

NSX Administration Guide

Update 17

Modified on 25 AUGUST 2022

VMware NSX Data Center for vSphere 6.4

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 © 2010 - 2022 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

NSX Administration Guide	13
1 System Requirements for NSX Data Center for vSphere	14
2 Ports and Protocols Required by NSX Data Center for vSphere	17
3 Overview of NSX Data Center for vSphere	18
NSX Data Center for vSphere Components	20
Data Plane	20
Control Plane	21
Management Plane	22
Consumption Platform	22
NSX Edge	23
NSX Services	25
4 Overview of Cross-vCenter Networking and Security	28
Benefits of Cross-vCenter NSX	28
How Cross-vCenter NSX Works	29
Support Matrix for Services in Cross-vCenter NSX	30
Universal NSX Controller Cluster	32
Universal Transport Zone	32
Universal Logical Switches	32
Universal Logical (Distributed) Routers	33
Universal Firewall Rules	33
Universal Network and Security Objects	34
Cross-vCenter NSX Topologies	34
Multi-Site and Single Site Cross-vCenter NSX	34
Local Egress	37
Modifying NSX Manager Roles	38
5 Transport Zones	40
Understanding Replication Modes	42
Add a Transport Zone	45
Edit a Transport Zone	46
Expand a Transport Zone	47
Contract a Transport Zone	48
Controller Disconnected Operation (CDO) Mode	48

6 Logical Switches 49

- Add a Logical Switch 51
 - Add a Logical Switch 52
 - Connect a Logical Switch to an NSX Edge 54
 - Deploy Services on a Logical Switch 55
- Connect Virtual Machines to a Logical Switch 55
- Test Logical Switch Connectivity 56
- Prevent Spoofing on a Logical Switch 56
- Edit a Logical Switch 56
- Logical Switch Scenario 57
 - John Admin Assigns Segment ID Pool to NSX Manager 59
 - John Admin Configures VXLAN Transport Parameters 60
 - John Admin Adds a Transport Zone 61
 - John Admin Creates a Logical Switch 61

7 Configuring Hardware Gateway 62

- Scenario: Hardware Gateway Sample Configuration 63
 - Set Up the Replication Cluster 64
 - Connect the Hardware Gateway to the NSX Controllers 65
 - Add Hardware Gateway Certificate 66
 - Bind the Logical Switch to the Physical Switch 68

8 L2 Bridges 70

- Add L2 Bridge 71
- Add L2 Bridge to a Logically Routed Environment 72
- Improving Bridging Throughput 73
 - Enable Software Receive Side Scaling 74

9 Routing 75

- Add a Distributed Logical Router 75
- Add an Edge Services Gateway 88
- Specify Global Configuration 96
- NSX Edge Configuration 98
 - Working with Certificates 98
 - FIPS Mode 103
 - Managing NSX Edge Appliances 105
 - Managing NSX Edge Appliance Resource Reservations 107
 - Working with Interfaces 110
 - Add a Sub Interface 113
 - Change Auto Rule Configuration 117
 - Change CLI Credentials 117

- About High Availability 117
- Force Sync NSX Edge with NSX Manager 121
- Configure Syslog Servers for NSX Edge 122
- View the Status of NSX Edge Services 123
- Redeploy NSX Edge 123
- Download Tech Support Logs for NSX Edge 126
- Add a Static Route 126
- Configure OSPF on a Logical (Distributed) Router 128
- Configure OSPF on an Edge Services Gateway 133
- Configure BGP 138
- Configure Route Redistribution 143
- View the NSX Manager Locale ID 144
- Configure Locale ID on a Universal Logical (Distributed) Router 144
- Configure Locale ID on a Host or Cluster 145
- Multicast Routing Support, Limitations, and Topology 146
 - Configure Multicast on a Logical (Distributed) Router 147
 - Configure Multicast on an Edge Services Gateway 150
 - Multicast Topology 152

10 Logical Firewall 154

- Distributed Firewall 154
 - Context-Aware Firewall 156
 - Session Timers 166
 - IP Discovery for Virtual Machines 168
 - Exclude Virtual Machines from Firewall Protection 170
 - View Firewall CPU and Memory Threshold Events 171
 - Distributed Firewall Resource Utilization 172
- Edge Firewall 172
 - Working with NSX Edge Firewall Rules 172
- Working with Firewall Rule Sections 192
 - Add a Firewall Rule Section 192
 - Merge Firewall Rule Sections 193
 - Delete a Firewall Rule Section 194
 - Lock Firewall Rule Sections 194
 - Unlock Firewall Rule Sections 195
- Working with Firewall Rules 196
 - Add a Firewall Rule 196
 - Edit the Default Distributed Firewall Rule 204
 - Force Sync Distributed Firewall Rules 204
 - Firewall Rules with a Custom Layer 3 Protocol 205
 - Save an Unpublished Configuration 205

- Load a Saved Firewall Configuration 206
- Filter Firewall Rules 207
- Change the Order of a Firewall Rule 207
- Firewall Rule Behavior in Security Groups 208
- Firewall Rule Hit Count and Reset 208
- Firewall Logs 209

- 11 Firewall Scenarios 214**
 - Context-Aware Firewall Scenarios 214
 - Configuring Application Identification 216

- 12 Identity Firewall Overview 217**
 - Identity Firewall Workflow 218
 - Identity Firewall Tested and Supported Configurations 219

- 13 Working with Active Directory Domains 224**
 - Register a Windows Domain with NSX Manager 224
 - Synchronize a Windows Domain with Active Directory 226
 - Edit a Windows Domain 227
 - Enable Security Read-Only Log Access on Windows 2008 227
 - Verifying Directory Privileges 228

- 14 Using SpoofGuard 230**
 - Create a SpoofGuard Policy 231
 - Approve IP Addresses 232
 - Change an IP Address 233
 - Clear an IP Address 234

- 15 Virtual Private Networks (VPN) 236**
 - SSL VPN-Plus Overview 236
 - Configure Network Access SSL VPN-Plus 238
 - Install SSL VPN-Plus Client 248
 - Configure Proxy Server Settings in SSL VPN-Plus Client 251
 - SSL VPN-Plus Logs 252
 - Edit Client Configuration 253
 - Edit General Settings 253
 - Edit Web Portal Design 254
 - Working with IP Pools for SSL VPN 254
 - Working with Private Networks 256
 - Working with Installation Packages 258
 - Working with Users 258

Working with Login and Logoff Scripts	259
IPSec VPN Overview	260
Policy-Based IPSec VPN	261
Route-Based IPSec VPN	262
Configure Policy-Based IPSec VPN Site	264
Configure Route-Based IPSec VPN Site	275
Edit IPSec VPN Site	276
Disable IPSec VPN Site	276
Delete IPSec VPN Site	277
IPsec Terminology	277
IKEv1 Phase 1 and Phase 2	277
Configure Policy-Based IPSec VPN Site Example	280
Using a Cisco 2821 Integrated Services Router	282
Using a Cisco ASA 5510	285
Using a Cisco CSR 1000V Appliance	287
Configuring a WatchGuard Firebox X500	290
L2 VPN Overview	291
L2 VPN Best Practices	293
L2 VPN Over SSL	299
L2 VPN Over IPSec	306
Standalone Edge as L2 VPN Client	307
Scenario: Add a Stretched VLAN or VXLAN Network	315
Scenario: Remove a Stretched VLAN or VXLAN Network	319
16 Logical Load Balancer	321
Setting Up Load Balancing	325
Configure Load Balancer Service	327
Create a Service Monitor	328
Add a Server Pool	336
Create an Application Profile	339
Add an Application Rule	347
Add Virtual Servers	354
Managing Application Profiles	356
Edit an Application Profile	356
Configure SSL Termination for a Load Balancer	356
Delete an Application Profile	357
Managing Service Monitors	357
Edit a Service Monitor	357
Delete a Service Monitor	358
Managing Server Pools	358
Edit a Server Pool	358

- Configure a Load Balancer to Use Transparent Mode 359
 - Delete a Server Pool 359
 - Show Pool Status 360
 - Managing Virtual Servers 360
 - Edit a Virtual Server 360
 - Delete a Virtual Server 361
 - Managing Application Rules 361
 - Edit an Application Rule 361
 - Delete an Application Rule 362
 - Load Balance Web Servers using NTLM Authentication 362
 - Load Balancer HTTP Connection Modes 362
 - Scenarios for NSX Load Balancer Configuration 365
 - Scenario: Configure a One-Armed Load Balancer 365
 - Scenario: Configure an Inline Load Balancer 369
 - Scenario: Configure NSX Load Balancer for Platform Service Controller 372
 - Scenario: SSL Offloading 375
 - Scenario: Import SSL Certificate 380
 - Scenario: SSL Passthrough 383
 - Scenario: SSL Client and Server Authentication 385
- 17 Other Edge Services 388**
 - Managing DHCP Service 388
 - Add a DHCP IP Pool 389
 - Start the DHCP Service 390
 - Edit DHCP IP Pool 391
 - Add a DHCP Static Binding 391
 - Edit DHCP Binding 393
 - Configuring DHCP Relay 393
 - Add DHCP Relay Server 394
 - Add DHCP Relay Agent 395
 - Configure a DNS Server 395
- 18 Service Composer 397**
 - Using Service Composer 399
 - Create a Security Group in Service Composer 400
 - Global Settings 403
 - Create a Security Policy 405
 - Apply a Security Policy to a Security Group 409
 - Service Composer Canvas 410
 - Working with Security Tags 412
 - Unique ID Selection 413

View Applied Security Tags	414
Create a Security Tag	414
Assign a Security Tag	415
Edit a Security Tag	415
Delete a Security Tag	415
Viewing Effective Services	416
View Effective Services on a Security Policy	416
View Service Failures for a Security Policy	416
View Effective Services on a Virtual Machine	417
Working with Security Policies	417
Manage Security Policy Priority	417
Edit a Security Policy	418
Delete a Security Policy	418
Importing and Exporting Security Policy Configurations	419
Export a Security Policy Configuration	419
Import a Security Policy Configuration	420
Service Composer Scenarios	421
Quarantining Infected Machines Scenario	421
Backing up Security Configurations	425
19 Guest Introspection	428
Guest Introspection Architecture	429
Install Guest Introspection on Host Clusters	431
Install the Guest Introspection Thin Agent on Windows Virtual Machines	433
Install the Guest Introspection Thin Agent on Linux Virtual Machines	435
View Guest Introspection Status	437
Guest Introspection Audit Messages	437
Guest Introspection Events	438
Uninstall a Guest Introspection Module	439
Uninstall Guest Introspection for Linux	440
20 Network Extensibility	441
Distributed Service Insertion	442
Edge-Based Service Insertion	442
Integrating Third Party Services	442
Deploy a Partner Service	443
Consuming Vendor Services through Service Composer	445
Redirecting Traffic to a Vendor Solution through Logical Firewall	445
Using a Partner Load Balancer	446
Remove Third-Party Integration	447

21 User Management 448

- NSX Users and Permissions by Feature 448
- Configure Single Sign-On 457
- Managing User Rights 459
- Managing the Default User Account 460
- Assign a Role to a vCenter User 460
- Group-Based Role Assignments 461
- Create a User with Web Interface Access Using CLI 464
- Edit a User Account 466
- Change a User Role 466
- Disable or Enable a User Account 467
- Delete a User Account 467

22 Network and Security Objects 468

- Working with IP Address Groups 468
 - Create an IP Address Group 469
 - Edit an IP Address Group 469
 - Delete an IP Address Group 470
- Working with MAC Address Groups 470
 - Create a MAC Address Group 470
 - Edit a MAC Address Group 471
 - Delete a MAC Address Group 471
- Working with IP Pools 472
 - Create an IP Pool 472
 - Edit an IP Pool 472
 - Delete an IP Pool 473
- Working with Security Groups 473
 - Firewall Rule Behavior in Security Groups 474
 - Create a Security Group 474
 - Edit a Security Group 478
 - Delete a Security Group 478
- Working with Services and Service Groups 479
 - Create a Service 479
 - Create a Service Group 480
 - Edit a Service or Service Group 480
 - Delete a Service or Service Group 481

23 Operations and Management 482

- Add and Assign a License 482
- Using the Dashboard 484
 - Custom Widget 485

- System Scale Dashboard 486
- Check Communication Channel Health 487
- NSX Controller Management 487
 - Change NSX Controller Name 487
 - Change Controller Password 488
 - Configure DNS, NTP, and Syslog for the NSX Controller Cluster 488
 - Download Technical Support Logs for NSX Controller 490
- Controller Disconnected Mode for Multiple Sites 491
 - Enable Controller Disconnected Operation (CDO) Mode 492
 - Disable Controller Disconnected Operation (CDO) Mode 493
 - Resync CDO Configuration 493
- Change VXLAN Port 494
- Customer Experience Improvement Program 496
 - Edit the Customer Experience Improvement Program Option 496
- About NSX Logs 497
- Audit Logs 498
 - Using NSX Ticket Logger 498
 - View the Audit Log 499
- System Events 499
 - View the System Event Report 499
 - Format of a System Event 499
 - Alarms 500
 - Format of an Alarm 501
 - Working with SNMP Traps 501
- Management System Settings 505
 - Log In to the NSX Manager Virtual Appliance 505
 - Edit the NSX Manager Date and Time 506
 - Change NSX Manager Appliance IP Address 506
 - Configure a Syslog Server for NSX Manager 508
 - Change FIPS Mode and TLS Settings on NSX Manager 509
 - Edit DNS Servers 511
 - Edit Lookup Service Details 511
 - Edit vCenter Server 511
 - Download Technical Support Logs for NSX 512
 - NSX Manager SSL Certification 512
- NSX Backup and Restore 515
 - Back Up and Restore NSX Manager 516
 - Back Up vSphere Distributed Switches 522
 - Back Up vCenter 522
- NSX Monitoring And Diagnostic Tools 522
 - Flow Monitoring 522

- Configure IPFIX 529
- Application Rule Manager 548
- Host Health Status Monitoring 558
- Network Latency Monitoring 560
- Endpoint Monitoring Data Collection 561
- Traceflow 564
- Packet Capture 567
- Support Bundle Collection Tool 570

24 Disaster Recovery Scenarios with Cross-vCenter NSX 575

- Scenario 1: Scheduled Full Site Failure 579
- Scenario 2: Unscheduled Full Site Failure 583
- Scenario 3: Full Failback to Primary Site 586

NSX Administration Guide

The *NSX Administration Guide* describes how to configure, monitor, and maintain the VMware NSX® Data Center for vSphere® system by using the VMware NSX® Manager™ user interface, the VMware vSphere® Web Client, and the VMware vSphere® Client™. The information includes step-by-step configuration instructions, and suggested best practices.

Important NSX for vSphere is now known as NSX Data Center for vSphere.

Intended Audience

This manual is intended for anyone who wants to install or use NSX Data Center for vSphere in a VMware vSphere® environment. The information in this manual is written for experienced system administrators who are familiar with virtual machine technology and virtual datacenter operations. This manual assumes familiarity with vSphere, including VMware ESXi™, VMware vCenter Server®, and the vSphere Web Client.

Task Instructions

Task instructions in this guide are based on the vSphere Web Client. You can also perform some of the tasks in this guide by using the new vSphere Client. The new vSphere Client user interface terminology, topology, and workflow are closely aligned with the same aspects and elements of the vSphere Web Client.

Note Not all functionality of the NSX plug-in for the vSphere Web Client has been implemented for the vSphere Client in NSX 6.4. For an up-to-date list of supported and unsupported functionality, see <https://docs.vmware.com/en/VMware-NSX-for-vSphere/6.4/rn/nsx-vsphere-client-65-functionality-support.html>.

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 <http://www.vmware.com/support/pubs>.

System Requirements for NSX Data Center for vSphere

1

Before you install or upgrade NSX Data Center for vSphere, consider your network configuration and resources. You can install one NSX Manager per vCenter Server, one instance of Guest Introspection per ESXi host, and multiple NSX Edge instances per datacenter.

Hardware

This table lists the hardware requirements for NSX Data Center for vSphere appliances.

Table 1-1. Hardware Requirements for Appliances

Appliance	Memory	vCPU	Disk Space
NSX Manager	16 GB (24 GB for larger deployments)	4 For large deployments: Do not exceed 8 (6.4.0 to 6.4.6) Do not exceed 32 (starting in 6.4.7)	60 GB
NSX Controller	4 GB	4	28 GB
NSX Edge (Distributed logical router is deployed as compact appliance)	Compact: 512 MB Large: 1 GB Quad Large: 4 GB (starting in 6.4.9), 2 GB (6.4.0 to 6.4.8) X-Large: 8 GB	Compact: 1 Large: 2 Quad Large: 4 X-Large: 6	Compact, Large: 1 disk 640 MB + 1 disk 512 MB Quad Large: 1 disk 640 MB + 2 disks 512 MB X-Large: 1 disk 640 MB + 1 disk 2 GB + 1 disk 512 MB
Guest Introspection	2 GB	2	5 GB (Provisioned space is 6.26 GB)

As a general guideline, if your NSX-managed environment contains more than 256 hypervisors, increase NSX Manager resources to at least 8 vCPU and 24 GB of RAM. Do not exceed 32 vCPU. For more information on configuration maximums, see the NSX Data Center for vSphere section of the [VMware Configuration Maximums](#) tool. The documented configuration maximums all assume the large NSX Manager appliance size. For specific sizing details contact VMware support.

For information about increasing the memory and vCPU allocation for your virtual appliances, see "Allocate Memory Resources", and "Change the Number of Virtual CPUs" in *vSphere Virtual Machine Administration*.

The provisioned space for a Guest Introspection appliance shows as 6.26 GB for Guest Introspection. This is because vSphere ESX Agent Manager creates a snapshot of the service VM to create fast clones, when multiple hosts in a cluster shares a storage. For more information on how to disable this option via ESX Agent Manager, refer to the *ESX Agent Manager* documentation.

Network Latency

You should ensure that the network latency between components is at or below the maximum latency described.

Table 1-2. Maximum network latency between components

Components	Maximum latency
NSX Manager and NSX Controller nodes	150 ms RTT
NSX Manager and ESXi hosts	150 ms RTT
NSX Manager and vCenter Server system	150 ms RTT
NSX Manager and NSX Manager in a cross-vCenter NSX environment	150 ms RTT
NSX Controller and ESXi hosts	150 ms RTT

Software

For the latest interoperability information, see the Product Interoperability Matrixes at http://partnerweb.vmware.com/comp_guide/sim/interop_matrix.php.

For recommended versions of NSX Data Center for vSphere, vCenter Server, and ESXi, see the release notes for the version of NSX Data Center for vSphere to which you are upgrading. Release notes are available at the NSX Data Center for vSphere documentation site: <https://docs.vmware.com/en/VMware-NSX-for-vSphere/index.html>.

For an NSX Manager to participate in a cross-vCenter NSX deployment the following conditions are required:

Component	Version
NSX Manager	6.2 or later
NSX Controller	6.2 or later
vCenter Server	6.0 or later
ESXi	ESXi 6.0 or later Host clusters prepared with NSX 6.2 or later VIBs

To manage all NSX Managers in a cross-vCenter NSX deployment from a single vSphere Web Client, you must connect your vCenter Server systems in Enhanced Linked Mode. See Using Enhanced Linked Mode in *vCenter Server and Host Management*.

To verify the compatibility of partner solutions with NSX, see the VMware Compatibility Guide for Networking and Security at <http://www.vmware.com/resources/compatibility/search.php?deviceCategory=security>.

Client and User Access

The following items are required to manage your NSX Data Center for vSphere environment:

- Forward and reverse name resolution. This is required if you have added ESXi hosts by name to the vSphere inventory, otherwise NSX Manager cannot resolve the IP addresses.
- Permissions to add and power on virtual machines.
- Access to the datastore where you store virtual machine files, and the account permissions to copy files to that datastore.
- Cookies must be enabled on your Web browser to access the NSX Manager user interface.
- Port 443 must be open between the NSX Manager and the ESXi host, the vCenter Server, and the NSX Data Center for vSphere appliances to be deployed. This port is required to download the OVF file on the ESXi host for deployment.
- A Web browser that is supported for the version of vSphere Client or vSphere Web Client you are using. See the list of supported Web browsers at <https://docs.vmware.com/en/VMware-NSX-Data-Center-for-vSphere/6.4/rn/nsx-vsphere-client-65-functionality-support.html>.

Note that Windows 32-bit and 64-bit Microsoft Internet Explorer browser is not supported with NSX 6.4.x.

- For information about using the vSphere Client (HTML5) on vSphere 6.5 with NSX Data Center for vSphere 6.4, see <https://docs.vmware.com/en/VMware-NSX-for-vSphere/6.4/rn/nsx-vsphere-client-65-functionality-support.html>.

Ports and Protocols Required by NSX Data Center for vSphere

2

NSX Data Center for vSphere requires multiple ports to be open for it to operate properly.

- If you have a cross-vCenter NSX environment and your vCenter Server systems are in Enhanced Linked Mode, each NSX Manager appliance must have the required connectivity to each vCenter Server system in the environment. In this mode, you can manage any NSX Manager from any vCenter Server system.
- When you are upgrading from an earlier NSX version to version 6.4.x, Guest Introspection and host clusters must be upgraded for the Remote Desktop Session Host (RDSH) policies to be created and enforced successfully.

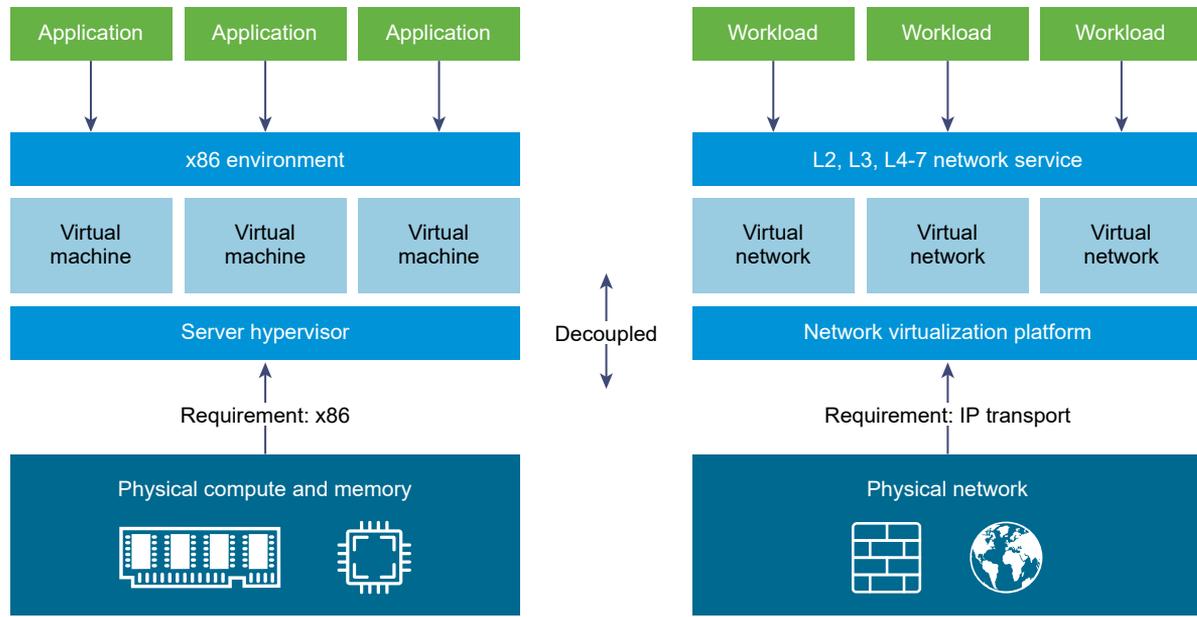
For the list of all supported ports and protocols in NSX 6.4.0 and later, see the VMware Ports and Protocols portal at <https://ports.vmware.com/home/NSX-Data-Center-for-vSphere>. You can use this portal to download the list in either CSV, Excel, or PDF file formats.

Overview of NSX Data Center for vSphere

3

IT organizations have gained significant benefits as a direct result of server virtualization. Server consolidation reduced physical complexity, increased operational efficiency and the ability to dynamically repurpose underlying resources to quickly and optimally meet the needs of increasingly dynamic business applications.

VMware's Software Defined Data Center (SDDC) architecture is now extending virtualization technologies across the entire physical data center infrastructure. NSX Data Center for vSphere is a key product in the SDDC architecture. With NSX Data Center for vSphere, virtualization delivers for networking what it has already delivered for compute and storage. In much the same way that server virtualization programmatically creates, snapshots, deletes, and restores software-based virtual machines (VMs), NSX Data Center for vSphere network virtualization programmatically creates, snapshots, deletes, and restores software-based virtual networks. The result is a transformative approach to networking that not only enables data center managers to achieve orders of magnitude better agility and economics, but also allows for a vastly simplified operational model for the underlying physical network. With the ability to be deployed on any IP network, including both existing traditional networking models and next-generation fabric architectures from any vendor, NSX Data Center for vSphere is a non-disruptive solution. In fact, with NSX Data Center for vSphere, the physical network infrastructure you already have is all you need to deploy a software-defined data center.



The figure above draws an analogy between compute and network virtualization. With server virtualization, a software abstraction layer (server hypervisor) reproduces the familiar attributes of an x86 physical server (for example, CPU, RAM, Disk, NIC) in software, allowing them to be programmatically assembled in any arbitrary combination to produce a unique VM in a matter of seconds.

With network virtualization, the functional equivalent of a network hypervisor reproduces the complete set of Layer 2 through Layer 7 networking services (for example, switching, routing, access control, firewalling, QoS, and load balancing) in software. As a result, these services can be programmatically assembled in any arbitrary combination, to produce unique, isolated virtual networks in a matter of seconds.

With network virtualization, benefits similar to server virtualization are derived. For example, just as VMs are independent of the underlying x86 platform and allow IT to treat physical hosts as a pool of compute capacity, virtual networks are independent of the underlying IP network hardware and allow IT to treat the physical network as a pool of transport capacity that can be consumed and repurposed on demand. Unlike legacy architectures, virtual networks can be provisioned, changed, stored, deleted, and restored programmatically without reconfiguring the underlying physical hardware or topology. By matching the capabilities and benefits derived from familiar server and storage virtualization solutions, this transformative approach to networking unleashes the full potential of the software-defined data center.

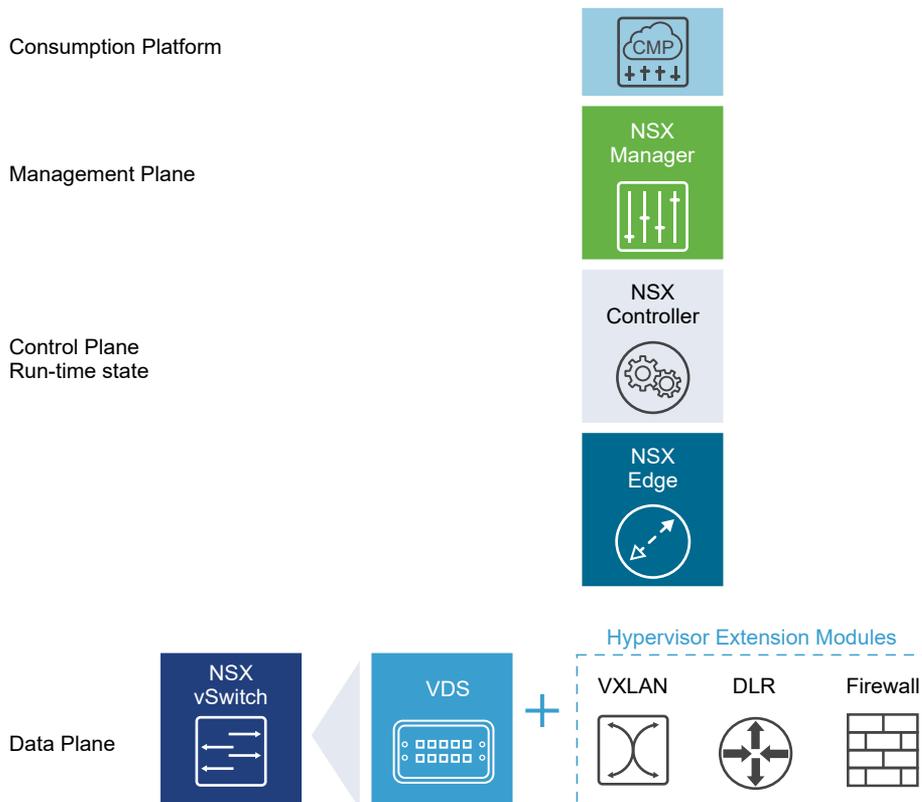
NSX Data Center for vSphere can be configured through the vSphere Web Client, a command-line interface (CLI), and a REST API.

This chapter includes the following topics:

- [NSX Data Center for vSphere Components](#)
- [NSX Edge](#)
- [NSX Services](#)

NSX Data Center for vSphere Components

This section describes the components of the NSX Data Center for vSphere solution.



Note that a cloud management platform (CMP) is not an NSX Data Center for vSphere component, but NSX Data Center for vSphere provides integration into virtually any CMP via the REST API and out-of-the-box integration with VMware CMPs.

Data Plane

The data plane consists of the NSX Virtual Switch, which is based on the vSphere Distributed Switch (VDS) with additional components to enable services. Kernel modules, userspace agents, configuration files, and install scripts are packaged in VIBs and run within the hypervisor kernel to provide services such as distributed routing and logical firewall and to enable VXLAN bridging capabilities.

The NSX Virtual Switch (vDS-based) abstracts the physical network and provides access-level switching in the hypervisor. It is central to network virtualization because it enables logical networks that are independent of physical constructs, such as VLANs. Some of the benefits of the vSwitch are:

- Support for overlay networking with protocols (such as VXLAN) and centralized network configuration. Overlay networking enables the following capabilities:
 - Reduced use of VLAN IDs in the physical network.

- Creation of a flexible logical Layer 2 (L2) overlay over existing IP networks on existing physical infrastructure without the need to re-architect any of the data center networks
- Provision of communication (east–west and north–south), while maintaining isolation between tenants
- Application workloads and virtual machines that are agnostic of the overlay network and operate as if they were connected to a physical L2 network
- Facilitates massive scale of hypervisors
- Multiple features—such as Port Mirroring, NetFlow/IPFIX, Configuration Backup and Restore, Network Health Check, QoS, and LACP—provide a comprehensive toolkit for traffic management, monitoring, and troubleshooting within a virtual network

The logical routers can provide L2 bridging from the logical networking space (VXLAN) to the physical network (VLAN).

The gateway device is typically an NSX Edge virtual appliance. NSX Edge offers L2, L3, perimeter firewall, load balancing, and other services such as SSL VPN and DHCP.

Control Plane

The control plane runs in the NSX Controller cluster. The NSX Controller is an advanced distributed state management system that provides control plane functions for logical switching and routing functions. It is the central control point for all logical switches within a network and maintains information about all hosts, logical switches (VXLANs), and distributed logical routers.

The NSX Controller cluster is responsible for managing the distributed switching and routing modules in the hypervisors. The controller does not have any dataplane traffic passing through it. Controller nodes are deployed in a cluster of three members to enable high-availability and scale. Any failure of the controller nodes does not impact any data-plane traffic.

NSX Controllers work by distributing network information to hosts. To achieve a high level of resiliency the NSX Controller is clustered for scale out and HA. NSX Controller cluster must contain three nodes. The three virtual appliances provide, maintain, and update the state of all network functioning within the NSX domain. NSX Manager is used to deploy NSX Controller nodes.

The three NSX Controller nodes form a control cluster. The controller cluster requires a quorum (also called a majority) in order to avoid a "split-brain scenario." In a split-brain scenario, data inconsistencies originate from the maintenance of two separate data sets that overlap. The inconsistencies can be caused by failure conditions and data synchronization issues. Having three controller nodes ensures data redundancy in case of failure of one NSX Controller node.

A controller cluster has several roles, including:

- API provider
- Persistence server
- Switch manager
- Logical manager

- Directory server

Each role has a master controller node. If a master controller node for a role fails, the cluster elects a new master for that role from the available NSX Controller nodes. The new master NSX Controller node for that role reallocates the lost portions of work among the remaining NSX Controller nodes.

NSX Data Center for vSphere supports three logical switch control plane modes: multicast, unicast and hybrid. Using a controller cluster to manage VXLAN-based logical switches eliminates the need for multicast support from the physical network infrastructure. You don't have to provision multicast group IP addresses, and you also don't need to enable PIM routing or IGMP snooping features on physical switches or routers. Thus, the unicast and hybrid modes decouple NSX from the physical network. VXLANs in unicast control-plane mode do not require the physical network to support multicast in order to handle the broadcast, unknown unicast, and multicast (BUM) traffic within a logical switch. The unicast mode replicates all the BUM traffic locally on the host and requires no physical network configuration. In the hybrid mode, some of the BUM traffic replication is offloaded to the first hop physical switch to achieve better performance. Hybrid mode requires IGMP snooping on the first-hop switch and access to an IGMP querier in each VTEP subnet.

Management Plane

The management plane is built by the NSX Manager, the centralized network management component of NSX Data Center for vSphere. It provides the single point of configuration and REST API entry-points.

The NSX Manager is installed as a virtual appliance on any ESXi host in your vCenter Server environment. NSX Manager and vCenter have a one-to-one relationship. For every instance of NSX Manager, there is one vCenter Server. This is true even in a cross-vCenter NSX environment.

In a cross-vCenter NSX environment, there is both a primary NSX Manager and one or more secondary NSX Manager appliances. The primary NSX Manager allows you to create and manage universal logical switches, universal logical (distributed) routers and universal firewall rules. Secondary NSX Managers are used to manage networking services that are local to that specific NSX Manager. There can be up to seven secondary NSX Managers associated with the primary NSX Manager in a cross-vCenter NSX environment.

Consumption Platform

The consumption of NSX Data Center for vSphere can be driven directly through the NSX Manager user interface, which is available in the vSphere Web Client. Typically end users tie network virtualization to their cloud management platform for deploying applications. NSX Data Center for vSphere provides rich integration into virtually any CMP through REST APIs. Out-of-the-box integration is also available through VMware vRealize Automation Center, vCloud Director, and OpenStack with the Neutron plug-in.

NSX Edge

You can install NSX Edge as an edge services gateway (ESG) or as a distributed logical router (DLR).

Edge Services Gateway

The ESG gives you access to all NSX Edge services such as firewall, NAT, DHCP, VPN, load balancing, and high availability. You can install multiple ESG virtual appliances in a data center. Each ESG virtual appliance can have a total of ten uplink and internal network interfaces. With a trunk, an ESG can have up to 200 subinterfaces. The internal interfaces connect to secured port groups and act as the gateway for all protected virtual machines in the port group. The subnet assigned to the internal interface can be a publicly routed IP space or a NATed/routed RFC 1918 private space. Firewall rules and other NSX Edge services are enforced on traffic between network interfaces.

Uplink interfaces of ESGs connect to uplink port groups that have access to a shared corporate network or a service that provides access layer networking. Multiple external IP addresses can be configured for load balancer, site-to-site VPN, and NAT services.

Distributed Logical Router

The DLR provides East-West distributed routing with tenant IP address space and data path isolation. Virtual machines or workloads that reside on the same host on different subnets can communicate with one another without having to traverse a traditional routing interface.

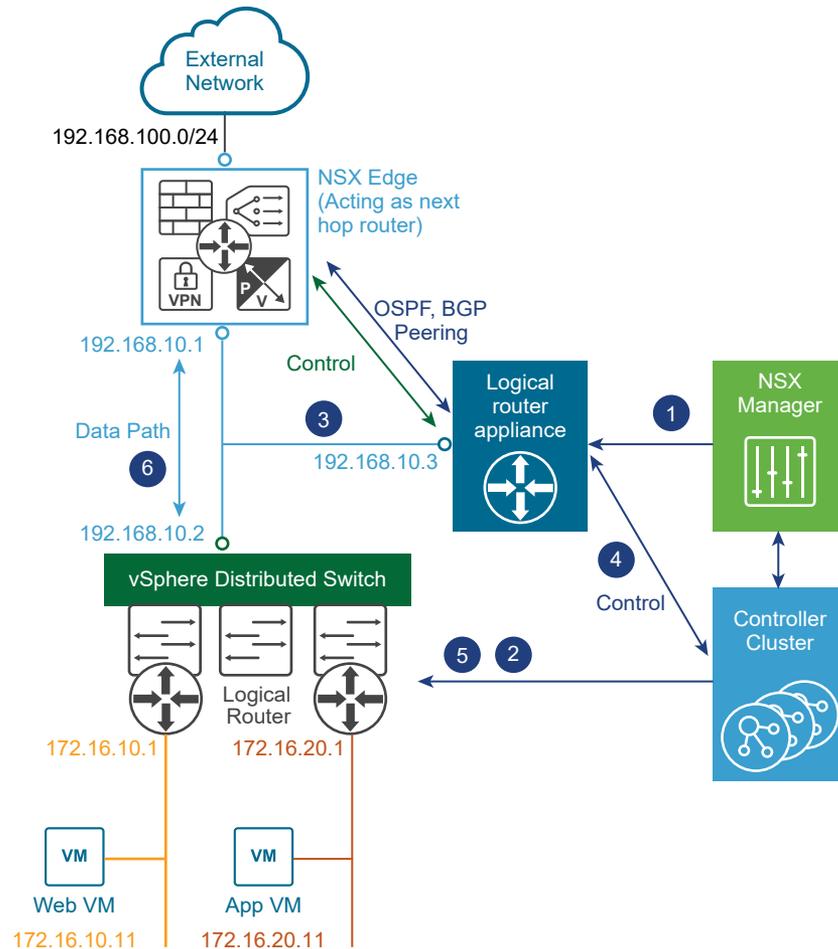
A logical router can have eight uplink interfaces and up to a thousand internal interfaces. An uplink interface on a DLR generally peers with an ESG, with an intervening Layer 2 logical transit switch between the DLR and the ESG. An internal interface on a DLR peers with a virtual machine hosted on an ESXi hypervisor with an intervening logical switch between the virtual machine and the DLR.

The DLR has two main components:

- The DLR control plane is provided by the DLR virtual appliance (also called a control VM). This VM supports dynamic routing protocols (BGP and OSPF), exchanges routing updates with the next Layer 3 hop device (usually the edge services gateway) and communicates with the NSX Manager and the NSX Controller cluster. High-availability for the DLR virtual appliance is supported through active-standby configuration: a pair of virtual machines functioning in active/standby modes are provided when you create the DLR with HA enabled.
- At the data-plane level, there are DLR kernel modules (VIBs) that are installed on the ESXi hosts that are part of the NSX domain. The kernel modules are similar to the line cards in a modular chassis supporting Layer 3 routing. The kernel modules have a routing information base (RIB) (also known as a routing table) that is pushed from the controller cluster. The data plane functions of route lookup and ARP entry lookup are performed by the kernel modules. The kernel modules are equipped with logical interfaces (called LIFs) connecting to

the different logical switches and to any VLAN-backed port-groups. Each LIF has assigned an IP address representing the default IP gateway for the logical L2 segment it connects to and a vMAC address. The IP address is unique for each LIF, whereas the same vMAC is assigned to all the defined LIFs.

Figure 3-1. Logical Routing Components



- 1 A DLR instance is created from the NSX Manager UI (or with API calls), and routing is enabled, using either OSPF or BGP.
- 2 The NSX Controller uses the control plane with the ESXi hosts to push the new DLR configuration including LIFs and their associated IP and vMAC addresses.
- 3 Assuming a routing protocol is also enabled on the next-hop device (an NSX Edge [ESG] in this example), OSPF or BGP peering is established between the ESG and the DLR control VM. The ESG and the DLR can then exchange routing information:
 - The DLR control VM can be configured to redistribute into OSPF the IP prefixes for all the connected logical networks (172.16.10.0/24 and 172.16.20.0/24 in this example). As a consequence, it then pushes those route advertisements to the NSX Edge. Notice that the

next hop for those prefixes is not the IP address assigned to the control VM (192.168.10.3) but the IP address identifying the data-plane component of the DLR (192.168.10.2). The former is called the DLR "protocol address," whereas the latter is the "forwarding address".

- The NSX Edge pushes to the control VM the prefixes to reach IP networks in the external network. In most scenarios, a single default route is likely to be sent by the NSX Edge, because it represents the single point of exit toward the physical network infrastructure.
- 4 The DLR control VM pushes the IP routes learned from the NSX Edge to the controller cluster.
 - 5 The controller cluster is responsible for distributing routes learned from the DLR control VM to the hypervisors. Each controller node in the cluster takes responsibility of distributing the information for a particular logical router instance. In a deployment where there are multiple logical router instances deployed, the load is distributed across the controller nodes. A separate logical router instance is usually associated with each deployed tenant.
 - 6 The DLR routing kernel modules on the hosts handle the data-path traffic for communication to the external network by way of the NSX Edge.

NSX Services

The NSX Data Center for vSphere components work together to provide the following functional VMware NSX[®] Services™.

Logical Switches

A cloud deployment or a virtual data center has a variety of applications across multiple tenants. These applications and tenants require isolation from each other for security, fault isolation, and non-overlapping IP addresses. NSX Data Center for vSphere allows the creation of multiple logical switches, each of which is a single logical broadcast domain. An application or tenant virtual machine can be logically wired to a logical switch. This allows for flexibility and speed of deployment while still providing all the characteristics of a physical network's broadcast domains (VLANs) without physical Layer 2 sprawl or spanning tree issues.

A logical switch is distributed and can span across all hosts in vCenter (or across all hosts in a cross-vCenter NSX environment). This allows for virtual machine mobility (vMotion) within the data center without limitations of the physical Layer 2 (VLAN) boundary. The physical infrastructure is not constrained by MAC/FIB table limits, because the logical switch contains the broadcast domain in software.

Logical Routers

Routing provides the necessary forwarding information between Layer 2 broadcast domains, thereby allowing you to decrease the size of Layer 2 broadcast domains and improve network efficiency and scale. NSX Data Center for vSphere extends this intelligence to where the workloads reside for East-West routing. This allows more direct VM-to-VM communication without the costly or timely need to extend hops. At the same time, logical routers provide North-South connectivity, thereby enabling tenants to access public networks.

Logical Firewall

Logical Firewall provides security mechanisms for dynamic virtual data centers. The Distributed Firewall component of Logical Firewall allows you to segment virtual datacenter entities like virtual machines based on VM names and attributes, user identity, vCenter objects like datacenters, and hosts, as well as traditional networking attributes like IP addresses, VLANs, and so on. The Edge Firewall component helps you meet key perimeter security requirements, such as building DMZs based on IP/VLAN constructs, and tenant-to-tenant isolation in multi-tenant virtual data centers.

The Flow Monitoring feature displays network activity between virtual machines at the application protocol level. You can use this information to audit network traffic, define and refine firewall policies, and identify threats to your network.

Logical Virtual Private Networks (VPNs)

SSL VPN-Plus allows remote users to access private corporate applications. IPsec VPN offers site-to-site connectivity between an NSX Edge instance and remote sites with NSX Data Center for vSphere or with hardware routers/VPN gateways from 3rd-party vendors. L2 VPN allows you to extend your datacenter by allowing virtual machines to retain network connectivity while retaining the same IP address across geographical boundaries.

Logical Load Balancer

The NSX Edge load balancer distributes client connections directed at a single virtual IP address (VIP) across multiple destinations configured as members of a load balancing pool. It distributes incoming service requests evenly among multiple servers in such a way that the load distribution is transparent to users. Load balancing thus helps in achieving optimal resource utilization, maximizing throughput, minimizing response time, and avoiding overload.

Service Composer

Service Composer helps you provision and assign network and security services to applications in a virtual infrastructure. You map these services to a security group, and the services are applied to the virtual machines in the security group using a Security Policy.

NSX Data Center for vSphere Extensibility

3rd-party solution providers can integrate their solutions with the NSX Data Center for vSphere platform, thus enabling customers to have an integrated experience across VMware products and partner solutions. Data center operators can provision complex, multi-tier virtual networks in seconds, independent of the underlying network topology or components.

Overview of Cross-vCenter Networking and Security

4

NSX Data Center for vSphere allows you to manage multiple environments from a single primary NSX Manager.

This chapter includes the following topics:

- Benefits of Cross-vCenter NSX
- How Cross-vCenter NSX Works
- Support Matrix for Services in Cross-vCenter NSX
- Universal NSX Controller Cluster
- Universal Transport Zone
- Universal Logical Switches
- Universal Logical (Distributed) Routers
- Universal Firewall Rules
- Universal Network and Security Objects
- Cross-vCenter NSX Topologies
- Modifying NSX Manager Roles

Benefits of Cross-vCenter NSX

NSX environments containing more than one vCenter Server system can be managed centrally.

There are many reasons multiple vCenter Server systems may be required, for example:

- To overcome scale limits of vCenter Server
- To accommodate products that require dedicated or multiple vCenter Server systems, such as Horizon View or Site Recovery Manager
- To separate environments, for example by business unit, tenant, organization, or environment type

In NSX Data Center for vSphere 6.2 and later you can create universal objects on the primary NSX Manager, which are synchronized across all vCenter Servers systems in the environment.

Cross-vCenter NSX includes these features:

- Increased span of logical networks. The same logical networks are available in the cross-vCenter environment, so it's possible for VMs on any cluster on any vCenter Server system to be connected to the same logical network.
- Centralized security policy management. Firewall rules are managed from one centralized location, and apply to the VM regardless of location or vCenter Server system.
- Support of mobility boundaries in vSphere 6, including cross vCenter and long distance vMotion across logical switches.
- Enhanced support for multi-site environments, from metro distance to 150ms RTT. This includes both active-active and active-passive datacenters.

Cross-vCenter NSX environments have many benefits:

- Centralized management of universal objects, reducing administration effort.
- Increased mobility of workloads - VMs can be migrated using vMotion across vCenter Servers without having to reconfigure the VM or change firewall rules.
- Enhanced multi-site and disaster recovery capabilities.

Note Cross-vCenter NSX functionality is supported with vSphere 6.0 and later.

How Cross-vCenter NSX Works

In a cross-vCenter NSX environment, you can have multiple vCenter Servers, each of which must be paired with its own NSX Manager. One NSX Manager is assigned the role of primary NSX Manager, and the others are assigned the role of secondary NSX Manager.

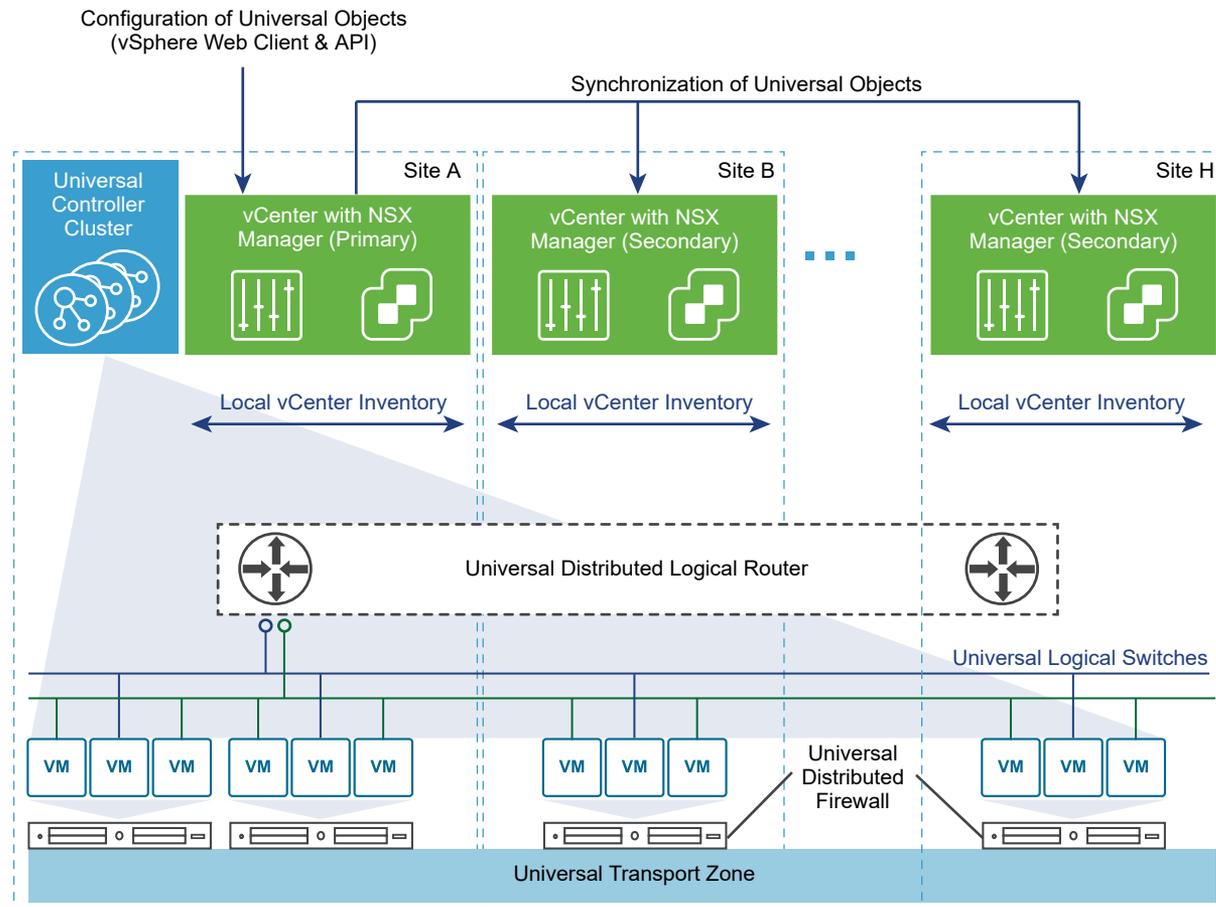
The primary NSX Manager is used to deploy a universal controller cluster that provides the control plane for the cross-vCenter NSX environment. The secondary NSX Managers do not have their own controller clusters.

The primary NSX Manager can create universal objects, such as universal logical switches. These objects are synchronized to the secondary NSX Managers by the NSX Universal Synchronization Service. You can view these objects from the secondary NSX Managers, but you cannot edit them there. You must use the primary NSX Manager to manage universal objects.

On both primary and secondary NSX Managers, you can create objects that are local to that specific environment, such as logical switches, and logical (distributed) routers. They exist only within the environment in which they were created. They are not visible on the other NSX Managers in the cross-vCenter NSX environment.

NSX Managers can be assigned the standalone role. A standalone NSX Manager manages an environment with a single NSX Manager and single vCenter. A standalone NSX Manager cannot create universal objects.

Note If you change the role of a primary NSX Manager to standalone and any universal objects exist in the NSX environment, the NSX Manager is assigned the transit role. The universal objects remain, but they cannot be changed, and no other universal objects can be created. You can delete universal objects from the transit role. Use the transit role temporarily, for example, when changing which NSX Manager is the primary.



Support Matrix for Services in Cross-vCenter NSX

A subset of NSX Data Center for vSphere services are available for universal synchronization in cross-vCenter NSX. Services that are not available for universal synchronization can be configured for use local to the NSX Manager.

Table 4-1. Support matrix for NSX Data Center for vSphere Services in cross-vCenter NSX

NSX Data Center for vSphere Service	Details	Supports cross-vCenter NSX synchronization?
Logical switch	Transport zone	Yes
	Logical switch	Yes
L2 bridges		No
Routing	Logical (distributed) router	Yes
	Logical (distributed) router appliance	No by design. Appliances must be created on each NSX Manager if multiple appliances are required per universal logical router. This allows for different configurations per appliance, which may be required in an environment with local egress configured.
	NSX Edge services gateway	No
Logical firewall	Distributed firewall	Yes
	Exclude list	No
	SpoofGuard	No
	Flow monitoring for aggregate flows	No
	Network service insertion	No
	Edge firewall	No
VPN		No
Logical load balancer		No
Other edge services		No
Service composer		No
Network extensibility		No
Network and security objects	IP address groups (IP sets)	Yes
	MAC address groups (MAC sets)	Yes
	IP pools	No
	Security groups	Yes, but membership configuration differs from non-universal security groups membership. See "Create a Security Group" in the <i>NSX Administration Guide</i> for details.
	Services	Yes

Table 4-1. Support matrix for NSX Data Center for vSphere Services in cross-vCenter NSX (continued)

NSX Data Center for vSphere Service	Details	Supports cross-vCenter NSX synchronization?
	Service groups	Yes
	Security tags	Yes
	Hardware Gateway (also known as Hardware VTEP)	No. See "Hardware Gateway Sample Configuration" in the <i>NSX Administration Guide</i> for details.

Universal NSX Controller Cluster

Each cross-vCenter NSX environment has one universal controller cluster associated with the primary NSX Manager. Secondary NSX Managers do not have a controller cluster.

As the universal controller cluster is the only controller cluster for the cross-vCenter NSX environment, it maintains information about universal logical switches and universal logical routers as well as logical switches and logical routers that are local to each NSX Manager.

In order to avoid any overlap in object IDs, separate ID pools are maintained for universal objects and local objects.

Universal Transport Zone

In a cross-vCenter NSX environment, there can be only one universal transport zone.

The universal transport zone is created on the primary NSX Manager, and is synchronized to the secondary NSX Managers. Clusters that need to participate in universal logical networks must be added to the universal transport zone from their NSX Managers.

Universal Logical Switches

Universal logical switches allow layer 2 networks to span multiple sites.

When you create a logical switch in a universal transport zone, you create a universal logical switch. This switch is available on all clusters in the universal transport zone. The universal transport zone can include clusters in any vCenter in the cross-vCenter NSX environment.

The segment ID pool is used to assign VNIs to logical switches, and the universal segment ID pool is used to assign VNIs to universal logical switches. These pools must not overlap.

You must use a universal logical router to route between universal logical switches. If you need to route between a universal logical switch and a logical switch, you must use an Edge Services Gateway.

Universal Logical (Distributed) Routers

Universal Logical (Distributed) Routers offer centralized administration and a routing configuration that can be customized at the universal logical router, cluster, or host level.

When you create a universal logical router you must choose whether to enable local egress, as this cannot be changed after creation. Local egress allows you to control what routes are provided to ESXi hosts based on an identifier, the locale ID.

Each NSX Manager is assigned a locale ID, which is set to the NSX Manager UUID by default. You can override the locale ID at the following levels:

- Universal logical router
- Cluster
- ESXi host

If you do not enable local egress the locale ID is ignored and all ESXi hosts connected to the universal logical router will receive the same routes. Whether or not to enable local egress in a cross-vCenter NSX environment is a design consideration, but it is not required for all cross-vCenter NSX configurations.

Universal Firewall Rules

Distributed Firewall in a cross-vCenter NSX environment allows centralized management of rules that apply to all vCenter Servers in your environment. It supports cross-vCenter vMotion which enables you to move workloads or virtual machines from one vCenter Server to another and seamlessly extends your software defined datacenter security.

As your datacenter needs scale out, the existing vCenter Server may not scale to the same level. This may require you to move a set of applications to newer hosts that are managed by a different vCenter Server. Or you may need to move applications from staging to production in an environment where staging servers are managed by one vCenter Server and production servers are managed by a different vCenter Server. Distributed Firewall supports these cross-vCenter vMotion scenarios by replicating firewall policies that you define for the primary NSX Manager on up to seven secondary NSX Managers.

From the primary NSX Manager you can create distributed firewall rule sections that are marked for universal synchronization. You can create more than one universal L2 rule section and more than one universal L3 rule section. Universal sections are always listed at the top of primary and secondary NSX Managers. These sections and their rules are synchronized to all secondary NSX Managers in your environment. Rules in other sections remain local to the appropriate NSX Manager.

The following Distributed Firewall features are not supported in a cross-vCenter NSX environment:

- Exclude list
- SpoofGuard

- Flow monitoring for aggregate flows
- Network service insertion
- Edge Firewall

Service Composer does not support universal synchronization, so you cannot use it to create distributed firewall rules in the universal section.

Universal Network and Security Objects

You can create custom network and security objects to use in Distributed Firewall rules in the universal section.

Universal Security Groups (USGs) can have the following:

- Universal IP Sets
- Universal MAC Sets
- Universal Security Groups
- Universal Security Tags
- Dynamic criteria

Universal network and security objects are created, deleted, and updated only on the primary NSX Manager, but are readable on the secondary NSX Manager. Universal Synchronization Service synchronizes universal objects across vCenters immediately, as well as on demand using force synchronization.

Universal security groups are used in two types of deployments: multiple live cross-vCenter NSX environments, and cross-vCenter NSX active standby deployments, where one site is live at a given time and the rest are on standby. Only active standby deployments can have universal security groups with dynamic membership based on VM name static membership based on universal security tag. Once a universal security group is created it cannot be edited to be enabled or disabled for the active standby scenario functionality. Membership is defined by included objects only, you cannot use excluded objects.

Universal security groups cannot be created from Service Composer. Security groups created from Service Composer will be local to that NSX Manager.

Cross-vCenter NSX Topologies

You can deploy cross-vCenter NSX in a single physical site, or across multiple sites.

Multi-Site and Single Site Cross-vCenter NSX

A cross-vCenter NSX environment allows you to use the same logical switches and other network objects across multiple NSX Data Center for vSphere environments. The corresponding vCenter Server systems can be located in the same site, or in different sites.

Whether the cross-vCenter NSX environment is contained within a single site or crosses multiple sites, a similar configuration can be used. These two example topologies consist of the following:

- A universal transport zone that includes all clusters in the site or sites.
- Universal logical switches attached to the universal transport zone. Two universal logical switches are used to connect VMs and one is used as a transit network for the router uplink.
- VMs added to the universal logical switches
- A universal logical router with an NSX Edge appliance to enable dynamic routing. The universal logical router appliance has internal interfaces on the VM universal logical switches and an uplink interface on the transit network universal logical switch.
- Edge Services Gateways (ESGs) connected to the transit network and the physical egress router network.

For more information about cross-vCenter NSX topologies, see the *Cross-vCenter NSX Design Guide* at <https://communities.vmware.com/docs/DOC-32552>.

Figure 4-1. Cross-vCenter NSX in a single site

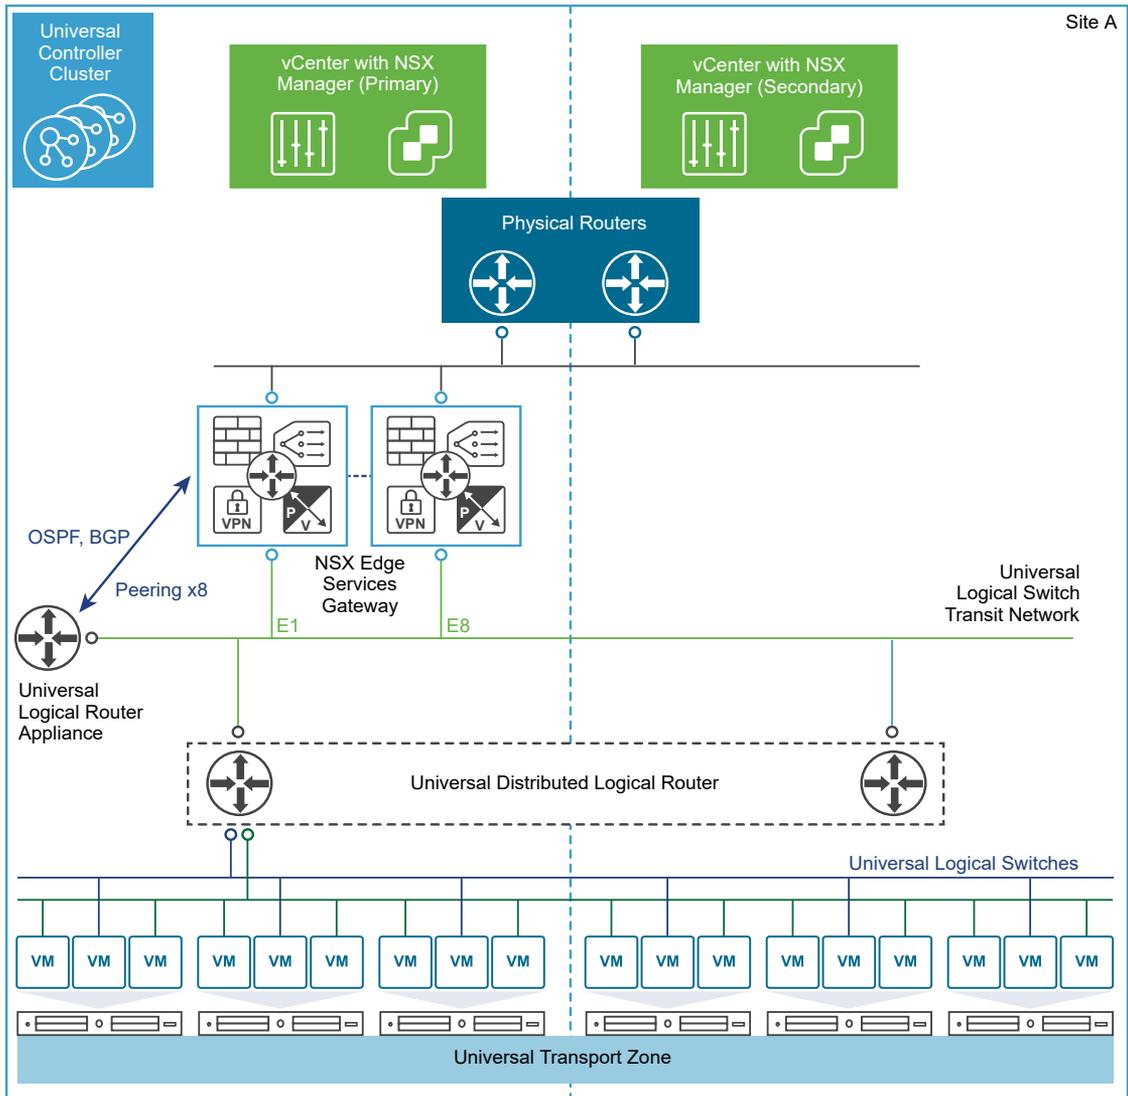
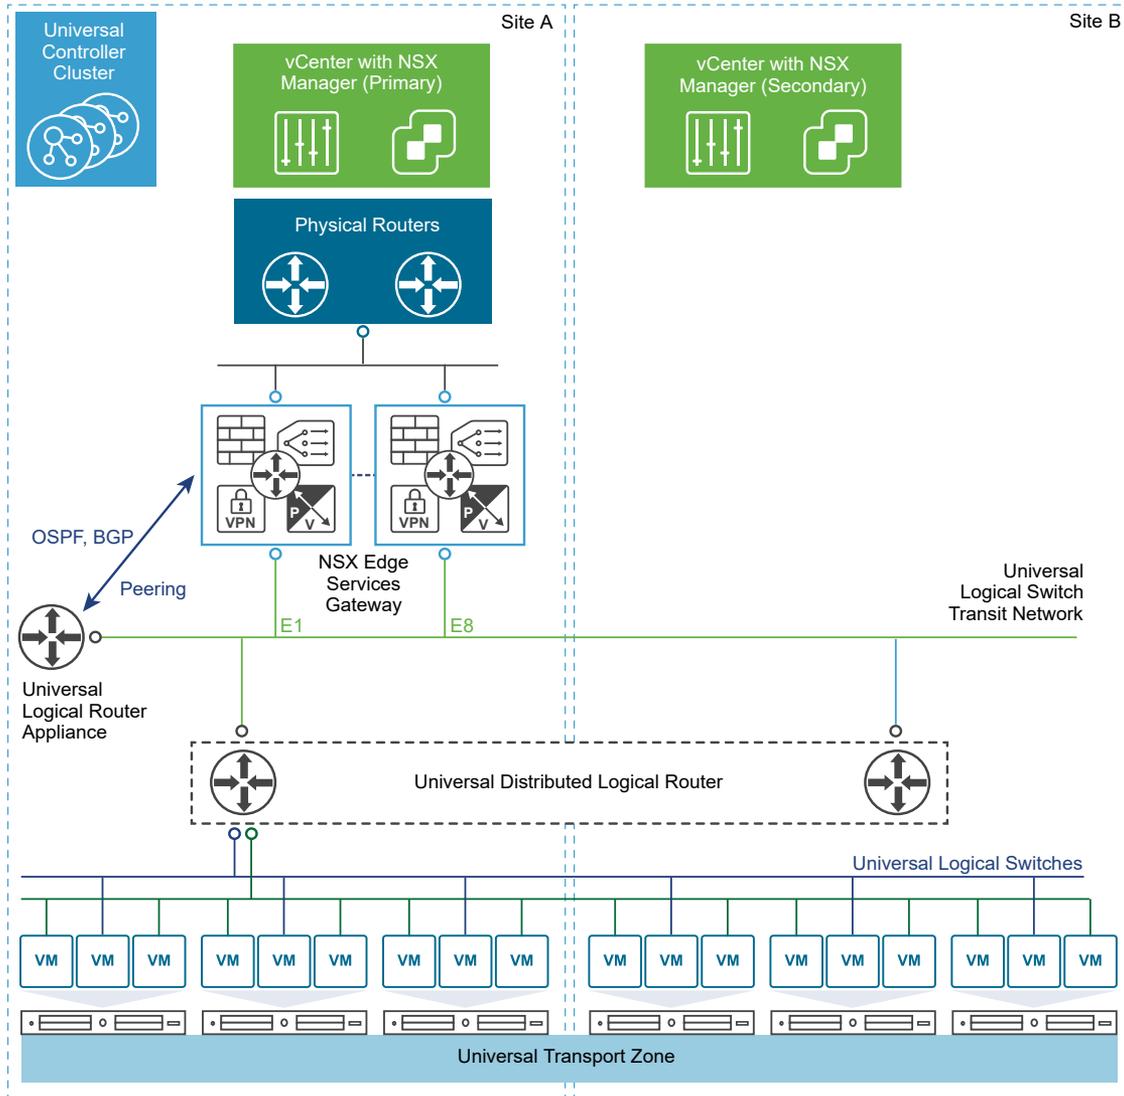


Figure 4-2. Cross-vCenter NSX spanning two sites



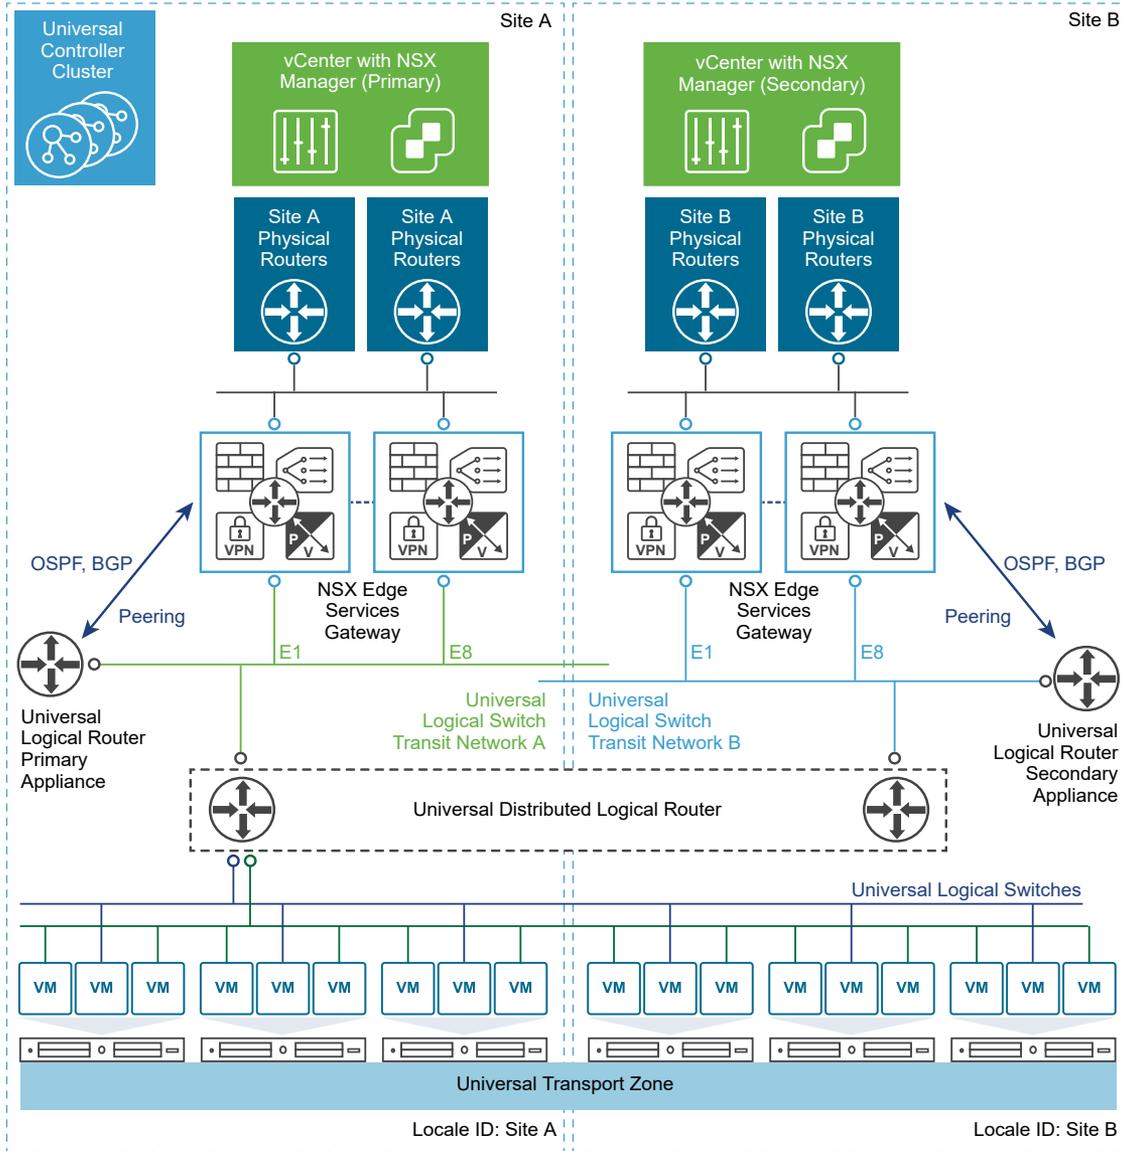
Local Egress

All sites in a multi-site cross-vCenter NSX environment can use the same physical routers for egress traffic. However, if egress routes need to be customized, the local egress feature must be enabled when the universal logical router is created.

Local egress allows you to customize routes at the universal logical router, cluster, or host level. This example of a cross-vCenter NSX environment in multiple sites has local egress enabled. The edge services gateways (ESGs) in each site have a default route that sends traffic out through that site's physical routers. The universal logical router is configured with two appliances, one in each site. The appliances learn routes from their site's ESGs. The learned routes are sent to the

universal controller cluster. Because local egress is enabled, the locale ID for that site is associated with those routes. The universal controller cluster sends routes with matching locale IDs to the hosts. Routes learned on the site A appliance are sent to the hosts in site A, and routes learned on the site B appliance are sent to the hosts in site B.

For more information about local egress, see the *Cross-vCenter NSX Design Guide* at <https://communities.vmware.com/docs/DOC-32552>.



Modifying NSX Manager Roles

An NSX Manager can have roles, such as primary, secondary, standalone, or transit. Special synchronization software runs on the primary NSX Manager, synchronizing all universal objects to secondary NSX Managers.

It is important to understand what happens when you change an NSX Manager's role.

Set as primary

This operation sets the role of an NSX Manager to primary and starts the synchronization software. This operation fails if the NSX Manager is already the primary or already a secondary.

Set as standalone (from secondary)

This operation sets the role of NSX Manager to standalone or transit mode. This operation might fail if the NSX Manager already has the standalone role.

Set as standalone (from primary)

This operation resets the primary NSX Manager to standalone or transit mode, stops the synchronization software, and unregisters all secondary NSX Managers. This operation might fail if the NSX Manager is already standalone or if any of the secondary NSX Managers are unreachable.

Disconnect from primary

When you run this operation on a secondary NSX Manager, the secondary NSX Manager is unilaterally disconnected from the primary NSX Manager. This operation should be used when the primary NSX Manager has experienced an unrecoverable failure, and you want to register the secondary NSX Manager to a new primary. If the original primary NSX Manager does come up again, its database continues to list the secondary NSX Manager as registered. To resolve this issue, include the **force** option when you disconnect or unregister the secondary from the original primary. The **force** option removes the secondary NSX Manager from the original primary NSX Manager's database.

Transport Zones

5

A transport zone controls to which hosts a logical switch can reach. It can span one or more vSphere clusters. Transport zones dictate which clusters and, therefore, which VMs can participate in the use of a particular network. In a cross-vCenter NSX environment you can create a universal transport zone, which can include clusters from any vCenter in the environment. You can create only one universal transport zone.

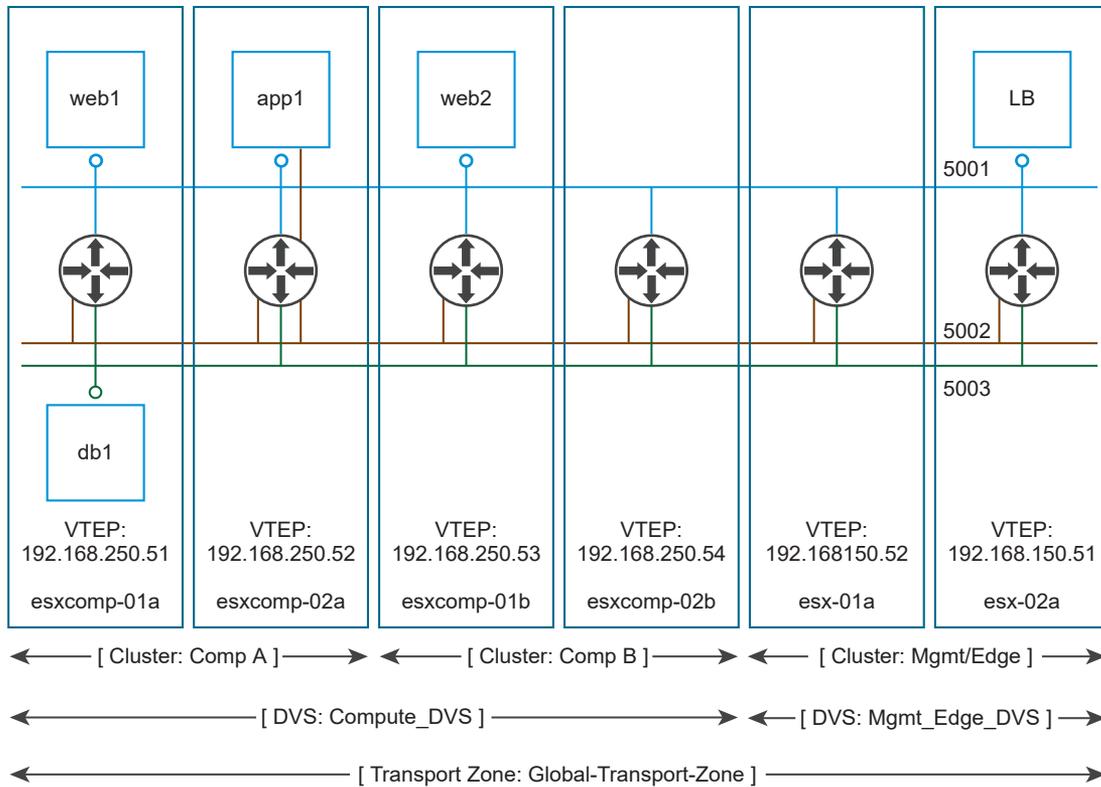
An NSX Data Center for vSphere environment can contain one or more transport zones based on your requirements. A host cluster can belong to multiple transport zones. A logical switch can belong to only one transport zone.

NSX Data Center for vSphere does not allow connection of VMs that are in different transport zones. The span of a logical switch is limited to a transport zone, so virtual machines in different transport zones cannot be on the same Layer 2 network. A distributed logical router cannot connect to logical switches that are in different transport zones. After you connect the first logical switch, the selection of further logical switches is limited to those that are in the same transport zone.

The following guidelines are meant to help you design your transport zones:

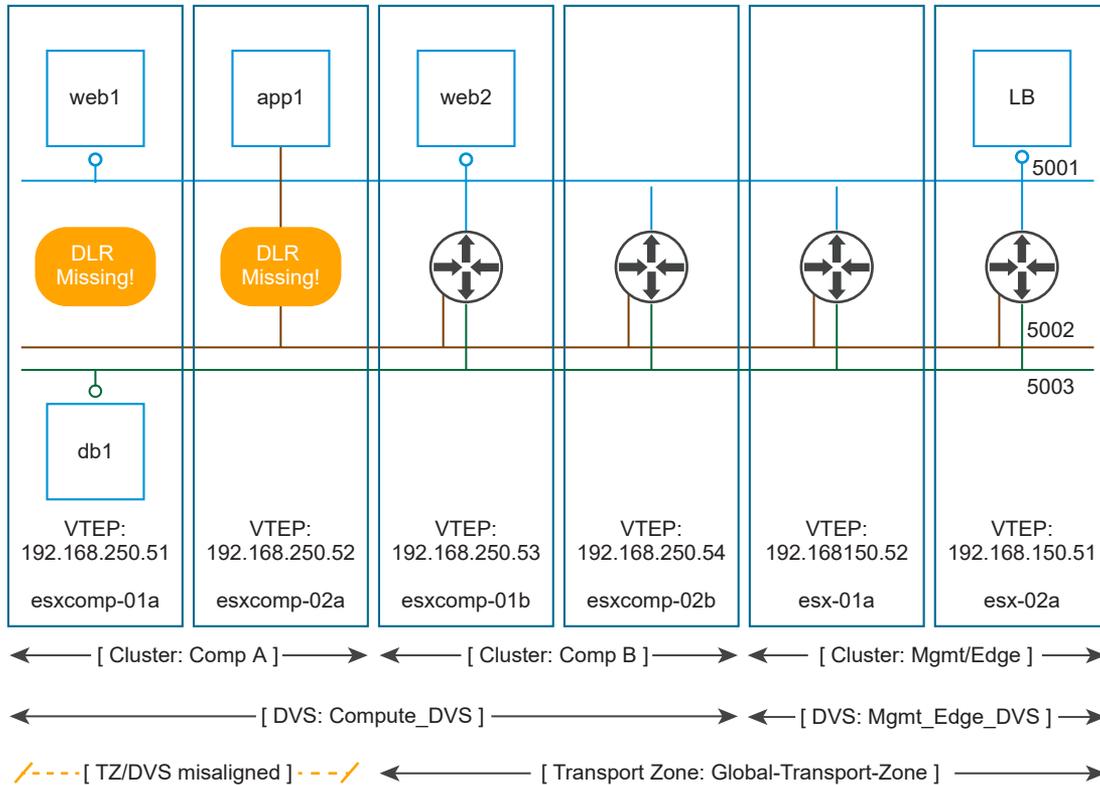
- If a cluster requires Layer 3 connectivity, the cluster must be in a transport zone that also contains an edge cluster, meaning a cluster that has Layer 3 edge devices (distributed logical routers and edge services gateways).
- Suppose you have two clusters, one for web services and another for application services. To have VXLAN connectivity between the VMs in these two clusters, both of the clusters must be included in the transport zone.
- Keep in mind that all logical switches included in the transport zone will be available and visible to all VMs within the clusters that are included in the transport zone. If a cluster includes secured environments, you might not want to make it available to VMs in other clusters. Instead, you can place your secure cluster in a more isolated transport zone.
- The span of the vSphere distributed switch (VDS or DVS) should match the transport zone span. When creating transport zones in multi-cluster VDS configurations, make sure all clusters in the selected VDS are included in the transport zone. This is to ensure that the DLR is available on all clusters where VDS dvPortgroups are available.

The following diagram shows a transport zone correctly aligned to the VDS boundary.



If you do not follow this best practice, keep in mind that if a VDS spans more than one host cluster and the transport zone includes only one (or a subset) of these clusters, any logical switch included within this transport zone can access VMs within all clusters spanned by the VDS. In other words, the transport zone will not be able to constrain the logical switch span to a subset of the clusters. If this logical switch is later connected to a DLR, you must ensure that the router instances are created only in the cluster included in the transport zone to avoid any Layer 3 issues.

For example, when a transport zone is not aligned to the VDS boundary, the scope of the logical switches (5001, 5002 and 5003) and the DLR instances that these logical switches are connected to becomes disjointed, causing VMs in cluster Comp A to have no access to the DLR logical interfaces (LIFs).



This chapter includes the following topics:

- [Understanding Replication Modes](#)
- [Add a Transport Zone](#)
- [Edit a Transport Zone](#)
- [Expand a Transport Zone](#)
- [Contract a Transport Zone](#)
- [Controller Disconnected Operation \(CDO\) Mode](#)

Understanding Replication Modes

When you create a transport zone or a logical switch, you must select a replication mode. Understanding the different modes can help you decide which is most appropriate for your environment.

Each ESXi host prepared for NSX is configured with a VXLAN tunnel endpoint (VTEP). Each VXLAN tunnel endpoint has an IP address. These IP addresses can be in the same subnet or in different subnets.

When two VMs on different ESXi hosts communicate directly, unicast-encapsulated traffic is exchanged between the two VTEP IP addresses without any need for flooding. However, as with any layer 2 network, sometimes traffic from a VM must be flooded, or sent to all other VMs belonging to the same logical switch. Layer 2 broadcast, unknown unicast, and multicast traffic are known as BUM traffic. BUM traffic from a VM on a given host must be replicated to all other hosts that have VMs connected to the same logical switch. NSX Data Center for vSphere supports three different replication modes:

- Unicast Replication Mode
- Multicast Replication Mode
- Hybrid Replication Mode

Summary of Replication Modes

Table 5-1. Summary of Replication Modes

Replication Mode	Method of BUM Replication to VTEPs on the Same Subnet	Method of BUM Replication to VTEPs on a Different Subnet	Physical Network Requirements
Unicast	Unicast	Unicast	<ul style="list-style-type: none"> ■ Routing between VTEP subnets
Multicast	Layer 2 multicast	Layer 3 multicast	<ul style="list-style-type: none"> ■ Routing between VTEP subnets ■ Layer 2 multicast, IGMP ■ Layer 3 multicast, PIM ■ Assignment of multicast groups to logical switches
Hybrid	Layer 2 multicast	Unicast	<ul style="list-style-type: none"> ■ Routing between VTEP subnets ■ Layer 2 multicast, IGMP

Unicast Replication Mode

Unicast replication mode does not require the physical network to support layer 2 or layer 3 multicast to handle the BUM traffic within a logical switch. Using unicast mode completely decouples logical networks from the physical network. Unicast mode replicates all the BUM traffic locally on the source host and forwards the BUM traffic in a unicast packet to the remote hosts. In unicast mode, you can have all VTEPs in one subnet, or in multiple subnets.

One subnet scenario: If all host VTEP interfaces belong to a single subnet, the source VTEP forwards the BUM traffic to all remote VTEPs. This is known as head-end replication. Head-end replication might result in unwanted host overhead and higher bandwidth usage. The impact depends on the amount BUM traffic and the number of hosts and VTEPs within the subnet.

Multiple subnet scenario: If the host VTEP interfaces are grouped into multiple IP subnets, the source host handles the BUM traffic in two parts. The source VTEP forwards the BUM traffic to each VTEP in the same subnet (the same as the one subnet scenario). For VTEPs in remote subnets, the source VTEP forwards the BUM traffic to one host in each remote VTEP subnet and sets the replication bit to mark this packet for local replication. When a host in the remote subnet receives this packet and finds the replication bit set, it sends the packet to all the other VTEPs in its subnet where the logical switch exists.

Therefore, unicast replication mode scales well in network architectures with many VTEP IP subnets as the load is distributed among multiple hosts.

Multicast Replication Mode

Multicast replication mode requires that both layer 3 and layer 2 multicast is enabled in the physical infrastructure. To configure multicast mode, the network administrator associates each logical switch with an IP multicast group. For ESXi hosts that are hosting VMs on a specific logical switch, the associated VTEPs join the multicast group using IGMP. The routers track the IGMP joins and create a multicast distribution tree between them using a multicast routing protocol.

When hosts replicate BUM traffic to VTEPs in the same IP subnet, they use layer 2 multicast. When hosts replicate BUM traffic to VTEPs in different IP subnets, they use layer 3 multicast. In both cases, the replication of BUM traffic to remote VTEPs is handled by the physical infrastructure.

Even though IP multicast is a well-known technology, the deployment of IP multicast in the data center is often considered a roadblock for different technical, operational, or administrative reasons. The network administrator must be careful about the maximum supported multicast states in the physical infrastructure to enable the one-to-one mapping between the logical switch and the multicast group. One of the benefits of virtualization is that it allows scaling the virtual infrastructure without exposing additional states to the physical infrastructure. Mapping logical switches to "physical" multicast groups breaks this model.

Note In multicast replication mode, the NSX Controller cluster is not used for logical switching.

Hybrid Replication Mode

Hybrid mode is a hybrid between unicast and multicast replication modes. In hybrid replication mode, host VTEPs use layer 2 multicast to distribute BUM traffic to peer VTEPs in the same subnet. When host VTEPs replicate BUM traffic to VTEPs in different subnets, they forward the traffic as unicast packets to one host per VTEP subnet. This receiving host in turn uses layer 2 multicast to send the packets to other VTEPs in its subnet.

Layer 2 multicast is more common in customer networks than Layer 3 multicast as it is typically easy to deploy. The replication to different VTEPs in the same subnet is handled in the physical network. Hybrid replication can be a significant relief for the source host for BUM traffic if there are many peer VTEPs in the same subnet. With hybrid replication, you can scale up a dense environment with little or no segmentation.

Add a Transport Zone

A transport zone controls which hosts a logical switch can reach and can span one or more vSphere clusters. Transport zones dictate which clusters and, therefore, which VMs can participate in the use of a particular network. Universal transport zones can span vSphere cluster across a cross-vCenter NSX environment.

You can have only one universal transport zone in a cross-vCenter NSX environment.

Prerequisites

Determine the appropriate NSX Manager on which to make your changes.

- In a standalone or single vCenter NSX environment there is only one NSX Manager so you do not need to select one.
- Universal objects must be managed from the primary NSX Manager.
- Objects local to an NSX Manager must be managed from that NSX Manager.
- In a cross-vCenter NSX environment that does not have Enhanced Linked Mode enabled, you must make configuration changes from the vCenter linked to the NSX Manager that you want to modify.
- In a cross-vCenter NSX environment in Enhanced Linked Mode, you can make configuration changes to any NSX Manager from any linked vCenter. Select the appropriate NSX Manager from the NSX Manager drop-down menu.

Procedure

- 1 Navigate to logical network settings.
 - ◆ In NSX 6.4.1 and later, navigate to **Networking & Security > Installation and Upgrade > Logical Network Settings**.
 - ◆ In NSX 6.4.0, navigate to **Networking & Security > Installation and Upgrade > Logical Network Preparation**.
- 2 Click **Transport Zones**, and then click **Add**.
- 3 (Optional) If you want to configure this transport zone as a universal transport zone, make the following selection.
 - In NSX 6.4.1 and later, click the **Universal Synchronization** button to turn the setting on.
 - In NSX 6.4.0, select **Mark this object for Universal Synchronization**.
- 4 Select the replication mode:
 - **Multicast**: Multicast IP addresses in the physical network are used for the control plane. This mode is recommended only when you are upgrading from older VXLAN deployments. Requires PIM/IGMP in the physical network.

- **Unicast:** The control plane is handled by an NSX Controller. All unicast traffic leverages optimized head-end replication. No multicast IP addresses or special network configuration is required.
- **Hybrid:** Offloads local traffic replication to the physical network (L2 multicast). This requires IGMP snooping on the first-hop switch and access to an IGMP querier in each VTEP subnet, but does not require PIM. The first-hop switch handles traffic replication for the subnet.

Important If you create a universal transport zone and select hybrid as the replication mode, you must ensure that the multicast address used does not conflict with any other multicast addresses assigned on any NSX Manager in the environment.

5 Select the clusters to add to the transport zone

Results

Transport-Zone is a transport zone local to the NSX Manager on which it was created.

Universal-Transport-Zone is a universal transport zone which is available on all NSX Managers in a cross-vCenter NSX environment.

Name	Description	Control Plane Mode	Logical Switches
Transport-Zone		Unicast	1
Universal-Transport-Zone		Unicast	4

What to do next

If you added a transport zone, you can add logical switches.

If you added a universal transport zone, you can add universal logical switches.

If you added a universal transport zone, you can select the secondary NSX Managers and add their clusters to the universal transport zone.

Edit a Transport Zone

You can edit the name, description, and the replication mode of a transport zone.

Procedure

- ◆ To edit a transport zone, complete these steps.

NSX Version	Procedure
NSX 6.4.1 and later	<ol style="list-style-type: none"> Navigate to Networking & Security > Installation and Upgrade > Logical Network Settings > Transport Zones. Select the transport zone and click Edit. Edit the name, description, or replication mode of the transport zone. <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p>Note If you change the transport zone replication mode, select Migrate existing Logical Switches to the new control plane mode to change the replication mode for existing logical switches linked to this transport zone. If you do not select this check box, only the logical switches linked to this transport zone after the edit is done will have the new replication mode.</p> </div> Click SAVE.
NSX 6.4.0	<ol style="list-style-type: none"> Navigate to Networking & Security > Installation and Upgrade > Logical Network Preparation > Transport Zones. Select the transport zone, and click Actions > All NSX User Interface Plugin Actions > Edit Settings. Edit the name, description, or replication mode of the transport zone. <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p>Note If you change the transport zone replication mode, select Migrate existing Logical Switches to the new control plane mode to change the replication mode for existing logical switches linked to this transport zone. If you do not select this check box, only the logical switches linked to this transport zone after the edit is done will have the new replication mode.</p> </div> Click OK.

Expand a Transport Zone

You can add clusters to a transport zone. All existing transport zones become available on the newly added clusters.

Prerequisites

The clusters you add to a transport zone have the network infrastructure installed and are configured for VXLAN. See the *NSX Installation Guide*.

Procedure

- Navigate to the transport zone.
 - ◆ In NSX 6.4.1 and later, navigate to **Networking & Security > Installation and Upgrade > Logical Network Settings > Transport Zones**.
 - ◆ In NSX 6.4.0, navigate to **Networking & Security > Installation and Upgrade > Logical Network Preparation > Transport Zones**.
- Click the transport zone that you want to expand.
- Click **Connect Clusters** ( or )

- 4 Select the clusters that you want to add to the transport zone and click **OK** or **Save**.

Contract a Transport Zone

You can remove clusters from a transport zone. The size of existing transport zones is reduced to accommodate the contracted scope.

Procedure

- 1 Navigate to the transport zone.
 - ◆ In NSX 6.4.1 and later, navigate to **Networking & Security > Installation and Upgrade > Logical Network Settings > Transport Zones**.
 - ◆ In NSX 6.4.0, navigate to **Networking & Security > Installation and Upgrade > Logical Network Preparation > Transport Zones**.
- 2 Click the transport zone that you want to contract.
- 3 Click **Disconnect Clusters** ( or ).
- 4 Select the clusters that you want to remove.
- 5 Click **OK** or **Save**.

Controller Disconnected Operation (CDO) Mode

Controller Disconnected Operation (CDO) mode functionality is now extended for multiple sites and is available at NSX Manager level.

For details, refer to [Controller Disconnected Mode for Multiple Sites](#).

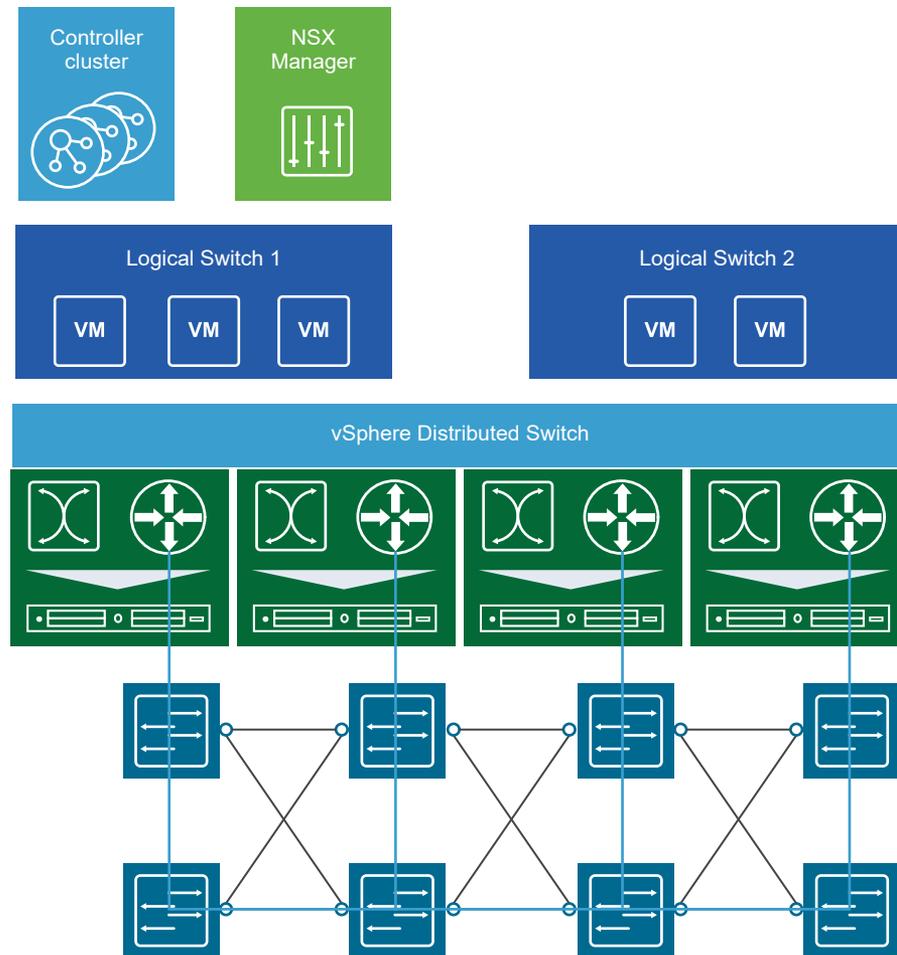
Logical Switches

6

A cloud deployment or a virtual data center has a variety of applications across multiple tenants. These applications and tenants require isolation from each other for security, fault isolation, and avoiding overlapping IP addressing issues. The NSX logical switch creates logical broadcast domains or segments to which an application or tenant virtual machine can be logically wired. This allows for flexibility and speed of deployment while still providing all the characteristics of a physical network's broadcast domains (VLANs) without physical Layer 2 sprawl or spanning tree issues.

A logical switch is distributed and can span arbitrarily large compute clusters. This allows for virtual machine mobility (vMotion) within the datacenter without limitations of the physical Layer 2 (VLAN) boundary. The physical infrastructure does not have to deal with MAC/FIB table limits since the logical switch contains the broadcast domain in software.

A logical switch is mapped to a unique VXLAN, which encapsulates the virtual machine traffic and carries it over the physical IP network.



The NSX controller is the central control point for all logical switches within a network and maintains information of all virtual machines, hosts, logical switches, and VXLANs. The controller supports two new logical switch control plane modes, Unicast and Hybrid. These modes decouple NSX from the physical network. VXLANs no longer require the physical network to support multicast in order to handle the Broadcast, Unknown unicast, and Multicast (BUM) traffic within a logical switch. The unicast mode replicates all the BUM traffic locally on the host and requires no physical network configuration. In the hybrid mode, some of the BUM traffic replication is offloaded to the first hop physical switch to achieve better performance. This mode requires IGMP snooping to be turned on the first hop physical switch. Virtual machines within a logical switch can use and send any type of traffic including IPv6 and multicast.

You can extend a logical switch to a physical device by adding an L2 bridge. See [Chapter 8 L2 Bridges](#).

You must have the Super Administrator or Enterprise Administrator role permissions to manage logical switches.

This chapter includes the following topics:

- [Add a Logical Switch](#)

- [Connect Virtual Machines to a Logical Switch](#)
- [Test Logical Switch Connectivity](#)
- [Prevent Spoofing on a Logical Switch](#)
- [Edit a Logical Switch](#)
- [Logical Switch Scenario](#)

Add a Logical Switch

Prerequisites

- You have the Super Administrator or Enterprise Administrator role permission to configure and manage logical switches.
- VXLAN UDP port is opened on firewall rules (if applicable). The VXLAN UDP port can be configured through the API.
- Physical infrastructure MTU is at least 50 bytes more than the MTU of the virtual machine vNIC.
- Managed IP address is set for each vCenter Server in the vCenter Server Runtime Settings. See *vCenter Server and Host Management*.
- DHCP is available on VXLAN transport VLANs if you are using DHCP for IP assignment for VMKNics.
- A consistent distributed virtual switch type (vendor, and so on) and version is being used across a given transport zone. Inconsistent switch types can lead to undefined behavior in your logical switch.
- You have configured an appropriate LACP teaming policy and connected physical NICs to the ports. For more information on teaming modes, refer to the VMware vSphere documentation.
- 5-tuple hash distribution is enabled for Link Aggregation Control Protocol (LACP).
- Verify that for every host where you want to use LACP, a separate LACP port channel exists on the distributed virtual switch.
- For multicast mode, multicast routing is enabled if VXLAN traffic is traversing routers. You have acquired a multicast address range from your network administrator.
- Port 1234 (the default controller listening port) is opened on firewall for the ESXi host to communicate with controllers.
- (Recommended) For multicast and hybrid modes, you have enabled IGMP snooping on the L2 switches to which VXLAN participating hosts are attached. If IGMP snooping is enabled on L2, IGMP querier must be enabled on the router or L3 switch with connectivity to multicast enabled networks.

Add a Logical Switch

An NSX logical switch reproduces switching functionality (unicast, multicast, broadcast) in a virtual environment that is decoupled from the underlying hardware. Logical switches are similar to VLANs, in that they provide network connections to which you can attach virtual machines. Logical switches are local to a single vCenter NSX deployment. In a cross-vCenter NSX deployment, you can create universal logical switches, which can span all vCenters. The transport zone type determines whether the new switch is a logical switch or a universal logical switch.

When you create a logical switch, in addition to selecting a transport zone and replication mode, you configure two options: IP discovery, and MAC learning.

IP discovery minimizes ARP traffic flooding within individual VXLAN segments---in other words, between VMs connected to the same logical switch. IP discovery is enabled by default.

Note You cannot disable IP discovery when you create a universal logical switch. You can disable IP discovery via the API after the universal logical switch is created. This setting is managed separately on each NSX Manager. See the *NSX API Guide*.

MAC learning builds a VLAN/MAC pair learning table on each vNIC. This table is stored as part of the dfilter data. During vMotion, dfilter saves and restores the table at the new location. The switch then issues RARPs for all the VLAN/MAC entries in the table. You might want to enable MAC learning if you are using virtual NICs that are trunking VLANs.

Prerequisites

Table 6-1. Prerequisites for Creating a Logical Switch or Universal Logical Switch

Logical Switch	Universal Logical Switch
<ul style="list-style-type: none"> ■ vSphere distributed switches must be configured. ■ NSX Manager must be installed. ■ Controllers must be deployed. ■ Host clusters must be prepared for NSX. ■ VXLAN must be configured. ■ A segment ID pool must be configured. ■ A transport zone must be created. 	<ul style="list-style-type: none"> ■ vSphere distributed switches must be configured. ■ NSX Manager must be installed. ■ Controllers must be deployed. ■ Host clusters must be prepared for NSX. ■ VXLAN must be configured. ■ A primary NSX Manager must be assigned. ■ A universal segment ID pool must be configured. ■ A universal transport zone must be created.

Determine the appropriate NSX Manager on which to make your changes.

- In a standalone or single vCenter NSX environment there is only one NSX Manager so you do not need to select one.
- Universal objects must be managed from the primary NSX Manager.
- Objects local to an NSX Manager must be managed from that NSX Manager.
- In a cross-vCenter NSX environment that does not have Enhanced Linked Mode enabled, you must make configuration changes from the vCenter linked to the NSX Manager that you want to modify.

- In a cross-vCenter NSX environment in Enhanced Linked Mode, you can make configuration changes to any NSX Manager from any linked vCenter. Select the appropriate NSX Manager from the NSX Manager drop-down menu.

Procedure

- 1 Navigate to **Home > Networking & Security > Logical Switches**.
- 2 Select the NSX Manager on which you want to create a logical switch. To create a universal logical switch, you must select the primary NSX Manager.
- 3 Click **Add** or the **New Logical Switch (+)** icon.
- 4 Type a name and optional description for the logical switch.
- 5 Select the transport zone in which you want to create the logical switch. If you select a universal transport zone, a universal logical switch is created.

By default, the logical switch inherits the control plane replication mode from the transport zone. You can change it to one of the other available modes. The available modes are unicast, hybrid, and multicast.

If you create a universal logical switch and select hybrid as the replication mode, you must ensure that the multicast address used does not conflict with other multicast addresses assigned on any NSX Manager in the cross-vCenter NSX environment.

- 6 (Optional) Enable IP Discovery to enable ARP suppression.
- 7 (Optional) Enable MAC Learning.

Example: Logical Switch and Universal Logical Switch

`App` is a logical switch connected to a transport zone. It is available only on the NSX Manager on which it was created.

`Universal-App` is a universal logical switch connected to a universal transport zone. It is available on any of the NSX Managers in the cross-vCenter NSX environment.

The logical switch and the universal logical switch have segment IDs from different segment ID pools.

Virtual Wire ID	Segment ID	Name	Status	Transport Zone
virtualwire-1	5000	App	Normal	Transport-Zone
universalwire-2	900000	Universal-App	Normal	Universal-Transport-Zone

What to do next

Add VMs to a logical switch or universal logical switch.

Create a logical router and attach it to your logical switches to enable connectivity between VMs that are connected to different logical switches.

Create a universal logical router and attach it to your universal logical switches to enable connectivity between VMs that are connected to different universal logical switches.

Connect a Logical Switch to an NSX Edge

Connecting a logical switch to an NSX Edge services gateway or an NSX Edge logical router provides East-West traffic routing (among the logical switches) or North-South traffic routing to the external world or to provide advanced services.

Procedure

Procedure

- 1 In Logical Switches, select the logical switch to which you want to connect an NSX Edge.
- 2 Click **Actions > Connect Edge**.
- 3 Select the NSX Edge to which you want to connect the logical switch.
- 4 Select the interface that you want to connect to the logical switch.

A logical network is typically connected to an internal interface.

- 5 Specify the interface details of the NSX Edge.
 - a Enter a name for the NSX Edge interface.
 - b To indicate whether this interface is an internal or an external (uplink) interface, click **Internal** or **Uplink**.
 - c Select the connectivity status of the interface.
 - d In **Configure Subnets**, click **Add** to add a subnet for the interface.

An interface can have multiple non-overlapping subnets. Enter one primary IP address and a comma-separated list of multiple secondary IP addresses. NSX Edge considers the primary IP address as the source address for locally generated traffic. You must add an IP address to an interface before using it on any feature configuration.

If the NSX Edge to which you are connecting the logical switch has **Manual HA Configuration** selected, specify two management IP addresses in CIDR format.

- e Enter the subnet prefix length or subnet mask for the interface.
- f If you are using NSX 6.4.4 or later, click the **Advanced** tab, and then continue with the remaining steps in this procedure. If you are using NSX 6.4.3 or earlier, directly go to the next step.

- g Change the default MTU, if necessary.
- h Under **Options**, specify the following options.

Option	Description
Proxy ARP	Supports overlapping network forwarding between different interfaces.
Send ICMP Redirect	Conveys routing information to hosts.
Reverse Path Filter	Verifies the reachability of the source address in packets being forwarded. In enabled mode, the packet must be received on the interface that the router might use to forward the return packet. In loose mode, the source address must appear in the routing table.

- 6 Enter the fence parameters.

Configure fence parameters when you want to reuse IP and MAC addresses across different fenced environments. For example, in a cloud management platform (CMP), fencing allows you to run several cloud instances simultaneously with the same IP and MAC addresses isolated or "fenced".

- 7 Click **Add** or **Finish**.

Deploy Services on a Logical Switch

You can deploy third party services on a logical switch.

Prerequisites

One or more third party virtual appliances must have been installed in your infrastructure.

Procedure

- 1 In **Logical Switches**, select the logical switch on which you want to deploy services.
- 2 Click the **Add Service Profile** () icon.
- 3 Select the service and service profile that you want to apply.
- 4 Click **OK**.

Connect Virtual Machines to a Logical Switch

You can connect virtual machines to a logical switch or a universal logical switch.

Procedure

- 1 In **Logical Switches**, select the logical switch to which you want to add virtual machines.
- 2 Click the **Add Virtual Machine** ( or ) icon.
- 3 Select one or more virtual machines you want to add to the logical switch.
- 4 Select a vNIC for each VM that you connected to the logical switch.

- 5 Review the VMs and vNICs that you selected, and then click **Finish**.

Test Logical Switch Connectivity

A ping test checks whether two hosts in a VXLAN transport network can reach each other.

- 1 In **Logical Switches**, double-click the logical switch that you want to test for connectivity.
- 2 Click the **Monitor** tab.
- 3 Click the **Ping** tab or the **Hosts-Ping** tab.
- 4 Select the source host machine.
- 5 Select the size of the test packet.
 - **VXLAN Standard**: The standard size is 1550 bytes. This packet size matches the physical infrastructure MTU without fragmentation. This packet size allows NSX to check connectivity and verify that the infrastructure is prepared for VXLAN traffic.
 - **Minimum**: This packet size allows fragmentation. Hence, with the packet size minimized, NSX can check connectivity, but not whether the infrastructure is ready for a larger frame size.
- 6 Select the destination host machine.
- 7 Click **Start Test**.

The host-to-host ping test results are displayed.

Prevent Spoofing on a Logical Switch

After synchronizing with the vCenter Server, NSX Manager collects the IP addresses of all vCenter guest virtual machines from VMware Tools on each virtual machine, or from IP discovery if it is enabled. NSX does not trust all IP addresses provided by VMware Tools or IP discovery. If a virtual machine has been compromised, the IP address can be spoofed and malicious transmissions can bypass firewall policies.

SpoofGuard allows you to authorize the IP addresses reported by VMware Tools or IP discovery, and alter them if necessary to prevent spoofing. SpoofGuard inherently trusts the MAC addresses of virtual machines collected from the VMX files and vSphere SDK. Operating separately from the Firewall rules, you can use SpoofGuard to block traffic identified as spoofed.

For more information, see [Chapter 14 Using SpoofGuard](#).

Edit a Logical Switch

You can edit the name, description, and control plane mode of a logical switch.

Procedure

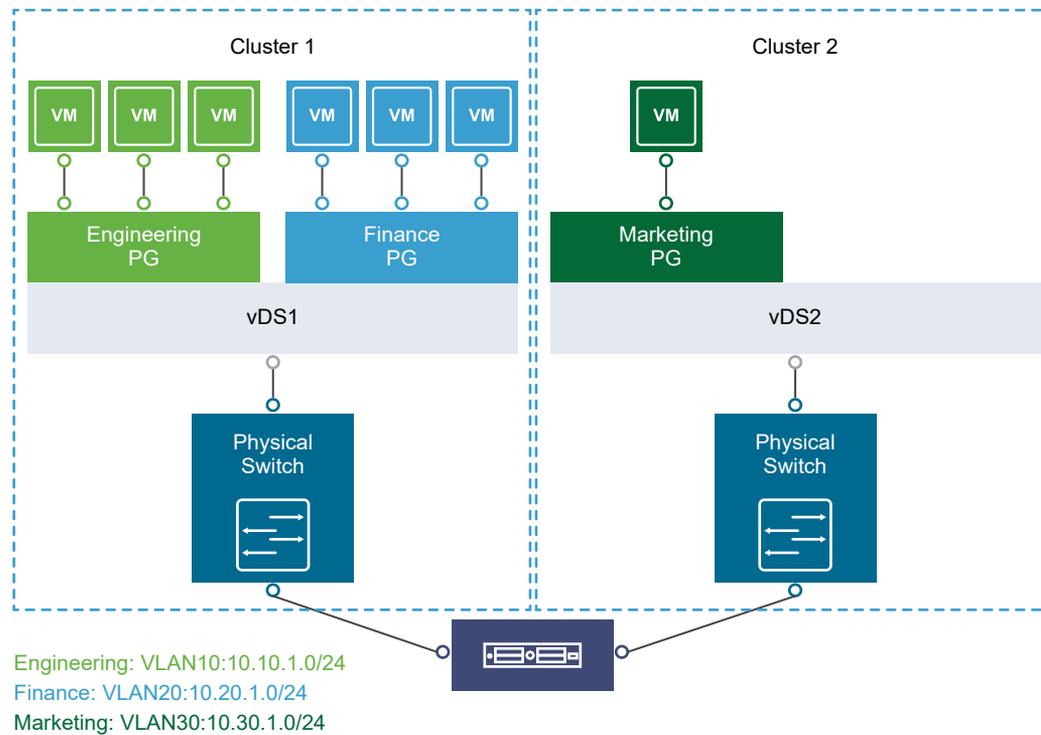
- 1 In **Logical Switches**, select the logical switch that you want to edit.

- 2 Click the **Edit** icon.
- 3 Make the desired changes.
- 4 Click **Save** or **OK**.

Logical Switch Scenario

This scenario presents a situation where company ACME Enterprise has several ESXi hosts on two clusters in a datacenter, ACME_Datacenter. The Engineering (on port group PG-Engineering) and Finance departments (on port group PG-Finance) are on Cluster1. The Marketing department (PG-Marketing) is on Cluster2. Both clusters are managed by a single vCenter Server.

Figure 6-1. ACME Enterprise network before implementing logical switches

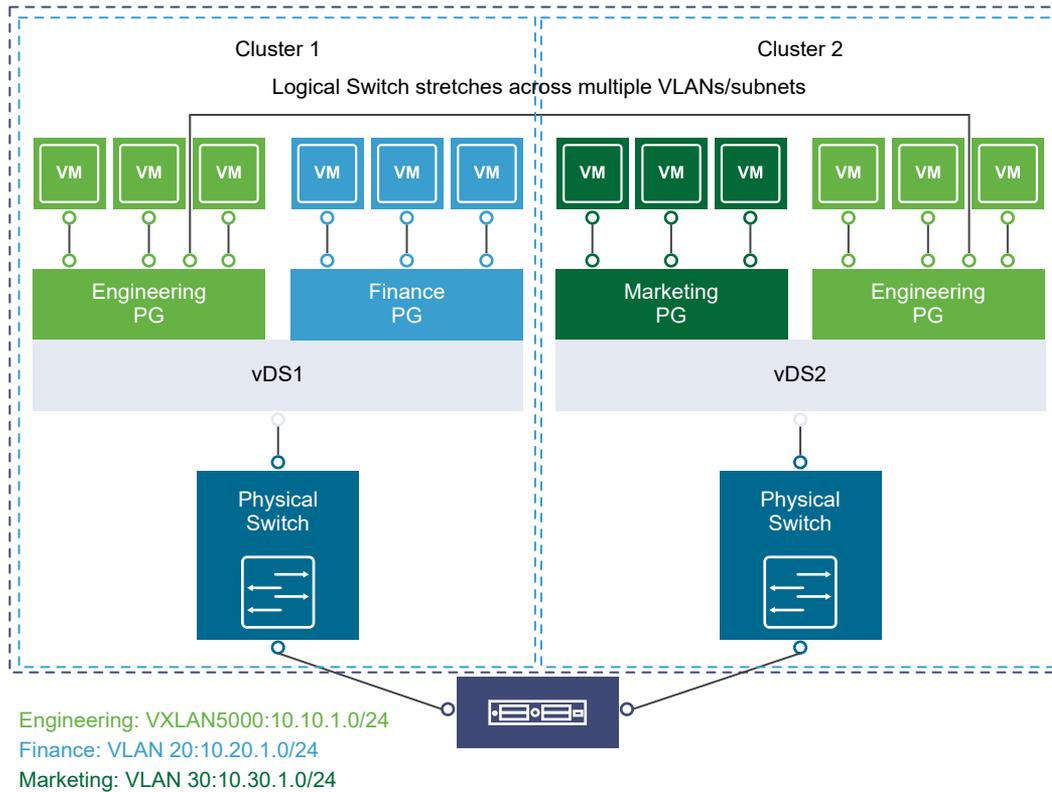


ACME is running out of compute space on Cluster1 while Cluster2 is under-utilized. The ACME network supervisor asks John Admin (ACME's virtualization administrator) to figure out a way to extend the Engineering department to Cluster2 in a way that virtual machines belonging to Engineering on both clusters can communicate with each other. This would enable ACME to utilize the compute capacity of both clusters by stretching ACME's L2 layer.

If John Admin were to do this the traditional way, he would need to connect the separate VLANs in a special way so that the two clusters can be in the same L2 domain. This might require ACME to buy a new physical device to separate traffic, and lead to issues such as VLAN sprawl, network loops, and administration and management overhead.

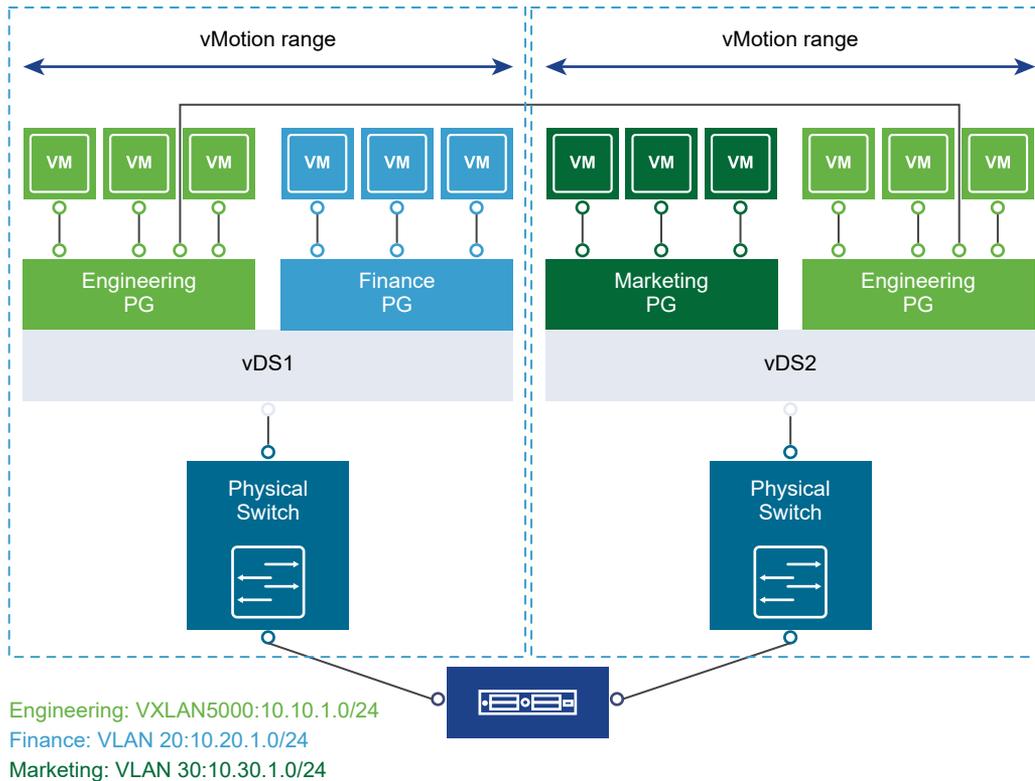
John Admin remembers seeing a logical network demo at VMworld, and decides to evaluate NSX. He concludes that building a logical switch across dvSwitch1 and dvSwitch2 will allow him to stretch ACME's L2 layer. Since John can leverage the NSX controller, he will not have to touch ACME's physical infrastructure as NSX works on top of existing IP networks.

Figure 6-2. ACME Enterprise implements a logical switch



Once John Admin builds a logical switch across the two clusters, he can vMotion virtual machines from one cluster to another while keeping them attached to the same logical switch.

Figure 6-3. vMotion on a logical network



Let us walk through the steps that John Admin follows to build a logical network at ACME Enterprise.

John Admin Assigns Segment ID Pool to NSX Manager

John Admin must specify the segment ID pool he received to isolate Company ABC's network traffic.

Prerequisites

- 1 John Admin verifies that dvSwitch1 and dvSwitch2 are vSphere Distributed Switches.
- 2 John Admin sets the Managed IP address for the vCenter Server.
 - a Select **Administration > vCenter Server Settings > Runtime Settings**.
 - b In vCenter Server Managed IP, type **10.115.198.165**.
 - c Click **OK**.
- 3 John Admin installs the network virtualization components on Cluster1 and Cluster 2. See *NSX Installation Guide*.
- 4 John Admin gets a segment ID pool (5000 - 5250) from ACME's NSX Manager administrator. Since he is leveraging the NSX controller, he does not require multicast in his physical network.

- 5 John Admin creates an IP pool so that he can assign a static IP address to the VXLAN VTEPs from this IP pool. See [Add an IP Pool](#).

Procedure

- 1 In the vSphere Web Client, click **Networking & Security > Installation and Upgrade**.
- 2 Click the **Logical Network Preparation** tab and then click **Segment ID**.
- 3 Click **Edit**.
- 4 In Segment ID pool, type **5000 - 5250**.
- 5 Do not select **Enable multicast addressing**.
- 6 Click **OK**.

John Admin Configures VXLAN Transport Parameters

John Admin configures VXLAN on Cluster 1 and Cluster 2, where he maps each cluster to a vSphere Distributed Switch. When he maps a cluster to a switch, each host in that cluster is enabled for logical switches.

Procedure

- 1 Click the **Host Preparation** tab.
- 2 For Cluster1, select **Configure** in the VXLAN column.
- 3 In the Configuring VXLAN networking dialog box, select dvSwitch1 as the vSphere Distributed Switch for the cluster.
- 4 Type **10** for dvSwitch1 to use as the ACME transport VLAN.
- 5 In Specify Transport Attributes, leave 1600 as the Maximum Transmission Units (MTU) for dvSwitch1.

MTU is the maximum amount of data that can be transmitted in one packet before it is divided into smaller packets. John Admin knows that VXLAN logical switch traffic frames are slightly larger in size because of the encapsulation, so the MTU for each switch must be set to 1550 or higher.

- 6 In **VMKNic IP Addressing**, select **Use IP Pool** and select an IP pool.
- 7 For **VMKNic Teaming Policy**, select **Failover**.

John Admin wants to maintain the quality of service in his network by keeping the performance of logical switches the same in normal and fault conditions. Hence, he chooses Failover as the teaming policy.

- 8 Click **Add**.
- 9 Repeat steps 4 through step 8 to configure VXLAN on Cluster2.

Results

After John Admin maps Cluster1 and Cluster2 to the appropriate switch, the hosts on those clusters are prepared for logical switches:

- 1 A VXLAN kernel module and vmknic is added to each host in Cluster 1 and Cluster 2.
- 2 A special dvPortGroup is created on the vSwitch associated with the logical switch and the VMKNic is connected to it.

John Admin Adds a Transport Zone

The physical network backing a logical network is called a transport zone. A transport zone is the compute diameter spanned by a virtualized network.

Procedure

- 1 Click **Logical Network Preparation** and then click **Transport Zones**.
- 2 Click the **New Transport Zone** icon.
- 3 In Name, type **ACME Zone**.
- 4 In Description, type **Zone containing ACME's clusters**.
- 5 Select Cluster 1 and Cluster 2 to add to the transport zone.
- 6 In **Control Plane Mode**, select **Unicast**.
- 7 Click **OK**.

John Admin Creates a Logical Switch

After John Admin configures VXLAN transport parameters, he is ready to create a logical switch.

Procedure

- 1 Click **Logical Switches** and then click the **New Logical Network** icon.
- 2 In Name, type **ACME logical network**.
- 3 In Description, type **Logical Network for extending ACME Engineering network to Cluster2**.
- 4 In **Transport Zone**, select ACME Zone.
- 5 Click **OK**.

NSX creates a logical switch providing L2 connectivity between dvSwitch1 and dvSwitch2.

What to do next

John Admin can now connect ACME's production virtual machines to the logical switch, and connect the logical switch to an NSX Edge services gateway or Logical Router.

Configuring Hardware Gateway

7

Hardware gateway configuration maps physical networks to virtual networks. The mapping configuration allows NSX to leverage the Open vSwitch Database (OVSDB).

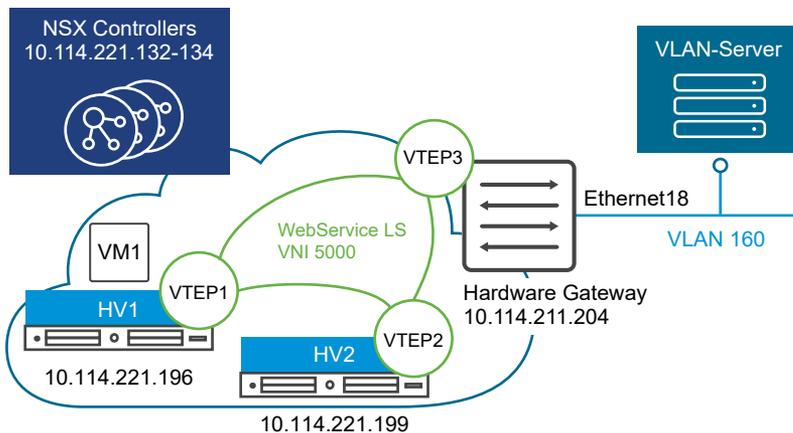
The OVSDB database contains information about the physical hardware and the virtual network. The vendor hardware hosts the database server.

The hardware gateway switches in the NSX logical networks terminate VXLAN tunnels. To the virtual network, the hardware gateway switches are known as hardware VTEP. For more information about VTEPs, see the *NSX Installation* guide and *NSX Network Virtualization Design* guide.

A minimal topology with a hardware gateway includes the following components:

- Physical server
- Hardware gateway switch (L2 port)
- IP network
- Hypervisors a minimum of four, including two replication clusters with VMs
- Controller cluster with at least three nodes

The sample topology with a hardware gateway shows HV1 and HV2 as the two hypervisors. The VM1 virtual machine is on HV1. VTEP1 is on HV1, VTEP2 is on HV2, and VTEP3 is on the hardware gateway. The hardware gateway is located in a different subnet 211 compared to the two hypervisors that are located in the same subnet 221.



The hardware gateway underlying configuration can have any one of the following components:

- Single switch
- Multiple physical bus switches with different IP addresses
- Hardware switch controller with multiple switches

The NSX Controller communicates with the hardware gateway using its IP address on port 6640. This connection is used to send and receive OVSDB transactions from hardware gateways.

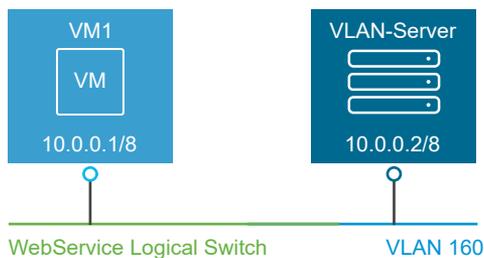
This chapter includes the following topics:

- [Scenario: Hardware Gateway Sample Configuration](#)

Scenario: Hardware Gateway Sample Configuration

This scenario describes typical tasks used to configure a hardware gateway switch with an NSX deployment. The sequence of tasks show how to connect the virtual machine VM1 to the physical server and to connect the WebService logical switch to the VLAN-Server VLAN 160 using the hardware gateway.

The sample topology shows that virtual machine VM1 and VLAN-Server are configured with an IP address in the subnet 10. VM1 is attached to WebService logical switch. The VLAN-Server is attached to VLAN 160 on the physical server.



Important In a cross-vCenter NSX environment, hardware gateway switch configurations are supported on the primary and secondary NSX Manager and multiple replication clusters. However, only hardware gateways on a primary NSX manager can use the default replication cluster. On a secondary NSX manager, new replication clusters for hardware gateways should be created. Hardware gateway switches must be bound to non-universal logical switches.

Prerequisites

- Read the vendor documentation to meet the physical network requirements.
- Verify that you meet the NSX system and hardware requirements for hardware gateway configuration. See [Chapter 1 System Requirements for NSX Data Center for vSphere](#).
- Verify that the logical networks are set up properly. See the *NSX Installation* guide.
- Verify that the transport parameter mappings in the VXLAN are accurate. See the *NSX Installation* guide.

- Retrieve the vendor certificate for your hardware gateway.
- Verify that the VXLAN port value is set to 4789. See [Change VXLAN Port](#).

Procedure

1 Set Up the Replication Cluster

A replication cluster is a set of hypervisors responsible for forwarding traffic sent from the hardware gateway. The traffic can be broadcast, unknown-unicast, and multicast traffic.

2 Connect the Hardware Gateway to the NSX Controllers

You must configure the the OVSDB manager table on the physical switch to connect the hardware gateway to the NSX Controller.

3 Add Hardware Gateway Certificate

Hardware gateway certificate must be added to the hardware device for the configuration to work.

4 Bind the Logical Switch to the Physical Switch

The WebService logical switch attached to the virtual machine VM1 must communicate with the hardware gateway on the same subnet.

Set Up the Replication Cluster

A replication cluster is a set of hypervisors responsible for forwarding traffic sent from the hardware gateway. The traffic can be broadcast, unknown-unicast, and multicast traffic.

Note Hypervisors including the replication nodes and the hardware gateway switches must not be on the same IP subnet. This restriction is due to the limitation of the chipset used in most hardware gateways. Most hardware gateways, if not all, use the Broadcom Trident II chipset, which has a limitation that a layer 3 underlay network is required between the hardware gateway and the hypervisors.

In a cross-vCenter NSX environment, hardware gateway switch configurations are supported on the primary and secondary NSX Manager and multiple replication clusters. However, only hardware gateways on a primary NSX Manager can use the default replication cluster. On a secondary NSX Manager, new replication clusters for hardware gateways must be created.

