a Gateway Firewall .	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](#).

Table 22-11. Summary of Tag Types and Limitations

Tag type	Tag scope or predetermined prefix	Limitations	Enterprise Administrator Privileges	Auditor Privileges
System-defined	Complete system tags: <ul style="list-style-type: none"> ■ azure:subscription_id ■ azure:region ■ azure:vm_rg ■ azure:vnet_name ■ azure:vnet_rg ■ aws:vpn ■ aws:availability zone 	Scope (key): 20 characters Tag (value): 65 characters Maximum possible: 5	Read only	Read only
Discovered	Prefix for Microsoft Azure tags that are imported from your VNet: dis:azure: Prefix for AWS tags that are imported from your VPC: dis:aws:	Scope (key): 20 characters Tag (value): 65 characters Maximum allowed: unlimited Note The limits on characters excludes the prefix dis:<public cloud name> . Tags that exceed these limits are not reflected in NSX Manager. Tags with the prefix nsx 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"> ■ the scope (key) prefix dis:azure: or dis:aws: ■ the same scope (key) as system tags 	Scope (key): 30 characters Tag (value): 65 characters 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 22-12.

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=AbcdefghijklmnopqrstuvwxyzAb23690hgjgjuytreswqacvbcdefghijklmnopqrstuvwxyz	No (value exceeds 65 characters)	none
nsx.name=Tester	No (key has the prefix nsx)	none

How to use Tags in NSX Manager

- See [Manage Tags for a VM](#).
- 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 Public Cloud

NSX Cloud supports the use of third-party services in your public cloud for NSX-managed workload VMs.

To utilize service insertion for your public cloud workload VMs, you must host the service appliance in the public cloud, not in NSX-T Data Center. It is recommended to host the service appliance in a Transit VPC/VNet.

Before you can enable service insertion, you must have the PCG deployed in a Transit VPC or VNet.

Here is an overview of the one-time configurations to allow service insertion for your NSX-managed workload VMs.

Table 22-13. Overview of configurations required for service insertion for NSX-managed workload VMs in the public cloud

How often?	Task	Instructions
Once for the initial setup	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 Create the Service Definition and a Corresponding Virtual Endpoint
	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 Create the Service Definition and a Corresponding Virtual Endpoint
	Create an IPSec VPN tunnel between the service appliance and the PCG.	See Set up an IPSec VPN Session
	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 Configure BGP and Route Redistribution
	Note In the current release, service insertion is only supported for north-south traffic.	
As and when required	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.	See Set up Redirection Rules .

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](#)

Redirection rules can be adjusted according to your requirements.

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
Name	The name of this VPN service is used to set up the local endpoint and the IPSec VPN sessions. Make a note of it.
Service Type	Confirm that this value is set to IPSec.
Tier-0 Gateway	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 Route Based .
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. Note 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.

Option	Description
Tunnel Interface	<p>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.</p> <p>Note You configure BGP on this tunnel interface. See Configure BGP and Route Redistribution .</p>
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
- 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

What to do next

[Set up Redirection Rules](#)

Set up Redirection Rules

Redirection rules can be adjusted according to your requirements.

After the initial setup is completed, you can create and edit redirection rules as required for rerouting different types of traffic for your NSX-managed workload VMs through the service appliance.

Prerequisites

You must have all the Service Insertion setup completed before you can create redirection rules.

Procedure

- 1 Navigate to **Security > North South Firewall > Network Introspection (N-S)**
- 2 Click **Add Policy**.

Option	Description
Name:	Provide a descriptive name for the policy, for example North-south Service Insertion for Azure VMs .
Redirect To:	Select the name of the Virtual Endpoint you created for this service appliance when registering the service. See Create the Service Definition and a Corresponding Virtual Endpoint .
Apply To:	Select the PCG's tier-0 gateway.

- 3 Select the new policy and click **Add Rule**. Note the following values specific to service insertion:

Option	Description
Sources	Select a group of subnets whose traffic must be redirected, for example, a group of your NSX-managed workload VMs.
Destinations	Select a list of destination IP addresses or services, such as Google , that you want to route through the service appliance.
Applied To	Select the uplink port of the active and standby PCG.
Action	Select Redirect .

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 22-14.

Key	Value
<code>nsx.publicip</code>	public IP address from your public cloud, 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 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 5 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 it 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><VPC-ID>-AWS_VPN</code> or <code><VNet-ID>-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><VPC-ID>-PCG-preferred-LE</code> or <code><VNET-ID>-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><VPC--ID>-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
	Note Remote IP can be a private IP address if you are able to reach the private IP address, for example, using DirectConnect or ExpressRoute.
Tunnel Interface	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

Admin Status: Enabled

IKE Log Level: Info

Tags: 0

Session sync: Enabled

Step 2b.

LOCAL ENDPOINTS

Name	VPN Service	IP Address	Site Certificate	Se
<VPC-ID>-PCG1-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>-PCG1-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: [copy] [paste]

IP address of VPN Peer (pointing to 3.213.92.220)

- 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 ⓘ
ROUTE RE-DISTRIBUTION				
Route Re-distribution		Route Re-distribution Status		On

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

Step 4.

Frequently Asked Questions (FAQs)

This topic lists some frequently asked questions.

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.

VMware NSX® Intelligence™ provides a visualization of the security posture of your on-premises NSX-T Data Center environment. The visualization is based on the network traffic flows aggregated within a specific time period. NSX Intelligence also assists you with micro-segmentation planning by making recommendations that are based on analytics with enforcement on security policies.

Important You must have an Enterprise Administrator role to have permission to install, configure, and use NSX Intelligence.

Before you can begin using the NSX Intelligence features, you must first install and configure the NSX Intelligence appliance. See "Installing and Configuring the NSX Intelligence Appliance" in the *NSX-T Data Center Installation Guide*.

This chapter includes the following topics:

- [Getting Started with NSX Intelligence](#)
- [Understanding NSX Intelligence Views and Flows](#)
- [Working with NSX Intelligence Recommendations](#)
- [Backing Up and Restoring NSX Intelligence](#)
- [Troubleshooting NSX Intelligence Issues](#)

Getting Started with NSX Intelligence

To get started using the NSX Intelligence features, familiarize yourself with the NSX Intelligence graphical user interface.

After the NSX Intelligence appliance is installed and configured, the NSX Intelligence features are enabled in the **Plan & Troubleshoot** tab of the NSX Manager UI. In the **Discover & Plan** section, you use **Discover & Take Action** to visualize your NSX-T data center entities and **Recommendations** to obtain recommendations for micro-segmentation planning.

Tour of the NSX Intelligence Home Page

You access the NSX Intelligence home page by clicking **Plan & Troubleshoot > Discover & Take Action** in the NSX Manager user interface.