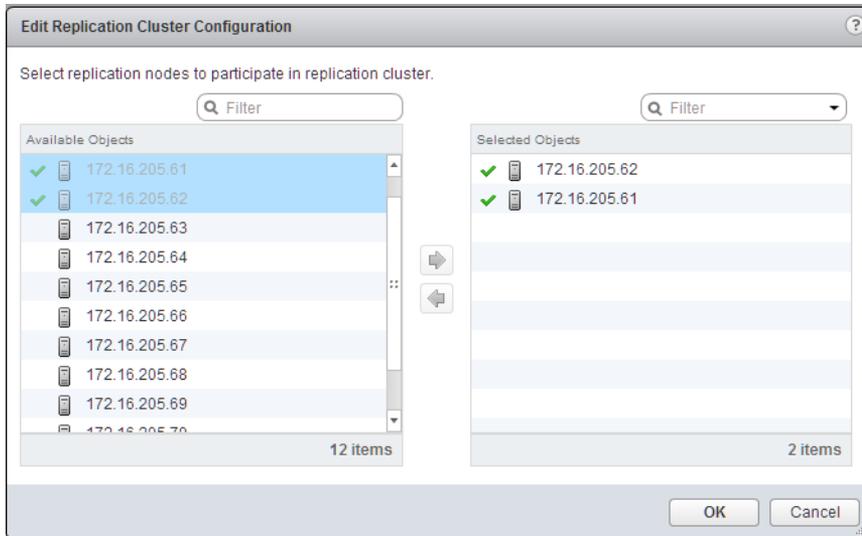
Important Through the NSX user interface, you can view and manage a single default replication cluster, but not multiple replication clusters. Support for multiple replication clusters is available through the API. See *Working With a Specific Hardware Gateway Replication Cluster* in the *NSX API Guide*.

Prerequisites

Verify that you have hypervisors to serve as replication nodes available.

Procedure

- 1 Log in to the vSphere Web Client.
- 2 Select **Networking & Security > Service Definitions**.
- 3 Click the **Hardware Devices** tab.
- 4 Click **Edit** in the Replication Cluster section to select hypervisors to serve as replication nodes in this replication cluster.
- 5 Select hypervisors and click the blue arrow.



The selected hypervisors move to the selected objects column.

- 6 Click **OK**.

Results

The replication nodes are added to the replication cluster. At least one host must exist in the replication cluster.

Connect the Hardware Gateway to the NSX Controllers

You must configure the the OVSDb manager table on the physical switch to connect the hardware gateway to the NSX Controller.

The Controller passively listens to the connection attempt from the physical switch. Therefore, the hardware gateway must use the OVSDb manager table to initiate connection.

Prerequisites

Controllers must be deployed before any hardware gateway instances are configured. If controllers are not deployed first, the error message "Failed to do the Operation on the Controller" is shown.

Procedure

- 1 Use the commands that apply to your environment to connect the hardware gateway to the NSX Controller.

Sample commands to connect hardware gateway and NSX Controller.

```
prmh-nsx-tor-7050sx-3#enable
prmh-nsx-tor-7050sx-3#configure terminal
prmh-nsx-tor-7050sx-3(config)#cvx
prmh-nsx-tor-7050sx-3(config-cvx)#service hsc
prmh-nsx-tor-7050sx-3(config-cvx-hsc)#manager 172.16.2.95 6640
prmh-nsx-tor-7050sx-3(config-cvx-hsc)#no shutdown
prmh-nsx-tor-7050sx-3(config-cvx-hsc)#end
```

- 2 Set the OVSDB manager table on the hardware gateway.
- 3 Set the OVSDB port number value as 6640.
- 4 (Optional) Verify that the hardware gateway is connected to the NSX Controller through the OVSDB channel.
 - Check that the connection status is UP.
 - Ping the VM1 and VLAN 160 to verify that the connection succeeds.
- 5 (Optional) Verify that the hardware gateway is connected to correct NSX Controller.
 - a Log in to the vSphere Web Client.
 - b Select **Networking & Security > Installation and Upgrade > Management > NSX Controller nodes**.

Add Hardware Gateway Certificate

Hardware gateway certificate must be added to the hardware device for the configuration to work.

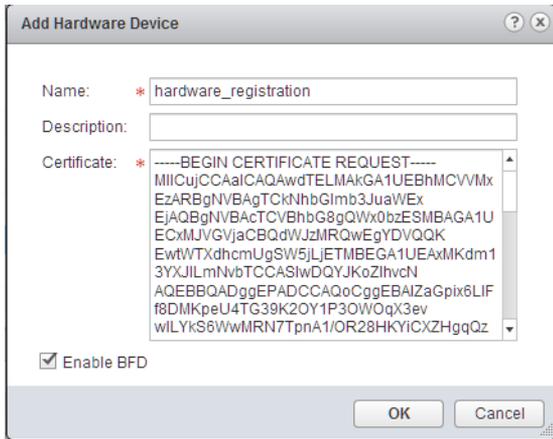
Prerequisites

Verify that the hardware gateway certificate from your environment is available.

Procedure

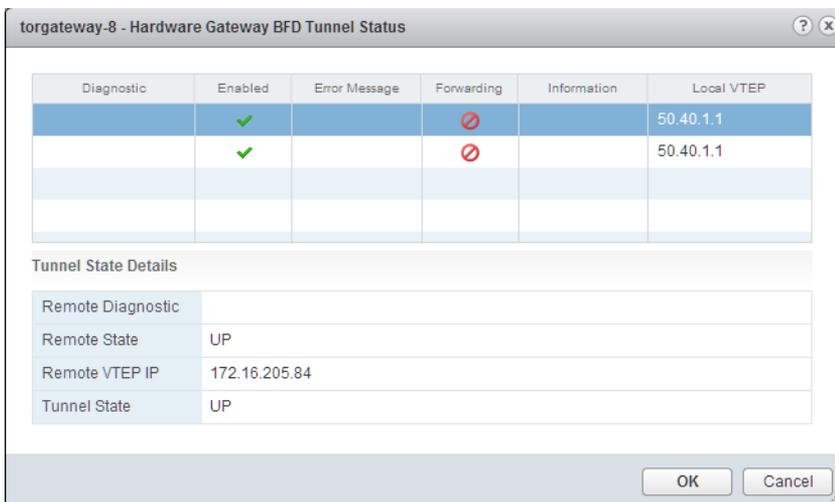
- 1 Log in to the vSphere Web Client.
- 2 Select **Networking & Security > Service Definitions**.
- 3 Click the **Hardware Devices** tab.

- Click the Add (+) icon to create the hardware gateway profile details.



Option	Description
Name and Description	Specify a hardware gateway name. You can add details of the profile in the description section.
Certificate	Paste the certificate that you extracted from your environment.
Enable BFD	Bidirectional Forwarding Detection (BFD) protocol is enabled by default . The protocol is used to synchronize the hardware gateway configuration information.

- Click **OK**.
A profile that represents the hardware gateway is created.
- Refresh the screen to verify that the hardware gateway is available and running.
The connectivity should be UP.
- (Optional) Click the hardware gateway profile and right-click to select **View the BFD Tunnel Status** from the drop-down menu.



The dialog box shows diagnostic tunnel status details for troubleshooting.

Bind the Logical Switch to the Physical Switch

The WebService logical switch attached to the virtual machine VM1 must communicate with the hardware gateway on the same subnet.

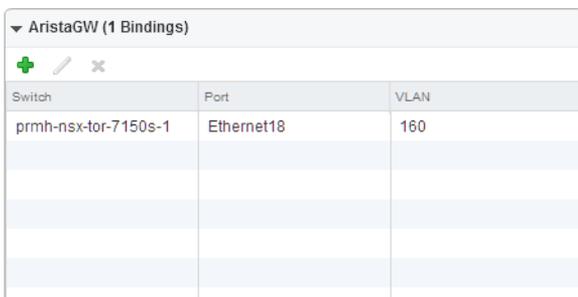
Note If you bind multiple logical switches to hardware ports, you must apply these steps for each logical switch.

Prerequisites

- Verify that the WebService logical switch is available. See [Add a Logical Switch](#).
- Verify that a physical switch is available.

Procedure

- 1 Log in to the vSphere Web Client.
- 2 Select **Networking & Security > Logical Switches**.
- 3 Locate the WebService logical switch and right-click to select **Manage Hardware Bindings** from the drop-down menu.
- 4 Select the hardware gateway profile.
- 5 Click the Add (+) icon and select the physical switch from the drop-down menu.
For example, AristaGW.
- 6 Click **Select** to choose a physical port from the Available Objects list.
For example, Ethernet 18.
- 7 Click **OK**.
- 8 Specify the VLAN name.



Switch	Port	VLAN
prmh-nsx-tor-7150s-1	Ethernet18	160

For example, 160.

- 9 Click **OK**.

Results

The binding is complete.

The NSX Controller synchronizes the physical and logical configuration information with the hardware gateway.

L2 Bridges



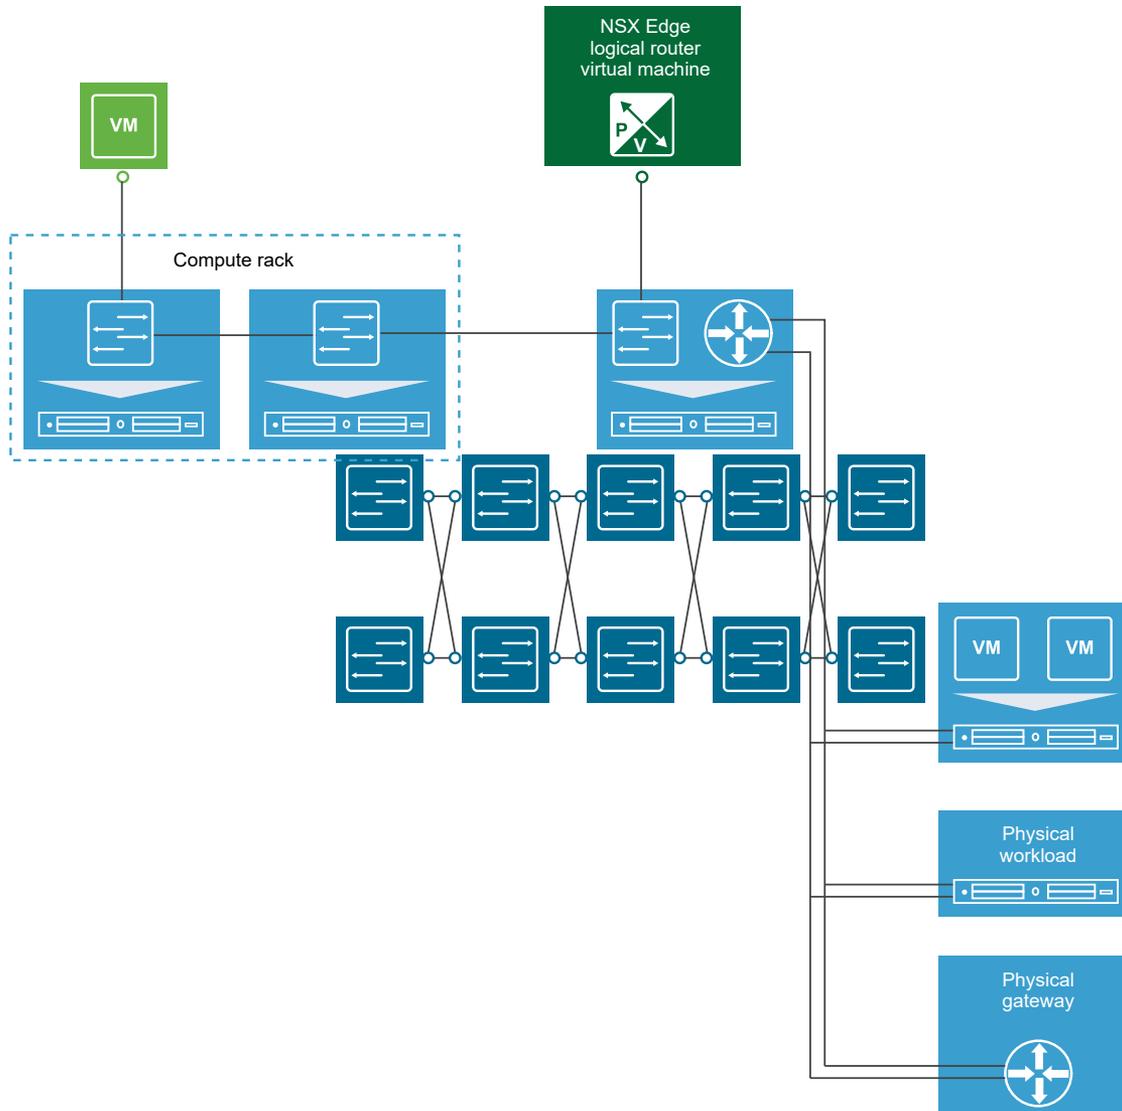
You can create an L2 bridge between a logical switch and a VLAN, which enables you to migrate virtual workloads to physical devices with no impact on IP addresses.

A Layer 2 bridge enables connectivity between the virtual and physical network by enabling virtual machines (VMs) to be connected to a physical server or network. Use cases include:

- Physical to virtual, or virtual to virtual migration. L2 bridging allows you to maintain connectivity between workloads inside NSX and outside NSX, without requiring IP re-addressing.
- Insertion into NSX of an appliance that cannot be virtualized, and that require L2 connectivity with its clients. This is common for some physical database servers.
- Service insertion. An L2 Bridge allows integrating transparently into NSX any physical appliance such as a router, load balancer or firewall

A logical network can leverage a physical L3 gateway and access existing physical networks and security resources by bridging the logical switch broadcast domain to the VLAN broadcast domain. The L2 bridge runs on the host that has the NSX DLR control virtual machine. An L2 bridge instance maps to a single VLAN, but there can be multiple bridge instances. The VLAN port group and VXLAN logical switch that is bridged must be on the same vSphere distributed switch (VDS) and both must share same physical NICs.

VXLAN (VNI) network and VLAN-backed port groups must be on the same distributed virtual switch (VDS).



Note that you should not use an L2 bridge to connect a logical switch to another logical switch, a VLAN network to another VLAN network, or to interconnect datacenters. Also, you cannot use a universal logical router to configure bridging and you cannot add a bridge to a universal logical switch.

This chapter includes the following topics:

- [Add L2 Bridge](#)
- [Add L2 Bridge to a Logically Routed Environment](#)
- [Improving Bridging Throughput](#)

Add L2 Bridge

You can add a bridge from a logical switch to a distributed virtual port group.

Prerequisites

A configured logical switch and a VLAN-backed distributed virtual port group.

The logical switch and the VLAN-backed distributed virtual port group that are to be bridged together must exist on the same Virtual Distributed Switch (VDS).

A DLR Control VM on a hypervisor where the VDS with the logical switch and VLAN-backed distributed virtual port group are instantiated must be deployed in your environment.

You cannot use a universal distributed logical router to configure bridging, and you cannot add a bridge to a universal logical switch.

Caution Bridged traffic enters and leaves an ESXi host through the uplink port on the dvSwitch that is used for the VXLAN traffic. VDS teaming or failover policy for VLAN is not used for the bridged traffic.

Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click a distributed logical router.
- 4 Click **Manage > Bridging**.
- 5 Click **Add**.
- 6 Enter a name for the bridge.

Caution Bridge name must not exceed 40 characters. Bridge configuration fails when the name exceeds 40 characters.

- 7 Select the logical switch that you want to create a bridge for.
- 8 Select the distributed virtual port group to which you want to bridge the logical switch.
- 9 Click **Publish** for the changes to take effect.

Add L2 Bridge to a Logically Routed Environment

You can bridge a given logical switch to a single VLAN with one active bridge instance. One logical router can have multiple bridging instances, however, the same VXLAN and VLAN cannot connect to more than one bridge instance.

You can use a logical switch to participate in both distributed logical routing and layer 2 bridging. Therefore, the traffic from the bridged logical switch does not need to flow through the centralized Edge VM. The traffic from the bridged logical switch can flow to the physical VLAN through the L2 bridge instance. The bridge instance gets enabled on the ESXi host where the DLR control VM is running.

For more information about L2 bridging in NSX, see the "NSX Distributed Routing and Layer 2 Bridging Integration" section in the *NSX Network Virtualization Design Guide* at <https://communities.vmware.com/docs/DOC-27683>.

Tip You can create multiple sets of layer 2 bridging instances and associate them with different DLRs. By following this practice, you can spread the bridging load across different ESXi hosts.

Prerequisites

- An NSX distributed logical router must be deployed in your environment.
- You cannot use a universal logical router to configure bridging, and you cannot add a bridge to a universal logical switch.

Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click the distributed logical router that you want to use for bridging.

Note The bridge instance must be created in the same routing instance to which the VXLAN is connected. One bridge instance can have one VXLAN and one VLAN, and the VXLAN and VLAN cannot overlap. The same VXLAN and VLAN cannot connect to more than one bridge instance.

- 4 Click **Manage > Bridging**.
- 5 Click **Add**.
- 6 Enter a name for the bridge.

Caution Bridge name must not exceed 40 characters. Bridge configuration fails when the name exceeds 40 characters.

- 7 Select the logical switch that you want to create a bridge for.
- 8 Select the distributed virtual port group to which you want to bridge the logical switch.
- 9 Click **Publish** for the changes to the bridging configuration to take effect.

The logical switch that is used for bridging appears with **Routing Enabled** specified. For more information, see [Add a Logical Switch](#) and [Connect Virtual Machines to a Logical Switch](#).

Improving Bridging Throughput

You can improve bridging throughput with Receive Side Scaling. Starting in NSX 6.4.2, you can also use Software Receive Side Scaling to improve bridging throughput.

With Receive Side Scaling (RSS) technology, you can spread incoming traffic across different receive descriptor queues. If you assign each queue to a different CPU core, the incoming traffic can be load balanced, improving performance.

However, RSS does not work well with unknown unicast and multicast traffic. These packets end up in the default queue processed by a single CPU core, which leads to low throughput. Most of the packets received by the ESXi host performing VLAN-VXLAN bridging belong to this category, so bridging throughput is low.

Some physical NIC vendors support a feature called Default Queue Receive Side Scaling (DRSS). Using DRSS, you can configure multiple hardware queues backing up the default RX queue, spreading VLAN-VXLAN flows across multiple CPU cores.

For physical NICs that do not support DRSS (for example, ixgbe, ixgben), you can use Software Receive Side Scaling (SoftRSS) to improve bridging network throughput.

SoftRSS offloads the handling of individual flows to one of the multiple kernel worlds, so the thread which pulls packets from the NIC can process more packets. Similar to RSS, network throughput improvement when using SoftRSS has a linear correlation with CPU utilization.

For more information, see [Enable Software Receive Side Scaling](#).

Enable Software Receive Side Scaling

You can improve VLAN-VXLAN bridging throughput using Software Receive Side Scaling.

Prerequisites

- Verify that you have NSX 6.4.2 or later installed.
- Determine which hosts should have SoftRSS enabled.
 - Enable SoftRSS on the ESXi hosts on which the active/standby bridge exists (where the DLR Control VMs are hosted). If the Control VMs might be migrated using vMotion, enable SoftRSS on all the hosts in the cluster.
 - If Default Queue Receive Side Scaling (DRSS) is supported on the host physical NIC, enable that and do not enable SoftRSS. Use `esxcli system module parameters list -m [nic_module]` to verify if DRSS is supported.

Procedure

- 1 Enable SoftRSS on a host.

The recommended value for *number of worlds* for a 10G uplink is 4.

You can increase the number of worlds up to a maximum of 16. You might want to increase the number if the uplink can support a higher rate (using link aggregation) or if you notice an uneven distribution of flows to the worlds based on your traffic pattern.

```
net-dvs -s com.vmware.net.vdr.softrss=[number of worlds] -p hostPropList [dvs name]
```

- 2 (Optional) If you want to disable SoftRSS, use this command.

```
net-dvs -u com.vmware.net.vdr.softrss -p hostPropList [dvs name]
```

You can specify static and dynamic routing for each NSX Edge.

Dynamic routing provides the necessary forwarding information between Layer 2 broadcast domains, thereby allowing you to decrease Layer 2 broadcast domains and improve network efficiency and scale. NSX extends this intelligence to where the workloads reside for doing East-West routing. This allows more direct virtual machine to virtual machine communication without the added cost or time needed to extend hops. At the same time, NSX also provides North-South connectivity, thereby enabling tenants to access public networks.

This chapter includes the following topics:

- [Add a Distributed Logical Router](#)
- [Add an Edge Services Gateway](#)
- [Specify Global Configuration](#)
- [NSX Edge Configuration](#)
- [Add a Static Route](#)
- [Configure OSPF on a Logical \(Distributed\) Router](#)
- [Configure OSPF on an Edge Services Gateway](#)
- [Configure BGP](#)
- [Configure Route Redistribution](#)
- [View the NSX Manager Locale ID](#)
- [Configure Locale ID on a Universal Logical \(Distributed\) Router](#)
- [Configure Locale ID on a Host or Cluster](#)
- [Multicast Routing Support, Limitations, and Topology](#)

Add a Distributed Logical Router

Distributed logical router (DLR) kernel modules in the host perform routing between VXLAN networks, and between virtual and physical networks. An NSX Edge Appliance provides dynamic routing ability if needed. Distributed logical routers can be created on both primary and secondary

NSX Managers in a cross-vCenter NSX environment, but universal distributed logical routers can be created only on the primary NSX Manager.

Note Starting in NSX Data Center 6.4.4, the term "Logical Router" is replaced with "Distributed Logical Router" in the vSphere Web Client. In the documentation, both terms are used interchangeably; however, they refer to the same object.

When deploying a new logical router, consider the following:

- NSX Data Center for vSphere 6.2 and later allows logical router-routed logical interfaces (LIFs) to be connected to a VXLAN that is bridged to a VLAN.
- Logical router interfaces and bridging interfaces cannot be connected to a dvPortgroup with the VLAN ID set to 0.
- A given logical router instance cannot be connected to logical switches that exist in different transport zones. This is to ensure that all logical switches and logical router instances are aligned.
- A logical router cannot be connected to VLAN-backed port groups if that logical router is connected to logical switches spanning more than one vSphere distributed switch (VDS). This is to ensure correct alignment of logical router instances with logical switch dvPortgroups across hosts.
- Logical router interfaces must not be created on two different distributed port groups (dvPortgroups) with the same VLAN ID if the two networks are in the same vSphere distributed switch.
- Logical router interfaces should not be created on two different dvPortgroups with the same VLAN ID if two networks are in different vSphere distributed switches, but the two vSphere distributed switches share identical hosts. In other words, logical router interfaces can be created on two different networks with the same VLAN ID if the two dvPortgroups are in two different vSphere distributed switches, as long as the vSphere distributed switches do not share a host.
- If VXLAN is configured, logical router interfaces must be connected to distributed port groups on the vSphere Distributed Switch where VXLAN is configured. Do not connect logical router interfaces to port groups on other vSphere Distributed Switches.

The following list describes feature support by interface type (uplink and internal) on the logical router:

- Dynamic routing protocols (BGP and OSPF) are supported only on uplink interfaces.
- Firewall rules are applicable only on uplink interfaces and are limited to control and management traffic that is destined to the Edge virtual appliance.

- For more information about the DLR Management Interface, see the Knowledge Base Article "Management Interface Guide: DLR Control VM - NSX" <http://kb.vmware.com/kb/2122060>.

Important If you enable high availability on an NSX Edge in a cross-vCenter NSX environment, both the active and standby NSX Edge Appliances must reside in the same vCenter Server. If you migrate one of the appliances of an NSX Edge HA pair to a different vCenter Server, the two HA appliances no longer operate as an HA pair, and you might experience traffic disruption.

Attention vSphere Fault Tolerance does not work with logical router control VM.

Prerequisites

- You must be assigned the **Enterprise Administrator** or **NSX Administrator** role.
- You must create a local segment ID pool, even if you have no plans to create logical switches.
- Make sure that the controller cluster is up and available before creating or changing a logical router configuration. A logical router cannot distribute routing information to hosts without the help of NSX controllers. A logical router relies on NSX controllers to function, while Edge Services Gateways (ESGs) do not.
- If a logical router is to be connected to VLAN dvPortgroups, ensure that all hypervisor hosts with a logical router appliance installed can reach each other on UDP port 6999. Communication on this port is required for logical router VLAN-based ARP proxy to work.
- Determine where to deploy the logical router appliance.
 - The destination host must be part of the same transport zone as the logical switches connected to the new logical router's interfaces.
 - Avoid placing it on the same host as one or more of its upstream ESGs if you use ESGs in an ECMP setup. You can use DRS anti-affinity rules to enforce this practice, reducing the impact of host failure on logical router forwarding. This guideline does not apply if you have one upstream ESG by itself or in HA mode. For more information, see the *NSX Network Virtualization Design Guide* at <https://communities.vmware.com/docs/DOC-27683>.
- Verify that the host cluster on which you install the logical router appliance is prepared for NSX Data Center for vSphere. See "Prepare Host Clusters for NSX" in the *NSX Installation Guide*.
- Determine the appropriate NSX Manager on which to make your changes.
 - In a standalone or single vCenter NSX environment there is only one NSX Manager so you do not need to select one.
 - Universal objects must be managed from the primary NSX Manager.
 - Objects local to an NSX Manager must be managed from that NSX Manager.
 - In a cross-vCenter NSX environment that does not have Enhanced Linked Mode enabled, you must make configuration changes from the vCenter linked to the NSX Manager that you want to modify.

- In a cross-vCenter NSX environment in Enhanced Linked Mode, you can make configuration changes to any NSX Manager from any linked vCenter. Select the appropriate NSX Manager from the NSX Manager drop-down menu.
- Determine which kind of logical router you want to add:
 - To connect a logical switch, add a logical router.
 - To connect a universal logical switch, add a universal logical router.
- If you are adding a universal logical router, determine if you need to enable local egress. Local egress allows you to selectively send routes to hosts. You may want this ability if your NSX deployment spans multiple sites. See [Cross-vCenter NSX Topologies](#) for more information. You cannot enable local egress after the universal logical router has been created.

Procedure

- 1 In the vSphere Web Client, navigate to **Home > Networking & Security > NSX Edges**.
- 2 Select the appropriate NSX Manager on which to make your changes. If you are creating a universal logical router, you must select the primary NSX Manager.
- 3 Click **Add**, and then select the type of logical router you want to add:
 - Select **Logical (Distributed) Router** to add a logical router local to the selected NSX Manager.
 - Select **Universal Logical (Distributed) Router** to add a logical router that can span a cross-vCenter NSX environment. This option is available only if you have assigned a primary NSX Manager, and are making changes from the primary NSX Manager. If you select **Universal Logical (Distributed) Router**, you can optionally enable local egress.
- 4 Enter name, description, and other details of the logical router.

Option	Description
Name	Enter a name for the logical router as you want it to appear in the vCenter inventory. Make sure that this name is unique across all logical routers within a single tenant.
Host Name	Optional. Enter a host name that you want to display for the logical router in the CLI. If you do not enter a host name, the Edge ID that is created automatically is displayed in the CLI.
Description	Optional. Enter a description of the logical router.
Deploy Edge Appliance	By default, this option is selected. An Edge appliance (also called a logical router virtual appliance) is required for dynamic routing and the logical router appliance's firewall, which applies to logical router pings, SSH access, and dynamic routing traffic. If you require only static routes, and do not want to deploy an Edge appliance, deselect this option. You cannot add an Edge appliance to the logical router after the logical router is created.

Option	Description
High Availability	Optional. By default, HA is disabled. Select this option to enable and configure HA on the logical router. If you are planning to do dynamic routing, HA is required.
HA Logging	Optional. By default, HA logging is disabled. When logging is enabled, the default log level is set to info. You can change it, if necessary.

5 Specify the CLI settings and other settings of the logical router.

Option	Description
User Name	Enter a user name that you want to use for logging in to the Edge CLI.
Password	Enter a password that is at least 12 characters and it must satisfy these rules: <ul style="list-style-type: none"> ■ Must not exceed 255 characters ■ At least one uppercase letter and one lowercase letter ■ At least one number ■ At least one special character ■ Must not contain the user name as a substring ■ Must not consecutively repeat a character 3 or more times.
Confirm password	Reenter the password to confirm.
SSH access	Optional. By default, SSH access is disabled. If you do not enable SSH, you can still access the logical router by opening the virtual appliance console. Enabling SSH causes the SSH process to run on the logical router. You must adjust the logical router firewall configuration manually to allow SSH access to the logical router's protocol address. The protocol address is configured when you configure dynamic routing on the logical router.
FIPS mode	Optional. By default, FIPS mode is disabled. When you enable FIPS mode, any secure communication to or from the NSX Edge uses cryptographic algorithms or protocols that are allowed by FIPS.
Edge control level logging	Optional. By default, the log level is info.

6 Configure deployment of the NSX Edge Appliance.

- ◆ If you did not select **Deploy Edge Appliance**, you cannot add an appliance. Click **Next** to continue with the configuration.
- ◆ If you selected **Deploy Edge Appliance**, enter the settings of the logical router virtual appliance.

For example:

Option	Value
Cluster/Resource Pool	Management & Edge
Datastore	ds-1
Host	esxmgt-01a.corp.local
Resource Reservation	System Managed

See "Managing NSX Edge Appliance Resource Reservations" in the *NSX Administration Guide* for more information on Resource Reservation.

7 Configure the HA interface connection, and optionally an IP address.

If you selected **Deploy Edge Appliance**, you must connect the HA interface to a distributed port group or a logical switch. If you are using this interface as an HA interface only, use a logical switch. A /30 subnet is allocated from the link local range 169.254.0.0/16 and is used to provide an IP address for each of the two NSX Edge appliances.

Optionally, if you want to use this interface to connect to the NSX Edge, you can specify an extra IP address and prefix for the HA interface.

Note Before NSX Data Center for vSphere 6.2, HA interface was called management interface. You cannot do an SSH connection to a HA interface from anywhere that is not on the same IP subnet as the HA interface. You cannot configure static routes that point out of the HA interface, which means that RPF will drop incoming traffic. However, you can, in theory, disable RPF, but this action is counterproductive to high availability. For SSH access, you can also use the logical router's protocol address, which is configured later when you configure dynamic routing.

In NSX Data Center for vSphere 6.2 and later, the HA interface of a logical router is automatically excluded from route redistribution.

For example, the following table shows a sample HA interface configuration where the HA interface is connected to a management dvPortgroup.

Option	Description
Connected To	Mgmt_VDS-Mgmt
IP Address	192.168.110.60*
Subnet Prefix Length	24

8 Configure interfaces of the NSX Edge.

- a Specify the name, type, and other basic interface details.

Option	Description
Name	Enter a name for the interface.
Type	Select either Internal or Uplink. The internal interfaces are for connections to switches that allow VM-to-VM (sometimes called East-West) communication. Internal interfaces are created as pseudo vNICs on the logical router virtual appliance. Uplink interfaces are for North-South communication, and they are created as vNICs on the logical router virtual appliance. A logical router uplink interface might connect to an Edge Services Gateway or a third-party router VM. You must have at least one uplink interface for dynamic routing to work.
Connected To	Select the distributed virtual port group or the logical switch to which you want to connect this interface to.

- b Configure the subnets of the interface.

Option	Description
Primary IP Address	On logical routers, only IPv4 addressing is supported. The interface configuration that you enter here is modifiable later. You can add, remove, and modify interfaces after a logical router is deployed.
Subnet Prefix Length	Enter the subnet mask of the interface.

- c (Optional) Edit the default MTU value, if necessary. The default value for both uplink and internal interface is 1500.

The following table shows an example of two internal interfaces (app and web) and one uplink interface (to-ESG).

Table 9-1. Example: NSX Edge Interfaces

Name	IP address	Subnet Prefix Length	Connected To
app	172.16.20.1*	24	app
web	172.16.10.1*	24	web
to-ESG	192.168.10.2*	29	transit

9 Configure the default gateway settings.

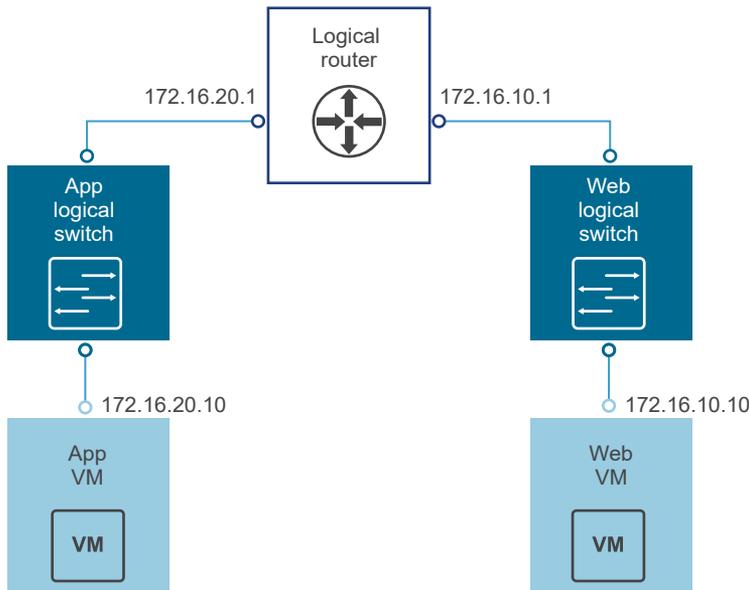
For example:

Option	Value
vNIC	Uplink
Gateway IP	192.168.10.1
MTU	1500

10 Make sure that the VMs connected to the logical switches have their default gateways set properly to the logical router interface IP addresses.

Results

In the following example topology, the default gateway of app VM is 172.16.20.1. The default gateway of web VM is 172.16.10.1. Make sure the VMs can ping their default gateways and each other.



Connect to the NSX Manager using SSH or the console, and run the following commands:

- List all logical router instance information.

```

nsxmgr-1-01a> show logical-router list all
Edge-id          Vdr Name          Vdr id           #Lifs
edge-1          default+edge-1    0x00001388      3
    
```

- List the hosts that have received routing information for the logical router from the controller cluster.

```
nsxmgr-1-01a> show logical-router list dlr edge-1 host
ID                HostName
host-25           192.168.210.52
host-26           192.168.210.53
host-24           192.168.110.53
```

The output includes all hosts from all host clusters that are configured as members of the transport zone that owns the logical switch that is connected to the specified logical router (edge-1 in this example).

- List the routing table information that is communicated to the hosts by the logical router. Routing table entries should be consistent across all the hosts.

```
nsx-mgr-1-01a> show logical-router host host-25 dlr edge-1 route

VDR default+edge-1 Route Table
Legend: [U: Up], [G: Gateway], [C: Connected], [I: Interface]
Legend: [H: Host], [F: Soft Flush] [!: Reject] [E: ECMP]
```

Destination	GenMask	Gateway	Flags	Ref	Origin	UpTime	Interface
0.0.0.0	0.0.0.0	192.168.10.1	UG	1	AUTO	4101	
138800000002							
172.16.10.0	255.255.255.0	0.0.0.0	UCI	1	MANUAL	10195	
13880000000b							
172.16.20.0	255.255.255.0	0.0.0.0	UCI	1	MANUAL	10196	
13880000000a							
192.168.10.0	255.255.255.248	0.0.0.0	UCI	1	MANUAL	10196	
138800000002							
192.168.100.0	255.255.255.0	192.168.10.1	UG	1	AUTO	3802	
138800000002							

- List additional information about the router from the point of view of one of the hosts. This output is helpful to learn which controller is communicating with the host.

```
nsx-mgr-1-01a> show logical-router host host-25 dlr edge-1 verbose

VDR Instance Information :
-----

Vdr Name:                default+edge-1
Vdr Id:                  0x00001388
Number of Lifs:          3
Number of Routes:        5
State:                   Enabled
Controller IP:           192.168.110.203
Control Plane IP:        192.168.210.52
```

```
Control Plane Active:    Yes
Num unique nexthops:    1
Generation Number:      0
Edge Active:            No
```

Check the Controller IP field in the output of the `show logical-router host host-25 dlr edge-1 verbose` command.

SSH to a controller, and run the following commands to display the controller's learned VNI, VTEP, MAC, and ARP table state information.

```
192.168.110.202 # show control-cluster logical-switches vni 5000
VNI      Controller      BUM-Replication ARP-Proxy Connections
5000     192.168.110.201 Enabled           Enabled          0
```

The output for VNI 5000 shows zero connections and lists controller 192.168.110.201 as the owner for VNI 5000. Log in to that controller to gather further information for VNI 5000.

```
192.168.110.201 # show control-cluster logical-switches vni 5000
VNI      Controller      BUM-Replication ARP-Proxy Connections
5000     192.168.110.201 Enabled           Enabled          3
```

The output on 192.168.110.201 shows three connections. Check additional VNIs.

```
192.168.110.201 # show control-cluster logical-switches vni 5001
VNI      Controller      BUM-Replication ARP-Proxy Connections
5001     192.168.110.201 Enabled           Enabled          3
```

```
192.168.110.201 # show control-cluster logical-switches vni 5002
VNI      Controller      BUM-Replication ARP-Proxy Connections
5002     192.168.110.201 Enabled           Enabled          3
```

Because 192.168.110.201 owns all three VNI connections, we expect to see zero connections on the other controller, 192.168.110.203.

```
192.168.110.203 # show control-cluster logical-switches vni 5000
VNI      Controller      BUM-Replication ARP-Proxy Connections
5000     192.168.110.201 Enabled           Enabled          0
```

- Before checking the MAC and ARP tables, ping from one VM to the other VM.

From app VM to web VM:

```
vmware@app-vm$ 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=64 time=2.605 ms
64 bytes from 172.16.10.10: icmp_req=2 ttl=64 time=1.490 ms
64 bytes from 172.16.10.10: icmp_req=3 ttl=64 time=2.422 ms
```

Check the MAC tables.

```
192.168.110.201 # show control-cluster logical-switches mac-table 5000
VNI      MAC                VTEP-IP          Connection-ID
5000     00:50:56:a6:23:ae 192.168.250.52  7
```

```
192.168.110.201 # show control-cluster logical-switches mac-table 5001
VNI      MAC                VTEP-IP          Connection-ID
5001     00:50:56:a6:8d:72 192.168.250.51  23
```

Check the ARP tables.

```
192.168.110.201 # show control-cluster logical-switches arp-table 5000
VNI      IP                MAC                Connection-ID
5000     172.16.20.10     00:50:56:a6:23:ae 7
```

```
192.168.110.201 # show control-cluster logical-switches arp-table 5001
VNI      IP                MAC                Connection-ID
5001     172.16.10.10     00:50:56:a6:8d:72 23
```

Check the logical router information. Each logical router instance is served by one of the controller nodes.

The `instance` subcommand of `show control-cluster logical-routers` command displays a list of logical routers that are connected to this controller.

The `interface-summary` subcommand displays the LIFs that the controller learned from the NSX Manager. This information is sent to the hosts that are in the host clusters managed under the transport zone.

The `routes` subcommand shows the routing table that is sent to this controller by the logical router's virtual appliance (also known as the control VM). Unlike on the ESXi hosts, this routing table does not include directly connected subnets because this information is provided by the LIF configuration. Route information on the ESXi hosts includes directly connected subnets because in that case it is a forwarding table used by ESXi host's datapath.

- List all logical routers connected to this controller.

```
controller # show control-cluster logical-routers instance all
LR-Id      LR-Name          Universal Service-Controller Egress-Locale
0x1388     default+edge-1  false      192.168.110.201  local
```

Note the LR-Id and use it in the following command.

- ```
controller # show control-cluster logical-routers interface-summary 0x1388
Interface Type Id IP[]
13880000000b vxlan 0x1389 172.16.10.1/24
13880000000a vxlan 0x1388 172.16.20.1/24
138800000002 vxlan 0x138a 192.168.10.2/29
```

- controller # show control-cluster logical-routers routes 0x1388

```

Destination Next-Hop[] Preference Locale-Id Source
192.168.100.0/24 192.168.10.1 110 00000000-0000-0000-0000-000000000000
CONTROL_VM
0.0.0.0/0 192.168.10.1 0 00000000-0000-0000-0000-000000000000
CONTROL_VM

```

```

[root@comp02a:~] esxcfg-route -l
VMkernel Routes:
Network Netmask Gateway Interface
10.20.20.0 255.255.255.0 Local Subnet vmk1
192.168.210.0 255.255.255.0 Local Subnet vmk0
default 0.0.0.0 192.168.210.1 vmk0

```

- Display the controller connections to the specific VNI.

```

192.168.110.203 # show control-cluster logical-switches connection-table 5000
Host-IP Port ID
192.168.110.53 26167 4
192.168.210.52 27645 5
192.168.210.53 40895 6

```

```

192.168.110.202 # show control-cluster logical-switches connection-table 5001
Host-IP Port ID
192.168.110.53 26167 4
192.168.210.52 27645 5
192.168.210.53 40895 6

```

These Host-IP addresses are vmk0 interfaces, not VTEPs. Connections between ESXi hosts and controllers are created on the management network. The port numbers here are ephemeral TCP ports that are allocated by the ESXi host IP stack when the host establishes a connection with the controller.

- On the host, you can view the controller network connection matched to the port number.

```

[root@192.168.110.53:~] #esxccli network ip connection list | grep 26167
tcp 0 0 192.168.110.53:26167 192.168.110.101:1234
ESTABLISHED 96416 newreno netcpa-worker

```

- Display active VNIs on the host. Observe how the output is different across hosts. Not all VNIs are active on all hosts. A VNI is active on a host if the host has a VM that is connected to the logical switch.

```

[root@192.168.210.52:~] # esxccli network vswitch dvs vmware vxlan network list --vds-name
Compute_VDS
VXLAN ID Multicast IP Control Plane Controller
Connection Port Count MAC Entry Count ARP Entry Count VTEP Count


```

|      |                           |                                      |                 |
|------|---------------------------|--------------------------------------|-----------------|
| 5000 | N/A (headend replication) | Enabled (multicast proxy, ARP proxy) | 192.168.110.203 |
| (up) | 1                         | 0                                    | 0               |
| 5001 | N/A (headend replication) | Enabled (multicast proxy, ARP proxy) | 192.168.110.202 |
| (up) | 1                         | 0                                    | 0               |

**Note** To enable the vxlan namespace in vSphere 6.0 and later, run the `/etc/init.d/hostd restart` command.

For logical switches in hybrid or unicast mode, the `esxcli network vswitch dvs vmware vxlan network list --vds-name <vds-name>` command contains the following output:

- Control Plane is enabled.
- Multicast proxy and ARP proxy are listed. AARP proxy is listed even if you disabled IP discovery.
- A valid controller IP address is listed and the connection is up.
- If a logical router is connected to the ESXi host, the port Count is at least 1, even if there are no VMs on the host connected to the logical switch. This one port is the `vdrPort`, which is a special `dvPort` connected to the logical router kernel module on the ESXi host.
- First ping from VM to another VM on a different subnet and then display the MAC table. Note that the Inner MAC is the VM entry while the Outer MAC and Outer IP refer to the VTEP.

```
~ # esxcli network vswitch dvs vmware vxlan network mac list --vds-name=Compute_VDS --
vxlan-id=5000
Inner MAC Outer MAC Outer IP Flags

00:50:56:a6:23:ae 00:50:56:6a:65:c2 192.168.250.52 00000111
```

```
~ # esxcli network vswitch dvs vmware vxlan network mac list --vds-name=Compute_VDS --
vxlan-id=5001
Inner MAC Outer MAC Outer IP Flags

02:50:56:56:44:52 00:50:56:6a:65:c2 192.168.250.52 00000101
00:50:56:f0:d7:e4 00:50:56:6a:65:c2 192.168.250.52 00000111
```

### What to do next

When you install an NSX Edge Appliance, NSX enables automatic VM startup/shutdown on the host if vSphere HA is disabled on the cluster. If the appliance VMs are later migrated to other hosts in the cluster, the new hosts might not have automatic VM startup/shutdown enabled. For this reason, when you install NSX Edge Appliances on clusters that have vSphere HA disabled, you must preferably check all hosts in the cluster to make sure that automatic VM startup/shutdown is enabled. See "Edit Virtual Machine Startup and Shutdown Settings" in *vSphere Virtual Machine Administration*.

After the logical router is deployed, double-click the logical router ID to configure additional settings, such as interfaces, routing, firewall, bridging, and DHCP relay.

## Add an Edge Services Gateway

You can install multiple NSX Edge services gateway virtual appliances in a data center. Each NSX Edge Appliance can have a total of ten uplink and internal network interfaces. The internal interfaces connect to secured port groups and act as the gateway for all protected virtual machines in the port group. The subnet assigned to the internal interface can be a publicly routed IP address space or a NATed/routed RFC 1918 private space. Firewall rules and other NSX Edge services are enforced on traffic between interfaces.

Uplink interfaces of an ESG connect to uplink port groups that have access to a shared corporate network or a service that provides access layer networking.

The following list describes feature support by interface type (internal and uplink) on an ESG.

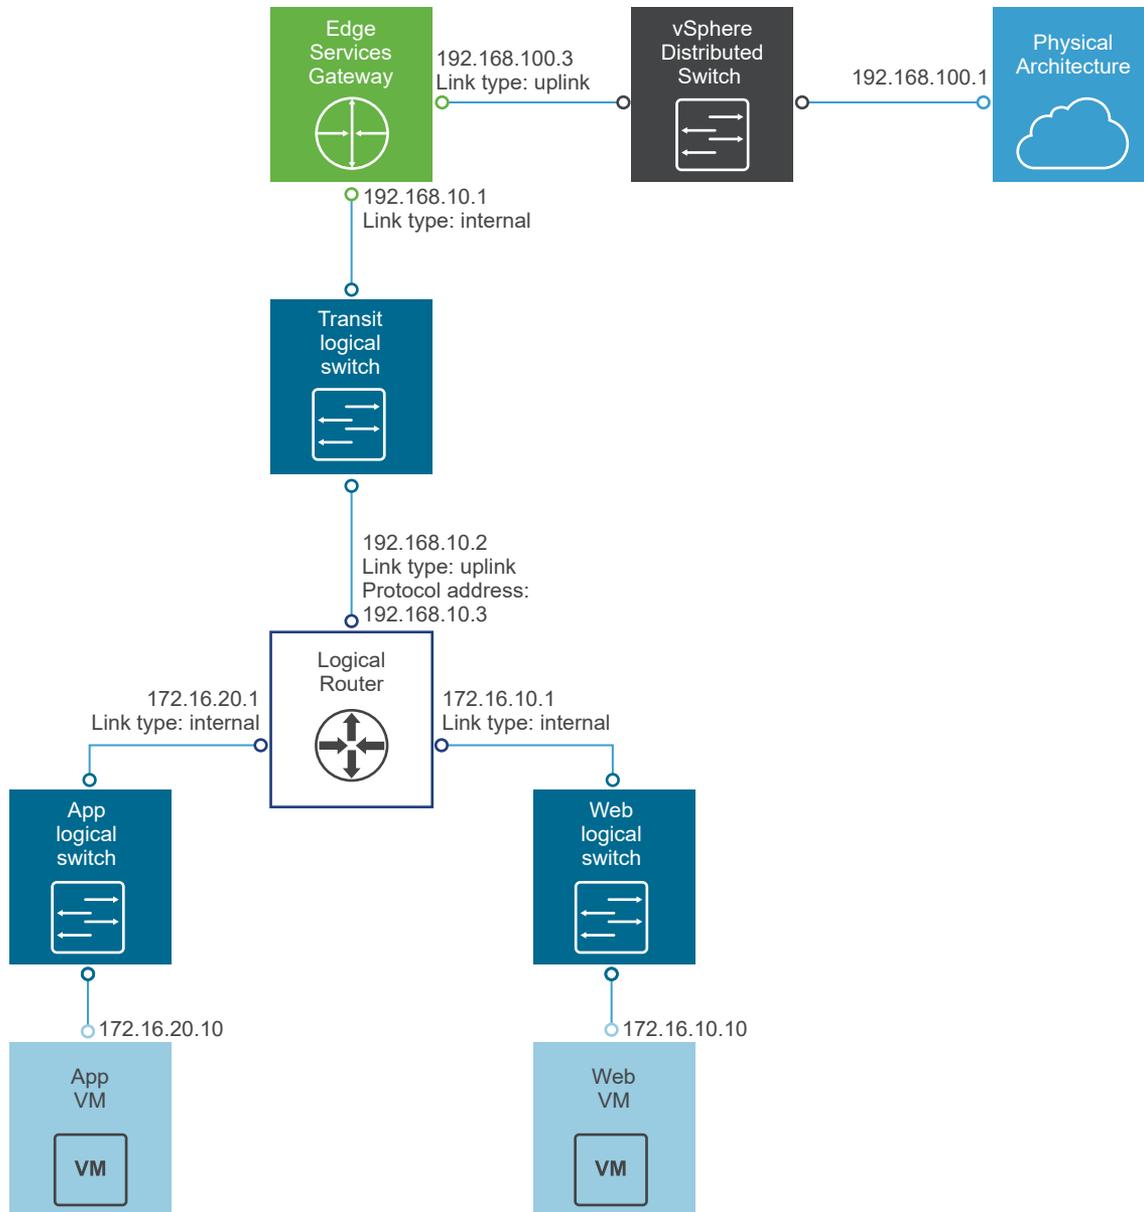
- DHCP: Not supported on uplink interfaces. See the note after this bulleted list.
- DNS Forwarder: Not supported on uplink interfaces.
- HA: Not supported on uplink interfaces, requires at least one internal interface.
- SSL VPN: Listener IP must belong to an uplink interface.
- IPsec VPN: Local site IP must belong to an uplink interface.
- L2 VPN: Only internal networks can be stretched.

---

**Note** By design, DHCP service is supported on the internal interfaces of an NSX Edge. However, in some situations, you may choose to configure DHCP on an uplink interface of the edge and configure no internal interfaces. In this situation, the edge can listen to the DHCP client requests on the uplink interface, and dynamically assign IP addresses to the DHCP clients. Later, if you configure an internal interface on the same edge, DHCP service stops working because the edge starts listening to the DHCP client requests on the internal interface.

---

The following figure shows a sample topology. The Edge Service Gateway uplink interface is connected to the physical infrastructure through the vSphere distributed switch. The Edge Service Gateway internal interface is connected to a logical router through a logical transit switch.



You can configure multiple external IP addresses for load balancing, site-to-site VPN, and NAT services.

**Important** If you enable high availability on an NSX Edge in a cross-vCenter NSX environment, both the active and standby NSX Edge Appliances must reside in the same vCenter Server. If you migrate one of the appliances of an NSX Edge HA pair to a different vCenter Server, the two HA appliances no longer operate as an HA pair, and you might experience traffic disruption.

#### Prerequisites

- You must be assigned the **Enterprise Administrator** or **NSX Administrator** role.

- Verify that the resource pool has enough capacity for the Edge Services Gateway (ESG) virtual appliance to be deployed. See [Chapter 1 System Requirements for NSX Data Center for vSphere](#) for the resources required for each size of appliance.
- Verify that the host clusters on which the NSX Edge Appliance will be installed are prepared for NSX. See "Prepare Host Clusters for NSX" in the *NSX Installation Guide*.
- Determine if you want to enable DRS. If you create an Edge Services Gateway with HA, and DRS is enabled, DRS anti-affinity rules are created to prevent the appliances from being deployed on the same host. If DRS is not enabled at the time the appliances are created, the rules are not created and the appliances might be deployed on or moved to the same host.

### Procedure

- 1 Log in to the vSphere Web Client, and navigate to **Home > Networking & Security > NSX Edges**.
- 2 Click **Add**, and then click **Edge Services Gateway**.
- 3 Enter name, description, and other details of the ESG.

| Option             | Description                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Name</b>        | Enter a name for the ESG as you want it to appear in the vCenter inventory. Make sure that this name is unique across all ESGs within a single tenant.                            |
| <b>Host Name</b>   | Optional. Enter a host name that you want to display for this ESG in the CLI. If you do not enter a host name, the Edge ID that is created automatically is displayed in the CLI. |
| <b>Description</b> | Optional. Enter a description of the ESG.                                                                                                                                         |

| Option                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Deploy NSX Edge</b>   | Optional. Select this option to create an NSX Edge Appliance virtual machine. If you do not select this option, the ESG will not operate until a VM is deployed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>High Availability</b> | <p>Optional. Select this option to enable and configure high availability on the ESG.</p> <ul style="list-style-type: none"> <li>■ If you need to run stateful services on an ESG, such as load balancer, NAT, DHCP, and so on, you can enable HA on the edge. HA helps in minimizing the failover time to a standby edge when an active edge fails. Enabling HA deploys a standalone edge on a different host in a cluster. So, you must ensure that you have enough resources in your environment.</li> <li>■ If you are not running stateful services on the ESG, and your ESG is used only for north-south routing, then enabling ECMP is recommended. ECMP uses a dynamic routing protocol to learn the next-hop towards a final destination and to converge during failures.</li> </ul> <p>ECMP configuration can substantially increase bandwidth by load-balancing traffic over multiple paths and providing fault tolerance for failed paths. In this configuration, data plane outage is limited to only a subset of the traffic. You also have the option of enabling HA on each ESG to provide a faster failover rather than relying on vSphere HA. In an ECMP configuration too, you must ensure that you have sufficient resources in your environment.</p> <p>You can enable ECMP on the edge while doing the global routing configuration, and not while deploying the edge in your network.</p> |

#### 4 Specify the CLI settings and other settings of the ESG.

| Option                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>User Name</b>        | Enter a user name that you want to use for logging in to the Edge CLI.                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Password</b>         | <p>Enter a password that is at least 12 characters and it must satisfy these rules:</p> <ul style="list-style-type: none"> <li>■ Must not exceed 255 characters</li> <li>■ At least one uppercase letter and one lowercase letter</li> <li>■ At least one number</li> <li>■ At least one special character</li> <li>■ Must not contain the user name as a substring</li> <li>■ Must not consecutively repeat a character 3 or more times.</li> </ul> |
| <b>Confirm password</b> | Reenter the password to confirm.                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>SSH access</b>       | Optional. Enable SSH access to the Edge. By default, SSH access is disabled. Usually, SSH access is recommended for troubleshooting purposes.                                                                                                                                                                                                                                                                                                        |
| <b>FIPS mode</b>        | <p>Optional. By default, FIPS mode is disabled.</p> <p>When you enable FIPS mode, any secure communication to or from the NSX Edge uses cryptographic algorithms or protocols that are allowed by FIPS.</p>                                                                                                                                                                                                                                          |

| Option                            | Description                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Auto rule generation</b>       | Optional. By default, this option is enabled. This option allows automatic creation of firewall rules, NAT, and routing configuration, which control traffic for certain NSX Edge services, including load balancing and VPN.<br><br>If you disable automatic rule generation, you must manually add these rules and configurations. Auto rule generation does not create rules for data-channel traffic. |
| <b>Edge control level logging</b> | Optional. By default, the log level is info.                                                                                                                                                                                                                                                                                                                                                              |

## 5 Configure the deployment of the NSX Edge Appliance.

- a Select the size of the appliance depending on your environment.

| Appliance Size    | Description                                                                                                            |
|-------------------|------------------------------------------------------------------------------------------------------------------------|
| <b>Compact</b>    | Suitable only for laboratory or PoC environments.                                                                      |
| <b>Large</b>      | Provides more CPU, memory, and disk space than Compact, and supports a larger number of concurrent SSL VPN-Plus users. |
| <b>Quad Large</b> | Suitable when you need a high throughput and a high connection rate.                                                   |
| <b>X-Large</b>    | Suitable for environments that have a load balancer with millions of concurrent sessions.                              |

See [Chapter 1 System Requirements for NSX Data Center for vSphere](#) for the resources required for each size of appliance.

- b Add an NSX Edge Appliance, and specify the resource details for the VM deployment.

For example:

| Option                       | Value                  |
|------------------------------|------------------------|
| <b>Cluster/Resource Pool</b> | Management & Edge      |
| <b>Datastore</b>             | ds-1                   |
| <b>Host</b>                  | esxmgmt-01a.corp.local |
| <b>Resource Reservation</b>  | System Managed         |

See "Managing NSX Edge Appliance Resource Reservations" in the *NSX Administration Guide* for more information on Resource Reservation.

If you enabled HA, you can add two appliances. If you add a single appliance, NSX Edge replicates its configuration for the standby appliance. For HA to work correctly, you must deploy both appliances on a shared datastore.

## 6 Configure interfaces of the ESG.

- a Specify the name, type, and other basic interface details.

| Option              | Description                                                                                                                   |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Name</b>         | Enter a name for the interface.                                                                                               |
| <b>Type</b>         | Select either Internal or Uplink. For High Availability to work, an Edge appliance must have at least one internal interface. |
| <b>Connected To</b> | Select the port group or the logical switch to which you want to connect this interface to.                                   |

- b Configure the subnets of the interface.

| Option                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Primary IP Address</b>     | <p>On an ESG, both IPv4 and IPv6 addresses are supported. An interface can have one primary IP address, multiple secondary IP addresses, and multiple non-overlapping subnets.</p> <p>If you enter more than one IP address for the interface, you can select the primary IP address.</p> <p>Only one primary IP address is allowed per interface and the Edge uses the primary IP address as the source address for locally generated traffic, for example remote syslog and operator-initiated pings.</p> |
| <b>Secondary IP Addresses</b> | Enter the secondary IP address. To enter multiple IP addresses, use a comma-separated list.                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Subnet Prefix Length</b>   | Enter the subnet mask of the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

- c Specify the following options for the interface.

| Option                     | Description                                                                                                                                                                                                                                                                                                                    |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>MAC Addresses</b>       | <p>Optional. You can enter a MAC address for each interface.</p> <p>If you change the MAC address using an API call later, you must redeploy the Edge after changing the MAC address.</p>                                                                                                                                      |
| <b>MTU</b>                 | The default value for uplink and internal interface is 1500. For trunk interface, the default value is 1600. You can modify the default value, if necessary. For sub-interfaces on the trunk, the default value is 1500. Make sure that the MTU for the trunk interface is equal to or more than the MTU of the sub interface. |
| <b>Proxy ARP</b>           | <p>Select this option if you want the ESG to answer ARP requests intended for other virtual machines.</p> <p>This option is useful, for example, when you have the same subnet on both sides of a WAN connection.</p>                                                                                                          |
| <b>Send ICMP Redirect</b>  | Select this option if you want the ESG to convey routing information to the hosts.                                                                                                                                                                                                                                             |
| <b>Reverse Path Filter</b> | <p>By default, this option is set to enabled. When enabled, it verifies the reachability of the source address in packets being forwarded.</p> <p>In enabled mode, the packet must be received on the interface that the router might use to forward the return packet.</p>                                                    |

| Option                  | Description                                                                                                                                                                                                                                                                                |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                         | In loose mode, the source address must appear in the routing table.                                                                                                                                                                                                                        |
| <b>Fence Parameters</b> | Configure fence parameters if you want to reuse IP and MAC addresses across different fenced environments.<br>For example, in a cloud management platform (CMP), fencing allows you to run several cloud instances simultaneously with the same IP and MAC addresses isolated or "fenced". |

The following table shows an example of two NSX Edge interfaces. The uplink interface attaches the ESG to the outside world through an uplink port group on a vSphere distributed switch. The internal interface attaches the ESG to a logical transit switch to which a distributed logical router is also attached.

**Table 9-2. Example: NSX Edge Interfaces**

| vNIC# | Name     | IP address     | Subnet Prefix Length | Connected To       |
|-------|----------|----------------|----------------------|--------------------|
| 0     | Uplink   | 192.168.100.30 | 24                   | Mgmt_VDS-HQ_Uplink |
| 1     | Internal | 192.168.10.1*  | 29                   | transit-switch     |

**Important** NSX 6.4.4 and earlier supports multicast on a single uplink interface of the ESG. Starting with NSX 6.4.5, multicast is supported on a maximum of two uplink interfaces of the ESG. In a multi-vCenter deployment scenario, if an NSX Edge is at version 6.4.4 or earlier, you can enable multicast only on a single uplink interface. To enable multicast on two uplink interfaces, you must upgrade the Edge to 6.4.5 or later.

## 7 Configure the default gateway settings.

For example:

| Option     | Value         |
|------------|---------------|
| vNIC       | Uplink        |
| Gateway IP | 192.168.100.2 |
| MTU        | 1500          |

**Note** You can edit the MTU value, but it cannot be more than the configured MTU on the interface.

## 8 Configure the default firewall policy.

**Caution** If you do not configure the firewall policy, the default policy is set to deny all traffic. However, the firewall is enabled on the ESG during deployment, by default.

## 9 Configure ESG logging and HA parameters.

- a Enable or disable logging on the NSX Edge Appliance.

By default, logs are enabled on all new NSX Edge appliances. The default logging level is Info. If logs are stored locally on the ESG, logging might generate too many logs and affect the performance of your NSX Edge. For this reason, you must preferably configure remote syslog servers, and forward all logs to a centralized collector for analysis and monitoring.

- b If you enabled high availability, configure the following HA parameters.

| Option                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>vNIC</b>              | <p>Select the internal interface for which you want to configure HA parameters. By default, HA automatically selects an internal interface and automatically assigns link-local IP addresses.</p> <p>If you select ANY for interface but there are no internal interfaces configured, the UI displays an error. Two Edge appliances are created but since there is no internal interface configured, the new NSX Edge remains in standby and HA is disabled. After an internal interface is configured, HA is enabled on the NSX Edge appliance.</p>        |
| <b>Declare Dead Time</b> | <p>Enter the period in seconds within which, if the backup appliance does not receive a heartbeat signal from the primary appliance, the primary appliance is considered inactive and the backup appliance takes over. The default interval is 15 seconds.</p>                                                                                                                                                                                                                                                                                              |
| <b>Management IPs</b>    | <p>Optional: You can enter two management IP addresses in CIDR format to override the local link IP addresses assigned to the HA virtual machines. Ensure that the management IP addresses do not overlap with the IP addresses used for any other interface and do not interfere with traffic routing. Do not use an IP address that exists somewhere else on your network, even if that network is not directly attached to the appliance. The management IP addresses must be in the same L2/subnet and must be able to communicate with each other.</p> |

## 10 Review all the ESG settings before deploying the appliance.

### Results

After the ESG is deployed, go to the Hosts and Clusters view and open the console of the NSX Edge virtual appliance. From the console, make sure that you can ping the connected interfaces.

### What to do next

When you install an NSX Edge Appliance, NSX enables automatic VM startup/shutdown on the host if vSphere HA is disabled on the cluster. If the appliance VMs are later migrated to other hosts in the cluster, the new hosts might not have automatic VM startup/shutdown enabled. For this reason, when you install NSX Edge Appliances on clusters that have vSphere HA disabled, you must preferably check all hosts in the cluster to make sure that automatic VM startup/shutdown is enabled. See "Edit Virtual Machine Startup and Shutdown Settings" in *vSphere Virtual Machine Administration*.

Now you can configure routing to allow connectivity from external devices to your VMs.

## Specify Global Configuration

You can configure the default gateway for static routes and specify dynamic routing details for an Edge Services Gateway or Distributed Logical Router.

You must have a working NSX Edge instance before you can configure routing on it. For information on setting up NSX Edge, see [NSX Edge Configuration](#).

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Routing** and then click **Global Configuration**.
- 5 To enable equal-cost multi-path routing (ECMP), next to **ECMP**, click **Start**.

ECMP is a routing strategy that allows next-hop packet forwarding to a single destination over multiple best paths. These best paths can be added as static routes or as a result of metric calculations by dynamic routing protocols like OSPF or BGP. Multiple paths for static routes can be added by providing multiple next hops separated by commas in the Static Routes dialog box. For more information, see [Add a Static Route](#).

The Edge Services Gateway uses the Linux network stack implementation, a round-robin algorithm with a randomness component. After a next hop is selected for a particular source and destination IP address pair, the route cache stores the selected next hop. All packets for that flow go to the selected next hop. The default IPv4 route cache timeout is 300 seconds (gc\_timeout). If an entry is inactive for this time, it is eligible to be removed from the route cache. The actual removal happens when garbage collection timer activates (gc\_interval = 60 seconds).

The Distributed Logical Router uses an XOR algorithm to determine the next hop from a list of possible ECMP next hops. This algorithm uses the source and destination IP address on the outgoing packet as sources of entropy.

Stateful services such as Load Balancing, VPN, NAT, and ESG firewall do not work with ECMP. However, from NSX 6.1.3 onwards, ECMP and Distributed Firewall can work together.

- 6 (Only for UDLR): To change the **Locale ID** on a universal distributed logical router, next to **Routing Configuration**, click **Edit** . Enter a locale ID and click **Save** or **OK**.

By default, the locale ID is set to the NSX Manager UUID. However, you can override the locale ID by enabling local egress at the time of creating the universal distributed logical router. Locale ID is used to selectively configure routes in a cross-vCenter NSX or multi-site environment. See [Cross-vCenter NSX Topologies](#) for more information.

The locale ID must be in UUID format. For example, XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX, where each X is replaced with a base 16 digit (0-F).

**7** To specify the default gateway, click **Edit** next to **Default Gateway**.

- a Select an interface from which the next hop towards the destination network can be reached.
- b Type the Gateway IP.
- c (Optional) Type the locale ID. Locale ID is available only on universal logical routers.
- d (Optional) Edit the MTU.
- e If prompted, type the **Admin Distance**.

Choose a value between 1 and 255. The admin distance is used to choose which route to use when there are multiple routes for a given network. The lower the admin distance, the higher the preference for the route.

**Table 9-3. Default Admin Distances**

| Route Source    | Default admin distance |
|-----------------|------------------------|
| Connected       | 0                      |
| Static          | 1                      |
| External BGP    | 20                     |
| OSPF Intra-Area | 30                     |
| OSPF Inter-Area | 110                    |
| Internal BGP    | 200                    |

- f (Optional) Type a Description for the default gateway.
- g Click **Save**.

**8** To configure dynamic routing, click **Edit** next to **Dynamic Routing Configuration**.

- a **Router ID** displays the first uplink IP address of the NSX Edge that pushes routes to the kernel for dynamic routing.
- b Do not enable any protocols here.
- c Select **Enable Logging** to save logging information and select the log level.

---

**Note** If you have IPSec VPN configured in your environment, you should not use dynamic routing.

---

**9** Click **Publish Changes**.

**What to do next**

To delete routing configuration, click **Reset**. This deletes all routing configurations (default, static, OSPF, and BGP configurations, as well as route redistribution).

# NSX Edge Configuration

Once you have installed a working NSX Edge (i.e. added one or more appliances and interfaces, and configured the default gateway, firewall policy, and high availability), you can begin using NSX Edge services.

## Working with Certificates

NSX Edge supports self-signed certificates, certificates signed by a Certification Authority (CA), and certificates generated and signed by a CA.

### Configure a CA Signed Certificate

You can generate a certificate signing request (CSR) and get it signed by a certification authority (CA). If you generate a CSR at the global level, it is available to all NSX Edges in your inventory.

#### Procedure

- 1 Do one of the following:
  - Generate a global certificate signing request for the NSX Manager.
    - 1 Log in to the NSX Manager virtual appliance.
    - 2 Click **Manage Appliance Settings**, and then click **SSL Certificates**.
    - 3 Click **Generate CSR**.
  - Generate a certificate signing request for an NSX Edge.
    - 1 Log in to the vSphere Web Client.
    - 2 Navigate to **Networking & Security > NSX Edges**.
    - 3 Double-click an NSX Edge.
    - 4 Click **Manage > Settings > Certificates**.
    - 5 Click **CSR Actions** or **Actions**, and then click **Generate CSR**.
- 2 Type your organization unit and name.
- 3 Type the locality, street, state, and country of your organization.
- 4 Select the encryption algorithm for communication between the hosts.

---

**Attention** SSL VPN-Plus only supports RSA certificates.

---

- 5 Edit the default key size, if necessary.
- 6 Type a description for the certificate.
- 7 Click **OK**.

The CSR is generated and displayed in the Certificates list.
- 8 Have an online Certification Authority sign this CSR.

**9** Do one of the following:

- Import certificate at the global level in the NSX Manager virtual appliance.
  - 1 Click the **Manage Appliance Settings**, and then click **SSL Certificates**.
  - 2 Click **Import**.
  - 3 In the **Import SSL Certificate** dialog box, click **Choose File**, and browse to the signed certificate file.
  - 4 Click **Import**.
- Import certificate for the NSX Edge.
  - 1 Copy the contents of the signed certificate that you received from the certification authority.
  - 2 In the vSphere Web Client, double-click the NSX Edge.
  - 3 Click **CSR Actions** or **Actions**, and then click **Import Certificate**.
  - 4 In the **Import Certificate** dialog box, paste the contents of the signed certificate.
  - 5 Click **OK**.

The CA-signed certificate appears in the certificates list.

## Add a CA Certificate

By adding a CA certificate, you can become an interim CA for your company. You then have the authority for signing your own certificates.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Navigate to **Manage > Settings > Certificates**.
- 5 Click **Add**, and then click **CA Certificate**.
- 6 Copy and paste the certificate contents in the **Certificate Contents** text box.
- 7 Enter a description for the CA certificate.
- 8 Click **Add** or **OK**.

You can now sign your own certificates.

## Add a Server Certificate

To add a server certificate, you paste the contents of the PEM certificate file and the private key contents of the server.

**Procedure**

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Navigate to **Manage > Settings > Certificates**.
- 5 Click **Add**, and then click **Certificate**.
- 6 In the **Certificates Contents** text box, paste the contents of the PEM certificate file.

Text must include "-----BEGIN xxx-----" and "-----END xxx-----". For example:

```
-----BEGIN CERTIFICATE-----
 Server cert
-----END CERTIFICATE-----
```

- 7 In the **Private Key** text box, paste the private key contents of the server.

Following is an example of the private key content:

```
-----BEGIN RSA PRIVATE KEY-----
XX
-----END RSA PRIVATE KEY-----
```

- 8 Enter the password of the private key file and reenter the password to confirm.
- 9 (Optional) Enter a description for the server certificate.
- 10 Click **Add** or **OK**.

**Add a Chained Certificate**

To add a server certificate that is chained with the intermediary and root CA certificates, you require a server certificate (PEM file), a private key for the server, an intermediate and a root certificate.

**Procedure**

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Navigate to **Manage > Settings > Certificates**.
- 5 Click **Add**, and then click **Certificate**.

- 6 In the **Certificates Contents** text box, paste the contents of the server cert.pem file, and then append the content of the intermediary certificates and the root certificate.

In the certificate chain, the order of certificates must be as follows:

- Server certificate
- Any number of intermediate CA certificates
- Root CA certificate

Each certificate must include the -----BEGIN CERTIFICATE----- and -----END CERTIFICATE----- lines, as shown in the following example:

```
-----BEGIN CERTIFICATE-----
 Server cert
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
 Intermediate cert
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
 Root cert
-----END CERTIFICATE-----
```

- 7 In the **Private Key** text box, paste the private key contents of the server.

Following is an example of the private key content:

```
-----BEGIN RSA PRIVATE KEY-----
XX
-----END RSA PRIVATE KEY-----
```

- 8 Enter the password for the private key of the server and reenter the password to confirm.
- 9 (Optional) Enter a description for the chained certificate.
- 10 Click **Add** or **OK**.

### Results

After the certificate is added, the server certificate that is chained with its intermediary certificates is displayed in the certificate details.

To view certificates details:

- In NSX 6.4.4 and later, in the Certificates table, click the text in the Issued To column. Certificate details are displayed in a pop-up window.
- In NSX 6.4.3 and earlier, select a certificate from the grid. The **Certificate Details** pane below the grid displays the details of the certificate.

## Configure a Self-Signed Certificate

You can create, install, and manage self-signed server certificates.

## Prerequisites

You must have a CA that can sign your certificate signing request (CSR).

## Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Navigate to **Manage > Settings > Certificates**.
- 5 Generate a certificate signing request (CSR) for an NSX Edge. For detailed information, see steps 1–7 in [Configure a CA Signed Certificate](#).
- 6 Make sure that the CSR you generated is selected.
- 7 Click **CSR Actions** or **Actions**, and then click **Self Sign Certificate**.
- 8 Type the number of days for which you want this self-signed certificate to be valid.
- 9 Click **OK**.

## Using Client Certificates

After generating a client certificate, you can distribute this certificate to your remote users, who can install the certificate on their web browser.

The main benefit of implementing client certificates is that the NSX Edge Load Balancer can ask the client for its client certificate, and validate it before forwarding its web requests to the backend servers. If a client certificate is revoked because it has been lost, or the client doesn't work in the company anymore, NSX Edge will validate the client certificate doesn't belong to the Certification Revocation List.

NSX Edge Client certificates are configured under Application Profile.

For more information on generating client certificates, refer to [Scenario: SSL Client and Server Authentication](#).

## Add a Certificate Revocation List

A Certificate Revocation List (CRL) is a list of subscribers and their status, which is provided and signed by Microsoft.

The list contains the following items:

- The revoked certificates and the reasons for revocation.
- The dates that the certificates are issued.
- The entities that issued the certificates.
- A proposed date for the next release.

When a potential user attempts to access a server, the server allows or denies access based on the CRL entry for that particular user.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Navigate to **Manage > Settings > Certificates**.
- 5 Click **Add**, and then click **CRL**.
- 6 In the **Certificate Contents** text box, paste the list.
- 7 (Optional) Enter a description.
- 8 Click **Add** or **OK**.

## FIPS Mode

When you enable the FIPS mode, any secure communication to or from the NSX Edge uses cryptographic algorithms or protocols that are allowed by United States Federal Information Processing Standards (FIPS). FIPS mode turns on the cipher suites that comply with FIPS.

If you configure components those are not FIPS compliant on a FIPS enabled edge, or if you enable FIPS on a edge which has ciphers or authentication mechanism that is not FIPS compliant, NSX Manager will fail the operation and provide a valid error message.

### Functionality Difference Between FIPS Mode And Non-FIPS Mode

| Component    | Functionality                           | FIPS Mode     | Non-FIPS Mode |
|--------------|-----------------------------------------|---------------|---------------|
| SSL VPN      | RADIUS Authentication                   | Not Available | Available     |
| SSL VPN      | RSA Authentication                      | Not Available | Available     |
| TLS Protocol | TLSv1.0                                 | Not Available | Available     |
| Routing      | OSPF, BGP - Password MD5 Authentication | Not Available | Available     |
| IPSec VPN    | PSK Authentication                      | Not Available | Available     |
| IPSec VPN    | DH2 and DH5 groups                      | Not Available | Available     |
| IPSec VPN    | DH14, DH15, and DH16 groups             | Available     | Available     |
| IPSec VPN    | AES-GCM Algorithm                       | Not Available | Available     |

## Change FIPS Mode on NSX Edge

Enabling the FIPS mode turns on the cipher suites that comply with FIPS. Thus, any secure communication to or from the NSX Edge uses cryptographic algorithms or protocols that are allowed by FIPS.

---

**Caution** Changing FIPS mode reboots the NSX Edge appliance causing temporary traffic disruption. This applies whether or not high availability is enabled.

---

Depending on your requirements, you can enable FIPS on some or all of your NSX Edge appliances. FIPS-enabled NSX Edge appliances can communicate with NSX Edge appliances that do not have FIPS enabled.

If a logical (distributed) router is deployed without an NSX Edge appliance, you cannot modify the FIPS mode. The logical router automatically gets the same FIPS mode as the NSX Controller cluster. If the NSX Controller cluster is NSX 6.3.0 or later, FIPS is enabled.

To change FIPS mode on a universal logical (distributed) router in a cross-vCenter NSX environment that has multiple NSX Edge appliances deployed in the primary and secondary NSX Managers, you must change FIPS mode on all the NSX Edge appliances associated with the universal logical (distributed) router on the primary NSX Manager.

If you change FIPS mode on an NSX Edge appliances with high availability enabled, FIPS will be enabled on both appliances, and the appliances will be rebooted one after the other.

If you want to change FIPS mode for a standalone edge, use the `fips enable` or `fips disable` command. For more information, refer to *NSX Command Line Interface Reference*.

### Prerequisites

- Verify that any partner solutions are FIPS mode certified. See the VMware Compatibility Guide at <http://www.vmware.com/resources/compatibility/search.php?deviceCategory=security>.
- If you have upgraded from an earlier version of NSX, do not enable FIPS mode until the upgrade to NSX 6.3.0 is complete. See Understand FIPS Mode and NSX Upgrade in the *NSX Upgrade Guide*.
- Verify that the NSX Manager is NSX 6.3.0 or later.
- Verify that the NSX Controller cluster is NSX 6.3.0 or later.
- Verify that all host clusters running NSX workloads are prepared with NSX 6.3.0 or later.
- Verify that all NSX Edge appliances on which you want to enable FIPS are version 6.3.0 or later.

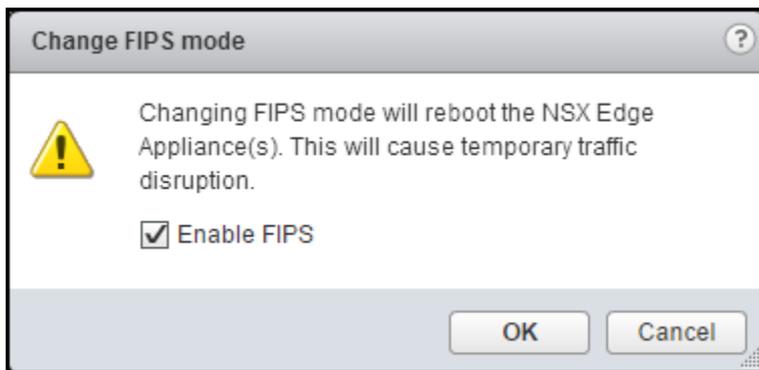
- Verify that the messaging infrastructure has status GREEN. Use the API method `GET /api/2.0/nwfabric/status?resource={resourceId}`, where `resourceId` is the vCenter Managed Object ID of a host or cluster. Look for the `status` corresponding to the `featureId` of `com.vmware.vshield.vsm.messagingInfra` in the response body:

```
<nwFabricFeatureStatus>
 <featureId>com.vmware.vshield.vsm.messagingInfra</featureId>
 <updateAvailable>>false</updateAvailable>
 <status>GREEN</status>
 <installed>>true</installed>
 <enabled>>true</enabled>
 <allowConfiguration>>false</allowConfiguration>
</nwFabricFeatureStatus>
```

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Select the required edge or router, click **Actions** (⚙️) and select **Change FIPS mode**.

The **Change FIPS mode** dialog box appears.



- 4 Select or deselect the **Enable FIPS** check box. Click **OK**.

The NSX Edge reboots, and FIPS mode is enabled.

### What to do next

Optionally, [Change FIPS Mode and TLS Settings on NSX Manager](#).

## Managing NSX Edge Appliances

You can add, edit, or delete NSX Edge appliances. An NSX Edge instance remains offline till at least one appliance has been added to it.

### Add an Edge Appliance

You must add at least one appliance to NSX Edge before deploying it.

**Procedure**

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Navigate to an NSX Edge Appliance.

Version	Procedure
NSX 6.4.4 and later	<ol style="list-style-type: none"> <li>a <b>Manage &gt; Settings &gt; Appliance Settings</b>.</li> <li>b Go to the <b>Edge Appliance VMs</b> section.</li> </ol>
NSX 6.4.3 and earlier	<ol style="list-style-type: none"> <li>a Click <b>Manage &gt; Settings &gt; Configuration</b>.</li> <li>b Go to the <b>NSX Edge Appliances</b> pane.</li> </ol>

- 5 Click **Add Edge Appliance VM** or the **Add (+)** icon.
- 6 Select the cluster or resource pool and datastore for the appliance.
- 7 (Optional) Select the host on which you want to add the appliance.
- 8 (Optional) Select the vCenter folder within which the appliance is to be added.
- 9 Click **Add**.

**Results**

- In NSX 6.4.4 or later, the NSX Edge Appliance details are displayed in a card view in the **Edge Appliance VMs** section. One card shows settings of one Edge Appliance VM.
- In NSX 6.4.3 or earlier, the NSX Edge Appliance details are displayed in a grid format in the **NSX Edge Appliances** pane.

**Edit an Edge Appliance**

You can edit a NSX Edge appliance.

**Procedure**

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Navigate to an NSX Edge Appliance.

Version	Procedure
NSX 6.4.4 and later	<ol style="list-style-type: none"> <li>a <b>Manage &gt; Settings &gt; Appliance Settings</b>.</li> <li>b Go to the <b>Edge Appliance VMs</b> section.</li> </ol>
NSX 6.4.3 and earlier	<ol style="list-style-type: none"> <li>a Click <b>Manage &gt; Settings &gt; Configuration</b>.</li> <li>b Go to the <b>NSX Edge Appliances</b> pane.</li> </ol>

## 5 Edit the NSX Edge Appliance settings.

Version	Procedure
NSX 6.4.4 and later	<ol style="list-style-type: none"> <li>In the <b>Edge Appliance VMs</b> section, go to the Edge Appliance VM that you want to edit.</li> <li>Click  and then click <b>Edit</b>.</li> <li>Make the appropriate changes and click <b>Save</b>.</li> </ol>
NSX 6.4.3 and earlier	<ol style="list-style-type: none"> <li>Select the appliance you want to edit, and click the <b>Edit</b> () icon.</li> <li>Make the appropriate changes and click <b>OK</b>.</li> </ol>

## Delete an Edge Appliance

You can delete an NSX Edge appliance.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Navigate to an NSX Edge Appliance.

Version	Procedure
NSX 6.4.4 and later	<ol style="list-style-type: none"> <li><b>Manage &gt; Settings &gt; Appliance Settings</b>.</li> <li>Go to the <b>Edge Appliance VMs</b> section.</li> </ol>
NSX 6.4.3 and earlier	<ol style="list-style-type: none"> <li>Click <b>Manage &gt; Settings &gt; Configuration</b>.</li> <li>Go to the <b>NSX Edge Appliances</b> pane.</li> </ol>

- 5 Delete an NSX Edge Appliance.
  - ◆ In NSX 6.4.4 and later, go to the NSX Edge Appliance, click , and then click **Delete**.
  - ◆ In NSX 6.4.3 and earlier, select an NSX Edge Appliance from the grid, and then click the **Delete** () icon.

## Managing NSX Edge Appliance Resource Reservations

NSX Data Center for vSphere uses vSphere resource allocation to reserve resources for NSX Edge appliances. Reserving CPU and memory resources for NSX Edge ensures that the appliance has enough resources to function correctly.

There are three methods of resource reservation: **System Managed**, **Custom**, or **No Reservation**

---

**Important** If you are using NSX 6.4.3 or earlier, and you select **Custom** or **No Reservation** reservations for an NSX Edge appliance, you cannot switch back to **System Managed**.

---

## System Managed Resource Reservation

If you select **System Managed**, the system reserves CPU and memory resources for the new NSX Edge appliance. The reserved resources are equal to the system requirements for the appliance size, modified by any percentages that are specified using the tuning configuration API.

When you install, upgrade, or redeploy an NSX Edge instance, the associated NSX Edge appliances are deployed. If an appliance has **System Managed** resource reservation, the reservation is applied on the resource pool after the appliance is powered on. If there are insufficient resources, the reservation fails and generates a system event, but the appliance deployment succeeds. The reservation is attempted the next time the appliance is deployed (during upgrade or redeploy).

With **System Managed** resource reservations, if you change the appliance size, the system updates the resource reservation to match the system requirements of the new appliance size.

## Custom Resource Reservation

If you select **Custom**, you determine the resource reservations for the NSX Edge appliance.

When you install, upgrade, or redeploy an NSX Edge, the associated NSX Edge appliances are deployed. If an appliance has **Custom** resource reservation, the reservation is applied on the resource pool before the appliance is powered on. If there are insufficient resources, the appliance fails to power on and the appliance deployment fails.

You can apply **Custom** reservations to an existing NSX Edge appliance. If the resource pool does not have sufficient resources, the configuration change fails.

With **Custom** resource reservations, the system does not manage resource reservations for the appliance. If you change the appliance size, the appliance system requirements change, but the system does not update the resource reservation. You should change the resource reservation to reflect the system requirements of the new appliance size.

## No Resource Reservation

If you select **No reservation**, no resources are reserved for the NSX Edge appliance. You can deploy NSX Edge appliances on hosts that do not have sufficient resources, but if there is a resource contention the appliances might not operate correctly.

## Configuring NSX Edge Appliance Resource Reservations

You set the resource reservation during the creation of an NSX Edge appliance. You can also update the reservation on an existing NSX Edge appliance. You can use the vSphere Web Client or the API for these tasks. See the *NSX API Guide* for more information about using the API.

Operation	vSphere Web Client	API
Create a new NSX Edge	Navigate to <b>Networking &amp; Security &gt; NSX Edges</b> and click <b>Add</b> . The wizard guides you through the steps of creating an NSX Edge. You can add an NSX Edge appliance in the <b>Configure Deployment</b> step. You select the reservation method from the <b>Resource Reservation</b> drop-down menu.	Use <code>POST /api/4.0/edges</code>
Update an existing NSX Edge	Navigate to <b>Networking &amp; Security &gt; NSX Edges &gt; NSX Edge Instance &gt; Manage &gt; Settings</b> and edit the appliance VM to select a different value for <b>Resource Reservation</b> .	Use <code>PUT /api/4.0/edges/{edgeId}/appliances</code>

Use the `cpuReservation > reservation` and `memoryReservation > reservation` parameters to configure the NSX Edge Appliance Resource Reservation using the API.

Resource Reservation Method	Values for Reservation Parameters
<b>System Managed</b>	Do not specify values for <code>cpuReservation &gt; reservation</code> and <code>memoryReservation &gt; reservation</code> .
<b>Custom</b>	Specify the values you want in <code>cpuReservation &gt; reservation</code> and <code>memoryReservation &gt; reservation</code> .
<b>No Reservation</b>	Set <code>cpuReservation &gt; reservation</code> and <code>memoryReservation &gt; reservation</code> to 0.

The system requirements for NSX Edge appliances depend on the appliance size: Compact, Large, Quad Large, or X-Large. These values are used for the default **System Managed** resource reservation.

Appliance Size	CPU Reservation	Memory Reservation
Compact	1000 MHz	512 MB
Large	2000 MHz	1 GB
Quad Large	4000 MHz	2 GB
X-Large	6000 MHz	8 GB

## Modifying the System Managed Resource Reservation Using Tuning Configuration

When facing lack of resources, you can temporarily disable the **System Managed** resource reservations or decrease the default value. You can change the reservation percentage by configuring values for the `edgeVCpuReservationPercentage` and `edgeMemoryReservationPercentage` parameters in the tuning configuration API, `PUT /api/4.0/`

`edgePublish/tuningConfiguration`. The default value for both parameters is 100. This change affects new NSX Edge appliance deployments, but not existing appliances. The percentages modify the default CPU and memory reserved for the relevant NSX Edge appliance size. To disable the resource reservation, set the values to 0. See the *NSX API Guide* for details.

## Changing Resource Reservation Method from Custom or No Reservation to System Managed

If you are using NSX 6.4.3 or earlier, and you select **Custom** or **No Reservation** reservations for an NSX Edge appliance, you cannot switch back to **System Managed** reservations.

Starting in NSX 6.4.4 you can use the API to switch back to **System Managed** reservations using `POST /api/4.0/edges/{edgeId}/appliances?action=applySystemResourceReservation`. See the *NSX API Guide* for details.

Starting in NSX 6.4.6, you can use the vSphere Web Client to edit the NSX Edge appliance VM and switch back to **System Managed** reservations.

## Working with Interfaces

An NSX Edge services gateway can have up to ten internal, uplink, or trunk interfaces. An NSX Edge router can have eight uplink interfaces and up to a thousand internal interfaces.

An NSX Edge must have at least one internal interface before it can be deployed.

### Configure an Interface

Internal interfaces are generally for East-West traffic, while uplink interfaces are for North-South traffic.

An NSX Edge Services Gateway (ESG) can have up to 10 internal, uplink, or trunk interfaces. These limits are enforced by the NSX Manager. When a logical router (DLR) is connected to an edge services gateway (ESG), the interface on the router is an uplink interface, while the interface on the ESG is an internal interface. An NSX trunk interface is for internal networks, not external networks. The trunk interface allows multiple internal networks (either VLAN or VXLAN) to be trunked.

An NSX Data Center deployment can have up to a 1,000 distributed logical router (DLR) instances on a single ESXi host. On a single logical router, you can configure up to eight uplink interfaces, and up to 991 internal interfaces. These limits are enforced by the NSX Manager. For more information about interface scaling in an NSX Data Center deployment, see the *NSX Network Virtualization Design Guide* at <https://communities.vmware.com/docs/DOC-27683>.

---

**Note** IPv6 multicast addresses are not supported on NSX ESG interfaces in NSX Data Center for vSphere 6.2.x, 6.3.x, and 6.4.x.

---

#### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.

- 3 Double-click an NSX Edge.
- 4 Navigate to NSX Edge interface settings by clicking **Manage > Settings > Interfaces**.
- 5 Select an interface and click the **Edit** (✎ or ✏) icon.
- 6 In the Edit Edge Interface dialog box, enter a name for the interface.
- 7 To indicate whether this interface is an internal or an external (uplink) interface, click **Internal** or **Uplink**.

Select **Trunk** when creating a sub interface. For more information, see [Add a Sub Interface](#).

- 8 Select the port group or logical switch to which you want to connect this interface to.
  - a Next to the **Connected To** text box, click ✎ or **Change**.
  - b Depending on what you want to connect to the interface, click the **Logical Switch**, **Standard Port Group**, or **Distributed Virtual Port Group** tab.
  - c Select the appropriate logical switch or port group, and click **OK**.
- 9 Select the connectivity status for the interface.

- 10 In **Configure Subnets**, click **Add** to add a subnet for the interface.

An interface can have multiple non-overlapping subnets. Enter one primary IP address and a comma-separated list of multiple secondary IP addresses. NSX Edge considers the primary IP address as the source address for locally generated traffic. You must add an IP address to an interface before using it on any feature configuration.

- 11 Enter the subnet prefix length or subnet mask for the interface.
- 12 If you are using NSX 6.4.4 or later, click the **Advanced** tab, and then continue with the remaining steps in this procedure. If you are using NSX 6.4.3 or earlier, go to the next step.
- 13 Change the default MTU, if necessary.
- 14 Under **Options**, specify the following options.

Option	Description
<b>Proxy ARP</b>	Supports overlapping network forwarding between different interfaces.
<b>Send ICMP Redirect</b>	Conveys routing information to hosts.
<b>Reverse Path Filter</b>	Verifies the reachability of the source address in packets being forwarded. In enabled mode, the packet must be received on the interface that the router might use to forward the return packet. In loose mode, the source address must appear in the routing table.

- 15 Enter the fence parameters.

Configure fence parameters if you want to reuse IP and MAC addresses across different fenced environments. For example, in a cloud management platform (CMP), fencing allows you to run several cloud instances simultaneously with the same IP and MAC addresses isolated or "fenced".

- 16 Click **Save** or **OK**.

## Delete an Interface

You can delete an NSX Edge interface.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Navigate to NSX Edge interface settings by clicking **Manage > Settings > Interfaces**.
- 5 Select an interface to delete.
- 6 Click the **Delete** (🗑️ or ❌) icon.

## Enable an Interface

An interface must be enabled or its status must be connected for an NSX Edge to isolate the virtual machines within that interface (port group or logical switch).

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Navigate to NSX Edge interface settings by clicking **Manage > Settings > Interfaces**.
- 5 Select an interface to connect.
- 6 Click the **Connect** (🔌 or ✅) icon.

## Disable an Interface

You can disable or disconnect an interface on an NSX Edge.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Navigate to NSX Edge interface settings by clicking **Manage > Settings > Interfaces**.
- 5 Select an interface to disable or disconnect.
- 6 Click the **Disconnect** (🔌 or 🚫) icon.

## Change Traffic Shaping Policy

You can change the traffic shaping policy on the vSphere Distributed Switch for an NSX Edge interface.

**Note** Starting with NSX Data Center 6.4.4, the terminology for some features in the UI has changed. The following table provides the list of modified terms.

**Table 9-4. Modified Terms**

NSX 6.4.3 or earlier	NSX 6.4.4 or later
Traffic Shaping Policy	Quality of Service (QoS)
In Shaping Policy	Ingress Shaping Policy
Out Shaping Policy	Egress Shaping Policy

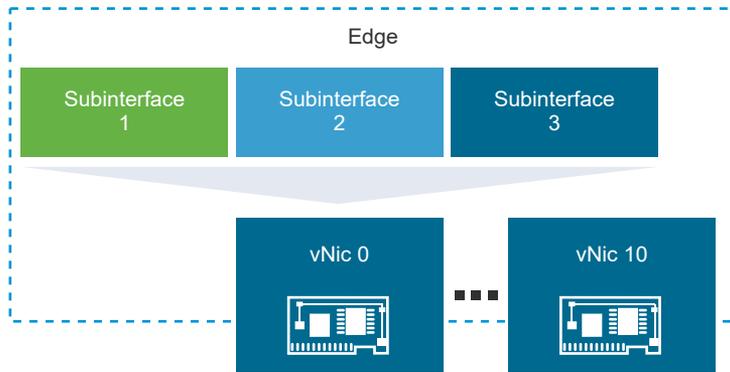
### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Navigate to NSX Edge interface settings by clicking **Manage > Settings > Interfaces**.
- 5 Select an interface for which you want to configure the quality of service.
- 6 Do one of the following:
  - In NSX 6.4.4 and later, click **Configure QoS**.
  - In NSX 6.4.3 and earlier, click **Actions > Configure Traffic Shaping Policy**.
- 7 Make the appropriate changes.
 

For more information about the traffic policy shaping options, see [Traffic Shaping Policy](#).
- 8 Click **Save** or **OK**.

## Add a Sub Interface

You can add a sub interface on a trunk vNIC, and use this sub interface in various NSX Edge services.



Trunk interfaces can be of the following types:

- VLAN trunk is standard and works with any version of ESXi. This type of interface is used to bring a tagged VLAN traffic into Edge.
- VXLAN trunk works with NSX version 6.1, and later. This type of interface is used to bring VXLAN traffic into Edge.

The following Edge services can use a sub interface:

- DHCP
- Routing (BGP and OSPF)
- Load Balancer
- IPsec VPN: You can configure IPsec VPN only as an uplink interface. Use sub interfaces when you want private traffic to traverse through the IPsec tunnel. If an IPsec policy is configured for private traffic, sub interface acts as a gateway for the private local subnet.
- L2 VPN
- NAT

. A sub interface cannot be used for HA or Logical Firewall. However, you can use the IP address of the sub interface in an edge firewall rule.

#### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Navigate to NSX Edge interface settings by clicking **Manage > Settings > Interfaces**.
- 5 Select an interface and click the **Edit** (🔗 or 🖋️) icon.
- 6 In the Edit Edge Interface dialog box, enter a name for the interface.
- 7 In Type, select **Trunk**.

- 8 Select the standard port group or distributed port group to which this interface must be connected.
  - a Next to the **Connected To** text box, click  or **Change**.
  - b Depending on what you want to connect to the interface, click the **Standard Port Group** or **Distributed Port Group** tab.
  - c Select the appropriate port group and click **OK**.

- 9 Select the connectivity status for the interface.

- 10 In Sub Interfaces, click **Add**.

- 11 Make sure that the sub interface is enabled, and enter a name for the sub interface.

- 12 In Tunnel ID, enter a number between 1 and 4094.

The tunnel ID is used to connect the networks that are being stretched. This value must be identical on both the client and server sites.

- 13 In Backing Type, select one of the following options to indicate the network backing for the sub interface.

Option	Description
<b>VLAN</b>	Enter the VLAN ID of the virtual LAN that your sub interface should use. VLAN IDs can range from 0 to 4094.
<b>Network</b>	Select the distributed port group or logical switch. NSX Manager extracts the VLAN ID and uses it for configuring the trunk.
<b>None</b>	Use this option to create a sub interface without specifying a network or VLAN ID. This sub interface is internal to an NSX Edge, and is used to route packets between a stretched network and an unstretched (untagged) network.

- 14 In Configure Subnets, click **Add** to add subnets to the sub interface.

- 15 Enter the IP address.

An interface can have multiple non-overlapping subnets. Enter one primary IP address and a comma-separated list of multiple secondary IP addresses. NSX Edge considers the primary IP address as the source address for locally generated traffic. You must add an IP address to an interface before using it on any feature configuration.

- 16 Enter the subnet prefix length.

- 17 Edit the default MTU value for the sub interface, if necessary.

The default MTU for a sub interface is 1500. The MTU for the sub interface should be equal to or less than the lowest MTU among all the trunk interfaces for the NSX Edge.

- 18 Enable the **Send Redirect** option to convey routing information to hosts.

- 19 Enable or disable the **Reverse Path Filter** option.

Reverse Path Filter verifies the reachability of the source address in packets being forwarded. In enabled mode, the packet must be received on the interface that the router can use to forward the return packet. In loose mode, the source address must appear in the routing table.

- 20 To return to the trunk interface settings, click **OK**.

- 21 If you are using NSX Data Center 6.4.4 or later, click the **Advanced** tab to continue with the remaining steps in this procedure.

- 22 Enter the MAC address for the interface, if needed. Enter two MAC addresses, if HA is enabled for the ESG.

If not needed, the MAC addresses are autogenerated.

- 23 Edit the default MTU of the trunk interface, if necessary.

The default MTU for a trunk interface is 1600, and the default MTU for a sub interface is 1500. The MTU for the trunk interface must be equal to or more than the MTU of the sub interface.

- 24 Click **Save** or **OK**.

### Results

You can now use the sub interface for the Edge services.

### What to do next

Configure a VLAN trunk if the sub interface added to a trunk vNic is backed by a standard port group. See [Configure VLAN Trunk](#) .

## Configure VLAN Trunk

When you add sub interfaces on the trunk vNic of an edge that is connected to a distributed portgroup, both VLAN trunk and VXLAN trunk are supported. When you add sub interfaces on the trunk vNic of an Edge that is connected to a standard portgroup, only VLAN trunk is supported.

### Prerequisites

Verify that a sub interface with a trunk vNic backed by standard portgroup is available. See [Add a Sub Interface](#).

### Procedure

- 1 Log in to the vCenter Web Client.
- 2 Click **Networking**.
- 3 Select the standard portgroup and click **Edit Settings**.
- 4 Click the **VLAN** tab.
- 5 In VLAN Type, select VLAN Trunking and type the VLAN IDs to be trunked.

- 6 Click **OK**.

## Change Auto Rule Configuration

If auto rule generation is enabled, NSX Edge adds firewall, NAT, and routing routes to enable control traffic to flow for these services. If auto rule generation is not enabled, you must manually add firewall, NAT, and routing configuration to allow control channel traffic for NSX Edge services such as Load Balancing, VPN, etc.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Select an NSX Edge.
- 4 Click **Actions > Change Auto Rule Configuration**.
- 5 Make the appropriate changes and click **OK**.

## Change CLI Credentials

You can edit the credentials to be used for logging in to the NSX Edge Command Line Interface (CLI).

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Select an NSX Edge.
- 4 Click **Actions > Change CLI Credentials**.
- 5 Enter and confirm the new password and click **OK**.

## About High Availability

High Availability (HA) ensures that the services provided by NSX Edge appliances are available even when a hardware or software failure renders a single appliance unavailable. NSX Edge HA minimizes failover downtime instead of delivering zero downtime, as the failover between appliances might require some services to be restarted.

For example, NSX Edge HA synchronizes the connection tracker of the stateful firewall, or the stateful information held by the load balancer. The time required to bring all services backup is not null. Examples of known service restart impacts include a non-zero downtime with dynamic routing when an NSX Edge is operating as a router.

Sometimes, the two NSX Edge HA appliances are unable to communicate and unilaterally decide to become active. This behavior is expected to maintain availability of the active NSX Edge services if the standby NSX Edge is unavailable. If the other appliance still exists, when the communication is re-established, the two NSX Edge HA appliances renegotiate active and standby status. If this negotiation does not finish and if both appliances declare they are active when the connectivity is re-established, an unexpected behavior is observed. This condition, known as split brain, is observed due to the following environmental conditions:

- Physical network connectivity issues, including a network partition.
- CPU or memory contention on the NSX Edge.
- Transient storage problems that might cause at least one NSX Edge HA VM to become unavailable.

For example, an improvement in NSX Edge HA stability and performance is observed when the VMs are moved off overprovisioned storage. In particular, during large overnight backups, large spikes in storage latency can impact NSX Edge HA stability.

- Congestion on the physical or virtual network adapter involved with the exchange of packets.

In addition to environmental issues, a split-brain condition is observed when the HA configuration engine falls into a bad state or when the HA daemon fails.

## Stateful High Availability

The primary NSX Edge appliance is in the active state and the secondary appliance is in the standby state. NSX Manager replicates the configuration of the primary appliance for the standby appliance or you can manually add two appliances. Create the primary and secondary appliances on separate resource pools and datastores. If you create the primary and secondary appliances on the same datastore, the datastore must be shared across all hosts in the cluster for the HA appliance pair to be deployed on different ESXi hosts. If the datastore is local storage, both virtual machines are deployed on the same host.

All NSX Edge services run on the active appliance. The primary appliance maintains a heartbeat with the standby appliance and sends service updates through an internal interface.

If a heartbeat is not received from the primary appliance within the specified time (default value is 15 seconds), the primary appliance is declared dead. The standby appliance moves to the active state, takes over the interface configuration of the primary appliance, and starts the NSX Edge services that were running on the primary appliance. When the switch over takes place, a system event is displayed in the **System Events** tab of Settings & Reports. Load Balancer and VPN services need to re-establish TCP connection with NSX Edge, so service is disrupted for a short while. Logical switch connections and firewall sessions are synched between the primary and standby appliances however, service is disrupted during the switch over while waiting for the standby appliance to become active and take over.

If the NSX Edge appliance fails and a bad state is reported, HA force syncs the failed appliance to revive it. When revived, it takes on the configuration of the now-active appliance and stays in a standby state. If the NSX Edge appliance is dead, you must delete the appliance and add a new one.

NSX Edge ensures that the two HA NSX Edge virtual machines are not on the same ESXi host even after you use DRS and vMotion (unless you manually vMotion them to the same host). Two virtual machines are deployed on vCenter in the same resource pool and datastore as the appliance you configured. Local link IPs are assigned to HA virtual machines in the NSX Edge HA so that they can communicate. You can specify management IP addresses to override the local links.

If syslog servers are configured, logs in the active appliance are sent to the syslog servers.

## High Availability in a Cross-vCenter NSX Environment

If you enable high availability on an NSX Edge in a cross-vCenter NSX environment, both the active and standby NSX Edge Appliances must reside in the same vCenter Server. If you migrate one of the appliances of an NSX Edge HA pair to a different vCenter Server, the two HA appliances no longer operate as an HA pair, and you might experience traffic disruption.

### vSphere High Availability

NSX Edge HA is compatible with vSphere HA. If the host on which a NSX Edge instance is running dies, the NSX Edge is restarted on the standby host, ensuring the NSX Edge HA pair is still available to take another failover.

If vSphere HA is not enabled, the active-standby NSX Edge HA pair will survive one fail-over. However, if another fail-over happens before the second HA pair was restored, NSX Edge availability can be compromised.

For more information on vSphere HA, see *vSphere Availability*.

## Change High Availability Configuration

You can change the HA configuration that you had specified while installing NSX Edge.

---

**Note** In NSX 6.2.3 and later, enabling high availability (HA) on an existing Edge will fail when sufficient resources cannot be reserved for the second Edge Appliance VM. The configuration will roll back to the last known good configuration.

---

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.

#### 4 Navigate to Edge HA configuration settings.

Version	Procedure
NSX 6.4.4 and later	<ol style="list-style-type: none"> <li>a Click <b>Manage &gt; Settings &gt; High Availability</b>.</li> <li>b Do these steps. <ul style="list-style-type: none"> <li>■ To edit HA configuration settings, next to High Availability Configuration, click <b>Edit</b>.</li> <li>■ To edit Management HA interface settings for a DLR appliance, next to Management/HA Interface, click <b>Edit</b>.</li> </ul> </li> </ol>
NSX 6.4.3 and earlier	<ol style="list-style-type: none"> <li>a Click <b>Manage &gt; Settings &gt; Configuration</b>.</li> <li>b Do these steps: <ul style="list-style-type: none"> <li>■ To edit HA configuration settings, go to the <b>HA Configuration</b> pane, and click <b>Change</b>.</li> <li>■ To edit Management HA interface settings for a DLR appliance, go to the <b>HA Interface Configuration</b> pane, and click <b>Change</b>.</li> </ul> </li> </ol>

#### 5 Change the HA configuration settings. See the following tables for a description of all HA configuration options.

**Table 9-5. Common HA Configuration Options**

Option	Description
Declare dead time	Enter the period in seconds within which, if the backup appliance does not receive a heartbeat signal from the primary appliance, the primary appliance is considered inactive and the backup appliance takes over. The default interval is 15 seconds.
Logging	Enable or disable logging on the appliance.
Log Level	Select the level of logging information you want to collect for the appliance.

**Table 9-6. HA Configuration Options For NSX Edge Services Gateway Appliance**

Option	Description
vNIC	Select the internal interface for which you want to configure HA parameters.  If you select ANY for interface but there are no internal interfaces configured, the UI displays an error. Two Edge appliances are created but since there is no internal interface configured, the new NSX Edge remains in standby and HA is disabled. After an internal interface is configured, HA is enabled on the NSX Edge appliance.
Management IPs	Optional: You can enter two management IP addresses in CIDR format to override the local link IP addresses assigned to the HA virtual machines. Ensure that the management IP addresses do not overlap with the IP addresses used for any other interface and do not interfere with traffic routing. Do not use an IP address that exists somewhere else on your network, even if that network is not directly attached to the appliance.

**Table 9-7. HA Configuration Options For DLR Appliance**

Option	Description
Connected to	Connect the HA interface to a distributed port group or logical switch. If you are using this interface as an HA interface only, use a logical switch. A /30 subnet is allocated from the link local range 169.254.0.0/16 and is used to provide an IP address for each of the two NSX Edge appliances
IP Address (available in NSX Data Center for vSphere 6.4.3 or earlier)	Optional: To use the HA interface to connect to the NSX Edge, you can specify an additional IP address and prefix for the HA interface.

**Note** If you configure L2 VPN on this Edge appliance before HA is enabled, you must have at least two internal interfaces set up. If there is a single interface configured on this Edge which is already being used by L2 VPN, HA is disabled on the Edge appliance.

6 Click **Save** or **OK**.

## Force Sync NSX Edge with NSX Manager

You can send a synchronization request from an NSX Manager to an NSX Edge.

Run force sync when you want to synchronize the edge configuration as known to the NSX Manager to all the components.

**Note** For NSX Data Center 6.2 or later, force sync avoids data loss for east-west routing traffic, however, north-south routing and bridging might experience an interruption.

In NSX 6.4.3 or earlier, NSX Manager takes the following actions during the force sync operation:

- Deletes the configuration of the Edge appliances, starting with Index 0, and then Index 1.
- Reboots the Edge appliances. Both Index 0 and Index 1 are rebooted simultaneously. This action results in high downtime.
- Publishes or applies the latest configuration to the Edge appliances.
- Closes the connection to the host.
- If the NSX Manager is primary or stand-alone, and the edge is a logical distributed router, the controller cluster is synced.
- Sends a message to all relevant hosts to sync the distributed router instance.

Starting with NSX 6.4.4, NSX Manager takes the following actions during the force sync operation:

- If the Edge appliances are in a Bad state, then NSX Manager deletes the Edge configuration, reboots the bad Edge appliances, and publishes the latest configuration to the Edge appliances.
- If the Edge appliances are not in a Bad state, then NSX Manager does not reboot the Edge appliances, and directly publishes the latest configuration to the Edge appliances. By eliminating unnecessary reboot of the Edge appliances, downtime is reduced.
- Closes the connection to the host.
- If the NSX Manager is primary or stand-alone, and the edge is a logical distributed router, the controller cluster is synced.
- Sends a message to all relevant hosts to sync the distributed router instance.

---

**Important** In a cross-vCenter NSX environment, you must first run force sync on an NSX Edge instance on the primary NSX Manager. When that is complete, force sync the NSX Edge instance on the secondary NSX Managers.

---

#### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Select an NSX Edge instance.
- 4 Click **Actions** (⚙️) and select **Force Sync**.

## Configure Syslog Servers for NSX Edge

You can configure one or two remote syslog servers. NSX Edge events and logs related to firewall events that flow from NSX Edge appliances are sent to the syslog servers.

#### Procedure

- 1 Log in to the vSphere Web Client.

- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Navigate to configure syslog server settings.

Version	Procedure
NSX 6.4.4 and later	<ol style="list-style-type: none"> <li>a Click <b>Manage &gt; Settings &gt; Appliance Settings</b>.</li> <li>b Next to Configuration, click , and then click <b>Change Syslog Configuration</b>.</li> </ol>
NSX 6.4.3 and earlier	<ol style="list-style-type: none"> <li>a Click <b>Manage &gt; Settings &gt; Configuration</b>.</li> <li>b In the <b>Details</b> pane, next to Syslog servers, click <b>Change</b>.</li> </ol>

- 5 Type an IP address for both remote syslog servers.
- 6 Select a protocol, and click **OK**.

## View the Status of NSX Edge Services

You can view the status of all services on the NSX Edge from a single location, and refresh the window to check the status at any time.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Navigate to view the status of all Edge services.

Version	Procedure
NSX 6.4.4 and later	Click <b>Manage &gt; Settings &gt; Services</b> .
NSX 6.4.3 and earlier	<ol style="list-style-type: none"> <li>a Click <b>Manage &gt; Settings &gt; Configuration</b>.</li> <li>b View the <b>Services</b> pane for the status of all Edge services.</li> </ol>

## Redeploy NSX Edge

If NSX Edge services do not work as expected after a force sync, you can redeploy the NSX Edge instance.

**Note** Redeploying is a disruptive action. First apply a force sync and check whether the problem is fixed. It is a good practice to download the tech support bundle for the Edge and troubleshoot the issue. If the problem is still not fixed, then redeploy.

Redeploying an NSX Edge instance results in the following actions:

- Edge appliances are deleted and freshly deployed with the latest configuration applied.

- Logical routers are deleted from the controller and then recreated with the latest configuration applied.
- Distributed logical router instances on hosts are deleted and then recreated with the latest configuration applied.

OSPF adjacencies are withdrawn during redeploy if graceful restart is not enabled.

The following good practices can help in preventing traffic loss when redeploying edges:

- Enable graceful restart when OSPF or BGP timers are large and high availability (HA) is enabled on both distributed logical routers (DLR) and edge services gateways (ESG).
- Use aggressive OSPF or BGP timer values and floating static routes when a DLR in HA is peered with multiple ESGs (ECMP).

---

**Important** In a cross-vCenter NSX environment, you must first redeploy the NSX Edge instance on the primary NSX Manager. After that is complete, redeploy the NSX Edge instance on the secondary NSX Managers. It is required that the NSX Edge instances on both the primary and the secondary NSX Managers are redeployed.

---

#### Prerequisites

- Verify that the hosts have enough resources to deploy additional NSX Edge Services Gateway appliances during the redeploy operation. See the [Chapter 1 System Requirements for NSX Data Center for vSphere](#) for the resources required for each NSX Edge size.
  - For a single NSX Edge instance, there are two NSX Edge appliances of the appropriate size in the poweredOn state during redeploy.
  - For an NSX Edge instance with high availability enabled, both replacement appliances are deployed before replacing the old appliances. This means that there are four NSX Edge appliances of the appropriate size in the poweredOn state during the upgrade of a given NSX Edge. After the NSX Edge instance is redeployed, either of the HA appliances can become active.
- Verify that the host clusters listed in the configured location and live location for the NSX Edge appliances you redeploy are prepared for NSX, and that their messaging infrastructure status is GREEN.

Verify that the host clusters listed in the configured location and live location for all NSX Edge appliances are prepared for NSX and that their messaging infrastructure status is GREEN. If the status is green, the hosts are using the messaging infrastructure to communicate with NSX Manager instead of VIX.

If the configured location is not available, for example, because the cluster has been removed since the NSX Edge appliance was created, then verify the live location only.

- Find the ID of the original configured location (*configuredResourcePool > id*) and the current live location (*resourcePoolId*) with the GET `https://NSX-Manager-IP-Address/api/4.0/edges/{edgeId}/appliances` API request.

- Find the host preparation status and the messaging infrastructure status for those clusters with the GET `https://NSX-Manager-IP-Address/api/2.0/nwfabric/status?resource={resourceId}` API request, where *resourceId* is the ID of the configured and live location of the NSX Edge appliances found previously.
- Look for the status corresponding to the *featureId* of `com.vmware.vshield.vsm.nwfabric.hostPrep` in the response body. The status must be GREEN.

```
<nwFabricFeatureStatus>
 <featureId>com.vmware.vshield.vsm.nwfabric.hostPrep</featureId>
 <featureVersion>6.3.1.5124716</featureVersion>
 <updateAvailable>>false</updateAvailable>
 <status>GREEN</status>
 <installed>>true</installed>
 <enabled>>true</enabled>
 <allowConfiguration>>false</allowConfiguration>
</nwFabricFeatureStatus>
```

- Look for the status corresponding to the *featureId* of `com.vmware.vshield.vsm.messagingInfra` in the response body. The status must be GREEN.

```
<nwFabricFeatureStatus>
 <featureId>com.vmware.vshield.vsm.messagingInfra</featureId>
 <updateAvailable>>false</updateAvailable>
 <status>GREEN</status>
 <installed>>true</installed>
 <enabled>>true</enabled>
 <allowConfiguration>>false</allowConfiguration>
</nwFabricFeatureStatus>
```

If the hosts are not prepared for NSX, do the following:

- Navigate to **Installation and Upgrade > Host Preparation** and prepare the hosts for NSX.
- Verify that the messaging infrastructure is GREEN.
- Redeploy the NSX Edges on the host.

#### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Select an NSX Edge instance.
- 4 Click **Actions > Redeploy**.

It is a good practice to download the tech support bundle for the Edge and troubleshoot the problem. If the problem persists, redeploy the Edge.

## Results

The NSX Edge virtual machine is replaced with a new virtual machine and all services are restored. If redeploy does not work, power off the NSX Edge virtual machine and redeploy NSX Edge again.

---

**Note** Redeploy might not work in the following cases.

- The resource pool on which the NSX Edge was installed is no longer in the vCenter inventory or its Managed Object ID (Moid) has changed.
- The datastore on which the NSX Edge was installed is corrupted/unmounted or in-accessible.
- The dvportGroups on which the NSX Edge interfaces were connected are no longer in the vCenter inventory or their Moid (identifier in vCenter Server) has changed.

If any of these cases is true, you must update the Moid of the resource pool, datastore, or dvPortGroup using a REST API call. See *NSX API Guide*.

---

If FIPS mode is enabled on NSX Edge and something goes wrong, NSX Manager does not allow you to redeploy the NSX Edge. You must resolve infrastructure issues for communication failures instead of redeploying the edge.

## Download Tech Support Logs for NSX Edge

You can download technical support logs for each NSX Edge instance. If high availability is enabled for the NSX Edge instance, support logs from both NSX Edge virtual machines are downloaded. You can also collect the support bundle data for NSX Edge using the **Support Bundle** collection tool. For details, refer to *NSX Administration Guide*.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Select an NSX Edge instance.
- 4 Click **Actions > Download Tech Support Logs**.
- 5 After the tech support logs are generated, click **Download**.

## Add a Static Route

You can add a static route for a destination subnet or host.

### Procedure

- 1 In the vSphere Client, navigate to **Networking & Security > NSX Edges**.
- 2 Click an NSX Edge.
- 3 Click **Routing > Static Routes**.
- 4 Click **Add**.

5 Enter the **Network** in CIDR notation.

6 Enter the IP address of the **Next Hop**.

The router must be able to reach the next hop directly. If ECMP is enabled, you can enter multiple next hops as a comma-separated list of IP addresses.

- In NSX 6.4.4 or earlier, next hop is mandatory. Starting in NSX 6.4.5, next hop is optional for ESG. You can specify either the next hop or the interface. When you specify the next hop, interface is unavailable for selection, and conversely.
- When multicast traffic is sent through a GRE tunnel interface on the ESG, you must specify the IP address of the remote GRE tunnel endpoint in the next hop when configuring the static routes.
- For DLR and UDLR, next hop is mandatory.

7 Select the **Interface** on which you want to add a static route.

The **Interface** drop-down menu does not display the GRE tunnel interfaces.

8 For **MTU**, edit the maximum transmission value for the data packets if necessary.

The MTU cannot be higher than the MTU set on the NSX Edge interface.

9 If prompted, enter the **Admin Distance**.

Choose a value between 1 and 255. The admin distance is used to choose which route to use when there are multiple routes for a given network. The lower the admin distance, the higher the preference for the route.

**Table 9-8. Default Admin Distances**

Route Source	Default admin distance
Connected	0
Static	1
External BGP	20
OSPF Intra-Area	30
OSPF Inter-Area	110
Internal BGP	200

An administrative distance of 255 causes the static route to be excluded from the routing table (RIB) and the data plane, so the route is not used.

10 (Optional) Enter the **Locale ID**.

By default, routes have the same locale ID as the NSX Manager. The locale ID specified here associates the route with this locale ID. These routes are sent only to those hosts that have a matching locale ID. See [Cross-vCenter NSX Topologies](#) for more information.

11 (Optional) Enter a **Description** for the static route.

12 Click **OK**.

## Configure OSPF on a Logical (Distributed) Router

Configuring OSPF on a logical router enables VM connectivity across logical routers and from logical routers to edge services gateways (ESGs).

OSPF routing policies provide a dynamic process of traffic load balancing between routes of equal cost.

An OSPF network is divided into routing areas to optimize traffic flow and limit the size of routing tables. An area is a logical collection of OSPF networks, routers, and links that have the same area identification.

Areas are identified by an Area ID.

### Prerequisites

A Router ID must be configured, as shown in [OSPF Configured on the Logical \(Distributed\) Router](#).

When you enable a router ID, the text box is populated by default with the uplink interface of the logical router.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click a logical router.
- 4 Click **Manage > Routing > OSPF**.
- 5 Enable OSPF.
  - a Next to **OSPF Configuration**, click **Edit**, and then click **Enable OSPF**
  - b In **Forwarding Address**, type an IP address that is to be used by the router datapath module in the hosts to forward datapath packets.
  - c In **Protocol Address**, type a unique IP address within the same subnet as the **Forwarding Address**. The protocol address is used by the protocol to form adjacencies with the peers.
  - d (Optional) Enable **Graceful Restart** for packet forwarding to be uninterrupted during restart of OSPF Services.
- 6 Configure the OSPF areas.
  - a (Optional) Delete the not-so-stubby area (NSSA) 51 that is configured by default.
  - b In **Area Definitions**, click **Add**.

- c Type an Area ID. NSX Edge supports an area ID in the form of a decimal number. Valid values are 0–4294967295.
- d In **Type**, select **Normal** or **NSSA**.

NSSAs prevent the flooding of AS-external link-state advertisements (LSAs) into NSSAs. They rely on default routing to external destinations. Hence, NSSAs must be placed at the edge of an OSPF routing domain. NSSA can import external routes into the OSPF routing domain, thereby providing transit service to small routing domains that are not part of the OSPF routing domain.

- 7 (Optional) Select the type of **Authentication**. OSPF performs authentication at the area level. All routers within the area must have the same authentication and corresponding password configured. For MD5 authentication to work, both the receiving and transmitting routers must have the same MD5 key.
  - a **None**: No authentication is required, which is the default value.
  - b **Password**: In this method of authentication, a password is included in the transmitted packet.
  - c **MD5**: This authentication method uses MD5 (Message Digest type 5 ) encryption. An MD5 checksum is included in the transmitted packet.
  - d For **Password** or **MD5** type authentication, type the password or MD5 key.

---

#### Important

- If NSX Edge is configured for HA with OSPF graceful restart enabled and MD5 is used for authentication, OSPF fails to restart gracefully. Adjacencies are formed only after the grace period expires on the OSPF helper nodes.
  - You cannot configure **MD5** authentication when FIPS mode is enabled.
  - NSX Data Center for vSphere always uses a key ID value of 1. Any device not managed by NSX Data Center for vSphere that peers with an Edge Services Gateway or Logical Distributed Router must be configured to use a key ID of value 1 when MD5 authentication is used. Otherwise an OSPF session cannot be established.
- 

- 8 Map interfaces to the areas.
  - a In **Area to Interface Mapping**, click **Add** to map the interface that belongs to the OSPF area.
  - b Select the interface that you want to map and the OSPF area that you want to map it to.
- 9 (Optional) Edit the default OSPF settings.

In most cases, it is recommended to retain the default OSPF settings. If you do change the settings, make sure that the OSPF peers use the same settings.

- a **Hello Interval** displays the default interval between hello packets that are sent on the interface.

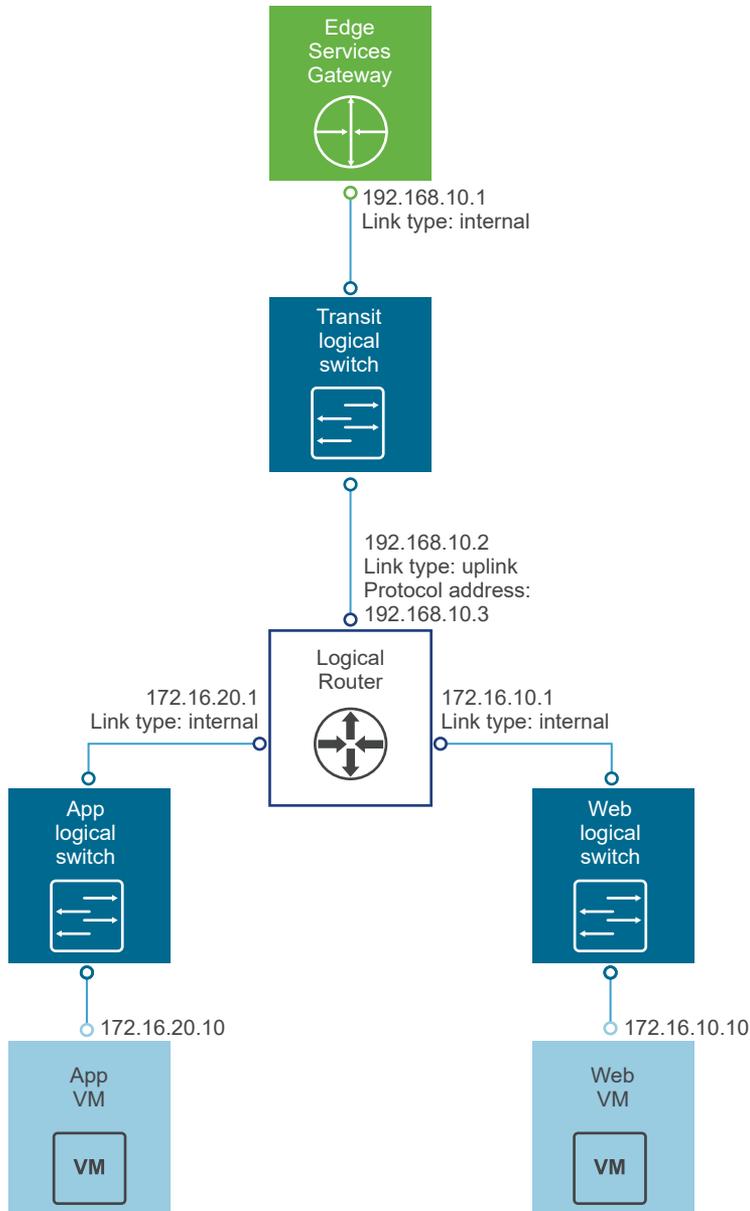
- b **Dead Interval** displays the default interval during which at least one hello packet must be received from a neighbor before the router declares that neighbor down.
- c **Priority** displays the default priority of the interface. The interface with the highest priority is the designated router.
- d **Cost** of an interface displays the default overhead required to send packets across that interface. The cost of an interface is inversely proportional to the bandwidth of that interface. The larger the bandwidth, the smaller the cost.

10 Click **Publish Changes**.

## Example: OSPF Configured on the Logical (Distributed) Router

One simple NSX scenario that uses OSPF is when a logical router (DLR) and an edge services gateway (ESG) are OSPF neighbors, as shown here.

Figure 9-1. NSX Data Center for vSphere Topology



On the **Global Configuration** page, the configuration settings are as follows:

- **Gateway IP:** 192.168.10.1. The logical router's default gateway is the ESG's internal interface IP address (192.168.10.1).
- **Router ID:** 192.168.10.2. The router ID is the uplink interface of the logical router. In other words, the IP address that faces the ESG.

On the **OSPF Configuration** page, the configuration settings are as follows:

- **Forwarding Address:** 192.168.10.2
- **Protocol Address:** 192.168.10.3. The protocol address can be any IP address that is in the same subnet and is not used anywhere else. In this case, 192.168.10.3 is configured.

- **Area Definition:**

- Area ID: 0
- Type: Normal
- Authentication: None

The uplink interface (the interface facing the ESG) is mapped to the area, as follows:

- Interface: To-ESG
- Area ID: 0
- Hello Interval (seconds): 10
- Dear Interval (seconds): 40
- Priority: 128
- Cost: 1

#### What to do next

Make sure the route redistribution and firewall configuration allow the correct routes to be advertised.

In this example, the logical router's connected routes (172.16.10.0/24 and 172.16.20.0/24) are advertised into OSPF. To verify the redistributed routes, on the left navigation panel, click **Route Redistribution**, and check the following settings:

- **Route Redistribution Status** shows that OSPF is enabled.
- **Route Redistribution Table** shows the following:
  - Learner: OSPF
  - From: Connected
  - Prefix: Any
  - Action: Permit

If you enabled SSH when you created the logical router, you must also configure a firewall filter that allows SSH to the logical router's protocol address. For example, you can create a firewall filter rule with the following settings:

- Name: ssh
- Type: User
- Source: Any
- Destination: Protocol address with value: 192.168.10.3
- Service: SSH

## Configure OSPF on an Edge Services Gateway

Configuring OSPF on an edge services gateway (ESG) enables the ESG to learn and advertise routes. The most common application of OSPF on an ESG is on the link between the ESG and a Logical (Distributed) Router. This allows the ESG to learn about the logical interfaces (LIFS) that are connected to the logical router. This goal can be accomplished with OSPF, IS-IS, BGP or static routing.

OSPF routing policies provide a dynamic process of traffic load balancing between routes of equal cost.

An OSPF network is divided into routing areas to optimize traffic flow and limit the size of routing tables. An area is a logical collection of OSPF networks, routers, and links that have the same area identification.

Areas are identified by an Area ID.

### Prerequisites

A Router ID must be configured, as shown in [OSPF Configured on the Edge Services Gateway](#).

When you enable a router ID, the text box is populated by default with the ESG's uplink interface IP address.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an ESG.
- 4 Click **Manage > Routing > OSPF**.
- 5 Enable OSPF.
  - a Next to **OSPF Configuration**, click **Edit**, and then click **Enable OSPF**
  - b (Optional) Click **Enable Graceful Restart** for packet forwarding to be uninterrupted during restart of OSPF services.
  - c (Optional) Click **Enable Default Originate** to allow the ESG to advertise itself as a default gateway to its peers.
- 6 Configure the OSPF areas.
  - a (Optional) Delete the not-so-stubby area (NSSA) 51 that is configured by default.
  - b In **Area Definitions**, click **Add**.

c Enter an Area ID. NSX Edge supports an area ID in the form of a decimal number. Valid values are 0–4294967295.

d In **Type**, select **Normal** or **NSSA**.

NSSAs prevent the flooding of AS-external link-state advertisements (LSAs) into NSSAs. They rely on default routing to external destinations. Hence, NSSAs must be placed at the edge of an OSPF routing domain. NSSA can import external routes into the OSPF routing domain, providing transit service to small routing domains that are not part of the OSPF routing domain.

7 (Optional) When you select the NSSA area type, the **NSSA Translator Role** appears. Select the **Always** check box to translate Type-7 LSAs to Type-5 LSAs. All Type-7 LSAs are translated into Type-5 LSAs by the NSSA.

8 (Optional) Select the type of **Authentication**. OSPF performs authentication at the area level.

All routers within the area must have the same authentication and corresponding password configured. For MD5 authentication to work, both the receiving and transmitting routers must have the same MD5 key.

a **None**: No authentication is required, which is the default value.

b **Password**: In this method of authentication, a password is included in the transmitted packet.

c **MD5**: This authentication method uses MD5 (Message Digest type 5 ) encryption. An MD5 checksum is included in the transmitted packet.

d For **Password** or **MD5** type authentication, enter the password or MD5 key.

---

### Important

- If NSX Edge is configured for HA with OSPF graceful restart enabled and MD5 is used for authentication, OSPF fails to restart gracefully. Adjacencies are formed only after the grace period expires on the OSPF helper nodes.
- You cannot configure **MD5** authentication when FIPS mode is enabled.
- NSX Data Center for vSphere always uses a key ID value of 1. Any device not managed by NSX Data Center for vSphere that peers with an Edge Services Gateway or Logical Distributed Router must be configured to use a key ID of value 1 when MD5 authentication is used. Otherwise an OSPF session cannot be established.

---

9 Map interfaces to the areas.

a In **Area to Interface Mapping**, click **Add** to map the interface that belongs to the OSPF area.

b Select the interface that you want to map and the OSPF area that you want to map it to.

**10** (Optional) Edit the default OSPF settings.

In most cases, it is preferable to retain the default OSPF settings. If you do change the settings, make sure that the OSPF peers use the same settings.

- a **Hello Interval** displays the default interval between hello packets that are sent on the interface.
- b **Dead Interval** displays the default interval during which at least one hello packet must be received from a neighbor before the router declares that neighbor down.
- c **Priority** displays the default priority of the interface. The interface with the highest priority is the designated router.
- d **Cost** of an interface displays the default overhead required to send packets across that interface. The cost of an interface is inversely proportional to the bandwidth of that interface. The larger the bandwidth, the smaller the cost.

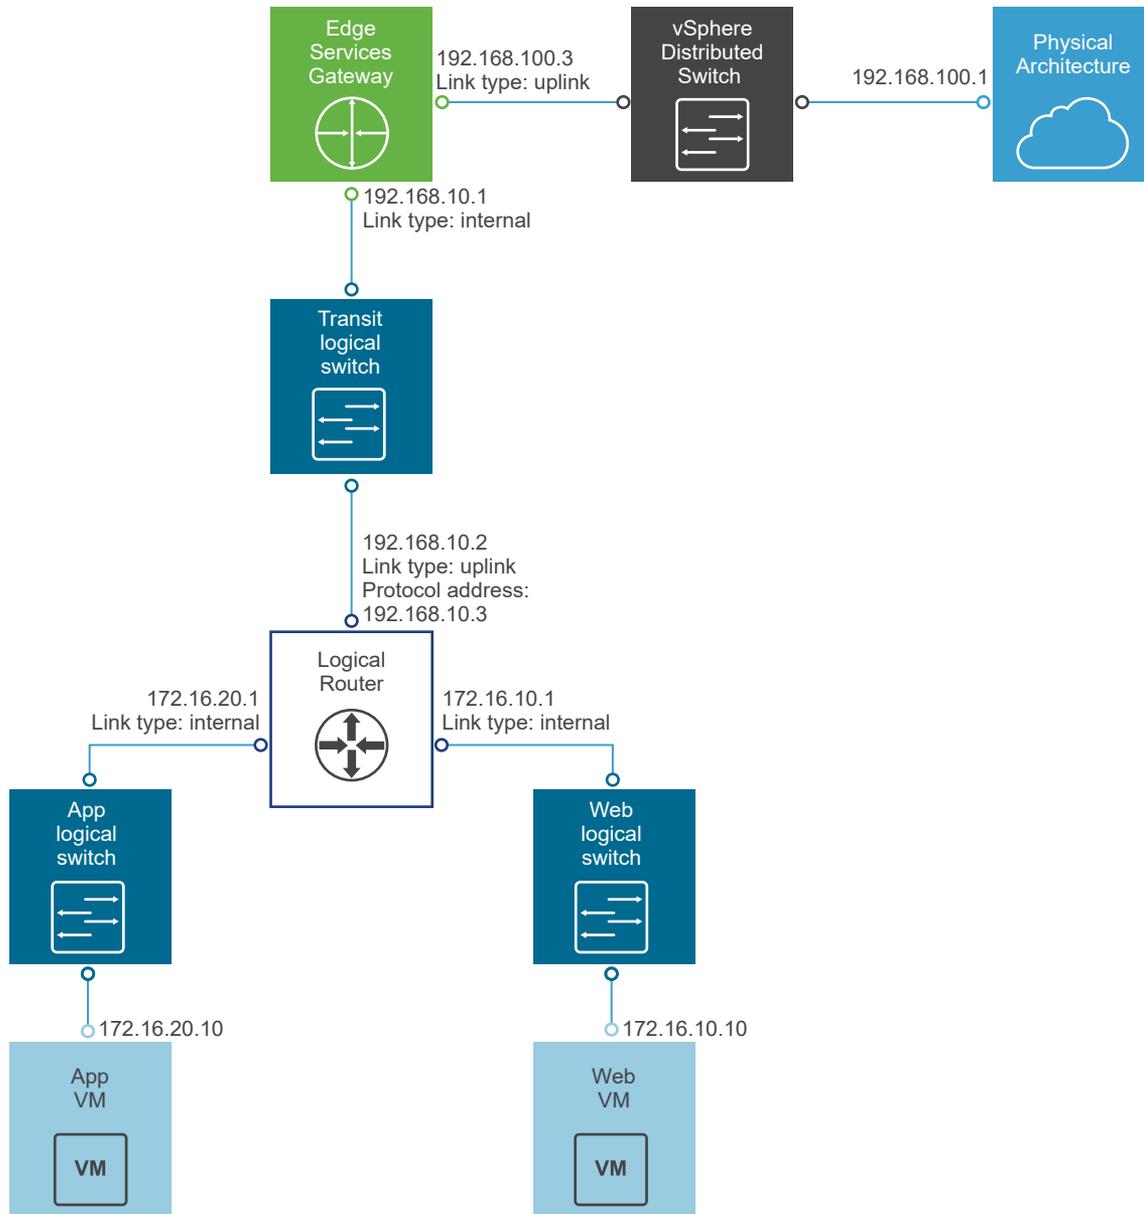
**11** Click **Publish Changes**.**12** Make sure that the route redistribution and firewall configuration allow the correct routes to be advertised.

## Example: OSPF Configured on the Edge Services Gateway

One simple NSX scenario that uses OSPF is when a logical router and an edge services gateway are OSPF neighbors, as shown here.

The ESG can be connected to the outside world through a bridge, a physical router, or through an uplink port group on a vSphere distributed switch, as shown in the following figure.

Figure 9-2. NSX Data Center for vSphere Topology



On the **Global Configuration** page, the configuration settings are as follows:

- **vNIC:** uplink
- **Gateway IP:** 192.168.100.1. The ESG's default gateway is the ESG's uplink interface to its external peer.
- **Router ID:** 192.168.100.3. The router ID is the uplink interface of the ESG. In other words, the IP address that faces its external peer.

On the **OSPF Configuration** page, the configuration settings are as follows:

- **Area Definition:**
  - Area ID: 0
  - Type: Normal
  - Authentication: None

The internal interface (the interface facing the logical router) is mapped to the area, as follows:

- vNIC: internal
- Area ID: 0
- Hello Interval (seconds): 10
- Dear Interval (seconds): 40
- Priority: 128
- Cost: 1

The connected routes are redistributed into OSPF so that the OSPF neighbor (the logical router) can learn about the ESG's uplink network. To verify the redistributed routes, on the left navigation panel, click **Route Redistribution**, and check the following settings:

- **Route Redistribution Status** shows that OSPF is enabled.
- **Route Redistribution Table** shows the following:
  - Learner: OSPF
  - From: Connected
  - Prefix: Any

- Action: Permit

**Note** OSPF can also be configured between the ESG and its external peer router, but more typically this link uses BGP for route advertisement.

Make sure that the ESG is learning OSPF external routes from the logical router.

```

NSX-edge-7-0> show ip route

Codes: O - OSPF derived, i - IS-IS derived, B - BGP derived,
C - connected, S - static, L1 - IS-IS level-1, L2 - IS-IS level-2,
IA - OSPF inter area, E1 - OSPF external type 1, E2 - OSPF external type 2,
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2

Total number of routes: 5

S 0.0.0.0/0 [0/0] via 192.168.100.1
O E2 172.16.10.0/24 [110/1] via 192.168.10.2
O E2 172.16.20.0/24 [110/1] via 192.168.10.2
C 192.168.10.0/29 [0/0] via 192.168.10.1
C 192.168.100.0/24 [0/0] via 192.168.100.3

```

To verify connectivity, make sure that an external device in the physical architecture can ping the VMs.

For example:

```

PS C:\Users\Administrator> ping 172.16.10.10

Pinging 172.16.10.10 with 32 bytes of data:
Reply from 172.16.10.10: bytes=32 time=5ms TTL=61
Reply from 172.16.10.10: bytes=32 time=1ms TTL=61

Ping statistics for 172.16.10.10:
 Packets: Sent = 2, Received = 2, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
 Minimum = 1ms, Maximum = 5ms, Average = 3ms

```

```

PS C:\Users\Administrator> ping 172.16.20.10

Pinging 172.16.20.10 with 32 bytes of data:
Reply from 172.16.20.10: bytes=32 time=2ms TTL=61
Reply from 172.16.20.10: bytes=32 time=1ms TTL=61

Ping statistics for 172.16.20.10:
 Packets: Sent = 2, Received = 2, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
 Minimum = 1ms, Maximum = 2ms, Average = 1ms

```

## Configure BGP

Border Gateway Protocol (BGP) makes core routing decisions. It includes a table of IP networks or prefixes, which designate network reachability among multiple autonomous systems.

An underlying connection between two BGP speakers is established before any routing information is exchanged. Keepalive messages are sent by the BGP speakers in order to keep this relationship alive. After the connection is established, the BGP speakers exchange routes and synchronize their tables.

#### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > Routing > BGP**.
- 5 Next to **BGP Configuration**, click **Edit**, and then click **Enable BGP**.
- 6 (Optional) Click **Enable Graceful Restart** for packet forwarding to be uninterrupted during restart of BGP services.
- 7 (Optional) Click **Enable Default Originate** to allow the ESG to advertise itself as a default gateway to its peers.
- 8 In **Local AS**, enter the router ID. The routes are advertised when BGP peers with routers in other autonomous systems (AS). The path of autonomous systems that a route traverses is used as one metric when selecting the best path to a destination.
- 9 In **Neighbors**, click **Add**.
- 10 Specify basic details of the BGP neighbor.
  - a Type the IP address of the neighbor.

When you configure BGP peering between an edge services gateway (ESG) and a logical router, use the protocol IP address of the logical router as the BGP neighbor address of the ESG.
  - b (On a logical router only) Type the forwarding address.

The forwarding address is the IP address that you assigned to the distributed logical router's interface facing its BGP neighbor (its uplink interface).
  - c (On a logical router only) Type the protocol address.

The protocol address is the IP address that the logical router uses to form a BGP neighbor relationship. It can be any IP address in the same subnet as the forwarding address, but this IP address must not be used anywhere else. When you configure BGP peering between an edge services gateway (ESG) and a logical router, use the protocol IP address of the logical router as the BGP neighbor IP address of the ESG.
  - d Type the remote AS.
  - e (Optional) Disable **Remove Private AS**. By default, it is enabled.
  - f Edit the default weight for the neighbor connection, if necessary. The default weight is 60.

- g **Hold Down Timer** displays a default value of 180 seconds, which is three times the value of keep alive timer. Edit if necessary.

When BGP peering is achieved between two neighbors, the NSX Edge starts a hold down timer. Each keep alive message it receives from the neighbor resets the hold down timer to 0. When the NSX Edge fails to receive three consecutive keep alive messages so that the hold down timer reaches 180 seconds, the NSX Edge considers the neighbor as down and deletes the routes from this neighbor.

---

**Note** The default time-to-live (TTL) value for eBGP neighbors is 1 and for iBGP neighbors is 64. This value cannot be modified.

---

- h **Keep Alive Timer** displays the default frequency of 60 seconds at which a BGP neighbor sends keep alive messages to its peer. Edit if necessary.
- i If authentication is required, enter an authentication password.

Password must be at least 12 characters and it must satisfy these rules:

- Must not exceed 255 characters
- At least one uppercase letter and one lowercase letter
- At least one number
- At least one special character
- Must not contain the user name as a substring
- Must not consecutively repeat a character 3 or more times.

Each segment sent on the connection between the neighbors is verified. MD5 authentication must be configured with the same password on both BGP neighbors, otherwise, the connection between them is not made. You cannot enter a password when FIPS mode is enabled.

## 11 Specify the BGP Filters.

- a Click **Add**.

---

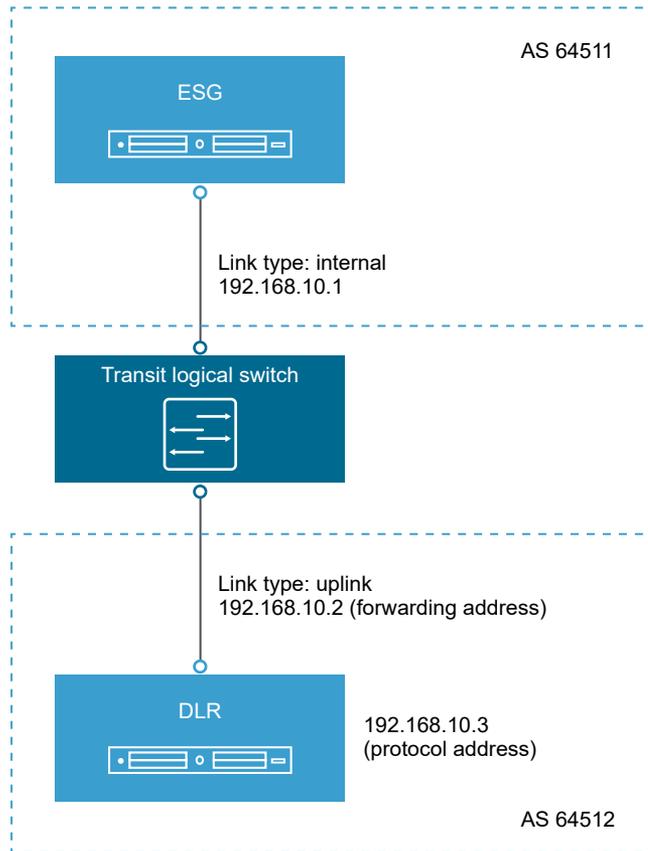
**Caution** A "block all" rule is enforced at the end of the filters.

---

- b Select the direction to indicate whether you are filtering traffic to or from the neighbor.
- c Select the action to indicate whether you are allowing or denying traffic.
- d Type the network in CIDR format that you want to filter to or from the neighbor.
- e Type the IP prefixes that are to be filtered and click **OK**.

## 12 Click **Publish Changes**.

## Example: Configure BGP Between an ESG and a Logical (Distributed) Router



In this topology, the ESG is in AS 64511. The logical router (DLR) is in AS 64512.

The forwarding address of the logical router is 192.168.10.2. This address is configured on the uplink interface of the logical router. The protocol address of the logical router is 192.168.10.3. The ESG uses this address to form a BGP peer relationship with the logical router.

On the **BGP Configuration** page of the logical router, the configuration settings are as follows:

- **Local AS:** 64512
- **Neighbor settings:**
  - Forwarding address: 192.168.10.2
  - Protocol address: 192.168.10.3
  - IP address: 192.168.10.1
  - Remote AS: 64511

On the **BGP Configuration** page of the ESG, the configuration settings are as follows:

- **Local AS:** 64511

- **Neighbor settings:**
  - IP address: 192.168.10.3. This IP address is the protocol address of the logical router.
  - Remote AS: 64512

On the logical router, run the `show ip bgp neighbors` command, and make sure that the BGP state is Established.

```

NSX-edge-6-0> show ip bgp neighbors

BGP neighbor is 192.168.10.1, remote AS 64511,
BGP state = Established, up
Hold time is 180, Keep alive interval is 60 seconds
Neighbor capabilities:
 Route refresh: advertised and received
 Address family IPv4 Unicast:advertised and received
 Graceful restart Capability:advertised and received
 Restart remain time: 0
Received 120 messages, Sent 125 messages
Default minimum time between advertisement runs is 30 seconds
For Address family IPv4 Unicast:advertised and received
 Index 1 Identifier 0x9aa20f3c
 Route refresh request:received 0 sent 0
 Prefixes received 0 sent 0 advertised 0
Connections established 1, dropped 5
Local host: 192.168.10.3, Local port: 179
Remote host: 192.168.10.1, Remote port: 43846

```

On the ESG, run the `show ip bgp neighbors` command, and make sure that the BGP state is Established.

```

NSX-edge-7-0> show ip bgp neighbors

BGP neighbor is 192.168.10.3, remote AS 64512,
BGP state = Established, up
Hold time is 180, Keep alive interval is 60 seconds
Neighbor capabilities:
 Route refresh: advertised and received
 Address family IPv4 Unicast:advertised and received
 Graceful restart Capability:advertised and received
 Restart remain time: 0
Received 121 messages, Sent 120 messages
Default minimum time between advertisement runs is 30 seconds
For Address family IPv4 Unicast:advertised and received
 Index 3 Identifier 0x40212c6c
 Route refresh request:received 0 sent 0
 Prefixes received 0 sent 0 advertised 0
Connections established 1, dropped 1
Local host: 192.168.10.1, Local port: 43846
Remote host: 192.168.10.3, Remote port: 179

```

## Configure Route Redistribution

By default, routers share routes with other routers that are running the same protocol. In a multi-protocol environment, you must configure route redistribution for cross-protocol route sharing.

You can exclude an interface from route redistribution by adding a deny criterion for its network. From NSX 6.2, the HA (management) interface of a logical (distributed) router is automatically excluded from route redistribution.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > Routing > Route Redistribution**.
- 5 Next to **Route Redistribution Status**, click **Edit**.
- 6 Select the protocols for which you want to enable route redistribution and click **OK**.
- 7 Add an IP prefix.

Entries in the IP Prefix list are processed sequentially.

a In **IP Prefixes**, click **Add**.

b Enter a name and IP address of the network.

The IP prefix entered is exactly matched, except if you include less-than-or-equal-to (LE) or greater-than-or-equal-to (GE) modifiers.

c LE and GE together specify a range of prefix lengths that the rule must match. You can add IP prefix GE as a minimum prefix length to be matched and IP prefix LE as a maximum prefix length to be matched.

You can use these two options individually or in conjunction. Values of LE and GE cannot be zero or greater than 32. GE value cannot be greater than LE value. For example,

- If you provide a prefix as 10.0.0.0/16 and LE = 28, then the redistribution rule matches all prefixes ranging from 10.0.0.0/16 to 10.0.0.0/28. It means that the rule matches all prefix lengths from 16 to 28. Prefix 10.0.2.0/24 is matched.
- If you provide a prefix as 10.0.0.0/16 and GE = 24, then the redistribution rule matches all prefixes ranging from 10.0.0.0/24 to 10.0.0.0/32. Prefix 10.0.0.16/28 is matched.
- If you provide GE = 24 and LE = 28, then the redistribution rule matches all prefixes ranging from 10.0.0.0/24 to 10.0.0.0/28. Prefix 10.0.0.32/27 is matched.

d Click **Add** or **OK**.

- 8 Specify a redistribution criteria for the IP prefix.
  - a In **Route Redistribution Table**, click **Add**.
  - b In **Prefix Name**, select the IP prefix that you added earlier.
  - c In **Learner Protocol**, select the protocol that has to learn routes from other protocols.
  - d In **Allow Learning From**, select the protocols from which routes must be learned.
  - e In **Action**, select **Permit** for exact subnet to redistribute, or select **Deny**.
  - f Click **Add** or **OK**.
- 9 Click **Publish Changes**.

## View the NSX Manager Locale ID

Each NSX Manager has a locale ID. By default, it is set to the NSX Manager UUID. This setting can be overwritten at the universal logical router, cluster, or host level.

### Procedure

- 1 Navigate to the **About NSX** or **NSX Home** page.
  - In NSX 6.4.6 and later, click **Networking & Security > About NSX**.
  - In NSX 6.4.5 and earlier, click **Networking & Security > NSX Home > Summary**.
- 2 From the **NSX Manager** drop-down menu, select the IP address of the NSX Manager. Observe that the **Locale ID** or **ID** field shows the UUID of the NSX Manager.

## Configure Locale ID on a Universal Logical (Distributed) Router

If local egress is enabled when a universal logical router is created, routes are sent to hosts only when the host locale ID matches the locale ID associated with the route. You can change the locale ID on a router, and this updated locale ID gets associated with all the static and dynamic routes on this router. The routes are sent to hosts and clusters with matching locale IDs.

See [Cross-vCenter NSX Topologies](#) for information on routing configurations for cross-vCenter NSX environments.

### Prerequisites

The universal logical (distributed) router must have been created with local egress enabled.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click a universal logical (distributed) router.

- 4 Click **Manage > Routing > Global Configuration**.
- 5 Next to **Routing Configuration**, click **Edit**.
- 6 Enter a new Locale ID.

---

**Important** The locale ID must be in UUID format. For example, XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX, where each X is replaced with a base 16 digit (0-F).

---

## Configure Locale ID on a Host or Cluster

If local egress is enabled when a universal logical router is created, routes are sent to hosts only when the host locale ID matches the locale ID associated with the route. You can selectively send routes to hosts by configuring the locale ID on a cluster of hosts, or a host.

### Prerequisites

The universal logical (distributed) router that performs routing for the hosts or clusters must have been created with local egress enabled.

### Procedure

- 1 Navigate to **Networking & Security > Installation and Upgrade > Host Preparation**.
- 2 Select the NSX Manager that manages the hosts or clusters you need to configure.
- 3 Complete the following steps to change the locale ID.

NSX Version	Procedure
NSX 6.4.1 and later	<ol style="list-style-type: none"> <li>a Click a cluster from the left pane. In the right pane, the hosts in the selected cluster are displayed in the Hosts table.</li> <li>b To change the locale ID for the cluster, click <b>Actions &gt; Change Locale ID</b>.</li> <li>c To change the locale ID for a host, click the three dots menu (  ) next to the host and click <b>Change Locale ID</b>.</li> <li>d Type a new locale ID and click <b>Save</b>.</li> </ol> <p><b>Note</b> The locale ID must be in UUID format. For example, XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX, where each X is replaced with a base 16 digit (0-F).</p>
NSX 6.4.0	<ol style="list-style-type: none"> <li>a Select the host or cluster you want to modify, expanding clusters to display hosts if needed.</li> <li>b Click <b>Actions &gt; Change Locale ID</b>.</li> <li>c Type a new locale ID and click <b>OK</b>.</li> </ol> <p><b>Note</b> The locale ID must be in UUID format. For example, XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX, where each X is replaced with a base 16 digit (0-F).</p>

### Results

The universal controller cluster will send only routes matching this new locale ID to the hosts.

## What to do next

Configure a static route with a locale ID specified.

# Multicast Routing Support, Limitations, and Topology

NSX Data Center 6.4.2 and later supports multicast routing.

NSX uses two multicast routing protocols: Internet Group Management Protocol (IGMPv2) and Protocol Independent Multicast (PIM). PIM sparse mode is supported (PIM-SM). PIM is used on ESGs, but not on the DLR.

- Multicast routing is supported between the ESG and the DLR.
- Receiving hosts advertise their group membership to a local multicast router, enabling them to join and leave multicast groups.
- Protocol-Independent Multicast (PIM) is used for a router-to-router signaling. It maintains the current IP multicast service mode of receiver-initiated membership.

After any routing protocol is first enabled (or disabled and re-enabled), traffic is not forwarded until the protocol has converged, and the routes corresponding to the traffic have been learned and installed. In a multicast network, traffic forwarding requires both the unicast and the multicast routing protocols to converge. PIM Sparse-mode also requires that the Rendezvous Point (RP) for a multicast group is known before any control or data traffic for the multicast group is processed. When the PIM Bootstrap mechanism is used to disseminate the RP information, the Candidate RPs are learned only after a Bootstrap message from a PIM neighbor is received. These messages have an RFC default periodicity of 60 secs. If a Static RP is configured, the RP information is available immediately and the delay associated with the Bootstrap mechanism is avoided.

Support and Limitations:

- IPv4 support.
- IGMPv2 support.
- PIM Sparse mode is supported.
- Rendezvous point information can be delivered through bootstrap messages, or statically configured.
- Replication Multicast Range should not overlap with a Transport Zone multicast range.
- An Edge Services Gateway (ESG) cannot be a Bootstrap Candidate Router.
- An ESG cannot be the Rendezvous Point (RP).
- The routes of multicast participating nodes must be learned explicitly either by the unicast routing protocol or through a static route. NSX does not use the default route for multicast reverse path forwarding (RPF) checks.
- During vMotion of virtual machines that are receivers of multicast, a 1–2 second multicast traffic loss can occur.

- Starting in NSX 6.4.7, distributed firewall (DFW) is supported for multicast traffic. However, IPFIX is not supported for multicast.
- Starting in NSX 6.4.7, edge firewall is supported for multicast traffic. Edge firewall supports filtering of IGMP packets on the basis of protocol in the IP header. The firewall cannot filter the type of IGMP packets, such as, membership report, leave group, and so on.

#### Topologies:

- In a Cross-VC environment, connecting two Edge Services Gateways with multicast to the same universal TLS is not supported.
- A single tier of Edge Services Gateway is supported.
- Single distributed logical router, that is, only one downlink per ESG.
- In NSX 6.4.5 or later, multicast is supported on a maximum of two uplink interfaces and one downlink interface per ESG. However, if an NSX Edge is at 6.4.4 or earlier, multicast is supported on a single uplink interface and a single downlink interface per ESG.
- On a DLR, multicast is supported on a single uplink interface and on multiple internal interfaces.
- Starting in NSX 6.4.7, PIM is supported on one GRE tunnel per ESG. PIM can be enabled either on a maximum of two uplink interfaces of the ESG or one GRE tunnel interface, but not on both simultaneously. To reach the sources, receivers, and RP outside the NSX network, static routes must be configured with the IP address of the GRE tunnel endpoint as the next hop.
- Active-standby high availability is supported by enabling ESG HA. Active-active high availability using ECMP is not supported.
- High availability failover time of 30 seconds.
- Cold standby, no sync of mroutes or mFIB.
- L2 bridging is not supported on a logical switch with multicast routing.
- Hardware VTEP gateways (ToR gateways) are not supported on a logical switch with multicast routing.

## Configure Multicast on a Logical (Distributed) Router

IP multicast routing enables a host (source) to send a single copy of data to a single multicast address. Data is then distributed to a group of recipients by using a special form of IP address called the IP multicast group address. In a multicast environment, any host, regardless of whether it is a member of a group, can send to a group. However, only the members of a group can receive packets sent to that group.

NSX uses two multicast routing protocols: Internet Group Management Protocol (IGMPv2) and Protocol Independent Multicast (PIM). PIM sparse mode is supported (PIM-SM). PIM is used on ESGs, but not on the DLR.

For more information about multicast support in NSX, see [Multicast Routing Support, Limitations, and Topology](#).

---

**Attention** During vMotion of virtual machines that are receivers of multicast, a 1–2 second multicast traffic loss can occur.

---

### Prerequisites

Transport Zones must have a multicast address range configured. See [Assign a Segment ID Pool and Multicast Address Range](#) in the NSX Installation Guide.

IGMP configuration must be the same across the Edge Services Gateway and the Logical (Distributed) Router.

Enable IGMP snooping on the L2 switches to which VXLAN participating hosts are attached. If IGMP snooping is enabled on L2, IGMP querier must be enabled on the router or L3 switch with connectivity to multicast enabled networks. See [Add a Logical Switch](#).

### Procedure

- 1 In the vSphere Client, navigate to **Networking & Security > NSX Edges**.
- 2 Click a logical (distributed) router.
- 3 Click **Routing > Multicast**.
- 4 Enable multicast.

Version	Procedure
NSX 6.4.2 to 6.4.4	In <b>Configuration</b> , click the toggle switch to enable multicast.
NSX 6.4.5 and later	<ol style="list-style-type: none"> <li>a Next to <b>Configuration</b>, click <b>Edit</b>.</li> <li>b In <b>Status</b>, click the toggle switch to enable multicast.</li> </ol>

- 5 Enter the replication multicast range.

Version	Procedure
NSX 6.4.2 to 6.4.4	In <b>Replication Multicast Range</b> , enter a range of multicast group addresses in the CIDR format.
NSX 6.4.5 and later	<ol style="list-style-type: none"> <li>a Next to <b>Configuration</b>, click <b>Edit</b>.</li> <li>b In <b>Replication Multicast Range</b>, enter a range of multicast group addresses in the CIDR format.</li> </ol>

Replication Multicast Range is a range of multicast group addresses (VXLAN outer destination IP) that is used to replicate workload/tenant multicast group addresses (VXLAN inner destination IP). Replication Multicast Range IP addresses should not overlap with the multicast address range, configured in **Networking & Security > Installation and Upgrade > Logical Network Settings**. For more information, see [Assign a Segment ID Pool and Multicast Address Range](#) in the *NSX Installation Guide*.

- 6 Configure IGMP parameters. IGMP messages are used primarily by multicast hosts to signal their interest in joining a specific multicast group, and to begin receiving group traffic. IGMP Parameters configured on the DLR must match those configured on the ESG, and have to be configured globally for the ESG and the DLR.

IGMP Parameter	Description
Query	Optional. Configures the frequency at which the designated router sends IGMP host-query messages. The default is 30 seconds. Maximum value is 3,744 seconds.
Query Max Response Time (sec)	Optional. Sets the maximum amount of time that can elapse between when the querier router sends a host-query message and when it receives a response from a host. The default is 10 seconds. Maximum value is 25 seconds.
Last Member Query Interval (sec)	Optional. Configures the interval at which the router sends IGMP group-specific query messages. The default is 1 second. Maximum value is 25 seconds.
Robustness Variable	Optional. The default value is 2. Maximum value is 255.

- 7 Under **Enabled Interfaces**, click **Configure Interfaces** and enable multicast on the uplink and internal interfaces.

#### Note

- Multicast must be enabled on all the DLRs that should receive IPv4 multicast packets.
- On a DLR, multicast is supported on a single uplink interface and on multiple internal interfaces.

- 8 Click **Publish** for the changes to take effect.

#### Results

To verify the multicast routing configurations on a given host and DLR, run the CLI command: `show logical-router host <host ID> dlr <DLR instance> mrouting-domain`

In the example output below, the host is host-19 and the DLR instance is edge-1:

```
cli>show logical-router host host-19 dlr edge-1 mrouting-domain
VDR Mcast Routing Domain configurations:
Vdr Name: edge-1
Vdr ID: 0x00002328
Multicast Routing Doman: Enabled
Replication Mcst Grp Start IP: 237.0.0.0
Replciation Mcast Grp Mask: 255.255.255.0
Control VNI: 9008
Uplink VNI: 9007
```

```
IGMP Query Interval: 30 sec
IGMP Query Response Interval: 10 sec
IGMP Robustness Variable: 2
Group membership Interval: 70 se
```

## Configure Multicast on an Edge Services Gateway

Using Multicast, a source can send a single copy of data to a single multicast address, which is then distributed to a group or recipients.

NSX uses two multicast routing protocols: Internet Group Management Protocol (IGMPv2) and Protocol Independent Multicast (PIM). PIM sparse mode is supported (PIM-SM). PIM is used on ESGs, but not on the DLR.

---

**Attention** During vMotion of virtual machines that are receivers of multicast, a 1–2 second multicast traffic loss can occur.

---

For more information about multicast support in NSX, see [Multicast Routing Support, Limitations, and Topology](#).

If an ESG is at 6.4.4 or earlier, PIM is supported on a single uplink interface of the Edge. Starting in NSX Data Center 6.4.5, PIM is supported on two uplink interfaces of the ESG.

Starting in NSX 6.4.7, PIM is also supported on one GRE virtual tunnel interface (VTI) per ESG. PIM can be enabled on a maximum of two uplink interfaces of the ESG or one GRE tunnel interface. However, you cannot enable PIM simultaneously on the GRE virtual tunnel interface and edge uplink interfaces.

To enable PIM on a GRE tunnel interface, you must first configure GRE tunnels on the ESG by using the NSX APIs. For more information about configuring GRE tunnels, see the *NSX API Guide*. After configuring GRE tunnels on the ESG, you can view the list of GRE tunnels in the vSphere Client UI.

The GRE virtual tunnel interface can be configured with either IPv4 address, or IPv6 address, or both. However, to enable PIM on the GRE tunnel interface, the tunnel interface must have an IPv4 address. If the GRE virtual tunnel interface is configured with only an IPv6 address, this GRE tunnel interface cannot be enabled as a PIM interface.

When PIM is enabled on a GRE tunnel interface, static routes must be added with the IP address of the GRE virtual tunnel endpoint as the next hop IP address. Static routes are required to reach the multicast sources, receivers, and rendezvous point (RP) outside the NSX network.

### Procedure

- 1 In the vSphere Client, navigate to **Networking & Security > NSX Edges**.
- 2 Click an NSX Edge.
- 3 Click **Routing > Multicast**.

#### 4 Enable multicast.

Version	Procedure
NSX 6.4.2 to 6.4.4	In <b>Configuration</b> , click the toggle switch to enable multicast.
NSX 6.4.5 and later	<ol style="list-style-type: none"> <li>Next to <b>Configuration</b>, click <b>Edit</b>.</li> <li>In <b>Status</b>, click the toggle switch to enable multicast.</li> </ol>

- 5 Configure IGMP parameters. IGMP messages are used primarily by multicast hosts to signal their interest in joining a specific multicast group, and to begin receiving group traffic. IGMP Parameters configured on the DLR must match those configured on the ESG, and have to be configured globally for the ESG and the DLR.

IGMP Parameter	Description
Query	Optional. Configures the frequency at which the designated router sends IGMP host-query messages. The default is 30 seconds. Maximum value is 3,744 seconds.
Query Max Response Time (sec)	Optional. Sets the maximum amount of time that can elapse between when the querier router sends a host-query message and when it receives a response from a host. The default is 10 seconds. Maximum value is 25 seconds.
Last Member Query Interval (sec)	Optional. Configures the interval at which the router sends IGMP group-specific query messages. The default is 1 second. Maximum value is 25 seconds.
Robustness Variable	Optional. The default value is 2. Maximum value is 255.

- 6 (Optional) Under **PIM Sparse Mode Parameters** (PIM-SM), enter the **Static Rendezvous Point Address**. The rendezvous point (RP) is the router in a multicast network domain that acts as the root of the multicast shared tree. Packets from the upstream source and join messages from the downstream routers “rendezvous” at this core router. RP can be configured statically and dynamically with use of a Bootstrap Router (BR). If a static RP is configured, it is applicable to all the multicast groups.

An ESG cannot be the rendezvous point or Bootstrap Candidate Router. PIM-SM configuration is done at the PIM global configuration level per edge.

- 7 Under **Enabled Interfaces**, click **Configure Interfaces**, and enable PIM on the interfaces.

You can enable PIM on the following interfaces:

- A maximum of two uplink interfaces of an ESG or on a single GRE virtual tunnel interface of an ESG, but not on both at the same time.
- A single internal interface of an ESG.

- 8 Click **Publish** for the changes to take effect.

### What to do next

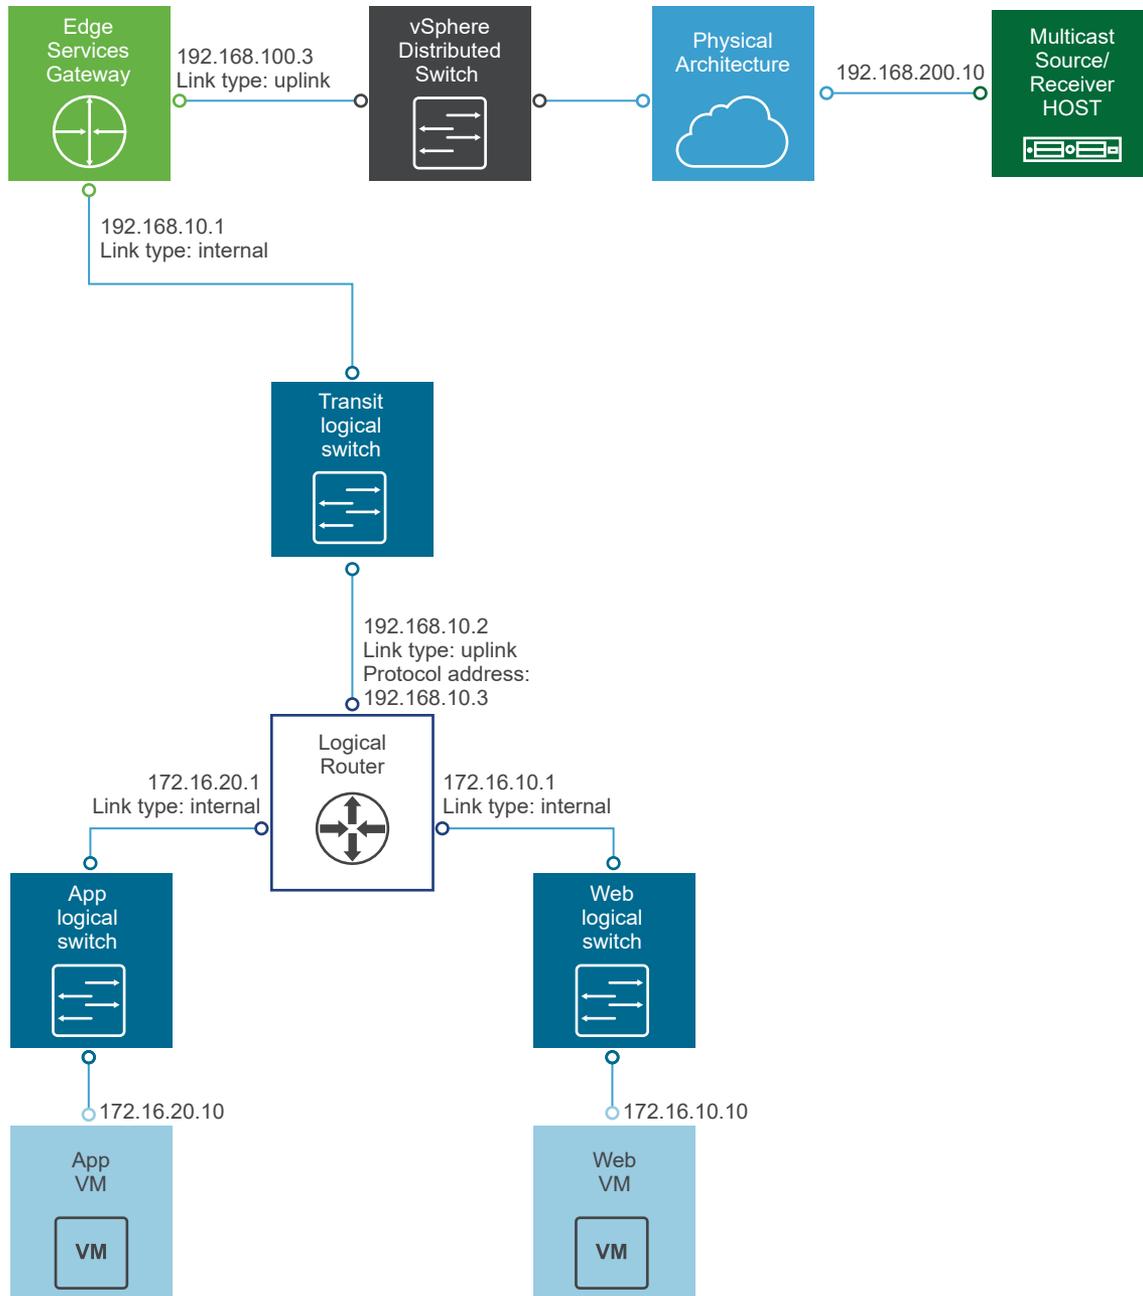
If you have enabled PIM on a GRE virtual tunnel interface, static routes are required to reach the multicast sources, receivers, and RP outside the NSX network. You must configure static routes with the IP address of the GRE virtual tunnel endpoint as the next hop IP address.

For detailed information about configuring static routes, see [Add a Static Route](#).

## Multicast Topology

The following figure shows a sample topology using Multicast.

The Edge Service Gateway uplink interface is connected to the physical infrastructure through the vSphere distributed switch. The Edge Service Gateway internal interface is connected to a logical router through a logical transit switch. The logical router's default gateway is the ESG's internal interface IP address (192.168.10.1). The router ID is the logical router's uplink interface---in other words, the IP address that faces the ESG (192.168.10.2). In this topology, Multicast traffic is replicated in an optimal way, across subnets, between sources and receivers inside or outside the NSX domain.



Logical Firewall provides security mechanisms for dynamic virtual data centers, and consists of two components to address different deployment use cases. Distributed Firewall focuses on East-West access controls, and Edge Firewall focuses on the North-South traffic enforcement at the tenant or datacenter perimeter. Together, these components address the end-to-end firewall needs of virtual datacenters. You can choose to deploy either of these technologies independently, or deploy both of them.

This chapter includes the following topics:

- [Distributed Firewall](#)
- [Edge Firewall](#)
- [Working with Firewall Rule Sections](#)
- [Working with Firewall Rules](#)
- [Firewall Logs](#)

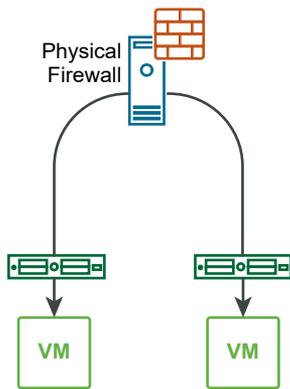
## Distributed Firewall

A Distributed Firewall (DFW) runs in the kernel as a VIB package on all the ESXi host clusters that are prepared for NSX. Host preparation automatically activates DFW on the ESXi host clusters.

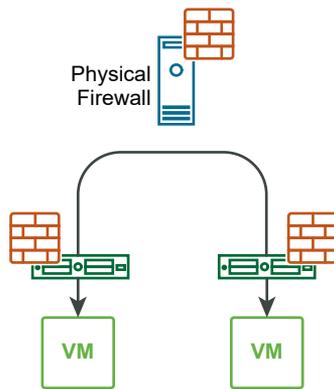
The fundamental constraints of traditional perimeter-centric security architecture impact both security posture and application scalability in modern data centers. For example, hair-pinning of traffic through physical firewalls at the perimeter of the network creates an extra latency for certain applications.

DFW complements and enhances your physical security by removing unnecessary hair-pinning from the physical firewalls and reduces the amount of traffic on the network. Rejected traffic is blocked before it leaves the ESXi host. There is no need for the traffic to traverse the network, only to be stopped at the perimeter by the physical firewall. Traffic destined to another VM on the same host or another host does not have to traverse through the network up to the physical firewall, and then go back down to the destination VM. Traffic is inspected at the ESXi level and delivered to the destination VM.

Security Without NSX DFW



Security With NSX DFW



NSX DFW is a stateful firewall, meaning it monitors the state of active connections and uses this information to determine which network packets to allow through the firewall. DFW is implemented in the hypervisor and applied to virtual machines on a per-vNIC basis. That is, the firewall rules are enforced at the vNIC of each virtual machine. Inspection of traffic happens at the vNIC of a VM just as the traffic is about to exit the VM and enter the virtual switch (egress). Inspection also happens at the vNIC just as the traffic leaves the switch but before entering the VM (ingress).

NSX Manager virtual appliance, NSX Controller VMs, and NSX Edge Service Gateways are automatically excluded from DFW. If a VM does not require DFW service, you can manually add it to the exclusion list.

As DFW is distributed in the kernel of every ESXi host, firewall capacity scales horizontally when you add hosts to the clusters. Adding more hosts increases the DFW capacity. As your infrastructure expands and you buy more servers to manage your ever-growing number of VMs, the DFW capacity increases.

## DFW Policy Rules

DFW policy rules are created by using the vSphere Web Client, and the rules are stored in the NSX Manager database. With DFW, you can create Ethernet rules (L2 rules) and General rules (L3 to L7 rules). The rules are published from NSX Manager to ESXi cluster and then from ESXi host down to VM level. All ESXi hosts in the same cluster have the same DFW policy rules.

A distributed firewall instance on an ESXi host contains the following two tables:

- Rules table to store all security policy rules.
- Connection Tracker table to cache flow entries for rules with an “allow” action.

DFW rules are run in a "top-down" order. Traffic that must go through a firewall is first matched against a firewall rules list. 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. The last rule in the table is the DFW default rule. Packets not matching any rule above the default rule are enforced by the default rule.

Each VM has its own firewall policy rules and context. During vMotion, when VMs move from one ESXi host to another host, the DFW context (Rules table, Connection Tracker table) moves with the VM. In addition, all active connections remain intact during vMotion. In other words, DFW security policy is independent of VM location.

## Micro-Segmentation Using DFW

Micro-segmentation makes the data center network more secure by isolating each related group of virtual machines onto a distinct logical network segment. Micro-segmentation allows the administrator to firewall traffic traveling from one logical segment of the data center to another logical segment (east-west traffic). So, firewalling of east-west traffic limits the attacker's ability to move laterally in the data center.

Micro-segmentation is powered by the Distributed Firewall (DFW) component of NSX. The power of DFW is that the network topology is no longer a barrier to security enforcement. The same degree of traffic access control can be achieved with any type of network topology.

For a detailed example of micro-segmentation use case, see the "Micro-Segmentation with NSX DFW and Implementation" section in the *NSX Network Virtualization Design Guide* at <https://communities.vmware.com/docs/DOC-27683>.

## DFW Policy Rules Based on User Identity

Distributed firewall can help in creating identity-based rules too. Security administrators can enforce access control based on the user identity and the user's group memberships as defined in the enterprise Active Directory. For example, identity-based distributed firewall rules can be used in the following scenarios:

- Users want to virtual applications using a laptop or mobile device where Active Directory is used for user authentication.
- Users want to access virtual applications using VDI infrastructure where the virtual machines are running Microsoft Windows operating system.

For more information about Active-Directory user-based DFW rules, see [Chapter 12 Identity Firewall Overview](#).

## Context-Aware Firewall

Context-aware firewall enhances visibility at the application level and helps to override the problem of application permeability. Visibility at the application layer helps you to monitor the workloads better from a resource, compliance, and security point of view.

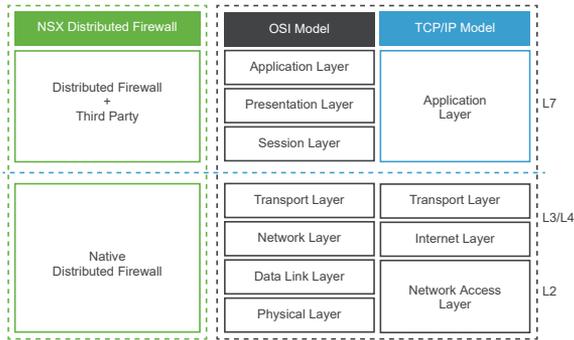
Firewall rules cannot consume application IDs. Context-aware firewall identifies applications and enforces a micro-segmentation for EAST-WEST traffic, independent of the port that the application uses. Context-aware or application-based firewall rules can be defined by defining Layer 7 service objects. After defining Layer 7 service objects in rules, you can define rules with specific protocol, ports, and their application definition. Rule definition can be based on more than 5-tuples. You can also use Application Rule Manager to create context-aware firewall rules.

Context-aware firewall is supported starting in NSX Data Center for vSphere 6.4.

All host clusters in an existing infrastructure managed by NSX Data Center for vSphere must be upgraded to NSX Data Center for vSphere 6.4.0 or later.

## Types of Firewall

Firewall takes action based on one or a combination of different L2, L3, L4, and L7 packet headers that are added to the data as it moves through each layer of the TCP/IP model.



In layer 3 or layer 4 firewall, the action is taken solely based on source/destination IP, port, and protocol. The activity of network connections is also tracked. This type of firewall is known as a stateful firewall.

Layer 7 or context-aware firewall can do everything that the layer 3 and layer 4 firewall do. Also, it can intelligently inspect the content of the packets. For example, a layer 7 firewall rule can be written to deny all HTTP requests from a specific IP address.

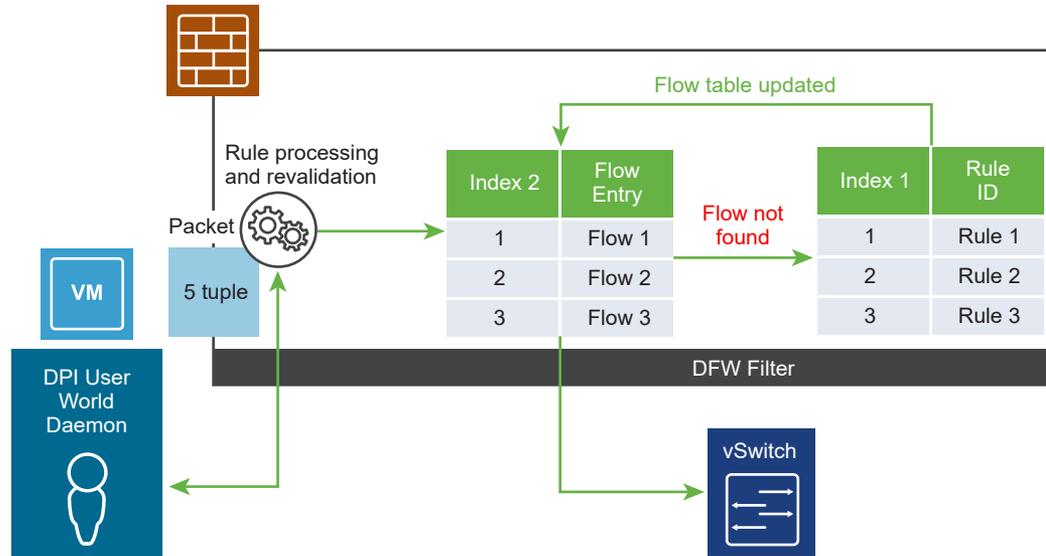
## Rule Definition and Packet Capture

You can create more than 5-tuples rules for context-aware firewall. You can keep adding as many tuples as required to create the correct rule. Protocol and user are the only attributes supported, and you can have either user or protocol per rule. These can be in the standard SRC/DEST fields, or added at the end for extra attributes.

There are 7-tuples in the following rule:

Source	Destination	Service	Direction	Action	Attributes
location-set-1	lport-set-1	HTTPS	INOUT	Allow	TLS-V10

Figure 10-1. Packet Processing in the Context-Aware Firewall



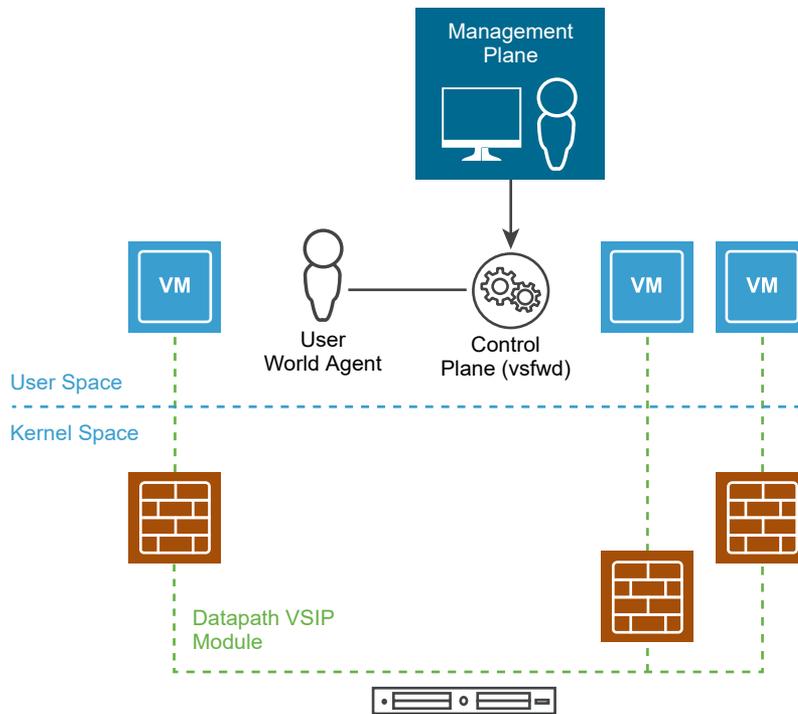
When a context-aware firewall is configured for the virtual machine, the Distributed Deep Packet Inspection (DPI) attributes must also be matched with the 5-tuples. This is where rules are processed and validated again and the correct rule is found. Depending on the action, a flow is created or dropped.

Here is how a rule is processed for an incoming packet:

- 1 Upon entering a DFW 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 Once 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 ensuring that no matching L4 rules take precedence. The appropriate action is taken (allow/deny) and the flow table entry is updated accordingly.

It is possible to have a context-aware firewall rule exactly like an L3 or L4 rule, without really defining the context. If that is the case, the validation step might be performed to apply the context, which might have more attributes.

## Context-Aware Firewall Workflow



## Application ID GUIDs

Layer 7 application identification identifies which application a particular packet or flow is generated by, independent of the port that is being used.

Enforcement based on application identity enables users to allow or deny applications to run on any port, or to force applications to run on their standard port. Deep Packet Inspection (DPI) enables matching packet payload against defined patterns, commonly referred to as signatures. Layer 7 service objects can be used for port-independent enforcement or to create new service objects that leverage a combination of Layer 7 application identity, protocol and port. Layer 7 based service objects can be used in the firewall rule table and Service Composer, and application identification information is captured in Distributed Firewall logs, and Flow Monitoring and Application Rule Manager (ARM) when profiling an application.

**Table 10-1. Application Identification GUIDs**

GUID	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
AD_BKUP	Microsoft Active Directory Backup Service	Networking
AD_NSP	Microsoft Active Directory Service Provider	Networking

Table 10-1. Application Identification GUIDs (continued)

GUID	Description	Type
AMQP	Advanced Message Queueing Protocol, is an application layer protocol which supports business message communication between applications or organizations	Networking
AVAST	Traffic generated by browsing Avast.com official website of Avast! Antivirus downloads	Web Services
AVG	AVG Antivirus/Security software download and updates	File Transfer
AVIRA	Avira Antivirus/Security software download and updates	File Transfer
BLAST	A remote access protocol that compresses, encrypts, and encodes a computing experiences at a data center and transmits it across any standard IP network for VMware Horizon desktops.	Remote Access
BDEFENDER	BitDefender Antivirus/Security software download and updates.	File Transfer
CA_CERT	Certification authority (CA) issues digital certificates which certifies the ownership of a public key for message encryption	Networking
CIFS	CIFS (Common Internet File System) is used to provide shared access to directories, files, printers, serial ports, and miscellaneous communications between nodes on a network	File Transfer
CLRCASE	A software tool for revision control of source code and other software development assets. It is developed by the Rational Software division of IBM. ClearCase forms the base of revision control for many large and medium sized businesses and can handle projects with hundreds or thousands of developers	Networking
CTRXCGP	Citrix Common Gateway Protocol	Remote Access
CTRXCOTO	Hosting Citrix GoToMeeting, or similar sessions based on the GoToMeeting platform. Includes voice, video, and limited crowd management functions	Collaboration
CTRXCICA	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
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

Table 10-1. Application Identification GUIDs (continued)

GUID	Description	Type
GITHUB	Web-based Git or version control repository and Internet hosting service	Collaboration
HTTP	(HyperText Transfer Protocol) the principal transport protocol for the World Wide Web	Web Services
HTTP2	Traffic generated by browsing websites that support the HTTP 2.0 protocol	Web Services
IMAP	IMAP (Internet Message Access Protocol) is an Internet standard protocol for accessing email on a remote server	Mail
KASPRSKY	Kaspersky Antivirus/Security software download and updates	File Transfer
KERBEROS	Kerberos is a network authentication protocol designed to provide strong authentication for client/server applications by using secret-key cryptography	Networking
LDAP	LDAP (Lightweight Directory Access Protocol) is a protocol for reading and editing directories over an IP network	Database
MAXDB	SQL connections and queries made to a MaxDB SQL server	Database
MCAFEE	McAfee Antivirus/Security software download and updates	File Transfer
MSSQL	Microsoft SQL Server is a relational database.	Database
NFS	Allows a user on a client computer to access files over a network in a manner similar to how local storage is accessed	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

Table 10-1. Application Identification GUIDs (continued)

GUID	Description	Type
RDP	RDP (Remote Desktop Protocol) provides users with a graphical interface to another computer	Remote Access
RTCP	RTCP (Real-Time Transport Control Protocol) is a sister protocol of the Real-time Transport Protocol (RTP). RTCP provides out-of-band control information for an RTP flow.	Streaming Media
RTP	RTP (Real-Time Transport Protocol) is primarily used to deliver real-time audio and video	Streaming Media
RTSP	RTSP (Real Time Streaming Protocol) is used for establishing and controlling media sessions between end points	Streaming Media
RTSPS	A secure network control protocol designed for use in entertainment and communications systems to control streaming media servers. The 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
SKIP	Simple Key Management for Internet Protocols (SKIP) is hybrid Key distribution protocol Simple Key Management for Internet Protocols (SKIP) is similar to SSL, except that it establishes a long-term key once, and then requires no prior communication in order to establish or exchange keys on a session-by-session basis.	Networking
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
SQLNET	Networking software that allows remote data-access between programs and the Oracle Database, or among multiple Oracle Databases.	Database
SQLSERV	SQL Services	Database
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
SVN	Managing content on a Subversion server.	Database
SYMUPDAT	Symantec LiveUpdate traffic, this includes spyware definitions, firewall rules, antivirus signature files, and software updates.	File Transfer
SYSLOG	Symantec LiveUpdate traffic, this includes spyware definitions, firewall rules, antivirus signature files, and software updates.	Network Monitoring

Table 10-1. Application Identification GUIDs (continued)

GUID	Description	Type
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

### Example: Create a Context-Aware Firewall Rule

You can configure a context-aware or an application-based firewall rule by defining layer 7 service objects. A layer 7 context-aware firewall rule can intelligently inspect the content of the packets.

This example explains the process of creating a layer 7 firewall rule with APP\_HTTP service object. This firewall rule allows HTTP requests from a virtual machine to any destination. After creating the firewall rule, you initiate some HTTP sessions on the source VM that passes this firewall rule, and turn on flow monitoring on a specific vNIC of the source VM. The firewall rule detects an HTTP application context and enforces the rule on the source VM.

#### Prerequisites

You must log in to the vSphere Web Client with an account that has any one of the following NSX roles:

- Security administrator
- NSX administrator
- Security engineer
- Enterprise administrator

---

**Note** Make sure that NSX Data Center for vSphere 6.4 or later is installed.

---

#### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Security > Firewall**.
- 2 (Optional) Add a firewall rule section to group context-aware firewall rules.
- 3 Click **Add Rule**.

- 4 Create the context-aware firewall rule.
  - a Enter a rule name to identify this rule. For example, enter **L7\_Rule\_HTTP\_Service**.
  - b In the Source column, click the **Edit** (✎) icon.  
The Specify Source page opens.
  - c From the **Object Type** drop-down menu, select **Virtual Machine**.
  - d From the **Available Objects** list, select the virtual machine. Move this object to the **Selected Objects** list, and then click **Save**.
  - e In the Destination column, retain the default value as Any.
  - f In the Service column, click the **Edit** (✎) icon.  
The Specify Service page opens.
  - g From the **Object Type** drop-down menu, select **Services**.
  - h From the **Available Objects** list, select **App\_HTTP** service. Move this service to the **Selected Objects** list, and then click **Save**.
  - i Make sure that the firewall rule is enabled, and the rule action is set to **Allow**.
  - j Click **Publish** to publish the firewall rule configuration.

The following figure shows the firewall rule that you created.

**Figure 10-2. Context-Aware Firewall Rule Definition**

#	Name	ID	Source	Destination	Service	Applied To	Action	Log
1	L7_Rule_HTTP_Service		l2vpn-client-vm	Any	APP_HTTP	Distributed Firewall	Allow	Off

- 5 Log in to the console of your source VM and initiate the `wget` Linux command to download files from the web using HTTP.
- 6 On the vNIC of the source VM, turn on live flow monitoring to monitor traffic flows on the source VM.
  - a Navigate to **Tools > Flow Monitoring**.
  - b Select a particular vNIC on the source VM. For example, select **l2vpn-client-vm-Network adapter 1**.
  - c Click **Start** to view the flow monitoring data.

- In the following figure, the flow monitoring data shows that the firewall rule has detected the application (HTTP) context. Rule 1005 is enforced on source VM (10.161.117.238) and traffic flows to destination IP addresses 151.101.129.67 and 151.101.53.67.

Figure 10-3. Traffic Flows on Source VM

Refresh Rate: 15 Seconds

● New active flows ● Flow with state change ● Terminated flows

Rule ID	Direction	Flow Type	Protocol	Source IP	Source Port	Destination IP	Destination Port	State	Incoming Bytes	Incoming Packets	Outgoing Bytes	Outgoing Packets	Application Context
1001	IN	ACTIVE	UDP	10.161.111.76	138	10.161.127.255	138		229 Bytes	1	0 Bytes	0	
1001	IN	ACTIVE	UDP	10.161.107.122	138	10.161.127.255	138		473 Bytes	2	0 Bytes	0	
1001	OUT	ACTIVE	UDP	10.161.117.238	34358	10.162.204.1	53		282 Bytes	2	212 Bytes	4	DNS
1001	IN	ACTIVE	UDP	10.161.101.97	138	10.161.127.255	138		229 Bytes	1	0 Bytes	0	
1001	OUT	ACTIVE	UDP	10.161.117.238	46089	10.162.204.1	53		236 Bytes	2	228 Bytes	4	DNS
1005	OUT	ACTIVE	TCP	10.161.117.238	60530	151.101.129.67	80	FINWAIT2	764 Bytes	6	642 Bytes	8	HTTP
1001	OUT	ACTIVE	TCP	10.161.117.238	44991	151.101.53.67	443	TIMEWAIT	2 MB	1316	19 KB	359	HTTPS_TLS_V12
1005	OUT	ACTIVE	TCP	10.161.117.238	38389	151.101.53.67	80	FINWAIT2	765 Bytes	6	766 Bytes	8	HTTP
1001	IN	ACTIVE	UDP	10.161.122.244	138	10.161.127.255	138		229 Bytes	1	0 Bytes	0	
1001	IN	ACTIVE	UDP	10.161.104.96	137	10.161.127.255	137		234 Bytes	3	0 Bytes	0	
1002	IN	ACTIVE	UDP	0.0.0.0	68	255.255.255.255	67		312 Bytes	1	0 Bytes	0	
1001	IN	INACTIVE	UDP	10.161.108.60	138	10.161.127.255	138		229 Bytes	1	0 Bytes	0	
1001	IN	INACTIVE	UDP	10.161.105.84	138	10.161.127.255	138		229 Bytes	1	0 Bytes	0	
1001	IN	INACTIVE	UDP	10.161.105.103	138	10.161.127.255	138		229 Bytes	1	0 Bytes	0	
1001	IN	ACTIVE	UDP	10.161.113.41	138	10.161.127.255	138		239 Bytes	1	0 Bytes	0	
1001	IN	ACTIVE	UDP	10.161.119.59	5353	224.0.0.251	5353		132 Bytes	1	0 Bytes	0	
1001	IN	ACTIVE	UDP	10.161.107.247	138	10.161.127.255	138		229 Bytes	1	0 Bytes	0	
1001	IN	ACTIVE	UDP	10.161.121.151	5353	224.0.0.251	5353		140 Bytes	1	0 Bytes	0	
1001	IN	ACTIVE	UDP	10.161.96.34	5353	224.0.0.251	5353		144 Bytes	1	0 Bytes	0	
1001	IN	ACTIVE	UDP	10.161.127.103	138	10.161.127.255	138		229 Bytes	1	0 Bytes	0	

- Return to the Firewall page, and change the rule action to **Block**.
- Go to the console of the source VM and run the `wget` command again.

Observe that the HTTP requests are now blocked on the source VM. You should see an error in the VM console that says something like this:

```
HTTP request sent, awaiting response ... Read error (Connection reset by peer) in headers
Retrying.
```

The following figure shows a flow with the application (HTTP) context detected and blocked on the vNIC of the source VM (10.161.117.238).

Figure 10-4. Traffic Flows on Source VM

Refresh Rate: 15 Seconds

● New active flows ● Flow with state change ● Terminated flows

Rule ID	Direction	Flow Type	Protocol	Source IP	Source Port	Destination IP	Destination Port	State	Incoming Bytes	Incoming Packets	Outgoing Bytes	Outgoing Packets	Application Context
1001	IN	ACTIVE	UDP	10.161.103.172	138	10.161.127.255	138		243 Bytes	1	0 Bytes	0	
1001	IN	ACTIVE	UDP	10.161.106.86	138	10.161.127.255	138		229 Bytes	1	0 Bytes	0	
1001	IN	ACTIVE	UDP	10.161.104.187	137	10.161.127.255	137		624 Bytes	8	0 Bytes	0	
1002	IN	ACTIVE	UDP	10.161.113.56	68	255.255.255.255	67		328 Bytes	1	0 Bytes	0	
1001	IN	ACTIVE	UDP	10.161.100.174	138	10.161.127.255	138		229 Bytes	1	0 Bytes	0	
1001	IN	ACTIVE	UDP	10.161.107.233	5353	224.0.0.251	5353		73 Bytes	1	0 Bytes	0	
1001	IN	ACTIVE	UDP	10.161.106.209	138	10.161.127.255	138		229 Bytes	1	0 Bytes	0	
1005	OUT	BLOCK	TCP	10.161.117.238	39598	151.101.167	80	EST	120 Bytes	2	329 Bytes	4	HTTP
1001	OUT	ACTIVE	UDP	10.161.117.238	35551	10.162.204.1	53		282 Bytes	2	212 Bytes	4	DNS
1003	IN	ACTIVE	IPv6-ICMP	fd01:0:101:2611:0:a:0:54e	0	ff02::1	0		72 Bytes	1	0 Bytes	0	
1005	OUT	BLOCK	TCP	10.161.117.238	39599	151.101.167	80	EST	120 Bytes	2	329 Bytes	4	HTTP
1003	IN	ACTIVE	IPv6-ICMP	fe80::250:56ff:feb1:6b...	0	ff02::1	0		72 Bytes	1	0 Bytes	0	
1001	IN	ACTIVE	UDP	10.161.125.35	138	10.161.127.255	138		229 Bytes	1	0 Bytes	0	
1005	OUT	BLOCK	TCP	10.161.117.238	39601	151.101.167	80	EST	120 Bytes	2	381 Bytes	5	HTTP
1005	OUT	BLOCK	TCP	10.161.117.238	39600	151.101.167	80	EST	120 Bytes	2	381 Bytes	5	HTTP
1005	OUT	BLOCK	TCP	10.161.117.238	39602	151.101.167	80	EST	120 Bytes	2	381 Bytes	5	HTTP
1001	IN	ACTIVE	UDP	10.161.114.246	138	10.161.127.255	138		229 Bytes	1	0 Bytes	0	
1001	IN	ACTIVE	UDP	10.161.97.126	137	10.161.127.255	137		78 Bytes	1	0 Bytes	0	
1001	IN	ACTIVE	UDP	10.161.112.214	137	10.161.127.255	137		3 KB	45	0 Bytes	0	
1003	IN	ACTIVE	IPv6-ICMP	fd01:0:101:2611:250:56f...	0	ff02::1	0		72 Bytes	1	0 Bytes	0	

What to do next

To know about other scenarios where you can use context-aware firewall rules, see [Context-Aware Firewall Scenarios](#).

## Session Timers

Session Timers can be configured for TCP, UDP, and ICMP sessions.

Session Timers define how long a session is maintained on the firewall after inactivity. When the session timeout for the protocol expires, the session closes.

On the firewall, a number of timeouts for TCP, UDP, and ICMP sessions can be specified to apply to a user-defined subset of virtual machines or vNICs. By default, any virtual machines or vNICs not included in the user-defined timer are included in the global session timer. All of these timeouts are global, meaning they apply to all of the sessions of that type on the host.

Default session values can be modified depending on your network needs. Note that setting a value too low could cause frequent timeouts, and setting a value too high could delay failure detection.

### Create a Session Timer

Session Timers define how long a session is maintained on the firewall after inactivity in the session.

On the firewall, you can define timeouts for TCP, UDP, and ICMP sessions for a set of user defined VMs or vNICs. The default timer is global, meaning that it applies to all virtual machines protected by firewall.

#### Procedure

- 1 Navigate to Timeout Settings.
  - ◆ In NSX 6.4.1 and later, navigate to **Networking & Security > Security > Firewall Settings > Timeout Settings**.
  - ◆ In NSX 6.4.0, navigate to **Networking & Security > Security > Firewall > Settings**.
- 2 If there is more than one NSX Manager available, select one from the drop-down list.
- 3 Click the **Add (+)** icon.
- 4 Enter a **name** (required) and a **description** (optional) for the session timer.
- 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.
Closing	The timeout value for the connection after the first FIN has been sent. The default is 120 seconds.
Open	The timeout value for the connection after a second packet has been transferred. The default is 30 seconds.
Fin Wait	The timeout value for the connection after both FINs have been exchanged and the connection is closed. The default is 45 seconds.
Established	The timeout value for the connection once the connection has become fully established.
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 will be 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 In NSX 6.1 and later, click **Next**.
- 7 Select the object type, **vNIC** or **VM**.  
The Available Objects list is automatically populated.
- 8 Select one or more objects and click the arrow to move them to the **Selected Objects** column.
- 9 Click **OK** or **Finish**.

#### Results

A timer has been created to apply to set of user defined hosts.

#### Edit a Session Timer

Configure timeout parameters for TCP, UDP and ICMP protocols.

After a session timer has been created it can be changed as needed. The default session timer can also be edited.

#### Procedure

- 1 Navigate to Timeout Settings.
  - ◆ In NSX 6.4.1 and later, navigate to **Networking & Security > Security > Firewall Settings > Timeout Settings**.
  - ◆ In NSX 6.4.0, navigate to **Networking & Security > Security > Firewall > Settings**.
- 2 If there is more than one NSX Manager available, select one from the drop-down list.
- 3 Select the timer you want to edit. Note that the default timer values can also be edited. Click the **pencil** icon.
- 4 Enter a **name** (required) and a **description** (optional) for the session timer.

- 5 Select the protocol. Edit any default values that you want to change.

TCP Variables	Description
First Packet	The timeout value for the connection after the first packet has been sent. The default is 120 seconds.
Closing	The timeout value for the connection after the first FIN has been sent. The default is 120 seconds.
Open	The timeout value for the connection after a second packet has been transferred. The default is 30 seconds.
Fin Wait	The timeout value for the connection after both FINs have been exchanged and the connection is closed. The default is 45 seconds.
Established	The timeout value for the connection once the connection has become fully established.
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 will be 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 In NSX 6.1 and later, click **Next**.
- 7 Select the object type, **vNIC** or **VM**.  
The Available Objects list is automatically populated.
- 8 Select one or more objects and click the arrow to move them to the **Selected Objects** column.
- 9 Click **OK** or **Finish**.

## IP Discovery for Virtual Machines

VMware Tools runs on a VM and provides several services. One service that is essential to distributed firewall is associating a VM and its vNICs with IP addresses. Before NSX 6.2, if VMware Tools was not installed on a VM, its IP address was not learned. In NSX 6.2 and later, you can configure clusters to detect virtual machine IP addresses with DHCP snooping, ARP snooping, or both. This allows NSX to detect the IP address if VMware Tools is not installed on the virtual machine. If VMware Tools is installed, it can work in conjunction with DHCP and ARP snooping.

VMware recommends that you install VMware Tools on each virtual machine in your environment. In addition to providing vCenter with the IP address of VMs, it provides the following functions:

- Allow copy and paste between VM and host or client desktop.

- Synchronize time with the host operating system.
- Allow shutdown or restart of the VM from vCenter.
- Collect network, disk, and memory usage from the VM and send it to the host.
- Determine VM availability by sending and collecting heartbeat.

Note that having two vNICs for a VM on the same network is not supported and can lead to unpredictable results around which traffic is blocked or allowed.

For those VMs that do not have VMware Tools installed, NSX will learn the IP address through ARP or DHCP snooping, if ARP and DHCP snooping is enabled on the VM's cluster.

IP addresses detected using ARP snooping are not removed automatically. In other words, there is no timeout for vNIC IP addresses that are detected using ARP snooping.

## Change IP Detection Type

The IP address of a virtual machine can be detected by VMware Tools, which is installed on the VM, or by DHCP snooping and ARP snooping. These IP discovery methods can be used together in the same NSX installation.

You can specify the IP detection types either at a global level or at the host cluster level. Typically, users with security administrator and security engineer roles might prefer to specify the IP detection type at a global level. They use the detected VM IP addresses to configure the SpoofGuard policies and the distributed firewall policies.

Users with an enterprise administrator role usually have a much wider view of the complete virtual network, and might prefer to control the IP detection type by editing the settings at the host cluster level. The IP detection settings at the host cluster level override the settings that are specified at the global level.

### Procedure

- 1 Navigate to the **Change IP Detection Type** page.

IP Detection Level	Steps
Global IP Detection	<ol style="list-style-type: none"> <li>a Navigate to <b>Networking &amp; Security &gt; Security &gt; SpoofGuard</b>.</li> <li>b Next to <b>IP Detection Type</b>, click the  icon.</li> </ol>
Host Cluster IP Detection	<ol style="list-style-type: none"> <li>a Navigate to <b>Networking &amp; Security &gt; Installation and Upgrade &gt; Host Preparation</b>.</li> <li>b Click the cluster for which you want to change the IP detection type, and then click <b>Actions &gt; Change IP Detection Type</b>.</li> </ol>

- 2 Select the desired IP detection types, and click **Save** or **OK**.

IP Detection Type	Description
DHCP Snooping	NSX detects the IP addresses of the VMs in the network by reading the DHCP snooping entries.
ARP Snooping	NSX detects the IP addresses of the VMs by using the ARP snooping mechanism.  <b>Recommendation</b> Configure SpoofGuard when you use ARP snooping to detect IP addresses. SpoofGuard helps you to defend your network against ARP poison attacks.

- 3 If you selected ARP snooping, enter the maximum ARP IP addresses that must be detected per vNIC, per VM. The default value is 1.

ARP snooping can detect a maximum of 128 IP addresses per vNIC, per VM. The valid range of values are 1 through 128. For example, if you specify a value of 5, it means that a maximum of first five IP addresses are detected per vNIC per VM.

IP addresses detected using ARP snooping are not removed automatically. In other words, there is no timeout for vNIC IP addresses that are detected using ARP snooping.

#### What to do next

- If you enabled ARP snooping, consider the option to configure SpoofGuard to defend your network against ARP poison attacks.
- Configure the default firewall rule.

## Exclude Virtual Machines from Firewall Protection

You can exclude a set of virtual machines from distributed firewall protection.

NSX Manager, NSX Controller, and NSX Edge virtual machines are automatically excluded from distributed firewall protection. In addition, place the following service virtual machines in the Exclusion List to allow traffic to flow freely.

- vCenter Server. It can be moved into a cluster that is protected by Firewall, but it must already exist in the exclusion list to avoid connectivity issues.

---

**Note** It is important to add the vCenter Server to the exclusion list before changing the "any any" default rule from allow to block. Failure to do so will result in access to the vCenter Server being blocked after creating a Deny All rule (or modifying default rule to block action). If this occurs, use the API to change the default rule from deny to allow. For example, use `GET /api/4.0/firewall/globalroot-0/config` to retrieve the current configuration, and use `PUT /api/4.0/firewall/globalroot-0/config` to change the configuration. See "Working with Distributed Firewall Configuration" in the *NSX API Guide* for more information.

---

- Partner service virtual machines.

- Virtual machines that require promiscuous mode. If these virtual machines are protected by distributed firewall, their performance may be adversely affected.
- The SQL server that your Windows-based vCenter uses.
- vCenter Web server, if you are running it separately.

#### Procedure

- 1 Navigate to Exclusion List settings.
  - ◆ In NSX 6.4.1 and later, navigate to **Networking & Security > Security > Firewall Settings > Exclusion List**.
  - ◆ In NSX 6.4.0, navigate to **Networking & Security > Security > Firewall > Exclusion List**.
- 2 Click **Add**.
- 3 Move the VMs that you want to exclude to **Selected Objects**.
- 4 Click **OK**.

#### Results

If a virtual machine has multiple vNICs, all of them are excluded from protection. If you add vNICs to a virtual machine after it has been added to the Exclusion List, Firewall is automatically deployed on the newly added vNICs. To exclude the new vNICs from firewall protection, you must remove the virtual machine from the Exclusion List and then add it back to the Exclusion List. An alternative workaround is to power cycle (power off and then power on) the virtual machine, but the first option is less disruptive.

## View Firewall CPU and Memory Threshold Events

When a cluster is prepared for network virtualization, the Firewall module is installed on all hosts of that cluster. This module allocates three heaps, a module heap for module parameters; a rule heap for rules, containers, and filters; and a state heap for traffic flows. Heap size allocation is determined by the available host physical memory. Depending on the number of rules, container sets, and the connections, the heap size may grow or shrink over time. The Firewall module running in the hypervisor also uses the host CPUs for packet processing.

Knowing the host resource utilization at any given time can help you in better organizing your server utilization and network designs.

The default CPU threshold is 100, and the memory threshold is 100. You can modify the default threshold values through REST API calls. The Firewall module generates system events when the memory and CPU usage crosses the thresholds. For information on configuring default threshold values, see Working with Memory and CPU Thresholds in the *NSX API Guide*.

#### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > System > Events**.
- 2 Ensure that you are in the **Monitor** tab.

- 3 Click the **System Events** tab.

## Distributed Firewall Resource Utilization

Memory is used by distributed firewall internal data structures, and can be configured for CPU, RAM and connections per second.

Each ESXi host is configured with the following threshold parameters for DFW resource utilization:

CPU utilization, heap memory, process memory, connections per second (CPS), and maximum connections. An alarm is raised if the respective threshold is crossed 20 consecutive times during a 200-second period. A sample is taken every 10 seconds.

The memory is used by distributed firewall internal data structures, which include filters, rules, containers, connection states, discovered IPs, and drop flows. These parameters can be manipulated using the following API call: `PUT /api/4.0/firewall/stats/thresholds`. See the *NSX API Guide* for more information.

## Edge Firewall

Edge Firewall monitors the North-South traffic to provide perimeter security functionality including firewall, Network Address Translation (NAT), and site-to-site IPsec and SSL VPN functionality. This solution is available in the virtual machine form factor and can be deployed in a High Availability mode.

Firewall support is limited on the Logical Router. Only the rules on management or uplink interfaces work, however, the rules on internal interfaces do not work.

---

**Note** The Edge Services Gateway (ESG) is vulnerable to SYN flood attacks, where an attacker fills the firewall state tracking table by flooding SYN packets. This DOS/DDOS attack creates a service disruption to genuine users. The NSX Edge can defend itself from SYN flood attacks by using the SYN cookie mechanism in a smart way to detect bogus TCP connections and stop them without consuming firewall state tracking resources. Before the SYN queue is not full, the incoming connections pass normally. After the SYN queue is full, the SYN cookie mechanism takes effect.

However, for the servers behind the NSX Edge, the SYN flood protection feature is disabled by default. The NSX Edge uses SYNPROXY to do the SYN flood protection.

---

For detailed information about SYNPROXY behavior when `SynFloodProtection` is enabled on an NSX Edge, see the VMware knowledge base article at <https://kb.vmware.com/s/article/54527>.

## Working with NSX Edge Firewall Rules

You can navigate to an NSX Edge to see the firewall rules that apply to it.

Firewall rules applied to a Logical Router only protect control plane traffic to and from the Logical Router control virtual machine. They do not enforce any data plane protection. To protect the data plane traffic, create Logical Firewall rules for East-West protection or rules at the NSX Edge Services Gateway level for North-South protection.

Rules are displayed and enforced in the following order:

- 1 Predefined distributed firewall rules that are applied to the edge.
  - These rules are defined on the Firewall user interface (**Networking & Security > Security > Firewall**)
  - These rules are displayed in **read-only** mode on the NSX Edge Firewall user interface.
- 2 Internal rules that enable the control traffic to flow for Edge services. For example, internal rules include the following auto-plumbed rules:
  - a SSL VPN auto-plumb rule: The Edge Firewall tab displays the sslvpn auto-plumb rule when server settings are configured and SSL VPN service is enabled.
  - b DNAT auto-plumb rule: The Edge NAT tab displays the DNAT auto-plumb rule as part of the default SSL VPN configuration.
- 3 User-defined rules that are added on the NSX Edge Firewall user interface.
- 4 Default rule.

## Edit the Default NSX Edge Firewall Rule

Default firewall settings apply to traffic that does not match any of the user-defined firewall rules. The default Edge firewall policy blocks all incoming traffic. You can change the default action and logging settings.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > NSX Edges**.
- 2 Double-click an NSX Edge.
- 3 Click **Manage > Firewall**.
- 4 Select the **Default Rule**, which is the last rule in the firewall table.

You can edit rule action or enable or disable logging of all sessions that match the default rule. Enabling logging can affect performance.

NSX Version	Procedure
6.4.6 and later	<ol style="list-style-type: none"> <li>a Edit the rule action, if necessary.</li> <li>b In the <b>Log</b> column, click the toggle switch to enable or disable logging.</li> </ol>
6.4.5 and earlier	<ol style="list-style-type: none"> <li>a Point to the <b>Action</b> cell of the default rule and click .</li> <li>b Edit the rule action, if necessary.</li> <li>c Click <b>Log</b> or <b>Do not log</b> as necessary.</li> </ol>

## 5 Click **Publish Changes**.

### Add an NSX Edge Firewall Rule

You can add user-defined edge firewall rules on the NSX Edge Service Gateways to accept, reject, or deny specific types of traffic. However, you cannot add user-defined firewall rules on a distributed logical router.

The Edge Firewall interface provides the following methods to add an edge firewall rule:

- Add a rule either above or below an existing rule in the firewall table.
- Add a rule by copying an existing rule.
- Add a rule by clicking the **Add** icon.

---

**Remember** If you have created distributed firewall rules and applied them to the edge, these firewall rules are displayed in a **read-only** mode on the Edge Firewall user interface. However, the edge firewall rules that you create using the Edge Firewall user interface are not displayed on the Firewall interface that you used to create the distributed firewall rules (**Networking & Security > Security > Firewall**).

---

#### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > Firewall**.
- 5 Use any of the following three methods to start the process of adding an edge firewall rule.

Method #1: Add a rule either above or below an existing rule in the firewall table.

NSX sets the source, destination, and service columns of the newly added rule as "any". If the system-generated default rule is the only rule in the firewall table, the new rule is added above the default rule. The new rule is enabled by default.

NSX Version	Steps
6.4.6 and later	<ol style="list-style-type: none"> <li>a Select a rule.</li> <li>b Click  and select <b>Add Above</b> or <b>Add Below</b>.</li> </ol>
6.4.5 and earlier	<ol style="list-style-type: none"> <li>a Select a rule.</li> <li>b In the No. column, click , and then select <b>Add Above</b> or <b>Add Below</b>.</li> </ol>

Method #2: Add a rule by copying an existing rule.

In NSX 6.4.5 and earlier, you can create a rule by copying one rule at a time. Starting in NSX 6.4.6, you can select multiple rules to copy simultaneously. The copied rules are enabled by default, and you can edit the rule properties, as necessary.

**Note** When you copy and paste system-generated "internal" rules and "default" rule, the newly created rules are automatically assigned the rule type as "user".

NSX Version	Steps
6.4.6 and later	<ol style="list-style-type: none"> <li>Select the check box next to the rules that you want to copy.</li> <li>Click <b>More &gt; Copy Selected Rule(s)</b>.</li> <li>Select the rule where you want the copied rules to be pasted.</li> <li>Click , and select <b>Paste Rule(s) Above</b> or <b>Paste Rule(s) Below</b>.</li> </ol>
6.4.5 and earlier	<ol style="list-style-type: none"> <li>Select a rule.</li> <li>Click the Copy () icon or , and then select <b>Copy</b>.</li> <li>Select a rule where you want the copied rule to be pasted.</li> <li>In the No. column, click , and select <b>Paste Above</b> or <b>Paste Below</b>.</li> </ol>

Method #3: Add a rule by clicking the **Add** ( or ) icon.

A new row is added in the firewall table. NSX sets the source, destination, and service columns of the newly added rule as "any". If the system-generated default rule is the only rule in the firewall table, the new rule is added above the default rule. The new rule is enabled by default.

## 6 (Optional) Specify a rule name.

- In NSX 6.4.6 and later, click in the **Name** column of the new rule, and enter a rule name.
- In NSX 6.4.5 and earlier, point to the **Name** column of the new rule, and click . Enter a rule name, and click **OK**.

## 7 (Optional) Specify the source of the firewall rule.

You can add IP addresses, vCenter objects, and grouping objects as the source. If no source is added, the source is set to "any". You can add multiple NSX Edge interfaces and IP address groups as the source for firewall rules.

You can choose to create a new IP set or a new security group. After the IP set or security group is created, it is automatically added in the **Source** column of the rule.

- a Select one or more objects to use as sources in the firewall rule.

NSX Version	Steps
6.4.6 and later	<p>To select objects:</p> <ol style="list-style-type: none"> <li>1 Point to the <b>Source</b> column of the rule, and click .</li> <li>2 In the <b>Objects</b> tab, select an object type from the <b>Object Type</b> drop-down menu.</li> <li>3 Select the objects from the <b>Available Objects</b> list and move them to the <b>Selected Objects</b> list.</li> </ol>
6.4.5 and earlier	<p>To select objects:</p> <ol style="list-style-type: none"> <li>1 Point to the <b>Source</b> column of the rule, and click .</li> <li>2 Select an object type from the <b>Object Type</b> drop-down menu.</li> <li>3 Select the objects from the <b>Available Objects</b> list and move them to the <b>Selected Objects</b> list.</li> </ol>

For example, in the following two situations, you can use the "vNIC Group" object type as the source:

#### Select all traffic generated by the NSX Edge

In this situation, select **vNIC Group** from the **Object Type** drop-down menu, and from the **Available Objects** list, select **vse**.

#### Select all traffic originating from any internal or uplink (external) interface of the selected NSX Edge

In this situation, select **vNIC Group** from the **Object Type** drop-down menu, and from the **Available Objects** list, select **internal** or **external**.

The rule is automatically updated when you configure additional interfaces on the edge.

---

**Remember** Firewall rules defined on the internal interfaces do not work on a distributed logical router.

---

- b Enter IP address to use as a source for the firewall rule.

You can enter multiple IP addresses by using a comma-separated list or enter an IP address range. Both IPv4 and IPv6 addresses are supported.

- In NSX 6.4.6 and later, click . Click the **IP Addresses** tab, and then click **Add** to enter the IP addresses.
- In NSX 6.4.5 and earlier, click  and enter the IP addresses.

- c (Optional) Negate the sources defined in your firewall rule.

- If the **Negate Source** option is turned on or selected, the rule is applied to traffic coming from all sources except for the sources defined in this rule.

- If the **Negate Source** option is turned off or not selected, the rule is applied to traffic coming from the sources in this rule.

**8** (Optional) Specify the destination of the firewall rule.

You can add IP addresses, vCenter objects, and grouping objects as the destination. If no destination is added, the destination is set to "any". You can add multiple NSX Edge interfaces and IP address groups as the destination for firewall rules.

The procedure to add objects and IP addresses in the rule destination remains the same as explained in the substeps for adding the rule source.

---

**Tip** Starting in NSX 6.4.6, you can drag objects and IP addresses from the **Source** column to the **Destination** column and conversely. In addition, you can drag objects and IP addresses from one rule to another rule.

---

## 9 (Optional) Specify the service to use in the firewall rule.

- a Add one or more services or service groups in the firewall rule.

You can add either a predefined service or a service group in the rule, or create a new service or a service group to use in the rule. NSX Edge supports services defined only with L3 protocols.

NSX Version	Steps
6.4.6 and later	<ol style="list-style-type: none"> <li>1 Point to the <b>Service</b> column of the new rule and click .</li> <li>2 In the <b>Service/Service Groups</b> tab, select either a service or a service group from the <b>Object Type</b> drop-down menu.</li> <li>3 Select the objects from the <b>Available Objects</b> list and move them to the <b>Selected Objects</b> list.</li> </ol>
6.4.5 and earlier	<ol style="list-style-type: none"> <li>1 Point to the <b>Service</b> column of the new rule and click .</li> <li>2 From the <b>Object Type</b> drop-down menu, select a service or a service group.</li> <li>3 Select the objects from the <b>Available Objects</b> list and move them to the <b>Selected Objects</b> list.</li> </ol>

**Tip** In NSX 6.4.6 and later, you can drag service and service group objects from one user-defined rule to another user-defined rule.

- b Add one or more services in the firewall rule as a port-protocol combination.

**Restriction** Stream Control Transmission Protocol (SCTP) protocol is not supported on an Edge Firewall.

NSX Version	Steps
6.4.6 and later	<ol style="list-style-type: none"> <li>1 Point to the <b>Service</b> column of the new rule and click .</li> <li>2 Click the <b>Raw Port-Protocol</b> tab, and then click <b>Add</b>.</li> <li>3 Select a protocol.</li> <li>4 In the <b>Source Port</b> column, and enter the port numbers.</li> </ol>
6.4.5 and earlier	<ol style="list-style-type: none"> <li>1 Point to the <b>Service</b> column of the new rule and click .</li> <li>2 Select a protocol.</li> <li>3 Expand <b>Advanced Options</b>, and enter the source port numbers.</li> </ol>

## 10 Specify the rule action.

- In NSX 6.4.6 and later, selection an action from the drop-down menu.
- In NSX 6.4.5 and earlier, point to the **Action** column of the rule, and click . Select an action and click **OK**.

The following table describes the rule actions.

Action	Description
Accept or Allow	Allows traffic from or to the specified sources, destinations, and services. By default, action is set to accept traffic.
Deny or Block	Blocks traffic from or to the specified sources, destinations, and services.
Reject	Sends reject message for unaccepted packets. <ul style="list-style-type: none"> <li>■ RST packets are sent for TCP connections.</li> <li>■ ICMP messages with administratively prohibited code are sent for UDP, ICMP, and other IP connections.</li> </ul>

11 (Optional) Specify whether sessions that match this new firewall rule must be logged.

By default, logging is disabled for the rule. Enabling logging can affect performance.

- In NSX 6.4.6 and later, click the toggle switch in the **Log** column to enable logging.
- In NSX 6.4.5 and earlier, point to the **Action** column of the new rule, and click . Select **Log** or **Do not log**.

12 (Optional) Specify the advanced settings of the firewall rule.

- In NSX 6.4.6 and later, click the **Advanced Settings** () icon.
- In NSX 6.4.5 and earlier, point to the **Action** column of the new rule, and click . Expand the **Advanced** options.

The following table describes the advanced options.

Option	Description
Direction	<p>Select whether the rule must be applied on incoming traffic or outgoing traffic or both. The default value is "In/Out", which means that rule is applied symmetrically across both source and destination.</p> <p>VMware does not recommend specifying the direction of firewall rules because "in" or "out" direction can cause the rules to become asymmetric. For example, consider that you have created a firewall rule to "allow" traffic from source A to destination B, and the rule direction is set to "out".</p> <ul style="list-style-type: none"> <li>■ When A sends a packet to B, a state is created based on this rule on A because the direction of traffic is "out" on A.</li> <li>■ When the packet is received on B, the actual traffic direction is "in". Because the rule direction is set to accepting only "outgoing traffic", the rule does not hit this packet on B.</li> </ul> <p>This example shows that setting the "out" direction in the rule causes the rule to become asymmetric.</p>
Match on	<p>Use this option to specify when the firewall rule must be applied.</p> <ul style="list-style-type: none"> <li>■ Select <b>Original</b> when you want the rule to be applied on original IP address and services before network address translation is performed.</li> <li>■ Select <b>Translated</b> when you want the rule to be applied on translated IP address and services after network address translation is performed.</li> </ul>

13 Click **Publish Changes** to push the new rule to the NSX Edge.

**Example: Sample Firewall Rules**

Figure 10-5. Firewall rule for traffic to flow from an NSX Edge interface to an HTTP server

No.	Name	Type	Source	Destination	Service	Action
1	firewall	Internal	vse	any	any	Accept
2	Traffic to HTTP server	User	vnic-index-0:any	HTTP Address Group	For HTTP server	Accept
3	Default Rule	Default	any			Deny

**HTTP Address Group**

Value:  
10.20.222.34

**For HTTP server**

Value:  
TCP:8080

Figure 10-6. Firewall rule for traffic to flow from all internal interfaces (subnets on portgroups connected to internal interfaces) of a NSX Edge to an HTTP Server

No.	Name	Type	Source	Destination	Service	Action
1	firewall	Internal	vse	any	any	Accept
2	Traffic to HTTP server	User	internal	HTTP Address Group	For HTTP server	Accept
3	Default Rule	Default	any			Deny

**HTTP Address Group**

Value:  
10.20.222.34

**For HTTP server**

Value:  
TCP:8080

Figure 10-7. Firewall rule for traffic to allow SSH into a machine on internal network

No.	Name	Type	Source	Destination	Service	Action
1	firewall	Internal	vse	any	any	Accept
2	Traffic to internal network	User	any	VM in internal network	Internal VM	Accept
3	Default Rule	Default	any			Deny

**VM in internal network**

Value:  
192.168.0.10

**Internal VM**

Value:  
TCP:22

**What to do next**

While working with edge firewall rules, you can perform several additional tasks in the firewall table. For example:

- Filter the list of rules in the table by hiding the system-generated default and internal rules, or by hiding the predefined distributed firewall rules that were applied on the edge.
- Search rules that match a specific string by using the **Search** text box. For instance, if you want to search all the rules that contain the string "133", type **133** in the **Search** text box.
- View statistics of the published rules.
  - In NSX 6.4.6 and later, click the **Statistics** (📊) icon.
  - In NSX 6.4.5 and earlier, make sure that the **Stats** column is displayed in the firewall table. If the **Stats** column is not displayed, click  and select the **Stats** column. To view the rule statistics, click .

- Change the order of user-defined rules by clicking the **Move Up** (⬆️ or 📄⬆️) or **Move Down** (⬆️ or 📄⬇️) icons. In NSX 6.4.6 and later, you can drag user-defined rules to change the order.

Point to the user-defined rule that you want to drag. A drag handle (☰) icon appears to the left of the rule. Click and drag this handle to move the rule to a valid location in the firewall table.

---

**Important** You cannot change the order of system-generated internal rules and the default rule.

---

- Disable a rule.
  - In NSX 6.4.6 and later, click the toggle switch to the left of the rule name.
  - In NSX 6.4.5 and earlier, click 🟢 in the **No.** column.
- Undo and redo rule changes until the rule is published. This feature is available in NSX 6.4.6 and later. After the rule is published, the history of rule changes is lost, and you cannot undo or redo the changes.

## Edit an NSX Edge Firewall Rule

You can edit only the user-defined edge firewall rules and make limited changes to the system-generated default firewall rule.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > Firewall**.
- 5 Select the rule to edit.

---

**Note** You cannot edit the following types of rules in the NSX Edge Firewall user interface:

- Internal rules (for example, auto-plumbed rules that enable the control traffic to flow for Edge services.)
  - Predefined distributed firewall rules that are applied to the edge. These firewall rules are defined in the Firewall user interface (**Networking & Security > Security > Firewall**).
- 

- 6 Make the required changes and click **Save** or **OK**.
- 7 Click **Publish Changes**.

## Change the Order of an NSX Edge Firewall Rule

You can change the order of user-defined firewall rules that were added in the Edge Firewall tab to customize traffic flowing through the NSX Edge. For example, suppose you have a rule to allow load balancer traffic. You can now add a rule to deny the load balancer traffic from a specific IP address group, and position this rule above the LB allow traffic rule.

**Procedure**

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > Firewall**.
- 5 Select the rule for which you want to change the order.

---

**Important** You cannot change the order of system-generated internal rules and the default rule.

---

- 6 Click the **Move Up** ( or ) or **Move Down** ( or ) icon.

---

**Tip** In NSX 6.4.6 and later, you can drag user-defined rules to change the order. Point to the user-defined rule that you want to drag. A drag handle () icon appears to the left of this rule. Click and drag this handle to move the rule to a valid location in the firewall table.

---

- 7 Click **Publish Changes**.

**Delete an NSX Edge Firewall Rule**

You can delete only user-defined firewall rules that are added in the NSX Edge Firewall tab.

**Procedure**

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > Firewall**.
- 5 Select a user-defined rule to delete.

---

**Restriction** You cannot delete the following types of firewall rules:

- Default rule
  - System-generated internal (auto-plumbed) rules
  - Predefined distributed firewall rules that are applied to the edge by using the Firewall user interface (**Networking & Security > Security > Firewall**).
- 

- 6 Click the **Delete** ( or ) icon.
- 7 Click **Publish Changes**.

## Mark an Edge Firewall Rule as Valid

An edge firewall rule becomes invalid when grouping objects, services, or service groups that are used in the rule are deleted. Invalid firewall rules cannot be published.

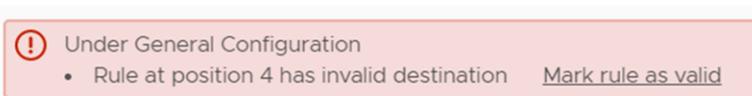
Consider that you have created an edge firewall rule that uses an IP set object in the destination of the rule, as shown in the following figure.

	<input type="checkbox"/>	#	Name	Id	Type	Source	Destination	Service	Action
⋮	<input checked="" type="checkbox"/>	1	firewall	131074	Internal	vse	Any	Any	Accept
⋮	<input checked="" type="checkbox"/>	2	routing	131083	Internal	Any	Any	tcp:179:a...	Accept
⋮	<input checked="" type="checkbox"/>	3	ipsec	133132	Internal	10.108.161... 10.108.161... 10.108.161... 15 more...	10.108.161... 10.108.161... 10.108.161... 15 more...	udp:500... esp:any:a...	Accept
⋮	<input checked="" type="checkbox"/>	4	edge-object	133133	User	Any	FW-IPset	Any	Accept

In the following procedure, you will delete the "FW-IPset" object on the edge, and then return to the firewall table to observe that NSX Edge detects the invalid rule. You will mark the rule as valid and republish the rule.

### Procedure

- 1 Force delete the IP set object on the edge.
  - a Log in to the vSphere Web Client.
  - b Click **Networking & Security > NSX Edges**
  - c Double-click the edge, and navigate to **Manage > Grouping Objects**.
  - d Click the **IP Sets** tab, and then select the **FW-IPSet** object.
  - e Click the **Delete** (🗑️ or ✖️) icon, and then select the **Proceed to force delete** check box.
- 2 Click the **Firewall** tab to return to the edge firewall table.
- 3 Observe that NSX displays the following error message above the firewall table.



NSX Edge detects that the destination of the firewall rule at position 4 is invalid, and therefore the rule becomes invalid. The empty object in the destination column of the rule is enclosed in a red box, as shown in the following figure.

⋮	<input checked="" type="checkbox"/>	<input type="checkbox"/>	4	edge-object	133133	User	Any	sys-gen-empty-ipset-edge-fw	
---	-------------------------------------	--------------------------	---	-------------	--------	------	-----	-----------------------------	--

## 4 (Required) Remove the empty object.

NSX Version	Steps
6.4.6 and later	Point to the empty <code>sys-gen-empty-ipset-edge-fw</code> object, click  , and then select <b>Remove</b> .
6.4.5 and earlier	Point to the empty <code>sys-gen-empty-ipset-edge-fw</code> object and click  .

## 5 (Optional) Edit the rule destination to make the rule configuration valid.

NSX Version	Steps
6.4.6 and later	<ol style="list-style-type: none"> <li>Point to the <b>Destination</b> column of the rule, click , and select <b>Edit Rule Destination</b>.</li> <li>Add objects or IP addresses, as necessary.</li> </ol>
6.4.5 and earlier	<ol style="list-style-type: none"> <li>Point to the <b>Destination</b> column of the rule, click .</li> <li>Add objects or IP addresses, as necessary.</li> </ol>

6 (6.4.6 and later only) Click the **Mark the rule as valid** link in the error message.

NSX Edge displays the following warning message:

```
This action will mark rule as valid.
Please ensure that all elements in the rule are valid objects before performing this
action.
Do you want to continue?
```

- To confirm that the rule can be marked as valid, click **Yes**. The error message is removed.
- To close the warning message and return to the firewall table to verify and edit the rule destination, click **No**.

**Note** In NSX 6.4.5 and earlier, the error message above the firewall table does not show the **Mark rule as valid** link. After you remove the empty object, and optionally edit the rule destination, NSX Edge removes the error message when it detects that the rule configuration has become valid.

7 Click **Publish Changes** for the rule changes to take effect.

## Edge Objects

You can create Edge-level grouping objects to limit the scope of objects to an Edge. Edge grouping objects differ from network grouping objects, which have a global scope and can be used across Edges and other objects.

For example, if you want to create an IP set (IP address group) for a specific Edge and ensure that this IP set is not available for reuse in other contexts, you can create an Edge IP set.

At the scope of an NSX Edge, you can create and manage the following grouping objects:

- IP Sets

- Services
- Service Groups

To create Edge objects in the vSphere Web Client, select an Edge, and navigate to **Manage > Grouping Objects**.

For more information about working with IP sets (IP address groups), services, and service groups, see [Chapter 22 Network and Security Objects](#).

## Managing NAT Rules

NSX Edge provides network address translation (NAT) service to assign a public address to a computer or a group of computers in a private network. NAT technology limits the number of public IP addresses that an organization or company must use for economy and security purposes. You must configure NAT rules on the Edge appliance to provide access to services running on virtual machines inside a company's private network.

The NAT service configuration is separated into source NAT (SNAT) and destination NAT (DNAT) rules.

### Add an SNAT Rule

You can create a source NAT (SNAT) rule to change the source IP address from a public to a private IP address or the reverse.

#### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > NSX Edges**.
- 2 Double-click an NSX Edge.
- 3 Click **Manage > NAT**.
- 4 Click **Add**, and then click **Add SNAT Rule**.
- 5 Select the interface on which to add the rule.
- 6 Select the required protocol.
- 7 Type the original source (public) IP address in one of the following formats.

Format	Example
IP address	192.0.2.0
IP address range	192.0.2.0-192.0.2.24
IP address/subnet	192.0.2.0/24
any	

- 8 Type the original source port or port range.

Format	Example
Port number	80
Port range	80–85
any	

- 9 Type the destination IP address in one of the following formats.

Format	Example
IP address	192.0.2.0
IP address range	192.0.2.0 -192.0.2.24
IP address/subnet	192.0.2.0 /24
any	

- 10 Type the destination port or port range.

Format	Example
Port number	80
Port range	80–85
any	

- 11 Type the translated source IP address in one of the following formats.

Format	Example
IP address	192.0.2.0
IP address range	192.0.2.0-192.0.2.24
IP address/subnet	192.0.2.0/24
any	

- 12 Enable the rule.
- 13 (Optional) Enable logging to log the address translation.
- 14 Click **Add** or **OK** to add the SNAT rule.
- 15 Click **Publish Changes**.

### Add a DNAT Rule

You can create a destination NAT (DNAT) rule to change the destination IP address from a public to private IP address or the reverse.

#### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > NSX Edges**.

- 2 Double-click an NSX Edge.
- 3 Click **Manage > NAT**.
- 4 Click **Add**, and then click **Add DNAT Rule**.
- 5 Select the interface on which to apply the DNAT rule.
- 6 Select the required protocol.
- 7 Type the source IP address in one of the following formats.

Format	Example
IP address	192.0.2.0
IP address range	192.0.2.0 -192.0.2.24
IP address/subnet	192.0.2.0 /24
any	

- 8 Type the source port or port range.

Format	Example
Port number	80
Port range	80–85
any	

- 9 Type the original (public) IP address in one of the following formats.

Format	Example
IP address	192.0.2.0
IP address range	192.0.2.0 -192.0.2.24
IP address/subnet	192.0.2.0 /24
any	

- 10 Type the original port or port range.

Format	Example
Port number	80
Port range	80–85
any	

- 11 Type the translated IP address in one of the following formats.

Format	Example
IP address	192.0.2.0
IP address range	192.0.2.0 -192.0.2.24

Format	Example
IP address/subnet	192.0.2.0 /24
<i>any</i>	

12 Type the translated port or port range.

Format	Example
Port number	80
Port range	80–85
<i>any</i>	

13 Enable the rule.

14 (Optional) Enable logging to log the address translation.

15 Click **Add** or **OK** to add the DNAT rule.

16 Click **Publish Changes**.

### Add a NAT64 Rule

Using NAT64 rules, an NSX Edge performs network address translation to allow traffic from external IPv6 subnetworks to internal IPv4 subnetworks.

NAT64 supports communications initiated by the IPv6-only node towards an IPv4-only node only.

NAT64 supports the following layer 4 protocols:

- TCP
- UDP
- ICMP
  - ICMP echo request and reply only.
  - ICMPv4 errors are supported, ICMPv6 errors are not supported.

All other protocol type packets are discarded.

The translation of IPv4options, IPv6 routing headers, hop-by-hop extension headers, destination option headers, and source routing headers is not supported. FTP is not supported. Fragmented packets are not supported.

NSX Edge high availability is not supported with NAT64. NAT64 sessions are not synced between active and standby appliances, so if a failover occurs, connectivity is interrupted.

If you have dynamic routing protocols configured, IPv4 prefixes are redistributed.

The following timers apply to NAT64 traffic:

**Table 10-3. NAT64 Timers**

Protocol		Timeout
TCP	Incoming TCP-SYNC	6 seconds
	TCP-ESTABLISHED	2 hours
	TCP-Trans	4 minutes
UDP		5 minutes
ICMP		1 minute

**Prerequisites**

- Configure an uplink interface of the Edge Services Gateway with an address on the IPv6 network.
- Configure an internal interface of the Edge Services Gateway with an address on the IPv4 network.
- Ensure that these addresses are not duplicated anywhere else in your environment.

**Procedure**

- 1 In the vSphere Web Client, navigate to **Networking & Security > NSX Edges**.
- 2 Double-click an NSX Edge.
- 3 Click **Manage > NAT**.
- 4 From the **View** drop-down menu, select **NAT64**.

5 Click **Add** and enter the NAT64 parameters.

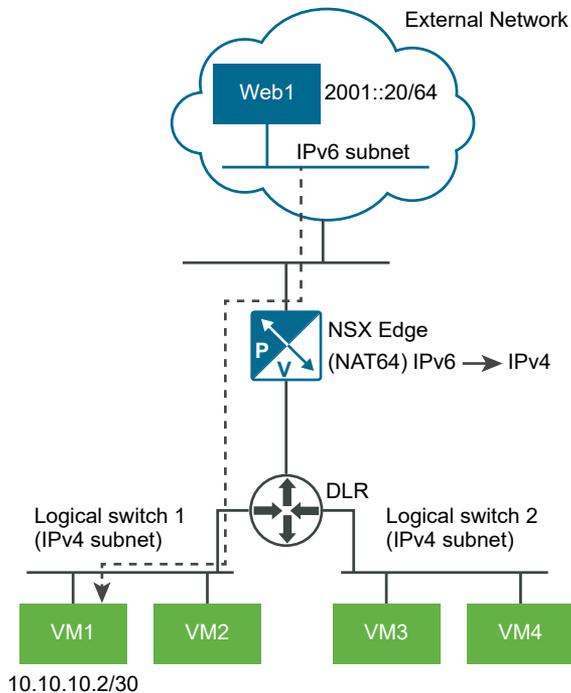
Option	Description
<b>Match IPv6 Destination Prefix</b>	<p>Enter an IPv6 network prefix (network address) or a specific IPv6 address in CIDR notation.</p> <p>As NAT64 provides connectivity from IPv6 subnets to IPv4 subnets, in most situations, you might want to enter an IPv6 network prefix instead of a specific IPv6 address.</p> <p>NAT64 uses the IPv6 network prefix that you specify in this text box to map the IPv4 destination addresses to IPv6 destination addresses. Prefix length must be any one of the following: 32, 40, 48, 56, 64, or 96.</p> <p>For example, if you use the /96 network prefix, NAT64 appends the hexadecimal equivalent of the IPv4 destination address to the IPv6 network prefix. See the sample NAT64 rule after this procedure for an example.</p> <p><b>Note</b> You can use the well-known 64:ff9b::/96 prefix defined in RFC 6052, or use any other IPv6 prefix that is not already used in your environment.</p>
<b>Translated IPv4 Source Prefix</b>	<p>Optional: Enter an IPv4 network prefix (network address) or a specific IPv4 address in CIDR notation.</p> <p>Ensure that the IPv4 network prefix or the IPv4 address is not already used in your environment.</p> <p>As NAT64 provides connectivity from IPv6 subnets to IPv4 subnets, in most situations, you might want to enter an IPv4 network prefix instead of a specific IPv4 address.</p> <p>NAT64 uses an IP address from the IPv4 network prefix to translate the IPv6 source address to an IPv4 source address. See the sample NAT64 rule after this procedure for an example.</p> <p><b>Note</b></p> <ul style="list-style-type: none"> <li>■ The 100.64.0.0/16 IPv4 shared address space is reserved for NAT64. You can use this reserved address space.</li> <li>■ If you keep this text box empty, NAT64 rule automatically uses the reserved address space when you publish the rule.</li> </ul>
<b>Description</b>	Optional description for the rule.
<b>Enabled or Status</b>	Enable the NAT64 rule.
<b>Enable logging or Logging</b>	Enable logging for the NAT64 rule.

6 Click **Add** to save the rule.

7 Click **Publish** for the rule to take effect.

### Example: Sample NAT64 Rule

You want the NSX Edge to allow traffic from Web 1 computer (2001::20/64) that is on an external IPv6 network to VM 1 (10.10.10.2/30), which is on the internal IPv4 subnet.



The NAT64 rule in this example uses the following sample values:

- Match IPv6 Destination Prefix: **64 : ff90 : : /96**
- Translated IPv4 Source Prefix: **30 . 30 . 30 . 0 /24**

The following screen capture shows the published rule. The Rule ID is autogenerated and it might vary in your environment.

**Figure 10-8. NAT64 Rule Definition**

Status	Order	RuleID	Match IPv6 Destination Prefix	Translated IPv4 Source Prefix	Logging
<input checked="" type="checkbox"/>	1	196609	64:ff90::/96	30.30.30.0/24	<input type="checkbox"/>

The NAT64 rule takes the hex equivalent of the destination IPv4 address (10.10.10.2) and appends it to the IPv6 network prefix (64:ff90::) to form the IPv6 destination address: 64:ff90::a0a:a02.

The rule picks up any IP address from the Translated IPv4 Source prefix (30.30.30.0/24). Let us say, the rule picks up 30.30.30.32. NAT64 uses this IPv4 source address to translate the 64:ff90::a0a:a02 destination address to the actual IPv4 destination address (10.10.10.2)

After the rule is published, do the following steps:

- 1 Log in to the command prompt of Web1 computer and issue a `ping` command to the IPv6 destination address 64:ff90::a0a:a02. A nat64 session is established.

- 2 Log in to the NSX Edge CLI and view the nat64 session by running the `show nat64 sessions` command.

```

Protocol IPv6-SA IPv6-DA SPort DPort IPv4_SA IPv4-
DA SPort DPort
TCP 2001::20 64:ff90::a0a:a02 2055 22 30.30.30.32
10.10.10.2 2055 22

```

## Working with Firewall Rule Sections

You can add a section to separate firewall rules. For example, you might want to have the rules for sales and engineering departments in separate sections.

You can create multiple firewall rule sections for L2 and L3 rules. Because multiple users can log in to the web client and simultaneously make changes to firewall rules and sections, users can lock sections that they are working on so that no one else will be able to modify the rules in the section they are working on.

Cross-vCenter NSX environments can have multiple universal rule sections. Multiple universal sections allow rules to be easily organized per tenant and application. If rules are modified or edited within a universal section, only the universal distributed firewall rules for that section are synced to the secondary NSX Managers. You must manage universal rules on the primary NSX Manager, and you must create the universal section there before you can add universal rules. Universal sections are always listed above local sections on both primary and secondary NSX Managers.

Rules outside the universal sections remain local to the primary or secondary NSX Managers on which they are added.

### Add a Firewall Rule Section

You can add sections in the firewall table to organize your rules or to create a universal section for use in cross-vCenter NSX environments.

#### Prerequisites

Determine the appropriate NSX Manager on which to make your changes.

- In a standalone or single vCenter NSX environment there is only one NSX Manager so you do not need to select one.
- Universal objects must be managed from the primary NSX Manager.
- Objects local to an NSX Manager must be managed from that NSX Manager.
- In a cross-vCenter NSX environment that does not have Enhanced Linked Mode enabled, you must make configuration changes from the vCenter linked to the NSX Manager that you want to modify.

- In a cross-vCenter NSX environment in Enhanced Linked Mode, you can make configuration changes to any NSX Manager from any linked vCenter. Select the appropriate NSX Manager from the NSX Manager drop-down menu.

#### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Security > Firewall**.
- 2 If there is more than one NSX Manager available, select one. You must select the Primary NSX Manager to add a universal section.
- 3 Ensure that you are in the **Configuration > General** tab to add a section for L3, L4, or L7 rules. Click the **Ethernet** tab to add a section for L2 rules.
- 4 Click **Add Section** ( or ).
- 5 Enter a name for the section. Section names must be unique within NSX Manager.
- 6 (Optional) In a cross-vCenter NSX environment, you can configure the section as a universal firewall rule section.
  - In NSX 6.4.1 and later, click the **Universal Synchronization** button.
  - In NSX 6.4.0, select **Mark this section for Universal Synchronization**.
- 7 (Optional) Configure firewall rule properties for the firewall section by selecting the appropriate check boxes.

Firewall Rule Section Properties	Description
Enable User Identity at Source	When using Identity Firewall for RDSH, <b>Enable User Identity at Source</b> must be checked. Note that this disables the enable stateless firewall option because the TCP connection state is tracked for identifying the context.
Enable TCP Strict	TCP strict determines whether to drop an established TCP connection when the firewall does not see the initial three-way handshake. If set to true, the connection is dropped.
Enable Stateless Firewall	Enable stateless firewall for the firewall section.

- 8 Click **OK** and then click **Publish Changes**.

#### What to do next

Add rules to the section. See [Add a Firewall Rule](#).

## Merge Firewall Rule Sections

You can merge sections and consolidate the rules within those sections. You cannot merge a section with the Service Composer or Default sections. In a cross-vCenter NSX environment, you cannot merge a section with the universal section.

Merging and consolidating a complex firewall configuration can help with maintenance and readability.

#### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Security > Firewall**.
- 2 Merge the sections.
  - In NSX 6.4.1 and later, on the firewall rule section you want to merge, click the menu (  ) and select **Merge Section**.
  - In NSX 6.4.0, on the firewall rule section you want to merge, click **Merge section** (  ).
- 3 Select whether you want to merge this section with the section above or below.
 

Rules from both sections are merged. The new section keeps the name of the section with which the other section is merged.
- 4 Click **Publish Changes**.

## Delete a Firewall Rule Section

You can 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

- 1 In the vSphere Web Client, navigate to **Networking & Security > Security > Firewall**.
- 2 Ensure that you are in the **Configuration > General** tab to delete a section for L3 rules. Click the **Ethernet** tab to delete a section for L2 rules.
- 3 Click the **Delete section** (  ) icon for the section you want to delete.
- 4 Click **OK** and then click **Publish Changes**.

#### Results

The section as well as all rules in that section are deleted.

## Lock Firewall Rule Sections

Firewall rule sections can be locked while making modifications, to prevent multiple users from simultaneously making changes to the same sections.

Firewall rule sections can be locked to prevent multiple users from simultaneously modifying the same section. The Enterprise Administrator can view and override all locks.

Security Administrator, Security Engineer and Enterprise Administrator user roles are able to lock and unlock their sections. Enterprise Administrator user roles have override capability - to unlock a section locked by any user of any role. Enterprise Administrators are also able to unlock other Enterprise Administrators. For more on user roles see [Managing User Rights](#).

Locked firewall sections cannot:

- Be merged with another section by another user.
- Have new rules added by another user.
- Be deleted by another user.
- Have rules dragged and dropped into them by another user.

#### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Security > Firewall**.
- 2 Click the section lock icon, enter the **Lock Section** name and comments, and click **LOCK**.

The section now displays a closed lock to indicate that it is locked.

## Unlock Firewall Rule Sections

Firewall rule sections can be locked while making modifications, to prevent multiple users from simultaneously making changes to the same sections.

Firewall rule sections can be locked to prevent multiple users from simultaneously modifying the same section. The Enterprise Administrator can view and override all locks.

Security Administrator, Security Engineer and Enterprise Administrator user roles are able to lock and unlock their sections. Enterprise Administrator user roles have override capability - to unlock a section locked by any user of any role. Enterprise Administrators are also able to unlock other Enterprise Administrators. For more on user roles see [Managing User Rights](#).

Locked firewall sections cannot:

- Be merged with another section by another user.
- Have new rules added by another user.
- Be deleted by another user.
- Have rules dragged and dropped into them by another user.

#### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Security > Firewall**.
- 2 To unlock a section, do one of the following:
  - Click the section lock icon, and then click **UNLOCK**. The section now displays an unlocked lock icon to indicate that it is unlocked.

- The number of locked sections is displayed above the firewall rule table. To view all locked sections, click the hyperlinked number next to **Locked**. To find the sections locked by you, filter rules by your name. Select the rule you want to unlock and click **UNLOCK**.

## Working with Firewall Rules

Distributed Firewall rules and Edge Firewall rules can be managed in a centralized manner on the Firewall tab. In a multi-tenant environment, providers can define high-level traffic flow rules on the centralized Firewall user interface.

Each traffic session is checked against the top rule in the Firewall table before moving down the subsequent rules in the table. The first rule in the table that matches the traffic parameters is enforced. Rules are displayed in the following order:

- 1 Rules defined in the Firewall user interface by users have the highest priority, and are enforced in top-to-bottom ordering with a per-virtual NIC level precedence.
- 2 Auto-plumbed rules (rules that enable control traffic to flow for Edge services).
- 3 Rules defined in the NSX Edge interface by users.
- 4 Service Composer rules - a separate section for each policy. You cannot edit these rules in the Firewall table, but you can add rules at the top of a security policy firewall rules section. If you do so, you must re-synchronize the rules in Service Composer. For more information, see [Chapter 18 Service Composer](#).
- 5 Default Distributed Firewall rules

Note that firewall rules are enforced only on clusters on which you have enabled firewall. For information on preparing clusters, see the *NSX Installation Guide*.

## Add a Firewall Rule

You add firewall rules 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.

### Procedure

#### 1 [Create a Firewall Rule](#)

You add firewall rules 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.

#### 2 [Add a Firewall Rule Source or Destination](#)

You can use IP addresses, vCenter objects, and NSX grouping objects as sources. You can also define sources and destinations and negate them. If no sources or destinations are defined, the source or destination is set to "any".

### 3 Add a Firewall Rule Service

For firewall rules you can create a new service group or use a predefined service group.

### 4 Specify a Firewall Rule Action and Logging

Firewall rules can be set to allow, block, or reject traffic from a specified source, destination, or service.

### 5 Define the Firewall Scope

Using the Applied To field, you can narrow down the scope at which you want to apply the rule.

### 6 Publish a Firewall Rule

After creating a new firewall rule, you have to publish it for changes to take effect.

## Create a Firewall Rule

You add firewall rules 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.

### Prerequisites

If you are adding a rule based on a VMware vCenter object, ensure that VMware Tools is installed on the virtual machines. See *NSX Installation Guide*.

VMs that are migrated from 6.1.5 to 6.2.3 do not have support for TFTP ALG. To enable TFTP ALG support after migrating, add and remove the VM from the exclusion list or restart the VM. A new 6.2.3 filter is created, with support for TFTP ALG.

---

### Note Identity Based Firewall Rule Prerequisites:

---

- One or more domains have been registered with NSX Manager. NSX Manager gets group and user information as well as the relationship between them from each domain that it is registered with. See [Register a Windows Domain with NSX Manager](#).
- A security group based on Active Directory objects has been created which can be used as the source or destination of the rule. See [Create a Security Group](#).
- Active Directory Server must be integrated with NSX Manager.
- Hosts must have DFW enabled and be upgraded to NSX 6.4.0.
- Guest machines must run updated VMware Tools.
- The version of the GI SVM must be 6.4 or later.
- The rule must be created in a new section of Firewall Rules.
- The rule must have **Enable User Identity at Source** selected.
- The **Applied to** field is not supported for rules for remote desktop access.

- ICMP is not supported for IDFW for RDSH.

---

**Note** Universal Firewall Rule Prerequisites:

---

In a cross-vCenter NSX environment, universal rules refer to the distributed firewall rules defined on the primary NSX Manager in the universal rules section. These rules are replicated on all secondary NSX Managers in your environment, which enables you to maintain a consistent firewall policy across vCenter boundaries. The primary NSX Manager can contain multiple universal sections for universal L2 rules and multiple universal sections for universal L3 rules. Universal sections are on top of all local and service composer sections. Universal sections and universal rules can be viewed but not edited on the secondary NSX Managers. The placement of the universal section with respect to the local section does not interfere with rule precedence.

Edge firewall rules are not supported for vMotion between multiple vCenter Servers.

**Table 10-4. Objects supported for universal firewall rules**

Source and Destination	Applied To	Service
<ul style="list-style-type: none"> <li>■ universal MAC set</li> <li>■ universal IP set</li> <li>■ universal security group, which can contain a universal security tag, an IP set, MAC set, or universal security group</li> </ul>	<ul style="list-style-type: none"> <li>■ universal security group, which can contain a universal security tag, IP set, MAC set, or universal security group</li> <li>■ universal logical switch</li> <li>■ Distributed Firewall - applies rules on all clusters on which Distributed Firewall is installed</li> </ul>	<ul style="list-style-type: none"> <li>■ pre-created universal services and service groups</li> <li>■ user created universal services and services groups</li> </ul>

---

Note that other vCenter objects are not supported for universal rules.

Make sure the state of NSX distributed firewall is not in backward compatibility mode. To check the current state, use the REST API call GET <https://<nsxmgr-ip>/api/4.0/firewall/globalroot-0/state>. If the current state is backward compatibility mode, you can change the state to forward by using the RES API call PUT <https://<nsxmgr-ip>/api/4.0/firewall/globalroot-0/state>. Do not try to publish a distributed firewall rule while the distributed firewall is in backward compatibility mode.

**Procedure**

- 1 In the vSphere Web Client, navigate to **Networking & Security > Security > Firewall**.
- 2 Ensure that you are in the **Configuration > General** tab to add an L3, L4, or L7 rule. Click the **Ethernet** tab to add an L2 rule.

If creating a universal firewall rule, create the rule in a universal rule section.

- 3 Point to the **Name** cell of the new rule and click .
- 4 Type a name for the new rule.

## Add a Firewall Rule Source or Destination

You can use IP addresses, vCenter objects, and NSX grouping objects as sources. You can also define sources and destinations and negate them. If no sources or destinations are defined, the source or destination is set to "any".

The following vCenter objects can be specified as the source or destination for a firewall rule:

**Table 10-5. Objects supported for firewall rules**

Source or Destination	Applied To
<ul style="list-style-type: none"> <li>■ cluster</li> <li>■ datacenter</li> <li>■ distributed port group</li> <li>■ IP set</li> <li>■ legacy port group</li> <li>■ logical switch</li> <li>■ resource pool</li> <li>■ security group</li> <li>■ vApp</li> <li>■ virtual machine</li> <li>■ vNIC</li> <li>■ IP address (IPv4 or IPv6)</li> </ul>	<ul style="list-style-type: none"> <li>■ All clusters on which Distributed Firewall has been installed (in other words, all clusters that have been prepared for network virtualization)</li> <li>■ All Edge gateways installed on prepared clusters</li> <li>■ cluster</li> <li>■ datacenter</li> <li>■ distributed port group</li> <li>■ Edge</li> <li>■ legacy port group</li> <li>■ logical switch</li> <li>■ security group</li> <li>■ virtual machine</li> <li>■ vNIC</li> </ul>

### Procedure

1 (Optional) Select objects to use in the firewall rule.

- a Click **Edit** in the source or destination column.
- b Select the object type from the **Object Type** drop-down menu.

You can create a new security group or IP set. Once you create the new object, it is added to the source or destination column by default. For information on creating a new security group or IP set, see [Chapter 22 Network and Security Objects](#).

- c Select one or more objects and click the arrow to move them to the **Selected Objects** column.