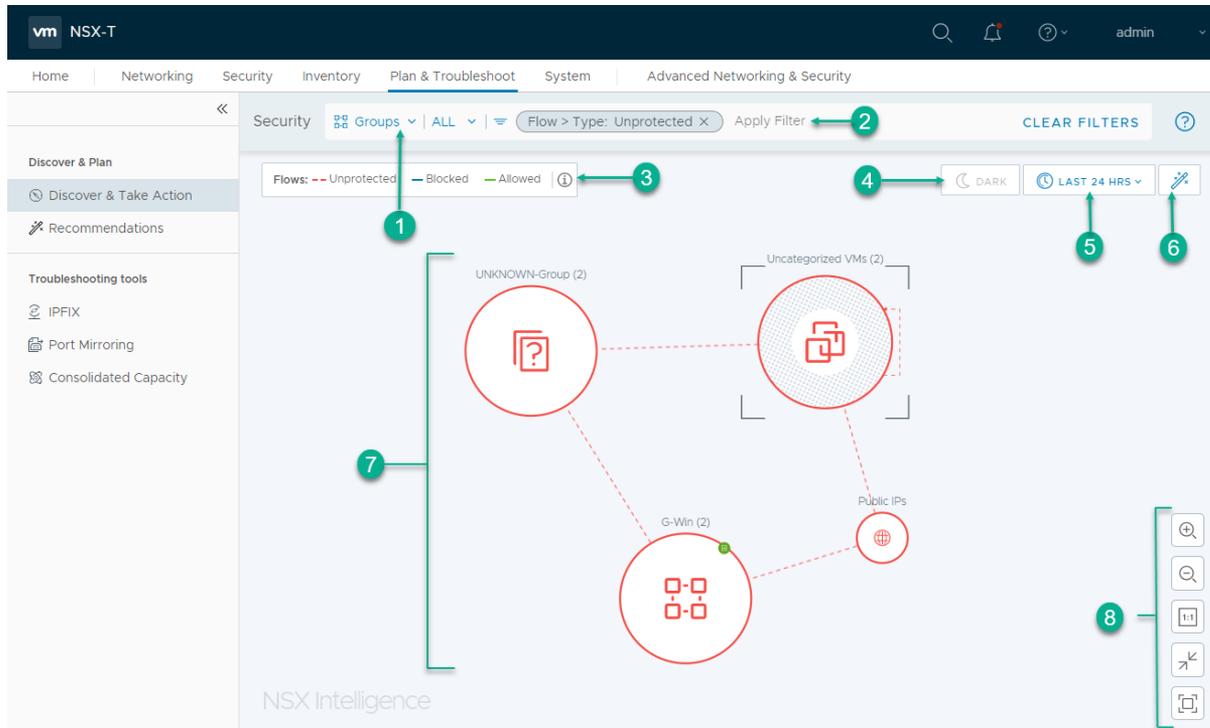
After you install and configure NSX Intelligence for the first time, when you click **Discover & Take Action** you might see the message, *No data found*. You might need to modify your filters above. The message appears because NSX Intelligence has yet to receive network traffic data to create a visualization. After some network traffic data has been received from NSX Manager, NSX Intelligence can begin to render some visualization.

By default, when you click **Discover & Take Action** you see the visualization of the security status of all the groups in your on-premises NSX-T Data Center that had unprotected traffic flows between their VM members in the last 24 hours. Unprotected network traffic flows are flows between VMs that do not have any micro-segmentation implemented. If there are no groups defined yet, there are no groups displayed. If there are VMs, but they do not belong to any group, you see the following icon for the Uncategorized VMs group.



If you already have defined groups and captured traffic data, you might see a visualization similar to the following screenshot. The table that follows describe the numbered sections in the screenshot.

Note NSX Intelligence categorizes an IP address belonging to one of the following CIDR notations as a private IP address: 192.168.0.0/16, 172.16.0.0/12, and 10.0.0.0/8. Any IP address that does not belong to any of these CIDR notations is classified as a public IP address. If your VM's IP address does not fall into one of these CIDR notations, consider adding your CIDR notation using the `PATCH /api/v1/intelligence/host-config` API in the *NSX-T Data Center API Guide*.



Section	Description
1	<p>The Security view selection area is where you select the type of security visualization to display. There are two types of Security views available: Groups and VMs. When you click Discover & Take Action, the default Security view displayed is the Groups view of the group objects in your NSX-T Data Center that had unprotected flow traffic within the last 24 hours.</p> <ul style="list-style-type: none"> ■ To select the VMs view, click the down arrow next to Groups and select VMs. ■ To select the specific groups or VMs to include in the view, click the down arrow next to ALL, and select from the list. ■ To clear your selection filters, click CLEAR FILTERS on the top right-side of the screen. When you click CLEAR FILTERS while in the VMs view, the selection filters are cleared and you are placed in the Groups view. <p>See Working with the Groups View and Working with the VMs View for more information on how to work with the two view types.</p>
2	<p>With the Apply Filter, you can refine the criteria used for the visualization. From the drop list, you can select the criteria to use for the visualization. You can select VM members, tags, flow types, source IP, destination IP, rule ID, or name. You can define multiple filters to apply by clicking Apply Filter again.</p>
3	<p>With this Flows section, you can select which traffic flow type to include in the visualization during the selected time period. The colors used in the visualization for the flow types are also shown in this section.</p> <ul style="list-style-type: none"> ■ Red-hued dashed line for Unprotected flows ■ Blue-hued solid line for Blocked flows ■ Green-hued solid line for the Allowed flows <p>By default, the Unprotected traffic flow type is selected for the current NSX Intelligence visualization. See Working with Traffic Flows for more information.</p>

Section	Description
4	<p>The display mode section defines what theme to use for the visualization. Light theme is the default mode used.</p> <ul style="list-style-type: none"> ■ To use the dark theme mode, click the DARK icon. You can use the Dark theme only when you are viewing the visualization in full screen mode. ■ To go into full screen mode, click  in the viewing control section.
5	<p>In this section, you select the time period to use to determine which network flow data is used to generate the desired visualization and recommendation. Your selection determines the historical data that is used in the Groups or VMs view. The time period is relative to the current time and some time period in the past.</p> <p>The last 24 hours is the default time range used. To change the selected time period, click the currently selected time period and select Last 1 hr, Last 12 hrs, Last 24 hrs, Last 1 week, or Last 1 month.</p>
6	<p>When you click this Recommendation wand  icon, the Recommendations dialog box displays the inventory summary for the current view. If you are in the VMs view, you can generate an NSX Intelligence recommendation by clicking Start New Recommendation. See Working with NSX Intelligence Recommendations.</p>
7	<p>This section is the visualization of the security status of the Groups or VMs in your on-premises NSX-T Data Center. It also includes the visualization of the network traffic flows that occurred during the selected time period. In this section, you can point to a specific node or flow arrow to obtain details about that specific entity.</p> <p>See Getting Familiar with NSX Intelligence Graphic Elements and Understanding NSX Intelligence Views and Flows for more information.</p>
8	<p>This section includes the viewing controls to zoom in, zoom out, apply 1:1 aspect ratio, resize-to-fit the view, and go into or out of full-screen viewing mode. You can also use keyboard hotkeys to manage your viewing controls. To display the Keyboard Shortcuts Help window, press Shift+.</p> <p>To navigate to a previously viewed visualization, use your Web browser's back button. When you are in full-screen mode, click Back (at the top left of the screen) to perform the same back button navigation.</p>

Getting Familiar with NSX Intelligence Graphic Elements

The NSX Intelligence user interface provides several graphic elements to help with the visualization of the data center entities, traffic flows, and certain activities in your NSX-T Data Center environment.

The following table lists a glossary of NSX-T Data Center graphic elements that you might see in a NSX Intelligence visualization.

Graphic Element	Description
	This icon represents a group, which is a collection of VMs where security policies, including East-West firewall rules, can be applied. See Working with the Groups View .
	This icon represents a virtual machine (VM) that is part of your NSX-T Data Center. A VM can belong to more than one group. See Working with the VMs View .

Graphic Element	Description
	<p>This icon represents the public IPs in the Internet. If at least one VM in your NSX-T Data Center environment communicated with a public IP during the selected time period, that traffic flow is included in the current visualization.</p>
	<p>An IP address, such as a unicast, broadcast, or multicast IP, that participated in the network traffic activities during the selected time period.</p>
	<p>This icon is used for the group of VMs that do not belong to a group.</p>
	<p>An arrow represents a network traffic flow that occurred between two VMs during a selected time period. There are three different types of arrows: a dashed red-hued arrow for an Unprotected flow, a solid blue-hued arrow for a Blocked flow, and solid green-hued arrow for an Allowed flow. See Working with Traffic Flows.</p>
	<p>A node that has been selected as the current node in focus is surrounded with a dashed circle. It is the pinned node during the selection mode and the current view being displayed.</p>
	<p>This icon appears on a group node's border if the group was added in the NSX-T Data Center inventory during the selected time period. If NSX-T Data Center discovered a VM during the selected time period, the icon appears on that VM node's border.</p>
	<p>This icon appears on the group node's border if the group was deleted during the selected time period and the VM members were not deleted. On a VM node's border, this icon indicates that the VM was deleted during the selected time period. Although, a VM or group has been deleted, it still appears in the current visualization to give a historical view that the VM or group was removed during the selected time period.</p>
	<p>This icon appears whenever we see group and VMs together. For example, in a deep dive groups view or related VMs of a group.</p> <p>The icon appears on a VM node's border in the following cases.</p> <ul style="list-style-type: none"> ■ if the VM was moved out of the currently viewed group during the selected time period ■ if, at some point during the selected time period, the VM was part of the group you are currently viewing, but it is no longer a member of that same group

Understanding NSX Intelligence Views and Flows

The NSX Intelligence visualization is composed of the groups or VMs and the network flows that occurred with those groups or VMs during the selected time period.