## 2 (Optional) Select IP addresses to use in the firewall rule.

Option	Description
NSX 6.4.1	<ol style="list-style-type: none"> <li>Click <b>Edit</b>  in the source or destination column, select <b>IP addresses</b>, and click <b>Add</b>.</li> <li>Enter one IP address. Both IPv4 and IPv6 addresses are valid.</li> <li>Click <b>Add</b> if you need to enter additional IP addresses.</li> </ol>
NSX 6.4.0	<ol style="list-style-type: none"> <li>Click IP () in the source column.</li> <li>Select <b>IPv4</b> or <b>IPv6</b>.</li> <li>Type the IP address.</li> </ol> <p>You can enter multiple IP addresses in a comma-separated list. The list can contain up to 255 characters.</p>

## 3 (Optional) Negate the sources or destinations defined in this rule.

If **Negate Source** is selected, the rule is applied to traffic coming from all sources except for the sources defined for this rule.

If **Negate Source** is not selected, the rule applies to traffic coming from the sources or destinations defined for this rule.

You can select **Negate Source** only if you have at least one source or destination defined.

Option	Description
NSX 6.4.1	<ol style="list-style-type: none"> <li>Click <b>Edit</b>  in the source column.</li> <li>Set <b>Negate Source</b> to On.</li> </ol>
NSX 6.4.0	<ol style="list-style-type: none"> <li>Click Edit () in the source column.</li> <li>Select the <b>Negate source</b> check box.</li> </ol>

## Add a Firewall Rule Service

For firewall rules you can create a new service group or use a predefined service group.

## Procedure

- 1 Select a pre-defined Service or Service Group to use in the firewall rule.

Option	Description
NSX 6.4.1	<ol style="list-style-type: none"> <li>a Point to the <b>Service</b> cell of the new rule and click .</li> <li>b Select the object type from the <b>Object Type</b> drop-down menu. You can create a new security group or IP set. Once you create the new object, it is added to the source or destination column by default. For information on creating a new security group or IP set, see <a href="#">Chapter 22 Network and Security Objects</a></li> <li>c Select one or more objects and click the arrow to move them to the <b>Selected Objects</b> column.</li> </ol>
NSX 6.4.0	<ol style="list-style-type: none"> <li>a Point to the <b>Service</b> cell of the new rule and click .</li> <li>b Select one or more objects and click . <p>You can create a new service or service group. Once you create the new object, it is added to the Selected Objects column by default.</p> </li> <li>c Click <b>OK</b>.</li> </ol>

- 2 Select a Port/Protocol to use in the firewall rule or define a new one.

Option	Description
NSX 6.4.1	<ol style="list-style-type: none"> <li>a Point to the <b>Service</b> cell of the new rule and click .</li> <li>b Select <b>Raw Port-Protocol</b>, and click <b>Add</b>.</li> <li>c Select the <b>Protocol</b> from the list and click <b>OK</b>.</li> </ol>
NSX 6.4.0	<ol style="list-style-type: none"> <li>a Point to the <b>Service</b> cell of the new rule and click .</li> <li>b Select the service protocol. <p>Distributed Firewall supports ALG (Application Level Gateway) for the following protocols: TFTP, FTP, ORACLE TNS, MS-RPC, and SUN-RPC.</p> <p>Edge supports ALG for FTP, TFTP, and SNMP_BASIC.</p> <p>Note: VMs that are migrated from 6.1.5 to 6.2.3 do not have support for TFTP ALG. To enable TFTP ALG support after migrating, add and remove the VM from the exclusion list or restart the VM. A new 6.2.3 filter is created, with support for TFTP ALG.</p> </li> <li>c Type the port number and click <b>OK</b>.</li> </ol>

In order to protect your network from ACK or SYN floods, you can set Service to TCP-all\_ports or UDP-all\_ports and set Action to Block for the default rule. For information on modifying the default rule, see [Edit the Default Distributed Firewall Rule](#).

## Specify a Firewall Rule Action and Logging

Firewall rules can be set to allow, block, or reject traffic from a specified source, destination, or service.

## Procedure

- 1 Point to the **Action** cell of the new rule and make appropriate selections as described in the table below.

Action	Results in
<b>Allow</b>	Allows traffic from or to the specified source(s), destination(s), and service(s).
<b>Block</b>	Blocks traffic from or to the specified source(s), destination(s), and service(s).
<b>Reject</b>	Sends reject message for unaccepted packets. RST packets are sent for TCP connections. ICMP messages with administratively prohibited code are sent for UDP, ICMP, and other IP connections.
<b>Log</b>	Logs all sessions matching this rule. Enabling logging can affect performance.
<b>Do not log</b>	Does not log sessions.

- 2 (Optional) Enable logging.

Option	Description
<b>NSX 6.4.1</b>	In the Logging column, click the <b>Log</b> button to on.
<b>NSX 6.4.0</b>	<ol style="list-style-type: none"> <li>a Point to the <b>Action</b> cell of the new rule and click </li> <li>b Select <b>Log</b> or <b>Do not Log</b>. Logging logs all sessions that match this rule and can affect performance.</li> </ol>

## Define the Firewall Scope

Using the Applied To field, you can narrow down the scope at which you want to apply the rule.

If the rule contains virtual machines/vNICs in the source and destination fields, you must add both the source and destination virtual machines/vNICs to **Applied To** for the rule to work correctly.

## Procedure

- ◆ In **Applied To**, define the scope at which this rule is applicable. Make appropriate selections as described in the table below and click **OK**. Note that if you adding a rule for remote desktop access, the **Applied To** field is not supported.

To apply a rule to	Do this
All prepared clusters in your environment	Select <b>Apply this rule on all clusters on which Distributed Firewall is installed</b> . After you click OK, the Applied To column for this rule displays <b>Distributed Firewall</b> .
All NSX Edge gateways in your environment	Select <b>Apply this rule on all the Edge gateways</b> . After you click OK or SAVE, the Applied To column for this rule displays <b>All Edges</b> . If both the above options are selected, the Applied To column displays <b>Any</b> .
One or more cluster, datacenter, distributed virtual port group, NSX Edge, network, virtual machine, vNIC, or logical switch	In the Available list, select one or more objects and click  .

## Publish a Firewall Rule

After creating a new firewall rule, you have to publish it for changes to take effect.

### Procedure

- ◆ Click **Publish** or **Publish Changes**. A new a rule is added at the top of the section. If the system-defined rule is the only rule in the section, the new rule is added above the default rule.

After a few moments, a message indicating whether the publish operation was successful is displayed. In case of any failures, the hosts on which the rule was not applied are listed. For additional details on a failed publish, navigate to **NSX Managers > NSX\_Manager\_IP\_Address > Monitor > System Events**.

If you want to add a rule at a specific place in a section, select a rule. In the No. column, click  and select **Add Above** or **Add Below**.

When you click **Publish Changes**, the firewall configuration is automatically saved. For information on reverting to an earlier configuration, see [Load a Saved Firewall Configuration](#).

### What to do next

- Deactivate a rule by clicking , or enable a rule by clicking .

- Display additional columns in the rule table by clicking  and selecting the appropriate columns.

Column Name	Information Displayed
Rule ID	Unique system generated ID for each rule
Log	Traffic for this rule is being logged or not
Stats	Clicking  shows the traffic related to this rule (traffic packets and size)
Comments	Comments for the rule

- Search for rules by typing text in the Search field.
- Move a rule up or down in the Firewall table.
- Merge sections by clicking the **Merge section** icon and selecting **Merge with above section** or **Merge with below section**.

## Edit the Default Distributed Firewall Rule

Default firewall settings apply to traffic that does not match any of the user-defined firewall rules. The Distributed Firewall default rule is displayed on the centralized firewall user interface, and the default rule for each NSX Edge is displayed at the NSX Edge level.

The default Distributed Firewall rule allows 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 or added to. However, you can change the Action element of the rule from Allow to Block or Reject, add comments for the rule, and indicate whether traffic for that rule should be logged.

In a cross-vCenter NSX environment the default rule is not a universal rule. Any changes to the default rule must be made on each NSX Manager.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Security > Firewall**.
- 2 Expand the Default Section and make the required changes.  
You can only edit **Action** and **Log**, or add comments to the default rule.

## Force Sync Distributed Firewall Rules

If you can't publish firewall rules to hosts, perform a force sync.

Force sync is used when you need to synchronize the firewall rules on an individual host with the NSX Manager.

### Procedure

- 1 In the vSphere Web client, navigate to **Networking & Security > Installation and Upgrade > Host Preparation**.

- 2 Select the cluster you want to force sync, then click **Actions** () > **Force Sync Services**.
- 3 Select **Firewall** from the services to force sync. Click **OK**.

The Firewall status changes to Busy while syncing.

## Firewall Rules with a Custom Layer 3 Protocol

Firewall rules can be created using a custom protocol number that is not listed in the protocols drop-down menu.

A firewall rule with a custom protocol number can be created on the distributed firewall or the NSX Edge firewall.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security** > **Security** > **Firewall**.
- 2 Ensure that you are in the **Configuration** > **General** tab to add an L3 rule. Click the **Add rule** () icon.
- 3 Point to the **Name** cell of the new rule and click .
- 4 Type a name for the new rule.
- 5 Specify the **Source** of the new rule. See [Add a Firewall Rule Source or Destination](#) for details.
- 6 Specify the **Destination** of the new rule. See [Add a Firewall Rule Source or Destination](#) for details.
- 7 Point to the **Service** cell of the new rule. Click the **Add Service** () icon.
- 8 Click **New Service** on the bottom left of the **Specify Service** window.
- 9 Enter the **Name** of the new protocol (such as OSPF).
- 10 From the Protocols drop-down menu select **L3\_OTHERS**.  
A **Protocol Number** field appears under the drop-down menu.
- 11 Enter the **Protocol Number** (such as 89 for OSPF).
- 12 Click **OK**.
- 13 Publish firewall rule. See [Publish a Firewall Rule](#) for details.

### Results

A firewall rule has been created using a custom protocol number.

## Save an Unpublished Configuration

You can add a rule and save the configuration without publishing it. You can then load and publish the saved configuration at a later time.

**Procedure**

- 1 Add a firewall rule. See [Add a Firewall Rule](#).
- 2 Click **Save Changes** or **Save**.
- 3 Enter a name and description to create a new configuration.
- 4 Click **Preserve Configuration** to preserve this change.

NSX can save up to 100 configurations. After this limit is exceeded, saved configurations marked with **Preserve Configuration** are preserved, while older non-preserved configurations are deleted to make room for preserved configurations.

- 5 Click **Save**
- 6 To make changes to a saved configuration:

Option	Description
In NSX 6.4.1 and later	<ol style="list-style-type: none"> <li>a Add another firewall rule.</li> <li>b Click <b>Save</b>.</li> <li>c Select <b>Update existing configuration</b></li> <li>d Click <b>Save</b>.</li> </ol>
In NSX 6.4.0	<ol style="list-style-type: none"> <li>a Add another firewall rule.</li> <li>b Click <b>Update Changes</b>.</li> <li>c Click <b>Save Changes</b>.</li> </ol>

- Click **Revert Changes** to go back to the configuration that existed before you added the rule. When you want to publish the rule you just added, click the **Load Configuration** icon, select the rule that you saved in step 3 and click **OK**.
- Click **Update Changes** to continue adding rules.

- 7 To revert changes you've made:

Option	Description
In NSX 6.4.1 and later	<ol style="list-style-type: none"> <li>a Click <b>Undo</b>.</li> </ol>
In NSX 6.4.0	<ol style="list-style-type: none"> <li>a Click <b>Revert Changes</b> to go back to the configuration that existed before you added the rule.</li> </ol>

## Load a Saved Firewall Configuration

You can load an autosaved or imported firewall configuration. If your current configuration contains rules managed by Service Composer, these are overridden after the import.

**Procedure**

- 1 In the vSphere Web Client, navigate to **Networking & Security > Security > Firewall**.
- 2 Ensure that you are in the **Configuration > General** tab to load an L3 firewall configuration. Click the **Ethernet** tab to load an L2 firewall configuration.

### 3 Load the saved configuration

Option	Description
NSX 6.4.1 and later	<ol style="list-style-type: none"> <li>Click <b>More</b>, then select <b>Load Saved Configuration</b>.</li> <li>Select the configuration to load and click <b>LOAD</b>.</li> </ol>
NSX 6.4.0	<ol style="list-style-type: none"> <li>Click the <b>Load configuration</b> () icon.</li> <li>Select the configuration to load and click <b>OK</b>.</li> </ol>

The current configuration is replaced by the selected configuration.

#### What to do next

If Service Composer rules in your configuration were overridden by the loaded configuration, click **Actions > Synchronize Firewall Rules** in the Security Policies tab within Service Composer.

## Filter Firewall Rules

You can use a wide number of criteria to filter your ruleset, which allows for easy rule modification.

Rules can be filtered by source or destination virtual machines or IP address, rule action, logging, rule name, comments, and rule ID. You can also filter rules based on a specific service, application, or a protocol.

#### Procedure

- In the Firewall tab, click the **Apply Filter** ( or ) icon.
- Enter the filtering criteria as appropriate.
- Click **Apply**.

Rules matching your filtering criteria are displayed.

#### What to do next

To display all rules again, clear the filters.

- In NSX 6.4.1 and later, click **Clear** in the **Filter** dialog box.
- In NSX 6.4.0, click the **Remove applied filter** () icon.

## Change the Order of a Firewall Rule

Firewall rules are applied in the order in which they exist in the rule table.

Rules are displayed (and enforced) in the following order:

- User-defined pre rules have the highest priority and are enforced in top-to-bottom ordering with a per-virtual NIC level precedence.
- Auto-plumbed rules.
- Local rules defined at an NSX Edge level.

- 4 Service Composer rules - a separate section for each policy. You cannot edit these rules in the Firewall table, but you can add rules at the top of a security policy firewall rules section. If you do so, you must re-synchronize the rules in Service Composer. For more information, see [Chapter 18 Service Composer](#).
- 5 Default Distributed Firewall rule

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 In the **Firewall** tab, select the rule that you want to move.
- 2 Click the **Move rule up** (≡↑) or **Move rule down** (≡↓) icon.
- 3 Click **Publish Changes**.

## Firewall Rule Behavior in Security Groups

Firewall rule behavior varies with different Security Groups.

**Table 10-6. Firewall Rule Behavior with RDSH and Non-RDSH Sections**

Enable User Identity Security Group (RDSH Section)	Identity Security Group (RDSH Section)	Any Security Group (Non-RDSH Section)
Source - SID based rules are preemptively pushed to hypervisor. Rule enforcement is on the first packet.	Source - IP based rules	Source - IP based rules
Destination - IP based rules	Destination - IP based rules	Destination - IP based rules
Applied To with Identity based Security Group - Applied to all hosts		User based Applied To
Applied To with Non-Identity based Security Group - User based Applied to		User based Applied to

## Firewall Rule Hit Count and Reset

The top of the firewall screen displays the rule summary panel, and is present on all three of the firewall tabs.

The rule summary panel displays:

- Total number of rules.
- Number of unpublished rules.
- Number of disabled rules.
- Total number of sections.
- Total number of locked sections.

To reset rule hit counts:

## Procedure

- 1 Click **MORE** in the upper right hand corner of the screen.
- 2 Click **Reset Rule Hit Count**, and then click **RESET**.
- 3 Click **RESET**.

All of the rule counts are reset to zero.

## Firewall Logs

Firewall generates and stores log files, such as audit logs, rules message logs, and system event logs. You must configure a syslog server for each cluster that has enabled the firewall . The syslog server is specified in the `Syslog.global.logHost` attribute.

---

**Recommendation** To collect firewall audit logs on a syslog server, ensure that you have upgraded the syslog server to the recent version. Preferably, configure a remote syslog-ng server to collect the firewall audit logs.

---

Firewall generates logs as described in the following table.

**Table 10-7. Firewall Logs**

Log Type	Description	Location
Rules message logs	Include all access decisions such as permitted or denied traffic for each rule if logging was enabled for that rule. Contains DFW packet logs for the rules where logging has been enabled.	<code>/var/log/dfwptlogs.log</code>
Audit logs	Include administration logs and Distributed Firewall configuration changes.	<code>/home/secureall/secureall/logs/vsm.log</code>
System event logs	Include Distributed Firewall configuration applied, filter created, deleted, or failed, and virtual machines added to security groups, and so on.	<code>/home/secureall/secureall/logs/vsm.log</code>
Data Plane/VMKernel logs	Capture activities related to a firewall kernel module (VSIP). It includes log entries for messages generated by the system.	<code>/var/log/vmkernel.log</code>
Message Bus Client/VSFWD logs	Capture activities of a firewall agent.	<code>/var/log/vsfwd.log</code>

---

**Note** The `vsm.log` file can be accessed by running the `show log manager` command from the NSX Manager Command Line Interface (CLI) and performing `grep` for the keyword `vsm.log`. This file is accessible only to the user or user group having the `root` privilege.

---

## Rules Message Logs

Rules message logs include all access decisions such as permitted or denied traffic for each rule, if logging was enabled for that rule. These logs are stored on each host in `/var/log/dfwpktlogs.log`.

Here are examples of firewall log message:

```
more /var/log/dfwpktlogs.log
2015-03-10T03:22:22.671Z INET match DROP domain-c7/1002 IN 242 UDP 192.168.110.10/138-
>192.168.110.255/138

more /var/log/dfwpktlogs.log
2017-04-11T21:09:59.877Z ESXi_FQDN dfwpktlogs: 50047 INET TERM domain-c1/1001 IN TCP RST
10.1.2.3/33491->10.4.5.6/10001 22/14 7684/1070
```

More examples:

```
2017-10-19T22:38:05.586Z 58734 INET match PASS domain-c8/1006 OUT 84 ICMP 172.18.8.121-
>172.18.8.119 RULE_TAG
2017-10-19T22:38:08.723Z 58734 INET match PASS domain-c8/1006 OUT 60 TCP 172.18.8.121/36485-
>172.18.8.119/22 S RULE_TAG
2017-10-19T22:38:18.785Z 58734 INET TERM domain-c8/1006 OUT ICMP 8 0 172.18.8.121-
>172.18.8.119 2/2 168/168 RULE_TAG
2017-10-19T22:38:20.789Z 58734 INET TERM domain-c8/1006 OUT TCP FIN 172.18.8.121/36484-
>172.18.8.119/22 44/33 4965/5009 RULE_TAG
```

In the following example:

- 1002 is the distributed firewall rule ID.
- domain-c7 is cluster ID in the vCenter managed object browser (MOB).
- 192.168.110.10/138 is the source IP address.
- 192.168.110.255/138 is the destination IP address.
- *RULE\_TAG* is an example of the text that you add in the **Tag** text box while adding or editing the firewall rule.

The following example shows the results of a ping 192.168.110.10 to 172.16.10.12.

```
tail -f /var/log/dfwpktlogs.log | grep 192.168.110.10

2015-03-10T03:20:31.274Z INET match DROP domain-c27/1002 IN 60 PROTO 1 192.168.110.10-
>172.16.10.12
2015-03-10T03:20:35.794Z INET match DROP domain-c27/1002 IN 60 PROTO 1 192.168.110.10-
>172.16.10.12
```

The following tables explain the text boxes in the firewall log message.

Table 10-8. Components of a log File Entry

Component	Value in example
Timestamp	2017-04-11T21:09:59
Firewall-specific portion	877Z ESXi_FQDN dfwpktlogs: 50047 INET TERM domain-c1/1001 IN TCP RST 10.1.2.3/33491->10.4.5.6/10001 22/14 7684/1070

Table 10-9. Firewall-Specific Portion of log File Entry

Entity	Possible Values
Filter hash	A number that can be used to get the filter name and other information.
AF Value	INET, INET6
Reason	<ul style="list-style-type: none"> <li>■ match: Packet matches a rule.</li> <li>■ bad-offset: Datapath internal error while getting packet.</li> <li>■ fragment: The non-first fragments after they are assembled to the first fragment.</li> <li>■ short: Packet too short (for example, not even complete to include an IP header, or TCP/UDP header).</li> <li>■ normalize: Malformed packets that do not have a correct header or a payload.</li> <li>■ memory: Datapath out of memory.</li> <li>■ bad-timestamp: Incorrect TCP timestamp.</li> <li>■ proto-cksum: Bad protocol checksum.</li> <li>■ state-mismatch: TCP packets that do not pass the TCP state machine check.</li> <li>■ state-insert: Duplicate connection is found.</li> <li>■ state-limit: Reached the maximum number of states that a datapath can track.</li> <li>■ SpoofGuard: Packet dropped by SpoofGuard.</li> <li>■ TERM: A connection is terminated.</li> </ul>
Action	<ul style="list-style-type: none"> <li>■ PASS: Accept the packet.</li> <li>■ DROP: Drop the packet.</li> <li>■ NAT: SNAT rule.</li> <li>■ NONAT: Matched the SNAT rule, but cannot translate the address.</li> <li>■ RDR: DNAT rule.</li> <li>■ NORDR: Matched the DNAT rule, but cannot translate the address.</li> <li>■ PUNT: Send the packet to a service VM running on the same hypervisor of the current VM.</li> <li>■ REDIRECT: Send the packet to network service running out of the hypervisor of the current VM.</li> <li>■ COPY: Accept the packet and make a copy to a service VM running on the same hypervisor of the current VM.</li> <li>■ REJECT: Reject the packet.</li> </ul>
Rule set and rule ID	<i>rule set/rule ID</i>
Direction	IN, OUT
Packet length	<i>length</i>

Table 10-9. Firewall-Specific Portion of log File Entry (continued)

Entity	Possible Values
Protocol	TCP, UDP, ICMP, or PROTO (protocol number) For TCP connections, the actual reason that a connection is terminated is indicated after the keyword TCP. If TERM is the reason for a TCP session, then an extra explanation appears in the PROTO row. The possible reasons for terminating a TCP connection include: RST (TCP RST packet), FIN (TCP FIN packet), and TIMEOUT (idle for too long) In the example above, it is <i>RST</i> . So it means that there is a <i>RST</i> packet in the connection that must be reset. For non-TCP connections (UDP, ICMP or other protocols), the reason for terminating a connection is only TIMEOUT.
Source IP address and port	<i>IP address/port</i>
Destination IP address and port	<i>IP address/port</i>
TCP flags	S (SYN), SA (SYN-ACK), A (ACK), P (PUSH), U (URGENT), F (FIN), R (RESET)
Number of packets	Number of packets. 22/14 - in packets / out packets
Number of bytes	Number of bytes. 7684/1070 - in bytes/ out bytes

To enable a rules message, log in to vSphere Web Client:

- 1 Navigate to **Networking & Security > Security > Firewall**.
- 2 Ensure that you are in the **General** tab.
- 3 Enable logging.

NSX Version	Procedure
NSX 6.4.1 and later	Click <b>More&gt;Enable&gt;Enable Rule Logs</b>
NSX 6.4.0	<ol style="list-style-type: none"> <li>1 Enable the <b>Log</b> column on the page.</li> <li>2 Enable logging for a rule by hovering over the Log table cell and clicking the pencil icon.</li> </ol>

**Note** If you want customized text to be displayed in the firewall log message, you can enable the **Tag** column and add the required text by clicking the pencil icon.

## Audit and System Event Logs

Audit logs include administration logs and Distributed Firewall configuration changes. These are stored in `/home/secureall/secureall/logs/vsm.log`.

System event logs include Distributed Firewall configuration applied, filter created, deleted, or failed, and virtual machines added to security groups, and so on. These logs are stored in `/home/secureall/secureall/logs/vsm.log`.

To view the audit and system event logs in the vSphere Web Client, navigate to **Networking & Security > System > Events**. In the **Monitor** tab, select the IP address of the NSX Manager.

You can use the firewall scenarios to get an understanding of the required end-to-end workflow.

This chapter includes the following topics:

- [Context-Aware Firewall Scenarios](#)
- [Configuring Application Identification](#)

## Context-Aware Firewall Scenarios

Context-aware firewall is intended specifically for EAST-WEST cases and not for general Web browsing classification. Applications can be limited to specific applications used in the data center such as SSH, FTP, TFTP, SQL, DNS, PCoIP, and so on.

Following are few uses cases for a context-aware firewall:

- Use Case 1: Don, the IT director of a team instructs his NSX administrator to restrict ALL HTTP traffic for a particular VM. Don wants to restrict this traffic irrespective of the port it comes from.
- Use Case 2: Robert, the IT lead of a team wants to restrict the HTTP traffic to a particular VM on the condition that the traffic does not come from TCP port 8080.
- Use Case 3: Now that there is a context-aware firewall, it can be extended to identity-based logins as well, such that an Active Directory user when logged into his virtual desktop, will only be able to access HTTP requests from port 8080. A manager wants his employee John to be able to access HTTP only from port 8080, and only when John is logged in to the Active Directory.

### Scenario 1: Allow Web Traffic on a Specific Port

You want to allow Web traffic only on port 80.

To create a context-aware firewall rule, perform the following steps:

- 1 [Add a Firewall Rule Section](#), if required.
- 2 [Add a Firewall Rule](#), say *HTTP to Web Server*.
- 3 Select the required Web server as the **Destination**.

- 4 [Create a Service](#) for application identification with the following parameters:

Parameter	Option
Layer	Layer7
App ID	HTTP
Protocol	TCP
Destination port	80

- 5 Change the default firewall rule to **Block**.
- 6 Publish the changes.

With the context-aware firewall rule, only traffic that is allowed is *Web* traffic on port *80*.

## Scenario 2: Allow SSH Traffic on Any Port

You want to allow SSH traffic on any port.

Perform the following steps to create context-aware firewall rule:

- 1 [Add a Firewall Rule Section](#), if required.
- 2 [Add a Firewall Rule](#), say *SSH to SSH Server*.
- 3 Select the required SSH server as the **Destination**.
- 4 [Create a Service](#) for application identification with the following parameters:

Parameter	Option
Layer	Layer7
App ID	SSH
Protocol	TCP
Destination Port	Keep the text box blank

- 5 Change the default firewall rule to **Block**.
- 6 Publish the changes.

With the context-aware firewall rule, only traffic that is allowed is *SSH* traffic on any port.

## Example: Example

For detailed steps on creating a context-aware firewall rule by using the vSphere Web Client, see [Example: Create a Context-Aware Firewall Rule](#).

# Configuring Application Identification

Application and protocol identity enables visibility across a large number of applications and enforcement based on application tiers such as Active Directory, DNS, HTTPS or MySQL.

Layer 7 application identification identifies which application a particular packet or flow is generated by, independent of the port that is being used.

Enforcement based on application identity enables users to allow or deny applications to run on any port, or to force applications to run on their standard port. Deep Packet Inspection (DPI) enables matching packet payload against defined patterns, commonly referred to as signatures. Layer 7 service objects can be used for port-independent enforcement or to create new service objects that leverage a combination of Layer 7 application identity, protocol and port. Layer 7 based service objects can be used in the firewall rule table and Service Composer, and application identification information is captured in Distributed Firewall logs, and Flow Monitoring and Application Rule Manager (ARM) when profiling an application.

## Procedure

- 1 In the vSphere Web Client, click **Networking & Security > Groups and Tags**.
- 2 Create service and specify Layer 7, App ID, protocol, and port. For port independent enforcement, this step can be skipped. See [Application ID GUIDs](#) and [Create a Service](#) for more details.
- 3 Create a new distributed firewall rule. In the service field, select the Layer 7 service you created in step 2. For port independent enforcement, select an App ID, see [Application ID GUIDs](#). See [Add a Firewall Rule](#) for details.
- 4 Save and publish the firewall rule.

# Identity Firewall Overview

# 12

Identity Firewall features allows an NSX administrator to create Active Directory user-based DFW rules.

A high level overview of the IDFW configuration workflow begins with preparing the infrastructure. This 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 (AD) user logs onto in order to apply DFW rules. There are two methods IDFW uses for logon detection: Guest Introspection (GI) and/or the Active Directory Event Log Scraper. Guest Introspection is deployed on ESXi clusters where IDFW virtual machines are running. When network events are generated by a user, a guest agent installed on the VM forwards the information through the Guest Introspection framework to the NSX Manager. The second option is the Active Directory event log scraper. Configure the Active Directory event log scraper in the NSX Manager to point at an instance of your Active Directory domain controller. NSX Manager will then pull events from the AD security event log. You can use both in your environment, or one or the other. When both the AD log scraper and Guest Introspection are used, Guest Introspection will take precedence. Note that if both the AD event log scraper and Guest Introspection are used, the two are mutually exclusive: if one of these stops working, the other does not begin to work as a back up.

Once the infrastructure is prepared, the administrator creates NSX Security Groups and adds the newly available AD Groups (referred to as Directory Groups). The administrator can then create Security Policies with associated firewall rules and apply those policies to the newly created Security Groups. Now, when a user logs into a desktop, the system will detect that event along with the IP address which is being used, look up the firewall policy that is associated with that user, and push those rules down. This works for both physical and virtual desktops. For physical desktops, AD event log scraper is also required to detect that a user is logged into a physical desktop.

Identity firewall can be used for micro-segmentation with remote desktop sessions (RDSH), enabling simultaneous logins by multiple users, user application access based on requirements, and the ability to maintain independent user environments. Identity Firewall with remote desktop sessions requires Active Directory.

For supported Windows operating systems see [Identity Firewall Tested and Supported Configurations](#). Note that Linux based operating systems are not supported for Identity Firewall.

This chapter includes the following topics:

- [Identity Firewall Workflow](#)
- [Identity Firewall Tested and Supported Configurations](#)

## Identity Firewall Workflow

Identity Firewall (IDFW) allows user-based distributed firewall rules (DFW).

User-based distributed firewall rules are determined by membership in an Active Directory (AD) group membership. IDFW monitors where AD users are logged in, and maps the login to an IP Address, which is used by DFW to apply firewall rules. Identity Firewall requires either guest introspection framework or active directory event log scraping. You can use both in your environment, or one or the other. When both the AD log scraper and Guest Introspection are used, Guest Introspection will take precedence. Note that if both the AD event log scraper and Guest Introspection are used, the two are mutually exclusive: if one of these stops working, the other does not begin to work as a back up.

AD group membership changes do not immediately take effect for logged in users using RDSH Identity Firewall rules, this includes enabling and disabling users, and deleting users. For changes to take effect, users must log off and then log back on. We recommend AD administrators force a log off when group membership is modified. This behavior is a limitation of Active Directory.

The Northbound flow of IDFW:

- 1 A user logs in to a VM.
- 2 A user login event is received by the NSX management plane.
- 3 The NSX management plane looks at the user and receives all of the Active Directory (AD) groups the user belongs to. The NSX management plane then sends group modify events for all of the affected AD groups.
- 4 For each Active Directory group all of the Security Groups (SG) including this AD group are flagged, and a job is added to the queue to process this change. Because a single SG can include multiple Active Directory groups, a single user login event will often trigger multiple processing events for the same SG. To address this, duplicate Security Group processing requests are removed.

The Southbound flow of IDFW:

- 1 A Security Group processing request is received. When a SG is modified, NSX updates all affected entities and triggers actions per IDFW rules.
- 2 NSX receives all of the Active Directory groups for a SG.
- 3 From Active Directory, NSX receives all of the users belonging to the AD groups.
- 4 The Active Directory users are associated with their IP addresses.

- 5 The IP address are mapped to vNICs, and then the vNICs are mapped to virtual machines (VMs). The resulting list of VMs is result of Security Group to VM translation.

---

**Note** Identity Firewall for RDSH is only supported with Windows Server 2016, Windows 2012 with VMware Tools 10.2.5 and later, and Windows 2012 R2 with VMware Tools 10.2.5 and later.

---

#### Procedure

- 1 Configure Active Directory Sync in NSX, see [Synchronize a Windows Domain with Active Directory](#). This is required to use Active Directory groups in Service Composer.
- 2 Prepare the ESXi cluster for DFW. See Prepare the Host Cluster for NSX in the *NSX Installation Guide*.
- 3 Configure Identity Firewall logon detection options. One or both of these options must be configured.

---

**Note** If you have a multi-domain AD architecture, and the log scrapper isn't accessible due to security constraints, use Guest Introspection to generate login and logout events.

---

- Configure Active Directory event log access. See [Register a Windows Domain with NSX Manager](#).
- Windows Guest OS with guest agent installed. This comes with a complete installation of VMware Tools™. Deploy Guest Introspection service to protected clusters. See [Install Guest Introspection on Host Clusters](#).

## Identity Firewall Tested and Supported Configurations

Directory servers and log scraping servers supported with IDFW.

**Table 12-1. Directory Servers and Versions**

Server/Version	Supported?
Windows Server 2016	Yes
Windows Server 2012	Yes
Windows Server 2012 R2	Yes
Windows Server 2008 R2	No
Windows Server 2008	No
Windows Server 2003	No
LDAP Servers other than Microsoft AD	No

**Table 12-2. Windows OS for RDSH desktops**

Server/Version	Supported?
Windows 2016	Yes
Windows 2012 with VMware Tools 10.2.5 and later	Yes
Windows 2012 R2 with VMware Tools 10.2.5 and later	Yes

Note that Identity Firewall with RDSH support requires Guest Introspection network drivers be installed.

**Table 12-3. Domain Synchronization Options**

Server/Version	Supported?
Domain synchronization with LDAP and LDAPS	Yes
Event log addition with CIFs and WMI	Yes
Domain sync with single rootDN	Yes
Domain sync with multiple RootDN OUs	6.4.0 and later
Domain sync with single subtree of OUs with level hierarchy	6.4.0 and later
Domain sync with multiple subtree of OUs	6.4.0 and later
Delete and re-add same domain with selective OU	6.4.0 and later
Add new subtree under synced OU	6.4.0 and later
Sync with selective BaseDN	6.4.0 and later
Sync with ignoring disabled users	Yes
Delta sync with changes in AD domain	Yes

**Table 12-4. Log Scraping Servers and Versions**

Server/Version	Supported?
Windows Server 2016	Yes
Windows Server 2012	Yes
Windows Server 2012 R2	Yes
Windows Server 2008 R2	Yes
Linux or other LDAP Implementations	No

## Log Scraping Limitations

- VM requires rebooting for an incoming login event if the following occur:
  - users are disabled or enabled
  - VM IP address change
  - re-adding the same domain with NSX Manager
- The event log queue for incoming login events is limited, and login events are not received if the log is full.

For more information about domain synchronization see [Synchronize a Windows Domain with Active Directory](#).

**Table 12-5. OS with Guest Introspection**

Server/Version	Supported?
Win-7 (32-bit, 64 bit)	Yes
Win-8 (64 bit)	Yes
Win-10 (32-bit, 64 bit)	Yes
Windows Server 2016	Yes.
Windows Server 2012	Yes
Windows Server 2008 R2	Yes
Linux Support	No

## Guest Introspection Limitations

- GI framework must be deployed to every cluster where IDFW VMs are running.
- A complete installation of VMware Tools™ must be installed on all Guest VMs.
- UDP sessions are not supported. Networking events are not generated for UDP sessions on Guest VMs.
- Linux GOS integration with Active Directory Server is not supported.

## Supported Microsoft Active Directory Configurations

Based on the standard and best practices design guides from Microsoft, <https://msdn.microsoft.com/en-us/library/bb727085.aspx>, following configurations of Active Directory Forests, Domains, Domain-Trees, Groups/Users are supported and tested for Identity Firewall:

**Table 12-6. Single forest, single domain and nesting of Active Directory groups and user configurations**

Scenarios	Supported?
Change user membership within domain	Yes
Circular group membership	Yes, supported from 6.2.8 and later
Nested group membership	Yes
Add and modify group name	Yes
Add and modify user name	Yes
Delete group and user	Yes
Disable and enable user	Yes

**Table 12-7. Single forest, single domain, subdomain tree**

Scenarios	Supported?
Users created in parent domain and part of groups in parent domain	Yes
Users created in child domain but part of groups in parent domain	No
Users created in child domain1 and membership is in child domain2	
Change user membership between two different domain (root and child)	Yes
Circular group membership	Yes, supported from 6.2.8 and later
Nested group membership in single domain, (Not supported for Cross Domain)	Yes
Add and modify group and username	Yes
Delete group and user	Yes
Disable and enable user	Yes

**Table 12-8. Single forest, single domain, subdomain tree**

Scenarios	Supported?
Change Domain Password after sync	Yes
Change IP address after sync	Yes
Rename domain controllers	Yes

**Table 12-8. Single forest, single domain, subdomain tree (continued)**

Scenarios	Supported?
Disconnect and reconnect network of domain and event log server during domain sync	Yes
Disconnect and reconnect network of domain and event log server after domain sync	Yes

---

**Note** Rule Enforcement Flow and Assumptions
 

---

- A user login event is processed only when a TCP session is initiated from a guest VM.
- User log out events are not sent or processed. Enforced ruleset remains until an 8-hour time span elapses since a user's last network activity, or a different user generates a TCP connection from the same VM. The system processes this as a log out from the previous user and a log on from the new user.
- Multi-user support is available with IDFW with RDSH in NSX 6.4.0 and later.
- RDSH VM logins are primarily handled by the context engine for rule enforcement. RDSH logins are only matched to firewall rules created with **Enable User Identity at Source**, and the rule must be created in a new section of Firewall Rules. If a user belongs to a non user identity at source security group and logs in to an RDSH VM, the login won't trigger any translation on the non user identity at source security group. An RDSH VM never belongs to any non user identity at source security groups.

# Working with Active Directory Domains

# 13

You can register one or more Windows domains with an NSX Manager and associated vCenter server. NSX Manager gets group and user information as well as the relationship between them from each domain that it is registered with. NSX Manager also retrieves Active Directory (AD) credentials.

Once NSX Manager retrieves AD credentials, you can create security groups based on user identity, create identity-based firewall rules, and run Activity Monitoring reports.

AD group membership changes do not immediately take effect for logged in users using RDSH Identity Firewall rules, this includes enabling and disabling users, and deleting users. For changes to take effect, users must log off and then log back on. We recommend AD administrators force a log off when group membership is modified. This behavior is a limitation of Active Directory.

---

**Important** Any changes made in Active Directory will NOT be seen on NSX Manager until a delta or full sync has been performed.

---

This chapter includes the following topics:

- [Register a Windows Domain with NSX Manager](#)
- [Synchronize a Windows Domain with Active Directory](#)
- [Edit a Windows Domain](#)
- [Enable Security Read-Only Log Access on Windows 2008](#)
- [Verifying Directory Privileges](#)

## Register a Windows Domain with NSX Manager

### 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.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > System > Users and Domains**.

- 2 Click the **Domains** tab, and then click the **Add domain** (+) icon.
- 3 In the **Add Domain** dialog box, enter the fully qualified domain name (for example, `eng.vmware.com`) and netBIOS name for the domain.  
  
To retrieve the netBIOS name for your domain, type `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.
- 4 When adding a child domain, select **Auto Merge**.
- 5 During sync, to filter out users that no longer have active accounts click **Ignore disabled users**.
- 6 Click **Next**.
- 7 In the LDAP Options page, specify the domain controller that the domain is to be synchronized with and select the protocol. See [Identity Firewall Tested and Supported Configurations](#) for more information about supported domain synchronization options.
- 8 Edit the port number, if required.
- 9 Enter the user credentials for the domain account. This user must be able to access the directory tree structure.
- 10 Click **Next**.
- 11 (Optional) In the Security Event Log Access page, select either **CIFS** or **WMI** for the connection method to access security event logs on the specified AD server. Change the port number if required. This step is used by Active Directory Event Log Scraper. See [Identity Firewall Workflow](#).

---

**Note** The event log reader looks for events with the following IDs from the AD Security event log: Windows 2008/2012: 4624, Windows 2003: 540. The event log server has a limit of 128 MB. When this limit is reached you may see Event ID 1104 in the Security Log Reader. See <https://technet.microsoft.com/en-us/library/dd315518> for more information.

---

- 12 Select **Use Domain Credentials** to use the LDAP server user credentials. To specify an alternate domain account for log access, un-select **Use Domain Credentials** and specify the user name and password.

The specified account must be able to read the security event logs on the Domain Controller specified in step 10.

- 13 Click **Next**.
- 14 In the Ready to Complete page, review the settings you entered.

**15** Click **Finish**.**Attention**

- If an error message appears stating that the Adding Domain operation failed for the entity because of a domain conflict, select Auto Merge. The domains will be created and the settings displayed below the domain list.

**Results**

The domain is created and its settings are displayed below the domain list.

**What to do next**

Verify that login events on the event log server are enabled.

You can add, edit, delete, enable, or disable LDAP servers by selecting the **LDAP Servers** tab in the panel below the domain list. You can perform the same tasks for event log servers by selecting the **Event Log Servers** tab in the panel below the domain list. Adding more than one Windows server (Domain Controllers, Exchange servers, or File Servers) as an event log server improves the user identity association.

**Note** If using IDFW, only AD Servers are supported.

## Synchronize a Windows Domain with Active Directory

By default, all registered domains are automatically synchronized with Active Directory every 3 hours. You can also synchronize on demand.

Through the vSphere Web Client UI, you can perform a force sync for Active Directory domains. A periodic sync is automatically performed once a week, and a delta sync every 3 hours. It is not possible to selectively sync sub-trees through the UI.

With NSX 6.4 and later it is possible to selectively sync active directory sub trees using API calls. The root domain cannot have any parent-child relationships and must have a valid directory distinguished name.

- `/api/1.0/directory/updateDomain` has an options to specify the folder under root domain. And there is an option to perform a force update `private boolean forceUpdate` .
- `/api/directory/verifyRootDN`. Verify that the list of rootDN doesn't have any parent-child relationships. Verify each rootDN is a valid active directory distinguished name.

**Procedure**

- 1 In the vSphere Web Client, navigate to **Networking & Security > System > Users and Domains**.

- Click the **Domains** tab, and then select the domain to be synchronized.

---

**Important** Any changes made in Active Directory will NOT be seen on NSX Manager until a delta or full sync has been performed.

---

- Select one of the following:

Click	To
	Perform a delta synchronization, where local AD objects that changed since the last synchronization event are updated
	Perform a full synchronization, where the local state of all AD objects is updated

## Edit a Windows Domain

You can edit the name, netBIOS name, primary LDAP server, and account credentials of a domain.

### Procedure

- In the vSphere Web Client, navigate to **Networking & Security > System > Users and Domains**.
- Click the **Domains** tab.
- Select a domain, and then click the **Edit domain** icon.
- Make the desired changes and click **Finish**.

## Enable Security Read-Only Log Access on Windows 2008

Read-only security log access is used by event log scraper in IDFW.

After creating a new user account, you must enable read-only security log access on a Windows 2008 server-based domain section to grant the user read-only access.

---

**Note** You must perform these steps on one Domain Controller of the domain, tree, or forest.

---

### Procedure

- Navigate to **Start > Administrative Tools > Active Directory Users and Computers**.
- In the navigation tree, expand the node that corresponds to the domain for which you want to enable security log access.
- Under the node that you just expanded, select the **Builtin** node.
- Double-click on **Event Log Readers** in the list of groups.
- Select the **Members** tab in the Event Log Readers Properties dialog box.

- 6 Click the **Add...** button.

The Select Users, Contacts, Computers, or Groups dialog appears.

- 7 If you previously created a group for the “AD Reader” user, select that group in the Select Users, Contacts, Computers, or Groups dialog. If you created only the user and you did not create a group, select that user in the Select Users, Contacts, Computers, or Groups dialog.
- 8 Click **OK** to close the Select Users, Contacts, Computers, or Groups dialog
- 9 Click **OK** to close the Event Log Readers Properties dialog.
- 10 Close the Active Directory Users and Computers window.

#### What to do next

After you have enabled security log access, verify directory privileges by following the steps in [Verifying Directory Privileges](#).

## Verifying Directory Privileges

Verify that the user account has the required privileges to read the security logs.

After you have created a new account and enabled security log access, you must verify the ability to read the security logs.

#### Prerequisites

Enable security log access. See [Enable Security Read-Only Log Access on Windows 2008](#).

#### Procedure

- 1 From any workstation that is part of the domain, log on to the domain as an administrator.
- 2 Navigate to **Start > Administrative Tools > Event Viewer**.
- 3 Select **Connect to Another Computer...** from the **Action** menu. The Select Computer dialog box appears. (Note that you must do this even if you are already logged on to the machine for which you plan to view the event log.)
- 4 Select the **Another computer** radio button, if it is not already selected.
- 5 In the text field adjacent to the **Another computer** radio button, enter the name of the Domain Controller. Alternatively, click the **Browse...** button and then select the Domain Controller.
- 6 Select the **Connect as another user** check box.
- 7 Click the **Set User...** button. The Event Viewer dialog box appears.
- 8 In the **User name** field, enter the user name for the user that you created.
- 9 In the **Password** field, enter the password for the user that you created
- 10 Click **OK**
- 11 Click **OK** again.

- 12 Expand the **Windows Logs** node in the navigation tree.
- 13 Under the **Windows Logs** node, select the Security node. If you can see log events then the account has the required privileges.

SpoofGuard protects against IP spoofing by maintaining a reference table of VM names and IP addresses. SpoofGuard maintains this reference table by using the IP addresses that the NSX Manager retrieves from VMware Tools when a VM initially starts.

After synchronizing with the vCenter Server, NSX Manager collects the IP addresses of all vCenter guest virtual machines from VMware Tools on each virtual machine. If a virtual machine has been compromised, the IP address can be spoofed and malicious transmissions can bypass firewall policies.

SpoofGuard is inactive by default, and you must explicitly enable it on each logical switch or VDS port-group. When a VM IP address change is detected, the Distributed Firewall (DFW) blocks the traffic from or to this VM until you approve this new IP address.

You create a SpoofGuard policy for specific networks that allows you to authorize the IP addresses reported by VMware Tools and alter them if necessary to prevent spoofing. SpoofGuard inherently trusts the MAC addresses of virtual machines collected from the VMX files and vSphere SDK. Operating separately from Firewall rules, you can use SpoofGuard to block traffic determined to be spoofed.

---

**Important** SpoofGuard works only when Distributed Firewall (DFW) is enabled.

---

SpoofGuard supports both IPv4 and IPv6 addresses. The SpoofGuard policy supports multiple IP addresses assigned to a vNIC when using VMwareTools and DHCP snooping. ARP snooping supports up to 128 addresses discovered per VM, per vNIC. The SpoofGuard policy monitors and manages the IP addresses reported by your virtual machines in one of the following modes.

### **Automatically Trust IP Assignments On Their First Use**

This mode allows all traffic from your virtual machines to pass while building a table of vNIC-to-IP address assignments. You can review this table at your convenience and make IP address changes. This mode automatically approves all IPv4 and IPv6 addresses that are first seen on a vNIC.

### **Manually Inspect and Approve All IP Assignments Before Use**

This mode blocks all traffic until you approve each vNIC-to-IP address assignment. In this mode, multiple IPv4 addresses can be approved.

---

**Note** SpoofGuard inherently allows DHCP requests regardless of enabled mode. However, if in manual inspection mode, traffic does not pass until the DHCP-assigned IP address has been approved.

---

SpoofGuard includes a system-generated default policy that applies to port groups and logical networks not covered by the other SpoofGuard policies. A newly added network is automatically added to the default policy until you add the network to an existing policy or create a new policy for it.

SpoofGuard is one of the ways that an NSX distributed firewall policy can determine the IP address of a virtual machine. For information, see [IP Discovery for Virtual Machines](#).

This chapter includes the following topics:

- [Create a SpoofGuard Policy](#)
- [Approve IP Addresses](#)
- [Change an IP Address](#)
- [Clear an IP Address](#)

## Create a SpoofGuard Policy

You can create a SpoofGuard policy to specify the operation mode for specific networks. The system-generated (default) policy applies to port groups and logical switches not covered by existing SpoofGuard policies.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Security > SpoofGuard**.
- 2 Click **Add**.
- 3 Enter a name for the policy.
- 4 **Enable** or **Disable** the policy.
- 5 Select one of the following **Operation Mode**:

Option	Description
<b>Automatically Trust IP Assignments on Their First Use</b>	Select this option to trust all IP assignments upon initial registration with the NSX Manager.
<b>Manually Inspect and Approve All IP Assignments Before Use</b>	Select this option to require manual approval of all IP addresses. All traffic to and from unapproved IP addresses is blocked.

- Click **Allow local address as valid address in this namespace** to allow local IP addresses in your setup.

When you power on a virtual machine and it is unable to connect to the DHCP server, a local IP address is assigned to it. This local IP address is considered valid only if the SpoofGuard mode is set to **Allow local address as valid address in this namespace**. Otherwise, the local IP address is ignored.

- Click **Next**.
- Select the object type this policy should apply to, then select the objects you want.
  - In NSX 6.4.0, click the **Add** icon. Select the object type this policy should apply to, then select the objects you want.

A port group or logical switch can belong to only one SpoofGuard policy.

- Click **OK** or **Finish**.

#### What to do next

You can edit a policy by clicking the **Edit** icon and delete a policy by clicking the **Delete** icon.

## Approve IP Addresses

If you set SpoofGuard to require manual approval of all IP address assignments, you must approve IP address assignments to allow traffic from those virtual machines to pass.

#### Procedure

- In **SpoofGuard**, select a policy.  
Policy details are displayed below the policy table.
- In NSX 6.4.1 and later, select one of the option links from the drop-down menu, or **All**.

Option	Description
<b>Active vNICs</b>	List of all validated IP addresses
<b>Pending Approval vNICs</b>	IP address changes that require approval before traffic can flow to or from these virtual machines
<b>Inactive vNICs</b>	List of IP addresses where the current IP address does not match the published IP address
<b>vNICs with Duplicate IP</b>	IP addresses that are duplicates of an existing assigned IP address within the selected datacenter

- In NSX 6.4.0, select **View**, and click one of the option links.

Option	Description
<b>Active Virtual NICs</b>	List of all validated IP addresses
<b>Active Virtual NICs Since Last Published</b>	List of IP addresses that have been validated since the policy was last updated

Option	Description
<b>Virtual NICs IP Required Approval</b>	IP address changes that require approval before traffic can flow to or from these virtual machines
<b>Virtual NICs with Duplicate IP</b>	IP addresses that are duplicates of an existing assigned IP address within the selected datacenter
<b>Inactive Virtual NICs</b>	List of IP addresses where the current IP address does not match the published IP address
<b>Unpublished Virtual NICs IP</b>	List of virtual machines for which you have edited the IP address assignment but have not yet published

4 Do one of the following:

- To approve a single IP address, click **Approve** next to the IP address.
- To approve multiple IP addresses, select the appropriate vNICs and then click **Approve IPs**.

## Change an IP Address

You can change the IP address assigned to a MAC address to correct the assigned IP address.

**Note** SpoofGuard accepts a unique IP address from virtual machines. However, you can assign an IP address only once. An approved IP address is unique across NSX. Duplicate approved IP addresses are not allowed.

### Procedure

- 1 In **SpoofGuard**, select a policy.
- 2 In NSX 6.4.1 and later, select one of the option links from the drop-down menu, or **All**.

Option	Description
<b>Active vNICs</b>	List of all validated IP addresses
<b>Pending Approval vNICs</b>	IP address changes that require approval before traffic can flow to or from these virtual machines
<b>Inactive vNICs</b>	List of IP addresses where the current IP address does not match the published IP address
<b>vNICs with Duplicate IP</b>	IP addresses that are duplicates of an existing assigned IP address within the selected datacenter

- 3 For NSX 6.4.0, select **View**, and click one of the option links.

Option	Description
<b>Active Virtual NICs</b>	List of all validated IP addresses
<b>Active Virtual NICs Since Last Published</b>	List of IP addresses that have been validated since the policy was last updated

Option	Description
<b>Virtual NICs IP Required Approval</b>	IP address changes that require approval before traffic can flow to or from these virtual machines
<b>Virtual NICs with Duplicate IP</b>	IP addresses that are duplicates of an existing assigned IP address within the selected datacenter
<b>Inactive Virtual NICs</b>	List of IP addresses where the current IP address does not match the published IP address
<b>Unpublished Virtual NICs IP</b>	List of virtual machines for which you have edited the IP address assignment but have not yet published

#### 4 Add an IP address.

Option	Description
<b>NSX 6.4.1</b>	Click <b>Add IP</b> , and add an IP address.
<b>NSX 6.4.0</b>	Click the pencil icon next to an <b>Approved IP</b> address, then click + and add a new IP address.

#### 5 To delete an incorrect IP address, select **Clear**.

#### 6 Click **OK**.

## Clear an IP Address

You clear an approved IP address assignment from a SpoofGuard policy.

### Procedure

- 1 In **SpoofGuard**, select a policy.
- 2 In NSX 6.4.1 and later, select one of the option links, or **All**.

Option	Description
<b>Active vNICs</b>	List of all validated IP addresses
<b>Pending Approval vNICs</b>	IP address changes that require approval before traffic can flow to or from these virtual machines
<b>Inactive vNICs</b>	List of IP addresses where the current IP address does not match the published IP address
<b>vNICs with Duplicate IP</b>	IP addresses that are duplicates of an existing assigned IP address within the selected datacenter

#### 3 For NSX 6.4.0, select **View**, and click one of the option links.

Option	Description
<b>Active Virtual NICs</b>	List of all validated IP addresses
<b>Active Virtual NICs Since Last Published</b>	List of IP addresses that have been validated since the policy was last updated

Option	Description
<b>Virtual NICs IP Required Approval</b>	IP address changes that require approval before traffic can flow to or from these virtual machines
<b>Virtual NICs with Duplicate IP</b>	IP addresses that are duplicates of an existing assigned IP address within the selected datacenter
<b>Inactive Virtual NICs</b>	List of IP addresses where the current IP address does not match the published IP address
<b>Unpublished Virtual NICs IP</b>	List of virtual machines for which you have edited the IP address assignment but have not yet published

- 4 Select **Clear** or **Clear Approved IPs** to clear an IP address.

# Virtual Private Networks (VPN)

# 15

NSX Edge supports several types of VPNs. SSL VPN-Plus allows remote users to access private corporate applications. IPsec VPN offers site-to-site connectivity between an NSX Edge instance and remote sites. L2 VPN allows you to extend your datacenter by allowing virtual machines to retain network connectivity across geographical boundaries.

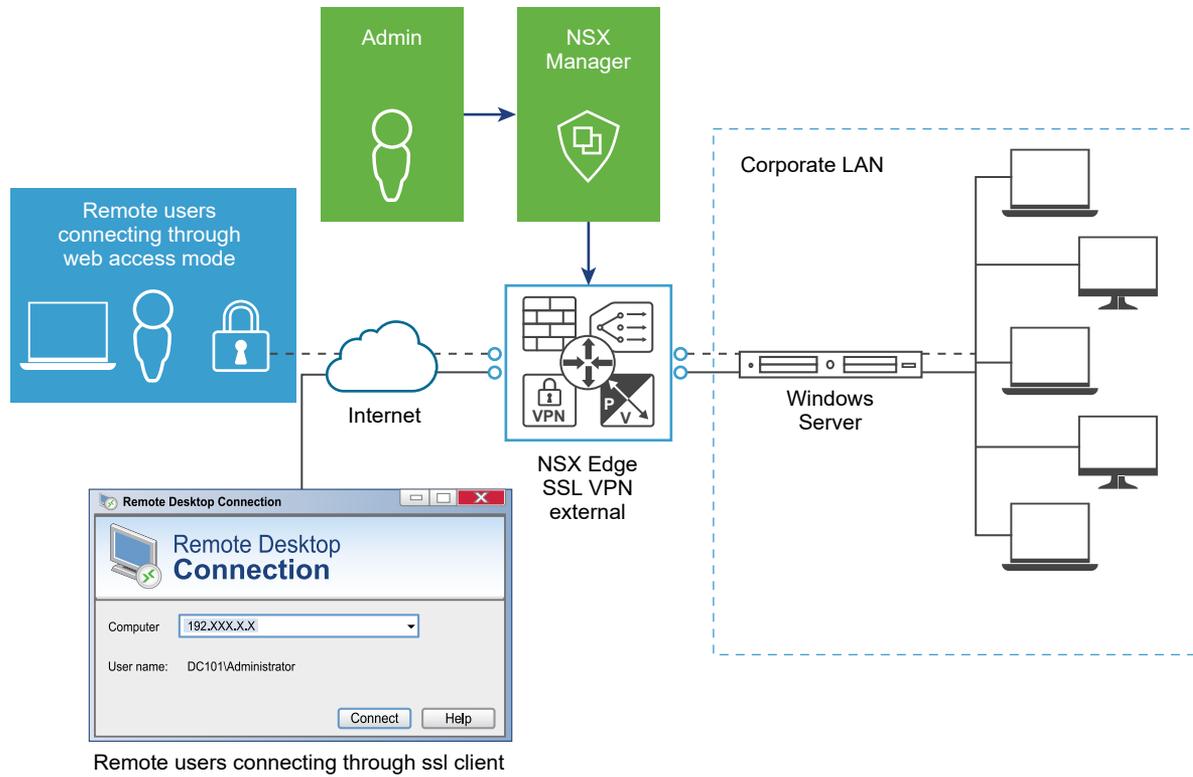
You must have a working NSX Edge instance before you can use VPN. For information on setting up NSX Edge, see [NSX Edge Configuration](#).

This chapter includes the following topics:

- [SSL VPN-Plus Overview](#)
- [IPsec VPN Overview](#)
- [L2 VPN Overview](#)

## SSL VPN-Plus Overview

With SSL VPN-Plus, remote users can connect securely to private networks behind a NSX Edge gateway. Remote users can access servers and applications in the private networks.



The following client operating systems are supported.

Operating System	Supported Versions
Windows	8 , 10 (including Windows 10 Secure Boot option enabled)
Mac OS Sierra	10.12.6
Mac OS High Sierra	10.13.4
Mac OS Mojave	10.14.2 to 10.14.6 (supported in NSX 6.4.4 and later)
Mac OS Catalina	10.15, 10.15.3, 10.15.4 (supported in NSX 6.4.7 and later) 10.15.6, 10.15.7 (supported starting in NSX 6.4.10)
Mac OS Big Sur	11.2 and later (supported starting in NSX 6.4.10)
Linux Fedora	26, 28

Operating System	Supported Versions
Linux CentOS	6.0, 7.5
Linux Ubuntu	18.04 (supported in NSX 6.4.6 and later)

### Important

- SSL VPN-Plus Client is not supported on computers that use ARM-based processors.
- In SSL VPN-Plus Client on Windows, the "auto-reconnect" feature does not work as expected when the Npcap loopback adapter is "enabled". This loopback adapter interferes with the function of the Npcap driver on a Windows computer. Make sure that the latest version of the Npcap driver (0.9983 or later) is installed on your Windows computer. This version of the driver does not require the loopback adapter for packet captures.
- Linux TCL, TK, and Network Security Services (NSS) libraries are required for the UI to work.

## Configure Network Access SSL VPN-Plus

In network access mode, a remote user can access private networks after downloading and installing an SSL client.

### Prerequisites

The SSL VPN gateway requires port 443 to be accessible from external networks and the SSL VPN client requires the NSX Edge gateway IP and port 443 to be reachable from client system.

### Add SSL VPN-Plus Server Settings

You must add SSL VPN server settings to enable SSL on a NSX Edge interface.

#### Procedure

- 1 In the **SSL VPN-Plus** tab, **Server Settings** from the left panel.
- 2 Click **Change**.
- 3 Select the IPv4 or IPv6 address.
- 4 Edit the port number if required. This port number is required to configure the installation package.
- 5 Select one or more encryption methods or ciphers.

**Note** If any of the following GCM ciphers is configured on the SSL VPN server, backward compatibility issue can occur in some browsers:

- AES128-GCM-SHA256
- ECDHE-RSA-AES128-GCM-SHA256
- ECDHE-RSA-AES256-GCM-SHA38

- (Optional) From the Server Certificates table, use the default server certificate, or deselect the **Use Default Certificate** check box and click the server certificate that you want to add.

---

**Restriction**

- SSL VPN-Plus service supports only RSA certificates.
  - SSL VPN-Plus service supports server certificate that is signed only by the Root CA. Server certificate signed by an Intermediate CA is not supported.
- 

- Click **OK**.

## Add an IP Pool

The remote user is assigned a virtual IP address from the IP pool that you add.

### Procedure

- In the **SSL Vpn-Plus** tab, select **IP Pools** from the left panel.
- Click the **Add** (+) icon.
- Type the begin and end IP address for the IP pool.
- Type the netmask of the IP pool.
- Type the IP address which is to add the routing interface in the NSX Edge gateway.
- (Optional) Type a description for the IP pool.
- Select whether to enable or disable the IP pool.
- (Optional) In the **Advanced** panel, type the DNS name.
- (Optional) Type the secondary DNS name.
- Type the connection-specific DNS suffix for domain based host name resolution.
- Type the WINS server address.
- Click **OK**.

## Add a Private Network

Add the network that you want the remote user to be able to access.

### Procedure

- In the **SSL VPN-Plus** tab, select **Private Networks** from the left panel.
- Click the **Add** (+) icon
- Type the private network IP address.
- Type the netmask of the private network.
- (Optional) Type a description for the network.

- 6 Specify whether you want to send private network and internet traffic over the SSL VPN-Plus enabled NSX Edge or directly to the private server by bypassing the NSX Edge.
- 7 If you selected **Send traffic over the tunnel**, select **Enable TCP Optimization** to optimize the internet speed.

Conventional full-access SSL VPNs tunnel sends TCP/IP data in a second TCP/IP stack for encryption over the internet. This results in application layer data being encapsulated twice in two separate TCP streams. When packet loss occurs (which happens even under optimal internet conditions), a performance degradation effect called TCP-over-TCP meltdown occurs. In essence, two TCP instruments are correcting a single packet of IP data, undermining network throughput and causing connection timeouts. TCP Optimization eliminates this TCP-over-TCP problem, ensuring optimal performance.

- 8 When optimization is enabled, specify the port numbers for which traffic should be optimized. Traffic for remaining ports for that specific network will not be optimized.

---

**Note** Traffic for all ports are optimized, if port numbers are not specified.

---

When TCP traffic is optimized, the TCP connection is opened by the SSL VPN server on behalf of the client. Because the TCP connection is opened by the SSLVPN server, the first automatically generated rule is applied, which allows all connections opened from the Edge to get passed. Traffic that is not optimized will be evaluated by the regular Edge firewall rules. The default rule is allow any any.

- 9 Specify whether you want to enable or disable the private network.
- 10 Click **OK**.

#### What to do next

Add a corresponding firewall rule to allow the private network traffic.

## Add Authentication

In addition to local user authentication, you can add an external authentication server (AD, LDAP, Radius, or RSA) which is bound to the SSL gateway. All users with accounts on the bound authentication server will be authenticated.

The maximum time to authenticate over SSL VPN is 3 minutes. This is because non-authentication timeout is 3 minutes and is not a configurable property. So in scenarios where AD authentication timeout is set to more than 3 minutes or there are multiple authentication servers in chain authorization and the time taken for user authentication is more than 3 minutes, you will not be authenticated.

#### Procedure

- 1 In the **SSL VPN-Plus** tab, select **Authentication** from the left panel.
- 2 Click the **Add** (+) icon.

- 3 Select the type of authentication server.
- 4 Depending on the type of authentication server you selected, complete the following fields.
  - ◆ AD authentication server

Table 15-1. AD Authentication Server Options

Option	Description
<b>Enable SSL</b>	Enabling SSL establishes an encrypted link between a web server and a browser.  <b>Note</b> There might be issues if you do not enable SSL and try to change password using SSL VPN-Plus tab or from client machine later.
<b>IP Address</b>	IP address of the authentication server.
<b>Port</b>	Displays default port name. Edit if required.
<b>Timeout</b>	Period in seconds within which the AD server must respond.
<b>Status</b>	Select <b>Enabled</b> or <b>Disabled</b> to indicate whether the server is enabled.
<b>Search base</b>	Part of the external directory tree to search. The search base may be something equivalent to the organizational unit (OU), domain controller (DC), or domain name (AD) of external directory. Examples: <ul style="list-style-type: none"> <li>■ OU=Users,DC=aslan,DC=local</li> <li>■ OU=VPN,DC=aslan,DC=local</li> </ul>
<b>Bind DN</b>	User on the external AD server permitted to search the AD directory within the defined search base. Most of the time, the bind DN is permitted to search the entire directory. The role of the bind DN is to query the directory using the query filter and search base for the DN (distinguished name) for authenticating AD users. When the DN is returned, the DN and password are used to authenticate the AD user. Example: CN=ldap.edge,OU=users,OU=Datacenter Users,DC=aslan,DC=local
<b>Bind Password</b>	Password to authenticate the AD user.
<b>Retype Bind Password</b>	Retype the password.
<b>Login Attribute Name</b>	Name against which the user ID entered by the remote user is matched with. For Active Directory, the login attribute name is <b>sAMAccountName</b> .
<b>Search Filter</b>	Filter values by which the search is to be limited. The search filter format is <i>attribute operator value</i> . If you need to limit the search base to a specific group in the AD and not allow searching across the entire OU, then <ul style="list-style-type: none"> <li>■ Do not put group name inside the search base, only put OU and DC.</li> <li>■ Do not put both <i>objectClass</i> and <i>memberOf</i> inside the same search filter string. Example of correct format for the search filter: memberOf=CN=VPN_Users,OU=Users,DC=aslan,DC=local</li> </ul>

Table 15-1. AD Authentication Server Options (continued)

Option	Description
Use this server for secondary authentication	If selected, this AD server is used as the second level of authentication.
Terminate Session if authentication fails	When selected, the session is ended if authentication fails.

## ◆ LDAP authentication server

Table 15-2. LDAP Authentication Server Options

Option	Description
Enable SSL	Enabling SSL establishes an encrypted link between a web server and a browser.
IP Address	IP address of the external server.
Port	Displays default port name. Edit if required.
Timeout	Period in seconds within which the AD server must respond.
Status	Select <b>Enabled</b> or <b>Disabled</b> to indicate whether the server is enabled.
Search base	Part of the external directory tree to search. The search base may be something equivalent to the organization, group, or domain name (AD) of external directory.
Bind DN	User on the external server permitted to search the AD directory within the defined search base. Most of the time, the bind DN is permitted to search the entire directory. The role of the bind DN is to query the directory using the query filter and search base for the DN (distinguished name) for authenticating AD users. When the DN is returned, the DN and password are used to authenticate the AD user.
Bind Password	Password to authenticate the AD user.
Retype Bind Password	Retype the password.
Login Attribute Name	Name against which the user ID entered by the remote user is matched with. For Active Directory, the login attribute name is <b>sAMAccountName</b> .
Search Filter	Filter values by which the search is to be limited. The search filter format is <i>attribute operator value</i> .
Use this server for secondary authentication	If selected, this server is used as the second level of authentication.
Terminate Session if authentication fails	When selected, the session is ended if authentication fails.

## ◆ RADIUS authentication server

RADIUS authentication is disabled in FIPS mode.

**Table 15-3. RADIUS authentication server options**

Option	Description
IP Address	IP address of the external server.
Port	Displays default port name. Edit if required.
Timeout	Period in seconds within which the AD server must respond.
Status	Select <b>Enabled</b> or <b>Disabled</b> to indicate whether the server is enabled.
Secret	Shared secret specified while adding the authentication agent in the RSA security console.
Retype secret	Retype the shared secret.
NAS IP Address	IP address to be configured and used as RADIUS attribute 4, NAS-IP-Address, without changing the source IP address in the IP header of the RADIUS packets.
Retry Count	Number of times the RADIUS server is to be contacted if it does not respond before the authentication fails.
Use this server for secondary authentication	If selected, this server is used as the second level of authentication.
Terminate Session if authentication fails	When selected, the session is ended if authentication fails.

◆ RSA-ACE authentication server

RSA authentication is disabled in FIPS mode.

**Table 15-4. RSA-ACE authentication server options**

Option	Description
Timeout	Period in seconds within which the AD server must respond.
Configuration File	Click <b>Browse</b> to select the <code>sdconf.rec</code> file that you downloaded from the RSA Authentication Manager.
Status	Select <b>Enabled</b> or <b>Disabled</b> to indicate whether the server is enabled.
Source IP Address	IP address of the NSX Edge interface through which the RSA server is accessible.

Table 15-4. RSA-ACE authentication server options (continued)

Option	Description
Use this server for secondary authentication	If selected, this server is used as the second level of authentication.
Terminate Session if authentication fails	When selected, the session is ended if authentication fails.

## ◆ Local authentication server

Table 15-5. Local authentication server options

Option	Description
Enable password policy	If selected, defines a password policy. Specify the required values.
Enable password policy	<p>If selected, defines an account lockout policy. Specify the required values.</p> <ol style="list-style-type: none"> <li>1 In Retry Count, type the number of times a remote user can try to access his or her account after entering an incorrect password.</li> <li>2 In Retry Duration, type the time period in which the remote user's account gets locked on unsuccessful login attempts.</li> </ol> <p>For example, if you specify Retry Count as 5 and Retry Duration as 1 minute, the remote user's account will be locked if he makes 5 unsuccessful login attempts within 1 minute.</p> <ol style="list-style-type: none"> <li>3 In Lockout Duration, type the time period for which the user account remains locked. After this time, the account is automatically unlocked.</li> </ol>
Status	Select <b>Enabled</b> or <b>Disabled</b> to indicate whether the server is enabled.
Use this server for secondary authentication	If selected, this server is used as the second level of authentication.
Terminate Session if authentication fails	When selected, the session is ended if authentication fails.

- 5 (Optional) Add client certificate authentication.
  - a Next to **Certificate Authentication**, click **Change**.
  - b Select the **Enable client certificate authentication** check box.
  - c Select a client certificate issued by the Root CA and click **OK**.

---

#### Restriction

- On the SSL VPN-Plus Web Portal and SSL VPN-Plus full access client (PHAT client), client or user certificate that is signed only by the Root CA is supported. Client certificate signed by an Intermediate CA is not supported.
  - Client certificate authentication is supported only on an SSL VPN-Plus client that is installed on a Windows computer. This authentication is not supported on an SSL VPN-Plus client that is installed on Linux and Mac computers.
- 

## Add Installation Package

Create an installation package of the SSL VPN-Plus client for the remote user.

### Procedure

- 1 In the **SSL VPN-Plus** tab, select **Installation Package** from the left panel.
- 2 Click the **Add (+)** icon.
- 3 Type a profile name for the installation package.
- 4 In **Gateway**, type the IP address or FQDN of the public interface of NSX Edge.  
This IP address or FQDN is binded to the SSL client. When the client is installed, this IP address or FQDN is displayed on the SSL client.
- 5 Type the port number that you specified in the server settings for SSL VPN-Plus. See [Add SSL VPN-Plus Server Settings](#).
- 6 (Optional) To bind additional NSX Edge uplink interfaces to the SSL client,
  - a Click the **Add (+)** icon.
  - b Type the IP address and port number.
  - c Click **OK**.
- 7 The installation package is created for Windows operating system by default. Select Linux or Mac to create an installation package for Linux or Mac operating systems as well.
- 8 (Optional) Enter a description for the installation package.
- 9 Select **Enable** to display the installation package on the Installation Package page.

10 Select the following options as appropriate.

Option	Description
<b>Start client on logon</b>	The SSL VPN client is started when the remote user logs on to his system.
<b>Allow remember password</b>	Enables the option.
<b>Enable silent mode installation</b>	Hides installation commands from remote user.
<b>Hide SSL client network adapter</b>	Hides the VMware SSL VPN-Plus Adapter, which is installed on the remote user's computer along with the SSL VPN installation package.
<b>Hide client system tray icon</b>	Hides the SSL VPN tray icon which indicates whether the VPN connection is active or not.
<b>Create desktop icon</b>	Creates an icon to invoke the SSL client on the user's desktop.
<b>Enable silent mode operation</b>	Hides the pop-up that indicates that installation is complete.
<b>Server security certificate validation</b>	The SSL VPN client validates the SSL VPN server certificate before establishing the secure connection.
<b>Block user on certificate validation failure</b>	If the certificate validation fails, then block the SSL VPN user.

11 Click **OK**.

## Add a User

Add a remote user to the local database.

### Procedure

- 1 In the **SSL VPN-Plus** tab, select **Users** from the left panel.
- 2 Click the **Add (+)** icon.
- 3 Type the user ID.
- 4 Type the password.
- 5 Retype the password.
- 6 (Optional) Type the first and last name of the user.
- 7 (Optional) Type a description for the user.
- 8 In Password Details, select **Password never expires** to always keep the same password for the user.
- 9 Select **Allow change password** to let the user change the password.
- 10 Select **Change password on next login** if you want the user to change the password the next time he logs in.
- 11 Set the user status.
- 12 Click **OK**.

## Enable the SSL VPN-Plus Service

After configuring the SSL VPN-Plus service, enable the service for remote users to begin accessing private networks.

### Procedure

- 1 In the **SSL VPN-Plus** tab, select **Dashboard** from the left panel.
- 2 Click **Start**.

The dashboard displays the status of the service, number of active SSL VPN sessions, and session statistics and data flow details. Click **Details** next to Number of Active Sessions to view information about the concurrent connections to private networks behind the NSX Edge gateway.

### What to do next

- 1 Add an SNAT rule to translate the IP address of the NSX Edge appliance to the VPN Edge IP address.
- 2 Using a web browser, navigate to the IP address of the NSX Edge interface by typing **https://NSXEdgeIPAddress**.
- 3 Login using the user name and password that you created in the [Add a User](#) section and download the installation package.
- 4 Enable port forwarding on your router for the port number used in [Add SSL VPN-Plus Server Settings](#).
- 5 Launch the VPN client, select your VPN server, and login. You can now navigate to the services on your network. SSL VPN-Plus gateway logs are sent to the syslog server configured on the NSX Edge appliance. SSL VPN-Plus client logs are stored in the following directory on the remote user's computer: %PROGRAMFILES%/VMWARE/SSLVPN Client/.

## Add a Script

You can add multiple login or logoff scripts. For example, you can bind a login script for starting Internet Explorer with gmail.com. When the remote user logs in to the SSL client, Internet Explorer opens up gmail.com.

### Procedure

- 1 In the **SSL Vpn-Plus** tab, select **Login/Logoff Scripts** from the left panel.
- 2 Click the **Add (+)** icon.
- 3 In **Script**, click **Browse** and select the script you want to bind to the NSX Edge gateway.

#### 4 Select the **Type** of script.

Option	Description
Login	Performs the script action when remote user logs in to SSL VPN.
Logoff	Performs the script action when remote user logs out of SSL VPN.
Both	Performs the script action both when remote user logs in and logs out of SSL VPN.

#### 5 Type a description for the script.

#### 6 Select **Enabled** to enable the script.

#### 7 Click **OK**.

## Install SSL VPN-Plus Client

Use the SSL VPN Full Access (PHAT) client to connect to a configured private network as a remote user. The client is supported on Windows, Mac, and Linux desktops.

The following topics explain the steps to install the SSL VPN-Plus Client on various operating systems.

- [Install SSL VPN-Plus Client on a Remote Windows Site](#)

Use the steps in this topic to install the SSL VPN-Plus client on a remote Windows site.

- [Install SSL VPN-Plus Client on a Remote Linux Site](#)

Use the steps in this topic to install the SSL VPN-Plus client on a remote Linux site.

- [Install SSL VPN-Plus Client on a Remote Mac Site](#)

Use the steps in this topic to install the SSL VPN-Plus client on a remote Mac computer.

### Install SSL VPN-Plus Client on a Remote Windows Site

Use the steps in this topic to install the SSL VPN-Plus client on a remote Windows site.

#### Procedure

- 1 On the remote client site, open a browser window, and type **https://*ExternalEdgeInterfaceIP*/sslvpn-plus/**, where *ExternalEdgeInterfaceIP* is the IP address of the Edge external interface where you enabled the SSL VPN-Plus service.
- 2 Log in to the portal using the credentials of the remote user.
- 3 Click the **Full Access** tab.
- 4 Click the name of the installer package from the list.
- 5 Click the **click here** link to download the installer package.  
The SSL client is downloaded.
- 6 Extract the downloaded files and run the `Installer.exe` file to install the client.

### What to do next

Log in to the SSL client with the credentials specified in the Users section. The SSL VPN-Plus client validates the SSL VPN server certificate.

Windows client is authenticated as the **Server security certificate validation** option is selected by default, when the installation package was created.

For Internet Explorer (IE) browser, add a trusted CA to the trust certificate store. If server certificate validation fails, you are prompted to contact your system administrator. If server certificate validation succeeds, a login prompt is displayed.

Adding a trusted CA to the trust store is independent of SSL VPN work flow.

## Install SSL VPN-Plus Client on a Remote Linux Site

Use the steps in this topic to install the SSL VPN-Plus client on a remote Linux site.

### Prerequisites

Install Linux TCL and TK packages on the remote computer.

You must have root privileges to install the SSL VPN-Plus client.

### Procedure

- 1 On the remote client site, open a browser window, and type **https://*ExternalEdgeInterfaceIP*/sslvpn-plus/**, where *ExternalEdgeInterfaceIP* is the IP address of the Edge external interface where you enabled the SSL VPN-Plus service.
- 2 Log in to the portal using the credentials of the remote user.
- 3 Click the **Full Access** tab.
- 4 Click the name of the installer package, and save the `linux_phat_client.tgz` compressed file on the remote computer.
- 5 Extract the compressed file. The `linux_phat_client` directory is created.
- 6 Open the Linux CLI and change to the `linux_phat_client` directory.
- 7 Run the `./install_linux_phat_client.sh` command.

### What to do next

Log in to the SSL VPN GUI with the credentials specified in the Users section.

---

### Attention

- Two-factor RSA authentication is not supported for logging in to the SSL VPN client on Linux operating systems.
  - SSL VPN Linux client CLI does not validate server certificates. If server certificate validation is required, use the SSL VPN GUI for connecting to the gateway.
-

The SSL VPN Linux client validates the server certificate against the browser's certificate store by default. If server certificate validation fails, you are prompted to contact your system administrator. If server certificate validation succeeds, a login prompt is displayed.

Adding a trusted CA to the trust store (for example, Firefox certificate store) is independent of the SSL VPN work flow.

## Install SSL VPN-Plus Client on a Remote Mac Site

Use the steps in this topic to install the SSL VPN-Plus client on a remote Mac computer.

### Prerequisites

You must have root privileges to install the SSL VPN-Plus client.

### Procedure

- 1 On the remote client site, open a browser window, and type **https://**  
***ExternalEdgeInterfaceIP/sslvpn-plus/***, where *ExternalEdgeInterfaceIP* is the IP address of the Edge external interface where you enabled the SSL VPN-Plus service.
- 2 Log in to the portal using the credentials of the remote user.
- 3 Click the **Full Access** tab.
- 4 Click the name of the installer package, and save the `mac_phat_client.tgz` compressed file on the remote computer.
- 5 Extract the compressed file. The `mac_phat_client` directory is created.
- 6 To install the SSL VPN-Plus client, double-click the `naclient.pkg` file .

Follow the steps in the wizard to finish the installation.

If your SSL VPN Client installation fails, check the installation log file at `/tmp/naclient_install.log`.

For troubleshooting installation problems on Mac OS High Sierra, see the *NSX Troubleshooting Guide*.

### What to do next

Log in to the SSL client with the credentials specified in the Users section.

---

**Attention** Two-factor authentication is not supported for logging in to the SSL VPN client on Mac operating systems.

---

The SSL VPN Mac client validates the server certificate against Keychain, a database that stores certificates on Mac OS, by default. If server certificate validation fails, you are prompted to contact your system administrator. If server certificate validation succeeds, a login prompt is displayed.

## Configure Proxy Server Settings in SSL VPN-Plus Client

Starting in NSX 6.4.6, proxy server configuration is supported on SSL VPN-Plus Client on Windows, Mac, and Linux computers. In NSX 6.4.5 and earlier, proxy server configuration is supported only on SSL VPN-Plus Client on a Windows computer.

---

**Caution** In NSX 6.4.5 and earlier:

- SSL VPN-Plus Client on Mac OS provides the facility to configure proxy server settings, but remote users must not configure the proxy server settings.
  - Remote Linux OS users must avoid configuring the proxy server settings on the SSL VPN-Plus Client through the Linux CLI.
- 

The following procedure explains the steps to configure the proxy server settings in an SSL VPN-Plus Client.

### Prerequisites

SSL VPN-Plus Client is installed on the remote computer.

### Procedure

- 1 Double-click the desktop icon of the SSL VPN-Plus Client on the remote computer.  
The **SSL VPN-Plus Client - Login** window opens.
- 2 Navigate to proxy settings.
  - On Windows and Mac computer, click **Settings**, and then click the **Proxy Settings** tab.
  - On a Linux computer, click **Proxy Settings**.

- 3 Specify the proxy server settings.
  - a Select the **Use Proxy** check box.
  - b Under **Type Of Proxy**, configure one of the proxy server types.

Option	Description
Use IE Settings	This option is available only in Windows SSL VPN-Plus Client. Use the proxy server configuration that is specified in your IE browser.
HTTP	Specify the following settings for an HTTP proxy server: <ul style="list-style-type: none"> <li>■ Proxy server name or an IP address of the proxy server.</li> <li>■ Proxy server port. The default port is 80, which you can edit.</li> </ul>
SOCKS ver 4	This option is available only in Windows SSL VPN-Plus Client. Specify the following settings for a SOCKS 4.0 proxy server: <ul style="list-style-type: none"> <li>■ Proxy server name or an IP address of the proxy server.</li> <li>■ Proxy server port. The default port is 1080, which you can edit.</li> </ul>
SOCKS ver 5	Specify the following settings for a SOCKS 5.0 proxy server: <ul style="list-style-type: none"> <li>■ Proxy server name or an IP address of the proxy server.</li> <li>■ Proxy server port. The default port is 1080, which you can edit.</li> <li>■ (Optional) User name and password to access the SOCKS 5.0 server.</li> </ul>

- 4 To save the proxy server settings, click **OK** .

## SSL VPN-Plus Logs

SSL VPN-Plus gateway logs are sent to the syslog server configured on the NSX Edge appliance.

The following table lists the locations on the remote user's computer where the SSL VPN-Plus client logs are stored.

Operating System	Location of Log File
Windows 8	C:\Users\username\AppData\Local\VMware\vpn\svp_client.log
Windows 10	C:\Users\username\AppData\Local\VMware\vpn\svp_client.log
Linux	System log files
Mac	Installation log file at /tmp/naclient_install.log System log files

## Change SSL VPN-Plus Client Logs and Log Level

- 1 In the **SSL VPN-Plus** tab, click **Server Settings** from the left panel.
- 2 Go to the Logging Policy section and expand the section to view the current settings.
- 3 Click **Change**.
- 4 Select **Enable logging** check box to enable logging.

OR

Deselect **Enable logging** check box to disable logging.

- 5 Select the required log level.

---

**Note** SSL VPN-Plus client logs are enabled by default and log level is set to NOTICE.

---

- 6 Click **OK**.

## Edit Client Configuration

You can change the way the SSL VPN client tunnel responds when the remote user logs in to SSL VPN.

### Procedure

- 1 In the **SSL VPN-Plus** tab, select **Client Configuration** from the left panel.
- 2 Select the **Tunneling Mode**.
 

In split tunnel mode, only the VPN flows through the NSX Edge gateway. In full tunnel, the NSX Edge gateway becomes the remote user's default gateway and all traffic (VPN, local, and internet) flows through this gateway.
- 3 If you selected the full tunnel mode:
  - a Select **Exclude local subnets** to exclude local traffic from flowing through the VPN tunnel.
  - b Type the IP address for the default gateway of the remote user's system.
- 4 Select **Enable auto reconnect** if you would like the remote user to automatically reconnect to the SSL VPN client after getting disconnected.
- 5 Select **Client upgrade notification** for the remote user to get a notification when an upgrade for the client is available. The remote user can then choose to install the upgrade.
- 6 Click **OK**.

## Edit General Settings

You can edit the default VPN settings.

### Procedure

- 1 In the **SSL VPN-Plus** tab, select **General Settings** from the left panel.
- 2 Make required selections.

Select	To
Prevent multiple logon using same username	Allow a remote user to login only once with a username.
Enable compression	Enable TCP based intelligent data compression and improve data transfer speed.
Enable logging	Maintain a log of the traffic passing through the SSL VPN gateway.

Select	To
<b>Force virtual keyboard</b>	Allow remote users to enter web or client login information only via the virtual keyboard.
<b>Randomize keys of virtual keyboard</b>	Make the virtual keyboard keys random.
<b>Enable forced timeout</b>	Disconnect the remote user after the specified timeout period is over. Type the timeout period in minutes.
<b>Session idle timeout</b>	If there is no activity on the user session for the specified period, end the user session after that period is over.  SSLVPN idle timeout considers all packets, including control packets sent by any Application and user data, towards timeout detection. As a result, even if there is no user data, the session won't timeout if there is an application which transmits a periodic control packet (such as MDNS).
<b>User notification</b>	Type a message to be displayed to the remote user after he logs in.

- 3 Click **OK**.

## Edit Web Portal Design

You can edit the client banner bound to the SSL VPN client.

### Procedure

- 1 In the **NSX Edges** tab, double-click an NSX Edge.
- 2 Click the **Manage** tab and then click the **SSL VPN-Plus** tab.
- 3 Select **Portal Customization** from the left panel.
- 4 Type the portal title.
- 5 Type the remote user's company name.
- 6 In **Logo**, click **Change** and preferably select a JPEG image for the company logo.  
  
There are no preferred dimensions for the logo size.
- 7 In **Colors**, click the color box next to numbered item for which you want to change the color, and select the desired color.
- 8 Change the client banner, if necessary. Select a BMP image for the banner.  
  
Preferred size for the client banner is 390 X 75 pixels.
- 9 Click **OK**.

## Working with IP Pools for SSL VPN

You can edit or delete an IP pool.

For information on adding an IP pool, see [Configure Network Access SSL VPN-Plus](#).

### Edit an IP Pool

You can edit an IP pool.

### Procedure

- 1 In the **SSL VPN-Plus** tab, click **IP Pool** in the left panel.
- 2 Select the IP pool that you want to edit.
- 3 Click the **Edit** () icon.  
The Edit IP Pool dialog box opens.
- 4 Make the required edits.
- 5 Click **OK**.

### Delete an IP Pool

You can delete an IP pool.

#### Procedure

- 1 In the **SSL VPN-Plus** tab, click **IP Pool** in the left panel.
- 2 Select the IP pool that you want to delete.
- 3 Click the **Delete** () icon.  
The selected IP pool is deleted.

### Enable an IP Pool

You can enable an IP pool if you want an IP address from that pool to be assigned to the remote user.

#### Procedure

- 1 In the **SSL VPN-Plus** tab, click **IP Pool** in the left panel.
- 2 Select the IP pool that you want to enable.
- 3 Click the **Enable** () icon.

### Disable an IP Pool

You can disable an IP pool if you do not want the remote user to be assigned an IP address from that pool.

#### Procedure

- 1 In the **SSL VPN-Plus** tab, select **IP Pool** from the left panel.
- 2 Select the IP pool that you want to disable.
- 3 Click the **Disable** () icon.

## Change the Order of an IP Pool

SSL VPN assigns an IP address to a remote user from an IP pool based on its order in the IP pool table.

### Procedure

- 1 In the **SSL VPN-Plus** tab, click **IP Pool** in the left panel.
- 2 Select the IP pool that you want to change the order for.
- 3 Click the **Move Up** () or **Move Down** () icon.

## Working with Private Networks

You can edit or delete a private network that a remote user can access.

For information on adding a private network, see [Configure Network Access SSL VPN-Plus](#).

### Delete a Private Network

You can delete a private network

#### Procedure

- 1 In the **SSL VPN-Plus** tab, click **Private Networks** in the left panel.
- 2 Select the network that you want to delete and click the **Delete** () icon.

### Enable a Private Network

When you enable a private network, the remote user can access it through SSL VPN-Plus.

#### Procedure

- 1 In the **SSL VPN-Plus** tab, click **Private Networks** in the left panel.
- 2 Click the network that you want to enable.
- 3 Click the **Enable** icon ()

The selected network is enabled.

### Disable a Private Network

When you disable a private network, the remote user cannot access it through SSL VPN-Plus.

#### Procedure

- 1 In the **SSL VPN-Plus** tab, click **Private Networks** in the left panel.
- 2 Click the network that you want to disable.
- 3 Click the **Disable** () icon.

The selected network is disabled.

## Change the Sequence of a Private Network

SSL VPN-Plus allows remote users to access private networks in the sequence in which they are displayed on the Private Networks panel.

If you select **Enable TCP Optimization** for a private network, some applications such as FTP in Active mode may not work within that subnet. To add an FTP server configured in Active mode, you must add another private network for that FTP server with TCP Optimization disabled. Also, the active TCP private network must be enabled, and must be placed above the subnet private network.

### Procedure

- 1 In the **SSL VPN-Plus** tab, click **Private Networks** in the left panel.
- 2 Click the **Change Order** (≡↓) icon.
- 3 Select the network that you want to change the order of.
- 4 Click the **Move Up** (≡↑) or **Move Down** (≡↓) icon.
- 5 Click **OK**.

### What to do next

To add an FTP server configured in Active mode, refer to [Configure Private Network for Active FTP Server](#).

## Configure Private Network for Active FTP Server

You can add an FTP server configured in Active mode to the private network. For active FTP, the control connection is initiated by the back end FTP server to the client machine which does not support TCP optimization.

### Prerequisites

FTP server is configured in Active mode.

### Procedure

- 1 In the **SSL VPN-Plus** tab, click **Private Networks** in the left panel.
- 2 Add the private network that you want to configure for active FTP. For more information, refer to [Add a Private Network](#).
- 3 Deselect the **Enable TCP Optimization** check box.
- 4 In the **Ports** field, add port number for the private network.
- 5 Select status as **Enabled** to enable the private network.
- 6 Place the private network that you want to configure for active FTP above other private networks that are configured. For more information, refer to [Change the Sequence of a Private Network](#).

### What to do next

Add a corresponding firewall rule to allow the private network traffic.

## Working with Installation Packages

You can delete or edit an installation package for the SSL client.

For information on creating an installation package, see [Configure Network Access SSL VPN-Plus](#).

### Edit an Installation Package

You can edit an installation package.

#### Procedure

- 1 In the **SSL VPN-Plus** tab, click **Installation Package** in the left panel.
- 2 Select the installation package that you want to edit.
- 3 Click the Edit (✎) icon.  
The Edit Installation Package dialog box opens.
- 4 Make the required edits.
- 5 Click **OK**.

### Delete an Installation Package

You can delete an installation package.

#### Procedure

- 1 In the **SSL VPN-Plus** tab, click **Installation Package** in the left panel.
- 2 Select the installation package that you want to delete.
- 3 Click the **Delete** (✖) icon.

## Working with Users

You can edit or delete users from the local database.

For information on adding a user, see [Configure Network Access SSL VPN-Plus](#).

### Edit a User

You can edit the details for a user except for the user ID.

#### Procedure

- 1 In the **SSL VPN-Plus** tab, click **Users** in the left panel.
- 2 Click the **Edit** (✎) icon.
- 3 Make the required edits.

- 4 Click **OK**.

## Delete a User

You can delete a user.

### Procedure

- 1 In the **SSL VPN-Plus** tab, click **Users** in the left panel.
- 2 **Users**In the **Configure** panel, click **Users**.
- 3 Select the user that you want to delete and click the **Delete** () icon.

## Change the Password for a User

You can change the password for a user.

### Procedure

- 1 In the **SSL VPN-Plus** tab, click **Users** in the left panel.
- 2 Click the **Change Password** icon.
- 3 Type and re-type the new password.
- 4 Click **Change password on next login** to change the password when the user logs in to his system next time.
- 5 Click **OK**.

## Working with Login and Logoff Scripts

You can bind a login or logoff script to the NSX Edge gateway.

### Edit a Script

You can edit the type, description, and status of a login or logoff script that is bound to the NSX Edge gateway.

### Procedure

- 1 In the **SSL VPN-Plus** tab, click **Login/Logoff Scripts** in the left panel.
- 2 Select a script and click the **Edit** () icon.
- 3 Make the appropriate changes.
- 4 Click **OK**.

### Delete a Script

You can delete a login or logoff script.

### Procedure

- 1 In the **SSL VPN-Plus** tab, click **Login/Logoff Scripts** in the left panel.
- 2 Select a script and click the **Delete** (✖) icon.

## Enable a Script

You must enable a script for it to work.

### Procedure

- 1 In the **SSL VPN-Plus** tab, click **Login/Logoff Scripts** in the left panel.
- 2 Select a script and click the **Enable** (✔) icon.

## Disable a Script

You can disable a login/logoff script.

### Procedure

- 1 In the **SSL VPN-Plus** tab, click **Login/Logoff Scripts** in the left panel.
- 2 Select a script and click the **Disable** (⊘) icon.

## Change the Order of a Script

You can change the order of a script. For example, suppose you have a login script for opening gmail.com in Internet Explorer placed above a login script for opening yahoo.com. When the remote user logs in to SSL VPN, gmail.com is displayed before yahoo.com. If you now reverse the order of the login scripts, yahoo.com is displayed before gmail.com.

### Procedure

- 1 In the **SSL VPN-Plus** tab, click **Login/Logoff Scripts** in the left panel.
- 2 Select the script that you want to change the order of and click the **Move Up** (≡↑) or **Move Down** (≡↓) icon.
- 3 Click **OK**.

## IPSec VPN Overview

NSX Edge supports site-to-site IPSec VPN between an NSX Edge instance and remote sites. Certificate authentication, preshared key mode, and IP unicast traffic are supported between the NSX Edge instance and remote VPN sites.

Starting with NSX Data Center 6.4.2, you can configure both policy-based IPsec VPN service and route-based IPsec VPN service. However, you can configure, manage, and edit route-based IPsec VPN parameters only by using REST APIs. You cannot configure or edit route-based IPsec VPN parameters in the vSphere Web Client. For more information about using APIs to configure route-based IPsec VPN, see the *NSX API Guide*.

In NSX 6.4.1 and earlier, you can configure only policy-based IPsec VPN service.

## Policy-Based IPsec VPN

In a policy-based IPsec VPN, you explicitly configure the subnets behind the NSX Edge on the local site that require secure and encrypted communication with the remote subnets on the peer site.

When the local IPsec VPN site originates traffic from unprotected local subnets to the protected remote subnets on the peer site, the traffic is dropped.

The local subnets behind an NSX Edge must have address ranges that do not overlap with the IP addresses on the peer VPN site. If the local and remote peer across an IPsec VPN tunnel has overlapping IP addresses, traffic forwarding across the tunnel might not be consistent.

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

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 tunnel. On both ends, static one-to-one NAT is required for the VPN address.

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

**Table 15-6. Number of Supported IPsec Tunnels**

ESG Size	Number of IPsec Tunnels
Compact	512
Large	1600
Quad-Large	4096
X-Large	6000

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

For a detailed example of configuring a policy-based IPsec tunnel between an NSX Edge and a remote VPN Gateway, see [Configure Policy-Based IPsec VPN Site Example](#).

## Route-Based IPSec VPN

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, virtual tunnel interfaces (VTI) are created on the ESG appliance. 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 interfaces.

### VPN Tunnel Redundancy

With route-based IPSec VPN service, you can configure VPN tunnel redundancy. Tunnel redundancy provides uninterrupted data path connectivity between the two sites when the ISP link fails, or when the remote VPN Gateway fails.

---

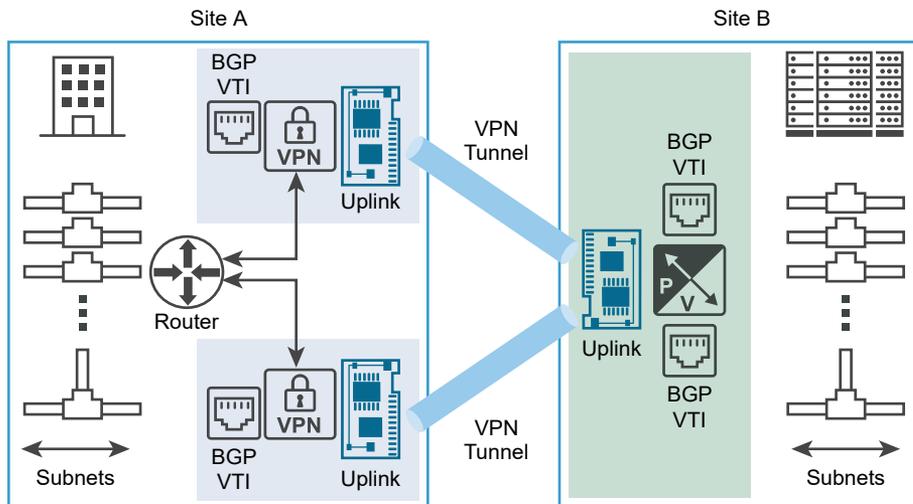
#### Important

- In NSX Data Center 6.4.2 and later, IPSec VPN tunnel redundancy is supported only using BGP. OSPF dynamic routing is not supported for routing through IPSec VPN tunnels.
- 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 Site A has Edge VPN Gateways that might not be managed by NSX, and Site B has an Edge Gateway virtual appliance that is managed by NSX.

Figure 15-1. 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. Both IPsec VPN tunnels remain in service if they are available. All the traffic destined from Site A to Site B through the ESG is routed through the VTI. The data traffic undergoes IPsec processing and goes out of its associated ESG uplink interface. All the incoming IPsec traffic received from Site B VPN Gateway on the ESG 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.

Some key points that you must remember about route-based IPsec VPN service are as follows:

- You can configure policy-based IPsec VPN tunnels and route-based IPsec tunnels on the same ESG appliance. However, you cannot configure a policy-based tunnel and a route-based tunnel with the same VPN peer site.
- NSX supports a maximum of 32 VTIs on a single ESG appliance. That is, you can configure a maximum of 32 route-based VPN peer sites.
- NSX does not support migration of existing policy-based IPsec VPN tunnels to route-based tunnels or conversely.

For information about configuring a route-based IPsec VPN site, see [Configure Route-Based IPsec VPN Site](#).

For a detailed example of configuring a route-based IPsec VPN tunnel between a local NSX Edge and a remote Cisco CSR 1000V VPN Gateway, see [Using a Cisco CSR 1000V Appliance](#).

## Configure Policy-Based IPsec VPN Site

You can set up policy-based IPsec VPN tunnels between local subnets and peer subnets.

---

**Note** If you connect to a remote site by using an IPsec VPN tunnel, dynamic routing on the edge uplink cannot learn the IP address of that site.

---

The task topics in this section explain the steps to configure a policy-based IPsec VPN site.

### Enable IPsec VPN Service

You must enable the IPsec VPN service for traffic to flow from the local subnet to the peer subnet.

#### Prerequisites

You must configure at least one IPsec VPN site on the NSX Edge before enabling the IPsec VPN service.

#### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > VPN > IPsec VPN**.
- 5 Next to **IPsec VPN Service Status**, click **Start**.

### Use OpenSSL to Generate CA-Signed Certificates for IPsec VPNs

To enable certificate authentication for IPsec, server certificates and corresponding CA-signed certificates must be imported. Optionally, you can use an open-source command-line tool such as OpenSSL to generate CA-signed certificates.

#### Prerequisites

OpenSSL must be installed.

#### Procedure

- 1 On a Linux or Mac machine where OpenSSL is installed, open the file: `/opt/local/etc/openssl/openssl.cnf` or `/System/Library/OpenSSL/openssl.cnf`.
- 2 Ensure that `dir = ..`
- 3 Run the following commands:

```
mkdir newcerts
mkdir certs
mkdir req
mkdir private
echo "01" > serial
touch index.txt
```

- 4 Run the command to generate a CA-signed certificate:

```
openssl req -new -x509 -newkey rsa:2048 -keyout private/cakey.pem -out cacert.pem -days 3650
```

- 5 On NSX Edge1, do these steps:

- a Generate a certificate signing request (CSR).

For detailed steps, see [Configure a CA Signed Certificate](#)

- b Copy the privacy-enhanced mail (PEM) file content, and save it in a file in `req/edge1.req`.

- 6 Run the following command to sign the CSR:

```
sudo openssl ca -policy policy_anything -out certs/edge1.pem -in req/edge1.req
```

- 7 On NSX Edge2, generate a CSR, copy the PEM file content, and save it in a file in `req/edge2.req`.

- 8 Run the following command to sign the CSR:

```
sudo openssl ca -policy policy_anything -out certs/edge2.pem -in req/edge2.req
```

- 9 Upload the PEM certificate at the end of the file `certs/edge1.pem` to Edge1.
- 10 Upload the PEM certificate at the end of the file `certs/edge2.pem` to Edge2.
- 11 Import the signed certificate (`cacert.pem`) to Edge1 and Edge2 as CA-signed certificates.
- 12 In the IPsec global configuration for Edge1 and Edge2, select the uploaded PEM certificate and the CA certificate and save the configuration.
- 13 Navigate to **Manage > Settings > Certificates**. Select the signed certificate that you imported and record the DN string.
- 14 Reverse the DN string to the format `C=IN,ST=ka,L=blr,O=bmware,OU=vmware,CN=edge2.eng.vmware.com` and save it for Edge1 and Edge2.
- 15 Create IPsec VPN sites on Edge1 and Edge2 with Local ID and Peer ID as the distinguished name (DN) string in the specified format.

## Results

Check the status by clicking **Show Statistics** or **Show IPsec Statistics**. Click the channel to see the tunnel status. The channel status should be enabled and the tunnel status should be Up.

## Specify Global IPsec VPN Configuration

Use the steps in this topic to enable IPsec VPN on the NSX Edge instance.

## Prerequisites

To enable certificate authentication, server certificates and corresponding CA-signed certificates must be imported. Optionally, you can use an open-source command-line tool such as OpenSSL to generate CA-signed certificates.

Self-signed certificates cannot be used for IPSec VPN. They can only be used in load balancing and SSL VPN.

## Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > VPN > IPSec VPN**.
- 5 Next to **Global Configuration**, click **Edit** or **Change**.
- 6 Enter a global pre-shared key for those sites whose peer endpoint is set to "any".

To view the pre-shared key, click the **Show Pre-Shared Key** (🔑) icon or select the **Display shared key** check box.

## 7 Configure the global extensions.

The following table describes the global extensions.

Extension	Description
<code>add_spd</code>	<p>Allowed values are <code>on</code> and <code>off</code>. The default value is <code>on</code>, even when you do not configure this extension.</p> <p>When <code>add_spd=off</code>:</p> <ul style="list-style-type: none"> <li>■ Security policies are installed only when the tunnel is up.</li> <li>■ If the tunnel is up, packets are sent encrypted through the tunnel.</li> <li>■ If the tunnel is down, packets are sent unencrypted, if a route is available.</li> </ul> <p>When <code>add_spd=on</code>:</p> <ul style="list-style-type: none"> <li>■ Security policies are installed regardless of whether the tunnel is established.</li> <li>■ If the tunnel is up, packets are sent encrypted through the tunnel.</li> <li>■ If the tunnel is down, packets are dropped.</li> </ul>
<code>ike_fragment_size</code>	<p>If the maximum transmission unit (MTU) is small, you can set the IKE fragment size by using this extension to avoid failures in the IKE negotiation. For example, <code>ike_fragment_size=900</code></p>
<code>ignore_df</code>	<p>Allowed values are <code>on</code> and <code>off</code>. Default value is <code>off</code>.</p> <ul style="list-style-type: none"> <li>■ When <code>ignore_df=off</code>, NSX Edge copies the value of the "don't fragment (DF)" bit from the clear text packet to the encrypted packet. This implies that if the clear text packet has the DF bit set, after encryption, the packet also has the DF bit set.</li> <li>■ When <code>ignore_df=on</code>, NSX Edge ignores the value of the DF bit in the clear text packet, and the DF bit is always 0 in the encrypted packet.</li> <li>■ Set this flag to <code>on</code> when the DF bit is set in the clear text packet and the size of the packet after encryption exceeds the MTU of the TCP packet. If the DF bit is set, the packet is dropped, but if the bit is cleared, the packet gets fragmented.</li> </ul>

8 Enable certificate authentication, and then select the appropriate Service certificate, CA certificate, and the certificate revocation list (CRL).

9 Click **Save** or **OK**, and then click **Publish Changes**.

## Enable Logging for IPSec VPN

You can enable logging of all IPSec VPN traffic.

By default, logging is enabled and is set to the WARNING level.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > VPN > IPSec VPN**.

- 5 Enable logging to log traffic flow between the local subnet and peer subnet.

NSX Version	Procedure
6.4.6 and later	<ol style="list-style-type: none"> <li>a Next to <b>Logging Configuration</b>, click <b>Edit</b>.</li> <li>b Click the toggle switch to enable logging, and then select the logging level.</li> <li>c Click <b>Save</b>.</li> </ol>
6.4.5 and earlier	<ol style="list-style-type: none"> <li>a Next to <b>Logging Policy</b>, click .</li> <li>b Select the <b>Enable logging</b> check box, and then select the logging level.</li> </ol>

- 6 Click **Publish Changes**.

## Configure IPSec VPN Parameters

You must configure at least one external IP address on the NSX Edge to provide IPSec VPN service.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > VPN > IPSec VPN**.
- 5 Click **Add**.
- 6 Enter a name for the IPSec VPN site.
- 7 Configure the endpoint parameters of the IPSec VPN site.
  - a Enter the local Id to identify the local NSX Edge instance. This local Id is the peer Id on the remote site.  
  
The local Id can be any string. Preferably, use the public IP address of the VPN or a fully qualified domain name (FQDN) for the VPN service as the local Id.
  - b Enter an IP address or an FQDN of the local endpoint.  
  
If you are adding an IP-to-IP tunnel using a pre-shared key, the local Id and local endpoint IP can be the same.
  - c Enter the subnets to share between the IPSec VPN sites in the CIDR format. Use a comma separator to enter multiple subnets.

- d Enter the Peer Id to identify the peer site.
- For peers using certificate authentication, this ID must be the distinguished name (DN) in the peer's certificate. Enter the DN of the certificate as a string of comma-separated values in the following order without spaces:  
`C=xxx, ST=xxx, L=xxx, O=xxx, OU=xxx, CN=xxx, E=xxx.`
  - For PSK peers, the peer Id can be any string. Preferably, use the public IP address of the VPN or an FQDN for the VPN service as the peer Id.

---

**Note** If the Edge has more than one uplink interface that can reach the remote IPSec peer, routing should be done in such a way that IPSec traffic goes out of the Edge interface, which is configured with a local peer IP.

---

- e Enter an IP address or an FQDN of the peer endpoint. The default value is **any**. If you retain the default value, you must configure the Global PSK.
- f Enter the internal IP address of the peer subnet in the CIDR format. Use a comma separator to type multiple subnets.

## 8 Configure the tunnel parameters.

- a (Optional) Select a security compliance suite to configure the security profile of the IPsec VPN site with predefined values defined by that suite.

The default selection is **none**, which means that you must manually specify the configuration values for authentication method, IKE profile, and tunnel profile. When you select a compliance suite, values that are predefined in that standard compliance suite are automatically assigned, and you cannot edit these values. For more information about compliance suites, see [Supported Compliance Suites](#).

### Note

- Compliance suite is supported in NSX Data Center 6.4.5 or later.
  - If FIPS mode is enabled on the Edge, you cannot specify a compliance suite.
- b Select one of the following Internet Key Exchange (IKE) protocols to set up a security association (SA) in the IPsec protocol suite.

Option	Description
IKEv1	When you select this option, IPsec VPN initiates and responds to IKEv1 protocol only.
IKEv2	When you select this option, IPsec VPN initiates and responds to IKEv2 protocol only.
IKE-Flex	When you select this option, and if the tunnel establishment fails with 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.

**Important** If you configure multiple sites with the same local and remote endpoints, make sure that you select the same IKE version and PSK across all these IPsec VPN sites.

- c From the **Digest Algorithm** drop-down menu, select one of the following secure hashing algorithms:
- SHA1
  - SHA\_256

- d From the **Encryption Algorithm** drop-down menu, select one of the following supported encryption algorithms:

- AES (AES128-CBC)
- AES256 (AES256-CBC)
- Triple DES (3DES192-CBC).
- AES-GCM (AES128-GCM)

---

#### Note

- AES-GCM encryption algorithm is not FIPS-compliant.
  - Starting in NSX 6.4.5, Triple DES cypher algorithm is deprecated in IPSec VPN service.
- 

The following table explains the encryption settings that are used on the peer VPN Gateway for the encryption settings that you select on the local NSX Edge.

**Table 15-7. Encryption Settings**

Encryption Settings on NSX Edge	IKE Settings on Peer VPN Gateway	IPSec Settings on Peer VPN Gateway
AES-256	AES-256	AES-256
AES-128	AES-128	AES-128
3DES	3DES	3DES
AES-GCM, IKEv1	AES-128	AES-GCM
AES-GCM, IKEv2	AES-128 or AES-GCM	AES-GCM

- e In Authentication Method, select one of the following options:

Option	Description
<b>PSK (Pre Shared Key)</b>	Indicates that the secret key shared between NSX Edge and the peer site is to be used for authentication. The secret key can be a string with a maximum length of 128 bytes.  PSK authentication is disabled in FIPS mode.
<b>Certificate</b>	Indicates that the certificate defined at the global level is to be used for authentication.

- f (Optional) Enter the pre-shared key of the peer IPSec VPN site.

- g To display the key on the peer site, click the **Show Pre-Shared Key** (🔑) icon or select the **Display Shared Key** check box.
- h From the **Diffie-Hellman (DH) Group** drop-down menu, select one of the following cryptography schemes that allows the peer site and the NSX Edge to establish a shared secret over an insecure communications channel.

- DH-2
- DH-5
- DH-14
- DH-15
- DH-16

DH14 is default selection for both FIPS and non-FIPS mode. DH2 and DH5 are not available when the FIPS mode is enabled.

- 9 Configure the advanced parameters.
  - a If the remote IPsec VPN site does not support PFS, disable the **Perfect forward secrecy (PFS)** option. By default, PFS is enabled.
  - b (Optional) To operate IPsec VPN in a responder-only mode, select the **Responder only** check box.  
  
In this mode, IPsec VPN never initiates a connection.
  - c (Optional) In the **Extension** text box, type one of the following:
    - `securelocaltrafficbyip=IPAddress` to redirect Edge local traffic over the IPsec VPN tunnel. IP address is the default value. For more information, see <http://kb.vmware.com/kb/20080007> .
    - `passthroughSubnets=PeerSubnet/IPAddress` to support overlapping subnets.

- 10 Click **Add** or **OK**, and then click **Publish Changes**.

The IPsec VPN configuration is saved on the NSX Edge.

## What to do next

Enable the IPsec VPN service.

**Tip** In the vSphere Web Client, you can follow these steps on the **IPsec VPN** page to generate the configuration script for the Peer VPN Gateway.

- In NSX 6.4.6 and later, select the IPsec VPN site, and then click **Actions > Generate Peer Configuration**.
- In NSX 6.4.5 and earlier, select the IPsec VPN site, and then click the **Generate Peer Configure** icon. In the dialog box that opens, click **Generate Peer Configure**.

The configuration script is generated. You can use this script as reference to configure the IPsec VPN parameters on the peer VPN Gateway.

## Supported Compliance Suites

Starting with NSX Data Center 6.4.5, you can specify a compliance suite to configure the various parameters in the security profile of an IPsec VPN site.

A security compliance suite has a predefined set of values for various security parameters. Think of a compliance suite as a predefined template to help you automatically configure the security profile of an IPsec VPN session according to a defined standard. For example, the National Security Agency in the US government publishes the CNSA suite, and this standard is used for national security applications. When you select a compliance suite, the security profile of an IPsec VPN site is automatically configured with predefined values, and you cannot edit these values. By specifying a compliance suite, you avoid the need of individually configuring each parameter in the security profile.

NSX supports seven security compliance suites. The following table lists the predefined values for various configuration parameters in each supported compliance suite.

**Table 15-8. Compliance Suites: Predefined Configuration Parameter Values**

Configuration Parameter	Compliance Suite						
	CNSA	Suite-B-GCM-128	Suite-B-GCM-256	Suite-B-GMAC-128	Suite-B-GMAC-256	Prime	Foundation
IKE Version	IKEv2	IKEv2	IKEv2	IKEv2	IKEv2	IKEv2	IKEv1
Digest Algorithm	SHA 384	SHA 256	SHA 384	SHA 256	SHA 384	SHA 256	SHA 256
Encryption Algorithm	AES 256	AES 128	AES 256	AES 128	AES 256	AES GCM 128	AES 128
Tunnel Encryption	AES 256	AES GCM 128	AES GCM 256	AES GMAC 128	AES GMAC 256	AES GCM 128	AES 128

Table 15-8. Compliance Suites: Predefined Configuration Parameter Values (continued)

Configuration Parameter	Compliance Suite							
	SHA 384	NULL	NULL	NULL	NULL	NULL	NULL	SHA 256
Tunnel Digest Algorithm								
Authentication	■ RSA Certificate (3072-bit key)	ECDSA Certificate (P-256 curve)	ECDSA Certificate (P-384 curve)	ECDSA Certificate (P-256 curve)	ECDSA Certificate (P-384 curve)	ECDSA Certificate (P-256 curve)	ECDSA Certificate (P-256 curve)	RSA Certificate (2048-bit key and SHA-256)
	■ ECDSA Certificate (P-384 curve)							
DH Group	DH15 and ECDH20	ECDH19	ECDH20	ECDH19	ECDH20	ECDH19	DH14	

**Caution** Starting in NSX 6.4.6, the "Suite-B-GMAC-128" and "Suite-B-GMAC-256" compliance suites are deprecated. If you configured IPsec VPN sites in NSX 6.4.5 with any of these two deprecated compliance suites, you can still upgrade the edges to 6.4.6. However, a warning message appears to inform you that the IPsec VPN sites are using a vulnerable compliance suite.

**Attention** When you configure an IPsec VPN site using "Prime" and "Foundation" compliance suites, you cannot configure `ikelifetime` and `salifetime` site extensions. These site extensions are pre-configured based on the standard.

When you select the "CNSA" compliance suite, both DH15 and ECDH20 DH groups are internally configured on the NSX Edge. However, the following caveats exist when you select this compliance suite:

- If the IPsec VPN service on an NSX Edge is configured as an initiator, NSX sends only ECDH20 to establish an IKE security association with the remote IPsec VPN site. By default, NSX uses ECDH20 because it is more secure than DH15. If a third-party responder IPsec VPN site is configured with only DH15, the responder sends an invalid IKE payload error message and asks the initiator to use the DH15 group. The initiator reinitiates IKE SA with the DH15 group, and a tunnel gets established between both the IPsec VPN sites. However, if the third-party IPsec VPN solution does not support an invalid IKE payload error, the tunnel is never established between both sites.
- If the IPsec VPN service on an NSX Edge is configured as a responder, the tunnel is always established depending on the DH group that is shared by the initiator IPsec VPN site.
- When both the initiator and responder IPsec VPN sites use an NSX Edge, the tunnel is always established with ECDH20.

## Configure Route-Based IPSec VPN Site

You want to configure a route-based IPSec tunnel between an NSX Edge on the local site and a remote VPN Gateway on the peer site.

Unlike a policy-based IPSec tunnel configuration where you configure local and remote subnets, in a route-based IPSec tunnel configuration, you do not define the local and peer subnets that want to communicate with each other. In a route-based IPSec tunnel configuration, you must define a VTI with a private IP address on both the local and peer sites. Traffic from the local subnets is routed through the VTI to the peer subnets. Use a dynamic routing protocol, such as BGP, to route traffic through the IPSec tunnel. The dynamic routing protocol decides traffic from which local subnet is routed using the IPSec tunnel to the peer subnet.

Following steps explain the procedure to configure a route-based IPSec tunnel between the two sites:

- 1 Configure the IPSec VPN parameters on the local NSX Edge. In NSX Data Center 6.4.2 and later, you can configure route-based IPSec VPN parameters only by using REST APIs. For more information, see the *NSX API Guide*.
  - Local endpoint IP address and local ID to identify the local NSX Edge Gateway.
  - Peer endpoint IP address and peer ID to identify the peer VPN Gateway.
  - IKE version to set up a security association between both the sites.
  - Digest algorithm.
  - Encryption algorithm.
  - Authentication mechanism (either pre-shared key or certificate).
  - Diffie-Hellman (DH) Group public key cryptography scheme.
  - Enable or disable perfect forward secrecy.
  - Enable or disable the Responder-only mode.
  - Virtual tunnel interface (VTI) on the NSX Edge. Provide a static private IP address for the VTI.

---

**Note** The VTI that you configure is a static VTI. Therefore, it cannot have more than one IP address. The best practice is to ensure that the IP address of the VTI on both the local and peer sites are on the same subnet.

---

- 2 Use the IPSec Config Download API to fetch the peer configuration for reference purposes and configure the peer VPN Gateway.

- 3 Configure BGP peering between the VTIs at both the sites. Peering ensures that BGP at the local site advertises the local subnets to the peer VPN gateway, and similarly BGP at the peer site advertises the remote subnets to the local VPN gateway. For more details about configuring BGP, see the Routing section in the *NSX Administration Guide*.

---

**Important** In NSX 6.4.2 and later, static routing and OSPF dynamic routing through an IPsec tunnel are not supported.

---

- 4 If you want to configure tunnel redundancy through more than one tunnel, configure BGP **Hold Down** timer and **Keep Alive** timer values. The timer values help in detecting loss of connectivity with the remote VPN gateway within the required failover time.

For a detailed example of configuring a route-based IPsec tunnel between a local NSX Edge and a remote Cisco CSR 1000V VPN Gateway, see [Using a Cisco CSR 1000V Appliance](#).

## Edit IPsec VPN Site

You can edit an IPsec VPN site.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > VPN > IPsec VPN**.
- 5 Select the IPsec VPN site that you want to edit.
- 6 Click the **Edit** (  or  ) icon.
- 7 Make the appropriate edits.
- 8 Click **Save** or **OK**, and then click **Publish Changes**.

## Disable IPsec VPN Site

You can disable an IPsec VPN site.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > VPN > IPsec VPN**.
- 5 Select the IPsec VPN site that you want to disable.

- 6 Disable the site.
  - In NSX 6.4.6 and later, click **Actions > Disable**.
  - In NSX 6.4.5 and earlier, click the **Disable** (  ) icon.

## Delete IPSec VPN Site

You can delete an IPSec VPN site.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > VPN > IPSec VPN**.
- 5 Select the IPSec VPN site that you want to delete.
- 6 Click the **Delete** (  or  ) icon.

## IPsec Terminology

IPSec is a framework of open standards. There are many technical terms in the logs of the NSX Edge and other VPN appliances that you can use to troubleshoot the IPSEC VPN.

The following terms are some of the standards that you might encounter:

- ISAKMP (Internet Security Association and Key Management Protocol) is a protocol defined by RFC 2408 for establishing Security Associations (SA) and cryptographic keys in an Internet environment. ISAKMP only provides a framework for authentication and key exchange and this framework is key exchange independent.
- Oakley is a key-agreement protocol that allows authenticated parties to exchange keying material across an insecure connection using the Diffie-Hellman key exchange algorithm.
- IKE (Internet Key Exchange) is a combination of ISAKMP framework and Oakley. NSX Edge provides IKEv1, IKEv2, and IKE-Flex.
- Diffie-Hellman (DH) key exchange is a cryptographic protocol that allows two parties that have no prior knowledge of each other to establish jointly a shared secret key over an insecure communications channel. VSE supports DH group 2 (1024 bits) and group 5 (1536 bits).

## IKEv1 Phase 1 and Phase 2

IKEv1 is a standard method used to arrange secure and authenticated communications.

## Phase 1 Parameters

Phase 1 sets up mutual authentication of the peers, negotiates cryptographic parameters, and creates session keys. The Phase 1 parameters used by NSX Edge are:

- Main mode.
- Triple DES, AES-128, AES-256 [Configurable]. AES-GCM is not supported in Phase 1, so AES-128 is used internally.
- SHA1, SHA\_256.
- MODP group 2, 5, 14, 15, and 16.
- Pre-shared secret key and certificate [Configurable].
- SA lifetime of 28800 seconds (eight hours) with no lifebytes rekeying.
- ISAKMP aggressive mode disabled

---

### Important

- IPsec VPN supports only time-based rekeying. You must disable lifebytes rekeying.
  - Starting in NSX 6.4.5, Triple DES cypher algorithm is deprecated in IPsec VPN service.
- 

## Phase 2 Parameters

IKE Phase 2 negotiates an IPsec tunnel by creating keying material for the IPsec tunnel to use (either by using the IKE phase 1 keys as a base or by performing a new key exchange). The IKE Phase 2 parameters supported by NSX Edge are:

- Triple DES, AES-128, AES-256, and AES-GCM [Matches the Phase 1 setting].
- SHA1, SHA\_256.
- ESP tunnel mode.
- MODP group 2, 5, 14, 15, and 16.
- Perfect forward secrecy for rekeying.
- SA lifetime of 3600 seconds (one hour) with no lifebytes rekeying.
- Selectors for all IP protocols, all ports, between the two networks, using IPv4 subnets

---

### Important

- IPsec VPN supports only time-based rekeying. You must disable lifebytes rekeying.
  - Starting with NSX 6.4.5, Triple DES cypher algorithm is deprecated in IPsec VPN service.
- 

## Transaction Mode Samples

NSX Edge supports Main Mode for Phase 1 and Quick Mode for Phase 2.

NSX Edge proposes a policy that requires PSK/Certificate, 3DES/AES128/AES256/AES-GCM, SHA1/SHA256, and DH Group 2/5/14/15/16. The peer must accept this policy. Otherwise, the negotiation phase fails.

## Phase 1: Main Mode Transactions

This example shows an exchange of Phase 1 negotiation initiated from a NSX Edge to a Cisco device.

The following transactions occur in a sequence between the NSX Edge and a Cisco VPN device in Main Mode.

- 1 NSX Edge to Cisco
  - Proposal: encrypt 3des-cbc, sha, psk, group5(group2)
  - DPD enabled
- 2 Cisco to NSX Edge
  - Contains proposal chosen by Cisco
  - If the Cisco device does not accept any of the parameters the NSX Edge sent in step 1, the Cisco device sends the message with flag NO\_PROPOSAL\_CHOSEN and ends the negotiation.
- 3 NSX Edge to Cisco
  - DH key and nonce
- 4 Cisco to NSX Edge
  - DH key and nonce
- 5 NSX Edge to Cisco (Encrypted)
  - Include ID (PSK).
- 6 Cisco to NSX Edge (Encrypted)
  - Include ID (PSK).
  - If the Cisco device finds that the PSK does not match, the Cisco device sends a message with flag INVALID\_ID\_INFORMATION, and Phase 1 fails.

## Phase 2: Quick Mode Transactions

The following transactions occur in a sequence between the NSX Edge and a Cisco VPN device in Quick Mode.

- 1 NSX Edge to Cisco

NSX Edge proposes Phase 2 policy to the peer. For example:

```
Aug 26 12:16:09 weiqing-desktop
ipsec[5789]:
"s1-c1" #2: initiating Quick Mode
PSK+ENCRYPT+TUNNEL+PFS+UP+SAREFTRACK
{using isakmp#1 msgid:d20849ac
proposal=3DES(3)_192-SHA1(2)_160
pfsgroup=OAKLEY_GROUP_MODP1024}
```

## 2 Cisco to NSX Edge

Cisco device sends back NO\_PROPOSAL\_CHOSEN if it does not find any matching policy for the proposal. Otherwise, the Cisco device sends the set of parameters chosen.

## 3 NSX Edge to Cisco

To facilitate debugging, you can enable IPsec logging on the NSX Edge and enable crypto debug on Cisco (debug crypto isakmp <level>).

# Configure Policy-Based IPsec VPN Site Example

This example contains a configuration scenario for a basic point-to-point policy-based IPsec VPN connection between an NSX Edge and a Cisco or WatchGuard VPN on the other end.

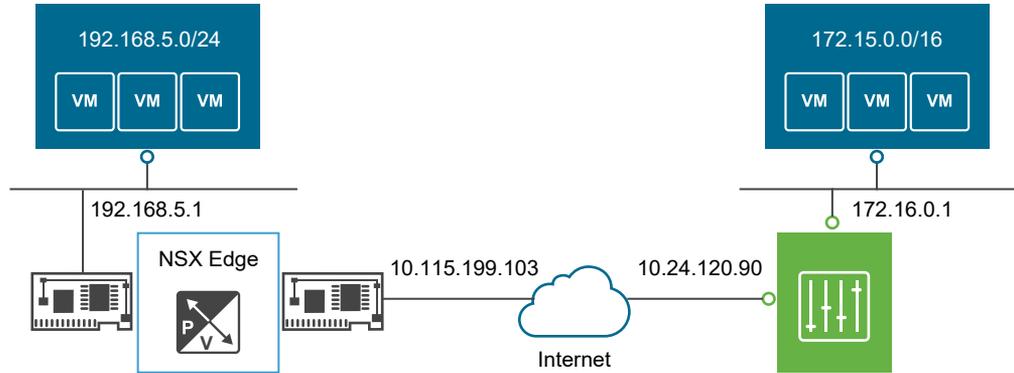
For this scenario, NSX Edge connects the internal network 192.168.5.0/24 to the Internet. NSX Edge interfaces are configured as follows:

- Uplink interface: 10.115.199.103
- Internal interface: 192.168.5.1

The VPN gateway at the remote site connects the 172.15.0.0/16 internal network to the Internet. The remote gateway interfaces are configured as follows:

- Uplink interface: 10.24.120.90
- Internal interface: 172.16.0.1

Figure 15-2. NSX Edge Connecting to a Remote VPN Gateway



**Note** For NSX Edge to NSX Edge IPsec tunnels, you can use the same scenario by setting up the second NSX Edge as the remote gateway.

#### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > VPN > IPsec VPN**.
- 5 Click **Add**.
- 6 In the **Name** text box, type a name for the IPsec VPN site.
- 7 In the **Local Id** text box, type **10.115.199.103** as the IP address of the NSX Edge instance. This local Id becomes the peer Id on the remote site.
- 8 In the **Local Endpoint** text box, type **10.115.199.103**.  
If you are adding an IP to IP tunnel using a pre-shared key, the local Id and local endpoint IP can be the same.
- 9 In the **Local Subnets** text box, type **192.168.5.0/24**.
- 10 In the **Peer Id**, type **10.24.120.90** to identify the peer site uniquely.
- 11 In the **Peer Endpoint** text box, type **10.24.120.90**.
- 12 In the **Peer Subnets** text box, type **172.15.0.0/16**.
- 13 Select the **IKE Version**. For example, select **IKEv2**.
- 14 Select the **Digest Algorithm**. For example, select **SHA\_256**.
- 15 Select the **Encryption Algorithm**. For example, select **AES**.
- 16 Select an **Authentication Method**. For example, select **PSK**.
- 17 Type the **Pre-shared Key**.