Important The visualization shown for a specific time period represents all the network flows and activities, such as addition, deletion, or movement of VMs and Groups, that occurred in your NSX-T data center during that time period. It is possible that a VM appears more than once in the visualization. For example, if a VM was attached to an ESXi host that was originally unmanaged and the host becomes managed by a VMware vCenter Server™ during the selected time period, the VM appears twice in the VMs view. Similarly, if an ESXi host is disconnected from vCenter Server and added back during the same selected time period, the VMs attached to the host appear as both deleted and new during the selected time period. In Groups view, if a VM was in the Uncategorized group and added in a Group during the same selected time period, the VM appears in both the Uncategorized group and in its new Group.

NSX Intelligence supports Groups with VM member types only. If you have Groups with any other types of members, the Groups view might show correlated flows between the Groups with VM members types instead of actual Groups in the security rule.

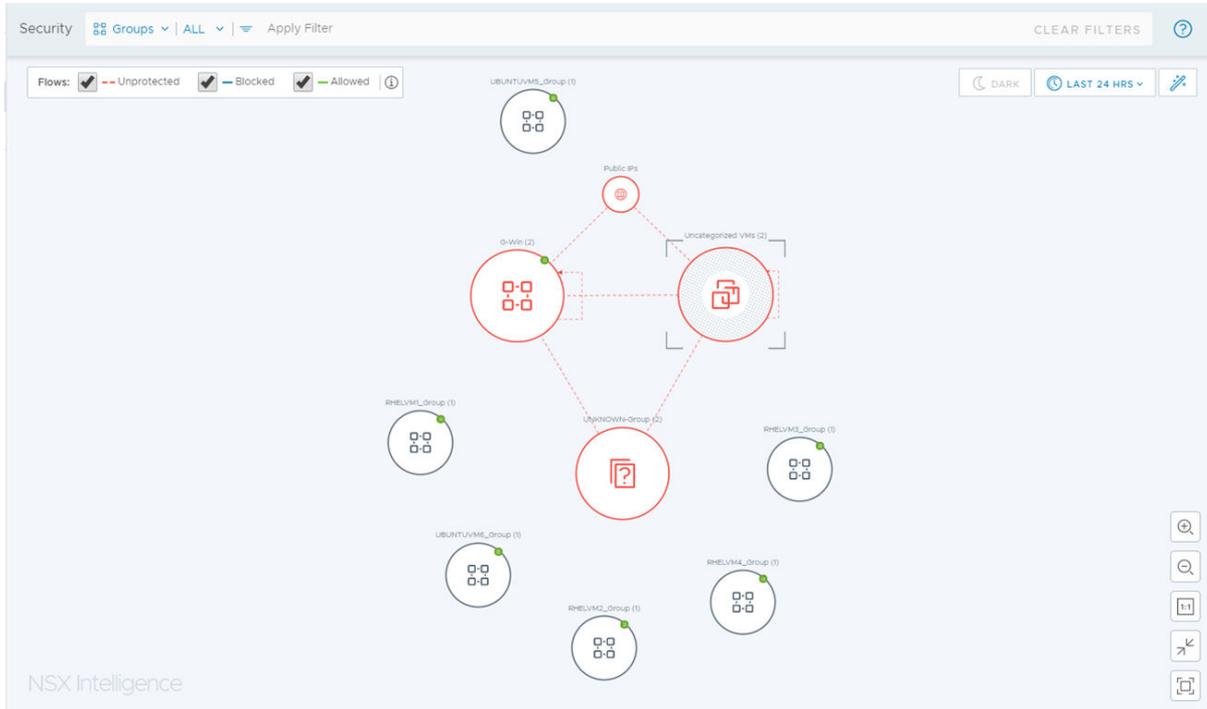
Use the information in this section to learn more about working with the Groups view, VMs view, and the different traffic flows.

Working with the Groups View

The default view that is shown in the NSX Intelligence home page is the Groups view. This Groups view is filtered to display all the groups that had unprotected traffic flow in the last 24 hours.

Nodes and Arrows in a Groups View

A node in a Groups view represents NSX objects, such as VMs, IP sets, and so on, in your NSX-T Data Center environment. The following screenshot is a sample of a Groups view.



The following table lists the types of Group nodes you might see in the Groups view.

Type of Group Node	Icon	Description
Regular Group		A Regular Group node in NSX Intelligence represents any collection of NSX objects in your NSX-T Data Center environment. For this release, those NSX objects are VMs only and so NSX Intelligence supports Regular Groups with only VM member types. An NSX object can belong to more than one Group and so a VM can appear in more than one Group node.
Uncategorized Group		An Uncategorized Group node represents a collection of VMs that do not belong to any Group.
Unknown Group		An Unknown Group node represents a set of miscellaneous objects that were not found in your NSX-T Data Center inventory. However, these objects are communicating to one or more NSX objects in your NSX-T Data Center environment.
Public IPs Group		A Public IPs Group node represents a collection of public IP addresses (IPv4 or IPv6) that are communicating to NSX objects in your NSX-T Data Center.

The size of a node in the Groups view is based on the number of NSX objects, such as VMs, that belong to that group. The bigger the group's node, the more VMs belong to that group, for example. The name of the group and the total number of member VMs it has are displayed above the node.

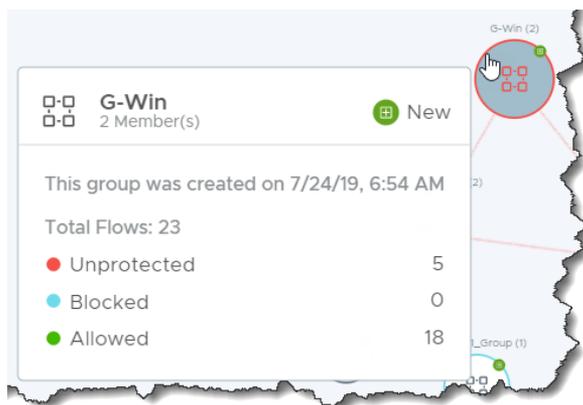
The arrows between the group nodes represent the traffic flows that have occurred between the VMs in those connected group nodes, during the selected time period. A self-referencing arrow on a group node indicates that at least one VM was communicating with another VM within that same group. See [Working with Traffic Flows](#) for more information.

A node with a red-hued border indicates that at least one unprotected flow occurred with a VM in the group, regardless of how many blocked or allowed flows were detected during the selected time period. A blue-hued border on a node means that no unprotected traffic flows were detected, but at least one blocked flow was detected, regardless of how many allowed flows were detected during the selected time period. A node with a green-hued border indicates that there were no unprotected or blocked flows detected during the selected time period, and at least one allowed flow was detected. A node with a gray-hued border means that there were no traffic flows detected for the VMs belonging to that group during the selected time period.

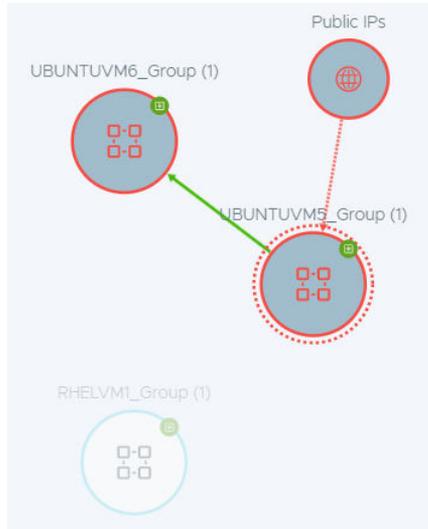
If you are not seeing the Groups view, click the down arrow next to **VMs** in the Security view selection area and select **Groups**. In the selection drop list displayed, you can select **All Groups** or specific groups from the list, and then click **Apply**. Use the **Search** text box to filter the selection list. If you click away from the selection drop list without making any selection or if you select **All Groups** in the drop list, the **All Groups** option is applied to the Groups view.

Node Selection in Groups View

When you point to a group's node, information about that group is displayed, as shown in the following example for the group G-Win. The number and types of flows detected during the selected time period are also listed. If the group was added during the selected time period, the New badge icon and the details of when the group was created are also displayed.



When you click a group's node, a dashed circle marks the selection as a pinned group node. The other groups that are connected to the selected group node are also made more prominent in the view. All other nodes become dimmed. For example, in the following screenshot, the node UBUNTUVM5_Group is selected and other groups that shared a traffic flow with UBUNTUVM5_Group during the selected time period are also highlighted. All the other groups that did not communicate with UBUNTUVM5_Group are faded out in the view.

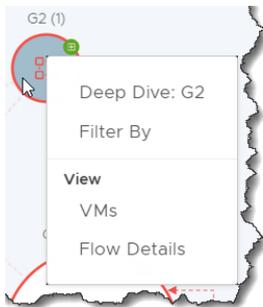


To clear the pinned selection, click in any empty area of the Groups view.

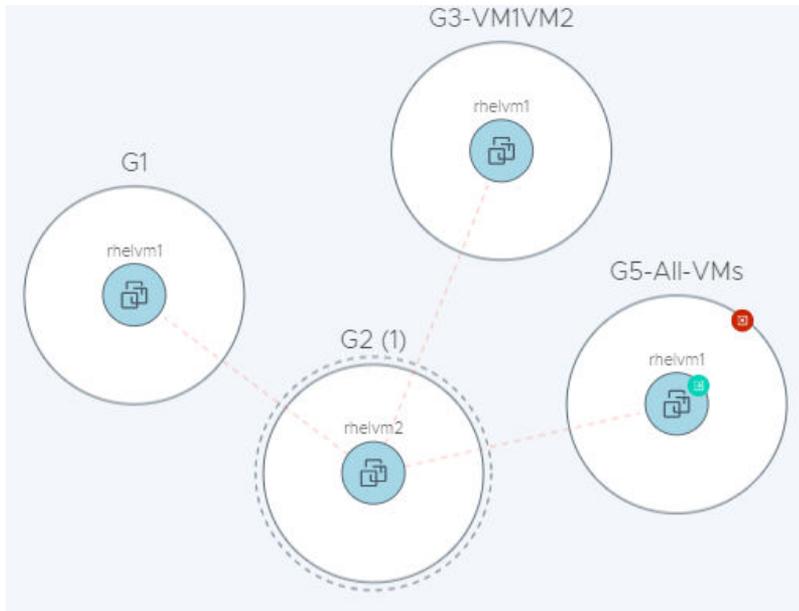
If you zoom out of the Groups view and the details on the nodes are no longer visible, point to any visible part of a node and its details are displayed.

Available Actions in the Groups View

A contextual menu of available actions is displayed when you right-click a group's node, as illustrated in the following image.



- Selecting **Deep Dive:Group_Name** surrounds the selected group's node with a dashed circle to mark it as the pinned group node or the current group in focus. The VMs that belong to the group are shown inside the group's node. All the groups that had traffic flows with the VMs in the pinned group during the selected time period are also placed in the Groups view. In the following example, group G2 is the pinned group and the other groups are in the view because their VM members had traffic flows with rhelvm2 in group G2 during the selected time period.



- When you select **Filter By**, the current group is added to the visualization filter that is used for the current Groups view.
- Selecting **VMs** displays a table of all the VMs that belonged to the current group during the selected time period. From that View VMs table, you can see the details about the VMs that belong to the selected group and the other groups to which each VM also belongs. To add the VM to the current visualization filter, click the filter icon.
- When you select **Flow Details**, the Flow Details table for the currently selected group is displayed, as shown in the following screenshot. It shows the details about the flows that have occurred and are currently active with the VMs that belong to the current group during the selected time period. The details include the flow type, the flow's source and destination groups, start and end time of the flow, and the services that were used. You can click some of the details to obtain more information. See [Working with Traffic Flows](#) for more information.

Flow Details | Last 24 hrs

Showing flow details for G-Win

Completed Flows | Active Flows

Search

Source	Source Group	Destination	Destination Group	Services	End Time	Latest Flow
Public IPs	UNKNOWN	vdnet-ad-Win2012R2-1-PA...	G-Win	2	7/27/19, 10:56 AM	Unprotected
vdnet-ad-Win2012R2-1-PA...	G-Win	Public IPs	UNKNOWN	2	7/27/19, 10:56 AM	Unprotected
Windows2012-02	G-Win	Public IPs	UNKNOWN	2	7/27/19, 10:56 AM	Unprotected
Public IPs	UNKNOWN	Windows2012-02	G-Win	2	7/27/19, 10:50 AM	Unprotected
Windows2012-02	G-Win	Public IPs	UNKNOWN	DHCPv6_Server	7/27/19, 10:56 AM	Allowed

Refresh | 1 - 22 of 22 Flow(s)

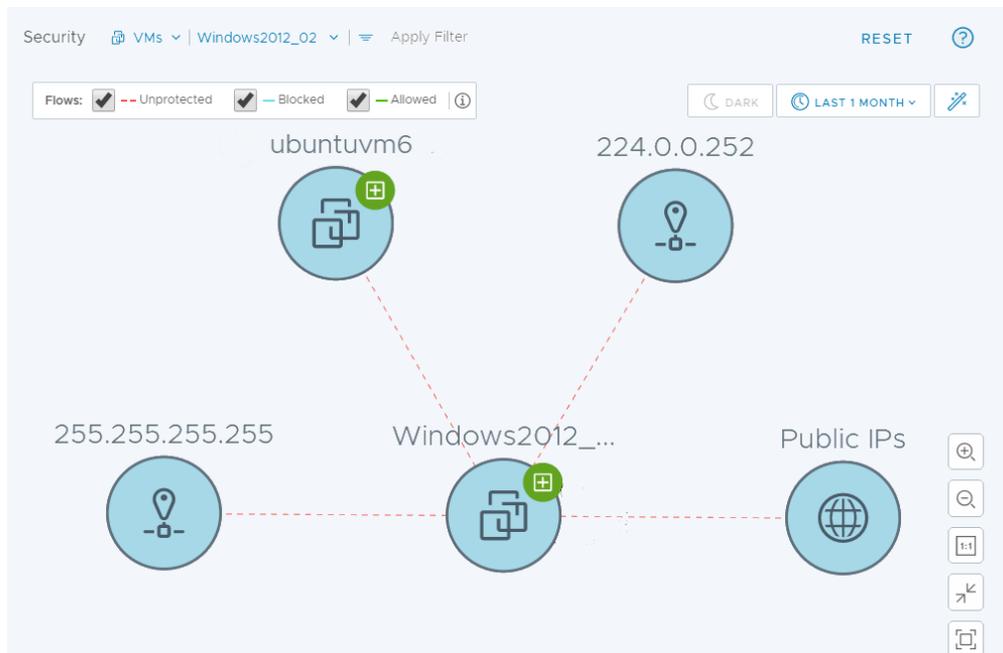
CLOSE

Working with the VMs View

A node in the VMs view represents a virtual machine (VM) in your on-premises NSX-T Data Center environment.

Nodes and Arrows in the VMs View

When you are in the VMs view, the group boundaries are not visible. Any node that is communicating with one of the VMs in your NSX-T Data Center environment, but was not identified as part of the NSX-T Data Center inventory, are also represented in the VMs view. The following illustrates a simple VMs view.



The following table lists the types of VM nodes you might see in the Views view.

Type of VM Node	Icon	Description
Regular VM		A Regular VM node represents a virtual machine (VM) that is part of your NSX-T Data Center environment. A VM can belong to more than one group.
Public IP		A Public IP node represents a public IP address, either an IPv4 or IPv6, that is communicating to or from your NSX-T Data Center environment.
IP		An IP node represents an IP address that participated in the network traffic activities during the selected time period. An IP address can be a unicast, broadcast, or multicast IP.