- 18 To display the pre-shared key on the peer site, click the **Show Pre-Shared Key** (🔑) icon or select the **Display Shared Key** check box.
- 19 Select the **Diffie-Hellman (DH) Group** cryptography scheme. For example, select **DH14**.
- 20 Click **Add** or **OK**.

The IPSec VPN site configuration is saved on the NSX Edge.

#### What to do next

Enable the IPSec VPN service.

## Using a Cisco 2821 Integrated Services Router

The following describes configurations performed using Cisco IOS.

### Procedure

#### 1 Configure Interfaces and Default Route

```
interface GigabitEthernet0/0
ip address 10.24.120.90 255.255.252.0
duplex auto
speed auto
crypto map MYVPN
!
interface GigabitEthernet0/1
ip address 172.16.0.1 255.255.0.0
duplex auto
speed auto
!
ip route 0.0.0.0 0.0.0.0 10.24.123.253
```

#### 2 Configure IKE Policy

```
Router# config term
Router(config)# crypto isakmp policy 1
Router(config-isakmp)# encryption 3des
Router(config-isakmp)# group 2
Router(config-isakmp)# hash sha
Router(config-isakmp)# lifetime 28800
Router(config-isakmp)# authentication
pre-share
Router(config-isakmp)# exit
```

#### 3 Match Each Peer with Its Pre-Shared Secret

```
Router# config term
Router(config)# crypto isakmp key vshield
address 10.115.199.103
Router(config-isakmp)# exit
```

#### 4 Define the IPSEC Transform

```
Router# config term
Router(config)# crypto ipsec transform-set
 myset esp-3des esp-sha-hmac
Router(config-isakmp)# exit
```

#### 5 Create the IPSEC Access List

```
Router# config term
Enter configuration commands, one per line.
End with CNTL/Z.
Router(config)# access-list 101 permit ip
 172.16.0.0 0.0.255.255 192.168.5.0 0.0.0.255
Router(config)# exit
```

#### 6 Bind the Policy with a Crypto Map and Label It

In the following example, the crypto map is labeled MYVPN

```
Router# config term
Router(config)# crypto map MYVPN 1
 ipsec-isakmp
% NOTE: This new crypto map will remain
 disabled until a peer and a valid
 access list have been configured.
Router(config-crypto-map)# set transform-set
 myset
Router(config-crypto-map)# set pfs group1
Router(config-crypto-map)# set peer
 10.115.199.103
Router(config-crypto-map)# match address 101
Router(config-crypto-map)# exit
```

### Example: Configuration

```
router2821#show running-config output
Building configuration...

Current configuration : 1263 bytes
!
version 12.4
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname router2821
!
boot-start-marker
boot-end-marker
!
! card type command needed for slot 0
! card type command needed for slot 1
enable password cisco
```

```
!
no aaa new-model
!
resource policy
!
ip subnet-zero
!
ip cef
!no ip dhcp use vrf connected
!
!
no ip ips deny-action ips-interface
!
crypto isakmp policy 1
encr 3des
authentication pre-share
group 2
crypto isakmp key vshield address 10.115.199.103
!
crypto ipsec transform-set myset esp-3des
 esp-sha-hmac
!
crypto map MYVPN 1 ipsec-isakmp
set peer 10.115.199.103
set transform-set myset
set pfs group1
match address 101
!
interface GigabitEthernet0/0
ip address 10.24.120.90 255.255.252.0
duplex auto
speed auto
crypto map MYVPN
!
interface GigabitEthernet0/1
ip address 172.16.0.1 255.255.0.0
duplex auto
speed auto
!
ip classless
ip route 0.0.0.0 0.0.0.0 10.24.123.253
!
ip http server
no ip http secure-server
!
access-list 101 permit ip 172.16.0.0
 0.0.255.255 192.168.5.0 0.0.0.255
!
control-plane
!
line con 0
line aux 0
line vty 0 4
password cisco
login
```

```

line vty 5 15
password cisco
login
!
scheduler allocate 20000 1000
!
end

```

## Using a Cisco ASA 5510

Use the following output to configure a Cisco ASA 5510.

```

ciscoasa# show running-config output
: Saved
:
ASA Version 8.2(1)18
!
hostname ciscoasa
enable password 2KFQnbNIdI.2KYOU encrypted
passwd 2KFQnbNIdI.2KYOU encrypted
names
!
interface Ethernet0/0
nameif untrusted
security-level 100
ip address 10.24.120.90 255.255.252.0
!
interface Ethernet0/1
nameif trusted
security-level 90
ip address 172.16.0.1 255.255.0.0
!
interface Ethernet0/2
shutdown
no nameif
no security-level
no ip address
!
interface Ethernet0/3
shutdown
no nameif
no security-level
no ip address
!
interface Management0/0
shutdown
no nameif
no security-level
no ip address
!
boot system disk0:/asa821-18-k8.bin
ftp mode passive
access-list ACL1 extended permit ip 172.16.0.0 255.255.0.0
 192.168.5.0 255.255.255.0

```

```

access-list ACL1 extended permit ip 192.168.5.0 255.255.255.0
 172.16.0.0 255.255.0.0
access-list 101 extended permit icmp any any
pager lines 24
mtu untrusted 1500
mtu trusted 1500
no failover
icmp unreachable rate-limit 1 burst-size 1
icmp permit any untrusted
icmp permit any trusted
no asdm history enable
arp timeout 14400
access-group 101 in interface untrusted
access-group 101 out interface untrusted
access-group 101 in interface trusted
access-group 101 out interface trusted
route untrusted 10.115.0.0 255.255.0.0 10.24.123.253 1
route untrusted 192.168.5.0 255.255.255.0 10.115.199.103 1
timeout xlate 3:00:00
timeout conn 1:00:00 half-closed 0:10:00
 udp 0:02:00 icmp 0:00:02
timeout sunrpc 0:10:00 h323 0:05:00 h225 1:00:00 mgcp 0:05:00
 mgcp-pat 0:05:00
timeout sip 0:30:00 sip_media 0:02:00 sip-invite 0:03:00
 sip-disconnect 0:02:00
timeout sip-provisional-media 0:02:00 uauth 0:05:00 absolute
timeout tcp-proxy-reassembly 0:01:00
dynamic-access-policy-record DfltAccessPolicy
no snmp-server location
no snmp-server contact
crypto ipsec transform-set MYSET esp-3des esp-sha-hmac
crypto ipsec security-association lifetime seconds 28800
crypto ipsec security-association lifetime kilobytes 4608000
crypto map MYVPN 1 match address ACL1
crypto map MYVPN 1 set pfs
crypto map MYVPN 1 set peer 10.115.199.103
crypto map MYVPN 1 set transform-set MYSET
crypto map MYVPN interface untrusted
crypto isakmp enable untrusted
crypto isakmp policy 1
authentication pre-share
encryption 3des
hash sha
group 2
lifetime 86400
telnet 10.0.0.0 255.0.0.0 untrusted
telnet timeout 5
ssh timeout 5
console timeout 0
no threat-detection basic-threat
no threat-detection statistics access-list
no threat-detection statistics tcp-intercept
username admin password f3UhLvUj1QsXsuK7 encrypted
tunnel-group 10.115.199.103 type ipsec-l2l
tunnel-group 10.115.199.103 ipsec-attributes

```

```
pre-shared-key *
!
!
prompt hostname context
Cryptochecksum:29c3cc49460831ff6c070671098085a9
: end
```

## Using a Cisco CSR 1000V Appliance

You want to configure route-based IPsec VPN tunnels between an NSX Edge and a Cisco CSR 1000V virtual appliance.

### NSX Edge Configuration

The following CLI output shows the route-based IPsec VPN configuration on the NSX Edge:

```
Edge IPsec VPN Config:
{
 "ipsec" : {
 "global" : {
 "extension" : null,
 "crlCertificates" : [],
 "pskForDynamicIp" : null,
 "id" : null,
 "caCertificates" : [],
 "serviceCertificate" : null
 },
 "logging" : {
 "logLevel" : "debug",
 "enable" : true
 },
 "disableEvent" : null,
 "enable" : true,
 "sites" : [
 {
 "name" : "VPN2 to edge-ext tun 2 192.168.14.2",
 "encryptionAlgorithm" : "3des",
 "psk" : "*****",
 "tunnelInterfaceId" : 1,
 "authenticationMode" : "psk",
 "peerIp" : "111.111.111.5",
 "ipsecSessionType" : "routebasedsession",
 "pskEncryption" : null,
 "digestAlgorithm" : "sha1",
 "enabled" : true,
 "localSubnets" : [
 "0.0.0.0/0"
],
 "description" : "VPN to edge subnet2",
 "mtu" : null,
 "peerId" : "111.111.111.5",
 "extension" : null,
 }
]
 }
}
```

```

 "ikeOption" : "ikev2",
 "localIp" : "51.51.51.1",
 "peerSubnets" : [
 "0.0.0.0/0"
],
 "responderOnly" : false,
 "certificate" : null,
 "dhGroup" : "dh2",
 "siteId" : "ipsecsite-53",
 "localId" : "51.51.51.1",
 "tunnelInterfaceLabel" : "vti-1",
 "enablePfs" : true
 },
 {
 "peerIp" : "71.71.71.5",
 "authenticationMode" : "psk",
 "ipsecSessionType" : "routebasedsession",
 "tunnelInterfaceId" : 2,
 "psk" : "*****",
 "name" : "VPN to edge-ext tun 1 192.168.13.2",
 "encryptionAlgorithm" : "3des",
 "description" : "VPN to edge subnet1",
 "localSubnets" : [
 "0.0.0.0/0"
],
 "enabled" : true,
 "pskEncryption" : null,
 "digestAlgorithm" : "sha1",
 "ikeOption" : "ikev2",
 "extension" : null,
 "peerSubnets" : [
 "0.0.0.0/0"
],
 "localIp" : "61.61.61.1",
 "peerId" : "71.71.71.5",
 "mtu" : null,
 "siteId" : "ipsecsite-54",
 "localId" : "61.61.61.1",
 "enablePfs" : true,
 "tunnelInterfaceLabel" : "vti-2",
 "responderOnly" : false,
 "certificate" : null,
 "dhGroup" : "dh2"
 }
]
}
}
}

```

The following CLI output shows the VTI configuration on the NSX Edge:

```

Edge VTI Tunnels Config:
{
 "vtiTunnels" : [
 {
 "name" : "vti-1",

```

```

 "mtu" : 1416,
 "label" : "vti-1",
 "sourceAddress" : "51.51.51.1",
 "destinationAddress" : "111.111.111.5",
 "tunnelAddresses" : [
 "192.168.14.2/24"
],
 "mode" : "VTI",
 "enabled" : true
 },
 {
 "enabled" : false,
 "tunnelAddresses" : [
 "192.168.13.2/24"
],
 "mode" : "VTI",
 "sourceAddress" : "61.61.61.1",
 "destinationAddress" : "71.71.71.5",
 "label" : "vti-2",
 "mtu" : 1416,
 "name" : "vti-2"
 }
]
}

```

## Cisco CSR 1000V Appliance Configuration

The following script configures the two matching route-based IPsec tunnels on the Cisco CSR 1000V appliance:

```

crypto ikev2 proposal PH1PROPOSAL
encryption 3des
integrity sha1
group 2
crypto ikev2 proposal PH2PROPOSAL
encryption 3des
integrity sha1
group 2
crypto ikev2 policy PH1POLICY
proposal PH1PROPOSAL
crypto ikev2 policy PH2POLICY
proposal PH2PROPOSAL
crypto ikev2 keyring PH1KEY
peer SITE1
 address 61.61.61.1
 pre-shared-key sharedvalue
!
crypto ikev2 keyring PH2KEY
peer SITE2
 address 51.51.51.1
 pre-shared-key sharedvalue
!
crypto ikev2 profile PH1PROFILE
match identity remote address 61.61.61.1 255.255.255.0

```

```

identity local address 71.71.71.5
authentication remote pre-share key sharedvalue
authentication local pre-share key sharedvalue
crypto ikev2 profile PH2PROFILE
match identity remote address 51.51.51.1 255.255.255.0
identity local address 111.111.111.5
authentication remote pre-share key sharedvalue
authentication local pre-share key sharedvalue
crypto ipsec transform-set TSET esp-3des esp-sha-hmac
mode tunnel
crypto ipsec profile IPSEC_PROF1
set transform-set TSET
set ikev2-profile PH1PROFILE
responder-only
crypto ipsec profile IPSEC_PROF2
set transform-set TSET
set ikev2-profile PH2PROFILE
responder-only

interface Tunnel1
ip address 192.168.13.1 255.255.255.0
tunnel source 71.71.71.5
tunnel mode ipsec ipv4
tunnel destination 61.61.61.1
tunnel protection ipsec profile IPSEC_PROF1
interface Tunnel2
ip address 192.168.14.1 255.255.255.0
tunnel source 111.111.111.5
tunnel mode ipsec ipv4
tunnel destination 51.51.51.1
tunnel protection ipsec profile IPSEC_PROF2
interface GigabitEthernet1
ip address dhcp
negotiation auto
interface GigabitEthernet2
no ip address
negotiation auto
interface GigabitEthernet2.2
encapsulation dot1Q 23
ip address 81.81.81.5 255.255.255.0
interface GigabitEthernet2.3
encapsulation dot1Q 19
ip address 111.111.111.5 255.255.255.0
interface GigabitEthernet2.4
encapsulation dot1Q 22
ip address 71.71.71.5 255.255.255.0

```

## Configuring a WatchGuard Firebox X500

You can configure your WatchGuard Firebox X500 as a remote gateway.

---

**Note** Refer to your WatchGuard Firebox documentation for exact steps.

---

### Procedure

- 1 In Firebox System Manager, select **Tools > Policy Manager**.
- 2 In Policy Manager, select **Network > Configuration**.
- 3 Configure the interfaces and click **OK**.
- 4 (Optional) Select **Network > Routes** to configure a default route.
- 5 Select **Network > Branch Office VPN > Manual IPSec** to configure the remote gateway.
- 6 In the IPSec Configuration dialog box, click **Gateways** to configure the IPSEC Remote Gateway.
- 7 In the IPSec Configuration dialog box, click **Tunnels** to configure a tunnel.
- 8 In the IPSec Configuration dialog box, click **Add** to add a routing policy.
- 9 Click **Close**.
- 10 Confirm that the tunnel is up.

## L2 VPN Overview

With L2 VPN, you can stretch multiple logical networks (both VLAN and VXLAN) across geographical sites. In addition, you can configure multiple sites on an L2 VPN server.

Virtual machines remain on the same subnet when they are moved between sites and their IP addresses do not change. Egress optimization enables Edge to route any packets sent towards the Egress Optimization IP address locally, and bridge everything else.

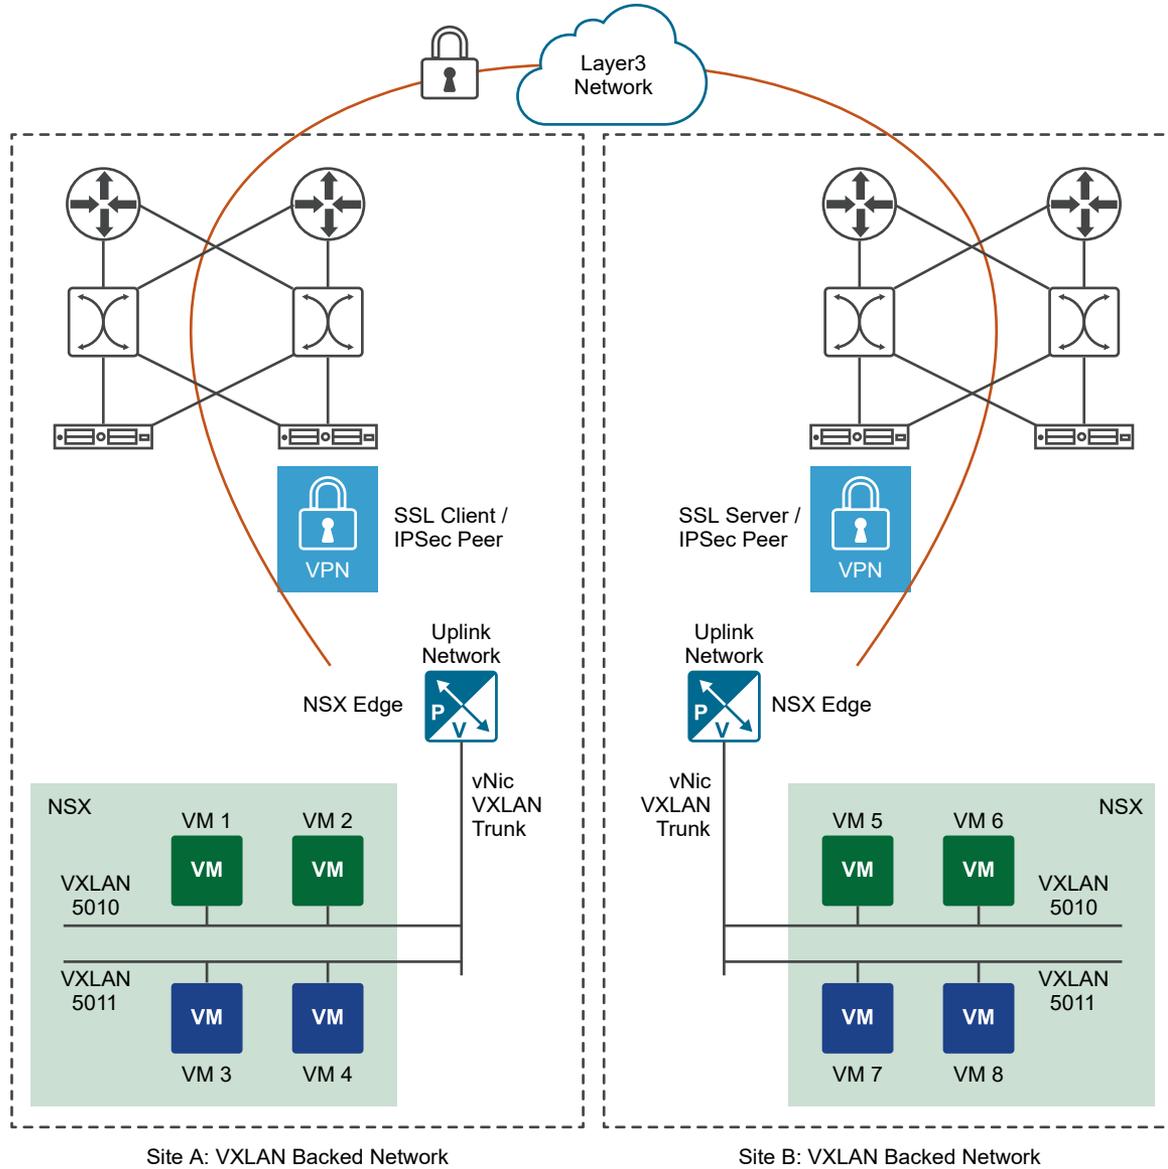
So, with L2 VPN service, enterprises can seamlessly migrate workloads between different physical sites. The workloads can run on either VXLAN-based networks or VLAN-based networks. For cloud service providers, L2 VPN provides a mechanism to on-board tenants without modifying existing IP addresses for workloads and applications.

---

### Note

- Starting in NSX Data Center 6.4.2, you can configure the L2 VPN service over both SSL and IPSec tunnels. However, you can configure the L2 VPN service over IPSec tunnels only by using REST APIs. For more information about configuring L2 VPN over IPSec, see the *NSX API Guide*.
  - With NSX 6.4.1 and earlier, you can configure the L2 VPN service only over SSL tunnels.
-

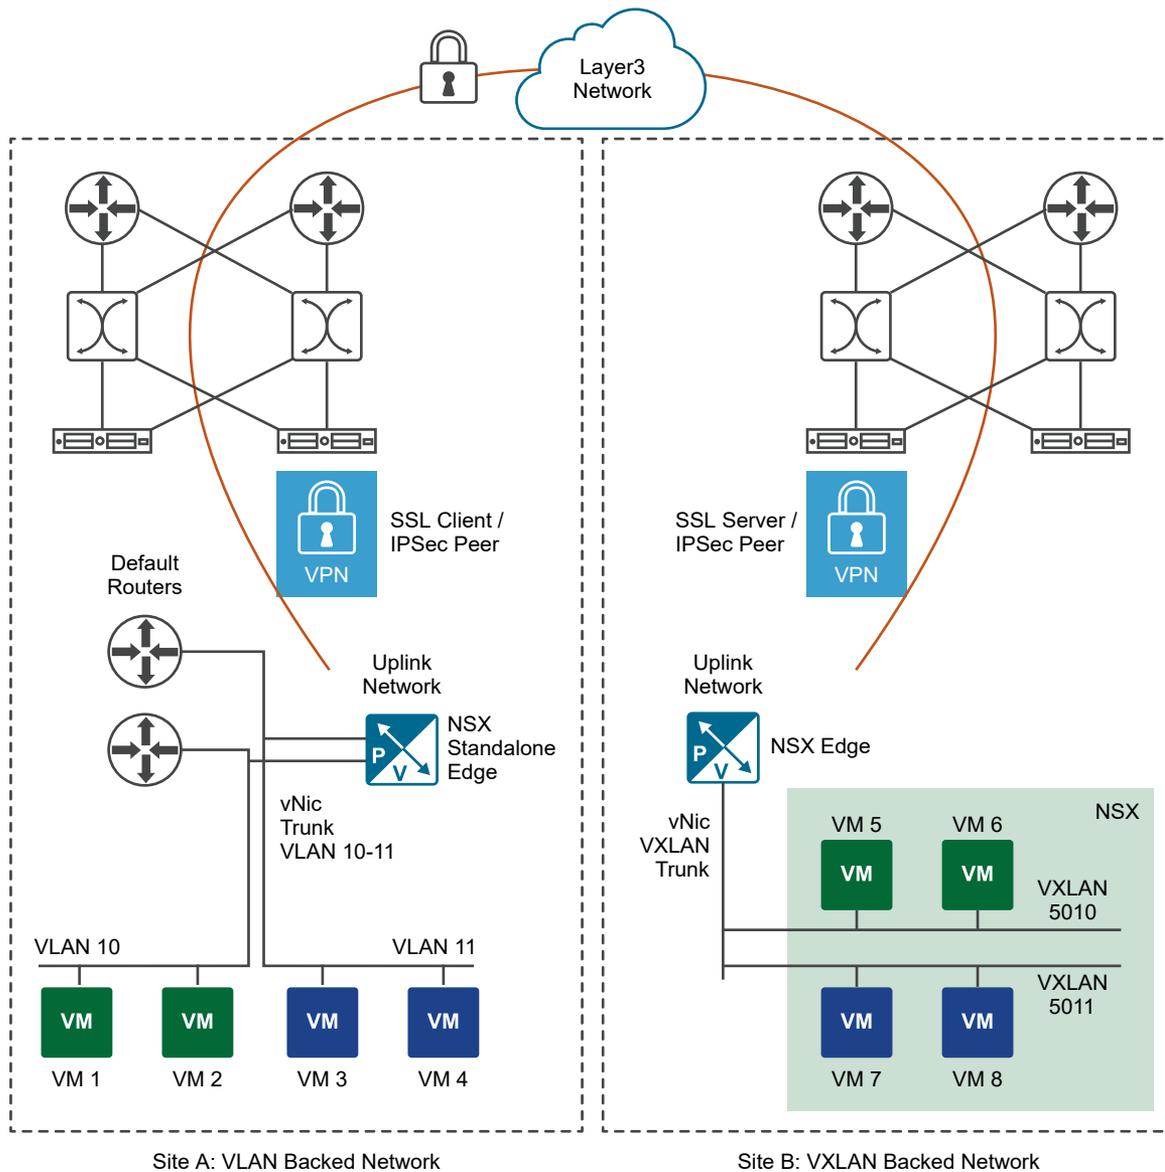
Figure 15-3. Extending VXLAN Across Multiple Sites Using L2 VPN



The L2 VPN client and server learn the MAC addresses on both local and remote sites based on the traffic flowing through them. Egress optimization maintains local routing because the default gateway for all virtual machines is always resolved to the local gateway using firewall rules. Virtual machines that have been moved to Site B can also access L2 segments that are not stretched on Site A.

If one of the sites does not have NSX deployed, a standalone Edge can be deployed on that site. In the following graphic, L2 VPN stretches network VLAN 10 to VXLAN 5010 and VLAN 11 to VXLAN 5011. So VM 1 bridged with VLAN 10 can access VMs 2, 5, and 6.

Figure 15-4. Extending non-NSX Site with VLAN-Based Networks to NSX-Site with VXLAN-Based Networks Using L2 VPN



## L2 VPN Best Practices

The best practices covered in this section apply to both L2 VPN over an SSL tunnel and L2 VPN over an IPsec tunnel.

**Recommendation** For an optimum L2 VPN performance, preferably deploy an NSX Edge with an XLarge form factor on both client and server for the following reasons:

- Increased TCP buffer size.
- CPU pinning is possible in vCPU 1, 3, and 5. CPU pinning is not possible in large and quad large form factors.

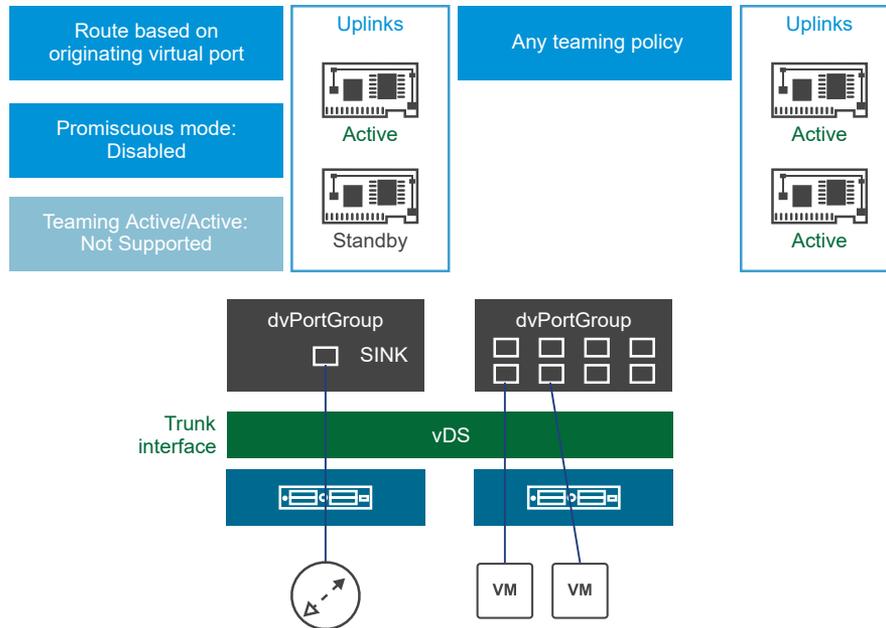
Configuring L2 VPN according to best practices can avoid problems such as looping and duplicate pings and responses.

## L2VPN Options to Mitigate Looping

There are two options to mitigate looping. Either the NSX Edges and VMs can be on different ESXi hosts, or the NSX Edges and VMs can be on the same ESXi host.

Option 1: Separate ESXi hosts for the L2VPN Edges and the VMs

### 1. Deploy L2VPN Edges and VMs on separate ESXi hosts



- 1 Deploy the Edges and the VMs on separate ESXi hosts.
- 2 Configure the Teaming and Failover Policy for the Distributed Port Group associated with the Edge's Trunk vNic as follows:
  - a Load balancing as "Route based on originating virtual port."
  - b Configure only one uplink as Active and the other uplink as Standby.
- 3 Configure the teaming and failover policy for the distributed port group associated with the VMs as follows:
  - a Any teaming policy is okay.
  - b Multiple active uplinks can be configured.

#### 4 Configure Edges to use sink port mode and disable promiscuous mode on the trunk vNic.

##### Note

- Disable promiscuous mode: If you are using vSphere Distributed Switch.
- Enable promiscuous mode: If you are using virtual switch to configure trunk interface.

If a virtual switch has promiscuous mode enabled, some of the packets that come in from the uplinks that are not currently used by the promiscuous port, are not discarded. You should enable and then disable `ReversePathFwdCheckPromisc` that will explicitly discard all the packets coming in from the currently unused uplinks, for the promiscuous port.

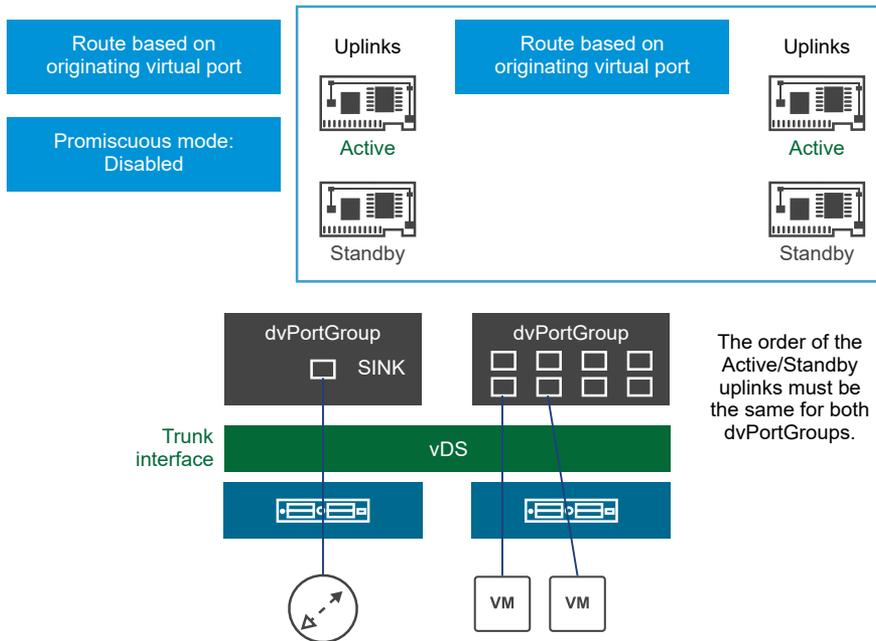
To block the duplicate packets, activate RPF check for the promiscuous mode from the ESXi CLI where NSX Edge is present:

```
esxcli system settings advanced set -o /Net/ReversePathFwdCheckPromisc -i 1
esxcli system settings advanced list -o /Net/ReversePathFwdCheckPromisc
Path: /Net/ReversePathFwdCheckPromisc
Type: integer
Int Value: 1
Default Int Value: 0
Max Value: 1
Min Value: 0
String Value:
Default String Value:
Valid Characters:
Description: Block duplicate packet in a teamed environment when the virtual switch is set to Promiscuous mode.
```

In **PortGroup** security policy, set **Promiscuous Mode** from **Accept** to **Reject** and back to **Accept** to activate the configured change.

- Option 2: Edges and VMs on the same ESXi host

## 2. Deploy L2VPN Edges and VMs on the same host



- a Configure the teaming and failover policy for the distributed port group associated with Edge's trunk vNic as follows:
  - 1 Load balancing as "Route based on originating virtual port."
  - 2 Configure one uplink as active and the other uplink as standby.
- b Configure the teaming and failover policy for the distributed port group associated with the VMs as follows:
  - 1 Any teaming policy is okay.
  - 2 Only one uplink can be active.
  - 3 The order of the active/standby uplinks must be the same for the VMs' distributed port group and the Edge's trunk vNic distributed port group.
- c Configure the client-side standalone edge to use sink port mode and disable promiscuous mode on the trunk vNic.

## Configure a Sink Port

When an NSX-managed Edge is set up as a L2 VPN client, some configuration is automatically done by NSX. When a standalone NSX Edge is set up as a L2 VPN client, these configuration steps must be done manually.

If one of the VPN sites does not have NSX deployed, you can configure an L2 VPN by deploying a standalone NSX Edge at that site. A standalone Edge is deployed using an OVF file on a host that is not managed by NSX. This deploys an Edge Services Gateway appliance to function as an L2 VPN client.

If a standalone edge trunk vNIC is connected to a vSphere Distributed Switch, either promiscuous mode or a sink port is required for L2 VPN function. Using promiscuous mode can cause duplicate pings and duplicate responses. For this reason, use sink port mode in the L2 VPN standalone NSX Edge configuration.

### Procedure

- 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:

```
<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.



The task topics in this section explain the steps to configure the L2 VPN service over SSL tunnels.

### Prerequisites

A sub interface must have been added on a trunk interface of the NSX Edge. See [Add a Sub Interface](#).

## Configure L2 VPN Server

The L2 VPN server is the destination NSX Edge to which the client is to be connected.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > VPN > L2 VPN**.
- 5 Next to **L2 VPN Mode**, select **Server**.
- 6 Next to **Global Configuration Details**, click **Edit** or **Change**.
- 7 In **Listener IP**, enter the primary or secondary IP address of an external interface of the NSX Edge.
- 8 The default port for the L2 VPN service is 443. Edit the port number, if necessary.
- 9 Select one or more encryption algorithms to encrypt the communication between the server and the client.
  - In NSX 6.4.6 and later, click the **Edit** (✎) icon. Select one or more encryption algorithms, and then click **Save**.
  - In NSX 6.4.5 and earlier, select an algorithm from the list box. To select multiple values, press Ctrl and click the algorithms in the list.
- 10 Select the certificate to be bound to SSL VPN server.

---

**Important** L2 VPN over SSL supports only RSA certificates.

---

- 11 Click **Save** or **OK**.

## Add Peer Sites

You can connect multiple sites to the L2 VPN server.

---

**Note** Changing site configuration settings causes the NSX Edge to disconnect and reconnect all existing connections.

---

### Procedure

- 1 Log in to the vSphere Web Client.

- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Next to **L2 VPN Mode**, select **Server**.
- 5 In **Site Configuration Details**, click **Add**.
- 6 Specify the configuration of the L2 VPN peer site.
  - a Enter a unique name for the peer site.
  - b Enter a user name and password with which the peer site is to be authenticated. User credentials on the peer site must be the same as those on the client side.
  - c In **Stretched Interfaces**, click  or **Select Sub Interfaces** to select the sub interfaces to be stretched with the client.
  - d Select the trunk interface for the Edge.

Sub interfaces configured on the trunk vNIC are displayed.
  - e Double-click the sub interfaces to be stretched.
  - f Click **Add** or **OK**.
  - g If the default gateway for virtual machines is the same across the two sites, enter the gateway IP addresses in the **Egress Optimization Gateway Address** text box. These IP addresses are the addresses for which the traffic is to be locally routed or for which the traffic is to be blocked over the tunnel.

- h (Optional) Enable **Unstretched Networks** when you want the VMs on the unstretched networks to communicate with the VMs that are behind the L2 VPN client edge on the stretched network. In addition, you want this communication to be routed through the same L2 VPN tunnel. Unstretched subnets can either be behind the L2 VPN server edge or the L2 VPN client edge or both.

For example, imagine that you have created an L2 VPN tunnel to stretch the 192.168.10.0/24 subnetwork between two data center sites using the NSX L2 VPN service.

Behind the L2 VPN server edge, you have two additional subnets (for example, 192.168.20.0/24 and 192.168.30.0/24). When unstretched networks are enabled, the VMs on 192.168.20.0/24 and 192.168.30.0/24 subnets can communicate with the VMs that are behind the L2 VPN client edge on the stretched network (192.168.10.0/24). This communication is routed through the same L2 VPN tunnel.

- i If you have enabled unstretched networks, do these steps depending on where the unstretched subnets are situated:
  - When unstretched subnets are behind the L2 VPN client edge, enter the network address of the unstretched network in the CIDR format while adding the peer (client) site on the L2 VPN server edge. To enter multiple unstretched networks, separate the network addresses by commas.
  - When unstretched subnets are behind the L2 VPN server edge, keep the **Unstretched Networks** text box blank. In other words, do not enter the network address of the unstretched networks while adding the client (peer) site on the L2 VPN server.

In the earlier example, because the unstretched subnets are behind the L2 VPN server edge, you must keep the **Unstretched Networks** text box blank in the **Add Peer Site** window.

- 7 Click **Add** or **OK**, and then click **Publish Changes**.

## Enable L2 VPN Service on Server

You must enable the L2 VPN service on the L2 VPN server (destination NSX Edge). If HA is already configured on this Edge appliance, ensure that Edge has more than one internal interface configured on it. If only a single interface is present and that has already been used by HA, L2 VPN configuration on the same internal interface fails.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click a destination NSX Edge, and navigate to **Manage > VPN > L2 VPN**.
- 4 Next to **L2 VPN Service Status**, click **Start**.

### What to do next

Create NAT or firewall rule on the Internet facing firewall side to enable the client and server to connect to each other.

## Configure L2 VPN Client

The L2 VPN client is the source NSX Edge that initiates a communication with the destination Edge (L2 VPN server).

You can also configure a standalone Edge as the L2 VPN client. See [Configure Standalone Edge as L2 VPN Client](#).

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an Edge that you want to configure as the L2 VPN client.
- 4 Click **Manage > VPN > L2 VPN**.
- 5 Next to **L2 VPN Mode**, select **Client**.
- 6 Next to **Global Configuration Details**, click **Edit** or **Change**.
- 7 Specify the L2 VPN client details.
  - a Enter the address of the L2 VPN server to which this client is to be connected. The address can be a host name or an IP address.
  - b Edit the default port to which the L2 VPN client must connect to, if necessary.
  - c Select the encryption algorithm for communicating with the server.
  - d In **Stretched Interfaces**, click  or **Select Sub Interfaces** to select the sub interfaces to be stretched to the server.
  - e Select the trunk interface for the Edge.

Sub interfaces configured on the trunk vNIC are displayed.
  - f Double-click the sub interfaces to be stretched and click **Add** or **OK**.
  - g In **Egress Optimization Gateway Address**, enter the gateway IP address of the sub interfaces or the IP addresses to which traffic should not flow over the tunnel.

- h (Optional) Select **Unstretched Networks** check box when you want the VMs on the unstretched networks to communicate with the VMs that are behind the L2 VPN server edge on the stretched network. In addition, you want this communication to be routed through the same L2 VPN tunnel. Unstretched subnets can either be behind the L2 VPN server edge or the L2 VPN client edge or both.

For example, imagine that you have created an L2 VPN tunnel to stretch the 192.168.10.0/24 subnetwork between two data center sites using the NSX L2 VPN service.

Behind the L2 VPN server edge, you have two additional subnets (for example, 192.168.20.0/24 and 192.168.30.0/24). When unstretched networks are enabled, the VMs on 192.168.20.0/24 and 192.168.30.0/24 subnets can communicate with the VMs that are behind the L2 VPN server edge on the stretched network (192.168.10.0/24). This communication is routed through the same L2 VPN tunnel.

- i If you have enabled unstretched networks, do these steps depending on where the unstretched subnets are situated:
  - When unstretched subnets are behind the L2 VPN server edge, enter the network address of the unstretched network in the CIDR format while configuring the L2 VPN client edge. To enter multiple unstretched networks, separate the network addresses by commas.
  - When unstretched subnets are behind the L2 VPN client edge, keep the **Unstretched Networks** text box blank. In other words, do not enter the network address of the unstretched networks on the L2 VPN client edge.

In the earlier example, because the unstretched subnets are behind the L2 VPN server edge, you must enter the unstretched networks as **192.168.20.0/24, 192.168.30.0/24** while configuring the L2 VPN client edge.

- j In **User Details**, type the user credentials to get authenticated at the server.
- 8 Click the **Advanced** tab and specify the other client details.
- a (Optional) Enable only secure proxy connections.
 

When a client Edge does not have direct access to the Internet and must reach the source (server) NSX Edge through a proxy server, you must specify proxy server settings.
  - b Enter the proxy server address, port, user name, and password.
  - c To enable server certificate validation, select **Validate Server Certificate** and select the appropriate CA certificate.
  - d Click **Save** or **OK**, and then click **Publish Changes**.

#### What to do next

Ensure that the Internet facing firewall allows traffic to flow from L2 VPN Edge to the Internet. The destination port is 443.

## Enable L2 VPN Service on Client

You must enable the L2 VPN service on the L2 VPN client (source NSX Edge).

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge that you configured in L2 VPN client mode.
- 4 Click **Manage > VPN > L2 VPN**.
- 5 Next to **L2 VPN Service Status**, click **Start**.

### What to do next

- To enable the client and server to connect to each other, create NAT or firewall rules on the Internet facing firewall side .
- If a trunk vNIC backed by a standard portgroup is being stretched, enable the L2 VPN traffic manually by doing the following steps:
  - a Set **Promiscuous mode** to **Accept**.
  - b Set **Forged Transmits** to **Accept**.

For more information about promiscuous mode operation and forged transmits, see *Securing vSphere Standard Switches* in the *VMware vSphere® Documentation*.

## View L2 VPN Statistics

You can view L2 VPN tunnel statistics, such as tunnel status, bytes sent and received, and other statistics, on both the L2 VPN server and client edges.

### Procedure

- 1 View statistics on the L2 VPN client edge.
  - a Double-click an NSX Edge that you configured in L2 VPN client mode.
  - b Click **Manage > VPN > L2 VPN**.
  - c Expand the **Tunnel Status** section, and click the **Refresh** icon to view the tunnel statistics.
- 2 View statistics on the L2 VPN server edge.
  - a Double-click an NSX Edge that you configured in L2 VPN server mode.
  - b Navigate to the **L2 VPN** page.
  - c In the **Site Configuration Details** section, click the **Show Statistics** or **Show L2VPN Statistics** link.

The statistics of all the peer sites that are configured on the L2 VPN server are displayed.

## What to do next

To see the networks configured on a trunk interface, navigate to **Manage > Settings > Interfaces** for the Edge and click **Trunk** in the Type column.

## L2 VPN Over IPsec

Starting with NSX Data Center 6.4.2, you can stretch your layer 2 networks between two sites with L2 VPN service over IPsec. Before configuring the L2 VPN service over IPsec, you must first create a route-based IPsec VPN tunnel. You then consume this route-based IPsec VPN tunnel to create an L2 VPN tunnel between the two sites.

You cannot create and edit a route-based IPsec VPN tunnel by using the vSphere Web Client. You must use the NSX REST APIs. For more information about creating route-based IPsec VPN tunnels, see the *NSX API Guide*.

## Workflow for Configuring L2 VPN Service Over IPsec

You must use NSX REST APIs to configure the L2 VPN service over IPsec on both the server Edge and the client Edge.

The steps in the workflow are supported only with NSX REST APIs. In this documentation, only the API URLs are mentioned. For a detailed information about the API parameters, sample requests and responses, see the *NSX API Guide*.

First, configure the L2 VPN service in the server (hub) mode on the NSX Edge by using the following steps. The Edge that you configure in the server mode must be an NSX Edge.

- 1 Create a route-based IPsec VPN tunnel with the Edge that you want to configure as the L2 VPN server (hub). A site ID is auto-generated when you create the tunnel.

```
PUT /api/4.0/edges/{edgeId}/ipsec/config
```

- 2 Create an L2 VPN tunnel for a client, and bind this L2 VPN tunnel with the site ID that was generated in step 1.

```
POST /api/4.0/edges/{edgeId}/l2t/config/l2tunnels
```

- 3 Retrieve the peer code for this client. This peer code becomes the input code (shared code) for configuring the L2 VPN service on the client Edge.

```
GET /api/4.0/edges/{edgeId}/l2t/config/l2tunnels/{l2tunnelId}/peercodes
```

- 4 Enable the L2 VPN over IPsec service.

```
POST /api/4.0/edges/{edgeId}/l2t/config
```

If you want to stretch the L2 network with other sites, repeat the preceding three steps on the server for the L2 VPN clients at other sites.

Now, configure the L2 VPN service in the client (spoke) mode on another Edge by using the following steps. This Edge can either be an NSX-managed Edge or a standalone Edge.

- 1 Create a route-based IPsec VPN tunnel with the same parameters that you used for configuring the route-based IPsec VPN tunnel on the server Edge.
- 2 Configure the Edge in spoke mode.

```
PUT /api/4.0/edges/{edgeId}/l2t/config/globalconfig
```

- 3 Create an L2 VPN tunnel by using the site ID that was generated on the server, and with the peer code that you retrieved from the server.
- 4 Enable the L2 VPN over IPsec service.

## Standalone Edge as L2 VPN Client

With NSX, you can set up L2 VPN connections with Edge devices or Edge virtual appliances on the client site regardless of whether NSX is deployed on the client sites.

Edge appliances deployed on a client site where NSX is not deployed are called standalone Edges. A standalone Edge is deployed using an OVF file on a host that is not managed by NSX.

### Configure Standalone Edge as L2 VPN Client

If one of the sites that you want to stretch is not backed by NSX, you can deploy a standalone Edge as the L2 VPN client on that site.

If you want to change FIPS mode for a standalone edge, use the `fips enable` or `fips disable` command. For more information, refer to *NSX Command Line Interface Reference*.

You can deploy a pair of standalone L2 VPN Edge clients and enable HA between them for VPN redundancy support. The two standalone L2 VPN Edge clients are called node 0 and node 1. It is not mandatory to specify the HA configuration settings on both standalone L2 VPN Edge appliance at the time of deployment. However, you must enable HA at the time of deployment.

The steps in the following procedure apply when you want to deploy the standalone Edge as a L2 VPN client for routing traffic either through an SSL tunnel or an IPsec VPN tunnel.

#### Prerequisites

You have created a trunk port group for the trunk interface of the standalone Edge to connect to. This port group requires some manual configuration:

- If the trunk port group is on a vSphere Standard Switch you must do the following:
  - Enable forged transmits.
  - Enable promiscuous mode.

See the *vSphere Networking Guide*.
- If the trunk port group is on a vSphere Distributed Switch you must do the following:
  - Enable forged transmits. See the *vSphere Networking Guide*.

- Enable sink port for the trunk vNic, or enable promiscuous mode. A good practice is to enable a sink port.

Sink port configuration must be done after the standalone Edge has been deployed, because you need to change the configuration of the port connected to the Edge trunk vNIC.

#### 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 standalone 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 standalone Edge OVF file and click **Next**.
- 5 On the OVF Template Details page, verify the template details and click **Next**.
- 6 On the Select name and folder page, type a name for the standalone Edge and select the folder or data center where you want to deploy. Then click **Next**.
- 7 On the Select storage page, select the location to store the files for the deployed template.
- 8 On the Select networks page, configure the networks the deployed template must use. Click **Next**.
  - The Public interface is the uplink interface.
  - The Trunk interface is used to create subinterfaces for the networks that will be stretched. Connect this interface to the trunk port group you created.
  - The HA interface is used to set up high availability on the standalone L2 VPN Edge appliances. Select a distributed port group for the HA interface.
- 9 On the Customize Template page, specify the following values.
  - 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 Type the uplink IP address and prefix length, and optionally default gateway and DNS IP address.
  - e Select the cipher to be used for authentication. The selected value must match the cipher used on the L2 VPN server.

---

**Note** Perform this step only when you want to configure L2 VPN over SSL.

---

  - f To enable Egress Optimization, type the gateway IP addresses for which traffic should be locally routed or for which traffic is to be blocked over the tunnel.

- g (Optional) Select the **Enable TCP Loose Setting** check box when you want the existing TCP connection (for example, an SSH session) to the VM over L2 VPN to remain active after the VM is migrated.

By default, this setting is not enabled. When this setting is disabled, the existing TCP connection to the VM over L2 VPN is lost after the VM is migrated. You must open a new TCP connection to the VM after the migration is done.

- h To enable high availability on the standalone L2 VPN Edge appliance, select the **Enable High Availability for this appliance** check box.
- i (Optional) Type the IP address of the first standalone L2 VPN Edge appliance (node 0). The IP address must be in the /30 IP subnet.
- j (Optional) Type the IP address of the second standalone L2 VPN Edge appliance (node 1). The IP address must be in the /30 IP subnet.
- k (Optional) On node 0 appliance, select 0 to assign the IP address of node 0 for the HA interface. Similarly, on node 1 appliance, select 1 so that IP address of node 1 is used for the HA interface.
- l (Optional) Specify an integer value for the dead interval time in seconds. For example, type **15**.

---

**Note** If you specify HA configuration settings during deployment of a standalone L2 VPN Edge client, the standalone Edge appliance VM gets a standby HA status until HA is fully configured. This means that the network interfaces on the Edge appliance VM are disabled. Make sure that you set the HA peer nodes to create the active and standby appliances. For detailed instructions, see [Configure HA on Standalone L2 VPN Clients](#).

---

- m Type the L2 VPN server address and port.

If you are configuring the L2 VPN client to route traffic through the IPsec VPN tunnel, you must specify the IP address of the peer site, and the peer code.

- n Type the user name and password with which the peer site is to be authenticated.

---

**Note** Perform this step only when you want to configure L2 VPN over SSL.

---

- o In Sub Interfaces VLAN (Tunnel ID), type VLAN ID(s) of the networks you want to stretch. You can list the VLAN IDs as a comma-separated list or range. For example, 2,3,10-20.

If you want to change the VLAN ID of the network before stretching it to the standalone Edge site, type the VLAN ID of the network, and then type the tunnel ID in brackets. For example, 2(100),3(200). The Tunnel ID is used to map the networks that are being stretched. However, you cannot specify the tunnel ID with a range. So this might not be allowed: 10(100)-14(104). You might need to rewrite this as 10(100),11(101),12(102),13(103),14(104).

- p If the standalone Edge does not have direct access to the Internet and must reach the source (server) NSX Edge through a proxy server, type the proxy address, port, user name, and password.
- q If a Root CA is available, you can paste it in the Certificate section.
- r Click **Next**.

**10** On the Ready to complete page, review the standalone Edge settings and click **Finish**.

#### What to do next

- Power on the standalone Edge appliance.
- Note the trunk vNIC port number and configure a sink port. See [Configure a Sink Port](#).
- If you have specified the HA configuration settings, such as HA IP address, HA index value, and the dead interval time while deploying the standalone L2 VPN Edge appliances, you can validate the HA configuration on the console of the deployed nodes with the `show configuration` command.
- If you have not specified the HA configuration settings during deployment, you can do it later from the NSX Edge Console by running the `ha set-config` command on each node.

Make any further configuration changes with the standalone Edge command-line interface. See the *NSX Command Line Interface Reference*.

## Configure HA on Standalone L2 VPN Clients

Log in to the NSX Edge console of standalone L2 VPN Edge appliances to specify the HA configuration settings and establish HA between both L2 VPN Edge appliances.

You have deployed two standalone L2 VPN Edge clients called L2VPN-Client-01 and L2VPN-Client-02 with the same configuration. The /30 IP subnet address of L2VPN-Client-01 is 192.168.1.1, and for L2VPN-Client-02 is 192.168.1.2. In this topic, Node-1 refers to L2VPN-Client-01 and Node-2 refers to L2VPN-Client-02.

#### Prerequisites

- Enable HA on both the L2 VPN Edge appliances.
- Make sure that HA configuration settings, such as HA IP address, HA index value, and dead interval time are configured on both nodes.
- Make sure that both standalone L2 VPN Edge clients have the same VPN configuration.

**Procedure**

- 1 Log in to each node, and run the `ha get-local node` command on both nodes individually to retrieve the MAC addresses of the three vNIC interface cards.

For example, on Node-1:

```
nsx-12vpn-edge(config)# ha get-localnode
00:50:56:90:12:ea 00:50:56:90:97:ca 00:50:56:90:d9:69
```

For example, on Node-2:

```
nsx-12vpn-edge(config)# ha get-localnode
00:50:56:90:1c:75 00:50:56:90:34:c1 00:50:56:90:36:80
```

- 2 Run the `ha set-peernode` command on both nodes individually to assign the MAC address of Node-1 to Node-2, and MAC address of Node-2 to Node-1.

For example, assign the MAC address of Node-2 on Node-1:

```
nsx-12vpn-edge(config)# ha set-peernode 00:50:56:90:1c:75 00:50:56:90:34:c1
00:50:56:90:36:80
```

For example, assign the MAC address of Node-1 to Node-2:

```
nsx-12vpn-edge(config)# ha set-peernode 0:50:56:90:12:ea 00:50:56:90:97:ca
00:50:56:90:d9:69
```

- 3 To start HA, run the `ha admin-state UP` command on each node.

For example, on both Node-1 and Node-2:

```
nsx-12vpn-edge(config)# ha admin-state UP
```

---

**Note** Make sure that you type UP in uppercase, as shown in the example.

---

- 4 Run the `commit` command on both nodes individually to save the HA configuration and establish HA between Node-1 and Node-2.

For example, on both Node-1 and Node-2:

```
nsx-12vpn-edge(config)# commit
High Availability Feature is enabled on this appliance. Please make sure to make
similar configuration on paired Standalone Edge appliance.
```

**Results**

HA configuration is saved on both nodes, and HA is established between the nodes.

**What to do next**

Log in to each node, and verify the HA status by running the `show service highavailability` command on both nodes.

The CLI on L2VPN-Client-01 node shows the following output:

```

HA Status of L2VPN-Client-1 is Active
Highavailability Service:
Highavailability Status: Active
Highavailability Status since: 2017-10-27 14:47:03:904
Highavailability Unit Id: 0
Highavailability State: Up
Highavailability Admin State: Up
Highavailability Running Nodes: 0, 1
Unit Poll Policy:
 Frequency: 3.75 seconds
 Deadtime: 15 seconds
Highavailability Status:
 Highavailability Config Channel: Up
 Highavailability Status Channel: Up
Highavailability Healthcheck Status:
 This unit [0]: Up Active:1
 Peer unit [1]: Up Active:0
 Session via vNIC_2:192.168.1.1:192.168.1.2 Up
Config Engine:
 HA Configuration: Enabled
 HA Admin State: Up
 Config Engine Status: Active
 ForceStandby Flag: Disabled
Highavailability Stateful Logical Status:
 File-Sync running
 Connection-Sync running

```

Reachability Status of Peer Node 192.168.1.2 is Up

The CLI on L2VPN-Client-02 node shows the following output:

```

HA Status of L2VPN-Client-2 is Standby
Highavailability Service:
Highavailability Status: Standby
Highavailability Status since: 2017-10-27 14:47:19:555
Highavailability Unit Id: 1
Highavailability State: Up
Highavailability Admin State: Up
Highavailability Running Nodes: 0, 1
Unit Poll Policy:
 Frequency: 3.75 seconds
 Deadtime: 15 seconds
Highavailability Status:
 Highavailability Config Channel: Up
 Highavailability Status Channel: Up
Highavailability Healthcheck Status:
 Peer unit [0]: Up Active:1
 Session via vNIC_2: 192.168.1.2:192.168.1.1 Up
 This unit [1]: Up Active:0
Config Engine:
 HA Configuration: Enabled
 HA Admin State: Up
 Config Engine Status: Standby
 ForceStandby Flag: Disabled
Highavailability Stateful Logical Status:
 File-Sync running
 Connection-Sync running

```

Reachability Status of Peer Node 192.168.1.1 is Up

At any stage, if you want to check the link status of the HA, run the `show service highavailability link` command on each node. This command lists the local and peer /30 IP subnet addresses that you configured while deploying the standalone L2 VPN nodes.

For example, on Node-1:

```

nsx-l2vpn-edge> show service highavailability link
Local IP address: 192.168.1.1/30
Peer IP Address: 192.168.1.2/30

```

## Approaches to Disable HA on Standalone L2 VPN Edges

NSX provides two approaches to disable HA on standalone L2 VPN Edge appliances.

Consider that you have deployed two standalone L2 VPN Edge appliances called L2VPN-Client-1 and L2VPN-Client-2, and you have established HA between both these appliances. L2VPN-Client-1 is the active appliance, and L2VPN-Client-2 is the standby appliance.

In the first approach, you can do the following steps to disable HA:

- 1 Delete or power off the standby appliance (L2-VPN-Client-2) directly.
- 2 Log in to the console of the active appliance (L2-VPN-Client-1), and run the `ha disable` command.

The advantage of this approach is that the HA failover time is minimum. However, the limitations of this approach are as follows:

- You have to power off the standby L2 VPN appliance manually.
- This approach might lead to a dual active state situation. For example, you might forget to power off the standby appliance and start up the standby appliance. This might cause both L2 VPN appliances to become active.

In the second approach, you can do the following steps to disable HA:

- 1 Delete or power off any one of the L2 VPN appliances, either active or standby appliance.
- 2 Log in to the console of the other appliance, and run the `ha disable` command to disable the HA feature on this appliance.

In the second approach too, the HA failover time is minimum. However, the limitations of the second approach are as follows:

- You have to power off the L2 VPN appliance manually.
- This approach might also lead to a dual active state situation. For example, you might forget to power off the appliance and start up the same appliance. This might cause both L2 VPN appliances to become active.
- Service is disrupted when you delete the active appliance because the HA failover might not happen immediately to make the standby appliance active.

## Replace a Failed Standalone L2 VPN Client

If a standalone L2 VPN Client has failed or crashed, you can replace this failed appliance by deploying a new standalone L2 VPN client, and configure this newly deployed appliance from the NSX Edge console.

Consider that you have already deployed two standalone L2 VPN client appliances called L2VPN-Client-01 and L2VPN-Client-02, and enabled HA on both the appliances. The L2VPN-Client-01 node has failed or crashed. To replace this failed node, you will deploy a new standalone L2 VPN client appliance called L2VPN-Client-Replace, and specify the same VPN configuration as the active node.

### Procedure

- 1 Log in to the console of the active node (L2VPN-Client-02), and check its HA index.
- 2 Run the `ha get-localnode` command on the L2VPN-Client-02 node to retrieve the vNIC MAC addresses, and copy the CLI output of this command.

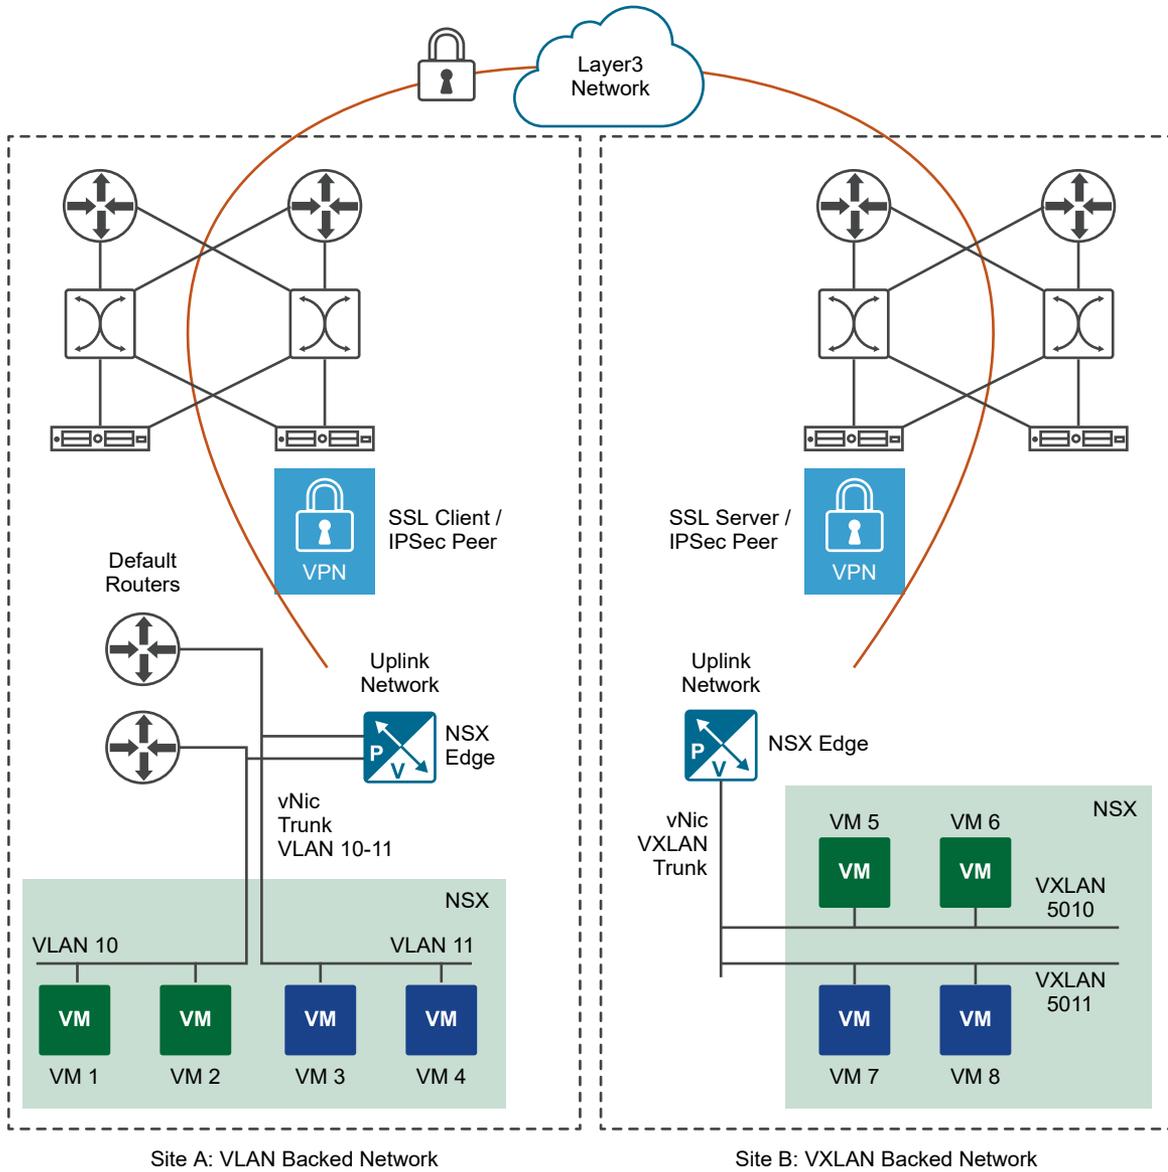
- 3 Deploy a new standalone L2 VPN client appliance and name it L2VPN-Client-Replace. During deployment, make sure that you specify the following details:
  - a Enable the HA feature.
  - b Type the correct HA IP address for both node 0 and node 1. The IP addresses must be in the /30 IP subnet.
  - c Select the HA index value.
    - If the node with HA index 1 has failed, then select 0 for the HA index of the L2VPN-Client-Replace appliance.
    - If the node with HA index 0 has failed, then select 1 for the HA index of the L2VPN-Client-Replace appliance.
- 4 Log in to the console of the new L2VPN-Client-Replace appliance, and do these steps:
  - a Run the `ha set-peernode` command and set the MAC address of the peer node (L2VPN-Client-02).
  - b Run the `ha get-localnode` command, and copy the CLI output of this command.
- 5 Log in to the console of the active node (L2VPN-Client-02), and run the `ha set-peernode` command to set the vNIC MAC addresses of the newly deployed appliance (L2VPN-Client-Replace).
- 6 Finally, run the `commit` command on both L2VPN-Client-02 and L2VPN-Client-Replace appliances.

## Scenario: Add a Stretched VLAN or VXLAN Network

Company ACME Enterprise has two private data centers "site A" and "site B". NSX Data Center is deployed at both data centers. As an NSX administrator, you want to migrate workloads (applications) from site A to site B by stretching or extending the VLAN networks on site A to the VXLAN networks on site B.

NSX L2 VPN supports egress optimization by using the same gateway IP address on both sites. This scenario uses the egress optimization feature and ensures that the IP addresses of the applications do not change after the migration.

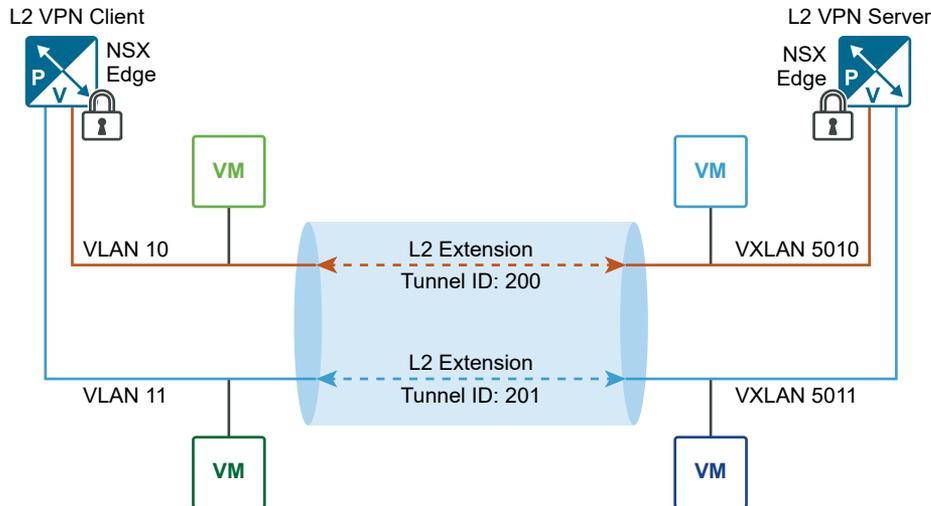
The following figure shows the logical topology of extending networks between two sites by using the L2 VPN service on the NSX Edges.



The L2 VPN service on the NSX Edge at site A is configured in "client" mode, and the L2 VPN service on NSX Edge at site B is configured in "server" mode. As an administrator, your objective is to create a L2 VPN tunnel and perform an L2 extension between sites A and B, such that:

- Tunnel ID 200 extends the VLAN 10 network on site A to the VXLAN 5010 network on site B.
- Tunnel ID 201 extends the VLAN 11 network on site A to the VXLAN 5011 network on site B.

The following figure shows the logical representation of the L2 extension between both sites.



**Remember** In this scenario, both sites have NSX-managed edges. To perform an L2 extension between two sites, the edge that is configured in "server" mode must be an NSX Edge. However, the edge that is configured in "client" mode can either be an NSX Edge or a standalone edge, which is not NSX-managed.

If the client site uses a standalone edge, you can stretch only VLAN networks on the client site with the VLAN or VXLAN networks on the server site.

You can perform an L2 extension either by configuring L2 VPN service over SSL, or by configuring L2 VPN over IPsec. The following procedure explains the steps for stretching L2 networks using L2 VPN over SSL.

#### Procedure

- 1 Navigate to the L2 VPN edge on site B and configure a vnic interface of type "trunk". Add two sub interfaces on this interface.

For detailed instructions about configuring an interface on the edge and adding sub interfaces, see [Configure an Interface](#).

For example, in this scenario, configure "vnic 1" on the server edge to connect to a distributed port group. Add sub interfaces that connect to logical switches with VNI 5010 and 5011. Each sub interface must have a unique tunnel ID. The following table shows the sub interface configuration on the L2 VPN server edge.

**Table 15-9. Sub Interfaces on vnic 1 of L2 VPN Server Edge**

Name	IP Addresses	Network	VNI	Tunnel ID	Status
sub_vxlan1	192.168.10.10/24	VXLAN-Network1	5010	200	Connected
sub_vxlan2	192.168.100.10/24	VXLAN-Network2	5011	201	Connected

- 2 Navigate to the L2 VPN edge on site A and configure a vnic interface of type "trunk". Add two sub interfaces on this interface.

For example, in this scenario, configure "vnic 2" on the client edge to connect to a standard port group. Add sub interfaces that connect to VLANs 10 and 11. The tunnel IDs on the client edge must match the tunnel IDs that you specified on the server edge. The following table shows the sub interface configuration on the L2 VPN client edge.

**Table 15-10. Sub Interfaces on vnic 2 of L2 VPN Client Edge**

Name	IP Addresses	VLAN	Tunnel ID	Status
sub_vlan1	192.168.10.10/24	10	200	Connected
sub_vlan2	192.168.100.10/24	11	201	Connected

- 3 Configure the L2 VPN edge at site B.

- a Set the L2 VPN mode as **Server**.
- b Specify the Global Configuration settings.

For detailed instructions about configuring the L2 VPN server, see [Configure L2 VPN Server](#).

- c In Site Configuration Details, click **Add** and specify the configuration of the L2 VPN client (peer) site.

For detailed instructions about adding L2 VPN peer sites, see [Add Peer Sites](#).

For example, in this scenario, do the following peer site configuration:

- Add a peer site with name "site-A". Select the "vnic 1" trunk interface on the server edge, and include the two sub interfaces "sub\_vxlan1" and "sub\_vxlan2" as the stretched networks. Ensure that you enable the peer site. The following table shows the sub interfaces (stretched interfaces) on peer site-A.

**Table 15-11. Sub Interfaces on Peer Site-A**

Name	Parent Index	Parent Name	IP Addresses	Network	VNI	Tunnel ID
sub_vxlan1	1	vnic1	192.168.10.1 0/24	VXLAN- Network1	5010	200
sub_vxlan2	1	vnic1	192.168.100. 10/24	VXLAN- Network2	5011	201

- In the **Egress Optimization Gateway Address** text box, enter **192.168.10.10,192.168.100.10**.
- d Start the L2 VPN service on the server edge.
  - e Publish the changes.

- 4 Configure the L2 VPN edge at site A.
  - a Set the L2 VPN mode as **Client**.
  - b Specify the Logging Configuration and Global Configuration settings.

For detailed instructions about configuring the L2 VPN client, see [Configure L2 VPN Client](#).

For example, in this scenario, do the following configuring on the L2 VPN Client:

- Select the "vnic 2" trunk interface on the client edge, and include the two sub interfaces "sub\_vlan1" and "sub\_vlan2" as the stretched networks. The following table shows the sub interfaces (stretched interfaces) on the L2 VPN client edge.

**Table 15-12. Stretched Interfaces on L2 VPN Client Edge**

Name	Parent Index	Parent Name	IP Addresses	VLAN	Tunnel ID
sub_vlan1	2	vnic2	192.168.10.10/24	10	200
sub_vlan2	2	vnic2	192.168.100.10/24	11	201

- In the **Egress Optimization Gateway Address** text box, enter **192.168.10.10,192.168.100.10**.
- c Start the L2 VPN service on the client edge.
  - d Publish the changes.

## Results

L2 VPN tunnel is established between site A and site B. You can now migrate workloads between the two sites by using the stretched L2 networks.

## What to do next

On the L2 VPN server edge and client edge:

- Verify that the L2 VPN tunnel status is "Up".
- View the tunnel statistics.

For detailed instructions, see [View L2 VPN Statistics](#).

Alternatively, you can log in to the CLI console of the L2 VPN server edge and client edge and verify the tunnel status by running the `show service l2vpn` command.

For more information about this command, see the *NSX Command Line Interface Reference Guide*.

## Scenario: Remove a Stretched VLAN or VXLAN Network

In this scenario, your objective is to remove the L2 stretching that extends the VLAN 10 network on site A to the VXLAN 5010 network on site B.

## Prerequisites

Make sure that you have configured L2 stretching between VLAN networks (VLAN ID: 10, 11) on site A to the VXLAN networks (VNI: 5010, 5011) on site B. See [Scenario: Add a Stretched VLAN or VXLAN Network](#).

## Procedure

- 1 Log in to the vSphere Web Client.
- 2 (Required) Navigate to the L2 VPN client edge on site A and stop the L2 VPN service.
- 3 (Required) Navigate to the L2 VPN server edge on site B and stop the L2 VPN service.
- 4 On the L2 VPN server edge, delete the "sub\_vxlan1" sub interface on the "vnic1" trunk interface.
  - a Navigate to the edge interface settings by clicking **Manage > Settings > Interfaces**.
  - b Select the **vnic1** interface and click the **Edit** (  or  ) icon.
  - c In **Sub Interfaces**, select **sub\_vxlan1**, and click the **Delete** (  or  ) icon.
  - d Click **Save** or **OK**.
- 5 On the L2 VPN client edge, delete the "sub\_vlan1" sub interface on the "vnic2" trunk interface.
  - a Navigate to the edge interface settings by clicking **Manage > Settings > Interfaces**.
  - b Select the **vnic2** interface and click the **Edit** (  or  ) icon.
  - c In **Sub Interfaces**, select **sub\_vlan1**, and click the **Delete** (  or  ) icon.
  - d Click **Save** or **OK**.

## Results

The sub interfaces that extend the VLAN 10 network on site A to the VXLAN 5010 network on site B are removed.

## What to do next

- On site A, navigate to the **L2 VPN** page of the client edge. Observe that in **Stretched Interfaces**, the "sub\_vlan1" interface is not shown. The other stretched interface "sub\_vlan2" still exists.
- On site B, navigate to the **L2 VPN** page of the server edge. Observe that in the **Stretched Interfaces** of the peer site (site-A), the "sub\_vxlan1" interface is not shown. The other stretched interface "sub\_vxlan2" still exists.

# Logical Load Balancer

# 16

The NSX Edge load balancer enables high-availability service and distributes the network traffic load among multiple servers. It distributes incoming service requests evenly among multiple servers in such a way that the load distribution is transparent to users. Load balancing thus helps in achieving optimal resource utilization, maximizing throughput, minimizing response time, and avoiding overload. NSX Edge provides load balancing up to Layer 7.

You map an external, or public, IP address to a set of internal servers for load balancing. The load balancer accepts TCP, UDP, HTTP, or HTTPS requests on the external IP address and decides which internal server to use. Port 80 is the default port for HTTP and port 443 is the default port for HTTPS.

You must have a working NSX Edge instance before you can configure load balancing. For information on setting up NSX Edge, see [NSX Edge Configuration](#).

For information on configuring an NSX Edge certificate, see [Working with Certificates](#).

NSX load balancing features are as follows:

- Protocols: TCP, UDP, HTTP, HTTPS
- Algorithms: Weighted round robin, IP hash, URI, least connection
- SSL termination with AES-NI acceleration
- SSL bridging (client-side SSL + server-side SSL)
- SSL certificates management
- X-header forwarding for client identification
- L4/L7 transparent mode
- Connection throttling
- Enable/disable individual servers (pool members) for maintenance
- Health check methods (TCP, UDP, HTTP, HTTPS)
- Enhanced health check monitor
- Persistence/sticky methods: SourceIP, MSRDP, COOKIE, SSLSESSIONID
- One-arm mode
- Inline mode

- URL rewrite and redirection
- Application Rules for advanced traffic management
- HA session sticky support for L7 proxy load balancing
- IPv6 support
- Enhanced load balancer CLI for troubleshooting
- Available on all flavors of an NSX edge services gateway, with a recommendatory of X-Large or Quad Large for production traffic

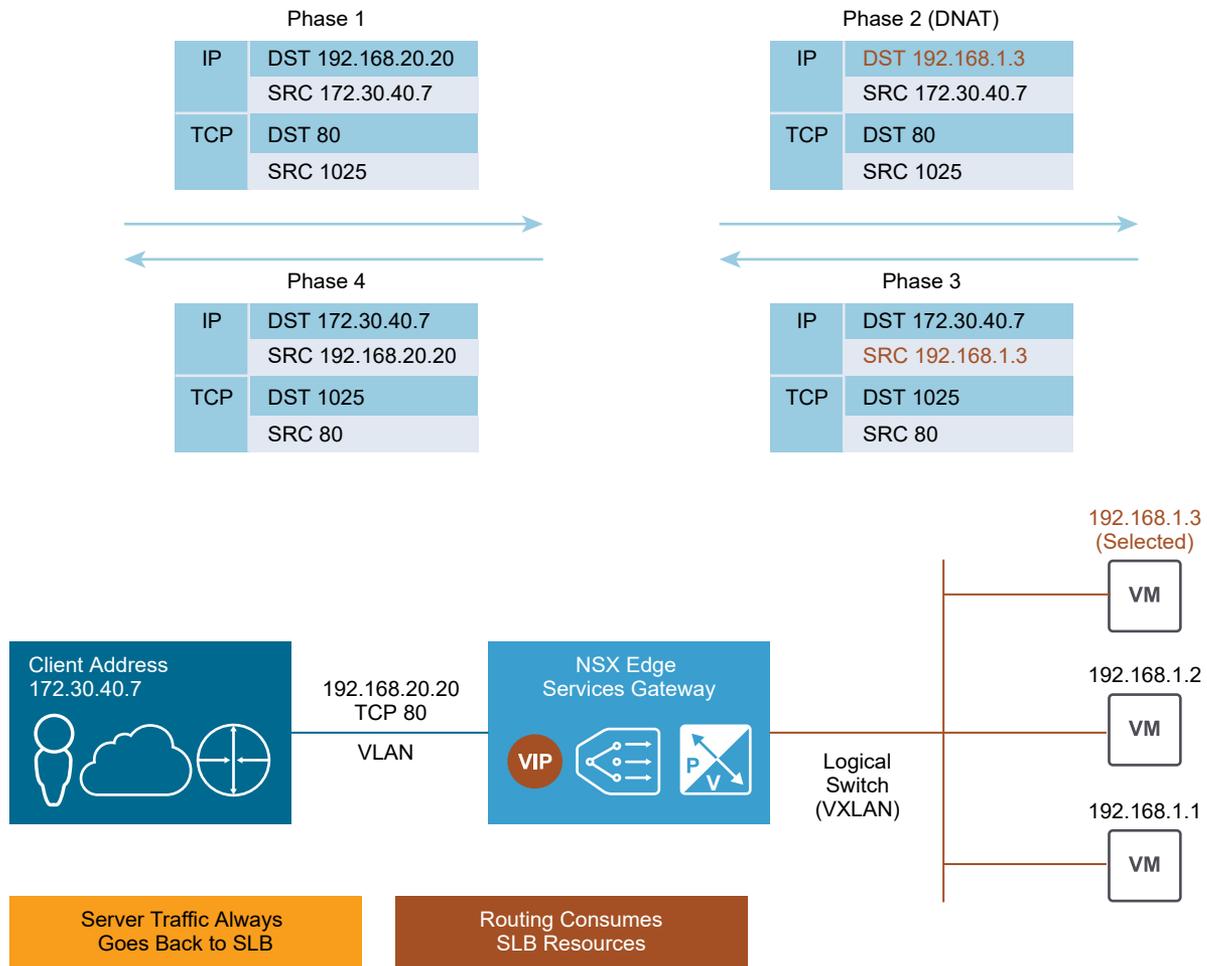
## Topologies

There are two types of load balancing services to configure in NSX, a one-armed mode, also known as a proxy mode, or the Inline mode, otherwise known as the transparent mode.

### NSX Logical Load Balancing: Inline Topology

Inline or Transparent mode deploys the NSX edge inline to the traffic destined to the server farm. Transparent mode traffic flow is processed as follows:

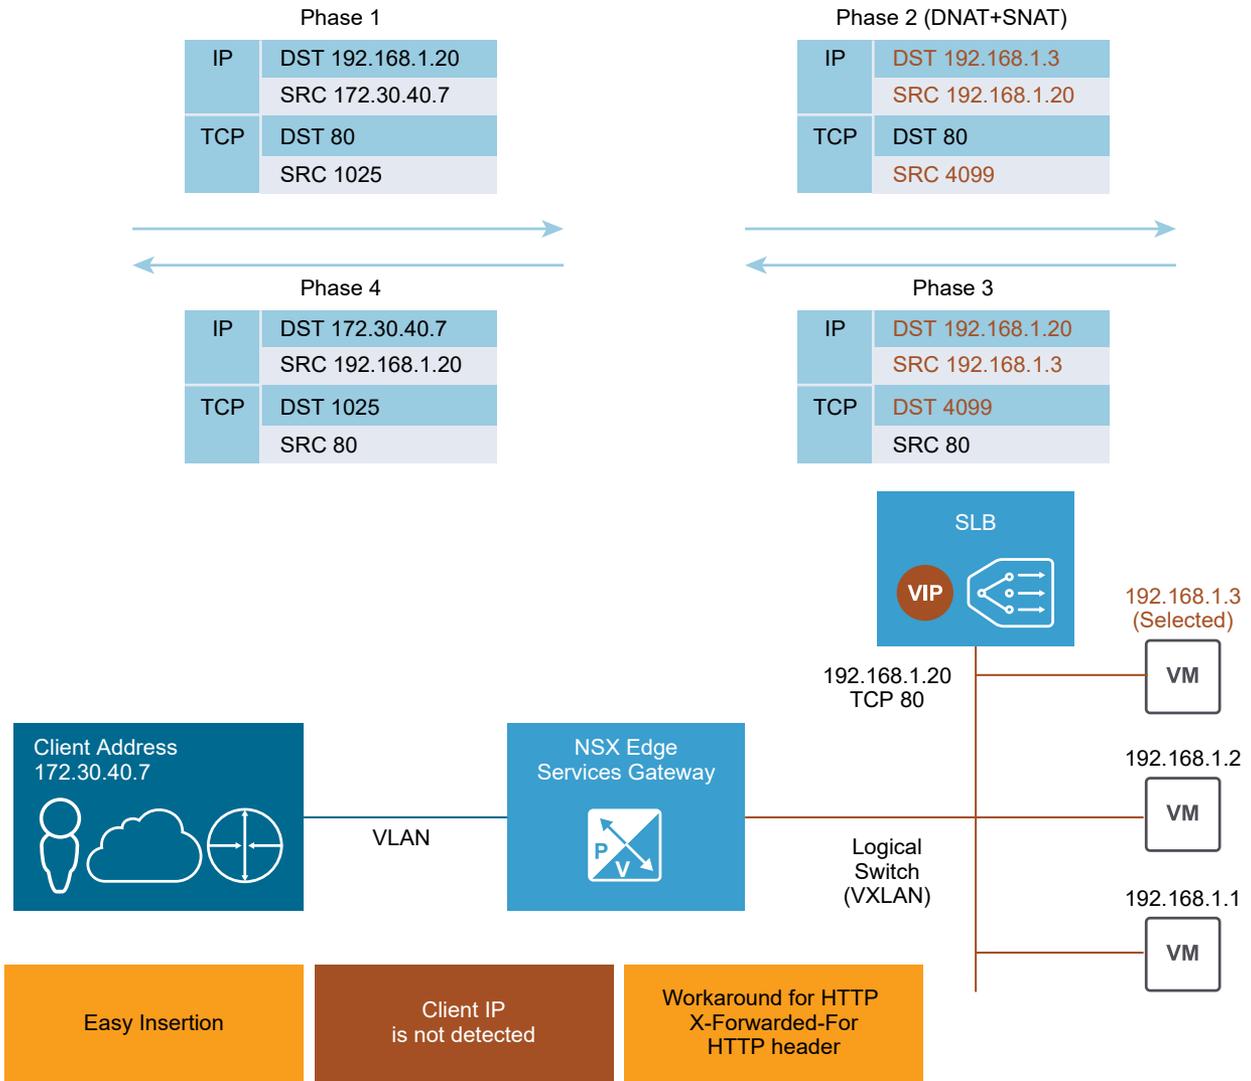
- The external client sends traffic to the virtual IP address (VIP) exposed by the load balancer.
- The load balancer – a centralized NSX edge – performs only destination NAT (DNAT) to replace the VIP with the IP address of one of the servers deployed in the server farm.
- The server in the server farm replies to the original client IP address. The traffic is received again by the load balancer since it is deployed inline, usually as the default gateway for the server farm.
- The load balancer performs source NAT to send traffic to the external client, leveraging its VIP as source IP address.



## NSX Logical Load Balancing: One-Armed Topology

One-Armed or Proxy mode consists of deploying an NSX edge directly connected to the logical network where load-balancing services are required.

- The external client sends traffic to the Virtual IP address (VIP) exposed by the load balancer.
- The load balancer performs two address translations on the original packets received from the client: destination NAT (DNAT) to replace the VIP with the IP address of one of the servers deployed in the server farm, and source NAT (SNAT) to replace the client IP address with the IP address identifying the load balancer itself. SNAT is required to force through the load balancer the return traffic from the server farm to the client.
- The server in the server farm replies by sending the traffic to the load balancer per SNAT functionality.
- The load balancer again performs a source and destination NAT service to send traffic to the external client, leveraging its VIP as source IP address.



This chapter includes the following topics:

- [Setting Up Load Balancing](#)
- [Managing Application Profiles](#)
- [Managing Service Monitors](#)
- [Managing Server Pools](#)
- [Managing Virtual Servers](#)
- [Managing Application Rules](#)
- [Load Balance Web Servers using NTLM Authentication](#)
- [Load Balancer HTTP Connection Modes](#)
- [Scenarios for NSX Load Balancer Configuration](#)

## Setting Up Load Balancing

The NSX Edge load balancer distributes network traffic across multiple servers to achieve optimal resource use, provide redundancy, and distribute resource utilization.

NSX load balancer supports layer 4 and layer 7 load balancing engines. The layer 4 load balancer is connection-based, providing fast path processing and the layer 7 load balancer is HTTP socket-based, allowing for advanced traffic manipulations and DDOS mitigation for back-end services.

Connection-based load balancing is implemented on the TCP and UDP layer. Connection-based load balancing does not stop the connection or buffer the whole request, it sends the packet directly to the selected server after manipulating the packet. TCP and UDP sessions are maintained in the load balancer so that packets for a single session are directed to the same server. Connection-based load balancing is done through Acceleration Disabled TCP and UDP virtual IP, or Acceleration Enabled TCP virtual IP.

Socket-based load balancing is implemented on top of the socket interface. Two connections are established for a single request, a client-facing connection and a server-facing connection. The server-facing connection is established after server selection. For HTTP socket-based implementation, the whole request is received before sending to the selected server with optional L7 manipulation. For HTTPS socket-based implementation, authentication information is exchanged either on the client-facing connection or on the server-facing connection. Socket-based load balancing is the default mode for TCP, HTTP, and HTTPS virtual servers.

Key concepts of the NSX load balancer include:

### Virtual Server

Abstract of an application service, represented by a unique combination of IP, port, protocol and application profile such as TCP or UDP.

### Server Pool

Group of backend servers.

## Server Pool Member

Represents the backend server as member in a pool.

## Service Monitor

Defines how to probe the health status of a backend server.

## Application Profile

Represents the TCP, UDP, persistence, and certificate configuration for a given application.

You begin by setting global options for the load balancer, then create a server pool of backend server members, and associate a service monitor with the pool to manage and share the backend servers efficiently.

Next, you create an application profile to define the common application behavior in a load balancer such as client SSL, server SSL, x-forwarded-for, or persistence. Persistence sends subsequent requests with similar characteristic such as source IP or cookie are required to be dispatched to the same pool member, without running the load balancing algorithm. Application profiles can be reused across virtual servers.

You then create an optional application rule to configure application-specific settings for traffic manipulation such as matching a certain URL or hostname so that different requests can be handled by different pools. Next, you create a service monitor that is specific to your application, or use a previously created service monitor.

Optionally, you can create an application rule to support advanced functionality of L7 virtual servers. Some use cases for application rules include content switching, header manipulation, security rules, and DOS protection.

Finally, you create a virtual server that connects your server pool, application profile, and any potential application rules together.

When the virtual server receives a request, the load balancing algorithm considers pool member configuration and runtime status. The algorithm then calculates the appropriate pool to distribute the traffic comprising one or more members. The pool member configuration includes settings such as, weight, maximum connection, and condition status. The runtime status includes current connections, response time, and health check status information. The calculation methods can be round-robin, weighted round-robin, least connection, source IP hash, weighted least connections, URL, URI, or HTTP header.

Each pool is monitored by the associated service monitor. When the load balancer detects a problem with a pool member, it is marked as DOWN. Only UP server is selected when choosing a pool member from the server pool. If the server pool is not configured with a service monitor, all the pool members are considered as UP.

---

**Note** For load balancer troubleshooting information, refer to *NSX Troubleshooting Guide*.

---

- [Configure Load Balancer Service](#)

- [Create a Service Monitor](#)

You create a service monitor to define health check parameters for a particular type of network traffic. When you associate a service monitor with a pool, the pool members are monitored according to the service monitor parameters.

- [Add a Server Pool](#)

You can add a server pool to manage and share back-end servers flexibly and efficiently. A server pool manages load balancer distribution methods and has a service monitor attached to it for health check parameters.

- [Create an Application Profile](#)

Use application profiles to enhance your control over managing network traffic, and make traffic-management tasks easier and more efficient.

- [Add an Application Rule](#)

You can write virtual server-side application rules by using the HAProxy syntax to manipulate and manage application traffic.

- [Add Virtual Servers](#)

Add an NSX Edge internal or uplink interface as a virtual server.

## Configure Load Balancer Service

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > Load Balancer > Global Configuration**.
- 5 Next to **Load Balancer Global Configuration**, click **Edit**.

## 6 Specify the global load balancer configuration parameters.

Option	Description
<b>Load Balancer</b>	Allows the NSX Edge load balancer to distribute traffic to internal servers for load balancing.
<b>Acceleration</b>	<p>When disabled, all virtual IP addresses (VIPs) use the L7 LB engine.</p> <p>When enabled, the virtual IP uses the faster L4 LB engine or L7 LB engine (based on the VIP configuration).</p> <p>The L4 VIP ("acceleration enabled" in the VIP configuration and no L7 setting such as AppProfile with cookie persistence or SSL-Offload) is processed before the edge firewall, and no edge firewall rule is required to reach the VIP. However, if the VIP is using a pool in non-transparent mode, the edge firewall must be enabled (to allow the auto-created SNAT rule).</p> <p>The L7 HTTP/HTTPS VIPs ("acceleration disabled" or L7 setting such as AppProfile with cookie persistence or SSL-Offload) are processed after the edge firewall, and require an edge firewall allow rule to reach the VIP.</p> <p>Note: To validate which LB engine is used for each VIP by the NSX Load Balancer, on the NSX Edge CLI (ssh or console), run the following command: "show service loadbalancer virtual" and look for "LB PROTOCOL [L4 L7]"</p>
<b>Logging</b>	<p>The NSX Edge load balancer collects traffic logs.</p> <p>You can select the log level from the drop-down menu. The logs are exported to the configured syslog server. You can also use the <code>show log follow</code> command to list the load balancing logs.</p> <p>The debug and info options log end-user requests. Warning, error, and critical options do not log end-users requests. If the NSX Edge Control-level log is set as debug or info, the load balancer logs lb, vip, pool, and pool member statistics every minute.</p> <p>Note that running in debug or info consumes CPU and edge log partition space, and might have a slight impact on the maximum traffic management capability.</p>
<b>Enable Service Insertion</b>	<p>Allows the load balancer to work with third-party vendor services.</p> <p>If you have a third party vendor load balancer service deployed in your environment, see <a href="#">Using a Partner Load Balancer</a>.</p> <p><b>Attention</b> Starting in NSX 6.4.5, support for integrating third-party services is deprecated.</p>

## 7 Click **OK**.

## Create a Service Monitor

You create a service monitor to define health check parameters for a particular type of network traffic. When you associate a service monitor with a pool, the pool members are monitored according to the service monitor parameters.

Following types of monitors are supported: ICMP, TCP, UDP, HTTP, HTTPS, DNS, MSSQL, and LDAP.

**Procedure**

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > Load Balancer > Service Monitoring**.
- 5 Click **Add**.
- 6 Enter a **Name** for the service monitor.

Interval, Timeout, and Max Retries are common parameters for all types of health checks.

- 7 Enter the **Interval** in seconds in which a server is to be tested.  
The interval is the period in seconds that the monitor sends requests to the back-end server.
- 8 Enter the **Timeout** value. In each health check, the timeout value is the maximum time in seconds within which a response from the server must be received.
- 9 Enter the **Max Retries**. This value is the number of times the server is tested before it is declared DOWN.

For example, if **Interval** is set as 5 seconds, **Timeout** as 15 seconds, and **Max Retries** as 3, it means that the NSX load balancer probes the back-end server every 5 seconds. In each probe, if the expected response is received from server within 15 seconds, then the health check result is OK. If not, the result is CRITICAL. If the recent three health check results are all DOWN, the server is marked as DOWN.

- 10 From the **Type** drop-down menu, select how to send the health check request to the server. Monitor types that are supported are ICMP, TCP, UDP, HTTP, HTTPS, DNS, MSSQL, and LDAP. Three predefined monitors are embedded in the system: default\_tcp\_monitor, default\_http\_monitor, and default\_https\_monitor.
- 11 If you select **ICMP** as the monitor type, no other parameters are applicable. Leave other parameters empty.

- 12 If you select **TCP** as the monitor type, three more parameters are available: Send, Receive, and Extension.
- a **Send** (optional) - The string sent to the back-end server after a connection is established. The maximum permitted string length is 256 characters.
  - b **Receive** (optional) Enter the string to be matched. This string can be a header or in the body of the response. When the received string matches this definition, the server is considered UP.
  - c **Extension:** Enter advanced monitor parameters as *key=value* pairs in the Extension section.

A sample extension, warning=10, indicates that if a server does not respond within 10 seconds, the status is set as warning.

All extension items must be separated with a carriage return character.

**Table 16-1. Extensions for TCP Protocol**

Monitor Extension	Description
escape	Can use \n, \r, \t, or \ in send or quit string. Must come before send or quit option. Default: nothing added to send, \r\n added to end of quit.
all	All expect strings need to occur in server response. Default is any.
quit= <i>STRING</i>	String to send to server to initiate a clean close of the connection.
refuse=ok warn crit	Accept TCP refusals with states: ok, warn, or crit. Default is crit.
mismatch=ok warn crit	Accept expected string mismatches with states: ok, warn, or crit. Default is warn.
jail	Hide output from TCP socket.
maxbytes= <i>INTEGER</i>	Close connection once more than the specified number of bytes are received.
delay= <i>INTEGER</i>	Seconds to wait between sending string and polling for response.
certificate= <i>INTEGER[,INTEGER]</i>	Minimum number of days a certificate has to be valid. The first value is #days for warning and the second value is critical (if not specified - 0).
warning= <i>DOUBLE</i>	Response time in seconds to result in warning status.
critical= <i>DOUBLE</i>	Response time in seconds to result in critical status.

**13** If you select **HTTP** or **HTTPS** as the monitor type, perform the following steps:

- a **Expected** (optional) - Enter the string that the monitor expects to match in the status line of HTTP response in the Expected section. This is a comma-separated list.

For example, 200,301,302,401.

- b **Method** (optional) - Select the method to detect server status from the drop-down menu: GET, OPTIONS, or POST.

- c **URL** (optional) - Enter the URL to GET or POST ("/" by default).

- d If you select the POST method, enter the data to be sent in the **Bold** section.

- e Enter the string to be matched in the response content in the **Receive** section. This string can be a header or in the body of the response.

If the string in the Expected section is not matched, the monitor does not try to match the Receive content.

Example of JSON format: Validate response contains `{"Healthy":true}`:  
`receive={"Healthy":true}`

- f **Extension:** Enter advanced monitor parameters as *key=value* pairs in the Extension section.

A sample extension, `warning=10`, indicates that if a server does not respond within 10 seconds, the status is set as warning.

All extension items should be separated with a carriage return character.

**Note** For *eregi*, *regex*, and *ereg*, if the string contains `{ }` and `"`, then you must add a character `|` before parsing the string for JSON format. Example of JSON format: Validate response contains `{"Healthy":true}`: `eregi="{\"Healthy\":true}\"`.

**Table 16-2. Extensions for HTTP/HTTPS Protocol**

Monitor Extension	Description
<code>no-body</code>	Do not wait for document body: stop reading after headers. Note that this still does an HTTP GET or POST, not a HEAD.
<code>ssl-version=3</code>	Force SSL handshake using sslv3. sslv3 and tlsv1 are deactivated in the health check option by default.
<code>ssl-version=10</code>	Force SSL handshake using tls 1.0.
<code>ssl-version=11</code>	Force SSL handshake using tls 1.1.
<code>ssl-version=12</code>	Force SSL handshake using tls 1.2.
<code>max-age=SECONDS</code>	Warn if document is more than SECONDS old. The number can also be in the form 10m for minutes, 10h for hours, or 10d for days.
<code>content-type=STRING</code>	Specify Content-Type header media type in POST calls.
<code>linespan</code>	Allow regex to span newlines (must precede <code>-r</code> or <code>-R</code> ).
<code>regex=STRING</code> or <code>ereg=STRING</code>	Search page for regex STRING.
<code>eregi=STRING</code>	Search page for case-insensitive regex STRING. For example: <ul style="list-style-type: none"> <li>■ Validate response contains "OK1" or "OK2": <code>eregi="(OK1 OK2)"</code></li> <li>■ Validate response contains <code>{"Healthy":true}</code>: <code>eregi="{\"Healthy\":true}"</code></li> </ul>

Table 16-2. Extensions for HTTP/HTTPS Protocol (continued)

Monitor Extension	Description
invert-regex	Return CRITICAL if found, OK if not.
proxy-authorization= <i>AUTH_PAIR</i>	Username:password on proxy-servers with basic authentication.
useragent= <i>STRING</i>	String to be sent in HTTP header as <i>User Agent</i> .
header= <i>STRING</i>	Any other tags to be sent in HTTP header. Use multiple times for additional headers. For example: header="Host: app1.xyz.com"
onredirect=ok warning critical follow sticky stickyport	How to handle redirected pages. <i>sticky</i> is like follow but stick to the specified IP address. <i>stickyport</i> also ensures port stays the same.
pagesize= <i>INTEGER:INTEGER</i>	Minimum page size required (bytes) : Maximum page size required (bytes).
warning= <i>DOUBLE</i>	Response time in seconds to result in warning status.
critical= <i>DOUBLE</i>	Response time in seconds to result in critical status.
expect = <i>STRING</i>	Comma-delimited list of strings, at least one of them is expected in the first (status) line of the server response (default: HTTP/1. If specified skips all other status line logic (ex: 3xx, 4xx, 5xx processing)
string = <i>STRING</i>	String to expect in the content.
url = <i>PATH</i>	URL to GET or POST (default: /).
post = <i>STRING</i>	URL to encode http POST data.
method = <i>STRING</i>	Set HTTP method (for example, HEAD, OPTIONS, TRACE, PUT, DELETE).
timeout = <i>INTEGER</i>	Seconds before connection times out (default is 10 seconds).
header=Host: <i>host_name</i> -H <i>host_name</i>	<i>host_name</i> is a valid host name or an FQDN of the host.

Table 16-3. Extensions for HTTPS Protocol

Monitor Extension	Description
certificate= <i>INTEGER</i>	Minimum number of days a certificate has to be valid. Port defaults to 443. When this option is used the URL is not checked.
authorization= <i>AUTH_PAIR</i>	Username:password on sites with basic authentication.
ciphers='ECDHE-RSA-AES256-GCM-SHA384'	Display ciphers used in HTTPS health check.

14 If you select **UDP** as the monitor type, perform the following steps:

- a **Send** (required): Enter the string to be sent to the back-end server after a connection is established.
- b **Receive** (required): Enter the string expected to receive from back-end server. Only when the received string matches this definition, is the server is considered as UP.

---

**Note** No extension is supported by the UDP monitor.

---

15 If you select **DNS** as the monitor type, perform the following steps:

- a **Send** (required): Enter the string to be sent to back-end server after a connection is established.
- b **Receive**: Enter the string expected to receive from the back-end server. Only when the received string matches this definition, the server is considered as UP.
- c **Extension**: Enter advanced monitor parameters as *key=value* pairs in the Extension section.

A sample extension, warning=10, indicates that if a server does not respond within 10 seconds, the status is set as warning. This monitor type supports only TCP protocol.

All extension items must be separated with a carriage return character.

**Table 16-4. Extensions for DNS Protocol**

Monitor Extension	Description
querytype=TYPE	Optional: DNS record query type where TYPE = <i>A</i> , <i>AAAA</i> , <i>SRV</i> , <i>TXT</i> , <i>MX</i> , <i>CNAME</i> , <i>ANY</i> . <ul style="list-style-type: none"> <li>■ A=IPv4 host address</li> <li>■ AAAA=Ipv6 host address</li> <li>■ SRV= Service locator</li> <li>■ TXT=Text record</li> <li>■ MX=Mail Exchange for the domain record</li> <li>■ CNAME=The Canonical name of an alias record</li> </ul> The default query type is <i>A</i> .
expect-authority	Optional: Expect the DNS server to be authoritative for the lookup.
accept-cname	Optional: Accept <i>cname</i> responses as a valid result to a query. It is used with <i>querytype=CNAME</i> together. The default is to ignore the <i>cname</i> responses as part of the result.
warning=seconds	Optional: Returns a WARNING message if the time elapse exceeds the provided value. Default is set to off.
critical=seconds	Optional: Returns a CRITICAL alert message if the time elapse exceeds the provided value. Default is set to off.

**16** If you select **MSSQL** as the monitor type, perform the following steps:

- a **Send:** Enter the string to be run on the back-end server after a connection is established.
- b **Receive:** Enter the string expected to receive from the back-end server. Only when the received string matches this definition, the server is considered as UP.
- c **User Name, Password and Confirm password** (required): Enter the required user name, password, and confirm the entered password. As monitor is associated with a pool, you must set MSSQL servers in the pool with the same user name and password that is specified here.
- d **Extension:** Enter advanced monitor parameters as *key=value* pairs in the Extension section.

A sample extension, `warning=10`, indicates that if a server does not respond within 10 seconds, the status is set as warning.

All extension items must be separated with a carriage return character.

**Table 16-5. Extensions for MSSQL Protocol**

Monitor Extension	Description
<code>database=DBNAME</code>	Optional: Database name to connect to. This extension is required when the parameter <i>Send</i> or <i>storedproc</i> is used.
<code>storedproc=STOREPROC</code>	Optional: Stored procedure to run against the MSSQL server.

**17** If you select **LDAP** as the monitor type, perform the following steps:

- a **Password and Confirm password** (optional): Enter the required password and confirm the entered password.
- b **Extension:** Enter advanced monitor parameters as *key=value* pairs in the Extension section.

A sample extension, `warning=10`, indicates that if a server does not respond within 10 seconds, the status is set as warning.

All extension items must be separated with a carriage return character.

**Table 16-6. Extensions for LDAP Protocol**

Monitor Extension	Description
<code>attr='ATTR'</code>	Optional: LDAP attribute to search (default: <code>'(objectclass=*)'</code> ). You must use <i>attr</i> together with <i>crit-entires</i> range.
<code>base='cn=admin,dc=example,dc=com'</code>	Required: LDAP base (For example, <code>ou=my unit, o=my org, c=at</code> ).

Table 16-6. Extensions for LDAP Protocol (continued)

Monitor Extension	Description
ver2 or ver3	Optional: <ul style="list-style-type: none"> <li>■ ver2: Use LDAP protocol version 2.</li> <li>■ ver3: Use LDAP protocol version 3.</li> </ul> Default protocol version is <i>ver2</i> .
bind=BINDDN	Optional: LDAP bind distinguished name (DN) (if necessary). For more information, refer to <a href="https://www.ldap.com/the-ldap-bind-operation">https://www.ldap.com/the-ldap-bind-operation</a> .
crit=DOUBLE	Optional: Response time to result in CRITICAL status (seconds).
crit-entries=low:high	Optional: Number of found entries to result in critical status. If the number of found entries is out of range [low, high], then the health check result is CRITICAL.

18 Click **OK**.

#### What to do next

Associate a service monitor with a pool.

## Add a Server Pool

You can add a server pool to manage and share back-end servers flexibly and efficiently. A server pool manages load balancer distribution methods and has a service monitor attached to it for health check parameters.

#### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > Load Balancer > Pools**.
- 5 Click **Add**.
- 6 Type a name and description for the load balancer pool.

## 7 Select the algorithm balancing method for each enabled service.

Option	Description
IP-HASH	<p>Selects a server based on a hash of the source IP address and the total weight of all the running servers.</p> <p>Algorithm parameters are disabled for this option.</p>
LEASTCONN	<p>Distributes client requests to multiple servers based on the number of connections already on the server.</p> <p>New connections are sent to the server with the fewest connections.</p> <p>Algorithm parameters are disabled for this option.</p>
ROUND_ROBIN	<p>Each server is used in turn according to the weight assigned to it.</p> <p>This is the smoothest and fairest algorithm when the server's processing time remains equally distributed.</p> <p>Algorithm parameters are disabled for this option.</p>
URI	<p>The left part of the URI (before the question mark) is hashed and divided by the total weight of the running servers.</p> <p>The result designates which server receives the request. This ensures that a URI is always directed to the same server if no server goes up or down.</p> <p>The URI algorithm parameter has two options <code>uriLength=&lt;len&gt;</code> and <code>uriDepth=&lt;dep&gt;</code>. The length parameter range should be <math>1 \leq \text{len} &lt; 256</math>. The depth parameter range should be <math>1 \leq \text{dep} &lt; 10</math>.</p> <p>Length and depth parameters are followed by a positive integer number. These options can balance servers based on the beginning of the URI only. The length parameter indicates that the algorithm should only consider the defined characters at the beginning of the URI to compute the hash.</p> <p>The depth parameter indicates the maximum directory depth to be used to compute the hash. One level is counted for each slash in the request. If both parameters are specified, the evaluation stops when either is reached.</p>
HTTPHEADER	<p>HTTP header name is looked up in each HTTP request.</p> <p>The header name in parentheses is not case-sensitive, which is similar to the ACL 'hdr()' function. If the header is absent or does not contain any value, the round robin algorithm is applied.</p> <p>The HTTPHEADER algorithm parameter has one option <code>headerName=&lt;name&gt;</code>. For example, you can use <code>host</code> as the HTTPHEADER algorithm parameter.</p>
URL	<p>URL parameter specified in the argument is looked up in the query string of each HTTP GET request.</p> <p>If the parameter is followed by an equal sign = and a value, then the value is hashed and divided by the total weight of the running servers. The result designates which server receives the request. This process is used to track user identifiers in requests and ensure that a same user ID is always sent to the same server as long as no server goes up or down.</p> <p>If no value or parameter is found, then a round robin algorithm is applied.</p> <p>The URL algorithm parameter has one option <code>urlParam=&lt;url&gt;</code>.</p>

8 (Optional) Select an existing default or custom monitor from the **Monitors** drop-down menu.

9 (Optional) Select the type of IP traffic for the pool. Default is any IP traffic.

- 10 To make client IP addresses visible to the back-end servers, enable the **Transparent** option. For more details, see [Chapter 16 Logical Load Balancer](#).

If Transparent is not selected (default value), back-end servers see the traffic source IP address as a load balancer internal IP address. If Transparent is selected, the source IP address is the real client IP address and NSX Edge must be set as the default gateway to ensure that return packets go through the NSX Edge device.

- 11 Add members to the pool.

- a Click **Add**.
- b Enter the name and IP address of the server member or click **Select** to assign grouping objects.

---

**Note** VMware Tools must be installed on each VM, or an enabled IP discovery method (DHCP snooping or ARP snooping, or both) must be available when using grouping objects instead of IP addresses. For more details, see [IP Discovery for Virtual Machines](#).

---

The grouping objects can either be vCenter Server or NSX.

- c Select the member state as **Enable**, **Disable**, or **Drain**.
  - **Drain** - Forces the server to shut down gracefully for maintenance. Setting the pool member as "drain" removes the back-end server from load balancing, while allowing it to be used for exiting connections and new connections from clients with persistence to that server. The persistence methods that work with a drain state are source IP persistence, cookie insert, and cookie prefix.

---

**Note** Drain state cannot be enabled on an NSX Edge load balancer that has been configured with **Enable Acceleration**. See [Configure Load Balancer Service](#) for more information.

---

**Note** Enabling and disabling High Availability configuration on the NSX Edge can break the persistence and drain state with source IP persistence method.

---

- **Enable** - Removes the server from maintenance mode and brings it back into operation. The pool member state should either be **Drain** or **Disabled**.
- **Disable** - The server remains in maintenance mode.

---

**Note** You cannot change the pool member state from **Disabled** to **Drain**.

---

- d Enter the port where the member is to receive traffic, and the monitor port where the member is to receive health monitor pings.

Port value should be null if the related virtual server is configured with a port range.

- e In Weight, enter the proportion of traffic this member can handle .

- f Enter the maximum number of concurrent connections that the member can handle.  
If the incoming requests go higher than the maximum, they are queued and wait for a connection to be released.
- g Enter the minimum number of concurrent connections that a member must always accept.
- h Click **OK**.

12 Click **Add** or **OK**.

## Create an Application Profile

Use application profiles to enhance your control over managing network traffic, and make traffic-management tasks easier and more efficient.

You create an application profile to define the behavior of a particular type of network traffic. After configuring a profile, you associate the profile with a virtual server. The virtual server then processes traffic according to the values specified in the profile.

The following topics explain the steps to create the various application profile types.

- [Create a TCP or UDP Application Profile](#)

To create an application profile that balances either TCP or UDP traffic type, specify the name and the type of persistence in the profile.

- [Create an HTTP Application Profile](#)

To create an application profile that balances the HTTP traffic type, specify the name, HTTP redirect URL, and the type of persistence in the profile. Optionally, you can also choose to insert the X-forwarded-for-HTTP header.

- [Create an HTTPS Application Profile](#)

You can create an HTTPS application profile for three HTTPS traffic types: SSL passthrough, HTTPS offloading, and HTTPS end-to-end. The workflow for creating the application profile varies for each HTTPS traffic type.

## Create a TCP or UDP Application Profile

To create an application profile that balances either TCP or UDP traffic type, specify the name and the type of persistence in the profile.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > Load Balancer > Application Profiles**.
- 5 Click **Add**.

The **New Application Profile** window opens.

- 6 Specify the application profile parameters.
  - a Select the type of profile (TCP or UDP).
  - b Enter a name for the application profile.
  - c Select a type of persistence.

Persistence tracks and stores session data, such as the specific pool member that serviced a client request. With persistence, client requests are directed to the same pool member throughout the life of a session or during subsequent sessions.

Persistence	Description
Source IP	This persistence type tracks sessions based on the source IP address. When a client requests a connection to a virtual server that supports a source IP address persistence, the load balancer checks whether that client was previously connected. If yes, the load balancer returns the client to the same pool member.
MSRDP	This persistence type maintains persistent sessions between Windows clients and servers that are running the Microsoft Remote Desktop Protocol (RDP) service.  For example, you can enable the MSRDP persistence to create a load balancing pool that consists of members running Windows Server 2003 or Windows Server 2008. In this scenario, all members belong to a Windows cluster and participate in a Windows session directory.  This persistence type is available only for a TCP application profile.

- 7 Click **Add** or **OK**.

## Create an HTTP Application Profile

To create an application profile that balances the HTTP traffic type, specify the name, HTTP redirect URL, and the type of persistence in the profile. Optionally, you can also choose to insert the X-forwarded-for-HTTP header.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > Load Balancer > Application Profiles**.
- 5 Click **Add**.  
The **New Application Profile** window opens.
- 6 Specify the properties of the application profile.
  - a Select the type of profile as **HTTP**.
  - b Enter a name for the application profile.

- c Enter the URL to which you want to redirect the HTTP traffic.

For example, you can direct traffic from `http://myweb.com` to `https://myweb.com`.

- d Select a type of persistence.

Persistence tracks and stores session data, such as the specific pool member that serviced a client request. With persistence, client requests are directed to the same pool member throughout the life of a session or during subsequent sessions.

Persistence	Description
Source IP	This persistence type tracks sessions based on the source IP address. When a client requests a connection to a virtual server that supports a source IP address persistence, the load balancer checks whether that client was previously connected. If yes, the load balancer returns the client to the same pool member.
Cookie	This persistence type inserts a unique cookie to identify a session the first time a client accesses the site. The cookie is referred in subsequent requests to persist the connection to the appropriate server.

- 7 If you selected the Cookie persistence type, enter the cookie name and select the mode of inserting the cookie; else, proceed to the next step.

Mode	Description
Insert	NSX Edge sends a cookie. If the server sends one or more cookies, the client receives an extra cookie (the server cookies and the Edge cookie). If the server does not send a cookie, the client receives the Edge cookie.
Prefix	Select this mode when your client does not support more than one cookie. All browsers accept multiple cookies. If you have a proprietary application using a proprietary client that supports only one cookie, the Web server sends its cookie as usual. NSX Edge injects its cookie information as a prefix in the server cookie value. This added cookie information is removed when the Edge sends it to the server.
App Session	In this mode, the application does not support a new cookie added by the virtual server (insert), and nor does it support a modified cookie (prefix). The virtual server learns the cookie injected by the back-end server. When the client presents that cookie, the virtual server forwards the client request to the same back-end server. It is not possible to see the App Session persistence table for troubleshooting.

- 8 (Optional) To identify the originating IP address of a client connecting to a Web server through the load balancer, enable the **Insert X-Forwarded-For HTTP header** option.
- 9 Click **Add** or **OK**.

## Create an HTTPS Application Profile

You can create an HTTPS application profile for three HTTPS traffic types: SSL passthrough, HTTPS offloading, and HTTPS end-to-end. The workflow for creating the application profile varies for each HTTPS traffic type.

### Note

- Starting in NSX 6.4.5, the **Application Profile Type** drop-down menu contains separate options to create a profile for each of the three HTTPS traffic types.
- In NSX 6.4.4 and earlier, the **Type** drop-down menu contains only a single HTTPS option. To create a profile for each of the three HTTPS traffic types, you must specify appropriate profile parameters.
- NSX load balancer does not support proxy SSL passthrough.

Starting in NSX 6.4.5, the UI terminologies for a couple of HTTPS profile parameters have changed. The following table lists the changes.

NSX 6.4.4 and Earlier	NSX 6.4.5 and Later
Virtual Server Certificates	Client SSL
Pool Certificates	Server SSL

The following table describes the three HTTPS traffic types.

**Table 16-7. HTTPS Traffic Types**

HTTPS Traffic Type	Description
SSL Passthrough	Application rules related to SSL attributes are allowed without requiring an SSL termination on the load balancer. The traffic pattern is: Client -> HTTPS-> LB (SSL passthrough) -> HTTPS -> Server.
HTTPS Offloading	HTTP-based load balancing occurs. SSL ends on the load balancer and HTTP is used between the load balancer and the server pool. The traffic pattern is: Client -> HTTPS -> LB (end SSL) -> HTTP -> Server.
HTTPS End-to-End	HTTP-based load balancing occurs. SSL ends on the load balancer and HTTPS is used between the load balancer and the server pool. The traffic pattern is: Client -> HTTPS -> LB (end SSL) -> HTTPS -> Server.

The following table describes the persistence supported in HTTPS traffic types.

**Table 16-8. Supported Persistence Types**

Persistence	Description
Source IP	This persistence type tracks sessions based on the source IP address. When a client requests a connection to a virtual server that supports a source IP address persistence, the load balancer checks whether that client was previously connected. If yes, the load balancer returns the client to the same pool member.
SSL Session ID	This persistence type is available when you create a profile for the SSL passthrough traffic type. SSL Session ID persistence ensures that repeat connections from the same client are sent to the same server. Session ID persistence allows the use of SSL session resumption, which saves processing time for both the client and the server.
Cookie	This persistence type inserts a unique cookie to identify a session the first time a client accesses the site. The cookie is referred in subsequent requests to persist the connection to the appropriate server.

For the **Source IP** and **SSL Session ID** persistence types, you can enter the persistence expiration time in seconds. The default value of persistence is 300 seconds (five minutes).

**Remember** The persistence table is of a limited size. If the traffic is heavy, a large timeout value might lead to the persistence table filling up quickly. When the persistence table fills up, the oldest entry is deleted to accept the newest entry.

The load balancer persistence table maintains entries to record that client requests are directed to the same pool member.

- 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.
- After the timeout period has expired, new connection requests will be sent to a pool member 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 a period, even if the existing connections are still alive.

The following table lists the approved cipher suites that can be used to negotiate security settings during an SSL or TLS handshake.

**Table 16-9. Approved Cipher Suites**

Cipher Value	Cipher Name
DEFAULT	DEFAULT
ECDHE-RSA-AES128-GCM-SHA256	TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
ECDHE-RSA-AES256-GCM-SHA384	TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384

**Table 16-9. Approved Cipher Suites (continued)**

Cipher Value	Cipher Name
ECDHE-RSA-AES256-SHA	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA
ECDHE-ECDSA-AES256-SHA	TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA
ECDH-ECDSA-AES256-SHA	TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA
ECDH-RSA-AES256-SHA	TLS_ECDH_RSA_WITH_AES_256_CBC_SHA
AES256-SHA	TLS_RSA_WITH_AES_256_CBC_SHA
AES128-SHA	TLS_RSA_WITH_AES_128_CBC_SHA
DES-CBC3-SHA	TLS_RSA_WITH_3DES_EDE_CBC_SHA
ECDHE-RSA-AES128-SHA	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA
ECDHE-RSA-AES128-SHA256	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
ECDHE-RSA-AES256-SHA384	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384
AES128-SHA256	TLS_RSA_WITH_AES_128_CBC_SHA256
AES128-GCM-SHA256	TLS_RSA_WITH_AES_128_GCM_SHA256
AES256-SHA256	TLS_RSA_WITH_AES_256_CBC_SHA256
AES256-GCM-SHA384	TLS_RSA_WITH_AES_256_GCM_SHA384

The following procedure explains the steps to create an application profile for each of the three HTTPS traffic types.

#### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > Load Balancer > Application Profiles**.
- 5 Click **Add**.

The **New Application Profile** window opens.

## 6 Specify the application profile parameters.

Steps for **SSL Passthrough** traffic type.

NSX Version	Procedure
6.4.5 and later	<ol style="list-style-type: none"><li>In the <b>Application Profile Type</b> drop-down menu, select <b>SSL Passthrough</b>.</li><li>Enter the name of the profile.</li><li>Select the type of persistence.</li><li>Enter the persistence expiration time.</li><li>Click <b>Add</b>.</li></ol>
6.4.4 and earlier	<ol style="list-style-type: none"><li>Enter the name of the profile.</li><li>In the <b>Type</b> drop-down menu, select <b>HTTPS</b>.</li><li>Select the <b>Enable SSL Passthrough</b> check box.</li><li>Select the type of persistence.</li><li>Enter the persistence expiration time.</li><li>Click <b>OK</b>.</li></ol>

Steps for **HTTPS Offloading** traffic type.

NSX Version	Procedure
6.4.5 and later	<ol style="list-style-type: none"> <li>a In the <b>Application Profile Type</b> drop-down menu, select <b>HTTPS Offloading</b>.</li> <li>b Enter the name of the profile.</li> <li>c Enter the URL to which you want to redirect the HTTP traffic.</li> <li>d Select the type of persistence. <ul style="list-style-type: none"> <li>■ For Cookie persistence, enter the cookie name and select the mode of inserting the cookie. For a description about each cookie mode, see <a href="#">Create an HTTP Application Profile</a>.</li> <li>■ For Source IP persistence, enter the persistence expiration time.</li> </ul> </li> <li>e Optional: To identify the originating IP address of a client connecting to a Web server through the load balancer, enable the <b>Insert X-Forwarded-For HTTP header</b> option.</li> <li>f Click the <b>Client SSL</b> tab.</li> <li>g Select one or more cipher algorithms or cipher suite to be used during the SSL handshake. Make sure that the approved cipher suite contains DH key length more than or equal to 1024 bits.</li> <li>h Specify whether client authentication is to be ignored or required. If necessary, then the client must provide a certificate after the request or the handshake is canceled.</li> <li>i Select the required service certificate, CA certificate, and CRL that the profile must use to end the HTTPS traffic from the client on the load balancer.</li> <li>j Click <b>Add</b>.</li> </ol>
6.4.4 and earlier	<ol style="list-style-type: none"> <li>a Enter the name of the profile.</li> <li>b In the <b>Type</b> drop-down menu, select <b>HTTPS</b>.</li> <li>c Follow the steps given in this table for NSX 6.4.5 and later to define the following application profile parameters: <ul style="list-style-type: none"> <li>■ HTTP Redirect URL</li> <li>■ Persistence</li> <li>■ Insert X-Forwarded-For HTTP header</li> </ul> </li> <li>d Click the <b>Virtual Server Certificates</b> tab.</li> <li>e Follow the steps given in this table for NSX 6.4.5 and later to define cipher algorithms and client authentication.</li> <li>f Click <b>Configure Service Certificates</b>, and select the required service certificate, CA certificate, and CRL that the profile must use to end the HTTPS traffic from the client on the load balancer.</li> <li>g Click <b>OK</b>.</li> </ol>

Steps for **HTTPS End-to-End** traffic type.

In this application profile type, you specify both the Client SSL (Virtual Server Certificates) parameters and the Server SSL (Pool Side SSL) parameters.

The Server SSL parameters are used to authenticate the load balancer from the server side. If the Edge load balancer has a CA certificate and a CRL already configured and the load balancer needs to verify a service certificate from the back-end servers, select the service certificate. You can also provide the load balancer certificate to the back-end server if the back-end server needs to verify the load balancer service certificate.

NSX Version	Procedure
6.4.5 and later	<ol style="list-style-type: none"> <li>a In the <b>Application Profile Type</b> drop-down menu, select <b>HTTPS End-to-End</b>.</li> <li>b Enter the name of the profile.</li> <li>c Follow the steps given in the table for creating an HTTPS offloading profile to define the following application profile parameters: <ul style="list-style-type: none"> <li>■ HTTP Redirect URL</li> <li>■ Persistence</li> <li>■ Insert X-Forwarded-For HTTP header</li> <li>■ Client SSL: cipher algorithms, client authentication, service certificate, CA certificate, and CRL</li> </ul> </li> <li>d Click the <b>Server SSL</b> tab, and select the cipher algorithms, the required service certificate, CA certificate, and CRL to authenticate the load balancer from the server side.</li> <li>e Click <b>Add</b>.</li> </ol>
6.4.4 and earlier	<ol style="list-style-type: none"> <li>a Enter the name of the profile.</li> <li>b In the <b>Type</b> drop-down menu, select <b>HTTPS</b>.</li> <li>c Follow the steps given in the table for creating an HTTPS offloading profile to define the following application profile parameters: <ul style="list-style-type: none"> <li>■ HTTP Redirect URL</li> <li>■ Persistence</li> <li>■ Insert X-Forwarded-For HTTP header</li> <li>■ Virtual Server Certificates: cipher algorithms, client authentication, service certificate, CA certificate, and CRL</li> </ul> </li> <li>d Select the <b>Enable Pool Side SSL</b> check box to enable the HTTPS communication between the load balancer and the back-end servers.</li> <li>e Select the cipher algorithms and the required service certificate, CA certificate, and CRL to authenticate the load balancer from the server side.</li> <li>f Click <b>OK</b>.</li> </ol>

## Add an Application Rule

You can write virtual server-side application rules by using the HAProxy syntax to manipulate and manage application traffic.

NSX Data Center supports only virtual server-side application rules. NSX load balancer internally uses HAProxy. It means that application rules that you add on the virtual server are internally inserted in that virtual server's "frontend" section of the HAProxy configuration file. Pool-side application rules (HAProxy "backend" section) are not supported.

For information about the application rule syntax, see the HAProxy documentation at <http://cbonte.github.io/haproxy-dconv/>

For examples of commonly used application rules, see [Application Rule Examples](#).

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > Load Balancer > Application Rules**.
- 5 Click **Add**.
- 6 Enter the name and script for the rule.
- 7 Click **Add** or **OK**.

## Application Rule Examples

Commonly used application rules.

### HTTP/HTTPS Redirection Based on Condition

An application profile allows you to specify HTTP/HTTPS redirection, which always redirects traffic regardless of the request URLs. You also have the flexibility to specify the conditions in which HTTP/HTTPS traffic should be redirected.

```
move the login URL only to HTTPS.
acl clear dst_port 80
acl secure dst_port 8080
acl login_page url_beg /login
acl logout url_beg /logout
acl uid_given url_reg /login?userid=[^&]+
acl cookie_set hdr_sub(cookie) SEEN=1
redirect prefix https://mysite.com set-cookie SEEN=1 if !cookie_set
redirect prefix https://mysite.com if login_page !secure
redirect prefix http://mysite.com drop-query if login_page !uid_given
redirect location http://mysite.com/ if !login_page secure
redirect location / clear-cookie USERID= if logout
```

### Routing by Domain Name

You can create an application rule to direct requests to a specific load balancer pool according to domain name. The following rule direct requests to foo.com to pool\_1, and requests to bar.com to pool\_2.

```
acl is_foo hdr_dom(host) -i foo
acl is_bar hdr_dom(host) -i bar
use_backend pool_1 if is_foo
use_backend pool_2 if is_bar
```

## Microsoft RDP Load Balancing and Protection

In the following sample scenario, the load balancer balances a new user to the less loaded server and resumes a broken session. The NSX Edge internal interface IP address for this scenario is 10.0.0.18, internal interface IP address is 192.168.1.1, and the virtual servers are 192.168.1.100, 192.168.1.101, and 192.168.1.102.

- 1 Create an application profile for TCP traffic with MSRDP persistence.
- 2 Create a TCP health monitor (tcp\_monitor).
- 3 Create a pool (named rdp-pool) with 192.168.1.100:3389, 192.168.1.101:3389 and 192.168.1.102:3389 as members.
- 4 Associate tcp\_monitor to rdp-pool.
- 5 Create the following application rule.

```
tcp-request content track-sc1 rdp_cookie(msthash) table rdp-pool
tcp-request content track-sc2 src table ipv4_ip_table

each single IP can have up to 2 connections on the VDI infrastructure
tcp-request content reject if { sc2_conn_cur ge 2 }

each single IP can try up to 5 connections in a single minute
tcp-request content reject if { sc2_conn_rate ge 10 }

Each user is supposed to get a single active connection at a time, block the second one
tcp-request content reject if { sc1_conn_cur ge 2 }

if a user tried to get connected at least 10 times over the last minute,
it could be a brute force
tcp-request content reject if { sc1_conn_rate ge 10 }
```

- 6 Create a virtual server (named rdp-vs).
- 7 Associate the application profile to this virtual server and add the application rule created in step 4.

The newly applied application rule on the virtual server protects the RDP servers.

## Advanced Logging

By default, NSX load balancer supports basic logging. You can create an application rule as follows to view more detailed logging messages for troubleshooting.

```
log the name of the virtual server
capture request header Host len 32

log the amount of data uploaded during a POST
capture request header Content-Length len 10
log the beginning of the referrer
capture request header Referer len 20

server name (useful for outgoing proxies only)
```

```
capture response header Server len 20

logging the content-length is useful with "option logasap"
capture response header Content-Length len 10

log the expected cache behaviour on the response
capture response header Cache-Control len 8

the Via header will report the next proxy's name
capture response header Via len 20

log the URL location during a redirection
capture response header Location len 20
```

After you associate the application rule to the virtual server, logs include detailed messages such as the following example.