If you are not seeing the VMs view, click the down arrow next to **Groups** in the Security view selection area and select **VMs**. In the selection drop list displayed, you can select **All VMs** or specific VMs from the list, and then click **Apply**. Use the **Search** text box to filter the selection list. If you click away from the drop list without making any selection or if you select **All VMs** in the drop list, the **All VMs** option is applied to the VMs view.

The arrows between the VM nodes represent the traffic flows that have occurred between the VMs during the selected time period. See [Working with Traffic Flows](#) for more information.

Node Selection in VMs View

When you point to a VMs node, information about the node is displayed, as shown in the following example. The number and types of flows to the VM that were detected during the selected time period are also listed. If the group was added during the selected time period, the New badge icon and the details of when the VM was added are also displayed.

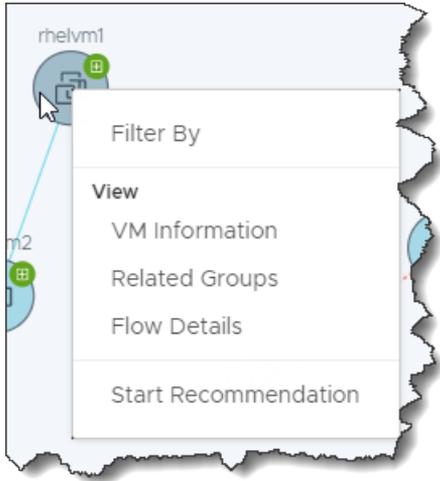


When you click a VM's node, a dashed circle marks the selection as a pinned VM node. Other VM nodes that had traffic flows with that pinned VM node are also made more prominent in the VMs view. All other nodes become dimmed to make them less visible. To clear the pinned selection, click in any empty area of the VMs view.

When you zoom out of the VMs view and the details in the VM nodes are no longer visible, point to any visible part of the VM node and its details are displayed.

Available Actions in the VMs View

A contextual menu of available actions is displayed when you right-click a VM's node, as illustrated in the following image.



Selection	Description
Filter By	The VM is added to the visualization filter that is used for the current VMs view.
VM Information	The details of the VM during the selected time period are displayed.
Related Groups	The Groups table with information about groups to which the VM belonged during the selected time period.
Flow Details	Shows the details about the flows that have occurred and are currently active with the VM during the selected time period. The details include the following. <ul style="list-style-type: none"> ■ flow type ■ flow's source and destination groups ■ start and end time of the flow ■ services that were used You can click some of the details to obtain more information. See Working with Traffic Flows for more information.
Start Recommendation	Displays the Start New Recommendations wizard. See Working with NSX Intelligence Recommendations for more details.

Working with Traffic Flows

The arrows between the Group or VM nodes represent the network traffic flows that have occurred between the VMs during the selected time period.

Network traffic flows are based on the L3 distributed firewall (DFW) rules in place and the traffic flows that occurred during the selected time period. All network traffic flows that matched a stateful L3 DFW rule using IPv4 or IPv6 with TCP, UDP, GRE, ESP, and SCTP protocols are included in the visualization and flow details. TCP and UDP flows have IP and port level details and others have IP level details only.

The traffic flows are categorized into the following types.

Flow Type	Graphic	Description
Unprotected		A dashed red-hued arrow indicates that the system detected that the traffic flow encountered a rule (Source: Any Destination: Any Action: Allow or Reject or Drop) and that more granular security policies are needed. This rule can be your default rule, or it can reside anywhere in the East-West distributed firewall.
Blocked		A solid blue-hued arrow indicates that the system detected that the traffic flow encountered a 'Reject' or 'Drop' rule that is more granular than the one mentioned in the 'Unprotected' flow definition.
Allowed		A solid green-hued arrow indicates that the system detected that the traffic flow encountered an 'Allow' rule that is more granular than the one mentioned in the 'Unprotected' flow definition.

To focus only on objects with certain types of traffic flows, use the Security view selection area to select which view type, and use the 'Flow Type' filter attribute to narrow down the selection.

If you deselect a flow type, the flow lines for that flow type are hidden from the displayed graph. Unless filters are in effect that exclude certain objects, all group or VM objects remain displayed regardless of the traffic flow types that have occurred with those objects during the selected time period. For example, if you deselect the 'Allowed' flow type, all the "Allowed" flow lines are hidden in the graph. However, all objects are still displayed, even those objects that only had 'Allowed' traffic flows during the selected time period.

A flow arrow's direction indicates the source and destination of the detected traffic flow. When in Groups view, a self-referencing arrow on a group node indicates that at least one VM was communicating with another VM within that same group. In a VMs view, a self-referencing arrow indicates that an NSX object in the VM communicated with another NSX object in the same VM.

When you point to a flow arrow, information about the flows involving the group or VM is displayed, as shown in the following example for Group G2.



When you click a flow arrow, the Flow Details dialog box is displayed. It shows the details about the completed and active flows that occurred during the selected time period. To get more detailed information about the flow's source, destination, type of service, and the type of flow, click the links in the table to see more details.

Working with NSX Intelligence Recommendations

NSX Intelligence can provide micro-segmentation recommendations that are based on the patterns of traffic flows that have occurred between the VMs in your NSX-T Data Center environment during a selected time period.

Understanding NSX Intelligence Recommendations

The recommendations that NSX Intelligence generates include security policies, policy security groups, and services for applications.

The recommendations are based on the network traffic flow patterns between VM workloads on ESXi hosts managed by a vCenter Server. They can assist you with enforcing a more dynamic security policy by correlating traffic patterns of communication that have occurred within your NSX-T Data Center environment.

The security policy recommendations are of the East-West Distributed Firewall Security Policies of Application category. The security group recommendations consist of a list of VMs that are seen in the network traffic flows that were analyzed for the time period and the VM boundary you had specified. The service recommendations are service objects that were used in certain ports by applications in the VMs you had specified, but the services are not yet defined in the NSX-T Data Center inventory.

There are multiple ways to request the recommendation, but the most straightforward one is by using **Plan & Troubleshoot > Recommendations** tab and clicking **Start New Recommendation**. You provide the virtual machines (VMs) that comprise the application boundaries and the time range in which the network traffic flows are to be analyzed for those specific VMs. Once the recommendation analysis is complete, you can view the details of the recommendation and, if necessary, modify the recommendation before publishing it. See [Generate a New NSX Intelligence Recommendation](#) for more information.

Generate a New NSX Intelligence Recommendation

The NSX Intelligence Recommendations feature can provide you with recommendations to help you micro-segment your applications.

Generating an NSX Intelligence recommendation involves recommendations of security policies, policy security groups, and services for the application. The recommendations are made based on the traffic pattern of communication between VMs in your NSX-T Data Center. There are multiple ways to generate a recommendation using the NSX Intelligence UI. The following procedure describes the three available methods to use.

Prerequisites

Install NSX Intelligence. See "Installing and Configuring NSX Intelligence" in the *NSX-T Data Center Installation Guide*.

Procedure

- 1 From your browser, log in with enterprise administrator privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Initiate the generation of a new recommendation.

Use the following table to decide which of the three available methods to use.

Method	Steps
Select Plan & Troubleshoot > Recommendations .	Click Start New Recommendation .
In the VMs view, select a VM and right-click.	From the contextual menu, select Start New Recommendations .
Select Plan & Troubleshoot > Discover & Take Action .	<ol style="list-style-type: none"> 1 In the Security Posture filter, click the down arrow and select VMs. 2 Select the VMs that comprise the application boundary and click Apply. 3 Click the Recommendations wand icon . 4 On the Recommendations dialog box, click Start New Recommendation.

- 3 In the Start New Recommendations wizard, optionally change the default value for the **Recommendation Name**.

- 4 Define or modify the VMs that are to be used as the boundary for the security policy recommendation.
 - a Click **Select VMs** or the number of **VMs Selected**.
 - b In the Select VMs dialog box, select the VMs that you want to use as the boundary for the analysis and deselect the ones you do not want included.

You can select up to 100 VMs to use for the recommendation boundary. You can also begin entering the name in the selection bar to filter the VMs to select.

- c Click **Save**.

The number of VMs selected is indicated on the Discover New Recommendation dialog box.