```
2013-04-25T09:18:17+00:00 edge-187 loadbalancer[18498]: [org1]: 10.117.7.117 - - [25/Apr/
2013:09:18:16 +0000] "GET /favicon.ico HTTP/1.1" 404 1440 "" "" 51656 856 "vip-http-complete"
"pool-http-complete" "m2" 145 0 1 26 172 --NI 1 1 0 0 0 0 0 "" "" "10.117.35.187"
"Mozilla/5.0 (Windows NT 6.1; WOW64) AppleWebKit/537.31
(KHTML, like Gecko) Chrome/26.0.1410.64 Safari/537.31" "Apache/2.2.15 (Linux" ""

2013-04-25T09:18:17+00:00 edge-187 loadbalancer[18498]: [org1]: 10.117.7.117 - - [25/Apr/
2013:09:18:16 +0000] "GET /favicon.ico HTTP/1.1" 404 1440 "" "" 51657 856 "vip-http-complete"
"pool-http-complete" "m2" 412 0 0 2 414 --NI 0 0 0 0 0 0 0 "" "" "10.117.35.187" "Mozilla/5.0
(Windows NT 6.1; WOW64) AppleWebKit/537.31
(KHTML, like Gecko) Chrome/26.0.1410.64 Safari/537.31" "Apache/2.2.15 (Linux" ""
```

To troubleshoot the HTTPS traffic, you may need to add more rules. Most web application use 301/302 responses with a location header to redirect the client to a page (most of the time after a login or a POST call) and also require an application cookie. So your application server may have difficulty in getting to know client connection information and may not be able to provide the correct responses: it may even stop the application from working.

To allow the web application to support SSL offloading, add the following rule.

```
See clearly in the log if the application is setting up response for HTTP or HTTPS
capture response header Location len 32
capture response header Set-Cookie len 32

Provide client side connection info to application server over HTTP header
http-request set-header X-Forwarded-Proto https if { ssl_fc }
http-request set-header X-Forwarded-Proto http if !{ ssl_fc }
```

The load balancer inserts the following header when the connection is made over SSL.

```
X-Forwarded-Proto: https
```

The load balancer inserts the following header when the connection is made over HTTP.

```
X-Forwarded-Proto: http
```

## Block Specific URLs

You can block requests that contain specific keywords in the URL. The following sample rule checks if the request starts with /private or /finance and blocks the requests that have those terms.

```
Check if the request starts with "/private" or "/finance" (case insensitive)
acl block_url_list path_beg -i /private /finance

If the request is part of the list forbidden urls, reply "Forbidden" (HTTP response code 403)
http-request deny if block_url_list
```

## Authentication HTTP Redirect If No Cookies

You can redirect a client request that does not have a cookie to get an authentication. The following sample rule checks if the HTTP request is authentic and has cookies in the header. If the request does not have cookies then the rule redirects the request to /authent.php for authentication.

```
acl authentic_url url /authent.php
acl cookie_present hdr_sub(cookie) cookie1=
redirect prefix /authent.php if !authentic_url !cookie_present
```

## Default Page Redirect

You can redirect the client request / to a default page. The following sample rule checks if the HTTP request is / and redirects the request to a default login page.

```
acl default_url url /
redirect location /login.php if default_url
```

## Redirect to Maintenance Site

When the primary pool is down, you can use a maintenance server pool and redirect the URL to the maintenance Web site.

```
redirect location http://maintenance.xyz.com/maintenance.htm
```

## NT LAN Manager (NTLM) Authentication

By default on the server side NSX closes the TCP connection after each request. When you do not want to close the server session after each request, you can keep the server session alive and secure with the NTLM protocol.

```
no option http-server-close
```

By default on the client side NSX keeps the TCP connection established between requests. However with option "X-Forwarded-For", the session is closed after each request. The following option keeps the client connection open between requests even if XFF is configured.

```
no option httpclose
```

## Replace Server Header

You can delete the existing response server header and replace it with another server. The following sample rule deletes the server header and replaces it with the NGINX Web server that can act as a reverse proxy server for HTTP, HTTPS, SMTP, POP3, and IMAP protocols, HTTP cache, and a load balancer.

```
rspidel Server
rspadd Server:\ nginx
```

## Rewrite Redirect

You can rewrite the Location header from HTTP to HTTPS. The following sample rule identifies the Location header and replaces the HTTP with HTTPS.

```
rspirop ^Location:\ http://(.*) Location:\ https://\1
```

## Select Specific Pool Based on Host

You can redirect requests with a specific host to defined pools. The following sample rule checks for the request for specific hosts app1.xyz.com, app2.xyz.com, and host\_any\_app3 and redirects these requests respectively to defined pools, pool\_app1, or pool\_app2, and pool\_app3. All other requests are redirected to existing pools defined in the Virtual Server.

```
acl host_app1 hdr(Host) -i app1.xyz.com
acl host_app2 hdr(Host) -i app2.xyz.com
acl host_any_app3 hdr_beg(host) -i app3
```

Use a specific pool for each hostname.

```
use_backend pool_app1 if host_app1
use_backend pool_app2 if host_app2
use_backend pool_app3 if host_any_app3
```

## Select Specific Pool Based on URLs

You can redirect requests with URL keywords to specific pools. The following sample rule checks if the request starts with /private or /finance and redirects these requests to defined pools, pool\_private or pool\_finance. All other requests are redirected to existing pools defined in the Virtual Server.

```
acl site_private path_beg -i /private
acl site_finance path_beg -i /finance
use_backend pool_private if site_private
use_backend pool_finance if site_finance
```

## Redirect When Primary Pool is Down

If your servers in the primary pool are down, you can redirect users to use the servers in the secondary pool. The following sample rule checks in the pool\_production is down and transfers users to pool\_sorry\_server.

```
acl pool_production_down nbsrv(pool_production) eq 0
use_backend pool_sorry_server if pool_production_down
```

## Whitelist TCP Connection

You can block client IP addresses from accessing your server. The following sample rule blocks the defined IP address and resets the connection if the client IP address is not in the whitelist.

```
acl whitelist src 10.10.10.0 20.20.20.0
tcp-request connection reject if !whitelist
```

## Enable sslv3 and tlsv1

By default, sslv3 and tlsv1 are disabled service monitor extensions. You can enable them using the following application rule.

```
sslv3 enable
tlsv1 enable
```

## Configure Client Session Timeout

Session timeout is the maximum connection inactivity time on the client side. The inactivity timeout applies when the client is expected to acknowledge or send data. In the HTTP mode, this timeout is particularly important to consider during the first phase, when the client sends the request, and during the response while the client is reading the data sent by the server. The default timeout value is five minutes.

The following sample rule sets the timeout period to 100 seconds.

```
timeout client 100s
```

Time can be set as an integer with milliseconds, seconds, minutes, hour, or days.

## Redirect to HTTPS Site

You can redirects the clients coming on HTTP to the same page on HTTPS.

```
Redirect all HTTP requests to same URI but HTTPS redirect scheme
https if !{ ssl_fc }
```

Other option is as follows:

```
rspirep ^Location:\ http://(.*) Location:\ https://\1
```

## Redirect Non Authentic Clients

Redirect the client requests to "/authentic.php" if they do not have a cookie.

```
Check the HTTP request if request is "/authentic.php"
acl authentic_url url /authentic.php
Check the cookie "cookie1" is present
acl cookie_present hdr_sub(cookie) cookie1=
If the request is NOT "/authentic.php" and there is no cookie, then redirect to "/authentic.php"
redirect prefix /authentic.php if !authentic_url !cookie_present
```

## HTTP Response Header Rewrite

Replace the response server header "Server" with the value "nginx".

```
Delete the existing Response Server header "Server"
rspidel Server
Add the Response Server header "Server" with the value "nginx"
rspadd Server:\ nginx
```

## Sorry Server

In case of servers in the primary pool are all dead, use servers in the secondary pool.

```
detect if pool "pool_production" is still up
acl pool_production_down nbsrv(pool_production) eq 0
use pool "pool_sorry_server" if "pool_production" is dead
use_backend pool_sorry_server if pool_production_down
Option 1: # Redirect everything to maintenance site
redirect location http://maintenance.xyz.com/maintenance.htm
Option 2: #Use a specific maintenance server pool and rewrite all URLs to maintenance.php
acl match_all always_true
use_backend maint_pool if match_all
requirep ^GET\(.*)\HTTP/(.*) GET\ /maintenance.php\ HTTP/\2
```

## Add Virtual Servers

Add an NSX Edge internal or uplink interface as a virtual server.

### Prerequisites

- An application profile must be available.
- If you want to enable acceleration to use a faster load balancer, ensure that acceleration is enabled in the Global Configuration settings of the load balancer. See [Configure Load Balancer Service](#).

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.

4 Click **Manage > Load Balancer > Virtual Servers**.

5 Click **Add**.

The **New Virtual Server** window opens.

6 Specify the virtual server details.

- a Enable the virtual server to make this virtual server available for use.
- b (Optional) Enable acceleration for the load balancer to use the faster L4 load balancer engine rather than L7 load balancer engine.

---

**Note** This configuration requires firewall enabled on the edge.

---

If a virtual server configuration, such as application rules, HTTP type, or cookie persistence, is using the L7 load balancer engine, the L7 load balancer engine is used regardless of whether acceleration is enabled. Acceleration must be selected under Global Configuration.

You can use the **show service loadbalancer virtual** CLI command to confirm the load balancer engine in use.

c Select the application profile to be associated with the virtual server.

You can associate only an application profile with the same protocol as the virtual server that you are adding. The services supported by the selected pool appear.

d Enter a name and description for the virtual server.

e Enter an IP address or click **Select IP Address** to set the IP address that the load balancer is listening on.

The **Select IP Address** window shows only the primary IP address. If you are creating a VIP using a secondary IP address, enter it manually.

f Select the protocol that the virtual server handles.

g Enter the port number that the load balancer listens on.

You can also enter a range of ports. For example, to share the virtual server configuration, such as server pool, application profile, and application rule, enter **80, 8001-8004, 443**.

To use FTP, the TCP protocol must have port 21 assigned to it.

h Select the application rule.

i In the **Connection Limit** text box, enter the maximum concurrent connections that the virtual server can process.

j In the **Connection Rate Limit** text box, enter the maximum incoming new connection requests per second section.

- k (Optional) Click the **Advanced** tab and add the application rule to associate it with the virtual server.
- l Click **Add** or **OK**.

## Managing Application Profiles

After you create an application profile and associate it with a virtual server, you can update the existing profile or delete it to save system resources.

### Edit an Application Profile

You can edit an application profile.

#### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > Load Balancer > Application Profiles**.
- 5 Select a profile and click the **Edit** (  or  ) icon.
- 6 Make the appropriate changes to traffic, persistence, certificate, or cipher configuration, and click **Save** or **OK**.

### Configure SSL Termination for a Load Balancer

Without SSL termination configured, HTTP requests are not inspected. The load balancer sees the source and destination IP addresses and encrypted data. If you want to inspect the HTTP requests, you can terminate the SSL session on the load balancer and then create a new SSL session towards the cell pool.

#### Prerequisites

Go to **Manage > Settings > Certificates** and ensure that a valid certificate is present. You can add a certificate for the load balancer in any one of the following ways:

- Import a PEM encoded file.
- Generate a CSR.
- Create a self-signed certificate.

#### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.

- 4 Click **Manage > Load Balancer > Application Profiles**.
- 5 Click **Add**, and specify the application profile parameters.

NSX Version	Procedure
6.4.5 and later	<ol style="list-style-type: none"> <li>a In the <b>Application Profile Type</b> drop-down menu, select <b>HTTPS Offloading</b>.</li> <li>b In the <b>Persistence</b> drop-down menu, select <b>None</b>.</li> <li>c Click <b>Client SSL &gt; Service Certificates</b>.</li> <li>d Select the service certificate that you added for the NSX Edge load balancer.</li> </ol>
6.4.4 and earlier	<ol style="list-style-type: none"> <li>a In the <b>Type</b> drop-down menu, select <b>HTTPS</b>.</li> <li>b Ensure that the <b>Enable SSL Passthrough</b> check box is not selected.</li> <li>c Go to <b>Virtual Server Certificates &gt; Service Certificates</b>, and click the <b>Configure Service Certificate</b> check box.</li> <li>d Select the service certificate that you added for the NSX Edge load balancer.</li> </ol>

- 6 Click **Add** or **OK**.

## Delete an Application Profile

You can delete an application profile.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > Load Balancer > Application Profiles**.
- 5 Select a profile and click the **Delete** (🗑 or ✖) icon.

## Managing Service Monitors

A service monitor defines health check parameters for the load balancer.

If you are using a load balancer service monitor with high availability (HA), HA must be enabled on a dedicated interface.

After you create a service monitor and associate it with a server pool, you can update the existing service monitor or delete it to save system resources.

For more information about service monitors, see [Create a Service Monitor](#).

## Edit a Service Monitor

You can edit a service monitor.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > Load Balancer > Service Monitoring**.
- 5 Select a service monitor and click the **Edit** ( or ) icon.
- 6 Make the appropriate changes and click **Save** or **OK**.

## Delete a Service Monitor

You can delete a service monitor.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > Load Balancer > Service Monitoring**.
- 5 Select a service monitor and click the **Delete** ( or ) icon.

## Managing Server Pools

After you add a server pool to manage load balancer distribution, you can update the existing pool or delete it to save system resources.

### Edit a Server Pool

You can edit a server pool.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > Load Balancer > Pools**.
- 5 Select the pool to edit.
- 6 Click the **Edit** ( or ) icon.
- 7 Make the appropriate changes and click **Save** or **OK**.

## Configure a Load Balancer to Use Transparent Mode

Transparent indicates whether the client IP addresses are visible to the back-end servers.

When you add a server pool, the transparent mode is disabled by default. When transparent is disabled, back-end servers see the traffic source IP as a Load balancer internal IP. When this mode is enabled, source IP is the real client IP and NSX Edge must be on the path of the server response. A typical design is to have the server default gateway be the NSX Edge. Transparent mode does not require SNAT.

For more information about inline or transparent mode, see [Chapter 16 Logical Load Balancer](#).

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > Load Balancer > Pools**.
- 5 Specify the following parameters for the pool:
  - Name
  - Description
  - Algorithm
  - Monitor
  - IP Filter

For information about the various algorithms, see [Add a Server Pool](#).

- 6 To enable the transport mode, click the toggle switch or select the **Transparent** check box.

## Delete a Server Pool

You can delete a server pool.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > Load Balancer > Pools**.
- 5 Select the pool to delete.
- 6 Click the **Delete** (🗑️ or ✖️) icon.

## Show Pool Status

You can view the latest health status of pool and associated pool members.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > Load Balancer > Pools**.
- 5 Click **Show Status** or **Show Pool Statistics**.

The Pool and Member Status window displays the status of all the server pools.

- 6 Select a pool from the Pool Status and Statistics table and view the status of all the members in that pool in the Member Status and Statistics table.

Pool status can be UP or DOWN. Pool is marked as DOWN when all the members in the pool are DOWN; else the pool is UP.

Member status can be one of the following:

- UP: Member is enabled, and the health status of the member is UP. Or monitor is not defined on the pool.
- DOWN: Member is enabled, and the health status of the member is DOWN.
- MAINT: Member is disabled.
- DRAIN: Member is in drain state.

---

**Tip** Starting in NSX 6.4.5, you can click the DOWN status in the Member Status and Statistics table to determine the cause for the member being down. However, in NSX 6.4.4 and earlier, you must run the `show service loadbalancer pool Edge` CLI command to determine the cause for the member being down.

---

## Managing Virtual Servers

After you add virtual servers, you can update the existing virtual server configuration or delete it.

### Edit a Virtual Server

You can edit a virtual server.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.

- 4 Click **Manage > Load Balancer > Virtual Servers**.
- 5 Select the virtual server to edit.
- 6 Click the **Edit** (  or  ) icon.
- 7 Make the appropriate changes and click **Save** or **Finish**.

## Delete a Virtual Server

You can delete a virtual server.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > Load Balancer > Virtual Servers**.
- 5 Select the virtual server to delete.
- 6 Click the **Delete** (  or  ) icon.

## Managing Application Rules

After you create application rules to configure application traffic, you can edit the existing rule or remove it.

### Edit an Application Rule

Use the HAProxy syntax to add or edit application rules for manipulating application traffic.

For information about the application rule syntax, see HAProxy documentation at <http://cbonte.github.io/haproxy-dconv/>

For examples of commonly used application rules, see [Application Rule Examples](#).

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > Load Balancer > Application Rules**.
- 5 Select a rule and click the **Edit** (  or  ) icon.
- 6 Make the appropriate changes and click **Save** or **OK**.

## Delete an Application Rule

You can delete an application rule.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > Load Balancer > Application Rules**.
- 5 Select an application rule and click the **Delete** (🗑 or ✖) icon.

## Load Balance Web Servers using NTLM Authentication

NSX Load Balancer and NTLM authentication require the server connection to be kept alive.

By default, NSX Load Balancer closes the server TCP connection after each client request, however, Windows NT LAN Manager (NTLM) authentication requires the same connection for the lifetime of the authenticated request, connections are kept alive for the duration of the requests.

To keep the server connection open between requests, add the following application rule on the Virtual IP load balancing the Web servers using NTLM authentication:

```
add # NTLM authentication and keep the server connection open between requests
no option http-server-close
```

## Load Balancer HTTP Connection Modes

In NSX Data Center for vSphere 6.1.5 and later, when you enable `x-forwarded-for`, the HTTP connection mode changes from passive close (`option httpclose`) to the default HTTP server-close (`option http-server-close`) mode. The server-close option keeps the client-facing connection open, while the server-facing connection is closed after receiving a response from the server. Previous to NSX Data Center for vSphere 6.1.5, the Load Balancer did not close the connection proactively, but inserted the `Connection:close` header in both directions to indicate to the client or server to close the connection. If an HTTP/HTTPS transaction fails on the Load Balancer after upgrading to NSX Data Center for vSphere 6.1.5 or later, add an application rule with the script `option httpclose` and associate it with the virtual server that is no longer working.

**HTTP Server Close (default)** - The server-facing connection is closed after the end of the response is received, and the client-facing connection remains open. HTTP Server Close provides latency on the client side (slow network) and the fastest session reuse on the server side to save server resources. It also permits non-keepalive capable servers to be served in keep-alive from a client perspective. This mode is suitable for most common use cases, especially for slow client-facing networks and fast server-facing networks.

HTTP Keep Alive - All requests and responses are processed and connections remain open but idle between responses and new requests. The advantages are reduced latency between transactions, and less processing power required on the server side. Memory requirements will increase to accommodate the number of active sessions, which will be higher because connections are no longer closed after each request. The client-facing idle timeout can be configured using the application rule `timeout http-keep-alive [time]`. By default the idle timeout is 1 second. This mode is mandatory when an application requires NTLM authentication.

HTTP Tunnel - Only the first request and response are processed and a tunnel is established between the client and the server, they can talk without further analysis of HTTP protocol. After the tunnel is established, the connection is persistent on both the client and server sides. To enable this mode, none of the following options should be set: `passive-close mode`, `server-close mode`, `force-close mode`.

HTTP tunnel mode impacts the following features, and applies to only the first request and response in a session:

- no logs are generated
- HTTP header parsing
- HTTP header manipulation
- cookie processing
- content switching
- insertion of `X-Forwarded-For` header

HTTP Passive Close - The same as tunnel mode, but with a `Connection: close` header added in both the client and server directions. Both ends close after the first request and response exchange. If option `httpclose` is set, the Load Balancer works in HTTP tunnel mode and checks if a `Connection: close` header is present in each direction. If the header is not present, a `Connection: close` header is added. Each end then actively closes the TCP connection after each transfer, resulting in a switch to the HTTP close mode. Any connection header other than `close` is removed. Applications that cannot process the second and subsequent requests properly, such as an inserted cookie by the Load Balancer then carried back by the following requests from client, can use tunnel mode or passive close mode.

Some HTTP servers do not necessarily close the connections when they receive the `Connection: close` set by option `httpclose`. If the client also does not close, the connection remains open until the timeout expires. This causes a high number of simultaneous connections on the servers, and shows high global session times in the logs. For this reason, they are not compatible with older HTTP 1.0 browsers. If this occurs, use the option `forceclose` which actively closes the request connection once the server responds. Option `forceclose` also releases the server connection earlier because it does not have to wait for the client to acknowledge it.

HTTP Force Close - Both the client and the server connections are actively closed by the Load Balancer after the end of a response. Some HTTP servers do not necessarily close the connections when they receive the `Connection: close` set by `option httpclose`. If the client also does not close, the connection remains open until the timeout expires. This causes a high number of simultaneous connections on the servers and shows high global session times in the logs. When this happens, `option forceclose` actively closes the outgoing server channel when the server has finished to respond and release some resources earlier than with `option httpclose`.

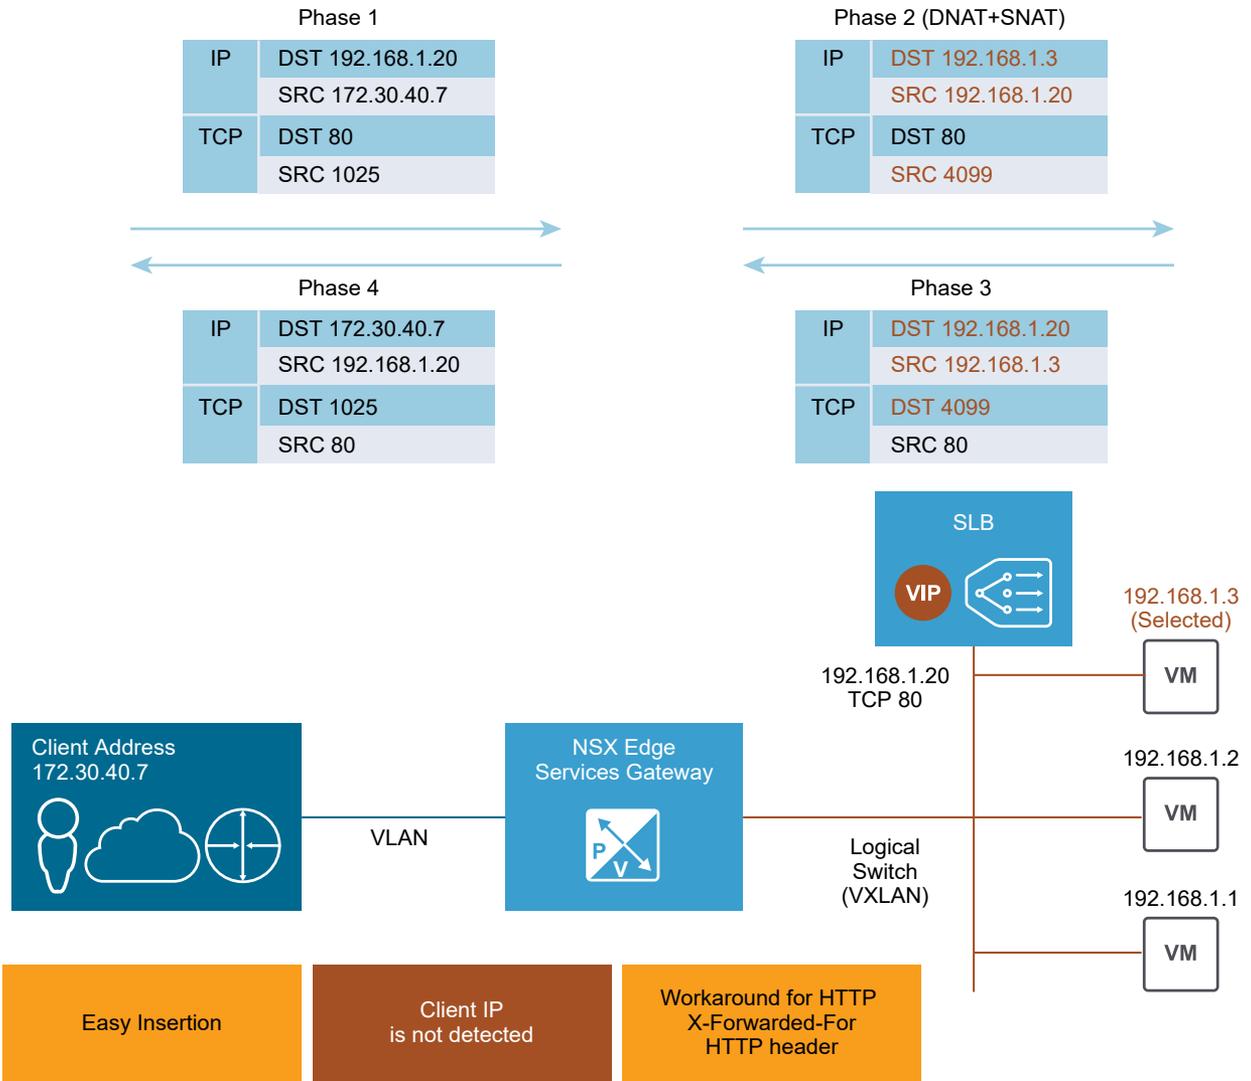
Version	Default Connection Mode	Connection Mode when	
		X-Forwarded-For is Enabled	Available Application Rules to Switch Connection Mode
6.0.x, 6.1.0, 6.1.1	HTTP Server Close	The <code>option httpclose</code> is automatically added to the virtual server to force <code>xff</code> to be added to each request as specified in the HAProxy document. The <code>xff</code> header is added into each request from the client when dispatching to a backend server.	No
6.1.2 - 6.1.4	HTTP Server Close	The HTTP Passive Close ( <code>option httpclose</code> is added automatically to the virtual server)	no <code>option http-server-close</code> <code>option httpclose</code> no <code>option httpclose</code>
6.1.5 - 6.1.x 6.2.0 - 6.2.2	HTTP Server Close	The HTTP Server Close <code>xff</code> header is added onto each request from the client when dispatching to the backend server.	no <code>option http-server-close</code> <code>option httpclose</code> no <code>option httpclose</code>
6.2.3-6.2.5	HTTP Server Close	The HTTP Server Close <code>xff</code> header is added onto each request from the client when dispatching to the backend server.	no <code>option http-server-close</code> <code>option httpclose</code> no <code>option httpclose</code>
6.2.3-6.2.5	HTTP Server Close	The HTTP Server Close <code>xff</code> header is added onto each request from the client when dispatching to the backend server.	no <code>option http-server-close</code> no <code>option httpclose</code> <code>option httpclose</code>
6.2.5 - 6.2.x	HTTP Server Close	The HTTP Server Close <code>xff</code> header is added onto each request from the client when dispatching to the backend server.	no <code>option http-server-close</code> <code>option http-keep-alive</code> <code>option http-tunnel</code> <code>option httpclose</code> <code>option forceclose</code>

## Scenarios for NSX Load Balancer Configuration

You can use the NSX load balancer configuration scenarios to get an understanding of the required end-to-end workflow.

### Scenario: Configure a One-Armed Load Balancer

The Edge Services Gateway (ESG) can be thought of as a proxy for the incoming client traffic.



In proxy mode, the load balancer uses its own IP address as the source address to send requests to a back-end server. The back-end server views all traffic as being sent from the load balancer and responds to the load balancer directly. This mode is also called SNAT mode or non-transparent mode. For more information, refer to *NSX Administration Guide*.

A typical NSX one-armed load balancer is deployed on the same subnet with its back-end servers, apart from the logical router. The NSX load balancer virtual server listens on a virtual IP for incoming requests from client and dispatches the requests to back-end servers. For the return traffic, reverse NAT is required to change the source IP address from the back-end server to a virtual IP (VIP) address and then send the virtual IP address to the client. Without this operation, the connection to the client can break.

After the ESG receives the traffic, it performs the following two operations:

- Destination Network Address Translation (DNAT) to change the VIP address to the IP address of one of the load balanced machines.
- Source Network Address Translation (SNAT) to exchange the client IP address with the ESG IP address.

Then the ESG server sends the traffic to the load balanced server and the load balanced server sends the response back to the ESG, and then back to the client. This option is much easier to configure than the inline mode, but has two potential caveats. The first is that this mode requires a dedicated ESG server, and the second is that the load balancer servers are not aware of the original client IP address. One workaround for HTTP or HTTPS applications is to enable the **Insert X-Forwarded-For** option in the HTTP application profile so that the client IP address is carried in the X-Forwarded-For HTTP header in the request that is sent to the back-end server.

If client IP address visibility is required on the back-end server for applications other than HTTP or HTTPS, you can configure the IP pool to be transparent. If clients are not on the same subnet as the back-end server, inline mode is recommended. Otherwise, you must use the load balancer IP address as the default gateway of the back-end server.

---

**Note** Usually, there are three methods to guarantee connection integrity:

- Inline/transparent mode
- SNAT/proxy/non-transparent mode (discussed above)
- Direct server return (DSR) - Currently, this is unsupported

In DSR mode, the back-end server responds directly to the client. Currently, NSX load balancer does not support DSR.

---

The following procedure explains the configuration of a one-armed load balancer with HTTPS offloading (SSL offloading) application profile type.

#### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.

- 3 Double-click an NSX Edge.
- 4 Click **Manage > Settings > Certificate**.  
For this scenario, add a self-signed certificate.
- 5 Enable the load balancer service.
  - a Click **Manage > Load Balancer > Global Configuration**.
  - b Click **Edit** and enable the load balancer.
- 6 Create an HTTPS application profile.
  - a Click **Manage > Load Balancer > Application Profiles**.
  - b Click **Add** and specify the application profile parameters.

Version	Procedure
NSX 6.4.5 and later	<ol style="list-style-type: none"> <li>1 In the <b>Application Profile Type</b> drop-down menu, select <b>HTTPS Offloading</b>.</li> <li>2 In the <b>Name</b> text box, enter the name of the profile. For example, enter <b>Web-SSL-Profile</b>.</li> <li>3 Click <b>Client SSL &gt; Service Certificates</b>.</li> <li>4 Select the self-signed certificate that you added earlier.</li> </ol>
NSX 6.4.4 and earlier	<ol style="list-style-type: none"> <li>1 In the <b>Type</b> drop-down menu, select <b>HTTPS</b>.</li> <li>2 In the <b>Name</b> text box, enter the name of the profile. For example, <b>Web-SSL-Profile</b>.</li> <li>3 Select the <b>Configure Service Certificate</b> check box.</li> <li>4 Select the self-signed certificate that you added earlier.</li> </ol>

- 7 (Optional) Click **Manage > Load Balancer > Service Monitoring**. Edit the default service monitoring to change it from basic HTTP or HTTPS to specific URL or URIs, as required.
- 8 Create a server pool.
  - a Click **Manage > Load Balancer > Pools**, and then click **Add**.
  - b In the **Name** text box, enter a name for the server pool. For example, enter **Web-Tier-Pool-01**.
  - c In the **Algorithm** drop-down menu, select **Round-Robin**.
  - d In the **Monitors** drop-down menu, select **default\_https\_monitor**.
  - e Add two members to the pool.

For example, specify the following configuration settings.

State	Name	IP Address	Weight	Monitor Port	Port	Max Connections	Min Connections
Enabled	web-01a	172.16.10.11	1	443	443	0	0
Enabled	web-02a	172.16.10.12	1	443	443	0	0

- f To use the SNAT mode, ensure that the **Transparent** option is not enabled.

- 9 Click **Show Status** or **Show Pool Statistics** and verify that the status of the Web-Tier-Pool-01 pool is UP.

Select the pool and ensure that the status of both members in this pool is UP.

- 10 Create a virtual server.

- a Click **Manage > Load Balancer > Virtual Servers**, and then click **Add**.
- b Specify the virtual server parameters.

For example, specify the following configuration settings.

Option	Description
<b>Virtual Server</b>	Enable the virtual server.
<b>Acceleration</b>	If you want to use the L4 load balancer for UDP or higher-performance TCP, enable acceleration. If you enable this option, ensure that the firewall status is enabled on the NSX Edge load balancer because a firewall is required for L4 SNAT.
<b>Application Profile</b>	Enter <b>OneArmWeb-01</b> .
<b>IP Address</b>	Select <b>172.16.10.110</b> .
<b>Protocol</b>	Select <b>HTTPS</b> .
<b>Port</b>	Enter <b>443</b> .
<b>Default Pool</b>	Select the <b>Web-Tier-Pool-01</b> server pool that you created earlier.
<b>Connection Limit</b>	Enter <b>0</b> .
<b>Connection Rate Limit</b>	Enter <b>0</b> .

- c (Optional) Click the **Advanced** tab, and associate an application rule with the virtual server.

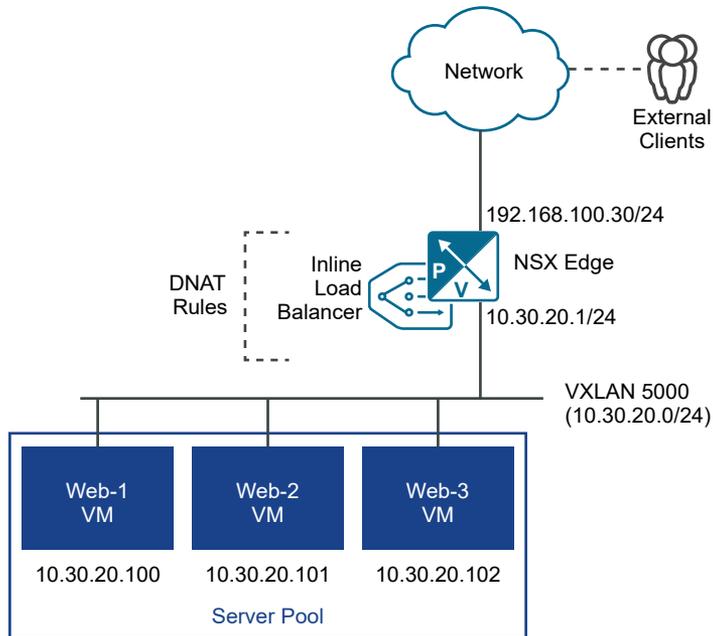
For supported examples, see: <https://communities.vmware.com/docs/DOC-31772>.

In non-transparent mode, the back-end server cannot see the client IP, but can see the load balancer internal IP address. As a workaround for HTTP or HTTPS traffic, select the **Insert X-Forwarded-For HTTP header** in the application profile. When this option is selected, the Edge load balancer adds the header "X-Forwarded-For" with the value of the client source IP address.

## Scenario: Configure an Inline Load Balancer

In this scenario, your goal is to configure an inline load balancer on the NSX Edge with an HTTP application profile type.

The following figure shows the logical topology of a network that uses an inline load balancer. The NSX Edge at the perimeter of the network does both north-south routing and the load balancing function.



For this scenario, consider that you have configured the following interfaces on the NSX Edge:

- Uplink interface: 192.168.100.30/24
- Internal interface: 10.30.20.1/24

The load balancer uses the uplink interface on the edge for the virtual IP address (VIP). The internal interface on the edge acts as the default gateway for the back-end web servers in the server pool.

You want to load balance the HTTP traffic coming from external clients on the NSX Edge and distribute the traffic to the Web servers that are connected to the VXLAN 5000 logical switch.

The following procedure explains the steps for configuring an inline load balancer on the NSX Edge.

### Prerequisites

You must have an NSX Edge Service Gateway deployed in your network.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Enable the load balancer service.
  - a Click **Manage > Load Balancer > Global Configuration**.
  - b Click **Edit** and enable the load balancer.

- 5 Create an HTTP application profile.
  - a Click **Manage > Load Balancer > Application Profiles**.
  - b Click **Add** and specify the application profile parameters.

For example:

Option	Description
Type	Select <b>HTTP</b> .
Name	Enter <b>Web-App-Profile</b> .
Persistence	Keep the default value (None).

- 6 Create a server pool.
  - a Click **Manage > Load Balancer > Pools**, and then click **Add**.
  - b Specify the pool parameters.

For example:

Option	Description
Name	Enter <b>Web-Server-Pool</b> .
Algorithm	Select <b>Round-Robin</b> .
Monitors	Select <b>default_http_monitor</b> .
Transparent	Enable this option to ensure that the source client IP addresses are visible to the back-end servers in the pool.

- c Add members to the server pool.

For example, specify the following settings for the pool members.

State	Name	IP Address	Weight	Monitor Port	Port	Max Connections	Min Connections
Enabled	Web-1	10.30.20.100	1	80	80	0	0
Enabled	Web-2	10.30.20.101	1	80	80	0	0
Enabled	Web-3	10.30.20.102	1	80	80	0	0

- 7 Click **Show Status** or **Show Pool Statistics** and verify that the status of the Web-Server-Pool is UP.

Select the pool and ensure that the status of all members in this pool is UP.

## 8 Create a virtual server.

- a Click **Manage > Load Balancer > Virtual Servers**, and then click **Add**.
- b Specify the virtual server parameters.

For example, specify the following configuration settings.

Option	Description
Virtual Server	Enable the virtual server.
Acceleration	Keep this option disabled.
Application Profile	Select the <b>Web-App-Profile</b> that you created earlier.
IP Address	Enter or select the IP address that you configured on the uplink (external) interface of the edge. For this scenario, select <b>192.168.100.30</b> .
Protocol	Select <b>HTTP</b> .
Default Pool	Select the <b>Web-Server-Pool</b> that you created earlier.
Connection Limit	Enter 0.
Connection rate Limit	Enter 0.

## Scenario: Configure NSX Load Balancer for Platform Service Controller

Platform Services Controller (PSC) provides infrastructure security functions, such as vCenter Single Sign-On, licensing, certificate management, and server reservation.

After configuring the NSX load balancer, provide the NSX Edge device uplink interface IP address for vCenter Single Sign-On.

---

**Note** The following procedure explains the steps for configuring an NSX Edge load balancer for use with Platform Services Controller 6.0. For configuring the Edge load balancer for use with Platform Services Controller 6.5, see the VMware knowledge base article at <https://kb.vmware.com/s/article/2147046>.

---

### Prerequisites

- Perform the PSC High Availability preparation tasks that are mentioned in the VMware knowledge base article at <http://kb.vmware.com/kb/2113315>.
- Save the `/ha/lb.crt` and `/ha/lb_rsa.key` from the first PSC node to configure certificates.
- Verify that an NSX Edge device is configured.
- Verify that you have at least one uplink for configuring VIP and one interface attached to an internal logical switch.

## Procedure

- 1 Add a PSC certificate to the NSX Edge.
  - a Save the PSC `root.cer` certificate, RSA, and passphrase that you generated with the OpenSSL command.
  - b Double-click the Edge and click **Manage > Settings > Certificates**.
  - c Click **Add > Certificate**.
  - d In the **Certificate Contents** text box, add the contents of the `root.cer` file.
  - e In the **Private key** text box, add the passphrase.
- 2 Enable the load balancer service.
  - a Click **Manage > Load Balancer > Global Configuration**.
  - b Click **Edit** and enable the load balancer.
- 3 Create application profiles with TCP and HTTPS protocols.
  - a Click **Manage > Load Balancer > Application Profiles**.
  - b Click **Add** and create a TCP application profile.

For example, specify the following parameters in the TCP profile.

Option	Description
Application Profile Type	Select <b>TCP</b> .
Name	For example, enter <code>sso_tcp_profile</code> .
Persistence	Select <b>Source IP</b>

- c Create an HTTPS application profile.

For example, specify the following parameters in the HTTPS profile.

Version	Procedure
NSX 6.4.5 and later	<ol style="list-style-type: none"> <li>1 In the <b>Application Profile Type</b> drop-down menu, select <b>HTTPS Offloading</b>.</li> <li>2 In the <b>Name</b> text box, enter the name of the profile. For example, enter <code>sso_https_profile</code>.</li> <li>3 Click <b>Client SSL &gt; Service Certificates</b>.</li> <li>4 Select the PSC certificate that you added earlier.</li> </ol>
NSX 6.4.4 and earlier	<ol style="list-style-type: none"> <li>1 In the <b>Type</b> drop-down menu, select <b>HTTPS</b>.</li> <li>2 In the <b>Name</b> text box, enter the name of the profile. For example, <code>sso_https_profile</code>.</li> <li>3 Select the <b>Configure Service Certificate</b> check box.</li> <li>4 Select the PSC certificate that you added earlier.</li> </ol>

- 4 Create server pools and add member PSC nodes.
  - a Click **Manage > Load Balancer > Pools**, and then click **Add**.
  - b Create a pool with the following configuration settings.

For example:

Option	Description
Name	Enter <code>sso_tcp_pool1</code> .
Algorithm	Select <b>Round-Robin</b> .
Monitors	Select <code>default_tcp_monitor</code> .

Add the following members to the `sso_tcp_pool1` pool with monitor port 443.

State	Name	IP Address	Weight	Monitor Port	Port	Max Connections	Min Connections
Enabled	PSC01	192.168.1.1	1	443		0	0
Enabled	PSC02	192.168.1.2	1	443		0	0

- c Create another pool with the following configuration settings.

For example:

Option	Description
Name	Enter <code>sso_tcp_pool2</code> .
Algorithm	Select <b>Round-Robin</b> .
Monitors	Select <code>default_tcp_monitor</code> .

Add the following members to the `sso_tcp_pool2` pool with monitor port 389.

State	Name	IP Address	Weight	Monitor Port	Port	Max Connections	Min Connections
Enabled	PSC01	192.168.1.1	1	389		0	0
Enabled	PSC02	192.168.1.2	1	389		0	0

- 5 Create virtual servers for the TCP and HTTPS protocols.
  - a Select **Manage > Load Balancer > Virtual Servers** , and then click **Add**.
  - b Create a virtual server for TCP VIP with the following configuration settings.

For example:

Option	Description
Virtual Server	Enable the virtual server.
Acceleration	Disable acceleration.
Application Profile	Enter <code>sso_tcp_profile</code> .
Name	Enter <code>sso_tcp_vip</code>
IP Address	Select <code>10.156.209.158</code> .
Protocol	Select <code>TCP</code> .
Port	Enter <code>389, 636, 2012, 2014, 2020</code> .
Default Pool	Select the <code>sso_tcp_pool2</code> server pool that you created earlier.
Connection Limit	Enter <code>0</code> .
Connection Rate Limit	Enter <code>0</code> .

- c Create a virtual server for HTTPS VIP with the following configuration settings.

For example:

Option	Description
Virtual Server	Enable the virtual server.
Acceleration	Disable acceleration.
Application Profile	Enter <code>sso_https_profile</code> .
Name	Enter <code>sso_https_vip</code>
IP Address	Select <code>10.156.209.158</code> .
Protocol	Select <code>HTTPS</code> .
Port	Enter <code>443</code> .
Default Pool	Select the <code>sso_tcp_pool1</code> server pool that you created earlier.
Connection Limit	Enter <code>0</code> .
Connection Rate Limit	Enter <code>0</code> .

## Scenario: SSL Offloading

This scenario uses an HTTPS offloading (SSL offloading) application profile type. Edge ends client HTTPS (SSL sessions). Edge load balances the clients on HTTP to the servers. L7 application rules can be applied.

## Procedure

- 1 Add a Web server certificate.
  - a Log in to the vSphere Web Client.
  - b Click **Networking & Security > NSX Edges**.
  - c Double-click an NSX Edge.
  - d Navigate to **Manage > Settings > Certificates**.
  - e Click **Add** and then click **Certificate**.

- f Copy and paste the certificate contents in the **Certificate Contents** text box. Text must include "-----BEGIN xxx-----" and "-----END xxx-----".

For chained certificates (server certificate and an intermediate CA certificate), select the **Certificate** option. Following is an example of a chained certificate content:

```
-----BEGIN CERTIFICATE-----
 Server cert
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
 Intermediate cert
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
 Root cert
-----END CERTIFICATE-----
```

- g In the **Private Key** text box, copy and paste the private key contents .

Following is an example of the private key content:

```
-----BEGIN RSA PRIVATE KEY-----
XX
-----END RSA PRIVATE KEY-----
```

Prefix the certificate content (PEM for certificate or private key) with one of the following strings:

```
-----BEGIN PUBLIC KEY-----
-----BEGIN RSA PUBLIC KEY-----
-----BEGIN CERTIFICATE REQUEST-----
-----BEGIN NEW CERTIFICATE REQUEST-----
-----BEGIN CERTIFICATE-----
-----BEGIN PKCS7-----
-----BEGIN X509 CERTIFICATE-----
-----BEGIN X509 CRL-----
-----BEGIN ATTRIBUTE CERTIFICATE-----
-----BEGIN RSA PRIVATE KEY-----
-----BEGIN DSA PRIVATE KEY-----
-----BEGIN EC PARAMETERS-----
-----BEGIN EC PRIVATE KEY-----
```

For complete examples of certificate and private key, see the [Example: Certificate and Private Key](#).

**Note** The following prefix is not supported in NSX Manager:

```
-----BEGIN ENCRYPTED PRIVATE KEY-----
```

- 2 Create an HTTPS application profile.
  - a Log in to the vSphere Web Client.
  - b Click **Networking & Security > NSX Edges**.
  - c Double-click an NSX Edge.
  - d Click **Manage > Load Balancer > Application Profiles**.
  - e Click **Add** and specify the application profile parameters.

Version	Procedure
NSX 6.4.5 and later	<ol style="list-style-type: none"> <li>1 In the <b>Application Profile Type</b> drop-down menu, select <b>HTTPS Offloading</b>.</li> <li>2 Click <b>Client SSL &gt; Service Certificates</b>.</li> <li>3 Select the web server certificate that you added in step 1.</li> </ol>
NSX 6.4.4 and earlier	<ol style="list-style-type: none"> <li>1 In the <b>Type</b> drop-down menu, select <b>HTTPS</b>.</li> <li>2 Select the <b>Configure Service Certificates</b> check box.</li> <li>3 Select the web server certificate that you added in step 1.</li> </ol>

- 3 Create a virtual server.
  - a Log in to the vSphere Web Client.
  - b Click **Networking & Security > NSX Edges**.
  - c Double-click an NSX Edge.
  - d Click **Manage > Load Balancer > Virtual Servers**.
  - e Click **Add** and specify the virtual server parameters.
    - 1 Enable the virtual server to make this virtual server available for use.
    - 2 Select the protocol as **HTTPS**.
    - 3 Select the default pool that is composed of HTTP servers (not HTTPS servers).
    - 4 Select the application profile that you created in step 2.

For information about specifying the other parameters in the **New Virtual Server** window, see [Add Virtual Servers](#).

## Example: Certificate and Private Key

Following are the examples of certificates and private key.

### Web Server Certificate

```
-----BEGIN CERTIFICATE-----
MIID0DCCArigAwIBAgIBATANBgkqhkiG9w0BAQUFADB/MQswCQYDVQQGEwJGUjET
MBEGA1UECAwKU29tZS1TdGF0ZTEOMAwwGA1UEBwwFUGFyaXMxDTALBgNVBAoM
BERpbWkxDTALBgNVBAsMBS5TQ1UxEDAQOBgNVBAMMB0RpbWkgQ0ExGzAZBgkqhkiG9w0B
CQEWDGRpbW1AZGltaS5mcjAeFw0xNDAxMjgyMDM2NTVaFw0yNDAxMjYyMDM2NTVa
MFsxCzAJBgNVBAYTAkZSMRMwEQYDVRQQIDApTb211LVN0YXR1MSEwHwYDVRQQDBhJ
```

```

bnRlcm5ldCBXaWRnaXRzIFB0eSBMdGQxFDASBgNVBAMMC3d3dy5kaWlpLmZyMIIIB
IjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEAvpnaPKLIKdvx98KW68Lz8pGa
RRcYersNGqPjpfMVjje8LuCoXgPU0HePnNTUjpShBnynKCvrtWhN+haKbSp+QWX
SxiTrW99HBfAl1MDQyWcukoEb9Cw6INctVUN4iRvkn9T8E6q174RbcnWA/7yTc7p
1NCvw+6B/aAN911G2pQXgRdYC/+G6o1IzEhtWhqze97nY5QKNUUVD0V09dc5CDYB
aKjgetwwv6DFk/GRdOSeD/6bW+20z0qSHpa3YNW6qSp+x5pyYmDrzRIR03os6Dau
ZkChSRyc/Whvurx6o85D6qpzywo8xwNaLZHxTQPgcIA5su9ZIyvtv9LH2E+1SwwID
AQABo3sweTAJBGNVHRMEAjAAMCwGCWCGSAGG+EIBDQQfFh1PcGVuU1NMIEdlBmVy
YXRlZCBZDZXJ0aWZpY2F0ZTADBgNVHQ4EFgQU+tugFtyN+cXe1wxUqeA7X+yS3bgw
HwYDVR0jBBGwFoAUhMwqkbBrGp87HxfvvgPnlGgVR64wDQYJKoZIhvcNAQEFBQAD
ggEBAIEEmqghEzeXZ4CKhE5UM9vCKzkj5Iv9TFs/a9CcQuepzplt7YVmevBFNOc0
+1ZyR4tXgi4+5MHGzhYCIvVHo4hKqYm+J+o5mwQInf1qoAHu07CLD3WNa1sKcVUV
vepIxc/1aHzrG+dPeEht0MdfFow13YdUc2FH6AgEdcEL4aV5PXq2eYR8hr4zKbc1
fBtuqUsvA8NWSIyzQ16fyGve+ANf6vXvUizyvwDrPrv/kfvLNa3ZPnLMMxU98Mvh
PXy3PkB8++6U4Y3vdK2Ni2WYY1Ils8yqbM4327IKmkDc2TimS8u60CT47mKU7aDY
cbTV5RDkrlaYwm5yq1TIg1vCv7o=
-----END CERTIFICATE-----

```

## Web Server Certificate with Chain (Including Root CA)

```

-----BEGIN CERTIFICATE-----
MIID0DCCArigAwIBAgIBATANBgkqhkiG9w0BAQUFADB/MQswCQYDVQQGEwJGUjET
MBEGA1UECAwKU29tZS1TdGF0ZTEOMAwwGA1UEBwwFUGFyaXNxDjEwMDALBgNVBAoM
BERp
bWkxDTALBgNVBAsMBE5TQ1UxEDA0BgNVBAMMB0RpbWkgQ00ExGzAZBgkqhkiG9w0B
CQEWDGRpbWlAZGltas5mcjAeFw0xNDAMjgyMDM2NTVaFw0yNDAMjYyMDM2NTVa
MFsxCzAJBgNVBAYTAkZSRMwEYDQVQIDApTb211LVN0YXRlMSEwHwYDVQQKDBhJ
bnRlcm5ldCBXaWRnaXRzIFB0eSBMdGQxFDASBgNVBAMMC3d3dy5kaWlpLmZyMIIIB
IjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEAvpnaPKLIKdvx98KW68Lz8pGa
RRcYersNGqPjpfMVjje8LuCoXgPU0HePnNTUjpShBnynKCvrtWhN+haKbSp+QWX
SxiTrW99HBfAl1MDQyWcukoEb9Cw6INctVUN4iRvkn9T8E6q174RbcnWA/7yTc7p
1NCvw+6B/aAN911G2pQXgRdYC/+G6o1IzEhtWhqze97nY5QKNUUVD0V09dc5CDYB
aKjgetwwv6DFk/GRdOSeD/6bW+20z0qSHpa3YNW6qSp+x5pyYmDrzRIR03os6Dau
ZkChSRyc/Whvurx6o85D6qpzywo8xwNaLZHxTQPgcIA5su9ZIyvtv9LH2E+1SwwID
AQABo3sweTAJBGNVHRMEAjAAMCwGCWCGSAGG+EIBDQQfFh1PcGVuU1NMIEdlBmVy
YXRlZCBZDZXJ0aWZpY2F0ZTADBgNVHQ4EFgQU+tugFtyN+cXe1wxUqeA7X+yS3bgw
HwYDVR0jBBGwFoAUhMwqkbBrGp87HxfvvgPnlGgVR64wDQYJKoZIhvcNAQEFBQAD
ggEBAIEEmqghEzeXZ4CKhE5UM9vCKzkj5Iv9TFs/a9CcQuepzplt7YVmevBFNOc0
+1ZyR4tXgi4+5MHGzhYCIvVHo4hKqYm+J+o5mwQInf1qoAHu07CLD3WNa1sKcVUV
vepIxc/1aHzrG+dPeEht0MdfFow13YdUc2FH6AgEdcEL4aV5PXq2eYR8hr4zKbc1
fBtuqUsvA8NWSIyzQ16fyGve+ANf6vXvUizyvwDrPrv/kfvLNa3ZPnLMMxU98Mvh
PXy3PkB8++6U4Y3vdK2Ni2WYY1Ils8yqbM4327IKmkDc2TimS8u60CT47mKU7aDY
cbTV5RDkrlaYwm5yq1TIg1vCv7o=
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
MIIDyTCCArGgAwIBAgIBADANBgkqhkiG9w0BAQUFADB/MQswCQYDVQQGEwJGUjET
MBEGA1UECAwKU29tZS1TdGF0ZTEOMAwwGA1UEBwwFUGFyaXNxDjEwMDALBgNVBAoM
BERp
bWkxDTALBgNVBAsMBE5TQ1UxEDA0BgNVBAMMB0RpbWkgQ00ExGzAZBgkqhkiG9w0B
CQEWDGRpbWlAZGltas5mcjAeFw0xNDAMjgyMDI2NDRAfW0yNDAMjYyMDI2NDRA
MH8xCzAJBgNVBAYTAkZSRMwEYDQVQIDApTb211LVN0YXRlMQ4wDAYDVQQHDAVQ
YXJpczENMAsGA1UECgWERGltaTENMAsGA1UECwwET1NCVTEQMA4GA1UEAwwHRGlta
aSBDbQEtbMkGCSqGSIb3DQEJARYMZGltatUBkaWlpLmZyMIIIBIjANBgkqhkiG9w0B
AQEFAAOCAQ8AMIIBCgKCAQEAuxuG4QeBIGXj/AB/YRLLtpgpTpGndntVlgsycZrL
3qqyOdBNlwnvcB9etfY5iWzjeq7YZRr6i0dIV4sFNBR2NoK+YvdD9j1TRi7njZg0
d6zth0x1sOhCsD1v/YCL1CTcYD1KA/QiKeIQa7GU3RhF0t/KnAkr6mwoDbdKBQX1
D5HGuXJiFdh5XRebXf1ZB3gH+0kCEaEzPrjFDAPkOXNxEARZdpBLpbvQ1jtvXtj

```

```

-----BEGIN CERTIFICATE-----
HMsvrIOc7QqUSOU3GcbBMSHjT8cgg8ssf492Go3bdQkIzTROz9QgDHaqDqTC9Hoe
v1IpTS+q/3BCY5AGWK13CCR6dDyK6honnOR/8srezaN4PwIDAQABO1AwTjAdBgNV
HQ4EFgQUhMwqkbBrGp87HxfvvgPnlGgVR64wHwYDVR0jBBgwFoAUhMwqkbBrGp87
HxfvvgPnlGgVR64wDAYDVR0TBAAUwAwEB/zANBgkqhkiG9w0BAQUFAAOCAQEAVqYq
vhm5wAEKmvKXRjeb5kiEIp7oZAFkYp6sKODuZ1VdkjMDD4wv46iqAe1QIIsfGwd
Dmv0oqS1+iPPy24ATMSZQbPLO5K64Hw7Q8KPos0yD8gHSg2d4SOUkj+FD2IjAH17
a8auMw7TThU6976JprQQKtPADRcfodGd5UFiz/6ZgLzUE23cktJMc2Bt18B9OZII
J9ef2PZxZirJg1OqF2KssDlJP5ECO9K3EmovC5M5Aly++s8ayjBnNivtklYL1VOT
ZrpPgcndTHUA5KS/Duf40dXm0snCxLAKNP28pMowDLSYc6IjVrD4+qqw3f1b7yGb
bJcFgxKDeg5YecQOSg==
-----END CERTIFICATE-----

```

## Private Key (No Passphrase)

```

-----BEGIN RSA PRIVATE KEY-----
MIIEowIBAAKCAQEAvpnaPKLIKdvx98KW681z8pGaRrcYersNGqPjpiFMVjjE8LuC
oXgPU0HePnNTUjPShBnynKCvrtWhN+haKbSp+QWXSxiTrW99HBfAl1MDQyWcukoE
b9Cw6INctVUN4iRvkn9T8E6q174RbcnwA/7yTc7p1NCvw+6B/aAN911G2pQXgRdY
C/+G6o1IZEHtWhqzE97nY5QKNuUVD0V09dc5CDYBaKjgetwvw6DFk/GRdOSeD/6b
W+20z0qSHPa3YNW6qSp+x5pyYmDrzRIR03os6DauZkChSRyc/Whvurx6o85D6qpz
ywo8xwNaLZHxTQPgcIA5su9ZiYtv9LH2E+1SwwIDAQABAoIBAFl8cD9a5pMq1W3
f9btTQz1sRL4Fvp7CmHSXhvjSjeHwhHckEe0ObkWTRsgkTsm1XLu5W8IITnhn0+1
iNr+78eB+rRGngdAXh8diOdkEy+8/Cee8tFI3jyutKdRlxMbwiKsouVviumoq3fx
OGQYwQ0Z21/PvCwy/Y82ffq3ysC5gAJsBbYsCrg14bQo44ulrELe4SDWs5HCjKYb
EI2b8cOMucqZSotxg9niLN/je2bo/I2HGSawibgcOdBms8k6TvsSrZMr3kJ506J+
77LGwKH37brVgbVYvbq6nWPL0xLG7dUv+7LWEo5qQaPy6aXb/zbcckqLqu6/EjOve
ydG5JQECgYEA9kKfTzd/WEVAreA0dzfeJRu8v1nwoagL7cJaoDxqXos4mcr5mPDT
kbWgFkLFFH/AyUnPB1K6BcJp1XK67B13ETUa3i9Q5t1WuZEobiKKBFLm9DDQJt43
uKZWJxBKFGSvFrYPTgzst719mZVcPct2CzPjEgN3H1pt6fyw3eOrnoECgYEAxiOu
jwXCOMuGaB7+OW2tR0PGEzbvVlEGdKAJ6TC/HoKM1A8r2u4hLTEJJCrLLTfw++4I
ddHE2dLeR4Q7058SfLphwgPmLDezN7WRLGr7Vyfuv7VmaHjGuC3Gv9agnhWD1A2Q
gBG9/R9oVfL0Dc7CgJgLeUtItCYC31bGT3yhV0MCgYEA4k3DG4L+RN4PXDpHvK9I
pA1jXAJHEifeHnaW1d3vWkbSkvJmgVf+9U5VeV+OwRHN1qzPZV4suRI6M/81K8rA
Gr4UnM4aqK4K/qkY4G05LKrik9Ev2CgqSLQDRA7CJQ+Jn3Nb50qg6hFnFPafN+J7
7juWln08wFYV4Atpdd+9XQECgYBxizkZFL+9IqkfOcONvWazGo+Dq1N0L3J4iTIk
w56CKWxyj88d4qB4eUU3yJ4uB4S9miaW/eLEwKZibWpUPFAn0db7i6h3ZmP5ZL8Q
qS3nQCb9DULmU2/tU641eRUKAmIoka1g9sndKAZuWo+o6fdkIb1RgObk9XNn8R4r
psv+aQKBgB+CIExR30vycv5bnZN9EF1IXNKaEMJUrYCXcRQNvrnUIUBvAO8+jAe
CdLygS5RtgOLZib0IVERqWsp3EI1ACGuLts0vQ9GFLQGaN1SaMS40C9kvns1mlDu
LhIhYpJ8UsCVt5snWo2N+M+6ANh5tpWdQnEK6zILh4tRbuzaiHgb
-----END RSA PRIVATE KEY-----

```

## Scenario: Import SSL Certificate

This scenario uses an HTTPS end-to-end (SSL end-to-end) application profile type. The NSX Edge closes client HTTPS (SSL sessions). Edge load balances the client on a new HTTPS connection to the servers. L7 application rules can be applied.

## Procedure

- 1 Add a Web server certificate.
  - a Log in to the vSphere Web Client.
  - b Click **Networking & Security > NSX Edges**.
  - c Double-click an NSX Edge.
  - d Navigate to **Manage > Settings > Certificates**.
  - e Click **Add**, and then click **Certificate**.

- f Copy and paste the certificate contents in the **Certificate Contents** text box. Text must include "-----BEGIN xxx-----" and "-----END xxx-----".

For chained certificates (server certificate and an intermediate CA certificate), select the **Certificate** option. Following is an example of a chained certificate content:

```
-----BEGIN CERTIFICATE-----
 Server cert
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
 Intermediate cert
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
 Root cert
-----END CERTIFICATE-----
```

- g In the **Private Key** text box, copy and paste the private key contents.

Following is an example of private key content:

```
-----BEGIN RSA PRIVATE KEY-----
XX
-----END RSA PRIVATE KEY-----
```

Prefix the certificate content (PEM for certificate or private key) with one of the following strings:

```
-----BEGIN PUBLIC KEY-----
-----BEGIN RSA PUBLIC KEY-----
-----BEGIN CERTIFICATE REQUEST-----
-----BEGIN NEW CERTIFICATE REQUEST-----
-----BEGIN CERTIFICATE-----
-----BEGIN PKCS7-----
-----BEGIN X509 CERTIFICATE-----
-----BEGIN X509 CRL-----
-----BEGIN ATTRIBUTE CERTIFICATE-----
-----BEGIN RSA PRIVATE KEY-----
-----BEGIN DSA PRIVATE KEY-----
-----BEGIN EC PARAMETERS-----
-----BEGIN EC PRIVATE KEY-----
```

For complete examples of certificates and private keys, see the [Example: Certificate and Private Key](#).

**Note** The following prefix is not supported on the NSX Manager:

```
-----BEGIN ENCRYPTED PRIVATE KEY-----
```

- 2 Create an HTTPS application profile.
  - a Log in to the vSphere Web Client.
  - b Click **Networking & Security > NSX Edges**.
  - c Double-click an NSX Edge.
  - d Click **Manage > Load Balancer > Application Profiles**.
  - e Click **Add** and specify the application profile parameters.

Version	Procedure
NSX 6.4.5 and later	<ol style="list-style-type: none"> <li>1 In the <b>Application Profile Type</b> drop-down menu, select <b>HTTPS End-to-End</b>.</li> <li>2 Click <b>Server SSL &gt; Service Certificates</b>.</li> <li>3 Select the web server certificate that you added in step 1.</li> </ol>
NSX 6.4.4 and earlier	<ol style="list-style-type: none"> <li>1 In the <b>Type</b> drop-down menu, select <b>HTTPS</b>.</li> <li>2 Select the <b>Enable Pool Side SSL</b> check box.</li> <li>3 Select the <b>Configure Service Certificates</b> check box.</li> <li>4 Select the web server certificate that you added in step 1.</li> </ol>

- 3 Create a virtual server.
  - a Log in to the vSphere Web Client.
  - b Click **Networking & Security > NSX Edges**.
  - c Double-click an NSX Edge.
  - d Click **Manage > Load Balancer > Virtual Servers**.
  - e Click **Add** and specify the virtual server parameters.
    - 1 Enable the virtual server to make this virtual server available for use.
    - 2 Select the protocol as **HTTPS**.
    - 3 Select the default pool that is composed of HTTPS servers.
    - 4 Select the application profile that you created in step 2.

For information about specifying the other parameters in the **New Virtual Server** window, see [Add Virtual Servers](#).

## Scenario: SSL Passthrough

This scenario uses an SSL passthrough application profile type. Edge does not close client HTTPS (SSL sessions). Edge load balances TCP sessions to the servers. Client SSL sessions are closed on the servers (not the edge). L7 application rules cannot be applied.

---

**Note** Certificates are not required for SSL passthrough application profiles.

---

## Procedure

- 1 Create an SSL passthrough application profile.
  - a Log in to the vSphere Web Client.
  - b Click **Networking & Security > NSX Edges**.
  - c Double-click an NSX Edge.
  - d Click **Manage > Load Balancer > Application Profiles**.
  - e Click **Add** and specify the application profile parameters.

Version	Procedure
NSX 6.4.5 and later	<ol style="list-style-type: none"> <li>1 In the <b>Application Profile Type</b> drop-down menu, select <b>SSL Passthrough</b>.</li> <li>2 In the <b>Persistence</b> drop-down menu, select <b>None</b>.</li> </ol>
NSX 6.4.4 and earlier	<ol style="list-style-type: none"> <li>1 In the <b>Type</b> drop-down menu, select <b>HTTPS</b>.</li> <li>2 Select the <b>Enable SSL Passthrough</b> check box.</li> <li>3 In the <b>Persistence</b> drop-down menu, select <b>None</b>.</li> </ol>

- 2 Create a virtual server.
  - a Log in to the vSphere Web Client.
  - b Click **Networking & Security > NSX Edges**.
  - c Double-click an NSX Edge.
  - d Click **Manage > Load Balancer > Virtual Servers**.
  - e Click **Add** and specify the virtual server parameters.
    - 1 Enable the virtual server to make this virtual server available for use.
    - 2 Select the protocol as **HTTPS**.
    - 3 Select the default pool that is composed of HTTPS servers.
    - 4 Select the application profile that you created in step 1.

For information about specifying the other parameters in the **New Virtual Server** window, see [Add Virtual Servers](#).

### Note

- If **Acceleration** is enabled and there are no L7 related configurations, Edge does not end the session.
- If **Acceleration** is disabled, the session might be treated as L7 TCP mode, and Edge ends it into two sessions.

## Scenario: SSL Client and Server Authentication

This scenario uses an HTTPS end-to-end application profile type with SSL client and server authentication.

### Client Authentication

Clients access the Web application through HTTPS. HTTPS session is closed on the Edge VIP and the session requests for a client certificate.

- 1 Add a Web server certificate that is signed by a root CA. For more information, see [Scenario: Import SSL Certificate](#).
- 2 Create an HTTPS application profile.

Version	Procedure
NSX 6.4.5 and later	<ol style="list-style-type: none"> <li>1 In the <b>Application Profile Type</b> drop-down menu, select <b>HTTPS End-to-End</b>.</li> <li>2 Click <b>Client SSL &gt; CA Certificates</b>.</li> <li>3 Select the web server certificate that you added in step 1.</li> <li>4 In the <b>Client Authentication</b> drop-down menu, select <b>Required</b>.</li> </ol>
NSX 6.4.4 and earlier	<ol style="list-style-type: none"> <li>1 In the <b>Type</b> drop-down menu, select <b>HTTPS</b>.</li> <li>2 Click <b>Virtual Server Certificates &gt; CA Certificates</b>. CA verifies the client certificate.</li> <li>3 Select the web server certificate that you added in step 1.</li> <li>4 In the <b>Client Authentication</b> drop-down menu, select <b>Required</b>.</li> </ol>

- 3 Create a virtual server. For information about specifying virtual server parameters, see [Scenario: Import SSL Certificate](#).

---

**Note** In NSX 6.4.4 and earlier, when the **Enable Pool Side SSL** option is disabled in the application profile, the pool selected is composed of HTTP servers. When the **Enable Pool Side SSL** option is enabled in the application profile, the pool selected is composed of HTTPS servers.

---

- 4 Add a client certificate that is signed by the root CA in the browser.
- 5 a Go to the website <https://www.sslshopper.com/ssl-converter.html>.

- b Convert certificate and private key to the *px* file. For complete examples of certificate and private key, see [Example: Certificate and Private Key](#) topic.

**Certificate File to Convert:**  client.crt  
**Private Key File:**  client.key  
**Chain Certificate File (optional):**  Dimi-CA.crt  
**Chain Certificate File 2 (optional):**  No file chosen  
**Type of Current Certificate:**  Detected type from file extension  
**Type To Convert To:**   
**PFX Password:**

- c Import the *px* file in the browser.

## Server Authentication

Clients access the Web application through HTTPS. HTTPS session is closed on the Edge VIP. Edge establishes new HTTPS connections to the servers, and it requests and verifies the server certificate.

NSX Edge accepts specific ciphers.

- 1 Add the Web server certificate that is chained with the root CA certificate for server certificate authentication. For more information, see [Scenario: Import SSL Certificate](#).
- 2 Create an HTTPS application profile.

Version	Procedure
NSX 6.4.5 and later	<ol style="list-style-type: none"> <li>1 In the <b>Application Profile Type</b> drop-down menu, select <b>HTTPS End-to-End</b>.</li> <li>2 Click <b>Server SSL &gt; CA Certificates</b>.</li> <li>3 Next to <b>Cypher</b>, click the <b>Edit</b>  icon, and select the ciphers.</li> <li>4 Enable the <b>Server Authentication</b> option.</li> <li>5 Select the CA certificate that you added in step 1.</li> </ol>
NSX 6.4.4 and earlier	<ol style="list-style-type: none"> <li>1 In the <b>Type</b> drop-down menu, select <b>HTTPS</b>.</li> <li>2 Select the <b>Enable Pool Side SSL</b> check box.</li> <li>3 Click <b>Pool Certificates &gt; CA Certificates</b>. CA verifies the client certificate from the back-end HTTPS server.</li> <li>4 Select the <b>Server Authentication</b> check box.</li> <li>5 Select the CA certificate that you added in step 1.</li> <li>6 From the <b>Cipher</b> list, select the required ciphers.</li> </ol> <p><b>Note</b> If your preferred cipher is not in the approved ciphers list, it resets to <code>Default</code>.</p> <p>After upgrading from an old NSX version, if the cipher is null or empty, or if the cypher is not in the approved ciphers list of the old version, it resets to <code>Default</code>.</p> <ol style="list-style-type: none"> <li>7 In the <b>Client Authentication</b> drop-down menu, select <b>Required</b>.</li> </ol>

- 3 Create a virtual server. For information about specifying virtual server parameters, see [Scenario: Import SSL Certificate](#).

---

**Note** In NSX 6.4.4 and earlier, when the **Enable Pool Side SSL** option is disabled in the application profile, the pool selected is composed of HTTP servers. When the **Enable Pool Side SSL** option is enabled in the application profile, the pool selected is composed of HTTPS servers.

---

An NSX services gateway offers IP address pooling and one-to-one static IP address allocation and external DNS server configuration.

You must have a working NSX Edge instance before you can use any of the above services. For information on setting up NSX Edge, see [NSX Edge Configuration](#).

This chapter includes the following topics:

- [Managing DHCP Service](#)
- [Configuring DHCP Relay](#)
- [Configure a DNS Server](#)

## Managing DHCP Service

NSX Edge supports IP address pooling and one-to-one static IP address allocation. Static IP address binding is based on the vCenter managed object ID and interface ID of the requesting client.

NSX Edge DHCP service adheres to the following guidelines:

- Listens on the NSX Edge internal interface for DHCP discovery.
- Uses the IP address of the internal interface on NSX Edge as the default gateway address for all clients (except for non-directly connected pools), and the broadcast and subnet mask values of the internal interface for the container network.

---

**Note** By design, DHCP service is supported on the internal interfaces of an NSX Edge. However, in some situations, you may choose to configure DHCP on an uplink interface of the edge and configure no internal interfaces. In this situation, the edge can listen to the DHCP client requests on the uplink interface, and dynamically assign IP addresses to the DHCP clients. Later, if you configure an internal interface on the same edge, DHCP service stops working because the edge starts listening to the DHCP client requests on the internal interface.

---

You must restart the DHCP service on client virtual machines in the following situations:

- You changed or deleted a DHCP pool, default gateway, or DNS server.
- You changed the internal IP address of the NSX Edge instance.

## Add a DHCP IP Pool

DHCP service requires a pool of IP addresses.

An IP pool is a sequential range of IP addresses within the network. Virtual machines protected by NSX Edge that do not have an address binding are allocated an IP address from this pool. An IP pool's range cannot intersect one another, as a result one IP address can belong to only one IP pool.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > DHCP**.
- 5 Click **Add**.
- 6 Configure the following general options for the new DHCP IP pool.

Option	Action
<b>Start IP</b>	Enter the starting IP address for the pool.
<b>End IP</b>	Enter the ending IP address for the pool.
<b>Domain Name</b>	Enter the domain name of the DNS server. This setting is optional.
<b>Auto Configure DNS</b>	Select to use the DNS service configuration for the DHCP binding.
<b>Primary Name Server</b>	If you did not select <b>Auto Configure DNS</b> , type the <b>Primary Nameserver</b> for the DNS service. You must enter the IP address of a DNS server for hostname-to-IP address resolution. This setting is optional.
<b>Secondary Name Server</b>	If you did not select <b>Auto Configure DNS</b> , type the <b>Secondary Nameserver</b> for the DNS service. You must enter the IP address of a DNS server for hostname-to-IP address resolution. This setting is optional.
<b>Default Gateway</b>	Enter the default gateway address. If you do not specify the default gateway IP address, the internal interface of the NSX Edge instance is taken as the default gateway. This setting is optional.
<b>Subnet Mask</b>	Specify the subnet mask. The subnet mask must be same as the subnet mask of the Edge interface or the DHCP Relay, in case of a distributed router.
<b>Lease Never Expires</b>	Select to bind the address to the MAC address of the virtual machine forever. If you select this, <b>Lease Time</b> is disabled.
<b>Lease Time</b>	Select whether to lease the address to the client for the default time (one day), or enter a value in seconds. If you selected <b>Lease never expires</b> , you cannot specify the lease time. This setting is optional.

## 7 (Optional) Configure the following DHCP options.

Option	Action
Next Server	Next boot TFTP server, used by the PXE boot or bootp.
TFTP server name (option 66)	Enter a unicast IPv4 address or a host name that the device will use to download the file specified in bootfile name (option 67).
TFTP server address (option 150)	Enter one or more TFTP server IPv4 addresses.
Bootfile name (option 67)	Enter the bootfile file name that is to be downloaded from the server specified in TFTP server name (option 66).
Interface MTU (option 26)	The Maximum Transmission Unit (MTU) is the maximum frame size that can be sent between two hosts without fragmentation. This option specifies the MTU size to be used on the interface. One MTU size (in bytes) can be set for each pool and static binding. The MTU minimum value is 68 bytes and the maximum value is 65535 bytes. If the interface MTU is not set on the DHCP server, DHCP clients will keep the OS default setting of the interface MTU.
Classless static route (option 121)	<p>Each classless static route option may have multiple routes with the same destination. Each route includes a destination subnet, subnet mask, next hop router . Note that 0.0.0.0/0 is an invalid subnet for a static route. For information about classless static routes and option 121, refer to RFC 3442.</p> <ol style="list-style-type: none"> <li>Click <b>Add</b>.</li> <li>Enter the destination and the next hop router IP address.</li> </ol> <p>In NSX 6.2.5 and later, if a DHCP pool is configured on an Edge Services Gateway with both classless static routes and a default gateway, the default gateway is added as a classless static route.</p>

## 8 Click **Add** or **OK**.

## Start the DHCP Service

Start the DHCP service to allow NSX Edge to automatically assign an IP address to a virtual machine from a defined IP pool.

### Prerequisites

A DHCP IP pool must be added.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > DHCP**.
- 5 Click **Start**.
- 6 (Optional) Enable logging policy and select the log level.
- 7 Click **Publish Changes**.

## Results

---

**Important** It is a good practice to create a firewall rule to prevent malicious users from introducing rogue DHCP servers. To do this, add a firewall rule that allows UDP traffic only on ports 67 and 68 when the traffic is going to or from a valid DHCP server IP address. For details, see [Working with Firewall Rules](#).

---

### What to do next

Create an IP pool and bindings.

## Edit DHCP IP Pool

You can edit the DHCP IP pool to add or remove IP addresses.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > DHCP**.
- 5 Select a DHCP pool and click the **Edit** icon.
- 6 Make the appropriate changes and click **Add** or **OK**.

## Add a DHCP Static Binding

If you have services running on a virtual machine and do not want the IP address to be changed, you can bind an IP address to the MAC address of a virtual machine. The IP address you bind must not overlap an IP pool.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > DHCP**.
- 5 Select **Bindings** from the left panel.
- 6 Click **Add**.
- 7 Select one of the two binding options.
  - **Use VM NIC Binding:** Select this option when you know the vNIC index of the VM. NSX determines the MAC address from the vNIC index and binds the IP address of the VM to the MAC address.

- **Use MAC Binding:** Select this option when you know the MAC address of the VM and want to use it for static binding with the IP address.

## 8 Configure the general DHCP binding settings.

a Do any one of the following:

- If you selected the **Use VM NIC Binding** option, select the interface to bind, the virtual machine, and the vNIC index of the VM to bind to the IP address.
- If you selected the **Use MAC Binding** option, enter the MAC address of the VM that you want to use for static binding.

b Specify the other general settings.

These options are common to VM NIC binding and MAC binding.

Option	Action
Host name	Type the host name of the DHCP client virtual machine.
IP Address	Enter the address to which to bind the MAC address of the selected virtual machine.
Subnet Mask	Specify the subnet mask. The subnet mask should be same as the subnet mask of the Edge interface or the DHCP Relay, in case of distributed router.
Domain Name	Enter the domain name of the DNS server.
Default Gateway	Enter the default gateway address. If you do not specify the default gateway IP address, the internal interface of the NSX Edge instance is taken as the default gateway.
Lease never expires	Select to bind the address to the MAC address of the virtual machine forever.
Lease Time	If you did not select <b>Lease never expires</b> , select whether to lease the address to the client for the default time (one day), or enter a value in seconds.

## 9 Configure the DNS settings.

Option	Action
Auto configure DNS	Select to use the DNS service configuration for the DHCP binding.
Primary Name Server	If you did not select <b>Auto Configure DNS</b> , enter the <b>Primary Nameserver</b> for the DNS service. You must enter the IP address of a DNS server for hostname-to-IP address resolution.
Secondary Name Server	If you did not select <b>Auto Configure DNS</b> , enter the <b>Secondary Nameserver</b> for the DNS service. You must enter the IP address of a DNS server for hostname-to-IP address resolution.

10 (Optional) Specify the DHCP options. For detailed information about configuring DHCP options, see step 7 in [Add a DHCP IP Pool](#) .

11 Click **Add**, and then click **Publish Changes**.

## Edit DHCP Binding

You assign a different static IP address that is bound to a MAC address of a virtual machine.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Click **Manage > DHCP**.
- 5 Select **Bindings** from the left panel and click the binding to edit.
- 6 Click the **Edit** icon.
- 7 Make the appropriate changes and click **Save** or **OK**.

## Configuring DHCP Relay

Dynamic Host Configuration Protocol (DHCP) relay enables you to leverage your existing DHCP infrastructure from within NSX without any interruption to the IP address management in your environment. DHCP messages are relayed from virtual machine(s) to the designated DHCP server(s) in the physical world. This enables IP addresses within NSX to continue to be in synch with IP addresses in other environments.

DHCP configuration is applied on the logical router port and can list several DHCP servers. Requests are sent to all listed servers. While relaying the DHCP request from the client, the relay adds a Gateway IP Address to the request. The external DHCP server uses this gateway address to match a pool and allocate an IP address for the request. The gateway address must belong to a subnet of the NSX port on which the relay is running.

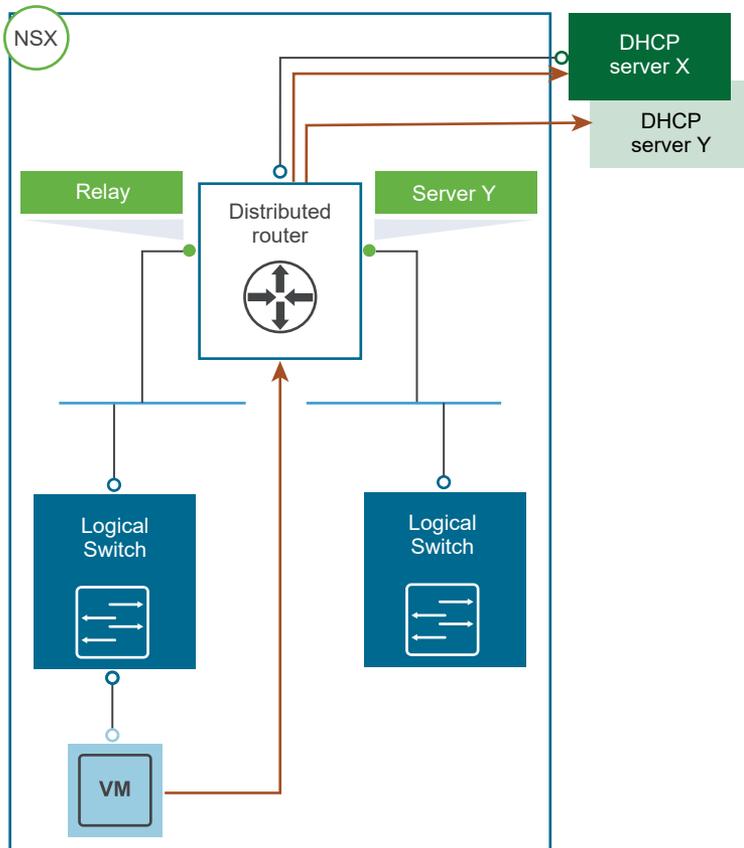
You can specify a different DHCP server for each logical switch and can configure multiple DHCP servers on each logical router to provide support for multiple IP domains.

---

**Note** If the DHCP Offer contains an IP address that doesn't match a logical interface (LIF), the DLR does not relay it back to the VM. The packet is dropped.

---

When configuring pool and binding at DHCP server, ensure that the subnet mask of the pool/binding for the relayed queries is same as the interface of the DHCP relay. Subnet mask information must be provided in API while DLR is acting as DHCP relay between VMs and Edge providing DHCP service. This subnet mask should match the one configured on gateway interface for VMs on DLR.



### Note

- DHCP relay does not support overlapping IP address space (option 82).
- DHCP Relay and DHCP service cannot run on a port/vNic at the same time. If a relay agent is configured on a port, a DHCP pool cannot be configured on the subnet(s) of this port.

## Add DHCP Relay Server

Add the external relay server(s) to which you want the DHCP messages to be relayed to. The relay server can be an IP set, IP address block, domain, or a combination of all of these. Messages are relayed to each listed DHCP server.

### Prerequisites

- DHCP relay does not support overlapping IP address space (option 82).
- DHCP Relay and DHCP service cannot run on a port/vNic at the same time. If a relay agent is configured on a port, a DHCP pool cannot be configured on the subnet(s) of this port.
- If the DHCP Offer contains an IP address that doesn't match a logical interface (LIF), the DLR does not relay it back to the VM. The packet is dropped.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > NSX Edges**.

- 2 Double-click an NSX Edge.
- 3 Click **Manage > DHCP > Relay**.
- 4 Next to **DHCP Relay Global Configuration**, click **Edit**.
- 5 Add a DHCP Relay Server by using one of these methods, or by using a combination of these methods.

Method	Description
Select IP Set	<p>In NSX 6.4.4 and later:</p> <ul style="list-style-type: none"> <li>■ Select an IP set from the <b>Available Objects</b> list, and move it to the <b>Selected Objects</b> list.</li> </ul> <p>In NSX 6.4.3 and earlier:</p> <ol style="list-style-type: none"> <li>a Click the <b>Add</b> icon.</li> <li>b Select an IP set from the <b>Available Objects</b> list, and move it to the <b>Selected Objects</b> list.</li> </ol>
Specify IP Addresses	Enter a comma-separated list of IP addresses.
Specify Domain Names	<p>Enter a comma-separated list of domain names. For example, <code>group1.vmware.com,group2.vmware.com</code>.</p> <p>Ensure that you have manually added the domain IP addresses to the firewall.</p>

- 6 Click **Save** or **OK**.

## Add DHCP Relay Agent

Add the Edge interfaces from which you want the DHCP requests to be relayed to the external DHCP relay servers.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > NSX Edges**.
- 2 Double-click an NSX Edge.
- 3 Click **Manage > DHCP > Relay**.
- 4 In the **DHCP Relay Agent** area, click **Add**.
- 5 In **vNIC**, ensure that an internal vNIC is selected.
 

The **Gateway IP Address** displays the primary IP address of the selected vNIC.
- 6 Click **Add** or **OK**.

## Configure a DNS Server

You can configure external DNS servers on an NSX Edge. The edge forwards DNS requests from client applications to the DNS servers to resolve a network name. The edge can also cache the response that it receives from the DNS servers. DNS service is supported on an Edge Services Gateway, DLR, and UDLR in a Cross-vCenter NSX environment.

**Procedure**

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security > NSX Edges**.
- 3 Double-click an NSX Edge.
- 4 Navigate to edit DNS configuration settings.

Version	Procedure
NSX 6.4.4 and later	<ol style="list-style-type: none"> <li>a Click <b>Manage &gt; DNS</b>.</li> <li>b Next to DNS Configuration, click <b>Change</b>.</li> </ol>
NSX 6.4.3 and earlier	<ol style="list-style-type: none"> <li>a Click <b>Manage &gt; Settings &gt; Configuration</b>.</li> <li>b In the <b>DNS Configuration</b> pane, click <b>Change</b>.</li> </ol>

- 5 Click **Enable DNS Service**.
- 6 Enter the IP address of one or both the DNS servers.
- 7 Change the default cache size, if necessary. The default size is 16 MB.
- 8 To log DNS traffic, click **Enable Logging**, and select the log level. Default log level is **info**.  
Generated logs are sent to the syslog server.
- 9 Click **Save** or **OK**.

Service Composer helps you provision and assign network and security services to applications in a virtual infrastructure. You map these services to a security group, and the services are applied to the virtual machines in the security group.

## Security Group

You begin by creating a security group to define assets that you want to protect. Security groups may be static (including specific virtual machines) or dynamic where membership may be defined in one or more of the following ways:

- vCenter containers (clusters, port groups, or datacenters)
- Security tags, IPset, MACset, or even other security groups. For example, you may include a criteria to add all members tagged with the specified security tag (such as AntiVirus.virusFound) to the security group.
- Directory Groups (if NSX Manager is registered with Active Directory)
- Regular expressions such as virtual machines with name *VM1*

Note that security group membership changes constantly. For example, a virtual machine tagged with the AntiVirus.virusFound tag is moved into the Quarantine security group. When the virus is cleaned and this tag is removed from the virtual machine, it again moves out of the Quarantine security group.

---

**Important** If a VM's VM-ID is regenerated due to move or copy, the security tags are not propagated to the new VM-ID.

---

## Security Policy

A security policy is a collection of the following service configurations.

**Table 18-1. Security services contained in a security policy**

Service	Description	Applies to
Firewall rules	Rules that define the traffic to be allowed to, from, or within the security group.	vNIC
Endpoint service	Third party solution provider services such as anti-virus or vulnerability management services.	virtual machines
Network introspection services	Services that monitor your network such as IPS.	virtual machines

During service deployment in NSX, the third party vendor selects the service category for the service being deployed. A default service profile is created for each vendor template.

When third party vendor services are upgraded to NSX 6.1, default service profiles are created for the vendor templates being upgraded. Existing service policies that include Guest Introspection rules are updated to refer to the service profiles created during the upgrade.

### Mapping Security Policy to Security Group

You map a security policy (say SP1) to a security group (say SG1). The services configured for SP1 are applied to all virtual machines that are members of SG1.

**Note** When you have many security groups to which you need to attach the same security policy, create an umbrella security group that includes all these child security groups, and apply the common security policy to the umbrella security group. This will ensure that the NSX distributed firewall utilises ESXi host memory efficiently.

**Figure 18-1. Service Composer overview**

If a virtual machine belongs to more than one security group, the services that are applied to the virtual machine depends on the precedence of the security policy mapped to the security groups.

Service Composer profiles can be exported and imported as backups or for use in other environments. This approach to managing network and security services helps you with actionable and repeatable security policy management.

This chapter includes the following topics:

- [Using Service Composer](#)
- [Service Composer Canvas](#)
- [Working with Security Tags](#)
- [Viewing Effective Services](#)
- [Working with Security Policies](#)
- [Importing and Exporting Security Policy Configurations](#)
- [Service Composer Scenarios](#)

## Using Service Composer

Service Composer helps you consume security services with ease.

Let us walk through an example to show how Service Composer helps you protect your network end-to-end. Let us say you have the followings security policies defined in your environment:

- An initial state security policy that includes a vulnerability scanning service (InitStatePolicy)
- A remediation security policy that includes a network IPS service in addition to firewall rules and an anti-virus service (RemPolicy)

Ensure that the RemPolicy has higher weight (precedence) than InitStatePolicy.

You also have the followings security groups in place:

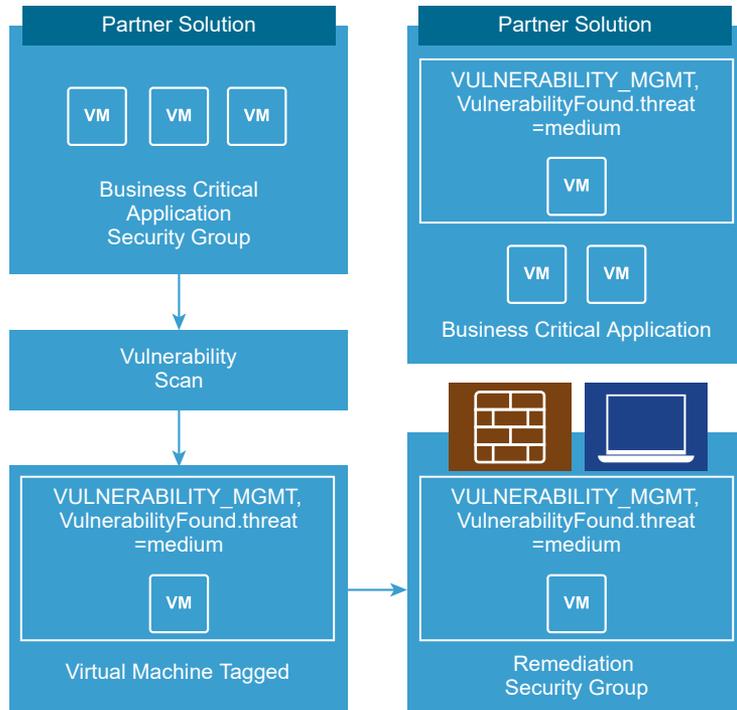
- An applications assets group that includes the business critical applications in your environment (AssetGroup)
- A remediation security group defined by a tag that indicates the virtual machine is vulnerable (`VULNERABILITY_MGMT.VulnerabilityFound.threat=medium`) named RemGroup

You now map the InitStatePolicy to AssetGroup to protect all business critical applications in your environment. You also map RemPolicy to RemGroup to protect vulnerable virtual machines.

When you initiate a vulnerability scan, all virtual machines in AssetGroup are scanned. If the scan identifies a virtual machine with a vulnerability, it applies the `VULNERABILITY_MGMT.VulnerabilityFound.threat=medium` tag to the virtual machine.

Service Composer instantly adds this tagged virtual machine to RemGroup, where a network IPS solution is already in place to protect this vulnerable virtual machine.

Figure 18-2. Service Composer in action



This topic will now take you through the steps required to consume the security services offered by Service Composer.

### Procedure

#### 1 Create a Security Group in Service Composer

You create a security group at the NSX Manager level.

#### 2 Global Settings

#### 3 Create a Security Policy

A security policy is a set of Guest Introspection, firewall, and network introspection services that can be applied to a security group. The order in which security policies are displayed is determined by the weight associated with the policy. By default, a new policy is assigned the highest weight so that it is at the top of the table. However, you can modify the default suggested weight to change the order assigned to the new policy.

#### 4 Apply a Security Policy to a Security Group

You can apply a security policy to a security group to secure your virtual desktops, business critical applications, and the connections between them. You can also view a list of the services that were not applied and the reason they failed to apply.

## Create a Security Group in Service Composer

You create a security group at the NSX Manager level.

## Prerequisites

If you are creating a security policy for use with RDSH, ensure that:

- Active Directory Server must be integrated with NSX Manager.
- Hosts must have DFW enabled and be upgraded to NSX 6.4.0.
- Guest machines must run updated VMware Tools.
- The version of the GI SVM must be 6.4 or later.
- The rule must be created in a new section of Firewall Rules.
- The rule must have **Enable User Identity at Source** selected.
- The **Applied to** field is not supported for rules for remote desktop access.
- ICMP is not supported for IDFW for RDSH.

## Procedure

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Security > Service Composer**.
- 2 Ensure that you are in the **Security Groups** tab.
- 3 Click the **Add Security Group** or the **Add** icon.

Security groups for use with Identity Firewall for RDSH, must use security policies that are marked **Enable User Identity at Source** when created. Security groups for use with Identity Firewall for RDSH can only contain Active Directory (AD) groups, and all nested security groups must also be AD groups.

- 4 Type a name and description for the security group and click **Next**.
- 5 On the Dynamic Membership page, define the criteria that an object must meet for it to be added to the security group you are creating.

For example, you may include a criteria to add all members tagged with the specified security tag (such as AntiVirus.virusFound) to the security group.

Or you can add all virtual machines containing the name `w2008` AND virtual machines that are in the logical switch `global_wire` to the security group.

Security tags are case sensitive.

---

**Note** If you define a security group by virtual machines that have a certain security tag applied to them, you can create a dynamic or conditional workflow. The moment the tag is applied to a virtual machine, the virtual machine is automatically added to that security group.

---

- 6 Click **Next**.

- 7 On the Select objects to include page, select the object type from the drop-down.

Note that security groups for use in remote desktop sessions can only contain Directory groups.

- 8 Select the object that you want to add to the include list. You can include the following objects in a security group.

- Other security groups to nest within the security group you are creating.
- Cluster
- Logical switch
- Network
- Virtual App
- Datacenter
- IP sets
- AD groups

---

**Note** The AD configuration for NSX security groups is different from the AD configuration for vSphere SSO. NSX AD group configuration is for end users accessing guest virtual machines while vSphere SSO is for administrators using vSphere and NSX.

---

- MAC Sets

---

**Note** Service Composer allows use of Security Groups that contain MAC Sets in Policy configurations, however, Service Composer fails to enforce rules for that specific MAC Set. Service Composer works on Layer 3 and does not support Layer 2 constructs.

---

- Security tag
- vNIC
- Virtual Machine
- Resource Pool
- Distributed Virtual Port Group

The objects selected here are always included in the security group regardless of whether or not they match the dynamic criteria.

When you add a resource to a security group, all associated resources are automatically added. For example, when you select a virtual machine, the associated vNIC is automatically added to the security group.

- 9 Click **Next** and double-click the objects that you want to exclude from the security group.

The objects selected here are always excluded from the security group even if they match the dynamic criteria or are selected in the include list .

- 10 Click **Finish**.

## Example

Membership of a security group is determined as follows:

{Expression result (derived from [step 5](#)) + Inclusions (specified in [step 8](#)) - Exclusion (specified in [step 9](#)) which means that inclusion items are first added to the expression result. Exclusion items are then subtracted from the combined result.

## Global Settings

### Edit Service Composer Firewall Applied To Setting

You can set the applied to setting for all firewall rules created through Service Composer to either Distributed Firewall or Policy's Security Groups. By default, the applied to is set to Distributed Firewall.

When Service Composer firewall rules have an applied to setting of distributed firewall, the rules are applied to all clusters on which distributed firewall is installed. If the firewall rules are set to apply to the policy's security groups, you have more granular control over the firewall rules, but may need multiple security policies or firewall rules to get the desired result.

#### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Security > Service Composer**.
- 2 Click the **Security Policies** tab.
- 3 To edit the global firewall settings:
  - ◆ In NSX 6.4.1 and later, next to Global Firewall Settings, click the edit (✎) icon.
  - ◆ In NSX 6.4.0, next to Global Settings: Firewall Rules Applied To, click **Edit**.
- 4 Select a default setting for Applied To and click **OK**. This value determines the vNICs on which the firewall rule will be applied.

Option	Description
<b>Distributed Firewall</b>	Firewall rules are applied to all clusters on which Distributed Firewall is installed.
<b>Policy's Security Groups</b>	Firewall rules are applied to security groups on which the security policy is applied.

The default Applied To setting can also be viewed and changed via the API. See the *NSX API Guide*.

Note that when using RDSH firewall rules the applied to setting is **Distributed Firewall**. **Policy's Security Groups** is not supported for the applied to setting for RDSH rules.

### Example: Applied To Behavior

In this example scenario, your default firewall rule action with service any, is set to block. You have two security groups: web-servers and app-servers, which contain VMs. You create a security policy, allow-ssh-from-web, which contains the following firewall rule, and apply it to the security group app-servers.

- Name: allow-ssh-from-web
- Source: web-servers
- Destination: Policy's Security Group
- Service: ssh
- Action: allow

If the firewall rule applies to Distributed Firewall, you will be able to ssh from a VM in the security group web-servers to a VM in the security group app-servers.

If the firewall rule applies to Policy's Security Group, you will not be able to ssh, as the traffic will be blocked from reaching the app servers. You will need to create an additional security policy to allow ssh to the app servers, and apply this policy to the security group web-servers.

- Name: allow-ssh-to-app
- Source: Policy's Security Group
- Destination: app-servers
- Service: ssh
- Action: allow

### Synchronize Firewall Configuration

Synchronize firewall configuration synchronizes the Service Composer configuration with the firewall configuration and recreates Service Composer related firewall sections and rules on the firewall side.

Synchronize firewall configuration allows users to re-sync all of the Service composer configuration with firewall configuration. It recreates Service composer related firewall sections and rules on the firewall side. This re-sync is applicable only for firewall and network introspection related policies configurations. This operation creates all the policy sections in order of precedence above the firewall default section.

---

**Note** This operation may take a long time to complete and should be triggered only when necessary.

---

#### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Security > Service Composer**.
- 2 Click the **Security Policies** tab.

3 To synchronize the firewall configuration:

- ◆ In NSX 6.4.1 and later, next to Sync Status, click **Synchronize**.
- ◆ In NSX 6.4.0, next to Global Settings, click **Synchronize**.

Global synchronization will occur with either distributed firewall or policy security groups. All service composer related firewall sections and rules will synchronize and change to honor the **Applied To** selection. See [Edit Service Composer Firewall Applied To Setting](#).

## Create a Security Policy

A security policy is a set of Guest Introspection, firewall, and network introspection services that can be applied to a security group. The order in which security policies are displayed is determined by the weight associated with the policy. By default, a new policy is assigned the highest weight so that it is at the top of the table. However, you can modify the default suggested weight to change the order assigned to the new policy.

### Prerequisites

Ensure that:

- the required VMware built in services (such as Distributed Firewall, and Guest Introspection) are installed.
- the required partner services have been registered with NSX Manager.
- the desired default applied to value is set for Service Composer firewall rules. See [Edit Service Composer Firewall Applied To Setting](#).

If you are creating a security policy framework for Identity Firewall for RDSH:

- Active Directory Server must be integrated with NSX Manager.
- Hosts must have DFW enabled and be upgraded to NSX 6.4.0.
- Guest machines must run updated VMware Tools.
- The version of the GI SVM must be 6.4 or later.
- The rule must be created in a new section of Firewall Rules.
- The rule must have **Enable User Identity at Source** selected.
- The **Applied to** field is not supported for rules for remote desktop access.
- ICMP is not supported for IDFW for RDSH.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Security > Service Composer**.
- 2 Click the **Security Policies** tab.

### 3 To create new security policy:

- ◆ In NSX 6.4.1 and later, click **Add**.
- ◆ In NSX 6.4.0, click the **Create Security Policy** () icon.

### 4 In the Create Security Policy or New Security Policy dialog box, type a name for the security policy.

### 5 Type a description for the security policy. The description must not exceed 255 characters.

NSX assigns a default weight (highest weight +1000) to the policy. For example, if the highest weight amongst the existing policy is 1200, the new policy is assigned a weight of 2200.

Security policies are applied according to their weight - a policy with the higher weight has precedence over a policy with a lower weight.

### 6 Select **Inherit security policy** if you want the policy that you are creating to receive services from another security policy. Select the parent policy.

All services from the parent policy are inherited by the new policy.

### 7 Click **Next**.

### 8 In the Guest Introspection Services page, click **Add** or the **Add Guest Introspection Service** () icon.

a In the Add Guest Introspection Service dialog box, type a name and description for the service.

b Specify whether you want to apply the service or block it.

When you inherit a security policy, you may choose to block a service from the parent policy.

If you apply a service, you must select a service and service profile. If you block a service, you must select the type of service to block.

c If you chose to block the service, select the type of service.

d If you chose to apply the Guest Introspection service, select the service name.

The default service profile for the selected service is displayed, which includes information about the service functionality types supported by the associated vendor template.

e In **State**, specify whether you want to enable the selected Guest Introspection service or disable it.

You can add Guest Introspection services as placeholders for services to be enabled at a later time. This is especially useful for cases where services need to be applied on-demand (for example, new applications).

- f Select whether the Guest Introspection service is to be enforced (i.e. it cannot be overridden). If the selected service profile supports multiple service functionality types, then this is set to **Enforce** by default and cannot be changed.

If you enforce a Guest Introspection service in a security policy, other policies that inherit this security policy would require that this policy be applied before the other child policies. If this service is not enforced, an inheritance selection would add the parent policy after the child policies are applied.

- g Click **OK**.

You can add additional Guest Introspection services by following the above steps. You can manage the Guest Introspection services through the icons above the service table.

In NSX 6.4.0, you can export or copy the services on this page by clicking the  icon on the bottom right side of the Guest Introspection Services page.

- 9 Click **Next**.

- 10 On the Firewall page, you are defining firewall rules for the security groups(s) that this security policy will be applied to.

When creating a security policy for Identity Firewall for RDSH, **Enable User Identity at Source** must be checked. Note that this disables the enable stateless firewall option because the TCP connection state is tracked for identifying the context. This flag cannot be changed while the policy is being updated. Once a security policy is created with **Enable User Identity at Source** inheritance is not supported.

- a Click the checkbox to enable the following optional parameters:

Option	Description
Enable User Identity at Source	When using Identity Firewall for RDSH, <b>Enable User Identity at Source</b> must be checked. Note that this disables the enable stateless firewall option because the TCP connection state is tracked for identifying the context.
Enable TCP Strict	Enables you to set TCP strict for each firewall section.
Enable Stateless Firewall	Enables stateless firewall for each firewall section.

- b Click **Add**, or the **Add Firewall Rule** () icon.
- c Type a name and description for the firewall rule you are adding.
- d Select **Allow**, **Block**, or **Reject** to indicate whether the rule needs to allow, block, or reject traffic to the selected destination.
- e Select the source for the rule. By default, the rule applies to traffic coming from the security groups to which this policy gets applied to. To change the default source, click **Select** or **Change** and select the appropriate security groups.

- f Select the destination for the rule.

---

**Note** Either the Source or Destination (or both) must be security groups to which this policy gets applied to.

---

Say you create a rule with the default Source, specify the Destination as Payroll, and select **Negate Destination**. You then apply this security policy to security group Engineering. This would result in Engineering being able to access everything except for the Payroll server.

- g Select the services and/or service groups to which the rule applies to.
- h Select **Enabled** or **Disabled** to specify the rule state.
- i Select **Log** to log sessions matching this rule.  
Enabling logging may affect performance.
- j Enter the text that you want to add in the **Tag** text box while adding or editing the firewall rule.
- k Click **OK**.

You can add additional firewall rules by following the above steps. You can manage the firewall rules through the icons above the firewall table.

In NSX 6.4.0, you can export or copy the rules on this page by clicking the  icon on the bottom right side of the Firewall page.

The firewall rules you add here are displayed on the Firewall table. VMware recommends that you do not edit Service Composer rules in the firewall table. If you must do so for an emergency troubleshooting, you must re-synchronize Service Composer rules with firewall rules as follows:

- In NSX 6.4.1 and later, select **Synchronize** on the Security Policies tab.
- In NSX 6.4.0, select the **Synchronize Firewall Rules** from the **Actions** menu on the Security Policies tab.

- 11 Click **Next**.

The Network Introspection Services page displays NetX services that you have integrated with your VMware virtual environment.

- 12 Click the checkbox to enable the following optional parameters:

Option	Description
Enable TCP Strict	Enables you to set TCP strict for each firewall section.
Enable Stateless Firewall	Enables stateless firewall for each firewall section.

13 Click **Add**, or the **Add Network Introspection Service** () icon.

- a Enter a name and description for the service you are adding.
- b Select whether or not to redirect to service.
- c Select the service name and profile.
- d Select the source and destination
- e Select the network service that you want to add.

You can make additional selections based on the service you selected.

- f Select whether to enable or disable the service.
- g Select Log to log sessions matching this rule.
- h Enter the text that you want to add in the **Tag** text box.
- i Click **OK**.

You can add additional network introspection services by following the above steps. You can manage the network introspection services through the icons above the service table.

In NSX 6.4.0, you can export or copy the services on this page by clicking the  icon on the bottom right side of the Network Introspection Service page.

---

**Note** Bindings created manually for the Service Profiles used in Service Composer policies will be overwritten.

---

14 Click **Finish**.

The security policy is added to the policies table. You can click the policy name and select the appropriate tab to view a summary of the services associated with the policy, view service errors, or edit a service.

#### What to do next

Map the security policy to a security group.

## Apply a Security Policy to a Security Group

You can apply a security policy to a security group to secure your virtual desktops, business critical applications, and the connections between them. You can also view a list of the services that were not applied and the reason they failed to apply.

#### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Security > Service Composer**.
- 2 Click the **Security Policies** tab.
- 3 Select a security policy, and click **Apply** or the **Apply Security Policy** () icon.

- 4 Select the security group that you want to apply the policy to.

Security groups for use with Identity Firewall for RDSH, must use security policies that are marked **Enable User Identity at Source** when created. Security groups for use with Identity Firewall for RDSH can only contain Active Directory (AD) groups, and all nested security groups must also be AD groups.

If you select a security group defined by virtual machines that have a certain security tag applied to them, you can create a dynamic or conditional workflow. The moment the tag is applied to a virtual machine, the virtual machine is automatically added to that security group.

Network Introspection rules and Endpoint rules associated with the policy will not take effect for security groups containing IPSet and/or MacSet members.

- 5 (Optional) (In NSX 6.4.0 only) Click the **Preview Service Status** icon to see the services that cannot be applied to the selected security group and the reason for the failure.

For example, the security group may include a virtual machine that belongs to a cluster on which one of the policy services has not been installed. You must install that service on the appropriate cluster for the security policy to work as intended.

- 6 Click **OK**.

## Service Composer Canvas

The Service Composer canvas tab offers a graphical view displaying all security groups within the selected NSX Manager. The view also displays details such as members of each security group as well as the security policy applied on it.

---

**Note** In NSX 6.4.1 and later, the **Service Composer > Canvas** tab is removed.

---

This topic introduces Service Composer by walking you through a partially configured system so that you can visualize the mappings between security groups and security policy objects at a high level from the canvas view.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Security > Service Composer**.
- 2 Click the **Canvas** tab.

**Synchronization Status**, displaying errors or warnings, and **Firewall Publish Status**, displaying the date and time stamp of the last successful publishing of firewall rules, are shown at the top of the screen.

All security groups within the selected NSX Manager (that are not contained within another security group) are displayed along with the policies applied on them. The **NSX Manager** drop-down lists all NSX Managers on which the currently logged in user has a role assigned.

## Results

Each rectangular box in the canvas represents a security group and the icons within the box represents security group members and details about the security policy mapped to the security group.

**Figure 18-3. Security group**



A number next to each icon indicates the number of instances - for example,  indicates that 1 security policy is mapped to that security group.

Icon	Click to display
	Security groups nested within the main security group.
	Virtual machines that are currently part of the main security group as well as nested security groups. Click the Errors tab to see virtual machines with service errors.
	Effective security policies mapped to the security group. <ul style="list-style-type: none"> <li>You can create a new security policy by clicking the <b>Create Security Policy</b> () icon. The newly created security policy object is automatically mapped to the security group.</li> <li>Map additional security policies to the security group by clicking the <b>Apply Security Policy</b> () icon.</li> </ul>
	Effective Endpoint services associated with the security policy mapped to the security group. Suppose you have two policies applied to a security group and both have the same category Endpoint service configured. The effective service count in this case will be 1 (since the second lower priority service is overridden). Endpoint service failures, if any, are indicated by the alert icon. Clicking the icon displays the error.
	Effective firewall rules associated with the security policy mapped to the security group. Service failures, if any, are indicated by the alert icon. Clicking the icon displays the error.
	Effective network introspection services associated with the security policy mapped to the security group. Service failures, if any, are indicated by the alert icon. Clicking the icon displays the error.

Clicking an icon displays a dialog box with appropriate details.

Figure 18-4. Details displayed when you click an icon in the security group

No.	Name	Source	Destination	Service	Action	Security
1	f1	Policy's S...	Any	Any	Allow	Sec...
2	f2	Policy's S...	Any	Any	Allow	Sec...

You can search for security groups by name. For example, if you type PCI in the search field in the top right corner of the canvas view, only the security groups with PCI in their names are displayed.

To see the security group hierarchy, click the **Top Level** (▼) icon at the top left of the window and select the security group you want to display. If a security group contains nested security groups, click ▶ to display the nested groups. The top bar displays the name of the parent security group and the icons in the bar display the total number of security policies, endpoint services, firewall services, and network introspection services applicable to the parent group. You can navigate back up to the top level by clicking the **Go up one level** (↶) icon in the top left part of the window.

You can zoom in and out of the canvas view smoothly by moving the zoom slider on the top right corner of the window. The Navigator box shows a zoomed out view of the entire canvas. If the canvas is much bigger than what fits on your screen, it will show a box around the area that is actually visible and you can move it to change the section of the canvas that is being displayed.

### What to do next

Now that we have seen how the mapping between security groups and security policies work, you can begin creating security policies to define the security services you want to apply to your security groups.

## Map Security Group to Security Policy

You can map the selected security group to a security policy.

### Procedure

- 1 Select the security policy that you want to apply to the security group.
- 2 To create a new policy, select **New Security Group**.  
See [Create a Security Policy](#).
- 3 Click **Save**.

## Working with Security Tags

You can view security tags applied on a virtual machine or create a user defined security tag.

Security tags are labels which can be associated with a Virtual Machine (VM). Numerous security tags can be created to identify a specific workload. The matching criteria of a Security Group can be a security tag, and a workload that is tagged can be automatically placed into a Security Group.

Adding or removing security tags to a VM can be done dynamically in response to various criteria such as antivirus or vulnerability scans, and intrusion prevention systems. Tags can also be added and removed manually by an administrator.

---

**Important** If a VM's VM-ID is regenerated due to move or copy, the security tags are not propagated to the new VM-ID.

---

In a cross-vCenter NSX environment, universal security tags are created on the primary NSX manager and are marked for universal synchronization with secondary NSX managers. Universal security tags can be assigned to VMs statically, based on unique ID selection.

## Unique ID Selection

The unique ID selection criteria is used when assigning tags to Virtual Machines on active standby deployments.

Unique ID is used by the NSX Manager when a Virtual Machine (VM) goes from standby to active deployment. The unique ID can be based on VM instance UUID, VM BIOS UUID, VM name, or a combination of these options. If the criteria changes (such as a VM name change) after universal security tags have been created and attached to VMs, the security tag must be detached and reattached to the VMs.

### Procedure

- 1 In the vSphere Web Client, navigate to **Home > Networking & Security > Installation and Upgrade**, click the **Management** tab.
- 2 Click the primary NSX Manager, then select **Actions > Unique ID Selection Criteria**.
- 3 Select one or more of the unique ID options:
  - Use Virtual Machine instance UUID (recommended) - The VM instance UUID is unique within a VC domain, however there are exceptions such as when deployments are made through snapshots. If the VM instance UUID is not unique, use the VM BIOS UUID in combination with the VM name.
  - Use Virtual Machine BIOS UUID - The BIOS UUID is not guaranteed to be unique within a VC domain, but it is always preserved in case of disaster. Use BIOS UUID in combination with VM name.
  - Use Virtual Machine Name - If all of the VM names in an environment are unique, then VM name can be used to identify a VM across vCenters. Use VM name in combination with VM BIOS UUID.
- 4 Click **OK**.

### What to do next

Next, create security tags.

## View Applied Security Tags

You can view the security tags applied to virtual machines in your environment.

### Prerequisites

An antivirus scan must have been run, and a tag applied to the appropriate virtual machine.

---

**Note** Refer to the third party solution documentation for details of the tags applied by those solutions.

---

### Procedure

- 1 In the vSphere Web Client, click **Networking & Security > Groups and Tags**.
- 2 Click the **Security Tags** tab.  
  
A list of tags applied in your environment is displayed along with details about the virtual machines to which those tags have been applied. Note down the exact tag name if you plan on adding a security group to include virtual machines with a specific tag.
- 3 Click the number in the **VM Count** column to view the virtual machines to which that tag in that row has been applied.

## Create a Security Tag

You can create a security tag and apply it to a virtual machine. In a cross-vCenter environment, security tags are synchronized between primary and secondary NSX managers.

### Prerequisites

If creating a universal security tag in an active standby deployment scenario, first set the unique ID selection criteria on the primary NSX Manager.

### Procedure

- 1 In the vSphere Web Client, click **Networking & Security > Groups and Tags**.
- 2 Click the **Security Tags** tab.
- 3 Click **Add**, or the **New Security Tag** icon.
- 4 (Optional) To create a universal security tag for use in cross-vCenter NSX environments:
  - In NSX 6.4.1 and later, click the **Universal Synchronization** toggle button to **On**.
  - In NSX 6.4.0, select **Mark this object for Universal Synchronization**.
- 5 Type a name and description for the tag and click **OK**.

## What to do next

Assign virtual machines to the security tag.

## Assign a Security Tag

In addition to creating a conditional workflow with a dynamic membership-based security tag, you can manually assign a security tag to virtual machines.

Security tags can be used as the matching criteria in security groups. In a cross-vCenter environment, security tags are synchronized between primary and secondary NSX managers.

### Procedure

1 In the vSphere Web Client, click **Networking & Security > Groups and Tags**.

2 Click the **Security Tags** tab.

3 To assign security tag to virtual machine:

- ◆ In NSX 6.4.1 and later, select the required security tag and click **Assign VM**.
- ◆ In NSX 6.4.0, right-click a security tag and select **Assign Security Tag**.

The **Assign Security Tag to Virtual Machine** window appears, populated with available VMs.

4 Select one or more virtual machines to move them to the **Selected Objects** column.

5 Click **OK**.

The **Security Tags** tab appears with an updated VM count for the security tag.

## Edit a Security Tag

You can edit a user-defined security tag. If a security group is based on the tag you are editing, changes to the tag may affect security group membership.

### Procedure

1 In the vSphere Web Client, click **Networking & Security > Groups and Tags**.

2 Click the **Security Tags** tab.

3 To edit a security tag:

- ◆ In NSX 6.4.1 and later, select the required security tag and click **Edit**.
- ◆ In NSX 6.4.0, right-click a security tag and select **Edit Security Tag**.

4 Make the appropriate changes, and click **Save** or **OK**.

## Delete a Security Tag

You can delete a user-defined security tag. If a security group is based on the tag you are deleting, changes to the tag may affect security group membership.

**Procedure**

- 1 In the vSphere Web Client, click **Networking & Security > Groups and Tags**.
- 2 Click the **Security Tags** tab.
- 3 Select a security tag, and click **Delete** or the **Delete Security Tag** (✖) icon.

## Viewing Effective Services

You can view the services that are effective on a security policy object or on a virtual machine.

### View Effective Services on a Security Policy

You can view the services effective on a security policy, including those services inherited from a parent policy.

**Procedure**

- 1 In the vSphere Web Client, navigate to **Networking & Security > Security > Service Composer**.
- 2 Click the **Security Policies** tab.
- 3 Click a security policy in the **Name** column.
- 4 The Security Policy window appears:

NSX Version	Procedure
NSX 6.4.1 and later	The left navigation displays <b>Summary</b> , <b>Firewall Rules</b> , <b>Guest Introspection Services</b> , <b>Network Introspection Services</b> , and <b>Child Policies</b> . You can edit, delete, or apply policy to the security group.
NSX 6.4.0	Go to the <b>Manage &gt; Information Security</b> tab. Each of the three tabs ( <b>Guest Introspection Services</b> , <b>Firewall Rules</b> , <b>Network Introspection Services</b> ) displays the corresponding services for the security policy. Services that are not effective are greyed out. The <b>Overridden</b> column displays the services that are actually applied on the security policy and the <b>Settings &gt; General &gt; Inherits from</b> column displays the security policy from which services are inherited.

### View Service Failures for a Security Policy

You can see the services associated with a security policy that failed to apply to the security group(s) mapped to the policy.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Security > Service Composer**.
- 2 Click the **Security Policies** tab.
- 3 Click a security policy in the **Name** column.
- 4 (Optional) (In NSX 6.4.0 only) Ensure that you are in the **Monitor > Service Errors** tab.  
Clicking the link in the **Status** column takes you to the Service Deployment page where you can correct service errors.

## View Effective Services on a Virtual Machine

You can view the services effective on a virtual machine. If multiple security policies are getting applied on a virtual machine (i.e. a virtual machine is part of multiple security groups that have policies mapped to them), then this view lists all effective services from all these policies, in the order in which they get applied. The service status column displays the status for each service.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **vCenter** and then click **Virtual Machines**.
- 3 Click a virtual machine in the **Name** column.
- 4 Ensure that you are in the **Monitor > Service Composer** tab.

## Working with Security Policies

A security policy is a group of network and security services.

The following network and security services can be grouped into a security policy:

- Endpoint services - anti-virus, and vulnerability management
- Distributed Firewall rules
- Network introspection services - network IPS and network forensics

## Manage Security Policy Priority

Security policies are applied according to their weight - a security policy with a higher weight has a higher priority. When you move a policy up or down in the table, its weight is adjusted accordingly.

Multiple security policies may be applied to a virtual machine either because the security group that contains the virtual machine is associated with multiple policies or because the virtual machine is part of multiple security groups associated with different policies. If there is a conflict between services grouped with each policy, the weight of the policy determines the services that will be

applied to the virtual machine. For example, say policy 1 blocks internet access and has a weight value of 1000 while policy 2 allows internet access and has a weight value of 2000. In this particular case, policy 2 has a higher weight and hence the virtual machine will be allowed internet access.

#### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Security > Service Composer**.
- 2 Click the **Security Policies** tab.
- 3 Click **Manage** or the **Manage Priority** () icon.
- 4 In the Manage Priority dialog box, select the security policy that you want to change the priority for and click **Move Up** () or **Move Down** () .
- 5 If you want to move object to a particular rank, click **Move To**.  
Enter the required rank, and click the green check mark or **OK**.  
Weight is recalculated accordingly to the new rank.
- 6 If you want to edit weight of the policy, click **Edit Weight** .  
Enter the required weight, and click the green check mark or **OK**.
- 7 Click **OK**.

## Edit a Security Policy

You can edit the name or description of a security policy, as well as the associated services and rules.

#### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Security > Service Composer**.
- 2 Click the **Security Policies** tab.
- 3 Select the security policy that you want to edit and click **Edit** or the **Edit Security Policy** () icon.
- 4 In the Edit Security Policy dialog box, make the appropriate changes and click **Finish**.

## Delete a Security Policy

You can delete a security policy.

#### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Security > Service Composer**.

- 2 Click the **Security Policies** tab.
- 3 Select the security policy that you want to delete and click **Delete** or the **Delete Security Policy** (✖) icon.

## Importing and Exporting Security Policy Configurations

You can use the Service Composer to export the security policy configuration to a specific file format from one NSX Manager and import the exported configuration in to another NSX Manager.

In the Service Composer, you cannot directly export security groups. You must first ensure that a security policy is assigned to a security group, and then export that security policy. All the contents of the security policy, such as DFW rules, guest introspection rules, network introspection rules, and the security groups that are bound to the security policy are exported.

When a container security group contains nested security groups, the nested security groups are not exported. While exporting, you can add a prefix to the policy. The prefix gets applied to policy name, policy actions name, and security group name.

When importing the configuration onto a different NSX Manager, you can specify a suffix. The suffix gets applied to policy name, policy actions name, and, security group name. If a security group or security policy with the same name exists on the NSX Manager where the import is happening, the import of the security policy configuration fails.

### Export a Security Policy Configuration

You can export a security policy configuration and save it to your desktop. The saved configuration can be used as a backup for situations where you may accidentally delete a policy configuration, or it can be exported for use in another NSX Manager environment.

#### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Security > Service Composer**.
- 2 Click the **Security Policies** tab.
- 3 Select the security policy that you want to export.
- 4 Click **More** or **Actions**, and then click **Export Configuration**.
- 5 Type a name and description for the configuration that you are exporting.
- 6 If necessary, type a prefix to be added to the security policies and security groups that are being exported.

If you specify a prefix, it is added to the target security policy names thus ensuring that they have unique names.

- 7 Click **Next**.

8 On the **Select Security Policies** page, select the security policy that you want to export and click **Next**.

9 The **Preview Selection** or **Ready to complete** page displays the security policies, endpoint services, firewall rules, and network introspection services to be exported.

This page also displays the security groups on which the security policies are applied.

10 Click **Finish**.

11 Select the directory on your computer where you want to export the blueprint file and click **Save**.

The security policy configuration file is saved at the specified location.

## Import a Security Policy Configuration

You can import a saved security policy configuration either as a backup or to restore a similar configuration on a different NSX Manager.

When importing the configuration, an empty security group is created. All the services, service profiles, applications, and application groups must exist in the destination environment, otherwise the import fails.

### Procedure

1 In the vSphere Web Client, navigate to **Networking & Security > Security > Service Composer**.

2 Click the **Security Policies** tab.

3 Click **More** or **Actions**, and then click the **Import Configuration** icon.

4 Select the security policy configuration file that you want to import.

5 If necessary, type a suffix to be added to the security policies and security groups that are being imported.

If you specify a suffix, it is added to the security policy names being imported thus ensuring that they have unique names.

6 Click **Next**.

Service Composer verifies that all services referred to in the configuration are available in the destination environment. If not, the **Manage Missing Services** page is displayed, where you can map missing services to available target services.

The **Ready to complete** page displays the security policies, endpoint services, firewall rules, and the network introspection services to be imported. This page also displays the security groups on which the security policies are applied.

## 7 Click **Finish**.

The imported security policy configuration is added to the top of the security policy table (above the existing policies) in the target NSX Manager. The original order of the imported rules and security services in the security policy is preserved.

## Service Composer Scenarios

This section illustrates some hypothetical scenarios for Service Composer. It is assumed that the Security Administrator role has been created and assigned to the administrator in each use case.

### Quarantining Infected Machines Scenario

Service Composer can identify infected systems on your network with 3rd party antivirus solutions and quarantine them to prevent further outbreaks.

Our sample scenario shows how you can protect your desktops end to end.

Figure 18-5. Configuring Service Composer

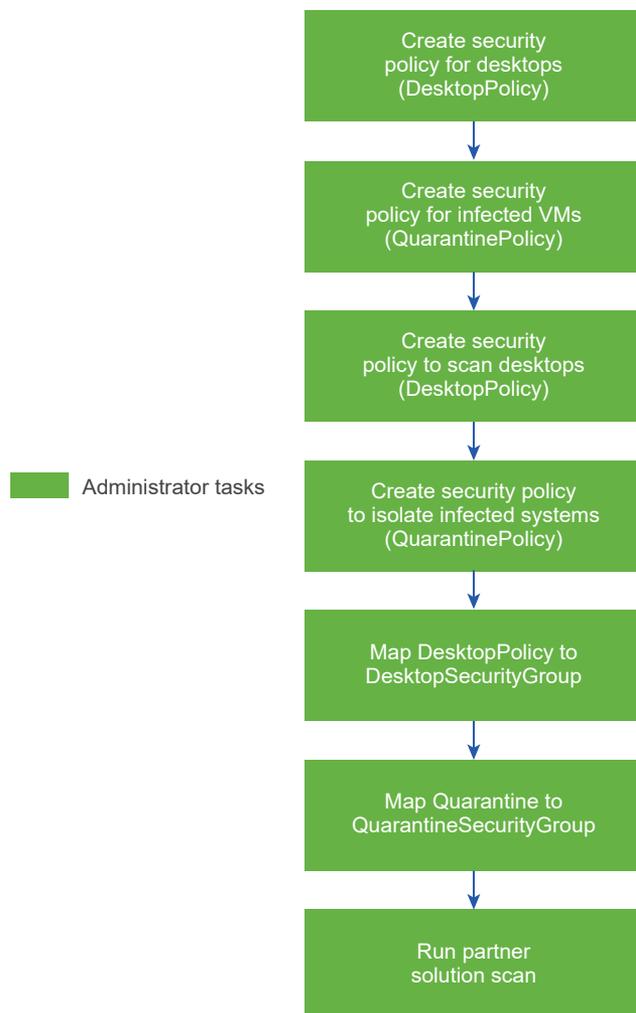
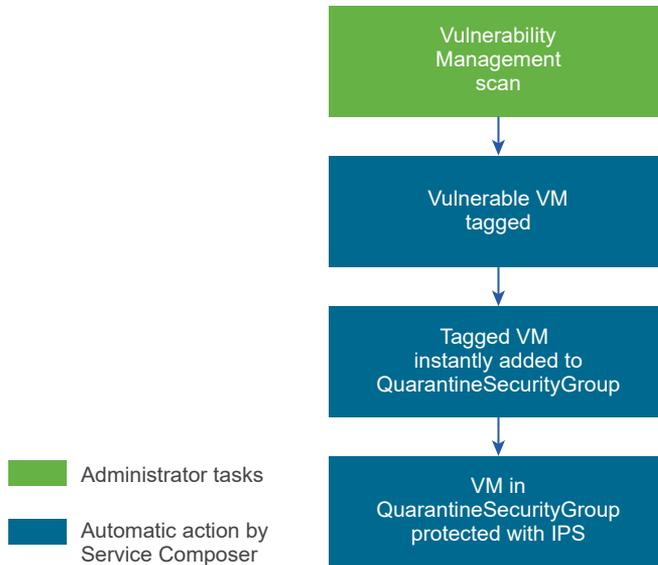


Figure 18-6. Service Composer Conditional Workflow



### Prerequisites

We are aware that Symantec tags infected virtual machine with the **AntiVirus.virusFound** tag.

### Procedure

- 1 Install, register, and deploy the Symantec Antimalware solution.
- 2 Create a security policy for your desktops.
  - a Click the **Security Policies** tab and click the **Add Security Policy** icon.
  - b In **Name**, type **DesktopPolicy**.
  - c In **Description**, type **Antivirus scan for all desktops**.
  - d Change the weight to 51000. The policy precedence is set very high so as to ensure that it is enforced above all other policies.
  - e Click **Next**.

- f On the Add Endpoint Service page, click  and fill in the following values.

Option	Value
Action	Do not modify the default value
Service Type	Anti Virus
Service Name	Symantec Antimalware
Service Configuration	Silver
State	Do not modify the default value
Enforce	Do not modify the default value
Name	Desktop AV
Description	Mandatory policy to be applied on all desktops

- g Click **OK**.
- h Do not add any firewall or network introspection services and click **Finish**.
- 3 Create a security policy for infected virtual machines.
- a Click the **Security Policies** tab and click the **Add Security Policy** icon.
- b In Name, type **QuarantinePolicy**.
- c In Description, type **Policy to be applied to all infected systems..**
- d Do not change the default weight.
- e Click **Next**.
- f On the Add Endpoint Service page, do not do anything and click **Next**.
- g In Firewall, add three rules - one rule to block all outgoing traffic, the next rule to block all traffic with groups, and the last rule to allow incoming traffic only from remediation tools.
- h Do not add any network introspection services and click **Finish**.
- 4 Move **QuarantinePolicy** to the top of the security policy table to ensure that it is enforced before all other policies.
- a Click the **Manage Priority** icon.
- b Select **QuarantinePolicy** and click the **Move Up** icon.
- 5 Create a security group for all desktops in your environment.
- a Log in to the vSphere Web Client.
- b Click **Networking & Security** and then click **Service Composer**.
- c Click the **Security Groups** tab and click the **Add Security Group** icon.
- d In Name, type **DesktopSecurityGroup**.
- e In Description, type **All desktops**.

- f Click **Next** on the next couple of pages.
  - g Review your selections on the Ready to Complete page and click **Finish**.
- 6 Create a Quarantine security group where the infected virtual machines are to be placed.
- a Click the **Security Groups** tab and click the **Add Security Group** icon.
  - b In **Name**, type `QuarantineSecurityGroup`.
  - c In **Description**, type `Dynamic group membership based on infected VMs identified by the antivirus scan`.
  - d On the Define membership Criteria page click  and add the following criteria.



- e Do not do anything on the Select objects to include or Select objects to exclude pages and click **Next**.
  - f Review your selections on the Ready to Complete page and click **Finish**.
- 7 Map the `DesktopPolicy` policy to the `DesktopSecurityGroup` security group.
- a On the Security Policies tab, ensure that the `DesktopPolicy` policy is selected.
  - b Click the **Apply Security Policy** () icon and select the SG\_Desktops group.
  - c Click **OK**.

This mapping ensures that all desktops (part of the `DesktopSecurityGroup`) are scanned when an antivirus scan is triggered.

- 8 Navigate to the canvas view to confirm that `QuarantineSecurityGroup` does not include any virtual machines yet.
- a Click the **Information Security** tab.
  - b Confirm that there are 0 virtual machines in the group ()

- 9 Map `QuarantinePolicy` to `QuarantineSecurityGroup`.

This mapping ensures that no traffic flows to the infected systems.

- 10 From the Symantec Antimalware console, trigger a scan on your network.

The scan discovers infected virtual machine and tags them with the security tag `AntiVirus.virusFound`. The tagged virtual machines are instantly added to `QuarantineSecurityGroup`. The `QuarantinePolicy` allows no traffic to and from the infected systems.

## Backing up Security Configurations

Service Composer can be effectively used to back up your security configurations and restore them at a later time.

### Procedure

- 1 Install, register, and deploy the Rapid 7 Vulnerability Management solution.
- 2 Create a security group for the first tier of the Share Point application - web servers.
  - a Log in to the vSphere Web Client.
  - b Click **Networking & Security** and then click **Service Composer**.
  - c Click the **Security Groups** tab and click the **Add Security Group** icon.
  - d In **Name**, type **SG\_Web**.
  - e In **Description**, type **Security group for application tier**.
  - f Do not do anything on the Define membership Criteria page and click **Next**.
  - g On the Select objects to include page, select the web server virtual machines.
  - h Do not do anything on the Select objects to exclude page and click **Next**.
  - i Review your selections on the Ready to Complete page and click **Finish**.
- 3 Now create a security group for your database and share point servers and name them **SG\_Database**, and **SG\_Server\_SharePoint** respectively. Include the appropriate objects in each group.
- 4 Create a top level security group for your application tiers and name it **SG\_App\_Group**. Add SG\_Web, SG\_Database, and SG\_Server\_SharePoint to this group.
- 5 Create a security policy for your web servers.
  - a Click the Security Policies tab and click the Add Security Policy icon.
  - b In Name, type **SP\_App**.
  - c In Description, type **SP for application web servers**.
  - d Change the weight to 50000. The policy precedence is set very high so as to ensure that it is enforced above most other policies (with the exception of quarantine).
  - e Click Next.

- f On the Endpoint Services page, click  and fill in the following values.

Option	Value
Action	Do not modify the default value
Service Type	Vulnerability Management
Service Name	Rapid 7
Service Configuration	Silver
State	Do not modify the default value
Enforce	Do not modify the default value

- g Do not add any firewall or network introspection services and click **Finish**.
- 6 Map SP\_App to SG\_App\_Group.
- 7 Navigate to the canvas view to confirm that the SP\_App has been mapped to SG\_App\_Group.
- Click the Information Security tab.
  - Click the number next to the  icon to see that the SP\_App is mapped.
- 8 Export the SP\_App policy.
- Click the Security Policies tab and then click the **Export Blueprint** () icon.
  - In **Name**, type **Template\_ App\_** and in **Prefix**, type **FromAppArchitect**.
  - Click Next.
  - Select the SP\_App policy and click Next.
  - Review your selections and click Finish.
  - Select the directory on your computer where you want to download the exported file and click Save.

The security policy as well as all the security groups to which this policy has been applied (in our case, the Application security group as well as the three security groups nested within it) are exported.

- 9 In order to demonstrate how the exported policy works, delete the SP\_App policy.
- 10 Now we will restore the Template\_ App\_ DevTest policy that we exported in step 7.
- Click **Actions** and then click the **Import Service Configuration** icon.
  - Select the **FromAppArchitect\_Template\_App** file from your desktop (you saved it in step 7).
  - Click **Next**.

- d The Ready to complete page displays the security policies along with associated objects (security groups on which these have been applied, as well as Endpoint services, firewall rules, and network introspection services) to be imported.
- e Click **Finish**.

The configuration and associated objects are imported to the vCenter inventory and are visible in the canvas view.

Guest Introspection offloads antivirus and anti-malware agent processing to a dedicated secure virtual appliance delivered by VMware partners. Since the secure virtual appliance (unlike a guest virtual machine) doesn't go offline, it can continuously update antivirus signatures thereby giving uninterrupted protection to the virtual machines on the host. Also, new virtual machines (or existing virtual machines that went offline) are immediately protected with the most current antivirus signatures when they come online.

Guest Introspection health status is conveyed by using alarms that show in red on the vCenter Server console. In addition, more status information can be gathered by looking at the event logs.

---

**Important** Your environment must be correctly configured for Guest Introspection security:

- All hosts in a resource pool containing protected virtual machines must be prepared for Guest Introspection so that virtual machines continue to be protected as they are vMotioned from one ESXi host to another within the resource pool. In NSX 6.4.1 and later, virtual machine hardware must be at v9.0 or above for Guest Introspection to support VM protection during migration (vMotion) of VMs from one host to another.
- Virtual machines must have the Guest Introspection thin agent installed to be protected by Guest Introspection security solution. Not all guest operating systems are supported. Virtual machines with non-supported operating systems are not protected by the security solution.

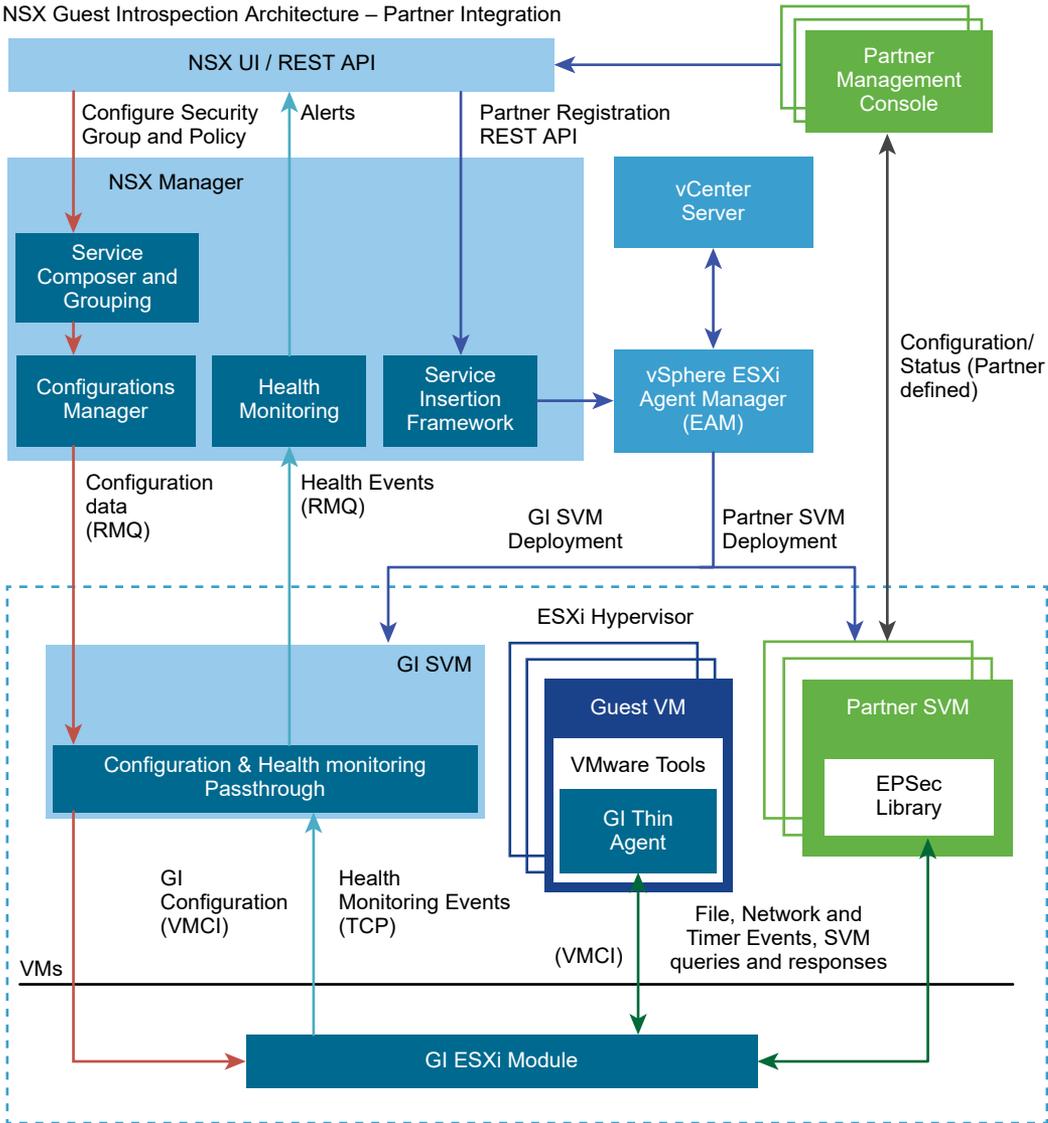
---

This chapter includes the following topics:

- [Guest Introspection Architecture](#)
- [Install Guest Introspection on Host Clusters](#)
- [Install the Guest Introspection Thin Agent on Windows Virtual Machines](#)
- [Install the Guest Introspection Thin Agent on Linux Virtual Machines](#)
- [View Guest Introspection Status](#)
- [Guest Introspection Audit Messages](#)
- [Guest Introspection Events](#)
- [Uninstall a Guest Introspection Module](#)

# Guest Introspection Architecture

NSX Guest Introspection Architecture – Partner Integration



**Legend**

- Configuration data flow
- Health Monitoring data flow
- Partner Registration data flow
- VM-SVM data flow
- Partner Configuration / Status flow

VMCI VMware Virtual Machine Communication Interface  
 RMQ RabbitMQ Message Bus

Guest Introspection is comprised of several related components:

- The partner management console is responsible for registering the service (e.g. agentless anti-virus) with NSX Data Center for vSphere, configuring and monitoring the deployed partner security virtual machines (Partner SVM) and sending VM tagging operations messages to NSX Manager.
- vCenter manages the ESX Agent Manager (EAM) which is responsible for deploying the Partner SVM and Guest Introspection security virtual machine (GI-SVM) to hosts on clusters that have the partner service configured.
- The NSX Manager is the central control for Guest Introspection and provides information to EAM regarding which hosts require a Partner SVM and GI-SVM to be deployed, sends GI configuration information to the GI SVM, receives GI health monitoring information from the host and executes tagging commands received from the Partner Management Console.

On the host, the Partner SVM receives activity events and information from the GI components through the EPSEC library, and performs security operations and analytics to detect potential threats or vulnerabilities. The Partner SVM communicates these events to the Partner Management Console to take NSX Data Center for vSphere actions, such as grouping and tagging. The GI ESX module in the hypervisor acts like a switch to pass relevant events from the thin agents installed on VM's, to the appropriate Partner SVM for analysis. The GI SVM uses configuration information received from NSX Manager to configure the GI ESX Module appropriately as VM's are instantiated or moved, generate Identity Firewall and Endpoint Monitoring context, and send GI-related health information back to NSX Manager.

Common questions asked about SVMs and GI SVMs:

Is there a difference between an SVM and a GI USVM? SVMs refer to third party (partners), such as Trend, and McAfee. USVM is GI SVM.

What are the key characteristics of SVMs and GI SVMs that make them different from a regular VM? Guest introspection offloads anti-virus and anti-malware agent processing to a dedicated secure virtual appliance. Since the secure virtual appliance (unlike a guest virtual machine) doesn't go offline, it can continuously update anti-virus signatures thereby giving uninterrupted protection to the virtual machines on the host.

The Guest Introspection Universal Service Virtual Machine (GI USVM), provides a framework for third-party anti-virus products to be run on guest virtual machines from the outside, removing the need for anti-virus agents in every virtual machine. SVMs contain specific binaries and applications added by the vendor of the SVM. The GI USVM vendor is NSX Data Center for vSphere.

Can any VM be deployed/managed as a SVM? No, SVMs are prebuilt and provided by the vendor.

Is there a SVM or GI USVM specification for logging? No, there is no specification.

Is there a public guide on SVM/USVM related events? No, the SVM logs are for internal troubleshooting purposes

## Install Guest Introspection on Host Clusters

Installing Guest Introspection automatically installs a new VIB and a service virtual machine on each host in the cluster. Guest Introspection is required for Activity Monitoring, and several third-party security solutions.

---

**Note** You cannot migrate a Service VM (SVM) using vMotion/SvMotion. SVMs must remain on the host on which they were deployed for a correct operation.

---

### Prerequisites

The installation instructions that follow assume that you have the following system:

- A data center with supported versions of vCenter Server and ESXi installed on each host in the cluster.
- Hosts in the cluster where you want to install Guest Introspection have been prepared for NSX. See "Prepare Host Clusters for NSX" in the *NSX Installation Guide*. Guest Introspection cannot be installed on standalone hosts. If you are deploying and managing Guest Introspection for anti-virus offload capability only, you do not need to prepare the hosts for NSX, and the NSX for vShield Endpoint license does not allow it.
- NSX Manager installed and running.
- Ensure the NSX Manager and the prepared hosts that run Guest Introspection services are linked to the same NTP server and that time is synchronized. Failure to do so might cause VMs to be unprotected by anti-virus services, although the status of the cluster will be shown as green for Guest Introspection and any third-party services.

If an NTP server is added, VMware recommends that you then redeploy Guest Introspection and any third-party services.

- If your network contains vSphere 7.0 or later, ensure that the vCenter clusters do not use a vSphere Lifecycle Manager (vLCM) image to manage ESXi host life-cycle operations. Guest introspection service cannot be installed on vCenter clusters that use a vLCM image.

To verify whether a vLCM image is used to manage hosts in the cluster, log in to the vSphere Client and go to **Hosts and Clusters**. In the navigation pane, click the cluster, and navigate to **Updates > Image**. If a vLCM image is not used for the cluster, you must see the **SetUp Image** button. If a vLCM image is used for the cluster, you can view the image details, such as ESXi version, vendor add-ons, image compliance details, and so on.

If you want to assign an IP address to the Guest Introspection service virtual machine from an IP pool, create the IP pool before installing Guest Introspection. See "Working with IP Pools" in the *NSX Administration Guide*.

---

**Caution** Guest Introspection uses the 169.254.x.x subnet to assign IP addresses internally for the GI service. If you assign the 169.254.1.1 IP address to any VMkernel interface of an ESXi host, the Guest Introspection installation will fail. The GI service uses this IP address for internal communication.

---

vSphere Fault Tolerance does not work with Guest Introspection.

Guest Introspection is not supported with vSphere Auto Deploy on stateless ESXi hosts.

#### Procedure

- 1 Navigate to **Networking & Security > Installation and Upgrade > Service Deployment**.
- 2 Click **Add**.
- 3 In the Deploy Network and Security Services dialog box, select **Guest Introspection**.
- 4 In **Specify schedule** (at the bottom of the dialog box), select **Deploy now** to deploy Guest Introspection immediately after it is installed or select a deployment date and time.
- 5 Click **Next**.
- 6 Select the datacenter and clusters where you want to install Guest Introspection, and click **Next**.
- 7 On the Select storage and Management Network Page, select the datastore on which to add the service virtual machines storage or select **Specified on host**. It is recommended that you use shared datastores and networks instead of "specified on host" so that deployment workflows are automated.

The selected datastore must be available on all hosts in the selected cluster.

If you selected **Specified on host**, complete the following substeps for each host in the cluster.

- a On the Home page, click **Hosts and Clusters**.
  - b Click a host in the **Navigator**, and then click **Configure**.
  - c In the left navigation pane, under **Virtual Machines** click **Agent VMs**, and then click **Edit**.
  - d Select the datastore and click **OK**.
- 8 If you set datastore as **Specified on host**, you must set the network also as **Specified on host**.
- If you selected **Specified on host**, follow the substeps in Step 7 to select a network on the host. When you add a host (or multiple hosts) to the cluster, the datastore and network must be set before each host is added to the cluster.

- 9 In IP assignment, select one of the following:

Select	To
DHCP	Assign an IP address to the Guest Introspection service virtual machine through Dynamic Host Configuration Protocol (DHCP). Select this option if your hosts are on different subnets.
Use IP Pool	Assign an IP address to the Guest Introspection service virtual machine from the selected IP pool.

- 10 Click **Next** and then click **Finish** on the Ready to complete page.

- 11 Monitor the deployment until the **Installation Status** column displays **Succeeded**.

In NSX 6.4.0 and later, the name of the GI SVM in vCenter Server displays the IP address of the host that it has been deployed on.

- 12 If the **Installation Status** column displays **Failed**, click the icon next to Failed. All deployment errors are displayed. Click **Resolve** to fix the errors. Sometimes, resolving the errors displays additional errors. Take the required action and click **Resolve** again.

**Caution** In a network that contains vSphere 7.0 or later, after the Guest Introspection service or any other third-party partner service is installed, you cannot use a vLCM image on the vCenter clusters. If you try to use a vLCM image on the vCenter clusters, warning messages are displayed in the vSphere Client to inform you that standalone VIBs are present on the hosts.

## 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.

- If you are using vSphere 6.0, see these instructions for installing VMware Tools: [http://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.vm\\_admin.doc%2FGUID-391BE4BF-89A9-4DC3-85E7-3D45F5124BC7.html](http://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.vm_admin.doc%2FGUID-391BE4BF-89A9-4DC3-85E7-3D45F5124BC7.html).
- If you are using vSphere 6.5 or later, see these instructions for installing VMware Tools: <https://www.vmware.com/support/pubs/vmware-tools-pubs.html>.

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.

For Linux instructions, see [Install the Guest Introspection Thin Agent on Linux Virtual Machines](#).

## 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) -- from 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) -- from vSphere 6.0 and later

## 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 will vary depending on the version of VMware Tools.

Driver	Description
vShield Endpoint Drivers	Installs File Introspection (vsepfilt) and Network Introspection (vnetflt) drivers.
Guest Introspection Drivers	Installs File Introspection (vsepfilt) and Network Introspection (vnetflt) drivers.
NSX File Introspection Driver and NSX Network Introspection Driver	Select NSX File Introspection Driver to install vsepfilt. Optionally select NSX Network Introspection Driver to install vnetflt (vnetWFP on Windows 10 or later).  <b>Note</b> Select NSX Network Introspection Driver only if you are using the Identity Firewall or Endpoint Monitoring features.

- 3 In the drop-down menu next to the drivers you want to add, select **This feature will be installed on the local hard drive**.
- 4 Follow the remaining steps in the procedure.

## What to do next

Check if 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 `vseflt`.

# 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 GI thin agent is available as part of the VMware Tools operating system-specific packages (OSPs). Installing VMware Tools is not required. GI thin agent installation and upgrade is not connected to NSX installation and upgrade. Also, Enterprise or Security Administrator (non-NSX Administrator) can install the agent on guest VMs outside of NSX.

To install the GI thin agent on RHEL, CentOS, and SLES Linux systems, use the *RPM* package. To install the GI thin agent on Ubuntu Linux systems, use the *DEB* package.

For Windows instructions, see [Install the Guest Introspection Thin Agent on Windows Virtual Machines](#).

## Prerequisites

- Ensure that the guest virtual machine has a supported version of Linux installed:
  - Red Hat Enterprise Linux (RHEL) 7.0–7.4 GA (64 bit).
  - CentOS 7.4 GA.
  - SUSE Linux Enterprise Server (SLES) 12 GA (64 bit).
  - Ubuntu 16.04.5 LTS GA (64 bit).
  - Ubuntu 14.04 LTS GA (64 bit).

---

**Note** Starting in NSX 6.4.6, support for Ubuntu 14.04 and RHEL 7.0–7.3 is deprecated. NSX 6.4.6 and later supports Ubuntu 16.04.5 and RHEL 7.4.

---

- Verify that GLib 2.0 is installed on the Linux VM.

## Procedure

- ◆ Based on your Linux operating system, perform the following steps with a root privilege:
  - 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 *vm.list* file under `/etc/apt/sources.list.d`.
- c Edit the file with the following content:

```
deb https://packages.vmware.com/packages/nsx-gi/latest/ubuntu/dists xenial main
```

- d Install the package:

```
apt-get update
apt-get install vmware-nsx-gi-file
```

- 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 *vm.repo* file under `/etc/yum.repos.d`.
- c Edit the file with the following content:

```
[vmware]
name = VMware
baseurl = https://packages.vmware.com/packages/nsx-gi/latest/rhel/x86_64
enabled = 1
gpgcheck = 1
metadata_expire = 86400
ui_repoid_vars = basearch
```

- d Install the package:

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

- 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/sles12/x86_64/"
VMware
```

- c Install the package:

```
zypper install vmware-nsx-gi-file
```

- 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

Check if the thin agent is running using the `service vsepd status` command with the administrative privileges. The status should be running.

## View Guest Introspection Status

Monitoring a Guest Introspection instance involves checking for status coming from the Guest Introspection components: the security virtual machine (SVM), the ESXi host-resident Guest Introspection module, and the protected virtual machine-resident thin agent.

### Procedure

- 1 In the vSphere Web Client, click **vCenter Inventory Lists**, and then click **Datacenters**.
- 2 In the **Name** column, click a data center.
- 3 Click **Monitor** and then click **Guest Introspection**.

The Guest Introspection Health and Alarms page displays the health of the objects under the data center you selected, and the active alarms. Health status changes are reflected within a minute of the actual occurrence of the event that triggered the change.

## Guest Introspection Audit Messages

Audit messages include fatal errors and other important audit messages and are logged to `vmware.log`.

The following conditions are logged as AUDIT messages:

- Thin agent initialization success (and version number.)

- Thin agent initialization failure.
- Established first time communication with SVM.
- Failure to establish communication with SVM (when first such failure occurs).

Generated log messages have the following substrings near the beginning of each log message: `vf-AUDIT`, `vf-ERROR`, `vf-WARN`, `vf-INFO`, `vf-DEBUG`.

## Guest Introspection Events

Events are used for logging and auditing conditions inside the Guest Introspection-based security system.

Events can be displayed without a custom vSphere plug-in. See the *vCenter Server Administration Guide* on events and alarms.

Events are the basis for alarms that are generated. Upon registering as a vCenter Server extension, the NSX Manager defines the rules that create and remove alarms.

Common arguments for all events are the event time stamp and the NSX Manager `event_id`.

Event Code	Event Severity	Alarm Triggered	Event Message	Description
26000	Informational			The Guest Introspection thin agent running in a guest VM is enabled. This event is reported when a VM is added to the VC inventory or when either the VM or NSX Manager is restarted. Action: Information-only event.
26001	Informational			The Guest Introspection service deployment process has succeeded, and the Service Status of the GI SVM (USVM) is up. Action: Information-only event.
26002	Medium			The thin agent version in the guest VM is not compatible with the EPsec library version in the partner SVM or the USVM. Action: Update the thin agent version in the guest VM.
26003	Medium			The ESX GI module (MUX) failed to establish a connection with the USVM. This event may be reported after the guest VM attempted to connect to the USVM to send health status reports. Action: Confirm the GI SVM (USVM) is powered on, and service status appears as up in the Service Deployment tab.

Event Code	Event Severity	Alarm Triggered	Event Message	Description
26004	Medium			The ESX GI module (MUX) failed to establish a connection with the USVM. This event may be reported after the guest VM attempted to connect to the USVM to send health status reports.  Action: Confirm the GI SVM (USVM) is powered on, and service status appears as up in the Service Deployment tab.
26005	Informational			The ESX GI module (MUX) was installed successfully.  Action: Information-only event for each host on which Guest Introspection is deployed.
26006	Informational			The ESX GI module (MUX) was uninstalled.  Action: Information-only event for each host on which the Guest Introspection host module is no longer running.
26007	Informational			NSX Manager failed to receive the host module's health status report for up to three minutes. The Service Status column for the cluster in the Service Deployment tab will be in a warning state.  Action: This event may be a transient event during service deployment. If it persists, collect the GI SVM and host technical support logs (vmware.log) and open a technical support request.

## Uninstall a Guest Introspection Module

Uninstalling guest introspection removes a VIB from the hosts in the cluster and removes the service virtual machine from each host in the cluster. Guest Introspection is required for Identity Firewall, Endpoint Monitoring, and several third-party security solutions. Uninstalling guest introspection can have wide ranging impacts.

**Caution** Before you uninstall a Guest Introspection module from a cluster, you must uninstall all third-party products that are using Guest Introspection from the hosts on that cluster. Use the instructions from the solution provider.

There is a loss of protection for VMs in the host cluster. You must vMotion the VMs out of the cluster before you uninstall.

To uninstall Guest Introspection:

- 1 Navigate to **Networking & Security > Installation and Upgrade > Service Deployment**.
- 2 Select a Guest Introspection instance and click the delete icon.
- 3 Either delete now or schedule the deletion for a later time.

## Uninstall Guest Introspection for Linux

You can uninstall Linux thin agent for Guest Introspection from guest virtual machine.

### Prerequisites

Guest Introspection for Linux is installed. You have root privileges on the Linux system.

### Procedure

- ◆ For uninstalling package from Ubuntu system, run the `apt-get remove vmware-nsx-gi-file` command.
- ◆ For uninstalling package from RHEL7 system, run the `yum remove vmware-nsx-gi-file` command.
- ◆ For uninstalling package from SLES system, run the `zypper remove vmware-nsx-gi-file` command.

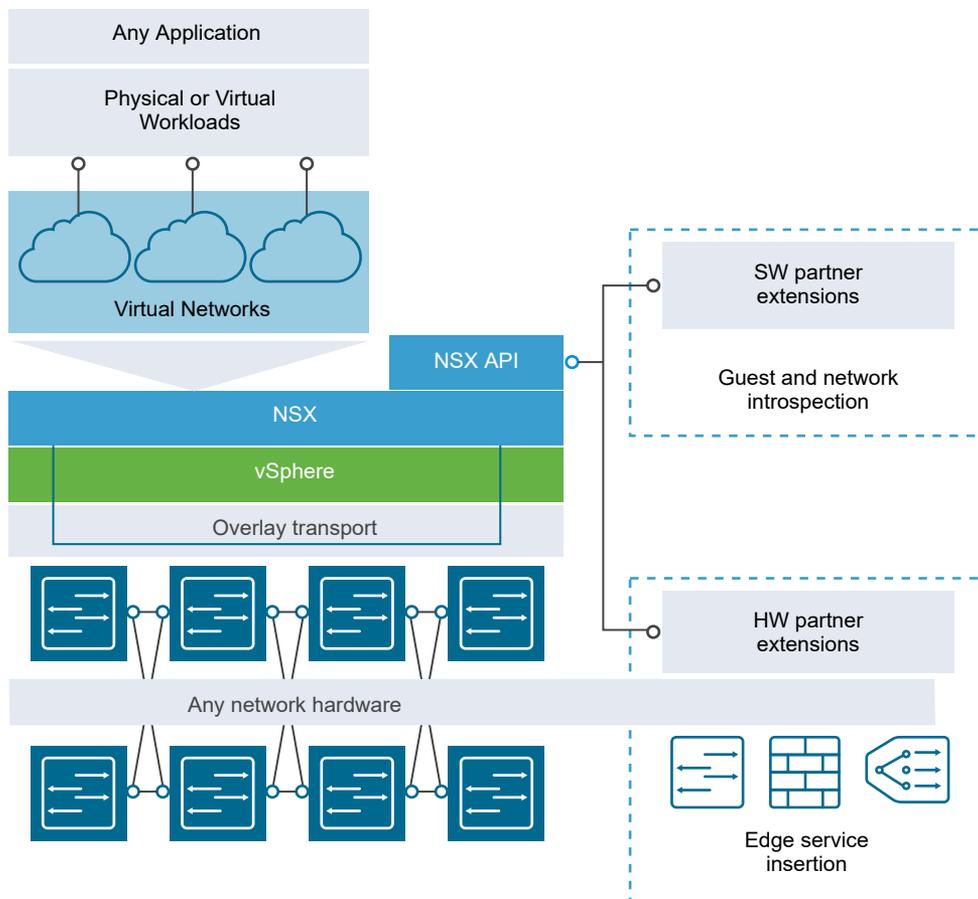
### Results

The thin agent installed on Linux virtual machine is uninstalled.

# Network Extensibility

# 20

Datacenter networks typically involve a wide range of network services, including switching, routing, firewalling, load balancing, and so on. In most cases, these services are delivered by different vendors. In the physical world, connecting these services in the network is a complicated exercise of racking and stacking physical network devices, establishing physical connectivity, and managing these services separately. NSX simplifies the experience of connecting the right services in the right traffic paths and can help you build complex networks within a single ESXi host or across multiple ESXi hosts for production, testing, or development purposes.



There are various deployment methods for inserting third party services into NSX.

This chapter includes the following topics:

- [Distributed Service Insertion](#)
- [Edge-Based Service Insertion](#)
- [Integrating Third Party Services](#)
- [Deploy a Partner Service](#)
- [Consuming Vendor Services through Service Composer](#)
- [Redirecting Traffic to a Vendor Solution through Logical Firewall](#)
- [Using a Partner Load Balancer](#)
- [Remove Third-Party Integration](#)

## Distributed Service Insertion

In distributed service insertion, a single host has all service modules, kernel modules, and virtual machine implementations on a single physical machine. All components of the system interact with components within the physical host. This allows for faster module-to-module communication and compact deployment models. The same configuration can be replicated on physical systems in the network for scalability, while control and data plane traffic to and from the service modules to the vmkernel stay on the same physical system. During vMotion of the protected virtual machines, the partner security machine moves the virtual machine state from the source to the destination host.

Vendor solutions that make use of this type of service insertion include Intrusion Prevention Service (IPS)/Intrusion Detection Service (IDS), Firewall, Anti Virus, File Identity Monitoring (FIM), and Vulnerability Management.

## Edge-Based Service Insertion

NSX Edge is deployed as a virtual machine in the Edge Services Cluster along with other network services. NSX Edge has the capability to redirect specific traffic to 3rd-party network services..

Vendor solutions that make use of this type of service insertion include ADC/Load Balancer devices.

## Integrating Third Party Services

This is a generic high-level workflow for inserting a third-party service into the NSX platform.

### Procedure

- 1 Register the third-party service with NSX Manager on the vendor's console.

You need NSX login credentials to register the service. For more information, refer to the vendor documentation.

- 2 Deploy the service in NSX. See [Deploy a Partner Service](#) .

Once deployed, the third-party service is displayed in the NSX Service Definitions window and is ready to be used. The procedure for using the service in NSX depends on the type of service inserted.

For example, you can enable a host-based firewall service by creating a security policy in Service Composer or creating a firewall rule to redirect traffic to the service. See [Consuming Vendor Services through Service Composer](#) or [Redirecting Traffic to a Vendor Solution through Logical Firewall](#). For information on using an Edge based service, see [Using a Partner Load Balancer](#).

## Deploy a Partner Service

If the partner solution includes a host-resident virtual appliance, you can deploy the service after the solution is registered with NSX Manager.

---

**Important** Guest VMs protected by a partner service temporarily lose protection if migrated to another cluster using vMotion. To avoid this, vMotion guest VMs only to hosts within the same cluster.

---

### Prerequisites

Ensure that:

- The partner solution is registered with NSX Manager.
- NSX Manager can access the partner solution's management console.
- The vCenter clusters in vSphere 7.0 or later do not use a vSphere Lifecycle Manager (vLCM) image to manage ESXi host life-cycle operations. Partner services cannot be installed on vCenter clusters that use a vLCM image.

To verify whether a vLCM image is used to manage hosts in the cluster, log in to the vSphere Client and go to **Hosts and Clusters**. In the navigation pane, click the cluster, and navigate to **Updates > Image**. If a vLCM image is not used for the cluster, you must see the **SetUp Image** button. If a vLCM image is used for the cluster, you can view the image details, such as ESXi version, vendor add-ons, image compliance details, and so on.

- The required license edition has been assigned. See <https://kb.vmware.com/kb/2145269>.

### Procedure

- 1 Navigate to **Networking & Security > Installation and Upgrade > Service Deployment**.
- 2 Click **Add**.
- 3 In the Deploy Network and Security Services dialog box, select **Guest Introspection**.
- 4 In **Specify schedule** (at the bottom of the dialog box), select **Deploy now** to deploy Guest Introspection immediately after it is installed or select a deployment date and time.

- 5 Click **Next**.
- 6 Select the datacenter and clusters where you want to install Guest Introspection, and click **Next**.
- 7 On the Select storage and Management Network Page, select the datastore on which to add the service virtual machines storage or select **Specified on host**. It is recommended that you use shared datastores and networks instead of "specified on host" so that deployment workflows are automated.

The selected datastore must be available on all hosts in the selected cluster.

If you selected **Specified on host**, complete the following substeps for each host in the cluster.

- a On the Home page, click **Hosts and Clusters**.
  - b Click a host in the **Navigator**, and then click **Configure**.
  - c In the left navigation pane, under **Virtual Machines** click **Agent VMs**, and then click **Edit**.
  - d Select the datastore and click **OK**.
- 8 If you set datastore as **Specified on host**, you must set the network also as **Specified on host**.  
If you selected **Specified on host**, follow the substeps in Step 7 to select a network on the host. When you add a host (or multiple hosts) to the cluster, the datastore and network must be set before each host is added to the cluster.
  - 9 In IP assignment, select one of the following:

Select	To
DHCP	Assign an IP address to the Guest Introspection service virtual machine through Dynamic Host Configuration Protocol (DHCP). Select this option if your hosts are on different subnets.
Use IP Pool	Assign an IP address to the Guest Introspection service virtual machine from the selected IP pool.

- 10 Click **Next** and then click **Finish** on the Ready to complete page.
- 11 Monitor the deployment until the **Installation Status** column displays **Succeeded**.  
In NSX 6.4.0 and later, the name of the GI SVM in vCenter Server displays the IP address of the host that it has been deployed on.
- 12 If the **Installation Status** column displays **Failed**, click the icon next to Failed. All deployment errors are displayed. Click **Resolve** to fix the errors. Sometimes, resolving the errors displays additional errors. Take the required action and click **Resolve** again.

---

**Caution** In a network that contains vSphere 7.0 or later, after the Guest Introspection service or any other third-party partner service is installed, you cannot use a vLCM image on the vCenter clusters. If you try to use a vLCM image on the vCenter clusters, warning messages are displayed in the vSphere Client to inform you that standalone VIBs are present on the hosts.

---

## What to do next

You can now consume the partner service through NSX UI or NSX API.

# Consuming Vendor Services through Service Composer

Third-party vendor services include traffic redirection, load balancer, and guest security services such as data loss prevention, anti virus, and so on. Service Composer enables you to apply these services to a set of vCenter objects.

A security group is a set of vCenter objects such as clusters, virtual machines, vNICs, and logical switches. A security policy is a set of Guest Introspection services, firewall rules, and network introspection services.

When you map a security policy to a security group, redirection rules are created on the appropriate third-party vendor service profile. As traffic flows from virtual machines belonging to that security group, it is redirected to registered third-party vendor services that determine how to process that traffic. For more information on Service Composer, see [Using Service Composer](#).

# Redirecting Traffic to a Vendor Solution through Logical Firewall

You can add firewall rules to redirect traffic to registered vendor solutions. Redirected traffic is then processed by the vendor service.

## Prerequisites

- The third party service must be registered with NSX Manager, and the service must be deployed in NSX.
- If the default firewall rule action is set to Block, you must add a rule to allow the traffic to be redirected.

## Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Firewall**.
- 2 Click the **Partner security services** tab.
- 3 In the section to which you want to add a rule, click the **Add rule** () icon.  
A new any any allow rule is added at the top of the section.
- 4 Point to the **Name** cell of the new rule, click , and type a name for the rule.
- 5 Specify the **Source**, **Destination**, and **Service** for the rule. For more information, see [Add a Firewall Rule](#)

- 6 Point to the **Action** cell of the new rule, and click .
  - a In **Action**, select **Redirect**.
  - b In **Redirect To**, select the service profile and the logical switch or security group to which you want to bind the service profile.

The service profile is applied to virtual machines connected to or contained in the selected logical switch or security group.
  - c Indicate whether the redirected traffic is to be logged and type comments, if any.
  - d Click **OK**.

The selected service profile is displayed as a link in the **Action** column. Clicking the service profile link displays the service profile bindings.
- 7 Click **Publish Changes**.

## Using a Partner Load Balancer

You can use a third-party load balancer to balance the traffic for a specific NSX Edge.

### Prerequisites

The third-party load balancer must be registered with NSX Manager, and it must be deployed in NSX.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > NSX Edges**.
- 2 Double-click an NSX Edge.
- 3 Click **Manage** and then click the **Load Balancer** tab.
- 4 Click **Edit** next to Load balancer global configuration.
- 5 Select **Enable Load Balancer** and **Enable Service Insertion**.
- 6 In **Service Definition**, select the appropriate partner load balancer.
- 7 In **Service Configuration**, select the appropriate service configuration.
- 8 Complete the remaining fields and set up the load balancer by adding a service monitor, server pool, application profile, application rules, and a virtual server. When adding a virtual server, select the template provided by the vendor. For more information, see [Setting Up Load Balancing](#).

### Results

Traffic for the specified Edge is load balanced by the third party vendor's management console.

## Remove Third-Party Integration

This example describes how to remove a third-party integration solution from NSX.

There is a correct order of software when removing any third-party software solution. If this sequence is not followed and specifically if the third-party solution is uninstalled or deleted before it is unregistered with NSX Manager, the removal operation will fail. See <https://kb.vmware.com/kb/2126678> for instructions on how to resolve this.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Service Composer**, and delete the rules (or security policies) that are redirecting traffic to the 3rd-party solution.
- 2 Navigate to **Service Definitions** and double-click the name of the third-party solution.
- 3 Click **Related Objects** and delete the related objects.
- 4 Navigate to **Installation and Upgrade > Service Deployments** and delete the third-party deployment.

This action uninstalls the associated VMs.

- 5 Return to **Service Definitions** and delete any sub-components of the definition.
- 6 In the service instance, delete the service profile.
- 7 Delete the service instance.
- 8 Delete the service definition.

### Results

The third-party integration solution is removed from NSX.

### What to do next

Make notes of the configuration settings, and then remove NSX from the third-party solution. For example, you may need to delete rules that reference other objects and then delete the objects.

In many organizations, networking and security operations are handled by different teams or members. Such organizations may require a way to limit certain operations to specific users. This topic describes the options provided by NSX to configure such access control.

NSX also supports Single Sign On (SSO), which enables NSX to authenticate users from other identity services such as Active Directory, NIS, and LDAP.

User management in the vSphere Web Client is separate from user management in the CLI of any NSX component.

This chapter includes the following topics:

- [NSX Users and Permissions by Feature](#)
- [Configure Single Sign-On](#)
- [Managing User Rights](#)
- [Managing the Default User Account](#)
- [Assign a Role to a vCenter User](#)
- [Group-Based Role Assignments](#)
- [Create a User with Web Interface Access Using CLI](#)
- [Edit a User Account](#)
- [Change a User Role](#)
- [Disable or Enable a User Account](#)
- [Delete a User Account](#)

## NSX Users and Permissions by Feature

To deploy and administer NSX Data Center for vSphere, certain vCenter permissions are required. NSX Data Center for vSphere provides extensive read and read/write permissions for various users and roles.

## Feature List with Roles and Permissions

### Note

- Security Engineer and Network Engineer roles are available in NSX 6.4.2 and later.
- Security & Role Administrator role is available in NSX 6.4.5 and later.

Feature	Description	Roles						
		Audit or	Securi ty Admin	Securi ty Engin eer	NSX Admin	Netw ork Engin eer	Securit y & Role Admin	Enterprise Admin
<b>Administrator</b>								
Configuration	vCenter and SSO Configuration with NSX	R	R	R	R	R	R	R, W
Update		No Acces s	No Acces s	No Acces s	R, W	R, W	No Access	R, W
System events	System Events	R	R, W	R, W	R, W	R, W	R, W	R, W
Audit Logs	Audit Logs	R	R	R	R	R	R	R
Debug		No Acces s	No Acces s	No Acces s	No Access	No Acces s	No Access	No Access
Housekeeping tasks		R	R	R	R, W	R, W	R	R, W
Basic auth disable		R	R	R	R	R	R	R, W
<b>User Account Management (URM)</b>								
User account management	User Management	R	No Acces s	No Acces s	R	R	R, W	R, W
Object access control		No Acces s	No Acces s	No Acces s	R	R	R	R
Feature access control		No Acces s	No Acces s	No Acces s	R	R	R	R
<b>Edge</b>								
System	System refers to general system parameters	R	R	R	R, W	R, W	R	R, W

Feature	Description	Roles						
		Audit or	Securi ty Admin	Securi ty Engin eer	NSX Admin	Netw ork Engin eer	Securit y & Role Admin	Enterprise Admin
Advanced services		R	R, W	R, W	R	R	R, W	R, W
Appliance	Different form factors of NSX Edge (Compact / Large/X-Large/ QuadLarge)	R	R	R	R, W	R, W	R	R, W
High availability		R	R	R	R, W	R, W	R	R, W
vNic	Interface configuration on NSX Edge	R	R, W	R	R, W	R, W	R	R, W
DNS		R	R, W	R	R	R, W	R	R, W
SSH	SSH configuration on NSX Edge	R	R, W	R	R, W	R, W	R	R, W
Auto plumbing		R	R, W	R, W	R	R	R, W	R, W
Statistics		R	R	R	R	R	R	R, W
NAT	NAT configuration on NSX Edge	R	R, W	R	R	R, W	R	R, W
DHCP		R	R, W	R	R	R, W	R	R, W
Load balance		R	R, W	R	R	R, W	R	R, W
L3 VPN	L3 VPN	R	R, W	R	R	R, W	R	R, W
VPN	L2 VPN, SSL VPN	R	R, W	R	R	R, W	R	R, W
Syslog	Syslog configuration on NSX Edge	R	R, W	R	R, W	R, W	R	R, W
Support Bundle		R (Down load access )	R, W	R, W	R, W	R, W	R, W	R, W

Feature	Description	Roles						
		Auditor	Security Admin	Security Engineer	NSX Admin	Network Engineer	Security & Role Admin	Enterprise Admin
Routing	All routing static and dynamic routing (BGP/OSPF) on NSX Edge	R	R, W	R	R	R, W	R	R, W
Firewall	Firewall configuration on NSX Edge	R	R, W	R, W	R	R	R, W	R, W
Bridging		R	R, W	R	R	R, W	R	R, W
Certificate		R	R, W	R, W	R	R	R, W	R, W
System control	System control refers to system kernel parameters such as maximum limits, IP forwarding, networking, and system settings. For example: ysctl.net.ipv4.conf.vNic_1.rp_filter sysctl.net.netfilter.nf_conntrack_tcp_timeout_established	R	R, W	R, W	R, W	R, W	R, W	R, W
<b>Distributed Firewall</b>								
Firewall config	<ul style="list-style-type: none"> <li>■ Layer 3 - 7 (General) firewall rules</li> <li>■ Layer 2 (Ethernet) firewall rules</li> </ul>	R	R, W	R, W	R, W	No Access	R, W	R, W
Flows	Flow monitoring is for monitoring traffic flows in the system. Live Flows can also be monitored	R	R, W	R, W	No Access	R, W	R, W	R, W

Feature	Description	Roles						
		Auditor	Security Admin	Security Engineer	NSX Admin	Network Engineer	Security & Role Admin	Enterprise Admin
IPFix config	IPFix enable/disable and assigning collectors	R	R, W	R, W	No Access	R, W	R, W	R, W
ForceSync	ForceSync does full sync from the <b>Installation and Upgrade &gt; Host Preparation</b> page	R	R	R, W	R, W	No Access	R, W	R, W
Install DFW (host preparation)	Install VIBs on clusters	R	R	R	R, W	R, W	R	R, W
Saved configurations (drafts)	Every publish will automatically save existing DFW configuration as a draft	R	R, W	R, W	No Access	No Access	R, W	R, W
Exclusion list	Add VMs to exclusion list to be NOT protected by DFW or to remove them	R	R, W	R, W	No Access	No Access	R, W	R, W
DFW tech support	Collecting DFW Tech Support bundle from a host (only NSX config shell)	No Access	R, W	R, W	R, W	No Access	R, W	R, W
DFW session timers	Configure TCP/UDP/Other protocol connection timeout configuration	R	R, W	R, W	No Access	No Access	R, W	R, W
IP Discovery (DHCP/ARP Snooping)	IP discovery when VMware Tools are not running on Guest VMs	R	R, W	R, W	R	No Access	R, W	R, W

Feature	Description	Roles						
		Audit or	Securi ty Admin	Securi ty Engin eer	NSX Admin	Netw ork Engin eer	Securit y & Role Admin	Enterprise Admin
Application Rule Manager	Flows are collected for selected set of applications. Firewall rules are then created based on the collected flows.	R	R, W	R, W	No Access	No Access	R, W	R, W
app.syslog		R	R	No Access	R, W	No Access	No Access	R, W
Packet capture		R	R, W	R, W	R, W	R, W	R, W	R, W
<b>NameSpace</b>								
Config		R	R	R	R, W	R,W	R	R, W
<b>SpoofGuard</b>								
Config	SpoofGuard publish in TOFU or Manual Mode	R	R, W	R, W	No Access	No Access	R, W	R, W
<b>Endpoint Security (EPSEC)</b>								
Reports		R	R	R	R, W	R	R	R, W
Registration	Manage [Register, Unregister, Query registered solutions, Activate] Solutions	R	No Access	No Access	R, W	R, W	No Access	R, W
Health monitoring	Retrieve health status of VM, SVM to the NSX Manager	No Access	R	R	R	R	R	R
Policy	Manage security policies [Create, Read, Update, Delete]	R	R, W	R, W	R, W	R	R, W	R, W

Feature	Description	Roles						
		Auditor	Security Admin	Security Engineer	NSX Admin	Network Engineer	Security & Role Admin	Enterprise Admin
Scan scheduling		R	No Access	R, W	R, W	R	R, W	R, W
<b>Library</b>								
Host preparation	Host preparation action on cluster	No Access	No Access	No Access	R, W	R, W	No Access	R, W
Grouping	IP Set, MAC Set, Security Group, Service, Service Group	R	R, W	R, W	R	R	R, W	R, W
Tagging	Security tag (for example, attach or detach VMs)	R	R, W	R, W	R	R	R, W	R, W
<b>Install</b>								
App		No Access	R	R	R, W	R, W	R	R, W
EPSEC		No Access	R	R	R, W	R, W	R	R, W
DLP		No Access	R	R	R, W	R, W	R	R, W
<b>VDN</b>								
Config NSM	Configure Network Security Manager	R	R	R	R, W	R, W	R	R, W
Provision		R	R	R	R, W	R, W	R	R, W
<b>ESX Agent Manager (EAM)</b>								
Install	ESX Agent Manager	No Access	R	R	R, W	R, W	R	R, W
<b>Service Insertion</b>								

Feature	Description	Roles						
		Audit or	Securi ty Admin	Securi ty Engin eer	NSX Admin	Netw ork Engin eer	Securit y & Role Admin	Enterprise Admin
Service		R	R, W	R, W	R, W	R	R, W	R, W
Service profile		R	R	R, W	R, W	R	R, W	R, W
<b>Trust Store</b>								
trustentity_ma nagement	NSX certificate management	R	R, W	R, W	R, W	R, W	R, W	R, W
<b>IP Address Management (IPAM)</b>								
Configuration	Configuration of IP pool	R	R, W	R	R, W	R, W	R	R, W
IP allocation	IP allocation and release	R	R, W	R	R, W	R, W	R	R, W
<b>Security Fabric</b>								
Deploy	Deploy service or security VM on cluster using the <b>Service Deployment</b> page	R	R	R	R, W	R	R	R, W
Alarms	From the <b>Service Deployment</b> page, manage alarms that are generated by security VM	R	R, W	R	R, W	R, W	R	R, W
Agent health status	Managing agent health status alarm over rest call, mainly used by partner VMs	R	R, W	R, W	R, W	R, W	R, W	R, W
<b>Messaging</b>								
Messaging	Messaging framework used by NSX Edge and Guest Introspection to communicate with NSX Manager	R	R, W	R, W	R, W	R, W	R, W	R, W

Feature	Description	Roles						
		Audit or	Security Admin	Security Engineer	NSX Admin	Network Engineer	Security & Role Admin	Enterprise Admin
<b>Replicator (Multi vCenter setup with secondary NSX Manager)</b>								
Configuration	Select or deselect Primary role for NSX Manager, and add or remove Secondary NSX Manager	R	R	R	R, W	R, W	R	R, W
<b>blueprint_sam.featurelist</b>								
blueprint_sam.ad_config	Used for Active Directory domain configuration	R	R	R	R, W	R, W	R	R, W
<b>Security Policy</b>								
Configuration	Configure security policy to create, update, edit, or delete	R	R, W	R, W	No Access	No Access	R, W	R, W
Security group binding	Associate security group with a security policy	R	R, W	R, W	No Access	No Access	R, W	R, W
Apply policy		R	R, W	R, W	No Access	No Access	R, W	R, W
Policy sync	Sync security policy with DFW	R	Can Sync	Can Sync	No Access	No Access	Can Sync	R, W
<b>NSX Appliance Management (In NSX 6.4 and later)</b>								
NSX Appliance Management	NSX Appliance Management	R	R	R	R	R	R	R, W
<b>IP Repository/IP Discovery</b>								
Configuration		R	R, W	R, W	R	No Access	R, W	R, W

Feature	Description	Roles						
		Audit or	Security Admin	Security Engineer	NSX Admin	Network Engineer	Security & Role Admin	Enterprise Admin
<b>Dashboard</b>								
Widget configuration		R	R, W	R	R, W	R	R	R, W
System configuration		R	R, W	R	R, W	R	R	R, W
<b>Upgrade Coordinator</b>								
Upgrade		No Access	No Access	R	R, W	R	R	R, W
Upgrade Plan		R	R	R	R, W	R	R	R, W
<b>Tech Support Bundle</b>								
Config	Endpoint	R, W	R, W	R, W	R, W	R, W	R, W	R, W
<b>Token Based Authentication</b>								
Invalidation		No Access	No Access	No Access	No Access	No Access	No Access	R, W
<b>Ops</b>								
Config		R	R	R	R, W	R	R	R, W

## Configure Single Sign-On

SSO makes vSphere and NSX Data Center for vSphere more secure by allowing the various components to communicate with each other through a secure token exchange mechanism, instead of requiring each component to authenticate a user separately.

You can configure lookup service on the NSX Manager and provide the SSO administrator credentials to register NSX Management Service as an SSO user. Integrating the single sign-on (SSO) service with NSX Data Center for vSphere improves the security of user authentication for vCenter users and enables NSX Data Center for vSphere to authenticate users from other identity services such as AD, NIS, and LDAP. With SSO, NSX Data Center for vSphere supports authentication using authenticated Security Assertion Markup Language (SAML) tokens from a trusted source using REST API calls. NSX Manager can also acquire authentication SAML tokens for use with other VMware solutions.

NSX Data Center for vSphere caches group information for SSO users. Changes to group memberships take up to 60 minutes to propagate from the identity provider (for example, active directory) to NSX Data Center for vSphere.

### Prerequisites

- To use SSO on NSX Manager, you must have vCenter Server 6.0 or later, and single sign-on (SSO) authentication service must be installed on the vCenter Server. Note that this is for embedded SSO. Instead, your deployment might use an external centralized SSO server.

For information about SSO services provided by vSphere, see the *Platform Services Controller Administration* documentation.

---

**Important** You must configure the NSX Manager appliance to use the same SSO configuration that is used on the associated vCenter Server system.

---

- NTP server must be specified so that the SSO server time and NSX Manager time are in sync.

For example:

Time Settings	
NTP Server	192.168.110.10
Timezone	UTC
Date/Time	12/28/2016 21:31:49

### Procedure

- 1 Log in to the NSX Manager virtual appliance.

In a Web browser, navigate to the NSX Manager appliance GUI at <https://<nsx-manager-ip>> or <https://<nsx-manager-hostname>>, and log in as **admin** or with an account that has the **Enterprise Administrator** role.

- 2 Log in to the NSX Manager virtual appliance.
- 3 From the home page, click **Manage Appliance Settings > NSX Management Service**.
- 4 Click **Edit** in the Lookup Service URL section.
- 5 Enter the name or IP address of the host that has the lookup service.
- 6 Enter the port number.

If you are using vSphere 6.0 or later, enter port 443.

The Lookup Service URL is displayed based on the specified host and port.

- 7 Enter the SSO Administrator user name and password, and click **OK**.

The certificate thumbprint of the SSO server is displayed.

- 8 Check that the certificate thumbprint matches the certificate of the SSO server.

If you installed a CA-signed certificate on the CA server, you are presented with the thumbprint of the CA-signed certificate. Otherwise, you are presented with a self-signed certificate.

- 9 Confirm that the Lookup Service status is **Connected**.

#### What to do next

See "Assign a Role to a vCenter User", in the *NSX Administration Guide*.

## Managing User Rights

A user's role defines the actions the user is allowed to perform on a given resource. The role determines the user's authorized activities on the given resource, ensuring that a user has access only to the functions necessary to complete applicable operations. This allows domain control over specific resources, or system-wide control if your right has no restrictions.

A user can have only one role. The following table lists the permissions of each user role.

**Table 21-1. NSX Manager User Roles**

Role	Permissions
<b>Enterprise Administrator</b>	Users in this role can perform all tasks related to deployment and configuration of NSX products and administration of this NSX Manager instance.
<b>NSX Administrator</b>	Users in this role can perform all tasks related to deployment and administration of this NSX Manager instance. For example, install virtual appliances, configure port groups.
<b>Security Administrator</b>	Users in this role can configure security compliance policies in addition to viewing the reporting and auditing information in the system. For example, define distributed firewall rules, configure NAT and load balancer services.
<b>Auditor</b>	Users in this role can only view system settings, auditing, events, and reporting information and cannot make any configuration changes.
<b>Security Engineer</b> (introduced in NSX Data Center for vSphere 6.4.2).	Users in this role can perform all security tasks, such as configuring policies, firewall rules. Users have read access to some networking features, but no access to host preparation and user account management.
<b>Network Engineer</b> (introduced in NSX Data Center for vSphere 6.4.2).	Users in this role can perform all networking tasks, such as routing, DHCP, bridging. Users have read access to endpoint security features, but no access to other security features.
<b>Security &amp; Role Administrator</b> (introduced in NSX Data Center for vSphere 6.4.5).	Users in this role have all the feature permissions that a <b>Security Engineer</b> has, and they can also perform user management tasks.

When you assign a role to an SSO user, access is granted in the following interfaces:

- The Networking and Security plug-in in the vSphere Web Client.

- The NSX Manager appliance, including the API. This access is available only in NSX 6.4 or later.

The **Enterprise Administrator** role gets the same access to the NSX Manager appliance and the API as the NSX Manager **admin** user. The other NSX roles get **read-only** access to the NSX Manager appliance and the API.

For example:

SSO users with any role other than the **Enterprise Administrator** role can access the NSX Manager UI and run API requests in **read-only** mode. Users can access NSX APIs with the GET API request, but they cannot run the PUT, POST, and DELETE API requests. In addition, these SSO users cannot perform actions such as stop, configure, edit, and so on, in the NSX Manager UI.

## Managing the Default User Account

The NSX Manager user interface includes a user account, which has access rights to all resources. You cannot edit the rights of or delete this user. The default user name is **admin** and the default password is **default** or the password you specified during NSX Manager installation.

You can manage NSX Manager appliance **admin** user only through CLI commands.

## Assign a Role to a vCenter User

When you assign a role to an SSO user, vCenter Server authenticates the user with the identity service configured on the SSO server. If the SSO server is not configured or is not available, the user is authenticated either locally or with Active Directory based on vCenter Server configuration.

When you assign a role to an SSO user, access is granted in the following interfaces:

- The Networking and Security plug-in in the vSphere Web Client.
- The NSX Manager appliance, including the API. This access is available only in NSX 6.4 or later.

The **Enterprise Administrator** role gets the same access to the NSX Manager appliance and the API as the NSX Manager **admin** user. The other NSX roles get **read-only** access to the NSX Manager appliance and the API.

Roles can be assigned individually or through a group membership. A user can be assigned an NSX role individually, and this user can also be a member of a group that is assigned a different NSX role. In such cases, the role that is assigned individually to the user is used for logging into the NSX Manager appliance.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > System > Users and Domains**.
- 2 Ensure that you are in the **Users** tab.

- 3 If multiple IP addresses are available in the **NSX Manager** drop-down menu, select an IP address, or keep the default selection.
- 4 Click the **Add** icon.

The **Assign Role** window opens.

- 5 Click **Specify a vCenter user** or **Specify a vCenter group**
- 6 Type the vCenter Server user details and group details.

For example:

Field	Example Value
Domain name	corp.vmware.com
Alias	corp
Group name	group1@corp.vmware.com
User name	user1@corp.vmware.com
User alias	user1@corp

**Note** When a group is assigned a role on the NSX Manager, any user from that group can log in to the NSX Manager UI.

- 7 Click **Next**.
- 8 Select the role for the user and click **Next**. For more information about available roles, see [Managing User Rights](#).
- 9 Click **Finish**.

The user account appears in the Users table.

## Group-Based Role Assignments

Organizations create user groups for proper user management. After integration with SSO, NSX Manager can get the details of groups to which a user belongs. Instead of assigning roles to individual users who can belong to the same group, NSX Manager assigns roles to groups. The following scenarios illustrate how NSX Manager assigns roles.

### Example: Role-Based Access Control Scenario

This scenario provides an IT network engineer (Sally Moore) access to NSX components in the following environment:

- AD domain: corp.local
- vCenter group: neteng@corp.local
- User name: smoore@corp.local

Prerequisites: vCenter Server must be registered with NSX Manager, and SSO must be configured. Note that SSO is required only for Groups.

- 1 Assign a role to Sally.
  - a Log in to the vSphere Web Client.
  - b Navigate to **Networking & Security > System > Users and Domains**.
  - c Ensure that you are in the **Users** tab.
  - d Click the **Add** icon.  
The **Assign Role** window opens.
  - e Click **Specify a vCenter group** and type `neteng@corp.local` in **Group**.
  - f Click **Next**.
  - g In **Select Roles**, click **NSX Administrator**, and then click **Next**.
- 2 Grant Sally permission to the data center.
  - a Click the **Home** icon and then click **Networking**.
  - b Select a data center and click **Actions > Add Permission**.
  - c Click **Add** and select the `corp.local` domain.
  - d In **Users and Groups**, select **Show Groups First**.
  - e Select **NetEng** and click **OK**.
  - f In **Assigned Role**, select **Read-only**, deselect **Propagate to children**, and click **OK**.
- 3 Log out of the vSphere Web Client and log in again as `smoore@corp.local`.  
Sally can perform NSX operations only. For example, install virtual appliances, create logical switches, and other operations tasks.

## Example: Inherit Permissions Through a User-Group Membership Scenario

In this scenario, John belongs to group G1, which is assigned the **auditor** role. John inherits the group role and resource permissions.

Group option	Example Value
Name	G1
Role assigned	Auditor (Read only)
Resources	Global root

User option	Example Value
Name	John
Belongs to group	G1
Role assigned	None

## Example: User Member of Multiple Groups Scenario

In this scenario, Joseph belongs to groups G1 and G2 and inherits a combination of the rights and permissions of the **auditor** and **security administrator** roles. For example, Joseph has the following permissions:

- Read, write (**security administrator** role) for Datacenter1
- Read only (**auditor** role) for global root

Group option	Example Value
Name	G1
Role assigned	Auditor (Read only)
Resources	Global root

Group option	Example Value
Name	G2
Role assigned	Security Administrator (Read and Write)
Resources	Datacenter1

User option	Example Value
Name	Joseph
Belongs to group	G1, G2
Role assigned	None

## Example: User Member of Multiple Roles Scenario

In this scenario, Bob is assigned the **security administrator** role, so he does not inherit the group role permissions. Bob has the following permissions:

- Read, write (**security administrator** role) for Datacenter1 and its child resources
- **Enterprise administrator** role on Datacenter1

Group option	Example Value
Name	G1
Role assigned	Enterprise Administrator
Resources	Global root

User option	Example Value
Name	Bob
Belongs to group	G1
Role assigned	Security Administrator (Read and Write)
Resources	Datacenter1

## Create a User with Web Interface Access Using CLI

You can create an NSX user having web interface access using CLI. You can use this user account to access and operate different plug-ins or use it for auditing purposes.

### Procedure

- 1 Create a CLI user account. You can create a CLI user account for each NSX virtual appliance. To create a CLI user account, perform the following steps:
  - a Log in to the vSphere Web Client, and select an NSX Manager virtual appliance.
  - b Click the **Console** tab to open a CLI session.
  - c Log in to the CLI session using the Administrator account and password that you specified while installing NSX Manager. For example,

```
nsx-mgr> enable
Password:
nsx-mgr>
```

- d Switch to Privileged mode from Basic mode using the `enable` command as follows:

```
nsx-mgr> enable
Password:
nsx-mgr#
```

- e Switch to Configuration mode from Privileged mode using the `configure terminal` command as follows:

```
nsx-mgr# configure terminal
nsx-mgr(config)#
```

- f Add a CLI user account using the `user username password (hash | plaintext) password` command. For example,

```
nsx-mgr(config)# user cliuser password plaintext abcd1234
```

---

**Note** User name with capital letter is not allowed .

---

- g Save the configuration as follows:

```
nsx-mgr(config)# write memory
Configuration saved
[OK]
```

- 2 Now provide web interface privilege which will enable the user to login to NSX Manager virtual appliance and allows the execution of appliance management REST APIs as follows:

- a Verify that you are in Configuration mode as follows:

```
nsx-mgr# configure terminal
nsx-mgr(config)#
```

- b Allow the created CLI user to run the REST API calls using the `user username` privilege `web-interface` command. For example:

```
nsx-mgr(config)# user userName privilege web-interface

nsx-mgr(config)# user cliuser privilege web-interface
```

- 3 (Optional) You can verify the running configuration as follows:

```
nsx-mgr# show running-config
Building configuration...

Current configuration:
!
user cliuser
!
ntp server 192.168.110.1
!
ip name server 192.168.110.10
!
hostname nsxmgr-01a
!
interface mgmt
 ip address 192.168.110.15/24
!
ip route 0.0.0.0/0 192.168.110.1
!
web-manager
```

- 4 Exit from the CLI session.

```
nsx-mgr#(config)# exit
nsx-mgr# exit
```

The created user is not listed in the **Networking & Security > System > Users and Domains > Users** tab. Also, no role is assigned to the user.

- 5 Assign the required role to the user using the REST API. You can assign auditor (Auditor), security\_admin (Security Administrator), or super\_user (System Administrator) role as follows:

```
POST - https://<NSX-IP>/api/2.0/services/usermgmt/role/<username>?isCli=true
<accessControlEntry>
<role>auditor</role> # Enter the required role #
<resource>
<resourceId>globalroot-0</resourceId>
</resource>
</accessControlEntry>
```

### Results

The NSX CLI user is created with web interface access.

### What to do next

You can log in to vSphere Web Client using the credentials provided while creating the user.

For more information on CLI, refer to *NSX Command Line Interface Reference*.

For more information on API, refer to *NSX API Guide*.

## Edit a User Account

You can edit a user account to change the role or scope. You cannot edit the **admin** account.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > System > Users and Domains**.
- 2 Ensure that you are in the **Users** tab.
- 3 If multiple IP addresses are available in the **NSX Manager** drop-down menu, select an IP address, or keep the default selection.
- 4 Select the user that you want to edit.
- 5 Click the **Edit** (  or  ) icon.
- 6 Make changes as necessary.
- 7 Click **Finish** to save your changes.

## Change a User Role

You can change the role assignment for all users, except for the **admin** user.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > System > Users and Domains**.

- 2 Ensure that you are in the **Users** tab.
- 3 If multiple IP addresses are available in the **NSX Manager** drop-down menu, select an IP address, or keep the default selection.
- 4 Select the user whose role you want to edit.
- 5 Click the **Edit** (  or  ) icon.
- 6 Make changes as necessary.
- 7 Click **Finish** to save your changes.

## Disable or Enable a User Account

You can disable a user account to prevent that user from logging in to the NSX Manager. You cannot disable the **admin** user or a user who is currently logged into the NSX Manager.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > System > Users and Domains**.
- 2 Ensure that you are in the **Users** tab.
- 3 If multiple IP addresses are available in the **NSX Manager** drop-down menu, select an IP address, or keep the default selection.
- 4 Select a user, and click the **Enable** or **Disable** icon.

## Delete a User Account

You can delete any created user account. You cannot delete the **admin** account. Audit records for deleted users are maintained in the database and can be referenced in an Audit Log report.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > System > Users and Domains**.
- 2 Ensure that you are in the **Users** tab.
- 3 If multiple IP addresses are available in the **NSX Manager** drop-down menu, select an IP address, or keep the default selection.
- 4 Select a user, and then click the **Delete** (  or  ) icon.
- 5 To confirm the deletion, click **Delete** or **Yes**.

If you delete a vCenter user account, only the role assignment for NSX Manager is deleted. The user account on the vCenter Server is not deleted.

This section describes custom network and security containers. These containers can be used in Distributed Firewall and Service Composer. In a cross-vCenter NSX environment, you can create universal network and security containers to be used in universal distributed firewall rules. You cannot use universal network and security objects in Service Composer.

---

**Note** Duplicate names are allowed when you create a group with a universal scope. You can provide duplicate names when you select the **Mark this object for Universal Synchronization** option for creating the following groups:

- IP Address Group (IP Set)
- MAC Address Group (MAC Set)
- Security Group
- Services and Service Group

---

This chapter includes the following topics:

- [Working with IP Address Groups](#)
- [Working with MAC Address Groups](#)
- [Working with IP Pools](#)
- [Working with Security Groups](#)
- [Working with Services and Service Groups](#)

## Working with IP Address Groups

IP address group (IP set) is a way of grouping a list of IP addresses or a range of IP addresses. You can use IP address groups while defining Edge firewall rules and distributed firewall rules.

You can create an IP address group either by manually entering the IP addresses or by importing a `.csv` or `.txt` file that contains a comma-separated list of IP addresses. In addition, you can export an existing IP address group to a `.txt` file that contains a comma-separated list of IP addresses.

## Create an IP Address Group

You can create an IP address group and then add this group as the source or destination in a firewall rule. Such a rule can help protect physical machines from virtual machines or the reverse.

### Prerequisites

- Install VMware Tools on each VM.
- If you plan to use grouping objects instead of IP addresses, enable an IP discovery method, such as DHCP snooping or ARP snooping, or both. For more information, see [IP Discovery for Virtual Machines](#).

### Procedure

- 1 In the vSphere Web Client, click **Networking & Security > Groups and Tags**.
- 2 Navigate to **IP Sets**:
  - In NSX 6.4.1 and later, ensure that you are in the **IP Sets** tab.
  - In NSX 6.4.0, ensure that you are in the **Grouping Objects > IP Sets** tab.
- 3 If multiple IP addresses are available in the **NSX Manager** drop-down menu, select an IP address, or keep the default selection.
  - ◆ You must select the primary NSX Manager if you want to manage universal IP address groups.
- 4 Click **Add** or the **Add (+)** icon.
- 5 Type a name for the address group.
- 6 (Optional) Type a description for the address group.
- 7 Type the IP addresses or a range of IP addresses to be included in the group.

---

**Caution** While entering IPv6 address ranges in the IP sets, ensure that you break the address ranges into /64. Otherwise, the publishing of the firewall rules fails.

---

- 8 (Optional) Select **Inheritance** or **Enable inheritance to allow visibility at underlying scopes**.

When inheritance is enabled, grouping objects created at the global scope are accessible from derived scopes, such as datacenter, Edge, and so on.

- 9 (Optional) To create a universal IP address group:
  - ◆ In NSX 6.4.1 and later, click the **Universal Synchronization** toggle button to **On**.
  - ◆ In NSX 6.4.0, select **Mark this object for Universal Synchronization**.
- 10 Click **Add** or **OK**.

## Edit an IP Address Group

### Procedure

- 1 In the vSphere Web Client, click **Networking & Security > Groups and Tags**.
- 2 Navigate to **IP Sets**:
  - In NSX 6.4.1 and later, ensure that you are in the **IP Sets** tab.
  - In NSX 6.4.0, ensure that you are in the **Grouping Objects > IP Sets** tab.
- 3 If multiple IP addresses are available in the **NSX Manager** drop-down menu, select an IP address, or keep the default selection.
  - ◆ You must select the primary NSX Manager if you want to manage universal IP address groups.
- 4 Select the group that you want to edit, and click the **Edit** (  or  ) icon.
- 5 Make the appropriate changes, and click **Save** or **OK**.

## Delete an IP Address Group

### Procedure

- 1 In the vSphere Web Client, click **Networking & Security > Groups and Tags**.
- 2 Navigate to **IP Sets**:
  - In NSX 6.4.1 and later, ensure that you are in the **IP Sets** tab.
  - In NSX 6.4.0, ensure that you are in the **Grouping Objects > IP Sets** tab.
- 3 If multiple IP addresses are available in the **NSX Manager** drop-down menu, select an IP address, or keep the default selection.
  - ◆ You must select the primary NSX Manager if you want to manage universal IP address groups.
- 4 Select the group that you want to delete, and click the **Delete** (  or  ) icon.

## Working with MAC Address Groups

### Create a MAC Address Group

You can create a MAC address group (MAC set) consisting of a range of MAC addresses and then add this group as the source or destination in a Distributed Firewall rule. Such a rule can help protect physical machines from virtual machines or the reverse.

### Procedure

- 1 In the vSphere Web Client, click **Networking & Security > Groups and Tags**.

**2** Navigate to **MAC Sets**:

- In NSX 6.4.1 and later, ensure that you are in the **MAC Sets** tab.
- In NSX 6.4.0, ensure that you are in the **Grouping Objects > MAC Sets** tab.

**3** If multiple IP addresses are available in the **NSX Manager** drop-down menu, select an IP address, or keep the default selection.

- ◆ To manage universal MAC address groups, the primary NSX Manager must be selected.

**4** Click **Add** or the **Add (+)** icon.**5** Type a name for the address group.**6** (Optional) Type a description for the address group.**7** Type the MAC addresses to be included in the group.**8** (Optional) Select **Inheritance** or **Enable inheritance to allow visibility at underlying scopes**.

When inheritance is enabled, grouping objects created at the global scope are accessible from derived scopes, such as datacenter, Edge, and so on.

**9** (Optional) Select **Universal Synchronization** or **Mark this object for Universal Synchronization** to create a universal MAC address group.**10** Click **Add** or **OK**.

## Edit a MAC Address Group

### Procedure

**1** In the vSphere Web Client, click **Networking & Security > Groups and Tags**.**2** Navigate to **MAC Sets**:

- In NSX 6.4.1 and later, ensure that you are in the **MAC Sets** tab.
- In NSX 6.4.0, ensure that you are in the **Grouping Objects > MAC Sets** tab.

**3** If multiple IP addresses are available in the **NSX Manager** drop-down menu, select an IP address, or keep the default selection.

- ◆ To manage universal MAC address groups, the primary NSX Manager must be selected.

**4** Select the group that you want to edit, and click the **Edit** (✎ or ✎) icon.**5** Make the appropriate changes, and click **Save** or **OK**.

## Delete a MAC Address Group

### Procedure

**1** In the vSphere Web Client, click **Networking & Security > Groups and Tags**.

- 2 Navigate to **MAC Sets**:
  - In NSX 6.4.1 and later, ensure that you are in the **MAC Sets** tab.
  - In NSX 6.4.0, ensure that you are in the **Grouping Objects > MAC Sets** tab.
- 3 If multiple IP addresses are available in the **NSX Manager** drop-down menu, select an IP address, or keep the default selection.
  - ◆ To manage universal MAC address groups, the primary NSX Manager must be selected.
- 4 Select the group that you want to delete, and click the **Delete** (🗑️ or ✖️) icon.

## Working with IP Pools

You can create an IP pool to specify a range of IP addresses.

### Create an IP Pool

You can add IP address ranges to be included in the IP pool. Make sure that the IP pool does not include the IP address range or IP addresses that are already used in the network.

#### Procedure

- 1 In the vSphere Web Client, click **Networking & Security > Groups and Tags**.
- 2 Navigate to **IP Pools**:
  - In NSX 6.4.1 and later, ensure that you are in the **IP Pools** tab.
  - In NSX 6.4.0, ensure that you are in the **Grouping Objects > IP Pools** tab.
- 3 If multiple IP addresses are available in the **NSX Manager** drop-down menu, select an IP address, or keep the default selection.
- 4 Click **Add** or the **Add New IP Pool** icon.
- 5 Type a name for the IP pool and type the default gateway and prefix length.
- 6 (Optional) Type the primary and secondary DNS and the DNS suffix.
- 7 Type the IP address ranges to be included in the pool and click **Add** or **OK**.

### Edit an IP Pool

You can edit the name and IP address range of an IP pool. However, you cannot edit the gateway and prefix length after an IP pool is used.

#### Procedure

- 1 In the vSphere Web Client, click **Networking & Security > Groups and Tags**.
- 2 Navigate to **IP Pools**:
  - In NSX 6.4.1 and later, ensure that you are in the **IP Pools** tab.

- In NSX 6.4.0, ensure that you are in the **Grouping Objects > IP Pools** tab.
- 3 If multiple IP addresses are available in the **NSX Manager** drop-down menu, select an IP address, or keep the default selection.
- 4 Select an IP pool and click the **Edit** (  or  ) icon.
- 5 Make the appropriate changes, and click **Save** or **OK**.

## Delete an IP Pool

You can delete the IP pool.

### Procedure

- 1 In the vSphere Web Client, click **Networking & Security > Groups and Tags**.
- 2 Navigate to **IP Pools**:
  - In NSX 6.4.1 and later, ensure that you are in the **IP Pools** tab.
  - In NSX 6.4.0, ensure that you are in the **Grouping Objects > IP Pools** tab.
- 3 If multiple IP addresses are available in the **NSX Manager** drop-down menu, select an IP address, or keep the default selection.
- 4 Select the IP pool that you want to delete, and click the **Delete** (  or  ) icon.

## Working with Security Groups

A security group is a collection of assets or grouping objects from your vSphere inventory.

Security Groups are containers that can contain multiple object types including logical switch, vNIC, IPset, and Virtual Machine (VM). Security groups can have dynamic membership criteria based on security tags, VM name or logical switch name. For example, all VM's that have the security tag "web" will be automatically added to a specific security group destined for Web servers. After creating a security group, a security policy is applied to that group.

---

**Important** If a VM's VM-ID is regenerated due to move or copy, the security tags are not propagated to the new VM-ID.

---

Security groups for use with Identity Firewall for RDSH, must use security policies that are marked **Enable User Identity at Source** when created. Security groups for use with Identity Firewall for RDSH can only contain Active Directory (AD) groups, and all nested security groups must also be AD groups.

Security groups used in Identity Firewall can contain only AD directory groups. Nested groups can be non-AD groups or other logical entities such as virtual machines.

See [Firewall Rule Behavior in Security Groups](#) for more information.

In a cross-vCenter NSX environment, universal security groups are defined on the primary NSX manager and are marked for universal synchronization with secondary NSX managers. Universal security groups cannot have dynamic membership criteria defined unless they are marked for use in an active standby deployment scenario.

In a cross-vCenter NSX environment with an active standby deployment scenario, the SRM creates a placeholder VM on the recovery site for every protected VM on the active site. The placeholder VMs are not active, and stay in the standby mode. When the protected VM goes down, the placeholder VMs on the recovery site are powered on and take over the tasks of the protected VM. Users create distributed firewall rules with universal security groups containing universal security tags on the active site. The NSX manager replicates the distributed firewall rule with the universal security groups containing universal security tags on the placeholder VMs and when the placeholder VMs are powered on the replicated firewall rules with the universal security groups and universal security tags are enforced correctly.

### Note

- Universal security groups created prior to 6.3 cannot be edited for use in active standby deployments.

## Firewall Rule Behavior in Security Groups

Firewall rule behavior varies with different Security Groups.

Table 22-1. Firewall Rule Behavior with RDSH and Non-RDSH Sections

Enable User Identity Security Group (RDSH Section)	Identity Security Group (RDSH Section)	Any Security Group (Non-RDSH Section)
Source - SID based rules are preemptively pushed to hypervisor. Rule enforcement is on the first packet.	Source - IP based rules	Source - IP based rules
Destination - IP based rules	Destination - IP based rules	Destination - IP based rules
Applied To with Identity based Security Group - Applied to all hosts		User based Applied To
Applied To with Non-Identity based Security Group - User based Applied to		User based Applied to

## Create a Security Group

You create a security group at the NSX Manager level.

Universal security groups are used in two types of deployments: active cross-vCenter NSX environments, and active standby cross-vCenter NSX environments, where one site is live at a given time and the rest are on standby.

- Universal security groups in an active environment can contain the following included objects only: security groups, IP sets, MAC sets. You cannot configure dynamic membership or excluded objects.

- Universal security groups in an active standby environment can contain the following included objects: security groups, IP sets, MAC sets, universal security tags. You can also configure dynamic membership using VM Name only. You cannot configure excluded objects.

---

**Note** Powered off VM's that are based on dynamic criteria such as Computer OS Name and Computer Name will not be included in Security Groups. Dynamic criteria is received by NSX only once when the VM is powered on. After being powered on, the guest details are synced to NSX Manager and remain with the NSX Manager even if the VM is later powered off.

---

**Note** Universal security groups created prior to 6.3 cannot be edited for use in active standby deployments.

---

### Prerequisites

If you are creating a security group based on Active Directory group objects, ensure that one or more domains have been registered with NSX Manager. NSX Manager gets group and user information as well as the relationship between them from each domain that it is registered with. See [Register a Windows Domain with NSX Manager](#).

### Procedure

- 1 In the vSphere Web Client, click **Networking & Security > Groups and Tags**.
- 2 Navigate to **Security Group**:
  - In NSX 6.4.1 and later, ensure that you are in the **Security Groups** tab.
  - In NSX 6.4.0, ensure that you are in the **Grouping Objects > Security Group** tab.
- 3 If multiple IP addresses are available in the **NSX Manager** drop-down menu, select an IP address, or keep the default selection.
  - ◆ To manage universal security groups, the primary NSX Manager must be selected.
- 4 Click **Add** or the **Add New Security Group** icon.
- 5 Type a name and optionally a description for the security group.
- 6 (Optional) If you are creating a universal security group, select **Universal Synchronization** or **Mark this object for universal synchronization**.
- 7 (Optional) If you are creating a universal security group for use in an active standby deployment, select both **Universal Synchronization / Mark this object for universal synchronization** and **Use for active standby deployments**. Dynamic membership for universal security groups with active standby deployment is based on virtual machine name
- 8 Click **Next**.

- On the Dynamic Membership page, define the criteria that an object must meet for it to be added to the security group you are creating. This gives you the ability to include virtual machines by defining a filter criteria with a number of parameters supported to match the search criteria.

**Note** If you are creating a universal security group, the **Define dynamic membership** step is not available in active active deployments. It is available in active standby deployments, based on virtual machine name only.

For example, you may include a criterion to add all virtual machines tagged with the specified security tag (such as AntiVirus.virusFound) to the security group. Security tags are case sensitive.

Or you can add all virtual machines containing the name w2008 and virtual machines that are in the logical switch global\_wire to the security group.

The screenshot shows the 'New Security Group' configuration interface. On the left, a navigation pane lists five steps: 1 Name and description, 2 Define dynamic membership (highlighted), 3 Select objects to include, 4 Select objects to exclude, and 5 Ready to complete. The main area is titled 'Define dynamic membership' and contains the instruction 'Specify dynamic membership criteria that objects must meet to be part of this security group.' Below this is a list of criteria:

- Members matching **All** of the criteria below
- Computer Name **Contains** w2008
- Entity **Belongs to** global\_wire

- Click **Next**.

- 11 On the Select objects to include page, select the tab for the resource you want to add and select one or more resources to add to the security group. You can include the following objects in a security group.

**Table 22-2. Objects that can be included in security groups and universal security groups.**

Security Group	Universal Security Group
<ul style="list-style-type: none"> <li>■ Other security groups to nest within the security group you are creating.</li> <li>■ Cluster</li> <li>■ Logical Switch</li> <li>■ Network</li> <li>■ Virtual App</li> <li>■ Datacenter</li> <li>■ IP sets</li> <li>■ Directory Groups</li> </ul> <hr/> <p><b>Note</b> The Active Directory configuration for NSX security groups is different from the AD configuration for vSphere SSO. NSX AD group configuration is for end users accessing guest virtual machines, while vSphere SSO is for administrators using vSphere and NSX. In order to consume these directory groups you must sync with Active Directory. See <a href="#">Chapter 12 Identity Firewall Overview</a>.</p> <ul style="list-style-type: none"> <li>■ MAC Sets</li> <li>■ Security tag</li> <li>■ vNIC</li> <li>■ Virtual Machine</li> <li>■ Resource Pool</li> <li>■ Distributed Virtual Port Group</li> </ul>	<ul style="list-style-type: none"> <li>■ Other universal security groups to nest within the universal security group you are creating.</li> <li>■ Universal IP sets</li> <li>■ Universal MAC sets</li> <li>■ Universal Security Tag (active standby deployments only)</li> </ul>

The objects selected here are always included in the security group regardless of whether or not they match the criteria that you defined earlier on the Dynamic Membership page.

When you add a resource to a security group, all associated resources are automatically added. For example, when you select a virtual machine, the associated vNIC is automatically added to the security group.

- 12 Click **Next** and select the objects that you want to exclude from the security group.

---

**Note** If you are creating a universal security group, the **Select objects to exclude** step is not available.

---

The objects selected here are always excluded from the security group regardless of whether or not they match the dynamic criteria.

- 13 Click **Next**.

The **Ready to Complete** window appears with a summary of the security group.

- 14 Click **Finish**.

## Example

Membership of a security group is determined as follows:

{Expression result (derived from **Define dynamic membership**) + Inclusions (specified in **Select objects to include**) - Exclusion (specified in **Select objects to exclude**)

This means that inclusion items are first added to the expression result. Exclusion items are then subtracted from the combined result.

## Edit a Security Group

### Procedure

- 1 In the vSphere Web Client, click **Networking & Security > Groups and Tags**.
- 2 Navigate to **Security Group**:
  - In NSX 6.4.1 and later, ensure that you are in the **Security Groups** tab.
  - In NSX 6.4.0, ensure that you are in the **Grouping Objects > Security Group** tab.
- 3 If multiple IP addresses are available in the **NSX Manager** drop-down menu, select an IP address, or keep the default selection.
  - ◆ To manage universal security groups, the primary NSX Manager must be selected.
- 4 Select the group that you want to edit and click the **Edit** (  or  ) icon.

---

**Note** Universal security groups created prior to 6.3 cannot be edited for use in active standby deployments.

---

- 5 In the Edit Security Group dialog box, make the appropriate changes.
- 6 Click **Finish** or **OK**.

## Delete a Security Group

### Procedure

- 1 In the vSphere Web Client, click **Networking & Security > Groups and Tags**.
- 2 Navigate to **Security Group**:
  - In NSX 6.4.1 and later, ensure that you are in the **Security Groups** tab.
  - In NSX 6.4.0, ensure that you are in the **Grouping Objects > Security Group** tab.
- 3 If multiple IP addresses are available in the **NSX Manager** drop-down menu, select an IP address, or keep the default selection.
  - ◆ To manage universal security groups, the primary NSX Manager must be selected.
- 4 Select the group that you want to delete and click the **Delete** (  or  ) icon.

# Working with Services and Service Groups

A service is a protocol-port combination, and a service group is a group of services or other service groups.

## Create a Service

You can use services in firewall rules. You can use pre-defined services, or create additional services.

You might need to create a service because your application is not already defined, or it is using a standard protocol, with a non-default port. For example:

- HTTP on a non-default port - TCP:8080
- FTP on a non-default port - FTP:8021
- NoMachine Server port - UDP:4000

---

**Note** SCTP protocol is not supported on Edge Firewall.

---

### Procedure

- 1 In the vSphere Web Client, click **Networking & Security > Groups and Tags**.
- 2 Navigate to **Services**:
  - In NSX 6.4.1 and later, ensure that you are in the **Services** tab.
  - In NSX 6.4.0, ensure that you are in the **Grouping Objects > Service** tab.
- 3 If multiple IP addresses are available in the **NSX Manager** drop-down menu, select an IP address, or keep the default selection.
  - ◆ To manage universal services, the primary NSX Manager must be selected.
- 4 Click **Add** or the **Add (+)** icon.
- 5 Enter a **Name** to identify the service.
- 6 (Optional) Enter a **Description** for the service.
- 7 Select a **Layer**.

If you select Layer 7, you are prompted to select an App ID.

- 8 Select a **Protocol**.

For example, TCP, UDP, or FTP.

Depending on the protocol selected, you might be prompted to enter further information, such as the destination port. Expand **Advanced Options** to enter a source port.

---

**Note** SCTP protocol is not supported on Edge Firewall.

---

- 9 (Optional) Select **Inheritance** or **Enable inheritance to allow visibility at underlying scopes**.

When inheritance is enabled, grouping objects created at the global scope are accessible from derived scopes, such as datacenter, Edge, and so on.

- 10 (Optional) Select **Universal Synchronization** or **Mark this object for Universal Synchronization** to create a universal service.

- 11 Click **Add** or **OK**.

### Results

The service appears in the Services table.

## Create a Service Group

You can create a service group and then define rules for that service group.

### Procedure

- 1 In the vSphere Web Client, click **Networking & Security > Groups and Tags**.
- 2 Navigate to **Service Groups**:
  - In NSX 6.4.1 and later, ensure that you are in the **Service Groups** tab.
  - In NSX 6.4.0, ensure that you are in the **Grouping Objects > Service Groups** tab.
- 3 If multiple IP addresses are available in the **NSX Manager** drop-down menu, select an IP address, or keep the default selection.
  - ◆ To manage universal security groups, the primary NSX Manager must be selected.
- 4 Click **Add** or the **Add (+)** icon.
- 5 Type a **Name** to identify the service group.
- 6 (Optional) Type a **Description** for the service group.
- 7 In **Members**, select the services or service groups that you want to add to the group.
- 8 (Optional) Select **Universal Synchronization** or **Mark this object for Universal Synchronization** to create a universal service group.
- 9 (Optional) Select **Inheritance** or **Enable inheritance to allow visibility at underlying scopes**.
 

When inheritance is enabled, grouping objects created at the global scope are accessible from derived scopes, such as datacenter, Edge, and so on.
- 10 Click **Add** or **OK**.

## Edit a Service or Service Group

You can edit services and service groups.

### Procedure

- 1 In the vSphere Web Client, click **Networking & Security > Groups and Tags**.
- 2 Click the **Service** tab or the **Service Groups** tab.
- 3 Select a custom service or service group and click the **Edit** ( or ) icon.
- 4 Make the appropriate changes and click **Save** or **OK**.

## Delete a Service or Service Group

You can delete services or service group.

### Procedure

- 1 In the vSphere Web Client, click **Networking & Security > Groups and Tags**.
- 2 Click the **Service** tab or the **Service Groups** tab.
- 3 Select a custom service or service group and click the **Delete** ( or ) icon.

This chapter includes the following topics:

- Add and Assign a License
- Using the Dashboard
- Check Communication Channel Health
- NSX Controller Management
- Controller Disconnected Mode for Multiple Sites
- Change VXLAN Port
- Customer Experience Improvement Program
- About NSX Logs
- Audit Logs
- System Events
- Management System Settings
- NSX Backup and Restore
- NSX Monitoring And Diagnostic Tools

## Add and Assign a License

You can add and assign a license using the vSphere Web Client.

Starting in NSX 6.4.0, you must be a member of the **LicenseService.Administrators** group to manage licenses.

The default license upon install is NSX for vShield Endpoint. This license enables use of NSX for deploying and managing vShield Endpoint for anti-virus offload capability only, and has hard enforcement to restrict usage of VXLAN, firewall, and Edge services, by blocking host preparation and creation of NSX Edges. To use other features, including logical switches, logical routers, Distributed Firewall, or NSX Edge, you must either purchase a license to use these features, or request an evaluation license for short-term evaluation of the features.

- NSX for vSphere Standard, Advanced, and Enterprise license keys are effective in NSX 6.2.2 and later.
- NSX Data Center Standard, Professional, Advanced, Enterprise Plus, and Remote Office Branch Office license keys are effective in NSX 6.4.1 and later.

For more information about the NSX Data Center, NSX, and NSX for vShield Endpoint licensing editions and associated features, see <https://kb.vmware.com/kb/2145269>.

For more information about VMware product licensing, see <http://www.vmware.com/files/pdf/vmware-product-guide.pdf>.

For more information about managing license in vSphere, see the *vCenter Server and Host Management* documentation for your version of vSphere.

### Prerequisites

Verify that all vCenter users who manage licenses are in the **LicenseService.Administrators** group.

If you have multiple vCenter Server systems using the same Platform Services Controller, and multiple NSX Managers registered with those vCenter Servers, you must combine the licenses for the NSX Manager appliances into one license. See <https://kb.vmware.com/s/article/53750> for more information. See <https://kb.vmware.com/s/article/2006973> for information about combining licenses.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Administration** and then click **Licenses**.
- 3 Click the **Assets** tab, then the **Solutions** tab.
- 4 Select NSX for vSphere in the Solutions list. From the **All Actions** drop-down menu, select **Assign license....**
- 5 Click the **Add (+)** icon. Enter a license key and click **Next**. Add a name for the license, and click **Next**. Click **Finish** to add the license.
- 6 Select the new license.
- 7 (Optional) Click the **View Features** icon to view what features are enabled with this license.
- 8 Click **OK** to assign the new license to NSX.

## Using the Dashboard

The dashboard provides visibility to the overall health of NSX components in one central view. The dashboard simplifies troubleshooting by displaying status of different NSX components such as NSX Manager, controllers, logical switches, host preparation, service deployment, backup as well as edge notifications.

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security**. The **Dashboard > Overview** page appears as your default homepage.

You can view existing system-defined widgets and the custom widgets.

Widget	Component
System Overview	<p>NSX Manager:</p> <ul style="list-style-type: none"> <li>■ CPU usage</li> <li>■ NSX Manager disk usage</li> <li>■ Service status for database service, message bus service, and replicator service. Replication errors on the secondary NSX Manager</li> <li>■ Controller synchronization status</li> </ul> <hr/> <p>Controller Nodes:</p> <ul style="list-style-type: none"> <li>■ Controller node status</li> <li>■ Controller peer connectivity status</li> <li>■ Controller VM status (powered off/deleted)</li> <li>■ Controller disk latency alerts</li> </ul> <hr/> <p>External Components:</p> <ul style="list-style-type: none"> <li>■ vSphere ESX Agent Manager (EAM) service status</li> </ul>
Firewall Publish Status	Number of hosts with Firewall Publish status as failed. Status is Red when any host does not successfully apply the published distributed firewall configuration
Logical Switch Status	Number of logical switches with status <i>Error</i> or <i>Warning</i> . Flags when the backed distributed virtual port group is deleted from vCenter Server
Service Deployment Status	<ul style="list-style-type: none"> <li>■ Installation status for the failed deployments</li> <li>■ Service status for all the failed services</li> </ul>
Host Notification	Security alerts for hosts. You can see this alert when the hardware address of the DHCP client is spoofed. A possible DHCP denial-of-service (DoS) attack is happening.

Widget	Component
Fabric Status	<p>Host preparation status:</p> <ul style="list-style-type: none"> <li>■ Deployment status like, clusters with installation failed status, pending upgrade, installation in-progress, and so on.</li> <li>■ Firewall: <ul style="list-style-type: none"> <li>■ Number of clusters with firewall disabled</li> <li>■ Status of the distributed firewall</li> </ul> </li> <li>■ VXLAN: <ul style="list-style-type: none"> <li>■ Number of clusters with VXLAN not configured</li> <li>■ VXLAN status</li> </ul> </li> </ul> <p>Communication channel health status</p>
Backup Status	<p>Backup status for NSX Manager:</p> <ul style="list-style-type: none"> <li>■ Backup schedule</li> <li>■ Last backup status (Failed/successful/not scheduled with the date and time)</li> <li>■ Last backup attempt (date and time with details)</li> <li>■ Last successful backup (date and time with details)</li> </ul>
Edge Notifications	<p>Highlights active alarms for certain services. It monitors the list of critical events that are listed and tracks them until the problem is not resolved. Alarms are auto resolved when the recovery event is reported, or edge is force synced, redeployed, or upgraded</p>
Tools	<ul style="list-style-type: none"> <li>■ Flow Monitoring status</li> <li>■ Endpoint Monitoring status</li> </ul>
System Scale Dashboard	<p>Shows a summary of warnings and alerts for scale. For detail listing of the parameters and scale numbers, click <b>Details</b> to go to the <a href="#">System Scale Dashboard</a></p>
Custom Widget	<p>You can view the custom widget created through API</p>

## Custom Widget

You can add custom widgets to the dashboard using the REST API. You can create five custom widgets for your personal viewing in the dashboard.

You can also share the custom widgets with other users by setting the `shared` parameter as `true` in the widget configuration. Maximum limit on shared widgets is 10. It means that the total number of widgets shared by all the users is limited to 10.

## Custom Widgets APIs

Use the following APIs to view, create, modify, and delete the custom widgets on your dashboard:

- To view information about a specific widget configuration: Use the `GET /api/2.0/services/dashboard/ui-views/dashboard/widgetconfigurations/<widgetconfiguration-id>` API .
- To create a custom widget: Use the `POST /api/2.0/services/dashboard/ui-views/dashboard/widgetconfigurations` API .

- To modify the existing widget configuration: Use the `PUT /api/2.0/services/dashboard/ui-views/dashboard/widgetconfigurations/<widgetconfiguration-id>` API .
- To delete the custom widget created earlier: Use the `DELETE /api/2.0/services/dashboard/ui-views/dashboard/widgetconfigurations/<widgetconfiguration-id>` API.

For more information on API, refer to *NSX API Guide*.

## System Scale Dashboard

The **System Scale** dashboard collects information about the current system scale and displays the configuration maximums for the supported scale parameters. You can configure a warning threshold for the alerts when the system scale exceeds the configured threshold value. When the threshold is crossed, system event is generated which is used to set up notifications. This information is also logged and included in the support bundle.

If the current value exceeds a specified threshold percentage, a warning indicator is displayed to alert that the maximum supported scale is approaching. A red indicator shows that configuration maximum is reached. The listing is sorted in descending order of the current scale percent value which ensures that the warning indicators are always displayed at the top.

The data is collected every hour to verify if the threshold values are exceeding the limits, and creates an indicator when the thresholds are exceeded. The information is logged twice a day to the NSX Manager technical support logs.

System events are generated when the following conditions occur:

- A warning event when a parameter crosses the threshold limit.
- A critical event when a parameter crosses the supported system scale configuration.
- An informational event when a parameter comes below the threshold value.

### Retrieve Current Scale Threshold Limits for All Parameters

You can find out the current and the supported system scale configuration using the `GET /api/2.0/capacity-parameters/report` API. The API output displays scale summary, current scale value, supported system scale value, and the threshold value for each parameter.

### Configure Scale Threshold for the System

You can set the scale threshold limit for your system.

To configure the threshold limit:

- 1 Retrieve the global system threshold using the `GET /api/2.0/capacity-parameters/thresholds` API. For example, the API output shows global threshold as *80*. It means the **System Scale** dashboard displays **Usage Warning Threshold** as 80%.

---

**Note** By default, the global threshold value is set at *80*.

---

- To change the system threshold, use the `PUT /api/2.0/capacity-parameters/thresholds` API. For example, you change the global threshold value to `70`. Now the **System Scale** dashboard displays **Usage Warning Threshold** as 70%.

For more information on APIs, refer to *NSX API Guide*.

For more information on system scales, refer to *NSX Recommended Configuration Maximums*.

## Check Communication Channel Health

NSX checks on the status of communication between NSX Manager and firewall agent, NSX Manager and control plane agent, and control plane agent and controllers.

You can also view Communication Channel Health on the dashboard at **Networking & Security > Dashboard** in the Fabric Status section.

### Procedure

- Navigate to **Networking & Security > Installation and Upgrade > Host Preparation**.
- Complete the following steps to view the health of the communication channels.

NSX Version	Procedure
NSX 6.4.1 and later	<ol style="list-style-type: none"> <li>Click a cluster from the left pane. In the right pane, the hosts in the selected cluster are displayed in the Hosts table.</li> <li>In the Communications Channels column of the Hosts table, click the status icon.</li> </ol>
NSX 6.4.0	<ol style="list-style-type: none"> <li>Expand the cluster that contains the host for which you want to view the communication channel health.</li> <li>Click the host, and then click <b>Actions &gt; Communication Channel Health</b>.</li> </ol>

A pop-up window displays the health status of the following communication channels:

- NSX Manager to Firewall Agent
- NSX Manager to Control Plane Agent
- Control Plane Agent to Controller

## NSX Controller Management

You can change the NSX Controller node configuration, including changing names, passwords, or configuring NTP. You can also download tech support information from the NSX Controller nodes.

For information about troubleshooting controller cluster problems, including deleting controllers safely, refer to the NSX Controller section in the *NSX Troubleshooting Guide*.

### Change NSX Controller Name

You can change the name of each NSX Controller node.

Starting in NSX Data Center for vSphere 6.4.0, you can change the controller name using the API. See the *NSX API Guide* for more information. Starting in NSX Data Center for vSphere 6.4.2, you can change the controller name using the vSphere Web Client or vSphere Client.

When you create a controller node, you are prompted to provide a name. The controller node is also assigned a controller ID, in the format `controller-X`, for example, `controller-5`. The controller name and ID are used to configure identifiers for the controller in a few locations:

- The name displayed in the Networking & Security UI is set to `name`
- The VM name in vSphere is set to `name-NSX-<controller_id>`
- The VM's hostname is set to `nsx-controller`

When you update the controller name, the following changes are made:

- The name displayed in the Networking & Security UI is changed to `newName`
- The VM name in vSphere is changed to `newName-NSX-<controller_id>`
- The VM's hostname is changed to `newName-NSX-<controller_id>`

---

**Note** The hostname is used in controller log entries. If you change the controller hostname, the log entries display the new hostname.

---

#### Procedure

- 1 Navigate to **Networking & Security > Installation and Upgrade > Management > NSX Controller Nodes**.
- 2 Select a controller node, then click **Actions > Change Controller Name**.
- 3 Enter the new name and click **Save**.

## Change Controller Password

To ensure security, you can change passwords for NSX controllers.

#### Procedure

- 1 Navigate to **Networking & Security > Installation and Upgrade > Management > NSX Controller Nodes**.
- 2 Select the controller for which you want to change the password.
- 3 Click **Actions > Change Controller Cluster Password**.
- 4 Type a new password and click **OK**.

The controller password is changed.

## Configure DNS, NTP, and Syslog for the NSX Controller Cluster

You can configure DNS, NTP, and syslog servers for the NSX Controller cluster. The same settings apply to all NSX Controller nodes in the cluster.

Starting in NSX Data Center for vSphere 6.4.2, you can make these changes using the vSphere Web Client or vSphere Client. In earlier 6.4 versions, you can change NTP, and syslog settings using the API only. See the *NSX API Guide* for more information.

---

**Important** If you have an invalid configuration (for example, unreachable NTP servers), and then deploy a controller, the controller node deployment fails. Verify and correct the configuration and deploy the controller node again.

---

The NSX Controller cluster DNS settings override any DNS settings configured on the controller IP pool.

#### Procedure

- 1 Navigate to **Networking & Security > Installation and Upgrade > Management > NSX Controller Nodes**.
- 2 Select the NSX Manager that manages the NSX Controller nodes you want to modify.
- 3 Click the Common Controller Attributes **EDIT** link.
- 4 (Optional) Enter a comma-separated list of DNS servers, and optionally DNS suffixes.

DNS Setting	Example Values
DNS Servers	192.168.110.10, 192.168.110.11
DNS Suffixes	eng.example.com, corp.example.com, example.com

- 5 (Optional) Enter a comma-separated list of NTP servers.

You can enter the NTP servers as IPv4 addresses or fully qualified domain names (FQDN). If an FQDN is used, you must configure DNS so that the names can be resolved.

- 6 (Optional) Configure one or more syslog servers.

- a In the Syslog Servers panel, click **ADD**.
- b Enter the syslog server name or address.

You can enter the syslog servers as IPv4 addresses or fully qualified domain names (FQDN). If an FQDN is used, you must configure DNS so that the names can be resolved.

- c Select the protocol.

If you select TLS, you must provide a PEM-encoded X.509 certificate.

---

**Important** Selecting TCP or TLS might result in extra consumption of memory for buffering that could negatively impact the performance of the controller. In extreme cases, this can stop controller processing until the buffered network log calls are drained.

---

#### Note

- If the syslog server is using a self-signed certificate, paste the contents of the syslog self-signed certificate in the **Certificate** text box.
- If the syslog server is using a CA-signed certificate, paste the contents of the intermediary certificates and the root certificate. In the certificate chain, the order of certificates must be as follows:
  - Any number of intermediate CA certificates
  - Root CA certificate

Each certificate must include the -----BEGIN CERTIFICATE----- and -----END CERTIFICATE----- lines, as shown in the following example:

```
-----BEGIN CERTIFICATE-----
 Intermediate cert
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
 Root cert
-----END CERTIFICATE-----
```

- d (Optional) Edit the port.

The default port for TCP and UDP syslog is 514. For TLS syslog, the default port is 6514.

- e (Optional) Select the log level.

INFO is selected by default.

## Download Technical Support Logs for NSX Controller

You can download technical support logs for each NSX Controller instance. These product specific logs contain diagnostic information for analysis. You can also collect the support bundle data for controllers using the **Support Bundle** collection tool. For details, refer to *NSX Administration Guide*.

To collect NSX Controller logs:

#### Procedure

- 1 Navigate to **Networking & Security > Installation and Upgrade > Management > NSX Controller Nodes**.

- 2 Select the controller for which you want to generate technical support logs.

---

**Caution** Generate support logs for one controller at a time. An error might occur if you try to generate support logs for multiple controllers simultaneously.

---

- 3 Click **Support Logs** (  or  ).

NSX starts collecting the technical support logs. It takes several minutes for the log files to be generated. You can click **Cancel** at any time to cancel the process and generate the support logs later.

- 4 After the support logs are generated, click **Download**.

The support logs are saved on your computer in a compressed file with the `.tgz` file extension.

### Results

You can now analyze the downloaded logs.

### What to do next

If you want to upload diagnostic information for VMware technical support, refer to the [Knowledge Base article 2070100](#).

## Controller Disconnected Mode for Multiple Sites

Controller Disconnected Operation (CDO) mode ensures that the data plane connectivity is unaffected in a multi-site environment, when the primary site loses connectivity. You can enable the CDO mode on the secondary site to avoid temporary connectivity issues related to the data plane, when the primary site is down or not reachable. You can also enable the CDO mode on the primary site for the control plane failure.

CDO mode avoids the connectivity issues during the following failure scenarios:

- The complete primary site of a cross-vCenter NSX environment is down.
- WAN is down.
- Control plane failure.

---

**Note** Starting in NSX 6.4.0, CDO is supported at NSX Manager level and not at the transport zone level.

---

The CDO mode is disabled by default.

When the CDO mode is enabled and host detects a control plane failure, the host waits for the configured time period and then enters the CDO mode. You can configure the time period for which you want the host to wait before entering the CDO mode. By default, the wait time is five minutes.

NSX Manager creates a special CDO logical switch (4999) on the controller. The VXLAN Network Identifier (VNI) of the special CDO logical switch is unique from all other logical switches. When the CDO mode is enabled, one controller in the cluster is responsible for collecting all the VTEP information reported from all transport nodes, and replicating the updated VTEP information to all other transport nodes. After detecting the CDO mode, broadcast packets like ARP/GARP and RARP is sent to the global VTEP list. This allows to vMotion the VMs across the vCenter Servers without any data plane connectivity issues.

When you disable the CDO mode, NSX Manager removes the CDO logical switch from the controller.

## Enable Controller Disconnected Operation (CDO) Mode

You can enable the Controller Disconnected Operation (CDO) mode for NSX Manager. CDO mode is disabled by default.

### Prerequisites

- After upgrading to NSX 6.4, NSX Manager disables CDO mode for the existing transport nodes. Use a pre-defined global VNI to configure vSphere Distributed Switch (VDS), if NSX Manager had one or more CDO enabled transport nodes before upgrade.

### Procedure

- 1 Navigate to **Networking & Security > Installation and Upgrade > Management > NSX Managers**.
- 2 Select the required NSX Manager.
- 3 Click **Actions > Enable CDO mode**.  
A confirmation dialog box appears.
- 4 Click **Yes**.  
CDO mode is enabled for the selected NSX Manager.

### Results

The CDO Mode column displays State as **Enabled** and Status as **Successful**.

NSX Manager creates a CDO logical switch on the controller. To view the details of the CDO logical switch, log in to the NSX Manager CLI as an **admin** user, and run the following command:

```
nsxmgr> show logical-switch controller controllerID host host_IP joined-vnis
```

For example:

```
nsxmgr> show logical-switch controller controller-1 host 192.168.110.54 joined-vnis
VNI Controller BUM-Replication ARP-Proxy Connections VTEPs Active
5000 192.168.110.31 Enabled Enabled 2 2 true
5004 192.168.110.31 Enabled Enabled 2 2 true
5006 192.168.110.31 Enabled Enabled 3 3 true
```

4999	192.168.110.31	Enabled	Enabled	5	5	true
5003	192.168.110.31	Enabled	Enabled	3	3	true
5005	192.168.110.31	Enabled	Enabled	1	1	true
5007	192.168.110.31	Enabled	Enabled	3	3	true

Observe that the CDO logical switch with VNI 4999 is created.

## Disable Controller Disconnected Operation (CDO) Mode

You can disable the already enabled Controller Disconnected Operation (CDO) mode when the connectivity issues with the secondary site are resolved. CDO mode remains disabled by default.

### Procedure

- 1 Navigate to **Networking & Security > Installation and Upgrade > Management > NSX Managers**.
- 2 Select the required NSX Manager.
- 3 Click **Actions > Disable CDO mode**.  
A confirmation dialog box appears.
- 4 Click **Yes**.  
CDO mode is disabled for the selected NSX Manager.

### Results

The CDO Mode column displays State as **Disabled** and Status as **Successful**.

NSX Manager removes the CDO logical switch from the controller. To verify whether the CDO logical switch is removed, log in to the NSX Manager CLI as an **admin** user, and run the following command:

```
nsxmgr> show logical-switch controller controllerID host host_IP joined-vnis
```

For example:

```
nsxmgr> show logical-switch controller controller-1 host 192.168.110.54 joined-vnis
VNI Controller BUM-Replication ARP-Proxy Connections VTEPs Active
5000 192.168.110.31 Enabled Enabled 2 2 true
5004 192.168.110.31 Enabled Enabled 2 2 true
5006 192.168.110.31 Enabled Enabled 3 3 true
5003 192.168.110.31 Enabled Enabled 3 3 true
5005 192.168.110.31 Enabled Enabled 1 1 true
5007 192.168.110.31 Enabled Enabled 3 3 true
```

Observe that the CDO logical switch with VNI 4999 is not available on the controller.

## Resync CDO Configuration

If the CDO mode operation fails, you can execute the same operation using the resync feature.

For example: A vSphere Distributed Switch (VDS) gets deleted from the NSX Manager database, when the last cluster associated with that VDS gets unconfigured from the VXLAN. Then, NSX Manager removes the CDO-related opaque property from VDS. If the removal of opaque property fails due to some reason, the opaque property remains in VDS. Now if you disable the CDO mode and configure VXLAN on a host associated with that VDS, the CDO gets enabled on that host as the opaque property is present and VNI is also present in the controller. In this case, use the resync feature to apply the CDO disabled mode again. Now, NSX Manager tries to remove the opaque property from VDS.

### Prerequisites

CDO mode enable or disable operation has failed.

### Procedure

- 1 Navigate to **Networking & Security > Installation and Upgrade > Management > NSX Managers**.
- 2 Select the secondary NSX Manager, and click **Actions > Enable CDO mode** to enable the CDO mode.
- 3 Now, select the required NSX Manager, and click **Actions > Resync CDO configuration mode**. A confirmation dialog box appears.
- 4 Click **Yes**.  
CDO mode is synchronized again for the selected NSX Manager.

## Change VXLAN Port

You can change the port used for VXLAN traffic.

In NSX 6.2.3 and later, the default VXLAN port is 4789, the standard port assigned by IANA. Before NSX 6.2.3, the default VXLAN UDP port number was 8472.

Any new NSX installations will use UDP port 4789 for VXLAN.

If you upgrade from NSX 6.2.2 or earlier to NSX 6.2.3 or later, and your installation used the old default (8472), or a custom port number (for example, 8888) before the upgrade, that port will continue to be used after the upgrade unless you take steps to change it.

If your upgraded installation uses or will use hardware VTEP gateways (ToR gateways), you must switch to VXLAN port 4789.

Cross-vCenter NSX does not require that you use 4789 for the VXLAN port, however, all hosts in a cross-vCenter NSX environment must be configured to use the same VXLAN port. If you switch to port 4789, this will ensure that any new NSX installations added to the cross-vCenter NSX environment are using the same port as the existing NSX deployments.

Changing the VXLAN port is done in a three phase process, and will not interrupt VXLAN traffic.

- 1 NSX Manager configures all hosts to listen for VXLAN traffic on both the old and new ports. Hosts continue to send VXLAN traffic on the old port.
- 2 NSX Manager configures all hosts to send traffic on the new port.
- 3 NSX Manager configures all hosts to stop listening on the old port, all traffic is sent and received on the new port.

In a cross-vCenter NSX environment you must initiate the port change on the primary NSX Manager. For each stage, the configuration changes are made on all hosts in the cross-vCenter NSX environment before proceeding to the next stage.

### Prerequisites

- Verify that the port you want to use for VXLAN is not blocked by a firewall.
- Verify that host preparation is not running at the same time as the VXLAN port change.

### Procedure

- 1 Navigate to VXLAN transport settings.
  - ◆ In NSX 6.4.1 and later, navigate to **Networking & Security > Installation and Upgrade > Logical Network Settings > VXLAN Settings**.
  - ◆ In NSX 6.4.0, navigate to **Networking & Security > Installation and Upgrade > Logical Network Preparation > VXLAN Transport**.
- 2 Next to **VXLAN Port**, click **Edit** or **Change**. Enter the port you want to switch to. 4789 is the port assigned by IANA for VXLAN.

It takes a short time for the port change to propagate to all hosts.

- 3 (Optional) Check the progress of the port change with the `GET /api/2.0/vdn/config/vxlan/udp/port/taskStatus` API request.

```
GET https://nsxmgr-01a/api/2.0/vdn/config/vxlan/udp/port/taskStatus
```

```
<?xml version="1.0" encoding="UTF-8"?>
<vxlanPortUpdatingStatus>
 <prevPort>8472</prevPort>
 <targetPort>4789</targetPort>
 <taskPhase>PHASE_TWO</taskPhase>
 <taskStatus>PAUSED</taskStatus>
</vxlanPortUpdatingStatus>
```

...

```
<?xml version="1.0" encoding="UTF-8"?>
<vxlanPortUpdatingStatus>
 <prevPort>8472</prevPort>
 <targetPort>4789</targetPort>
 <taskPhase>FINISHED</taskPhase>
 <taskStatus>SUCCEED</taskStatus>
</vxlanPortUpdatingStatus>
```

## Customer Experience Improvement Program

NSX 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, or edit program settings, see [Edit the Customer Experience Improvement Program Option](#).

### Edit the Customer Experience Improvement Program Option

When you install or upgrade NSX Manager, you can choose to join the CEIP. You can join or leave the CEIP at a later time. You can also define the frequency and the days the information is collected.

#### Prerequisites

- Verify that the NSX manager is connected and can sync with vCenter Server.
- Verify that DNS is configured on NSX Manager.
- Verify that NSX is connected to a public network for uploading data.

#### Procedure

- 1 In the vSphere Web Client, navigate to the Customer Experience Improvement Program settings.
  - In NSX 6.4.6 and later, click **Networking & Security > About NSX**.
  - In NSX 6.4.5 and earlier, click **Networking & Security > NSX Home > Summary**.
- 2 In the **Customer Experience Improvement Program** pane, click **Edit**.
- 3 Join the CEIP program.
  - In NSX 6.4.6 and later, next to **Program Status**, click **Join**.
  - In NSX 6.4.5 and earlier, click the **Join the VMware Customer Experience Improvement Program** check box.
- 4 (Optional) Configure the recurrence settings.

- 5 Click **Save** or **OK**.

## About NSX Logs

You can configure the syslog server and view technical support logs for each NSX component. Management plane logs are available through NSX Manager and data plane logs are available through vCenter Server. Hence, it is recommended that you specify the same syslog server for the NSX component and vCenter Server in order to get a complete picture when viewing logs on the syslog server.

For information on configuring a syslog server for hosts managed by a vCenter Server, see the appropriate version of vSphere documentation at <https://docs.vmware.com>.

**Note** Syslog or jump servers used to collect logs and access an NSX Distributed Logical Router (DLR) Control VM can't be on the logical switch that is directly attached to that DLR's logical interfaces.

**Table 23-1. NSX Logs**

Component	Description
ESXi Logs	These logs are collected as part of the VM support bundle generated from vCenter Server. For more information on ESXi log files, refer to vSphere documentation.
NSX Edge Logs	Use the <code>show log [follow   reverse]</code> command in the NSX Edge CLI. Download Technical Support Log bundle via NSX Edge UI.
NSX Manager Logs	Use the <code>show log</code> CLI command in the NSX Manager CLI. Download Technical Support Log bundle via the NSX Manager Virtual Appliance UI.
Routing Logs	See the <i>NSX Logging and System Events Guide</i> .
Firewall Logs	See <i>NSX Logging and System Events Guide</i> .
Guest Introspection Logs	See <i>NSX Logging and System Events Guide</i> .

## NSX Manager

To specify a syslog server, see [Configure a Syslog Server for NSX Manager](#).

To download technical support logs, see [Download Technical Support Logs for NSX](#).

## NSX Edge

To specify a syslog server, see [Configure Syslog Servers for NSX Edge](#).

To download technical support logs, see [Download Tech Support Logs for NSX Edge](#).

## NSX Controller

To specify a syslog server, see [Configure DNS, NTP, and Syslog for the NSX Controller Cluster](#).

To download technical support logs, see [Download Technical Support Logs for NSX Controller](#).

## Firewall

For more details, refer to [Firewall Logs](#).

## Audit Logs

Audit logs for operations tracked by a ticket includes the ticket ID. With the NSX ticket logger feature, you can track the changes you make with a ticket ID.

## Using NSX Ticket Logger

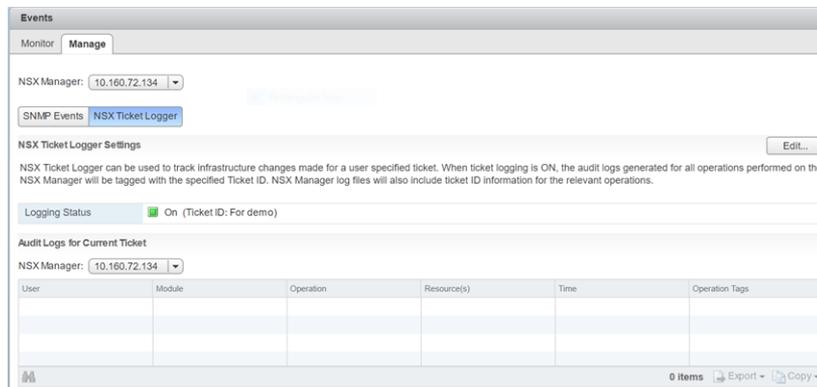
The NSX Ticket Logger allows you to track the infrastructure changes that you make. All operations are tagged with the specified ticket ID, and audit logs for these operations include the ticket ID. Log files for these operations are tagged with the same ticket ID.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > System > Events**.
- 2 Click the **Manage** tab, and then click **NSX Ticket Logger**.
- 3 Click **Edit** next to **NSX Ticket Logger Settings**.
- 4 Type a ticket ID and click **Turn On**.

The NSX Ticket Logging pane is displayed at the right side of the vSphere Web Client window. Audit logs for the operations that you perform in the current UI session include the ticket ID in the **Operation Tags** column.

Figure 23-1. NSX Ticket Logger pane



If multiple vCenter Servers are being managed by the vSphere Web Client, the ticket ID is used for logging on all applicable NSX Managers.

### What to do next

Ticket logging is session based. If ticket logging is on and you log out or if the session is lost, ticket logging will be turned off by default when you re-login to the UI. When you complete the operations for a ticket, you turn logging off by repeating steps 2 and 3 and clicking **Turn Off**.

## View the Audit Log

The **Audit Logs** tab provides a view into the actions performed by all NSX Manager users. The NSX Manager retains up to 100, 000 audit logs.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > System > Events**.
- 2 Ensure that you are in the **Monitor** tab.
- 3 If multiple IP addresses are available in the **NSX Manager** drop-down menu, select an IP address, or keep the default selection.

The audit log details are displayed in the **Audit Logs** tab.

- 4 When details are available for an audit log, the text in the **Operation** column for that log is clickable. To view details of an audit log, click the text in the **Operation** column.
- 5 In the **Audit Log Change Details**, select **Changed Rows** to display only those properties whose values have changed for this audit log operation.

## System Events

System events are events that are related to NSX operations. They are raised to detail every operational event. Events might relate to basic operation (Informational) or to a critical error (Critical).

### View the System Event Report

From vSphere Web Client you can view the system events for all the components that are managed by NSX Manager.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > System > Events**.
- 2 Ensure that you are in the **Monitor** tab.
- 3 Click the **System Events** tab.

You can click the arrows in the column headers to sort events, or use the **Filter** text box to filter events.

### Format of a System Event

If you specify a syslog server, NSX Manager sends all system events to the syslog server.

These messages have a format similar to the message displayed below:

```
Jan 8 04:35:00 NSXMGR 2017-01-08 04:35:00.422 GMT+00:00
INFO TaskFrameworkExecutor-18 SystemEventDaoImpl:133 -
[SystemEvent] Time:'Tue Nov 08 04:35:00.410 GMT+00:00 2016',
Severity:'High', Event Source:'Security Fabric', Code:'250024',
Event Message:'The backing EAM agency for this deployment could not be found.
It is possible that the VC services may still be initializing.
Please try to resolve the alarm to check existence of the agency.
In case you have deleted the agency manually, please delete the deployment
entry from NSX.', Module:'Security Fabric', Universal Object:'false
```

System event contains the following information.

```
Event ID and Time
Severity: Possible values include informational, low, medium, major, critical, high.
Event Source: Source where you should look to resolve the reported event.
Event Code: Unique identifier for the event.
Event Message: Text containing detailed information about the event.
Module: Event component. May be the same as event source.
Universal Object: Value displayed is True or False.
```

## Alarms

Alarms are notifications that are activated in response to an event, a set of conditions, or the state of an object. Alarms, along with other alerts, are displayed on the NSX Dashboard and other screens on the vSphere Web Client UI.

You can use the `GET api/2.0/services/systemalarms` API to view alarms on NSX objects.

NSX supports two methods for an alarm:

- Alarm corresponds to a system event and has an associated resolver that will attempt to resolve the issue that triggers the alarm. This approach is designed for network and security fabric deployment (for example, EAM, Message Bus, Deployment Plug-In), and is also supported by Service Composer. These alarms use the event code as the alarm code. For more details, refer to *NSX Logging and System Events* document.
- Edge notifications alarms are structured as a triggering and resolving alarm pair. This method is supported by several Edge functions, including IPsec VPN, load balancer, high availability, health check, edge file system, and resource reservation. These alarms use a unique alarm code which is not the same as the event code. For more details, refer to *NSX Logging and System Events* document.

Generally, an alarm gets automatically deleted by the system when the error condition is rectified. Some alarms are not auto cleared on a configuration update. Once the issue is resolved, you have to clear the alarms manually.

Here is an example of the API that you can use to clear the alarms.

You can get alarms for a specific source, for example, cluster, host, resource pool, security group, or NSX Edge. View alarms for a source by *sourceId*:

```
GET https://<<NSX-IP>>/api/2.0/services/alarms/{sourceId}
```

Resolve all alarms for a source by *sourceId*:

```
POST https://<<NSX-IP>>/api/2.0/services/alarms/{sourceId}?action=resolve
```

You can view NSX alarms, including Message Bus, Deployment Plug-In, Service Composer, and Edge alarms:

```
GET https://<<NSX-IP>>/api/2.0/services/systemalarms
```

You can view a specific NSX alarm by *alarmId*:

```
GET https://<<NSX-IP>>/api/2.0/services/systemalarms/<alarmId>
```

You can resolve a specific NSX alarm by *alarmId*:

```
POST https://<<NSX-IP>>/api/2.0/services/systemalarms/<alarmId>?action=resolve
```

For more information on API, refer to *NSX API Guide*.

## Format of an Alarm

You can view format of an alarm through API.

The format of an alarm contains the following information.

```
Event ID and Time
Severity: Possible values include informational, low, medium, major, critical, high.
Event Source: Source where you should look to resolve the reported event.
Event Code: Unique identifier for the event.
Message: Text containing detailed information about the event.
Alarm ID: ID of an alarm.
Alarm Code: Event code which uniquely identifies the system alarm.
Alarm Source: Source where you should look to resolve the reported event.
```

## Working with SNMP Traps

The NSX Manager receives system events that are informational, warning, and critical from for example, the NSX Edge and hypervisor. The SNMP agent forwards the SNMP traps with OIDs to the SNMP receiver.

SNMP traps must have the SNMPv2c version. The traps must be associated with a management information base (MIB) so that the SNMP receiver can process the traps with object identifiers (OID).

By default, the SNMP trap mechanism is disabled. Enabling the SNMP trap only activates the critical and high severity notifications so that the SNMP manager does not get inundated by a high volume of notifications. An IP address or a host name defines the trap destination. For the host name to work for the trap destination, the device must be set up to query a Domain Name System (DNS) server.

When you enable the SNMP service, a coldStart trap with OID 1.3.6.1.6.3.1.1.5.1 is sent out the first time. A warmStart trap with OID 1.3.6.1.6.3.1.1.5.2 is sent out later on each stop-start to the configured SNMP receivers.

If the SNMP service remains enabled, a heartbeat trap vmwHbHeartbeat with OID 1.3.6.1.4.1.6876.4.190.0.401 is sent out every five minutes. When you disable the service, a vmwNxmSnmPDisabled trap with OID 1.3.6.1.4.1.6876.90.1.2.1.0.1 is sent out. This process stops the vmwHbHeartbeat trap from running and disables the service.

When you add, modify, or delete a SNMP receiver value, a warmStart trap with OID 1.3.6.1.6.3.1.1.5.2 and vmwNxmSnmPManagerConfigUpdated trap with OID 1.3.6.1.4.1.6876.90.1.2.1.0.2 is sent to the new or updated set of SNMP receivers.

---

**Note** SNMP Polling is not supported.

---

## Configure SNMP Settings

You can enable the SNMP settings and configure destination receivers to send traps that are critical, high, or informational.

### Prerequisites

- Familiarize yourself with the SNMP trap mechanism. See [Working with SNMP Traps](#).
- Verify that an SNMP receiver is configured.
- Download and install the MIB module for the NSX Manager so that the SNMP receiver can process the traps with OID. See <http://kb.vmware.com/kb/1013445>.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > System > Events**.
- 2 Click the **Manage** tab. If multiple NSX Managers are available, select an IP address of an NSX Manager from the **NSX Manager** drop-down menu.
- 3 Ensure that the **SNMP Events** tab is selected.

- 4 Click **Edit** to configure the SNMP settings.