- 5 Expand **More Options** to change the default values for **Description** and **Time Range** that are used for the recommendation analysis. The default **Time Range** value is Last 1 Month, which means the network traffic flows that occurred in the last one month between the selected VMs are used during the recommendation analysis.

- 6 Click **Start Discovery**.

Recommendations are processed serially. On average, it can take anywhere from 3 to 4 minutes to finish each recommendation, depending on whether there are other recommendations that are pending to be processed. If there are many traffic flows between VMs that must be analyzed, the generation of a recommendation can take anywhere between 10–15 minutes. The status can be tracked from the **Recommendations** tab. The status progresses from `Waiting`, to `Analyzing`, and finally to `Ready to Publish`. The following screenshot shows the three different statuses of the generated recommendations.

	Name	Status	VMs	Create Time	Last Modified
⋮ >	REC 20190719 17:16:33	Waiting	4	7/19/19, 5:16 PM	7/19/19, 5:16 PM
⋮ >	REC 20190719 17:15:43	Analyzing	5	7/19/19, 5:15 PM	7/19/19, 5:15 PM
⋮ >	REC 20190719 15:59:02	Ready To Publish	3	7/19/19, 3:59 PM	7/19/19, 3:59 PM

After a recommendation is published successfully, the status is changed to `Published`.

What to do next

Review the generated recommendation and decide whether to publish it. See [Review and Publish a Generated Recommendation](#).

Review and Publish a Generated Recommendation

After the generated NSX Intelligence recommendation reaches the `Ready to Publish` status, you can review the recommendation, modify it if necessary, and decide whether to publish it.

Prerequisites

Generate a new recommendation. See [Generate a New NSX Intelligence Recommendation](#).

Procedure

- 1 From your browser, log in with enterprise administrator privileges to an NSX Manager at <https://<nsx-manager-ip-address>>.
- 2 Click **Plan & Troubleshoot > Recommendations**.
- 3 To help narrow down the list of recommendations being displayed, click **Filter by Name, Path or more** on the top right of the screen, and specify the filter criteria to be used.
- 4 If you decide not to use the recommendation, click the three-dot menu icon and select **Delete**.
- 5 To view the summary for a recommendation, click the arrowhead next to the recommendation's name to expand the row.

You see the number of rules generated and the number of groups affected.

- 6 Review and manage the details of the recommendation.
 - a Click the recommendation's name.

The **Recommendations** wizard is displayed, similar to the following image.

Name	Sources	Destinations	Services	Profiles	Applied To	Action
Policy-1 (REC 20190719 15:59:02)	(6)					
Rule-1 (REC 20190719 15:59:02)	Group-1 (REC 20190719 15:59:02)	Any	Win - RPC, DCOM, EP...	None	DFW	Allow <input checked="" type="checkbox"/>
Rule-2 (REC 20190719 15:59:02)	Group-1 (REC 20190719 15:59:02)	Any	NBDG-Broadcast-V1	None	DFW	Allow <input checked="" type="checkbox"/>
Rule-3 (REC 20190719 15:59:02)	Group-1 (REC 20190719 15:59:02)	Any	DHCP-Server	None	DFW	Allow <input checked="" type="checkbox"/>
Rule-4 (REC 20190719 15:59:02)	Group-1 (REC 20190719 15:59:02)	Any	DHCPv6 Server	None	DFW	Allow <input checked="" type="checkbox"/>
Rule-5 (REC 20190719 15:59:02)	Group-1 (REC 20190719 15:59:02)	Any	NBNS-Broadcast-V1	None	DFW	Allow <input checked="" type="checkbox"/>
Rule-6 (REC 20190719 15:59:02)	Group-2 (REC 20190719 15:59:02)	Group-3 (REC 20190719 15:59:02)	SSH	None	Group-2 (REC 20190719 15:59:02)	Allow <input checked="" type="checkbox"/>

- b In the **Recommended FW Rules** tab, review the firewall rule details. To modify any of the details, click the value in the appropriate column and select the edit (pencil) icon.
- c To define how the packets are to be handled, select **Allow**, **Drop**, or **Reject** in the **Action** column.
- d Toggle the button on the right-side to enable or disable the rule. By default the rule that was generated is set to be enabled when published, as shown in the image in the previous step.
- e Click **Recommended Groups**.

- f Click the link in the **Members** column to review the details about the VMs and IPs that were set for the group recommendation.
 - g Click the menu icon (three-dots) next to the group's name and select **Edit** to modify the group recommendation.
 - h Click **Recommended Services** and review the details.
 - i Click the menu icon (three-dots) next to the service's name and select **Edit** to modify the name or description. Before you delete a service, make sure that there are no rules using the service.
 - j Click **Next**.
- 7 In the **Place rules in FW context** pane, you can change the order in which the rule recommendation is to be applied with the existing firewall rules. Drag the highlighted section, or click the three-dot menu icon and select **Move Above selected section** or **Move Below selected section**.
 - 8 Click **Publish**.
 - 9 In the **Publish Recommendations** dialog box, click **Yes**.
 - 10 In the Enforcement Summary page, verify that the security policies have been published successfully and click **Close**.

The Status column for the recommendation is changed to Published in the table of Recommendations.

Results

Once the security policy recommendations have been published successfully, they are in read-only mode in the **Plan & Troubleshoot > Recommendations** tab. To view and manage the published rule recommendations, go to **Security > Distributed Firewall**.

Important After you have published the rule recommendations, the visualization continues to display the affected flows between the VMs as orange-hued arrows (Unprotected Flows) until new flows are generated between the affected VMs. The visualization only reports traffic flows based on the time when they occurred on the host and does not reflect the rule set published after those traffic flows occurred. After the rule set is published and new traffic flows are generated, the new flows are displayed as green-hued arrows (Allowed Flows).

Backing Up and Restoring NSX Intelligence

If your current NSX Intelligence configuration becomes inoperable or if you want to restore it to a previous state, you can restore your configuration from a backup. The backup and restore workflow is only supported using the NSX Intelligence CLI.

When you take a backup, NSX Intelligence only backs up the configuration files used by all the services that comprise the NSX Intelligence appliance. There is no visualization data included in the backup.

If data loss or corruption occurs in NSX Intelligence, all the existing data for the correlated flows and recommendations are also lost. Reinstalling NSX Intelligence restarts the collection of network traffic data and the visualization of those collected data is available from that point onwards.

After you finish the backup configuration, you can manually run the backup command on the NSX Intelligence appliance at any time. The backup is encrypted, compressed, and stored at the remote server defined during the backup configuration. When you create a backup, the date and time the backup is taken are appended to the backup filename so that each backup file is unique. For example, `config-backup-2019-06-21T21_06_07UTC.tar.gz`.

When you restore an NSX Intelligence backup, the configuration state when the backup was captured is restored. You must restore the backup to an NSX Intelligence appliance that is running the same version of the NSX Intelligence appliance from which the backup file was created. You can restore to an existing NSX Intelligence appliance or restore to a freshly installed NSX Intelligence appliance, but they must be the same version as the NSX Intelligence appliance you backed up.

Configure NSX Intelligence Backups

You must configure a backup file server before you can take a backup of your NSX Intelligence configuration. After a backup file server is configured, you can take a backup of NSX Intelligence at any time.

Prerequisites

- Verify that you have the CLI admin credentials to the NSX Intelligence CLI.
- Ensure that you have the user name and password for the remote server.
- Obtain the file path to where the backup files are to be stored in the remote server.

Procedure

- 1 From a command-line prompt, log in with admin privileges to the NSX Intelligence CLI host.

```
$ ssh admin@cli-ip-address
admin@cli-ip-address's password:
```

- 2 Configure the backup file server.

The command syntax is

```
set backup remote-host remote_host_address remote-path remote_folder_path remote-host-
username remote_host_username remote-host-password remote_host_password passphrase
pass_phrase
```

where the *remote_host_address* is the remote host IP or FQDN address of the backup file server and *remote_host_username* account must have the necessary privileges to create the backup files in the *remote_folder_path*. You must provide a strong value for the *passphrase* parameter. It must be at least eight characters long and has at least one uppercase, one lowercase, and one special character. For example,

```
set backup remote-host 10.11.22.33 remote-path /root remote-host-username root remote-host-
password MyRemotePassword passphrase MyPassPhra$e
```

3 Verify the configuration.

```
get configuration
```

From the output, verify that the line with `set backup` looks correct. Using the example in the previous step, the output must include the following line.

```
set backup remote-host 10.11.22.33 remote-path /root remote-host-username root
```

Back Up NSX Intelligence

You can back up your NSX Intelligence appliance configuration files using the CLI command.

Prerequisites

- Ensure that you have an admin access to the NSX Intelligence CLI.
- Configure a backup file server. See [Configure NSX Intelligence Backups](#).

Procedure

- 1 Log in with admin privileges to the NSX Intelligence CLI.
- 2 Create the backup.

```
backup intelligence configuration
```

If the backup is successful, you see a message similar to the following.

```
Backup Complete. Archived at: backup_file_server-IP_address:/root/backup_archives/
intelligence-config-backup-2019-07-18T07_00_26UTC.tar.gz
```

- 3 You can view the progress of the backup using another CLI session.
 - a Log in to another NSX Intelligence CLI session.
 - b Enter the following command.

```
get log-file node-mgmt.log follow
```

Restore NSX Intelligence Backups

When you restore a backup, you are restoring the state of the NSX Intelligence configuration at the time the backup was made. You can restore an NSX Intelligence backup using a CLI command.

You must restore a backup on an installation of the NSX Intelligence appliance that is the same version as the backup you are restoring. By default, the backup file restored is the most recently generated backup. If you are restoring a backup to a newly installed NSX Intelligence appliance, set the archive name before restoring the backup.

Prerequisites

- Verify that you have the admin login credentials and host info for the backup file server.
- Ensure that you have an admin access to the NSX Intelligence CLI.

Procedure

- 1 Log in with admin privileges to the new NSX Intelligence CLI server.
- 2 Configure the remote server where the backups are located.

The command syntax is

```
set restore remote-host backup_server_IP_address remote-path remote_folder_path remote-host-username remote_host_username remote-host-password remote_host_password passphrase pass_phrase
```

where the *backup_server_IP_address* is the remote host IP or FQDN address of the backup file server, *remote_host_username* account must have the necessary privileges to access the backup files in the *remote_folder_path*. For example,

```
set restore remote-host 10.11.22.33 remote-path /root remote-host-username root remote-host-password MyRemotePassword passphrase MyPassPhra$e
```

- 3 Verify the restore configuration.

```
get configuration
```

From the output, verify that the line with `set restore` looks correct. Using the example in the previous step, the output must include the following line.

```
set restore remote-host 10.11.22.33 remote-path /root remote-host-username root
```

- 4 Restore the backup with the following command.

```
restore intelligence configuration
```

If the restoration is successful, you see a message similar to the following.

```
NSX Intelligence Restore Complete.
```

- 5 You can view the progress of the backup restore using another CLI session.
 - a Log in to another NSX Intelligence CLI session.
 - b Enter the following command.

```
get log-file node-mgmt.log follow
```

Troubleshooting NSX Intelligence Issues

If the NSX Intelligence appliance becomes unresponsive or you need more details about an error message you received while using the appliance, you can run specific commands to get the state of the NSX Intelligence services.

You can also collect support bundles to assist you and VMware support personnel in debugging issues you might have encountered.

Check the Status of the NSX Intelligence Appliance

If the NSX Intelligence appliance becomes unresponsive, check the status of the NSX Intelligence services.

Problem

The NSX Intelligence appliance has become unresponsive or you received an error message that indicates the appliance is not functioning as expected.

Cause

It is possible that one or more of the underlying NSX Intelligence services has stopped or is not in a healthy state.

Solution

- 1 Log in to the NSX Intelligence appliance CLI host using an account with an Enterprise Administrator role.
- 2 Check the status of the NSX Intelligence services using the `get services` command.

If all the NSX Intelligence services are functioning properly, you see an output similar to the following example.

```
my_nsx-intel> get services
Service name:          druid
Service state:        running
Coordinator health:   good
Broker health:        good
Historical health:    good
Overlord health:      good
MiddleManager health: good

Service name:          http
Service state:        running
```

```
Session timeout:          1800
Connection timeout:      30
Redirect host:           (not configured)
Client API rate limit:   100 requests/sec
Client API concurrency limit: 40
Global API concurrency limit: 199

Service name:            kafka
Service state:           running
Service health:          good

Service name:            liagent
Service state:           stopped

Service name:            mgmt-plane-bus
Service state:           stopped

Service name:            node-mgmt
Service state:           running

Service name:            nsx-config
Service state:           running

Service name:            nsx-message-bus
Service state:           stopped

Service name:            nsx-upgrade-agent
Service state:           running

Service name:            ntp
Service state:           running
Start on boot:          True

Service name:            pace-server
Service state:           running

Service name:            postgres
Service state:           running
Service health:          good

Service name:            processing
Service state:           running

Service name:            snmp
Service state:           stopped
Start on boot:          False

Service name:            spark
Service state:           running
Service health:          good

Service name:            spark-job-scheduler
Service state:           running

Service name:            ssh
```

```

Service state:           running
Start on boot:          True

Service name:           syslog
Service state:          running

Service name:           ui-service
Service state:          running

Service name:           zookeeper
Service state:          running
Service health:         good

my_nsx-intel>

```

A service state can either be running or stopped. A service health can be good or degraded.

- 3 You can also view the `syslog` file and search for the output of the `pace-monitor.sh` health-check script that logs the health of the NSX Intelligence services to the `syslog` file.

If all the services are functioning as expected, you see an output similar to the following sample output after running the `get log-file syslog | find pace-monitor` command.

```

my_nsx-intel> get log-file syslog | find pace-monitor
<13>1 2019-08-30T03:19:20.409899+00:00 my_nsx-intel pace-monitor.sh - - -   "_self": {
<13>1 2019-08-30T03:19:20.410253+00:00 my_nsx-intel pace-monitor.sh - - -     "href": "/"
node/pace/appliance-health",
<13>1 2019-08-30T03:19:20.410623+00:00 my_nsx-intel pace-monitor.sh - - -     "rel":
"self"
<13>1 2019-08-30T03:19:20.410908+00:00 my_nsx-intel pace-monitor.sh - - -   },
<13>1 2019-08-30T03:19:20.411162+00:00 my_nsx-intel pace-monitor.sh - - -   "appliance-
health": {
<13>1 2019-08-30T03:19:20.411416+00:00 my_nsx-intel pace-monitor.sh - - -     "status":
"Following NSX Intelligence first boot services are either PENDING or FAILED - Token-
Registration",
<13>1 2019-08-30T03:19:20.411668+00:00 my_nsx-intel pace-monitor.sh - - -     "sub-system-
status": {
<13>1 2019-08-30T03:19:20.411923+00:00 my_nsx-intel pace-monitor.sh - - -       "app-
services": {
<13>1 2019-08-30T03:19:20.412280+00:00 my_nsx-intel pace-monitor.sh - - -     "services": [],
<13>1 2019-08-30T03:19:20.412528+00:00 my_nsx-intel pace-monitor.sh - - -     "status": ""
<13>1 2019-08-30T03:19:20.412807+00:00 my_nsx-intel pace-monitor.sh - - -   },
<13>1 2019-08-30T03:19:20.413075+00:00 my_nsx-intel pace-monitor.sh - - -     "base-
infra-services": {
<13>1 2019-08-30T03:19:20.413303+00:00 my_nsx-intel pace-monitor.sh - - -     "services": [
<13>1 2019-08-30T03:19:20.413613+00:00 my_nsx-intel pace-monitor.sh - - -       {
<13>1 2019-08-30T03:19:20.413848+00:00 my_nsx-intel pace-monitor.sh - - -     "druid-health": {
<13>1 2019-08-30T03:19:20.414146+00:00 my_nsx-intel pace-monitor.sh - - -     "broker": "good",
<13>1 2019-08-30T03:19:20.414473+00:00 my_nsx-intel pace-monitor.sh - - -

```

```

"coordinator": "good",
<13>1 2019-08-30T03:19:20.414717+00:00 my_nsx-intel pace-monitor.sh - - -
"historical": "good",
<13>1 2019-08-30T03:19:20.414979+00:00 my_nsx-intel pace-monitor.sh - - -
"middlemanager": "good",
<13>1 2019-08-30T03:19:20.415295+00:00 my_nsx-intel pace-monitor.sh - - -
"overlord": "good"
<13>1 2019-08-30T03:19:20.415533+00:00 my_nsx-intel pace-monitor.sh - - -           },
<13>1 2019-08-30T03:19:20.415762+00:00 my_nsx-intel pace-monitor.sh - - -
"service-name": "druid"
<13>1 2019-08-30T03:19:20.415982+00:00 my_nsx-intel pace-monitor.sh - - -           },
<13>1 2019-08-30T03:19:20.416269+00:00 my_nsx-intel pace-monitor.sh - - -           {
<13>1 2019-08-30T03:19:20.416539+00:00 my_nsx-intel pace-monitor.sh - - -
"health": "good",
<13>1 2019-08-30T03:19:20.416772+00:00 my_nsx-intel pace-monitor.sh - - -
"service-name": "kafka"
<13>1 2019-08-30T03:19:20.416991+00:00 my_nsx-intel pace-monitor.sh - - -           },
<13>1 2019-08-30T03:19:20.417204+00:00 my_nsx-intel pace-monitor.sh - - -           {
<13>1 2019-08-30T03:19:20.417510+00:00 my_nsx-intel pace-monitor.sh - - -
"health": "good",
<13>1 2019-08-30T03:19:20.417745+00:00 my_nsx-intel pace-monitor.sh - - -
"service-name": "postgres"
<13>1 2019-08-30T03:19:20.418133+00:00 my_nsx-intel pace-monitor.sh - - -           },
<13>1 2019-08-30T03:19:20.418389+00:00 my_nsx-intel pace-monitor.sh - - -           {
<13>1 2019-08-30T03:19:20.418626+00:00 my_nsx-intel pace-monitor.sh - - -
"health": "good",
<13>1 2019-08-30T03:19:20.418855+00:00 my_nsx-intel pace-monitor.sh - - -
"service-name": "spark"
<13>1 2019-08-30T03:19:20.419157+00:00 my_nsx-intel pace-monitor.sh - - -           },
<13>1 2019-08-30T03:19:20.419435+00:00 my_nsx-intel pace-monitor.sh - - -           {
<13>1 2019-08-30T03:19:20.419684+00:00 my_nsx-intel pace-monitor.sh - - -
"health": "good",
<13>1 2019-08-30T03:19:20.419928+00:00 my_nsx-intel pace-monitor.sh - - -
"service-name": "zookeeper"
<13>1 2019-08-30T03:19:20.420165+00:00 my_nsx-intel pace-monitor.sh - - -           }
<13>1 2019-08-30T03:19:20.420496+00:00 my_nsx-intel pace-monitor.sh - - -           ],
<13>1 2019-08-30T03:19:20.420786+00:00 my_nsx-intel pace-monitor.sh - - -
"status": ""
<13>1 2019-08-30T03:19:20.421022+00:00 my_nsx-intel pace-monitor.sh - - -           },
<13>1 2019-08-30T03:19:20.421255+00:00 my_nsx-intel pace-monitor.sh - - -           "first-
boot-services": {
<13>1 2019-08-30T03:19:20.421539+00:00 my_nsx-intel pace-monitor.sh - - -
"services": [
<13>1 2019-08-30T03:19:20.421777+00:00 my_nsx-intel pace-monitor.sh - - -           {
<13>1 2019-08-30T03:19:20.422010+00:00 my_nsx-intel pace-monitor.sh - - -
"health": "degraded",
<13>1 2019-08-30T03:19:20.422277+00:00 my_nsx-intel pace-monitor.sh - - -
"service-name": "token-registration"
<13>1 2019-08-30T03:19:20.422512+00:00 my_nsx-intel pace-monitor.sh - - -           }
<13>1 2019-08-30T03:19:20.422770+00:00 my_nsx-intel pace-monitor.sh - - -           ],
<13>1 2019-08-30T03:19:20.423012+00:00 my_nsx-intel pace-monitor.sh - - -
"status": "Following NSX Intelligence first boot, services are either PENDING or FAILED
- Token-Registration"
<13>1 2019-08-30T03:19:20.423354+00:00 my_nsx-intel pace-monitor.sh - - -           }
<13>1 2019-08-30T03:19:20.423601+00:00 my_nsx-intel pace-monitor.sh - - -           }

```

```
<13>1 2019-08-30T03:19:20.423882+00:00 my_nsx-intel pace-monitor.sh - - - }
<13>1 2019-08-30T03:19:20.424339+00:00 my_nsx-intel pace-monitor.sh - - - }
<13>1 2019-08-30T03:19:20.972629+00:00 my_nsx-intel pace-monitor.sh - - - NSX
Intelligence health OK.
<30>1 2019-08-30T03:19:20.973076+00:00 my_nsx-intel pace-monitor 20804 - - <13>Aug 30
03:19:19 pace-monitor.sh: NSX Intelligence health OK.
<182>1 2019-08-30T03:23:23.857Z my_nsx-intel NSX 21752 - [nsx@6876 comp="nsx-cli"
subcomp="node-mgmt" username="admin" level="INFO"] CMD: get log-file syslog | find pace-
monitor
```

If there is a problem with one of the services, you might see the following line when you run `get log-file syslog | grep pace-monitor`.

```
NSX Intelligence health DEGRADED. Return code not HTTP OK.
```

4 If you encounter one of the following outputs, restart the service using the `restart service service-name` command.

- After running the `get services` command, one of the services shows `Service state: stopped` or `Service health: degraded`.
- After running the `get log-file syslog | grep pace-monitor` command, the output shows something similar to the `PACE health DEGRADED. Return code not HTTP OK.` message.

For example, if the `postgres` service's state shows it is stopped, or if its state is running, but it has a degraded service health, run the following command.

```
restart service postgres
```

Important You must use the `restart service service-name` command to restart NSX Intelligence services. If you decide to use the `stop service service-name` and `start service service-name` commands instead, you have to also manually restart each of the services that depend on `service-name`. The following list shows the dependency order in which the NSX Intelligence services have to be restarted.

```
zookeeper > druid > kafka > spark > spark-job-scheduler > nsx-config > processing > pace-
server
```

For example, if the `nsx-config` service is stopped and then started using the `stop|start service service-name` command, you must also use the `restart service service-name` command to restart the `processing` and `pace-server` services.

In addition, if you use the `restart service service-name` command to restart any services shown in the dependency order list before the `spark-job-scheduler` service, you must also manually restart the `spark-job-scheduler` service using the `restart service spark-job-scheduler` command. Failure to do so results in the `spark-job-scheduler` service getting into a bad state.

Collect NSX Intelligence Support Bundles

You can collect a support bundle using the NSX Intelligence CLI.

The support bundle file contents do not include data. It includes files in the following directories.

- `/opt/vmware/*`
- `/var/log/*`
- `/etc/*`
- System state using `journalctl` and `systemctl`

Prerequisites

Ensure that you have an Enterprise Administrator access to the NSX Intelligence CLI.

Procedure

- 1 Log in to the NSX Intelligence CLI using an account with Enterprise Administrator role privileges.
- 2 Generate the support bundle.

The command syntax is as follows, where you provide the value for `support_filename.tgz`.

```
get support-bundle file support_filename.tgz
```

For example,

```
get support-bundle file support_bundle123.tgz
```

When the bundle file is created successfully, you receive the messages similar to the following example.

```
support_bundle123.tgz created, use the following command to transfer the file:
```

```
copy file support_bundle123.tgz url <url>
```

```
After transferring support_bundle123.tgz, extract it using:tar xvf support_bundle123.tgz
```

- 3 Verify that the support bundle exists using the following command.

```
get files
```

You receive an output similar to the following.

```
Directory of filestore:/
-rw- 21377586 June 29 05:29:12 UTC support_bundle123.tgz
```