Option	Description
<b>Service</b>	Sends out SNMP trap. By default, this option is disabled.
<b>Group Notification</b>	Predefined set of groups for some system events that are used to aggregate the events that are raised. By default, this option is enabled.  For example, if a system event belongs to a group, the trap for these grouped events are withheld. Every five minutes a trap is sent out detailing the number of system events that have been received from the NSX Manager. The fewer traps being sent out save the SNMP receiver resources.
<b>Receivers</b>	Configure up to four receivers for traps to be sent out to. You must complete the following sections when you add an SNMP receiver. Receiver Address - IP address or the fully qualified domain name of the receiver host. Receiver Port - SNMP receiver default UDP port is 162. Community String - Information to be sent out as part of the notification trap. Enabled - Indicates whether this receiver is sending a trap.

- 5 Click **OK**.

### Results

The SNMP service is enabled and traps are sent out to the receivers.

### What to do next

Check whether the SNMP configuration works. See [Verify SNMP Trap Configuration](#).

## Verify SNMP Trap Configuration

Before you start editing an existing system trap, you must check whether the newly enabled SNMP service or updated SNMP is working properly.

### Prerequisites

Verify that you have SNMP configured. See [Configure SNMP Settings](#).

### Procedure

- 1 Verify SNMP configuration and receiver connection.
  - a Click the **Manage > SNMP Events** tabs.
  - b Click **Edit** to configure the SNMP settings.  
Do not change the settings in the dialog box.
  - c Click **OK**.

A warmStart trap with OID 1.3.6.1.6.3.1.1.5.2 is sent out to all the SNMP receivers.

## 2 Debug SNMP configuration or receiver problems.

- a If the SNMP receiver does not receive the traps, verify that the SNMP receiver is running on a configured port.
- b Check the accuracy of the receiver details under the SNMP settings section.
- c If the SNMP receiver stops receiving a vmwHbHeartbeat trap with OID 1.3.6.1.4.1.6876.4.190.0.401 every five minutes, check whether the NSX Manager appliance or the NSX Manager SNMP agent is working.
- d If the Heartbeat trap stops, check whether the SNMP service is disabled or test whether the network connectivity between the NSX Manager and the SNMP receiver is working.

## Edit System Traps

You can edit a system trap to increase or decrease the severity and enablement of a trap so that traps are either sent out to receivers or withheld.

When the Module, SNMP OID, or SNMP trap enabled column value appears as --, it means that those events have not been allocated a trap OID. Therefore, a trap for these events is not going to be sent out.

A system trap has several columns that list different aspects of a system event.

Option	Description
Event Code	Static event code associated with an event.
Description	Summary describing the event.
Module	Sub component that triggers an event.
Severity	Level of an event can be informational, low, medium, major, critical, or high. By default when the SNMP service is enabled, traps are sent out for only critical and high severity events to highlight the traps that require immediate attention.
SNMP OID	Represents the individual OID and this OID is sent out when a system event is raised. Group notification is enabled by default. When group notifications is enabled, the events or traps under this group show the OID of the group the event or trap belongs to. For example, group notification OID categorized under configuration group has the OID 1.3.6.1.4.1.6876.90.1.2.0.1.0.1.
SNMP trap enabled	Shows whether sending out of the trap for this event is enabled or disabled. You can toggle the icon to individually an event or trap enablement. When group notification is enabled, you cannot toggle the trap enablement.
Filter	Search terms to filter the system traps.

### Prerequisites

Verify that the SNMP settings are available. See [Configure SNMP Settings](#).

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > System > Events**.
- 2 Click the **Manage** tab, and then select an NSX Manager IP address.

- 3 In the System Traps section, select a system event.
- 4 Click the **Edit** (✎) icon.  
Editing a trap enablement is not allowed when group notification is enabled. You can change the enablement of traps that do not belong to a group.
- 5 Change the severity of the system event from the drop-down menu.
- 6 If you change the severity from Informational to critical, check the **Enable as SNMP Trap** checkbox.
- 7 Click **OK**.
- 8 (Optional) Click the **Enable** (✔) icon or **Disable** (⊘) icon in the header to enable or disable sending a system trap.
- 9 (Optional) Click the **Copy** (📄) icon to copy one or more event rows to your clipboard.

## Management System Settings

You can edit the vCenter Server, DNS and NTP server, and Lookup server that you specified during initial login. NSX Manager requires communication with your vCenter Server and services such as DNS and NTP to provide details on your VMware Infrastructure inventory.

## Log In to the NSX Manager Virtual Appliance

After you have installed and configured the NSX Manager virtual machine, log in to the NSX Manager virtual appliance to review the settings specified during installation.

### Procedure

- 1 Open a Web browser window and type the IP address assigned to the NSX Manager. For example, `https://192.168.110.42`.

The NSX Manager user interface opens in a web browser window using SSL.

- 2 Accept the security certificate.

---

**Note** You can use an SSL certificate for authentication.

---

The NSX Manager login screen appears.

- 3 Log in to the NSX Manager virtual appliance by using the user name **admin** and the password you set during installation.
- 4 Click **Log In**.

## NSX Manager Virtual Appliance Events

The following events are specific to the NSX Manager virtual appliance.

Table 23-2. NSX Manager Virtual Appliance Events

	Power Off	Power On	Interface Down	Interface Up
Local CLI	Run <code>show log follow</code> command.			
GUI	NA	NA	NA	NA

Table 23-3. NSX Manager Virtual Appliance Events

	CPU	Memory	Storage
Local CLI	Run <code>show process monitor</code> command.	Run <code>show system memory</code> command.	Run <code>show filesystem</code> command.
GUI	NA	NA	NA

## Edit the NSX Manager Date and Time

You can change the NTP server specified during initial login.

### Procedure

- 1 Log in to the NSX Manager virtual appliance.
- 2 Under **Appliance Management**, click **Manage Appliance Settings**.
- 3 Click **Edit** next to **Time Settings**.
- 4 Make the appropriate changes.
- 5 Click **OK**.
- 6 Reboot the NSX Manager.

## Change NSX Manager Appliance IP Address

Starting in NSX 6.4.0, you can change the IP address of an NSX Manager appliance. You can change the IP address for all types of NSX Manager: standalone, primary, or secondary.

**Important** Changing the host name of NSX Manager is not supported.

If your partner solution configurations reference the NSX Manager IP address, and you change the IP address of NSX Manager, you must update your partner solution. See [Update Partner Solutions After NSX Manager IP Change](#).

If you change the network configuration of a secondary NSX Manager in a cross-vCenter NSX environment, you must update the secondary NSX Manager configuration on the primary NSX Manager. See [Update Secondary NSX Manager on Primary NSX Manager](#).

### Prerequisites

- Verify that the new IP address is added to DNS, and is resolvable from all related vCenter Server, Platform Services Controller, and ESXi systems.

## Procedure

- 1 Log in to the NSX Manager virtual appliance.

In a Web browser, navigate to the NSX Manager appliance GUI at <https://<nsx-manager-ip>> or <https://<nsx-manager-hostname>>, and log in as **admin** or with an account that has the **Enterprise Administrator** role.

- 2 From the home page, click **Manage Appliance Settings > Network**.
- 3 Click **Edit** in the General network settings pane and update the IPv4 or IPv6 configuration sections.
- 4 Reboot the NSX Manager appliance. Click Actions () and then **Reboot Appliance**.
- 5 After the NSX Manager appliance has rebooted, log in to the web interface, and navigate to **Home > View Summary**. Verify that the NSX Management Service has status of RUNNING. In a cross-vCenter NSX environment, also verify that the NSX Universal Synchronization Service has status of RUNNING.

If you have problems connecting to the NSX Manager appliance web interface once the NSX Manager appliance VM is running, your computer or browser might have cached the old IP address. Close all browser windows, start your browser again, and clear the cache. Or access the appliance using its new IP address instead of its hostname.

- 6 Log out of the vSphere Web Client and log back in.
- 7 Navigate to **Networking & Security > Installation and Upgrade**.

It can take several minutes for NSX Manager to connect to the vCenter Server system after the network configuration change and reboot. If you see an error message: "No NSX Managers found", log out of the vSphere Web Client, wait a few minutes, and log back in again.

- 8 Click **Host Preparation**. Each host cluster displays Not Ready. Click **Not Ready** and then click **Resolve all**.

The IP change causes the URL used to access VIBs to change. When you click **Resolve All**, the URL is updated.

### What to do next

If your partner solution configurations reference the NSX Manager IP address, update your partner solution. See [Update Partner Solutions After NSX Manager IP Change](#).

If you have changed the network configuration of a secondary NSX Manager, update the secondary NSX Manager configuration on the primary NSX Manager. See [Update Secondary NSX Manager on Primary NSX Manager](#).

## Update Partner Solutions After NSX Manager IP Change

If your partner solution configurations reference the NSX Manager IP address, and you change the IP address of NSX Manager, you must update your partner solution.

### Prerequisites

- Verify that the NSX Manager IP address has been changed. See [Change NSX Manager Appliance IP Address](#).
- Consult the partner documentation for instructions on updating the NSX Manager IP address in the partner solution configuration.

### Procedure

- 1 If the partner solution configuration references the NSX Manager IP address, update the partner solution with the new NSX Manager IP address.  
See the partner documentation for more information.
- 2 Log in to the vSphere Web Client.
- 3 Navigate to **Installation and Upgrade > Service Deployments**.
- 4 If any of your partner services use Guest Introspection, on the Guest Introspection Installation status columns, click **Not Ready** then click **Resolve all**.
- 5 On the partner solution Installation status columns, click **Not Ready** then click **Resolve all**.

## Update Secondary NSX Manager on Primary NSX Manager

If you change the IP address of a secondary NSX Manager, the primary NSX Manager displays sync errors until you update the secondary NSX Manager IP address information on the primary NSX Manager.

### Prerequisites

- Verify that you have changed the IP address on a secondary NSX Manager in a cross-vCenter NSX environment. See [Change NSX Manager Appliance IP Address](#).

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security** and then click **Installation and Upgrade**.
- 3 Click the **Management** tab. In the NSX Managers pane, click **Actions > Update Secondary NSX Manager**. Enter the new IP address assigned to the secondary NSX Manager.
- 4 Verify that the thumbprint displayed is the correct thumbprint for the secondary NSX Manager. Click **OK**.

## Configure a Syslog Server for NSX Manager

If you specify a syslog server, NSX Manager sends all audit logs and system events to the syslog server. NSX Manager supports five syslog servers.

Syslog data is useful for troubleshooting and reviewing data logged during installation and configuration.

## Procedure

- 1 Log in to the NSX Manager virtual appliance.

In a Web browser, navigate to the NSX Manager appliance GUI at `https://<nsx-manager-ip>` or `https://<nsx-manager-hostname>`, and log in as **admin** or with an account that has the **Enterprise Administrator** role.

- 2 From the home page, click **Manage Appliance Settings > General**.
- 3 Click **Edit** next to **Syslog Server**.
- 4 Specify the IP address or hostname, port, and protocol of the syslog server.

For example:

Syslog Server	Port	Protocol	
syslog-01a.corp.local	514	UDP	x

- 5 Click **OK**.

## Results

NSX Manager remote logging is enabled, and logs are stored in your syslog server. If you have configured multiple syslog servers, logs are stored in all the configured syslog servers.

## What to do next

For more details on API, refer to *NSX API Guide*.

## Change FIPS Mode and TLS Settings on NSX Manager

When you enable the FIPS mode, any secure communication to or from the NSX Manager will use cryptographic algorithms and protocols that are allowed by the United States Federal Information Processing Standards (FIPS).

- In a Cross-vCenter NSX environment, you should enable the FIPS mode on each NSX Manager separately.
- If one of the NSX Managers is not configured for FIPS, you must still ensure that it uses a secure communication method which complies with the FIPS standards.
- Both primary and secondary NSX Managers must be on the same TLS version for universal synchronization to work correctly.

---

**Important** Changing FIPS mode reboots the NSX Manager virtual appliance.

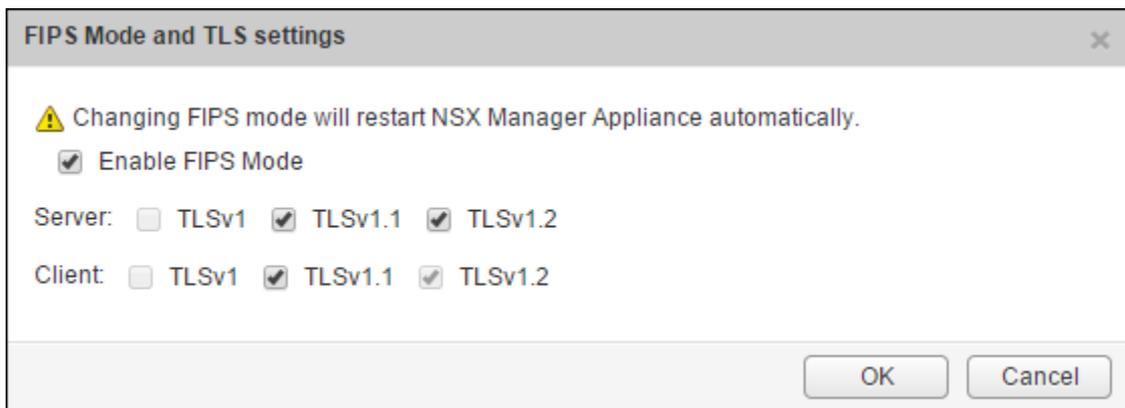
---

## Prerequisites

- Verify that any partner solutions are FIPS mode certified. See the VMware Compatibility Guide at <http://www.vmware.com/resources/compatibility/search.php?deviceCategory=security>.
- If you have upgraded from an earlier version of NSX, do not enable FIPS mode until the upgrade to NSX 6.3.0 is complete. See Understand FIPS Mode and NSX Upgrade in the *NSX Upgrade Guide*.
- Verify that the NSX Manager is NSX 6.3.0 or later.
- Verify that the NSX Controller cluster is NSX 6.3.0 or later.
- Verify that all host clusters running NSX workloads are prepared with NSX 6.3.0 or later.
- Verify that all NSX Edge appliances are version 6.3.0 or later, and that FIPS mode has been enabled on the required NSX Edge appliances. See [Change FIPS Mode on NSX Edge](#).

## Procedure

- 1 Log in to the NSX Manager virtual appliance.
- 2 Under **Appliance Management**, click **Manage Appliance Settings**.
- 3 From the Settings panel, click **General**.
- 4 Click **Edit** next to **FIPS Mode and TLS settings**.



- 5 To enable FIPS mode, select the **Enable FIPS Mode** check box.
- 6 For Server and Client, select the check boxes for the required TLS protocol version.

### Note

- When FIPS mode is enabled, NSX Manager disables the TLS protocols that are not compliant to the FIPS standards.
- In NSX 6.4.0 or later, TLS 1.0 is disabled by default.  
If you upgrade to NSX 6.4.0 or later, the TLS settings before upgrade remains unchanged.

- 7 Click **OK**.

The NSX Manager appliance reboots, and FIPS is enabled.

## Edit DNS Servers

You can change the DNS servers specified during Manager installation.

### Procedure

- 1 Log in to the NSX Manager virtual appliance.
- 2 Under **Appliance Management**, click **Manage Appliance Settings**.
- 3 From the Settings panel, click **Network**.
- 4 Click **Edit** next to **DNS Servers**.
- 5 Make the appropriate changes.
- 6 Click **OK**.

## Edit Lookup Service Details

You can change the Lookup Service details specified during initial login.

### Procedure

- 1 Log in to the NSX Manager virtual appliance.
- 2 Under **Appliance Management**, click **Manage Appliance Settings**.
- 3 From the Settings panel, click **NSX Management Service**.
- 4 Click **Edit** next to **Lookup Service**.
- 5 Make the appropriate changes.
- 6 Click **OK**.

## Edit vCenter Server

You can change the vCenter Server with which you registered NSX Manager during installation. You should do this only if you change the IP address of your current vCenter Server.

### Procedure

- 1 If you are logged in to the vSphere Web Client, log out.
- 2 Log in to the NSX Manager virtual appliance.
- 3 Under **Appliance Management**, click **Manage Appliance Settings**.
- 4 From the Settings panel, click **NSX Management Service**.
- 5 Click **Edit** next to **vCenter Server**.
- 6 Make the appropriate changes.
- 7 Click **OK**.

## Download Technical Support Logs for NSX

You can download NSX Manager system logs and Web Manager logs to your desktop. You can also collect the support bundle data for NSX Manager using the **Support Bundle** collection tool. For details, refer to *NSX Administration Guide*.

### Procedure

- 1 Log in to the NSX Manager virtual appliance.
- 2 Under Appliance Management, click **Manage Appliance Settings**.
- 3 Click  and then click **Download Tech Support Log**.
- 4 Click **Download**.
- 5 After the log is ready, click the **Save** to download the log to your desktop.

The log is compressed and has the file extension `.gz`.

### What to do next

You can open the log using a decompression utility by browsing for **All Files** in the directory where you saved the file.

## NSX Manager SSL Certification

NSX Manager requires a signed certificate to authenticate the identity of the NSX Manager web service and encrypt information sent to the NSX Manager web server. The process entails generating a certificate signing request (CSR), getting it signed by a CA, and importing the signed SSL certificate into NSX Manager. As a security best practice, you should use the generate certificate option to generate a private key and public key, where the private key is saved to the NSX Manager.

To obtain the NSX Manager certificate, you can use NSX Manager's built-in CSR generator or you can use another tool such as OpenSSL.

A CSR generated using NSX Manager's built-in CSR generator cannot contain extended attributes such as subject alternate name (SAN). If you wish to include extended attributes, you must use another CSR generation tool. If you are using another tool such as OpenSSL to generate the CSR, the process is 1) generate the CSR, 2) have it signed, and 3) proceed to the section [Convert the NSX Manager Certificate File to PKCS 12 Format](#).

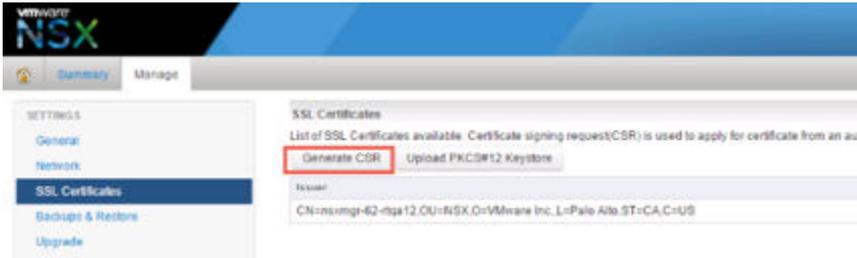
### Use the Built-In CSR Generator

One method of obtaining an SSL certificate for NSX Manager is use the built-in CSR generator.

This method is limited in that the CSR cannot contain extended attributes such as subject alternate name (SAN). If you wish to include extended attributes, you must use another CSR generation tool. If you are using another CSR generation tool, skip this procedure.

## Procedure

- 1 Log in to the NSX Manager virtual appliance.
- 2 Click **Manage Appliance Settings**.
- 3 From the Settings panel, click **SSL Certificates**.
- 4 Click **Generate CSR**.



- 5 Complete the form by filling in the following fields:

Option	Action
Key Size	Select the key length used in the selected algorithm.
Common Name	Type the IP address or fully qualified domain name (FQDN) of the NSX Manager. VMware recommends that you enter the FQDN.
Organization Unit	Enter the department in your company that is ordering the certificate.
Organization Name	Enter the full legal name of your company.
City Name	Enter the full name of the city in which your company resides.
State Name	Enter the full name of the state in which your company resides.
Country Code	Enter the two-digit code that represents your country. For example, the United States is <b>US</b> .

- 6 Click **OK**.
- 7 Send the CSR to your CA for signing.
  - a Download the generated request by clicking **Download CSR**.  
Using this method, the private key never leaves the NSX Manager.
  - b Submit this request to your CA.
  - c Get the Signed Certificate and Root CA and any intermediary CA certificates in PEM format.
  - d To convert CER/DER formatted certificates to PEM, use the following OpenSSL command:

```
openssl x509 -inform der -in Cert.cer -out 4-nsx_signed.pem
```

- e Concatenate all the certificates (server, intermediary and root certificates) in a text file.

- f In the NSX Manager UI, click **Import** and browse to the text file with all of the certificates.
- g Once the import is successful, the server certificate and all the CA certificates will be shown on the SSL Certificates page.

#### What to do next

Import the signed SSL certificate into NSX Manager.

## Convert the NSX Manager Certificate File to PKCS 12 Format

If you used another tool, such as OpenSSL, to create the NSX Manager certificate, make sure that the certificate and private key are in the PKCS 12 format. If the NSX Manager certificate and private key are not in the PKCS 12 format, you must convert them to PKCS 12 format and then import the PKCS 12 certificate file into NSX Manager.

#### Prerequisites

- Verify that OpenSSL is installed on the system. You can download openssl from <http://www.openssl.org>.
- Generate a public and private key pair. For example, run the following OpenSSL command:

```
openssl req -x509 -days [number of days] -newkey rsa:2048 -keyout my-key.pem -out my-cert.pem
```

#### Procedure

- ◆ After receiving the signed certificate from the authorized signer, run an OpenSSL command to generate a PKCS 12 (.pfx or .p12) keystore file from the public certificate file and your private key.

For example:

```
openssl pkcs12 -export -in my-cert.pem -inkey my-key.pem -out nsx-manager.p12
```

Where:

- `my-cert.pem` is the signed certificate.
- `my-key.pem` is the private key.
- `nsx-manager.p12` is the name of the generated output file after the conversion to PKCS 12 format.

#### What to do next

Import the PKCS 12 certificate file into NSX Manager.

## Import an SSL Certificate

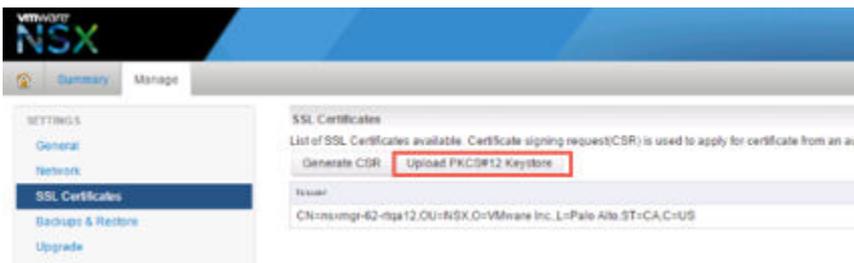
You can import a pre-existing or CA-signed SSL certificate for use by the NSX Manager.

## Prerequisites

When installing a certificate on NSX Manager, only the PKCS#12 keystore format is supported, and it must contain a single private key and its corresponding signed certificate or certificate chain.

## Procedure

- 1 Log in to the NSX Manager virtual appliance.
- 2 Click **Manage Appliance Settings**.
- 3 From the Settings panel, click **SSL Certificates**.
- 4 Click **Upload PKCS#12 Keystore**.



- 5 Click **Choose File** to locate the file.
- 6 Click **Import**.
- 7 To apply the certificate, reboot the NSX Manager appliance.

## Results

The certificate is stored in NSX Manager.

# NSX Backup and Restore

Proper backup of all NSX Data Center for vSphere components is crucial to restore the system to its working state in the event of a failure.

The NSX Manager backup contains all of the NSX Data Center for vSphere configuration, including controllers, logical switching and routing entities, security, firewall rules, and everything else that you configure within the NSX Manager UI or API. The vCenter database and related elements like the virtual switches need to be backed up separately.

At a minimum, we recommend taking regular backups of NSX Manager and vCenter. Your backup frequency and schedule might vary based on your business needs and operational procedures. We recommend taking NSX backups frequently during times of frequent configuration changes.

NSX Manager backups can be taken on demand or on an hourly, daily, or weekly basis. The backup checksum file is produced with the SHA256 algorithm.

We recommend taking backups in the following scenarios:

- Before an NSX Data Center for vSphere or vCenter upgrade.

- After an NSX Data Center for vSphere or vCenter upgrade.
- After Day Zero deployment and initial configuration of NSX Data Center for vSphere components, such as after the creation of the NSX Controller cluster, logical switches, logical routers, edge services gateways, security, and firewall policies.
- After infrastructure or topology changes.
- After any major Day 2 change.

To provide an entire system state at a given time to roll back to, we recommend synchronizing NSX Data Center for vSphere component backups (such as NSX Manager) with your backup schedule for other interacting components, such as vCenter, cloud management systems, operational tools, and so on.

## Back Up and Restore NSX Manager

NSX Manager backup and restore can be configured from the NSX Manager virtual appliance web interface or through the NSX Manager API. Backups can be scheduled on an hourly, daily or weekly basis.

The backup file is saved to a remote FTP or SFTP location that NSX Manager can access. NSX Manager data includes configuration, events, and audit log tables. Configuration tables are included in every backup.

Restore is only supported on the same NSX Manager version as the backup version. For this reason, it is important to create a backup file before and after performing an NSX upgrade, one backup for the old version and another backup for the new version.

### Back Up NSX Manager Data

You can back up NSX Manager data by performing an on-demand backup or a scheduled backup.

#### Procedure

- 1 Log in to the NSX Manager Virtual Appliance.
- 2 Click **Backup & Restore**.
- 3 To specify the backup location, click **Change** next to FTP Server Settings.
  - a Enter the IP address or host name of the backup system.
  - b From the **Transfer Protocol** drop-down menu, select either **SFTP** or **FTP**, based on what the destination supports.
  - c Edit the default port if necessary.
  - d Enter the user name and password required to log in to the backup system.

- e In the **Backup Directory** text box, enter the absolute path for storing the backups.

**Note** If you do not provide a backup directory, backup is stored to the default directory (home directory) of the FTP server.

To determine the absolute path, log in to the FTP server, navigate to the directory that you want to use, and run the present working directory command (`pwd`). For example:

```
PS C:\Users\Administrator> ftp 192.168.110.60
Connected to 192.168.110.60.
220 server-nfs FTP server ready.
User (192.168.110.60:(none)): admin
331 Password required for admin.
Password:
230 User admin logged in.
ftp> ls
200 PORT command successful.
150 Opening BINARY mode data connection for 'file list'.
datastore-01
226 Transfer complete.
ftp: 22 bytes received in 0.00Seconds 22000.00Kbytes/sec.
ftp> cd datastore-01
250 CWD command successful.
ftp> pwd
257 "/datastore-01" is current directory.
```

- f Enter a text string in **Filename Prefix**.

This text is prepended to each backup filename to help you easily recognize the file on the backup system. For example, if you enter `ppdb`, the resulting backup file is named as `ppdbHH_MM_SS_YYYY_Mon_Day`.

**Note** Files in the Backup Directory must be limited to 100. If the number of files in the directory exceeds this limit, a warning message is displayed.

- g Enter a pass phrase to secure the backup.

Pass phrase is required to restore the backup.

- h Click **OK**.

For example:

Option	Example
IP/Host name	192.168.110.60
Transfer Protocol	FTP
Port	21
User name	admin
Password	*****
Backup Directory	/datastore-01

Option	Example
Filename Prefix	nsxmgr-backup
Pass Phrase	*****

- 4 For an on-demand backup, click **Backup**.

A new file is added under **Backup History**.

- 5 (Required) For scheduled backups, click **Change** next to Scheduling.

- From the **Backup Frequency** drop-down menu, select **Hourly**, **Daily**, or **Weekly**. The Day of Week, Hour of Day, and Minute drop-down menus are disabled based on the selected frequency. For example, if you select Daily, the Day of Week drop-down menu is disabled because this drop-down menu is not applicable to a daily frequency.
- For a weekly backup, select the day of the week the data should be backed up.
- For a weekly or daily backup, select the hour at which the backup should begin.
- Select the minute to begin and click **Schedule**.

Option	Example
Backup Frequency	Weekly
Day of week	Friday
Hour of day	15
Minute	45

- 6 To exclude logs and flow data from being backed up, click **Change** next to Exclude.

- Select the items that you want to exclude from the backup.
- Click **OK**.

- 7 Save your FTP server IP/hostname, credentials, directory details, and pass phrase. This information is required to restore the backup.

## Restore an NSX Manager Backup

Restoring NSX Manager causes a backup file to be loaded on an NSX Manager appliance. The backup file must be saved to a remote FTP or SFTP location that NSX Manager can access. NSX Manager data includes configuration, events, and audit log tables.

---

**Important** Back up your current data before restoring a backup file.

---

### Prerequisites

Before restoring NSX Manager data, we recommend reinstalling the NSX Manager appliance. Running the restore operation on an existing NSX Manager appliance might work, too, but is not supported. The assumption is that the existing NSX Manager has failed, and therefore a new NSX Manager appliance is deployed.

The best practice is to take note of the current settings for the old NSX Manager appliance so that they can be used to specify IP information and backup location information for the newly deployed NSX Manager appliance.

### Procedure

1 Take note of all settings on the existing NSX Manager appliance. Also, note down FTP server settings.

2 Deploy a new NSX Manager appliance.

The version must be the same as the backed up NSX Manager appliance.

3 Log in to the new NSX Manager appliance.

4 Under Appliance Management, click **Backups & Restore**.

5 In FTP Server Settings, click **Change** and add the FTP server settings.

The **Host IP Address**, **User Name**, **Password**, **Backup Directory**, **Filename Prefix**, and **Pass Phrase** fields in the Backup Location screen must identify the location of the backup to be restored.

The **Backup History** section displays the backup folder.

---

**Note** If the backup folder does not appear in the **Backup History** section, verify the FTP server settings. Check if you can connect to FTP server and view the backup folder.

---

6 In the **Backup History** section, select the required backup folder to restore, and click **Restore**.

Restoring the NSX Manager data begins.

### Results

NSX configuration is restored to the NSX Manager.

---

**Caution** After restoring an NSX Manager backup, you might need to take additional action to ensure correct operation of NSX Edge appliances and logical switches. See [Restore NSX Edges](#) and [Resolve Out of Sync Errors on Logical Switches](#).

---

## Restore NSX Edges

All NSX Edge configurations (logical routers and edge services gateways) are backed up as part of NSX Manager data backup.

Taking individual NSX Edge backups is not supported.

If you have an intact NSX Manager configuration, you can recreate an inaccessible or failed Edge appliance VM by redeploying the NSX Edge. To redeploy an NSX Edge, select the NSX Edge, and click **Actions > Redeploy**. See "Redeploy NSX Edge" in the *NSX Administration Guide*.

---

**Caution** After restoring an NSX Manager backup, you might need to take additional action to ensure correct operation of NSX Edge appliances.

- Edge appliances created after last backup are not removed during restore. You must delete the VM manually.
- Edge appliances deleted after the last backup are not restored unless redeployed.
- If both the configured and current locations of an NSX Edge appliance saved in the backup no longer exist when the backup is restored, operations such as redeploy, migrate, enable or disable HA will fail. You must edit the appliance configuration and provide a valid location information. Use `PUT /api/4.0/edges/{edgeId}/appliances` to edit the appliance location configuration (*resourcePoolId*, *datastoreId*, *hostId* and *vmFolderId* as necessary). See "Working With NSX Edge Appliance Configuration" in the *NSX API Guide*.

If any of the following changes have occurred since the last NSX Manager backup, the restored NSX Manager configuration and the configuration present on the NSX Edge appliance will differ. You must **Force Sync** the NSX Edge to revert these changes on the appliance and ensure correct operation of the NSX Edge. See "Force Sync NSX Edge with NSX Manager" in the *NSX Administration Guide*.

- Changes made via Distributed Firewall for preRules for NSX Edge firewall.
- Changes in grouping objects membership.

If any of the following changes have occurred since the last NSX Manager backup, the restored NSX Manager configuration and the configuration present on the NSX Edge appliance will differ. You must **Redeploy** the NSX Edge to revert these changes on the appliance and ensure correct operation of the NSX Edge. See "Redeploy NSX Edge" in the *NSX Administration Guide*.

- Changes in Edge appliance settings:
  - HA enabled or disabled
  - appliance moved from deployed to undeployed state
  - appliance moved from undeployed to deployed state
  - resource reservation settings have been changed
- Changes in Edge appliance vNic settings:
  - add, remove, or disconnect vNic
  - port groups
  - trunk ports
  - fence parameters
  - shaping policy

---

**Attention** After upgrading to NSX 6.4.4 or 6.4.5, the default MTU value for newly added trunk interfaces on the edge is incorrectly set to 1500. This issue also occurs after you do a fresh installation of 6.4.4 or 6.4.5. The issue is fixed in 6.4.6. However, to resolve this issue in 6.4.4 or 6.4.5, you must manually change the default MTU value in all the trunk interfaces of the edge to 1600. For more information, see the VMware knowledge base article at [https://](https://kb.vmware.com/s/article/74878)

## Resolve Out of Sync Errors on Logical Switches

If logical switch changes have occurred between taking the NSX Manager backup and restoring the backup, logical switches might report being out of sync.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Navigate to **Networking & Security > Logical Switches**.
- 3 If present, click the **Out of sync** link in the Status column to display error details.
- 4 Click **Resolve** to recreate missing backing port groups for the logical switch.

## Back Up vSphere Distributed Switches

You can export vSphere Distributed Switch and distributed port group configurations to a file.

Export the vSphere Distributed Switch configuration to create a backup before preparing the cluster for VXLAN. For detailed instructions about exporting a vSphere Distributed Switch configuration, see <http://kb.vmware.com/kb/2034602>.

## Back Up vCenter

To secure your NSX deployment, it is important to back up the vCenter database and take snapshots of the VMs.

Refer to the vCenter documentation for your vCenter version for vCenter backup and restore procedures and best practices.

For vCenter Backups, see the following:

- The *vSphere Installation and Setup* documentation for your version of vSphere
- <http://kb.vmware.com/kb/2110294>

For VM snapshots, see <http://kb.vmware.com/kb/1015180>.

## NSX Monitoring And Diagnostic Tools

NSX provides different tools that can help you to monitor, and collect data to diagnose any problem with your system.

### Flow Monitoring

Flow monitoring is a traffic analysis tool that provides a detailed view of the traffic to and from protected virtual machines.

When flow monitoring is enabled, its output defines which machines are exchanging data and over which application. This data includes the number of sessions and packets transmitted per session. Session details include sources, destinations, applications, and ports being used. Session details can be used to create firewall to allow or block rules. You can view flow data for many

different protocol types, including TCP, UDP, ARP, ICMP, and so on. Flows can be excluded by specifying filters. You can live monitor TCP and UDP connections to and from a selected vNIC. Live flow monitoring provides visualization of flows as they traverse a specific vNIC, enabling quick troubleshooting. Application context is also captured in live flow monitoring for flows that match an L7 firewall rule.

## View Flow Monitoring Data

You can view traffic sessions on virtual machines within the specified time span. The last 24 hours of data is displayed by default, the minimum time span is one hour, and the maximum is two weeks.

---

### Caution

- When flow monitoring is enabled, the Dashboard shows a small yellow warning icon to indicate that the feature is turned on. Flow monitoring impacts performance and you must preferably turn it off after monitoring the flow data.
  - A critical alarm is displayed when the flow monitoring count exceeds a predefined maximum count (threshold value). This critical alarm does not impact your environment, and you can safely ignore the alarm. In NSX 6.4.4 and earlier, the maximum flow monitoring count is set to 2 million. Starting in NSX 6.4.5, the maximum flow monitoring count is increased to 5 million.
- 

### Prerequisites

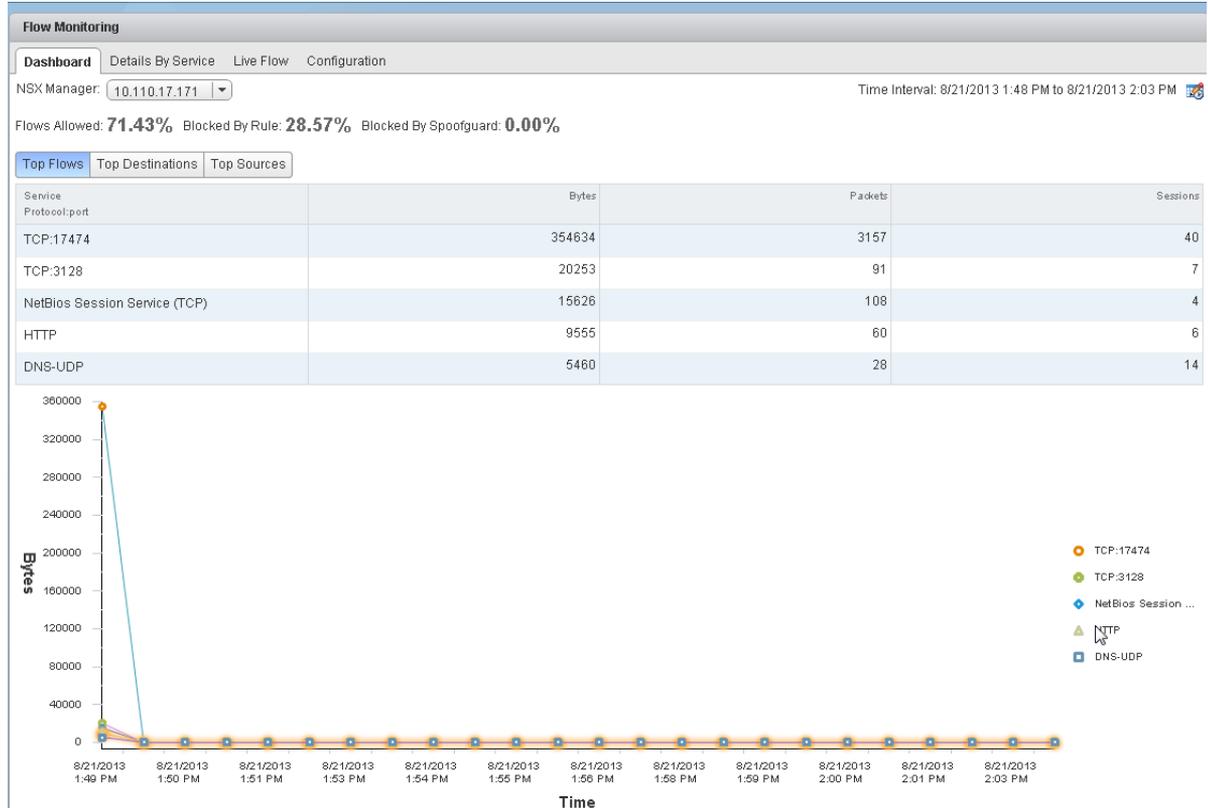
Flow monitoring data is only available for virtual machines in clusters that have the network virtualization components installed and firewall enabled. See the *NSX Data Center for vSphere Installation Guide*.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Tools > Flow Monitoring**.
- 2 Ensure that you are in the **Dashboard** tab.

### 3 Click **Flow Monitoring**.

The page might take several seconds to load. The top of the page displays the percentage of allowed traffic, traffic blocked by firewall rules, and traffic blocked by SpoofGuard. The multiple line graph displays data flow for each service in your environment. When you point to a service in the legend area, the plot for that service is highlighted.



Traffic statistics are displayed in three tabs:

- **Top Flows** displays the total incoming and outgoing traffic per service over the specified time period based on the total bytes value (not based on sessions/packets). The top five services are displayed. Blocked flows are not considered when calculating top flows.
- **Top Destinations** displays incoming traffic per destination over the specified time period. The top five destinations are displayed.
- **Top Sources** displays outgoing traffic per source over the specified time period. The top five sources are displayed.

### 4 Click the **Details by Service** tab.

Details about all traffic for the selected service is displayed. The **Allowed Flows** tab displays the allowed traffic sessions and the **Blocked Flows** tab displays the blocked traffic.

You can search on service names.

The screenshot shows the NSX Flow Monitoring interface. At the top, there are tabs for 'Dashboard', 'Details By Service', 'Live Flow', and 'Configuration'. Below the tabs, the NSX Manager is set to '10.110.17.171' and the Time Interval is '8/23/2013 6:10 AM to 8/23/2013 6:25 AM'. There are buttons for 'Allowed Flows' and 'Blocked Flows', and a search filter. The main table displays flow data:

Type	Service	Bytes	Sessions
UDP	DHCP-Server	4954	6
TCP	TCP:17474	2224	1
OTHER	IPv6-ICMP:0	1872	18
OTHER	ARP	1196	26
OTHER	0xffff	162	2
UDP	NTP Time Server	152	1

Below this table is a 'Find' search bar and a '6 items' indicator. A second table below shows rule details:

Rule Id	Time Stamp	Source	Source User(s)	Destination	Packets	Actions
1021	8/23/2013 6:15 AM	10.112.243.233	Unknown	10.112.192.5	2	Add Rule Edit Rule
1021	8/23/2013 6:15 AM	DB_server	Unknown	10.112.192.5	2	Add Rule Edit Rule
1021	8/23/2013 6:15 AM	win32rdpclone	Unknown	10.112.192.6	2	Add Rule Edit Rule
1021	8/23/2013 6:14 AM	10.112.243.214	Unknown	10.112.192.6	2	Add Rule Edit Rule
1021	8/23/2013 6:12 AM	win32rdpclone	Unknown	10.112.192.5	2	Add Rule Edit Rule
1021	8/23/2013 6:11 AM	10.112.243.229	Unknown	10.112.192.6	2	Add Rule Edit Rule
1021	8/23/2013 6:13 AM	win32rdpclone	Unknown	10.113.60.150	12	Add Rule Edit Rule

- 5 Click an item in the table to display the rules that allowed or blocked that traffic flow.
- 6 Click the **Rule Id** for a rule to display the rule details.

### Change the Date Range of the Flow Monitoring Charts

You can change the date range of the flow monitoring data for both the Dashboard and Details tabs.

#### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Tools > Flow Monitoring**.
- 2 Click  next to **Time interval**.
- 3 Select the time period or type a new start and end date.

The maximum time span for which you can view traffic flow data is the previous two weeks.

- 4 Click **OK**.

### Add or Edit a Firewall Rule from the Flow Monitoring Report

By drilling down into the traffic data, you can evaluate the use of your resources and send session information to Distributed Firewall to create a new allow or block rule at any level.

#### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Tools > Flow Monitoring**.

- 2 Click the **Details by Service** tab.
- 3 Click a service to view the traffic flow for it.  
Depending on the selected tab, rules that allowed or denied traffic for this service are displayed.
- 4 Click a rule ID to view rule details.
- 5 Do one of the following:
  - To edit a rule:
    - 1 Click **Edit Rule** in the **Actions** column.
    - 2 Change the name, action, or comments for the rule.
    - 3 Click **OK**.
  - To add a rule:
    - 1 Click **Add Rule** in the **Actions** column.
    - 2 Complete the form to add a rule. For information on completing the firewall rule form, see [Add a Firewall Rule](#).
    - 3 Click **OK**.

The rule is added at the top of the firewall rule section.

## View Live Flow

You can view UDP and TCP connections from and to a selected vNIC. In order to view traffic between two virtual machines, you can view live traffic for one virtual machine on one computer and the other virtual machine on a second computer. You can view traffic for a maximum of two vNICs per host and for 5 vNICs per infrastructure.

Viewing live flows can affect the performance of NSX Manager and the corresponding virtual machine.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Tools > Flow Monitoring**.
- 2 Click the **Live Flow** tab.
- 3 Select the required vNIC.
- 4 Click **Start** to begin viewing live flow.

The page refreshes every 5 seconds. You can select a different frequency from the **Refresh Rate** drop-down.

- 5 Click **Stop** when your debugging or troubleshooting is done to avoid affecting the performance of NSX Manager or the selected virtual machine.

## Configure Flow Monitoring Data Collection

After you have viewed and filtered the flow monitoring data that you want to collect, you can configure data collection.

You can filter the data being displayed by specifying exclusion criterion. For example, you may want to exclude a proxy server to avoid seeing duplicate flows. Or if you are running a Nessus scan on the virtual machines in your inventory, you may not want to exclude the scan flows from being collected. You can configure IPFix so that information for specific flows are exported directly from a firewall to a flow collector. The flow monitoring graphs do not include the IPFix flows. These are displayed on the IPFix collector's interface.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Tools > Flow Monitoring**.
- 2 Select the **Configuration** tab.
- 3 Ensure that **Global Flow Collection Status** is **Enabled**.

All firewall related flows are collected across your inventory except for the objects specified in **Exclusion Settings**.

4 To specify filtering criteria, click **Flow Exclusion** and follow the steps below.

- a Click the tab corresponding to the flows you want to exclude.

**Flow Monitoring**

Dashboard Details By Service Live Flow **Configuration**

NSX Manager: 10.110.8.93

Global Flow Collection Status: **Enabled** Disable

**Flow Exclusion** IPFix

**Exclusion Settings**  
System will not collect flows that match the specified condition

Filter	
Collect Blocked Flows	Yes
Collect Layer2 Flows	Yes
Source	
Destination	system-generated-broadcast-macset, 224.0.0.0/24,255.255.255.255
Destination ports	138,137
Service	

System is configured to collect all firewall related flows except those that match the conditions specified below

**Detail Collection Policy:** (Click Save to commit changes to settings)

Collect Blocked Flows:  Yes  No

Collect Layer2 Flows:  Yes  No

Save

- b Specify the required information.

If you selected	Specify the following information
<b>Collect Blocked Flows</b>	Select No to exclude blocked flows.
<b>Collect Layer2 Flows</b>	Select No to exclude Layer2 flows.
<b>Source</b>	Flows are not collected for the specified sources. <ol style="list-style-type: none"> <li>1 Click the <b>Add</b> icon.</li> <li>2 In View, select the appropriate container.</li> <li>3 Select the objects to exclude.</li> </ol>
<b>Destination</b>	Flows are not collected for the specified destinations. <ol style="list-style-type: none"> <li>1 Click the <b>Add</b> icon.</li> <li>2 In View, select the appropriate container.</li> <li>3 Select the objects to exclude.</li> </ol>
<b>Destination ports</b>	Excludes flows to the specified ports. Type the port numbers to exclude.
<b>Service</b>	Excludes flows for the specified services and service groups. <ol style="list-style-type: none"> <li>1 Click the <b>Add</b> icon.</li> <li>2 Select the appropriate services and/or service groups.</li> </ol>

- c Click **Save**.

5 To configure flow collection, click **IPFix** and follow the steps as described in [IPFIX for Distributed Firewall](#).

## 6 Click **Publish Changes**.

### Configure IPFIX

IPFIX (Internet Protocol Flow Information Export) is an IETF protocol that defines the standard of exporting flow information from an end device to a monitoring system. NSX supports IPFIX to export IP flow information to a collector.

You can enable IPFIX on:

- vSphere Distributed Switch (VDS)
- Distributed Firewall (DFW)

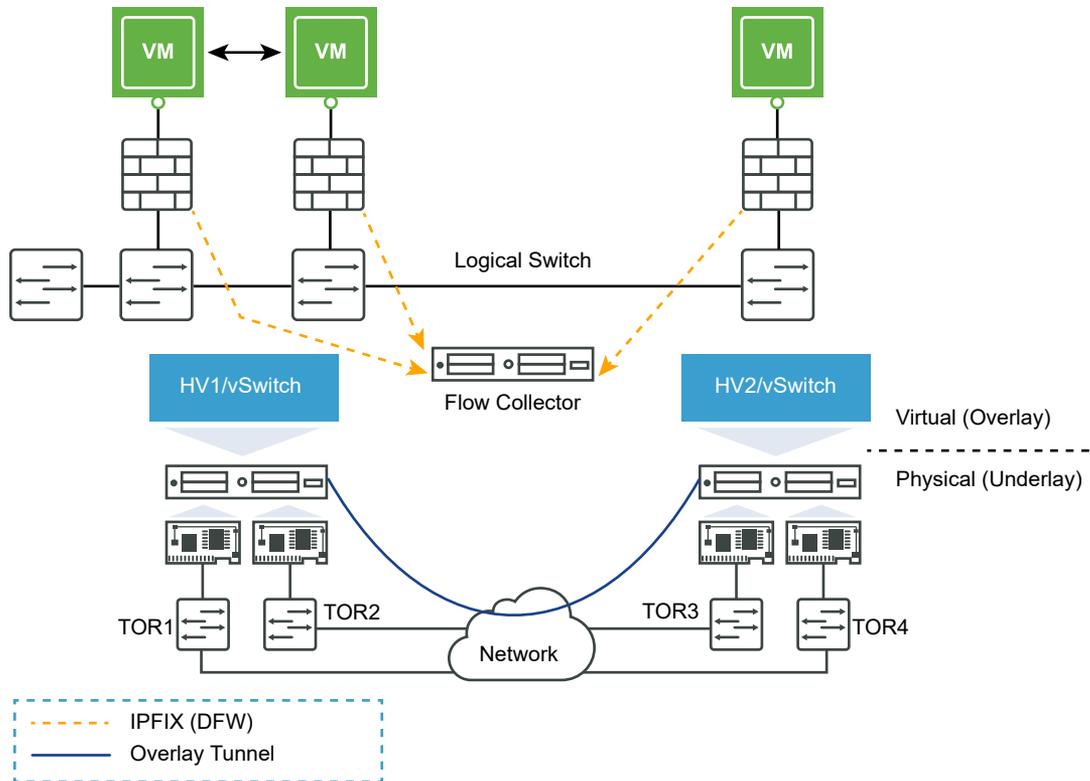
In the vSphere environment, vSphere Distributed Switch is the exporter and the collector is any monitoring tool available from any networking vendor.

The IPFIX standard specifies how IP flow information is presented and transferred from an exporter to a collector.

After you enable IPFIX on the vSphere Distributed Switch, it periodically sends messages to the collector tool. The contents of these messages are defined using the templates. For more information on templates, refer to [IPFIX Templates](#).

#### IPFIX for Distributed Firewall

You can enable IPFIX on a distributed firewall. Distributed firewall implements stateful tracking of flows and the tracked flows go through a set of state changes. IPFIX can be used to export data about the status of a flow. The tracked events include a flow creation, flow denial, flow update, and flow tear down.



You can enable flow export for IPFIX on a distributed firewall as follows:

- 1 In the vSphere Web Client, navigate to **Networking & Security > Tools > Flow Monitoring**.
- 2 Click the **Configuration** tab.
- 3 Ensure that **Global Flow Collection Status** is **Enabled**.
- 4 To configure flow collection, navigate to IPFIX:
  - In NSX 6.4.1 and later, navigate to **Networking & Security > Tools > IPFIX**.
  - In NSX 6.4.0, navigate to **Networking & Security > Tools > Flow Monitoring > Configuration > IPFix**.
- 5 Click **Edit** next to IPFIX Configuration, and then click **Enable IPFIX Configuration**.
- 6 In **Observation DomainID**, enter a 32-bit identifier that identifies the firewall exporter to the flow collector. Valid range is 0–65535.
- 7 In **Active Flow Export Timeout**, type the time (in minutes) after which active flows are to be exported to the flow collector. The default value is five. For example, if the flow is active for 30 minutes and the export timeout is five minutes, then the flow is exported seven times during its lifetime. One for each creation and deletion, and five times during the active period.
- 8 Click **Save**.
- 9 In **Collector IPs**, click **Add** and enter the IP address and UDP port of the flow collector. Refer to your NetFlow collector documentation to determine the port number.

10 Click **OK**.

11 Click **Publish Changes**.

## IPFIX Templates

Distributed firewall implements stateful tracking of flows and the tracked flows go through a set of state changes. You can use the IPFIX protocol to export data about the status of a flow. The tracked events include flow creation, flow denial, flow update, and flow teardown.

Because IPFIX is template-based, exporters must declare the format of the data before exporting any flow, so that the collector knows how to analyze incoming flow records. The format is declared in templates, which are sets of <type,length> that define the meaning and the length of each field in a record, one after the other.

The following table describes the information elements that are used in the IPFIX templates of the distributed firewall.

**Table 23-4. IPFIX Information Elements**

Name	Data Type	Size (Octet)	Description
sourceMacAddress	macAddress	6	The IEEE 802 source MAC address field.
destinationMacAddress	macAddress	6	The IEEE 802 destination MAC address field.
ethernetType	unsigned16	2	The Ethernet type field of an Ethernet frame that identifies the MAC client protocol carried in the payload.
sourceIPv4Address	ipv4Address	4	The IPv4 source address in the IP packet header.
destinationIPv4Address	ipv4Address	4	The IPv4 destination address in the IP packet header.
sourceIPv6Address	ipv6Address	16	The IPv6 source address in the IP packet header.
destinationIPv6Address	ipv6Address	16	The IPv6 destination address in the IP packet header.
sourceTransportPort	unsigned16	2	The source port identifier in the transport header.
destinationTransportPort	unsigned16	2	The destination port identifier in the transport header.
octetDeltaCount	unsigned64	8	The number of octets since the previous report (if any) in incoming packets for the flow at the observation point. The number of octets includes IP headers and IP payload.
packetDeltaCount	unsigned64	8	The number of incoming packets since the previous report (if any) for the flow at the observation point.

Table 23-4. IPFIX Information Elements (continued)

Name	Data Type	Size (Octet)	Description
flowId	unsigned64	8	A flow identifier that is unique within an observation domain. This information element helps to distinguish between different flows when flow keys, such as IP addresses and port numbers are not reported, or are reported in separate records.
flowStartSeconds	dateTimeSeconds	4	The absolute timestamp of the first packet of the flow.
flowEndSeconds	dateTimeSeconds	4	The absolute timestamp of the last packet of the flow.
protocolIdentifier	unsigned8	1	The value of the protocol number in the IP packet header.
firewallEvent	unsigned8	1	Valid values are: <ul style="list-style-type: none"> <li>■ 1 - Flow Created</li> <li>■ 2 - Flow Deleted</li> <li>■ 3 - Flow Denied</li> <li>■ 4 - Flow Alert (not used in this implementation)</li> <li>■ 5 - Flow Update</li> </ul>
direction	unsigned8	1	Valid values as applied to the filter at the observation point are: <ul style="list-style-type: none"> <li>■ 0x00 - ingres flow to VM</li> <li>■ 0x01 - egress flow from VM</li> </ul>
icmpTypeIPv4	unsigned8	1	Type of the IPv4 ICMP message.
icmpCodeIPv4	unsigned8	1	Code of the IPv4 ICMP message.
icmpTypeIPv6	unsigned8	1	Type of the IPv6 ICMP message.
icmpCodeIPv6	unsigned8	1	Code of the IPv6 ICMP message.
ruleId	unsigned32	4	firewall Rule Id - Enterprise specific IE.
vmUuid	string	16	VM UUID - Enterprise specific IE. Uniquely identifies the VM (octet array of 16).
vnicIndex	unsigned32	4	VNIC Index - Enterprise specific IE. Index of the VNIC for the specified VM.
sessionFlags	unsigned8	1	Session Flags - Enterprise specific IE. Valid values are: <ul style="list-style-type: none"> <li>■ 0 - unknown</li> <li>■ 0x1 - established</li> </ul>

Table 23-4. IPFIX Information Elements (continued)

Name	Data Type	Size (Octet)	Description
flowDirection	unsigned8	1	Flow Direction- Enterprise specific IE. Valid values are: <ul style="list-style-type: none"> <li>■ 0 - unknown</li> <li>■ 1 - forward</li> <li>■ 2 - reverse</li> </ul>
algControlFlowId	unsigned64	8	ALG Control Flow ID - Enterprise specific IE. Valid values are: <ul style="list-style-type: none"> <li>■ 0</li> <li>■ flowId of ALG control flow</li> </ul>
algType	unsigned8	1	ALG Control Flow ID - Enterprise specific IE. Valid values are: <ul style="list-style-type: none"> <li>■ 0 - none</li> <li>■ 1 - FTP</li> <li>■ 2 - Oracle</li> <li>■ 3 - SUNRPC</li> <li>■ 4 - DCERPC</li> <li>■ 5 - TFTP</li> </ul>
algFlowType	unsigned8	1	ALG Control Flow ID - Enterprise specific IE. Valid values are: <ul style="list-style-type: none"> <li>■ 0 - none</li> <li>■ 1 - control flow</li> <li>■ 2 - data flow</li> </ul>
averageLatency	unsigned32	4	Average TCP Latency - Enterprise specific IE Unit is in microseconds.
vifUuid	octetArray	16	VIF UUID - Enterprise specific IE. Uniquely identifies the VIF (octet array of 16).

The following IPFIX templates for a distributed firewall are supported only for UDP payloads.

### UDP IPV4 Template

The fields sent for this template are as follows:

```

IPFIX_TEMPLATE_FIELD(sourceMacAddress, 6)
IPFIX_TEMPLATE_FIELD(destinationMacAddress, 6)
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(ethernetType, 2)
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(vmUUIId,16)
IPFIX_TEMPLATE_FIELD(vnicIndex,4)
IPFIX_TEMPLATE_FIELD(sessionFlags,1) /* Introduced in 6.4.2 */
IPFIX_TEMPLATE_FIELD(flowDirection,1) /* Introduced in 6.4.2 */
IPFIX_TEMPLATE_FIELD(flowId,8) /* Introduced in 6.4.4 */
IPFIX_TEMPLATE_FIELD(algControlFlowId,8) /* Introduced in 6.4.4 */
IPFIX_TEMPLATE_FIELD(algType,1) /* Introduced in 6.4.4 */
IPFIX_TEMPLATE_FIELD(algFlowType,1) /* Introduced in 6.4.4 */
IPFIX_TEMPLATE_FIELD(averageLatency,4) /* Introduced in 6.4.4 */

```

## UDP IPV6 Template

The fields sent for this template are as follows:

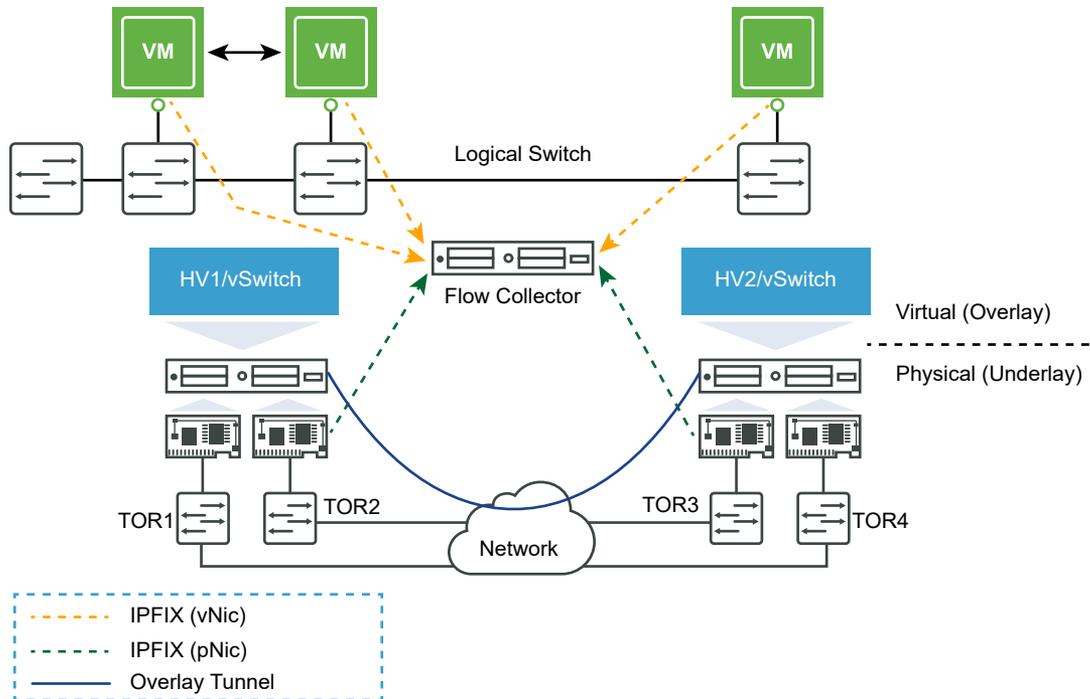
```

IPFIX_TEMPLATE_FIELD(sourceMacAddress,6)
IPFIX_TEMPLATE_FIELD(destinationMacAddress,6)
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(ethernetType,2)
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(vmUUIId,16)
IPFIX_TEMPLATE_FIELD(vnicIndex,4)
IPFIX_TEMPLATE_FIELD(sessionFlags,1) /* Introduced in 6.4.2 */
IPFIX_TEMPLATE_FIELD(flowDirection,1) /* Introduced in 6.4.2 */
IPFIX_TEMPLATE_FIELD(flowId,8) /* Introduced in 6.4.4 */
IPFIX_TEMPLATE_FIELD(algControlFlowId,8) /* Introduced in 6.4.4 */
IPFIX_TEMPLATE_FIELD(algType,1) /* Introduced in 6.4.4 */
IPFIX_TEMPLATE_FIELD(algFlowType,1) /* Introduced in 6.4.4 */
IPFIX_TEMPLATE_FIELD(averageLatency,4) /* Introduced in 6.4.4 */

```

## IPFIX for Logical Switch

You can enable IPFIX on vSphere Distributed Switch.



You can enable IPFIX for a logical switch as follows:

- 1 Configure the NetFlow collector on the vSphere Distributed Switch backing the NSX transport zone (Logical Switch). For more information on how to configure NetFlow collector, see "Configure the NetFlow Settings of a vSphere Distributed Switch" topic in the vSphere Networking Guide.
- 2 You can enable NetFlow monitoring on the distributed port group corresponding to the Logical Switch. If the NSX transport zone spans multiple vSphere Distributed Switches (VDS), then repeat these steps for each VDS/distributed port group. For more information on how to enable NetFlow monitoring, see "Enable or Disable NetFlow Monitoring on a Distributed Port Group or Distributed Port in the vSphere documentation.

In an NSX environment, the virtual machine data traffic on a logical switch traversing the NSX uplink of ESXi is VXLAN encapsulated. When NetFlow is enabled on the host uplink, the IP flow records are exported using a custom IPFIX flow-record template. The template includes the outer VXLAN UDP/IP header information and the information of the inner encapsulated IP packet. Such flow record, as a result provides visibility on the VTEP that is encapsulating the packet (outer header) and the details of the virtual machine that generated inter-host traffic (inner header) on a NSX logical switch (VXLAN).

For more details on the IPFIX templates for vSphere Distributed Switch, refer to [IPFIX Templates](#).

### IPFIX Templates

IPFIX templates provides the visibility into the VXLAN and non- VXLAN flows. The templates have additional parameters that provide more information regarding the encapsulated traffic.

The templates are supported in vSphere Distributed Switch (exporter). IPFIX support on vSphere Distributed Switch provides the required visibility into the virtual machine flows and VXLAN flows. If you are using any third party collector tool , you can use additional information available in the templates to provide correlation between the internal and external flows and the port connections.

The following table describes the information elements that are used in the IPFIX templates of the logical switch.

**Table 23-5. IPFIX Information Elements**

Name	Data Type	Size (Octet)	Description
sourceIPv4Address	ipv4Address	4	The IPv4 source address in the IP packet header.
destinationIPv4Address	ipv4Address	4	The IPv4 destination address in the IP packet header.
sourceIPv6Address	ipv6Address	16	The IPv6 source address in the IP packet header.
destinationIPv6Address	ipv^Address	16	The IPv6 destination address in the IP packet header.
octetDeltaCount	unsigned64	8	The number of octets since the previous report (if any) in incoming packets for the flow at the observation point. The number of octets includes IP headers and IP payload.
packetDeltaCount	unsigned64	8	The number of incoming packets since the previous report (if any) for the flow at the observation point.
flowStartSysUpTime	unsigned32	8	The relative timestamp of the first packet of the flow. It indicates the number of milliseconds since the last reinitialization of the IPFIX device (sysUpTime).
flowEndSysUpTime	unsigned32	8	The relative timestamp of the last packet of the flow. It indicates the number of milliseconds since the last reinitialization of the IPFIX device (sysUpTime).
sourceTransportPort	unsigned16	2	The source port identifier in the transport header
destinationTransportPort	unsigned16	2	The destination port identifier in the transport header
ingressInterface	unsigned32	4	The index of the IP interface where packets of the flow are received.
egressInterface	unsigned32	4	The index of the IP interface where packets of the flow are sent.

Table 23-5. IPFIX Information Elements (continued)

Name	Data Type	Size (Octet)	Description
vxlانId	unsigned64	8	<p>Identifier of a layer 2 network segment in an overlay network. The most significant byte identifies the layer 2 network overlay network encapsulation type:</p> <ul style="list-style-type: none"> <li>■ 0x00 reserved</li> <li>■ 0x01 VxLAN</li> <li>■ 0x02 NVGRE</li> </ul> <p>The lowest three significant bytes hold the value of the layer 2 overlay network segment identifier.</p> <p>For example:</p> <ul style="list-style-type: none"> <li>■ 24-bit segment ID VXLAN Network Identifier (VNI)</li> <li>■ 24-bit Tenant Network Identifier (TNI) for NVGRE</li> </ul>
protocolIdentifier	unsigned8	1	The value of the protocol number in the IP packet header.
flowEndReason	unsigned8	1	<p>The reason for terminating the flow. Valid values are:</p> <ul style="list-style-type: none"> <li>■ 0x01 - idle timeout</li> <li>■ 0x02 - active timeout</li> <li>■ 0x03 - end of Flow detected</li> <li>■ 0x04 - forced end</li> <li>■ 0x05 - lack of resources</li> </ul>
tcpFlags	unsigned8	1	<p>TCP control bits observed for the packets of the flow.</p> <p>This information is encoded as a bit field. For each TCP control bit, there is a bit in this set. The bit is set to 1 if any observed packet of the flow has the corresponding TCP control bit set to 1. The bit is cleared to 0 otherwise.</p>
IPv4TOS	unsigned8	1	The value of the TOS field in the IPv4 packet header.
IPv6TOS	unsigned8	1	The value of the Traffic Class field in the IPv6 packet header.
maxTTL	unsigned8	1	Maximum TTL value observed for any packet in the flow.
flowDir	unsigned8	1	<p>The direction of the flow observed at the observation point. Valid values are:</p> <ul style="list-style-type: none"> <li>■ 0x00 - ingress flow</li> <li>■ 0x01 - egress flow</li> </ul>

Table 23-5. IPFIX Information Elements (continued)

Name	Data Type	Size (Octet)	Description
ingressInterfaceAttr	unsigned16	2	The ingress interface attributes can take the following values based on type of the port: <ul style="list-style-type: none"> <li>■ IPFIX_UPLINK_PORT 0X01</li> <li>■ IPFIX_ACCESS_PORT 0X02</li> <li>■ IPFIX_VXLAN_TUNNEL_PORT 0X03</li> </ul>
egressInterfaceAttr	unsigned16	2	The egress interface attributes can take the following values based on type of the port: <ul style="list-style-type: none"> <li>■ IPFIX_UPLINK_PORT 0X01</li> <li>■ IPFIX_ACCESS_PORT 0X02</li> <li>■ IPFIX_VXLAN_TUNNEL_PORT 0X03</li> </ul>
vlanExportRole	unsigned8	1	Defines whether the exporter is an ESXi host or any other network device. IPFIX_END_POINT 0X01 means host is exporting the data. If other devices export the IPFIX templates, this field might have a different value (not defined yet).
paddingOctets	OctetArray	1	A sequence of 0x00 values.
tenantSourceIPv4	ipv4Address	4	The IPv4 source address in the IP header of the tenant packet, which is inside of the IP packet.
tenantDestIPv4	ipv4Address	4	The IPv4 destination address in the IP header of the tenant packet, which is inside of the IP packet.
tenantSourceIPv6	ipv6Address	16	The IPv6 source address in the IP header of the tenant packet, which is inside of the IP packet.
tenantDestIPv6	ipv6Address	16	The IPv6 destination address in the IP header of the tenant packet, which is inside of the IP packet.
tenantSourcePort	unsigned16	2	The source port identifier in the transport header of the tenant packet, which is inside of the IP packet.
tenantDestPort	unsigned16	2	The destination port identifier in the transport header of the tenant packet, which is inside of the IP packet.
tenantProtocol	unsigned8	1	The value of the protocol number in the IP header of the tenant packet, which is inside of the IP packet.

## IPv4 Template

```

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(vxlanId, 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(vxlanExportRole, 1)
IPFIX_TEMPLATE_PADDING(paddingOctets, 1)
IPFIX_TEMPLATE_END()

```

## IPv4 VXLAN Template

```

IPFIX_TEMPLATE_START(IPFIX_FLOW_TYPE_IPv4_VXLAN)
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(vxlanId, 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(vxlanExportRole, 1)
IPFIX_TEMPLATE_END()

```

## IPv4 ICMP VXLAN Template

```

IPFIX_TEMPLATE_START(IPFIX_FLOW_TYPE_IPv4_ICMP_VXLAN)
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(vxlanId, 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(vxlanExportRole, 1)
IPFIX_TEMPLATE_PADDING(paddingOctets, 1)
IPFIX_TEMPLATE_END()

```

## IPv4 ICMP Template

```

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(vxlanId, 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(vxlanExportRole, 1)
IPFIX_TEMPLATE_PADDING(paddingOctets, 2)
IPFIX_TEMPLATE_END()

```

## IPv6 ICMP VXLAN Template

```

IPFIX_TEMPLATE_START(IPFIX_FLOW_TYPE_IPv6_ICMP_VXLAN)
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)
//VXLAN Specific
IPFIX_TEMPLATE_FIELD(vxlanId, 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(vxlanExportRole, 1)
IPFIX_TEMPLATE_PADDING(paddingOctets, 1)
IPFIX_TEMPLATE_END()

```

## IPv6 ICMP Template

```

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(vxlanId, 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(vxlanExportRole, 1)
IPFIX_TEMPLATE_PADDING(paddingOctets, 2)
IPFIX_TEMPLATE_END()

```

## IPv6 Template

```

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(vxlanId, 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(vxlanExportRole, 1)
IPFIX_TEMPLATE_PADDING(paddingOctets, 1)
IPFIX_TEMPLATE_END()

```

## IPv6 VXLAN Template

```

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)

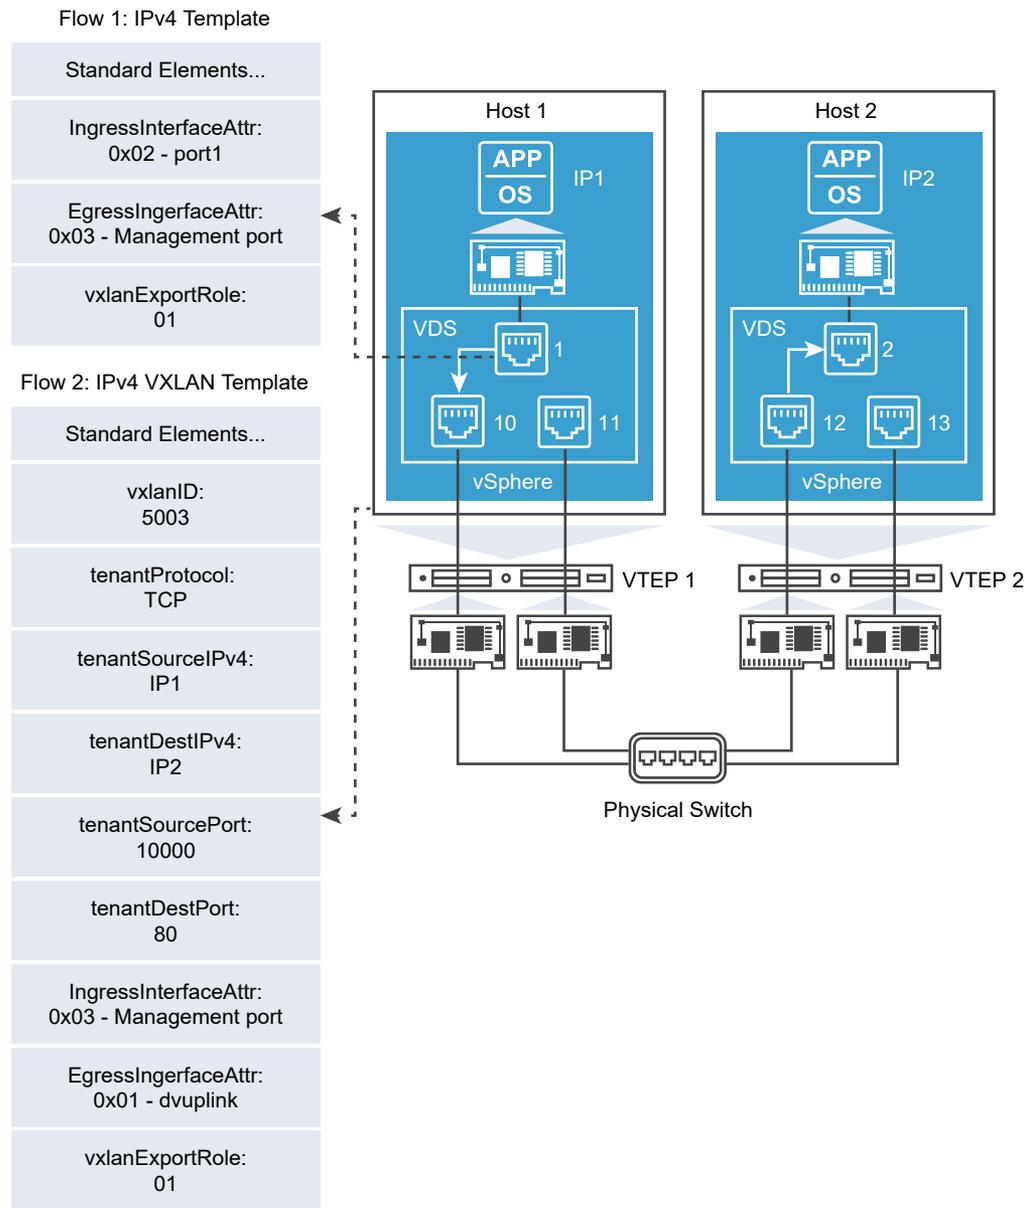
```

```
//VXLAN specific
IPFIX_TEMPLATE_FIELD(vxlanId, 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(vxlanExportRole, 1)
IPFIX_TEMPLATE_END()
```

### Flows Monitored by the IPFIX for vSphere Distributed Switch

The preceding diagrams show the communication between the two VMs running on two different hosts and the flows monitored by the IPFIX feature for vSphere Distributed Switch.

Figure 23-2. Flow on Host 1



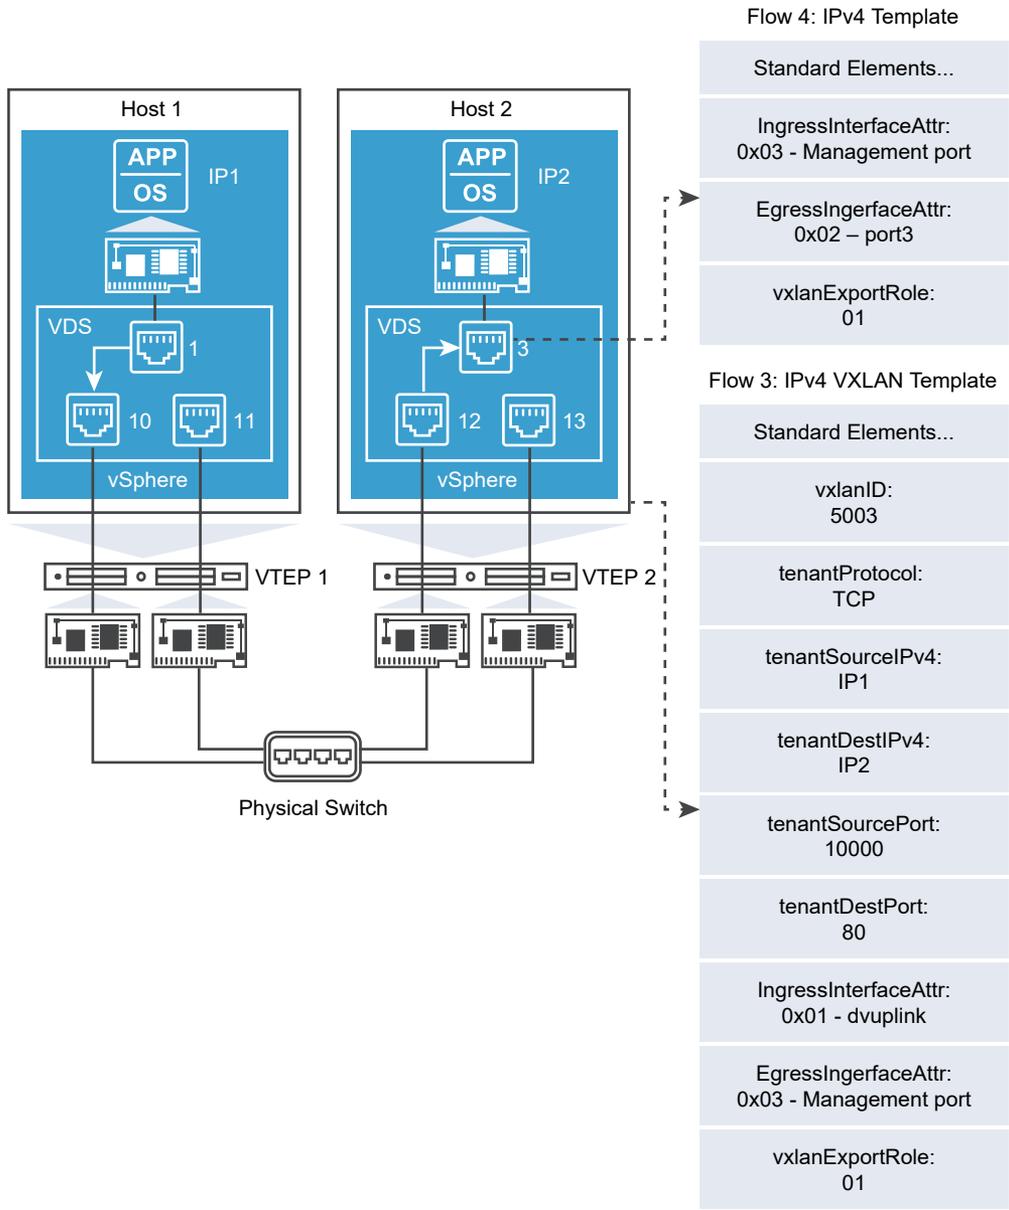
The [Figure 23-2. Flow on Host 1](#) shows the flows are collected from Host 1. The IPv4 template has additional information about the ingress and egress port and the standard elements.

The *ingressInterfaceAttr* text box *0x02* indicates it is an access port where the virtual machine is connected. The access port number is assigned to the *ingressInterface* parameter in the template.

The *egressInterfaceAttr* value of *0x03* shows that it is a VXLAN tunnel port and the port number associated with it is a management VMKNic port. This port number is assigned to the *egressInterface* parameter in the template.

The IPv4 VXLAN template on the other hand has additional information about the VXLAN ID, inner source, and destination IP/Port and protocol. The ingress and egress interfaces are *VXLAN tunnel port* and *dvuplink* port respectively.

Figure 23-3. Flow on Host 2



The Figure 23-2. Flow on Host 1 shows the flows on Host 2.

The templates in the Figure 23-2. Flow on Host 1 differs from the Figure 23-2. Flow on Host 1 only in the Ingress and egress attributes and port numbers.

The additional information provided through this template helps the collector tool vendors to correlate the external VXLAN flows and the internal virtual machine flows.

## Information Relevant to the Collector Tool Vendor

IPFIX support on vSphere Distributed Switch provides the required visibility into the virtual machine flows and VXLAN flows. If you are using any collector tool vendor, you can use additional information available in the templates to provide a correlation between the internal and external flows and the port connections.

The following section provides the details regarding how to decode the new parameters that are added in the VXLAN templates. IANA defines IPFIX information elements and their element IDs. You can find the list of standard element IDs at <http://www.iana.org/assignments/ipfix/ipfix.xml>.

All the new elements defined as part of VXLAN template have their new element IDs.

These custom parameters or elements provide additional information about the VXLAN and internal flows. The following are the new elements and their IDs:

**Table 23-6. Custom Parameters**

Element ID	Parameter Name	Data Type	Unit
880	tenantProtocol	unsigned8	1 byte
881	tenantSourceIPv4	ipv4Address	4 bytes
882	tenantDestIPv4	ipv4Address	4 bytes
883	tenantSourceIPv6	ipv6Address	16 bytes
884	tenantDestIPv6	ipv6Address	16 bytes
886	tenantSourcePort	unsigned16	2 bytes
887	tenantDestPort	unsigned16	2 bytes
888	egressInterfaceAttr	unsigned16	2 bytes
889	vxlanExportRole	unsigned8	1byte
890	ingressInterfaceAttr	unsigned16	2 bytes

**Note** The *Enterprise ID* is appended to all the custom elements defined above. The enterprise ID for VMware is *6876*.

The following table shows an example of complete list of element IDs. You can find data type and unit for standard element IDs at <http://www.iana.org/assignments/ipfix/ipfix.xml>.

Element ID	Parameter Name
1	octetDeltaCount
2	packetDeltaCount
4	protocolIdentifier
5	IPv4TOS

Element ID	Parameter Name
5	IPv6TOS
6	tcpFlags
7	sourceTransportPort
8	sourceIPv4Address
10	ingressInterface
11	destinationTransportPort
12	destinationIPv4Address
14	egressInterface
15	nextHopIPv4
27	sourceIPv6Address
28	destinationIPv6Address
53	maxTTL
61	flowDir
136	flowEndReason
152	flowStartSysUpTime
153	flowEndSysUpTime
210	paddingOctets
351	vxlanId
880	tenantProtocol
881	tenantSourceIPv4
882	tenantDestIPv4
883	tenantSourceIPv6
884	tenantDestIPv6
886	tenantSourcePort
887	tenantDestPort
888	egressInterfaceAttr
889	vxlanExportRole
890	ingressInterfaceAttr

## Application Rule Manager

The Application Rule Manager (ARM) tool simplifies the process of microsegmenting an application by creating security groups and firewall rules for existing applications.

Flow monitoring is used for long term data collection across the system, while the application rule manager is used for a targeted modeling of an application. During a flow monitoring phase, ARM learns about flows coming in and out of the application being profiled, as well as flows in between application tiers. It also learns about any Layer 7 Application Identity for the flows being discovered

There are three steps in the application rule manager workflow:

- 1 Select virtual machines (VM) that form the application and need to be monitored. Once configured, all incoming and outgoing flows for a defined set of vNICs (Virtualized Network Interface Cards) on the VMs are monitored. There can be up to five sessions collecting flows at a time.
- 2 Stop the monitoring to generate the flow tables. The flows are analyzed to reveal the interaction between VMs. The flows can be filtered to bring the flow records to a limited working set. After flow analysis, ARM automatically recommends:
  - Security Groups & IP set recommendation of the workloads based on the flow pattern and services used
  - Firewall policy based on the analyzed flow for a given ARM session
  - Layer 7 Application identity of the flow
- 3 Once a flow is analyzed with security group & policy recommendations, the policy for the given application can be published as a section in the firewall rule table. The recommended firewall rule also limits the scope of enforcement (applied to) to VM's associated with the application. Users can also modify the rules, especially naming of the groups and rule to make it more intuitive and readable.

### Create a Monitoring Session

A monitoring session collects all incoming and outgoing flows for up to 30 vNICs in a given session.

#### Prerequisites

Before starting a monitoring session, you need to define the VMs and vNICs that need to be monitored.

VMware Tools must be running and current on your Windows desktop VMs.

Selected VMs need to be in a cluster that has firewall enabled (they cannot be on the exclude list).

A default firewall rule of "any allow" that applies to the selected vNICs must be created for the duration of the monitoring session, so that flows to and from the vNICs are not dropped by any other firewall rule.

## Procedure

- 1 Log in to the vSphere Web Client, and navigate to Application Rule Manager.
  - In NSX 6.4.1 and later, navigate to **Networking & Security > Security > Application Rule Manager**.
  - In NSX 6.4.0, navigate to **Networking & Security > Tools > Flow Monitoring > Application Rule Manager**.
- 2 Click **Start New Session**.
- 3 In the **Start New Session** dialogue box, enter a name for the session.
- 4 Select either vNICs or VMs as the object type.
 

The **Available Objects** column is populated with the available objects.
- 5 Select the vNICs or VMs you want monitored. The selected vNICs or VMs move to the **Selected Objects** column.
- 6 Click **OK** to begin collecting flows.
 

The status is now **Collecting Data**. The latest set of flows collected is shown in the flow table.
- 7 Click **Stop** to end collecting flows.

## Results

A flow monitoring session has been created for the selected vNICs and VMs.

## What to do next

After collecting flows, analyze the flows.

## Analyze Flows and Auto Recommend

After a flow monitoring session has been collected, the results are analyzed and can be filtered for use in grouping objects and firewall rules. ARM automatically recommends firewall rules and security groups based on analyzed flows.

Analyzed flows can be filtered to limit the number of flows in a working set. The filter option icon is next to the Processed View drop-down menu on the right.

## Prerequisites

Before analysis, a flow monitoring session must have been collected from selected vNICs or VMs.

## Procedure

- 1 After flows have been collected, click **Analyze**.
 

Defined services are resolved, the IP address to VM translation begins, and duplicates are removed.

## 2 Once analyzed, the following data is provided for flows:

Field	Options
Direction	<p>IN - flow is coming into one of the VM and VNIC selected as part of the input seed.</p> <p>OUT - flow is generated from one of the VM and VNIC selected as part of the input seed.</p> <p>INTRA- flow is between VM- and VNIC selected as part of the input seed.</p>
Source	<p>VM Name, if the Source IP address of the flow record is resolved to one VM in the NSX inventory. Note that IP address can be resolved to VM, only if VM Tools has been enabled on those VMs.</p> <p>Raw IP, if there is no VM found for this source IP address in NSX Inventory. Note that multicast and broadcast IP addresses will not be resolved to VMs.</p> <p>Number of VMs (Ex:2 Virtual Machines) if the IP address is an overlapping IP address mapped to multiple VMs in different networks, the user needs to resolve Virtual machines to the correct Virtual Machine related to this flow record.</p>
Destination	Same values as Source field.
Service	<p>NSX defined service for protocol/port.</p> <p>Raw protocol/port, if there is no defined service in the NSX Manager.</p> <p>Number of services. If there is more than one service mapped to the same protocol/port and the user needs to resolve it to one service applicable to the flow record.</p>

## 3 Select the **Firewall Rules** tab to view the automatically recommended ARM grouped workflows and policy creation, and created firewall rules based on the selected flows. Users can modify the recommended rules, especially the naming of the groups and rules to make them more intuitive.

After flow analysis, ARM automatically recommends

- Grouping and IP set recommendations of the workflows based on the flow pattern and services. For example, with a 3-tier application, the outcome would be four recommended security groups - one for each of the application tiers and one groups for all the VMs in that application. ARM also recommends IP sets for destination based on services used by application VMs such as DNS/NTP servers if the destination IPs are outside of the vCenter domain.
- Security group recommendation based on analyzed flow data. A 3-tier application outcome could be four rules with LB to WEB on https, WEB to APP on http, APP to DB on MYSQL, and common rule for infra services such as DNS.
- Identify the Application Context (Layer 7) to the flow between application tiers. For example, L7 application running irrespective of TCP/UDP ports used and TLS version used for https.

- 4 Click **Publish** to publish the policy for the given application as a section in the firewall rule table. Or, modify the rules as needed. Note that the recommended firewall rule limits the scope of enforcement (applied to) to VM's associated with the application. Enter the firewall rule section name and click the checkbox to enable the following optional parameters:

Option	Description
Enable TCP Strict	Enables you to set TCP strict for each firewall section.
Enable Stateless Firewall	Enables stateless firewall for each firewall section.

## Flow Consolidation and Customization

After system analysis is complete, the analyzed flow table is available in the **Processed View**. Users can further consolidate the flows by changing the source, destination, and service fields. See [Customizing Services in Flow Records](#) and [Customizing Source and Destination in Flow Records](#).

### Processed View

Collected flows are displayed in a table with the following columns:

Field	Options
Direction	IN - flow is coming into one of the VMs or vNICs selected as part of the input seed. OUT - flow is generated from one of the VMs or vNICs selected as part of the input seed. INTRA- flow is between the VM or vNIC selected as part of the input seed.
Source	VM Name, if the Source IP address of the flow record is resolved to one VM in the NSX inventory. Raw IP if there is no VM found for this source IP address in the NSX Inventory. Note that multicast and broadcast IPs will not be resolved to VMs. Number of VMs if IP address is an overlapping IP address mapped to multiple VMs in different networks. The user needs to resolve multiple VMs to one VM related to this flow record.
Destination	Same values as Source field.
Service	NSX defined service for protocol/port. Raw protocol/port, if there is no defined service in the NSX Manager. Number of services. If there is more than one service mapped to the same protocol/port and the user needs to resolve it to one service applicable to the flow record.

Flow tables can be edited and the flows consolidated for easier rule creation. For example, the source field can be replaced with ANY. Multiple VMs receiving flows with HTTP and HTTPs can be replaced with "WEB-Service" service group, which includes both HTTP and HTTPs service. By doing so, Multiple flows may look similar and flow patterns may emerge that can be easily translated to a firewall rule.

Note that while each cell of the flow table can be modified, the cells are not auto-populated. For instance, if the IP address 196.1.1.1 is added to the DHCP-Server IPSet, the subsequent occurrences of that IP are not auto-populated to show the DHCP-Server group. There is a prompt asking if you want to replace all instances of the IP address with the IPSet. This allows the flexibility to make that IP part of multiple IPSet groups.

### Consolidated View

The consolidated view is accessed from the drop-down list in the right-hand corner. The consolidated view eliminates duplicate flows and displays the minimal number of flows. This view can be used to create firewall rules.

Clicking the arrow in the left hand corner of the Direction column shows the corresponding related raw flow information:

- for intra flows the corresponding IN and OUT flows with raw data are shown
- the original source IP, destination IP, port, and protocol information in all of the raw flows that were consolidated into the record
- for ALG flows, the corresponding data flow for the control flow is shown

## Customizing Services in Flow Records

Services flow cells can be customized on an individual cell basis by the user.

After flow analysis, users can associate any undefined protocol/port combinations and create a service. Service groups can be created for any of the services listed in the flows collected. For more information on modifying flow records see [Flow Consolidation and Customization](#).

### Prerequisites

Flow data must have been collected from a set of vNICs and VMs. See [Create a Monitoring Session](#).

## Procedure

- ◆ After the flow state is **Analysis Completed**, the flow table is populated with data, in the **Processed View**. To customize cell data, hover the cursor over a cell. A gear icon appears in the right-hand corner of the cell. Click the gear icon in the **Service** column and select one of the following options:

Option	Description
<b>Resolve Services</b>	If the port and protocol has been translated to multiple services, use this option to select the correct service.
<b>Create Services and Replace</b>	To add a service: <ol style="list-style-type: none"> <li>Enter a <b>name</b> for the service.</li> <li>From the drop-down list, select the protocol.</li> <li>Enter the destination ports for the service.</li> <li>Click <b>Advanced options</b> to enter the source ports of the service. The source port is used to track of new incoming connections and data streams.</li> <li>Optional - check <b>Enable inheritance to allow visibility at underlying scopes</b> to create a common group or criteria can be reused at the level of individual Edges.</li> <li>Click <b>OK</b> and a new service is created and populated in the Service column. Note that if there are other flow records with the same undefined port and protocol combination you will be asked for confirmation to replace all of them with the newly created Service. This occurs only for the flows with undefined services found in the Analysis phase.</li> </ol>
<b>Create Services Group and Replace</b>	You can create a new service group with the service from the flow included in it. Then, the new service group will replace the service. To add a service group: <ol style="list-style-type: none"> <li>Enter a <b>name</b> for the service group.</li> <li>Optional - enter a description of the Service Group.</li> <li>Select the <b>Object type</b>.</li> <li>Select the available objects you want to be added to the Service Group and click the arrow to move the object to the Selected Objects column.</li> <li>A new services group is created and populated in the Service column.</li> </ol>
<b>Replace Service with Any</b>	Replaces the specific service with any service.
<b>Replace Service with Service Group</b>	If the selected service is a member of multiple service groups, you select the specific service group you want to apply. <ol style="list-style-type: none"> <li>Click the desired Service Group from the list of available objects.</li> <li>Click <b>OK</b>.</li> </ol>
<b>Revert Protocol and Port</b>	Reverts any cell modifications back to the original data.

## Results

The changed flow record has a pink bar on the side. When the curser is hovered over any cell which has been modified there is a green checkmark. Clicking the checkmark displays a pop-up window with the previous and new values for that cell. The modified flow record is easier to translate into firewall rules.

## What to do next

Next, the flow record can be used to create firewall rules.

After flows have been modified, they can be further grouped together to get the smallest distinct working set. The **Processed View** is used to create Service Groups and IPSets and modify the flows. The **Consolidated view** further compresses these modified flows to make it easier to create firewall rules.

## Customizing Source and Destination in Flow Records

Source and destination flow cells can be customized on an individual cell basis by the user.

After flow analysis is complete, flow cells can be customized by the user.

### Prerequisites

Flow data must have been collected from a set of vNICs and VMs. See [Create a Monitoring Session](#)

### Procedure

- ◆ After the flow state shows **Analysis Completed**, the flow table is populated with data. To customize cell data, hover the cursor over a cell. A gear icon appears in the right-hand corner of the cell. Click the gear icon in the **Source** or **Destination** column and select one of the following options:

Option	Description
<b>Resolve VMs</b>	This option is available if multiple VMs have the same IP address. This option is used to chose the applicable VM name for the flow record.
<b>Replace with any</b>	If the source should be accessible to everyone then any source IP address is the correct option. In all other cases, you should specify the source address. Configuring a destination value of any for the destination IP address is discouraged.
<b>Replace with Membership</b>	If the VM is part of Security Groups they will be displayed here and can replace the VM name.

Option	Description
<b>Create Security Group</b>	<p>a Enter a Name and (optional) description of the security group.</p> <p>b Click <b>Next</b>.</p> <p>c Define the criteria that an object must meet for it to be added to the security group you are creating. This gives you the ability to include virtual machines by defining a filter criteria with a number of parameters supported to match the search criteria.</p> <p>d Select one or more resources to add to the security group. Note that when you add a resource to a security group, all associated resources are automatically added. For example, when you select a virtual machine, the associated vNIC is automatically added to the security group. You can include the following objects in a security group:</p> <p>Cluster</p> <p>Logical Switch</p> <p>Legacy Port Group</p> <p>vApp</p> <p>Datacenter</p> <p>e Click <b>Next</b>.</p> <p>f Select the objects to exclude from the security group. The objects selected here are always excluded from the security group, regardless of whether or not they match the dynamic criteria.</p> <p>g Click <b>Next</b>.</p> <p>h Review the Security Group details on the <b>Ready to complete</b> window. Click <b>Finish</b>.</p>
<b>Add to existing Security Group and Replace</b>	<p>For VMs, if the selected VM is a member of multiple security groups, select the specific security group you want to apply. This option is not available if the IP address is present in the source or destination field. For raw IP addresses, use Add to existing IPset and Replace option.</p> <p>a Click the desired Service Group from the list of available objects.</p> <p>b Click <b>OK</b>.</p>
<b>Create IPSet and Replace</b>	<p>An IPset allows you to apply a firewall rule to an entire set of IP addresses at once.</p> <p>a Enter a name for the IPSet.</p> <p>b Optional - enter a description.</p> <p>c Enter IP addresses or range of address in the new IP set.</p> <p>d Click <b>OK</b>.</p>
<b>Add to existing IPset and Replace</b>	<p>An IP address may be part of several IPsets. Use this option to replace the shown IP address and replace it with another.</p> <p>a Select the desire IPset from the Available Objects.</p> <p>b Click <b>OK</b>.</p>
<b>Revert to initial data</b>	Reverts any cell modifications back to the original data.

## What to do next

Create a firewall rule based on flow monitoring.

## Creating Firewall Rules from Application Rule Manager

Firewall rules can be edited, deleted, moved up and moved down as part of the Application Rule Manager.

### Prerequisites

After the flow record has been analyzed, ARM auto recommends firewall rules. You can modify the recommended rules, or create new firewall rules.

### Procedure

- 1 Open a flow session. If you are in the **Processed View**. Right-click on a single flow cell or shift + first cell > last cell to select several cells, and then right-click. If you are in the **Consolidated View** select a flow cell and click the **Action** icon. Select **Create Firewall rule**.

The **New Firewall Rule** pop-up window appears with all of the cells populated based on the selected row data. If several cells were selected, all the source, destination, service objects are added to the corresponding fields of the rule.

- 2 Enter a name for the new rule.
- 3 (Optional) To select a different source or destination click **Select** next to the Source or Destination box. Specify a new source or destination from the available objects and click **OK**.
- 4 (Optional) To select a different service click **Select** the Service box. Distributed Firewall supports ALG (Application Level Gateway) for the following protocols: FTP, CIFS, ORACLE TNS, MS-RPC, and SUN-RPC. Edge supports ALG for FTP only. Specify a new service from the available objects and click **OK**.
- 5 (Optional) To apply the rule to a different scope click **Select** next to the Applied To box. Make appropriate selections as described in the table below and click **OK**. By default, the rule is applied to the VNICS you originally right-clicked on.

To apply a rule to	Do this
All prepared clusters in your environment	Select <b>Apply this rule on all clusters on which Distributed Firewall is enabled</b> . After you click <b>OK</b> , the Applied To column for this rule displays <b>Distributed Firewall</b> .
One or more cluster, datacenter, distributed virtual port group, NSX Edge, network, virtual machine, vNIC, or logical switch	<ol style="list-style-type: none"> <li>1 In <b>Container type</b>, select the appropriate object..</li> <li>2 In the <b>Available</b> list, select one or more objects and click .</li> </ol>

If the rule contains virtual machines and vNICS in the source and destination fields, you must add both the source and destination virtual machines and vNICS to **Applied To** for the rule to work correctly.

6 Select the **Action** described in the table below.

Action	Results in
Allow	Allows traffic from or to the specified source(s), destination(s), and service(s).
Block	Blocks traffic from or to the specified source(s), destination(s), and service(s).
Reject	Sends reject message for unaccepted packets. RST packets are sent for TCP connections. ICMP messages with administratively prohibited code are sent for UDP, ICMP, and other IP connections.

7 Specify the rule **Direction** of the rule by clicking the drop-down arrow.

8 Click **OK**

#### What to do next

Publish the firewall rules. See [Publishing and Managing Firewall Rules From Application Rule Manager](#).

## Publishing and Managing Firewall Rules From Application Rule Manager

Firewall rules can be edited and published from the Application Rule Manager.

After firewall rules have been created they can be managed in the **Firewall Rules** tab of the Application Rule Manager.

#### Prerequisites

Analyze a flow session to create automatically recommended firewall rules, or create your own firewall rules from a flow monitoring session.

## Procedure

- ◆ Firewall rules appear in the **Firewall Rules** tab. Select one of the following options:

Option	Description						
<b>Publish</b>	<ol style="list-style-type: none"> <li>Click <b>Publish</b> to publish the created firewall rules. The rules are published as a new section.</li> <li>Enter <b>Section Name</b> for the firewall rule and click the checkbox to enable the following optional parameters: <table border="1" data-bbox="671 489 1430 709"> <thead> <tr> <th>Option</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>Enable TCP Strict</td> <td>Enables you to set TCP strict for each firewall section.</td> </tr> <tr> <td>Enable Stateless Firewall</td> <td>Enables stateless firewall for each firewall section.</td> </tr> </tbody> </table> </li> <li>Select where the new firewall section will be inserted in the existing firewall configuration.</li> <li>Click <b>OK</b>.</li> </ol>	Option	Description	Enable TCP Strict	Enables you to set TCP strict for each firewall section.	Enable Stateless Firewall	Enables stateless firewall for each firewall section.
Option	Description						
Enable TCP Strict	Enables you to set TCP strict for each firewall section.						
Enable Stateless Firewall	Enables stateless firewall for each firewall section.						
<b>Edit</b>	Select the pencil icon to edit the firewall rules.						
<b>Delete</b>	Select the red X icon to delete the firewall rule.						
<b>Down Arrow</b>	Select the down arrow icon to move the rule down						
<b>Up Arrow</b>	Select the up arrow icon to move the rule up.						

**Note** When firewall rules are published from **Application Rule Manager**, the section name is added to the **Publish** button. Any subsequent publishing from the **Application Rule Manager** overrides the existing section in the Firewall Configuration with the rules which are currently available in the **Application Rule Manager**.

## Host Health Status Monitoring

In NSX Data Center, you can diagnose the overall health status of the host. The overall host health status includes the host pNIC status, tunnel status, the connectivity status between the host and the control plane, and the connectivity status between the host and the management plane.

You can monitor the host health status only by using the NSX APIs. This diagnostic feature is not available in the vCenter UI.

The host health status includes the following substatuses:

- pNIC status
- Tunnel status
- Control plane status
- Management plane status

The following table describes each of these substatuses.

Substatus	Description
pNIC status	<p>This status is derived from the physical layer. When the pNICs belong to a link aggregation group (LAG), the status is either Up, Down, or Degraded.</p> <ul style="list-style-type: none"> <li>■ When all the pNICs in the LAG are up, the LAG status is Up.</li> <li>■ When all the pNICs in the LAG are down, the LAG status is Down.</li> <li>■ If any one of the pNICs in the LAG is down, the LAG status is Degraded.</li> </ul> <p>When the pNIC does not belong to a LAG, the status is either Up or Down.</p>
Tunnel status	<p>It is the connectivity status of the VTEP to VTEP tunnel between the hosts. The tunnel status is either Up, Down, or Degraded.</p> <ul style="list-style-type: none"> <li>■ When all the tunnels of the hosts are up, the tunnel status is Up.</li> <li>■ When all the tunnels of the hosts are down, the tunnel status is Down.</li> <li>■ When any one of the tunnels of the hosts is down, the tunnel status is Degraded.</li> </ul>
Control plane status	It is the connection status between the host and the NSX Controllers.
Management plane status	It is the connection status between the host and the NSX management plane.

The management plane determines the overall status of the host as follows:

- When all the substatuses are up, the overall host status is Up.
- When any substatus is down, the overall host status is Down.
- When at least one of the substatuses is degraded, and the other substatuses are up or down, the overall host status is Degraded.

## Activate Host Health Status Monitoring

To activate monitoring of the host health status, you must enable global pNIC status check on the hosts and global Bidirectional Forwarding Detection (BFD). Run the following PUT APIs:

### Enable global pNIC status check on hosts

```
PUT <NSX_Manager_IP>/api/2.0/vdn/pnic-check/configuration/global
```

### Enable global BFD

```
PUT <NSX_Manager_IP>/api/2.0/vdn/bfd/configuration/global
```

In NSX 6.4.6 or earlier, when you enable global BFD, the monitoring of tunnel latency and tunnel health is enabled simultaneously. You cannot separately turn on or turn off the monitoring of tunnel latency and tunnel health.

Starting in NSX 6.4.7, global BFD configuration API includes two additional parameters to enable or disable the monitoring of tunnel health and tunnel latency separately.

When BFD is disabled, tunnel latency and tunnel health monitoring cannot be turned on. When BFD is enabled, you can individually enable the monitoring of tunnel health and tunnel latency. This decoupling provides greater flexibility and avoids performance problems when the number of hosts scale in the network.

For a detailed information about configuring the global BFD parameters, see the *NSX API Guide*.

## View Host Health Status

To diagnose the host health status and the tunnel details, you can run the following APIs:

- GET <NSX\_Manager\_IP>/api/2.0/vdn/pnic-check/configuration/global
- GET <NSX\_Manager\_IP>/api/2.0/vdn/host/status
- GET <NSX\_Manager\_IP>/api/2.0/vdn/host/{hostId}/status
- GET <NSX\_Manager\_IP>/api/2.0/vdn/host/{hostId}/tunnel
- GET <NSX\_Manager\_IP>/api/2.0/vdn/host/{hostId}/remote-host-status

For a detailed information about each of these APIs, including the parameter description and the API response example, see the *NSX API Guide*.

## Network Latency Monitoring

Network administrators need the ability to monitor the latency of a virtualized network to diagnose and troubleshoot performance bottlenecks in the network.

For example, when NSX is installed and VXLAN-based networks are deployed in the network, the following types of latency exist:

- vNIC to pNIC (on the source hypervisor)
- pNIC to vNIC (on the destination hypervisor)
- vNIC to vNIC
- Tunnel latency (VTEP to VTEP)
- End-to-end latency of the data path

The network operations agent (netopa) on the ESXi host collects the network latency information from various sources, such as vSphere, NSX, and so on. Administrators can configure external collector tools, such as vRealize Network Insight (vRNI) to export the latency information to these collectors. Finally, they can run analytics on the latency information to troubleshoot network-specific problems.

---

**Note** The netopa agent can export the network latency information only to vRNI. Other collector tools are not supported currently.

---

You must use NSX REST APIs to configure NSX to calculate the latency metrics. For NSX to calculate the latency metrics correctly, ensure that the clocks on the different hosts are synchronized using the network time protocol (NTP).

## Tunnel Latency

To calculate the tunnel latency or VTEP-to-VTEP latency between ESXi hosts, NSX transmits Bidirectional Flow Detection (BFD) packets periodically in each tunnel. You must configure the BFD global configuration parameters by running the `PUT /api/2.0/vdn/bfd/configuration/global` API.

For more information about configuring the BFD global configuration parameters, see the *NSX API Guide*.

## End-to-End Latency

Starting in NSX 6.4.5, NSX can calculate the end-to-end latency of a data path as traffic moves between VMs that are either on the same ESXi host or on different ESXi hosts. However, both VMs must be attached to the same logical switch (subnet).

---

**Note** NSX cannot calculate the end-to-end latency information when data traffic is routed between VMs through a distributed logical router. That is, when VMs are attached to different logical switches or subnets.

---

To calculate the end-to-end latency of the data path, NSX uses the timestamp attribute of a data path packet inside the hypervisor. The end-to-end data path latency is calculated in terms of latency of the multiple segments in the data path: vNIC to pNIC and pNIC to vNIC.

For example, when traffic moves between VMs on the same host, vNIC to vNIC latency is calculated. When traffic moves between VMs on different ESXi hosts, vNIC to pNIC latency is calculated on the source hypervisor and pNIC to vNIC latency is calculated on the destination hypervisor. For traffic between the ESXi hosts, NSX calculates only the tunnel latency, if BFD global configuration parameters are configured.

For more information about configuring the latency parameters on a specific vSphere Distributed Switch and on a specific host, see the following sections in the *NSX API Guide*:

- Working with Latency Configuration of a Specific vSphere Distributed Switch
- Working with Latency Configuration of a Specific Host

## Endpoint Monitoring Data Collection

Endpoint monitoring enables users to map specific processes inside the guest OS to the network connections the processes are using.

After data is gathered, it is purged daily at 2:00 a.m. During the data purge the number of flow records across all sessions combined is checked, and any records above 20 million (or ~4GB) are deleted. Deletion begins with the oldest session, and continues until the number of flow records in the database is below 15 million records. If a session is in progress during the data purge, some records could be lost.

---

**Warning** When endpoint monitoring is enabled, the Dashboard shows a small yellow warning icon to indicate that the feature is turned on. Endpoint monitoring impacts performance and you must preferably turn it off after the data is collected.

---

### Prerequisites

- Endpoint Monitoring is supported on following Windows Operating systems:  
Windows Vista, Windows 7, Windows 8, Windows 8.1, Windows 2008, Windows 2008 R2, Windows 2012, Windows 10, and Windows 2016. It is not supported on Linux.
- Guest introspection must be installed on Virtual Machines (VMs).
- VMware Tools must be running and current on your Windows desktop VMs.
- Security Groups with 20 or fewer VMs are needed for data collection before Endpoint Monitoring can begin. See [Create a Security Group](#) for more information.
- Data collection must be enabled for one or more virtual machines on a vCenter Server before running an Endpoint Monitoring report. Before running a report, verify that the enabled virtual machines are active, and are generating network traffic.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Tools > Endpoint Monitoring**.
- 2 On the Summary tab, click **Start Collecting Data**.
- 3 On the Start Data Collection for Security Groups pop-up window, select the security groups for which you want to collect data. Click **OK**.

The VMs are listed in the field box.

- 4 Turn data collection **ON**.
- 5 Click **OK**.

The main Endpoint Monitoring Screen appears. In the left hand corner the status is Collecting Data.

- 6 Click **Stop Collecting Data** to end the data collection.

The EndPoint Monitoring screen appears with the Summary tab populated with data.

## Endpoint Monitoring

Endpoint Monitoring enables visibility into specific application processes and their associated network connections.

## Summary Tab

After data collection is completed, the summary screen displays the details of the NSX manager, the security group and the time slot of the collected data. The number of running virtual machines (VMs) and the total number of processes generating traffic is shown in the first box. Clicking the number of virtual machines running takes you to the VM Flows tab, described below. Clicking the number of processes generating traffic takes you to the Process Flows tab, described below.

The second box displays a donut with the total number of flows. A flow is any unique stream of network traffic as identified by its packet type, source and destination IP, and port. Hover the cursor over each section and the number of flows within the security group or outside the security group is shown.

## VM Flows Tab

This screen displays the details of the flows within the VMs including:

- VM name - Name of the VM that is being monitored
- Flows within security group - Traffic flowing between the VMs where the source or destination is inside the monitored security group
- Flows outside security group - Traffic flowing between the VMs where the source or destination is outside the monitored security group
- Shared service flows outside group - Shared service flows such as DHCP, LDAP, DNS, or NTP, outside the monitored security group
- Shared service flows inside security group - Shared service such as DHCP, LDAP, DNS, or NTP, inside the monitored security group

Clicking on a specific VM name in the table displays a bubble graph that shows the following:

- flows between VMs in the same security group
- flows that contain shared services
- flows between different security groups

Click on a bubble to view the details of the VM. The detailed flow view includes the process name, version and number of flows being generated by each process. If it contains shared services there is a special icon that is visible. Clicking on a line between two VM bubbles displays the process flow details of the flows between those two VMs including:

- Source process - Name of application/exe generating traffic and initiating the flow
- Source version - File version of source
- Protocol - TCP
- Destination process - Name of the server application/exe of the process that is the destination of the flow
- Destination port - Port number of the destination

## Process Flows Tab

This screen displays a list of all the applications that are generating flows. The table displays the following:

- Process Name - Name of application generating traffic
- VM name
- Flows within security group - Traffic flowing between the VMs where the source or destination is inside the monitored security group
- Flows outside security group - Traffic flowing between the VMs where the source or destination is outside the monitored security group
- Shared flows within security group - Shared flows, within the monitored security group
- Shared flows outside security group - Shared flows, outside the monitored security group

The bubble graph depicts the flows that are occurring with the process or application on the selected VM as the anchor. Click on any of the bubbles for the process name and version. Click on any line to display the following:

- Source VM - Name of client VM that is hosting the client process
- Source IP - IP address of the flow
- Protocol - TCP
- Destination VM - Name of the server VM that is hosting the server process
- Destination IP- IP address of the destination
- Destination port - Port number of the destination

## AD User Flows Tab

This screen displays the flows by all AD Users on AD joined VMs that are part of a security group. There are three tables:

- AD User Table -Lists all users that have initiated network flows from or to VMs that were part of the selected security group.
- AD Sessions Table - Lists all the sessions that were created by a user selected from the AD User Table. There are as many sessions as the number of unique pairs of users, source VM IPs.
- AD User Flows Table - When a user clicks on a session, this page appears, providing additional flow details

## Traceflow

Traceflow is a troubleshooting tool that provides the ability to inject a packet and observe where that packet is seen as it passes through the physical and logical network. The observations allow you to determine information about the network, such as identifying a node that is down or a firewall rule that is preventing a packet from being received by its destination.

## About Traceflow

Traceflow injects packets into a vSphere distributed switch (VDS) port and provides various observation points along the packet's path as it traverses physical and logical entities (such as ESXi hosts, logical switches, and logical routers) in the overlay and underlay networks. This 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.

Keep in mind that traceflow is not the same as a ping request/response that goes from guest-VM stack to guest-VM stack. What traceflow does is observe a marked packet as it traverses the overlay network. Each packet is monitored as it crosses the overlay network until it reaches and is deliverable to the destination guest VM. However, the injected traceflow packet is never actually delivered to the destination guest VM. This means that a traceflow can be successful even when the guest VM is powered down.

Traceflow supports the following traffic types:

- Layer 2 unicast
- Layer 3 unicast
- Layer 2 broadcast
- Layer 2 multicast

You can construct packets with custom header fields and packet sizes. The source for the traceflow is always a virtual machine virtual NIC (vNIC). 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 services gateway (ESG). The destination must be on the same subnet or must be reachable through NSX distributed logical routers.

The traceflow operation is considered Layer 2 if the source and destination vNICs are in the same Layer 2 domain. In NSX, this means that they are on the same VXLAN network identifier (VNI or segment ID). This happens, for example, when two VMs are attached to the same logical switch.

If NSX bridging is configured, unknown Layer 2 packets are always be sent to the bridge. Typically, the bridge forwards these packets to a VLAN and reports the traceflow packet as delivered. A packet reported as delivered does not necessarily mean that the trace packet was delivered to the specified destination.

For Layer 3 traceflow unicast traffic, the two end points are on different logical switches and have different VNIs, connected to a distributed logical router (DLR).

For multicast traffic, the source is a VM vNIC, and the destination is a multicast group address.

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.

---

**Caution** Depending on the number of logical ports in your deployment, multicast and broadcast traceflow operations might generate high traffic volume.

---

There are two ways to use traceflow: through the API and through the GUI. The API is the same API that the GUI uses, except the API allows you to specify the exact settings within the packet, while the GUI has more limited settings.

The GUI allows you to set the following values:

- Protocol---TCP, UDP, ICMP.
- Time-to-live (TTL). The default is 64 hops.
- TCP and UDP source and destination port numbers. The default values are 0.
- TCP flags.
- ICMP ID and sequence number. Both are 0 by default.
- An expiry timeout, in milliseconds (ms), for the traceflow operation. The default is 10,000 ms.
- Ethernet frame size. The default is 128 bytes per frame. The maximum frame size is 1000 bytes per frame.
- Payload encoding. The default is Base64.
- Payload value.

## Use Traceflow for Troubleshooting

There are multiple scenarios in which traceflow is useful.

Traceflow is useful in the following scenarios:

- Troubleshooting network failures to see the exact path that traffic takes
- Performance monitoring to see link utilization
- Network planning to see how a network will behave when it is in production

### Prerequisites

- Traceflow operations require communication among vCenter, NSX Manager, the NSX Controller cluster and the netcpa user world agents on the hosts.
- For Traceflow to work as expected, make sure that the controller cluster is connected and in healthy state.

## Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Tools > Traceflow**.
- 2 Select the traffic type: Unicast, Multicast, or L2 Broadcast.
- 3 Select the source VM vNIC.

If the VM is managed in the same vCenter Server where you are running the traceflow, you can select the VM and vNIC from a list.

---

**Note** When a logical switch is in multicast replication mode, the VMs connected to this logical switch will not be listed and cannot be chosen as the traceflow source or destination.

---

- 4 For a unicast traceflow, enter the destination vNIC information.

The destination can be a vNIC of any device in the NSX overlay or underlay, such as a host, a VM, a logical router, or an edge services gateway. If the destination is a VM that is running VMware Tools and is managed in the same vCenter Server from which you are running the traceflow, you can select the VM and vNIC from a list.

Otherwise, you must enter the destination IP address (and the MAC address for a unicast Layer 2 traceflow). You can gather this information from the device itself in the device console or in an SSH session. For example, if it is a Linux VM, you can get its IP and MAC address by running the `ifconfig` command in the Linux terminal. For a logical router or edge services gateway, you can gather the information from the `show interface` CLI command.

- 5 For a multicast traceflow, enter the multicast group address.

The packet is switched based on MAC address only.

Both the source and destination IP addresses are required to make the IP packet valid. In the case of multicast, the MAC address is deduced from the IP address.

- 6 For a Layer 2 broadcast traceflow, enter the subnet prefix length.

The packet is switched based on MAC address only. The destination MAC address is FF:FF:FF:FF:FF:FF.

Both the source and destination IP addresses are required to make the IP packet valid for firewall inspection.

- 7 Configure other required and optional settings.
- 8 Click **Trace**.

## Packet Capture

You can create a packet capture session for required hosts on the NSX Manager using the Packet Capture tool. After the packets are captured, the file is available to download. If your dashboard is indicating that a host is not in a healthy state, you can capture packets for that particular host for further troubleshooting.

For each session on the host, the packet capture file limit is 20 MB and the capture time limit is 10 minutes. A session remains active for 10 minutes or until the capture file reaches either 20 MB size or 20,000 packets, whichever limit is reached first. When any one of the limits is reached, the session is stopped. In the UI, you can create a maximum of 16 packet capture sessions. The NSX management plane limits the total captured file size across all sessions to 400 MB. When the combined packet file size of 400 MB is reached, a new capture session cannot be created. However, you can download the files of the previous packet capture sessions. If you want to start a new session after the limit of 400 MB file size is reached, you must clear the older sessions. An existing session is removed one hour after the session was created. If you restart NSX, all the existing sessions are cleared.

You cannot capture NSX VM interfaces.

You can perform the following tasks:

Options	Description
NSX Manager	Select the NSX Manager for which you want to create a packet capture session.
CREATE SESSION	To start a new packet capture session, click <b>CREATE SESSION</b> . For details, see <a href="#">Create a Packet Capture Session</a> .
STOP	Select the already started session, and click <b>STOP</b> . A confirmation dialog box appears. Click <b>YES</b> to stop the same session that is in-progress.
RESTART	Select the required session, and click <b>RESTART</b> . A confirmation dialog box appears. Click <b>YES</b> to restart the same session again. Restarting deletes the current session, clears the captured files, and starts a new session using the same configuration parameters.
CLEAR	Select the required session, and click <b>CLEAR</b> . A confirmation dialog box appears. Click <b>YES</b> to clear the session.
CLEAR ALL	If you want to clear all sessions, then click <b>CLEAR ALL</b> . A confirmation dialog box appears. Click <b>YES</b> to clear all the listed sessions.
DOWNLOAD	Select the required session, and click <b>DOWNLOAD</b> . You can also click the download icon. A confirmation dialog box appears. Click <b>YES</b> to download the captured session.

You can view the number of active sessions and total file size. The session status can be as follows:

- Started: Wait for the session to finish.
- Error: To view the error details, click the link.
- Finished: After the session is finished, you can download the session. You can also restart, and clear the session.

- Stopped: The session that was forced stopped. You can restart, or clear the session.

## Create a Packet Capture Session

If your dashboard is indicating that a host is not in a healthy state, you can capture packets for that particular host for further troubleshooting.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Tools > Packet Capture**.
- 2 To create new packet capture session, click **CREATE SESSION**.

The **Create Session** window appears. Enter the required details as explained further.

Parameter	Description
<b>General Tab</b>	The <b>General</b> tab is the default tab. Enter the required details as explained further.
Session Name	Type name for the session.
Host	Select the required host from the list.
Adapter	You can select either <b>Adapter</b> or <b>Filter</b> . If you select the adapter type, based on the selected adapter type, select the name of the adapter. If you select adapter type, you should not select <b>Filter Type</b> or <b>Filter Mode</b> . <i>dvPort</i> is specific to the <b>VdrPort</b> on the selected host. Other port is not supported.
Filter Type	Select the filter type from the list that is based on the selected host. The list depends on the configured firewall rules.
Filter Mode	Select filter mode as: <ul style="list-style-type: none"> <li>■ <b>Pre</b>: If you want to capture packets before applying the filter.</li> <li>■ <b>Post</b>: If you want to capture packets after applying the filter.</li> </ul>
Traffic Type	Select traffic type as <b>Incoming</b> or <b>Outgoing</b> . The packets are captured according to the direction of the flow with regard to the virtual switch.
<b>Advanced Tab</b>	To provide additional options to filter the capture packet, click the <b>Advanced</b> tab, and type the required details.

- 3 To create a session, click **SAVE**.

The packet capture session starts.

### Results

You can view the status of the session.

### What to do next

- You can download the captured session after it is finished, and can view the file using tools such as Wireshark.
- You can clear the session after you download the file.
- You can stop the session that is in-progress, if required.

## Support Bundle Collection Tool

You can collect the support bundle data for NSX components like NSX Manager, hosts, edges, and controllers. These support bundles are required to troubleshoot any problem related to NSX. You can either download this aggregate support bundle or can directly upload the bundle to a remote server. You can view the overall status of data collection and status for each component.

You can also use API to generate, download, delete, or cancel the bundle collection.

If the size limit in NSX is reached before all the requested logs are generated, the operation skips generation of remaining logs. The bundle gets generated with partial logs and is made available for the local download or FTP upload. The status of logs that are skipped is displayed.

When you request new log collection request, the old bundle gets deleted. The bundle also gets deleted after it has been uploaded to a remote server or the generate logs operation is canceled.

**Note** In an aggregated support bundle, the maximum number of nodes (including NSX Manager, NSX Edge, Host, and NSX Controller) should not be more than 200.

### Component and Log Types

Component	Log Type	Log
NSX Manager	N/A	NSX Manager logs
Host	N/A	<ul style="list-style-type: none"> <li>■ vmkernel logs</li> <li>■ vsfwd logs</li> <li>■ netcpa logs</li> <li>■ syslog</li> <li>■ General system information</li> </ul>
Host	Guest Introspection	<ul style="list-style-type: none"> <li>■ Guest Introspection Host Module (MUX)</li> <li>■ Guest Introspection SVM</li> </ul>
Host	Firewall	<ul style="list-style-type: none"> <li>■ vsfwd logs</li> <li>■ netcpa logs</li> </ul>
NSX Edge	N/A	Edge logs for selected edges
NSX Controller	N/A	Controller logs for selected controllers

Following logs are not collected:

- vSphere ESX Agent Manager (EAM) logs

- VM Guest logs

**Note** The support bundle does not contain all the host log files that you might need to troubleshoot host problems. To collect the complete host logs, use the vCenter Server. For information about collecting vSphere log files, see the *vSphere Monitoring and Performance* documentation.

## User Role and Permissions

User Role	Operation Permitted
NSX Administrator	All
Security Administrator	All
Enterprise Administrator	All
Auditor	View status and download bundle

## Create a Support Bundle

You can collect data for NSX components in form of a bundle. You can provide the bundle to VMware technical support for troubleshooting any issue with NSX. You can view the data collection status, and can download the bundle or collect it from the configured remote server.

### Prerequisites

You must be using NSX 6.4.0 or later to use this feature.

### Procedure

- 1 In the vSphere Web Client, navigate to **Networking & Security > Tools > Support Bundle**.
- 2 Select required NSX Manager from the list.
- 3 Select the components to be included in the support log as follows:
  - a To include NSX Manager logs, select the **Include NSX Manager logs** check box.
  - b Select the required object type from the list. You can select **Hosts**, **Edges**, and **Controllers**.
  - c Based on the selected object type, the available components are listed under the **Available Objects** column.
    - Select the check box next to the component that you want to include in the logs, and then click the arrow key (→) to move the component to the **Selected Objects** list.
    - To remove the selected objects, select the check box next to the component that you want to remove, and then click the arrow key (←) to move the component to the **Available Objects** list.
  - d For **Hosts**, you can select additional log options as **Guest Introspection** and **Firewall**.

- 4 Set the default timeout (in minutes) for the bundle collection process per selected object from the list.
- 5 Click **RESET** to clear all selections and start over again.

---

**Note** You can either download the support bundle, or upload it to a remote server. If you do not want to upload the bundle to a remote server, the bundle will be available for download after it is generated.

---

- 6 To upload the bundle to a remote server, perform the following steps:
  - a Select the **Upload support bundle to remote file server** check box.
  - b Select the transfer protocol as **SFTP** or **FTP** server.
  - c Enter the server details as follows:
    - **Hostname/IP**: Enter a hostname or IP address of the remote server.
    - **Username** and **Password**: Enter the user name and password of the remote server.
    - **Port**: Port number is added by default. You can change the port number, if necessary. To increase or decrease the number, click the arrow key.
    - **Backup Directory**: Enter the name of your backup directory.
- 7 Click **Start Bundle Collection**.

## Results

To view the bundle details, click the **View Bundle Details** link.

You can view overall completion status (in percent) and status for each component. You can view the completion status of bundle (in percent) being uploaded to a remote server. The status can be as follows:

- **Pending**: Wait for the process to start.
- **In Progress**: Wait for the process to complete.
- **Skipped**: This status can appear due to limited disk space. The bundle gets generated with partial logs and is made available for a local download, or is uploaded to a remote server. The status of the logs that are skipped is displayed.
- **Failed**: Log collection is failed due to various reasons like connectivity issues or timeout error. Click **START NEW** to start the data collection again.
- **Completed**: You can now download the bundle or view at the remote server.
- **Aborted**: The bundle generation process is canceled.
- **Partially Completed**: The logs are partially collected.

The bundle gets deleted after one of the following tasks:

- Bundle is uploaded to a remote server.

- You start a new log collection request.

#### What to do next

- You can [Download a Support Bundle](#) after it is generated, or you can view the filename that is uploaded to the configured remote server.
- Click **START NEW** to start the data collection again. A confirmation box appears. Click **Yes** to start new data generation request.

## Download a Support Bundle

You can download a support bundle after the data collection process is complete.

#### Prerequisites

Data collection process is 100% complete.

#### Procedure

- 1 Go to the **Support Bundle** page.
- 2 Click **DOWNLOAD**.
- 3 The support bundle filename ending with *tar.gz* is downloaded to your default Downloads folder. Some browsers might alter the file extension.

For example, the name of the support bundle file has a format similar to *VMware-NSX-TechSupport-Bundle-YYYY-MM-DD\_HH-MM-SS.tar.gz*.

- 4 Click the **DELETE SUPPORT BUNDLE** link.  
A confirmation box appears.
- 5 Click **Yes**. The log files generated in this particular request get deleted.

#### What to do next

You can provide the downloaded support bundle to VMware technical support.

## Cancel Support Bundle Generation Process

You can cancel the support bundle data collection process that is in-progress or pending status. You must delete the support bundle after canceling the process.

#### Prerequisites

Support bundle generation is in-progress.

#### Procedure

- 1 Go to the **Support Bundle** page.
- 2 If the log generation is in-progress, you can cancel the process. Click **ABORT GENERATION**.  
A confirmation box appears.

- 3 Click **OK** to cancel the process.

The **Support Bundle** page displays data collection as **Aborted**.

- 4 Click the **DELETE SUPPORT BUNDLE** link.

A confirmation box appears.

- 5 Click **Yes**. The log files generated in this particular request get deleted.

#### Results

The support bundle data collection process is canceled.

#### What to do next

Click **START NEW** to start data collection again. A confirmation box appears. Click **Yes** to start a new data generation request.

# Disaster Recovery Scenarios with Cross-vCenter NSX

# 24

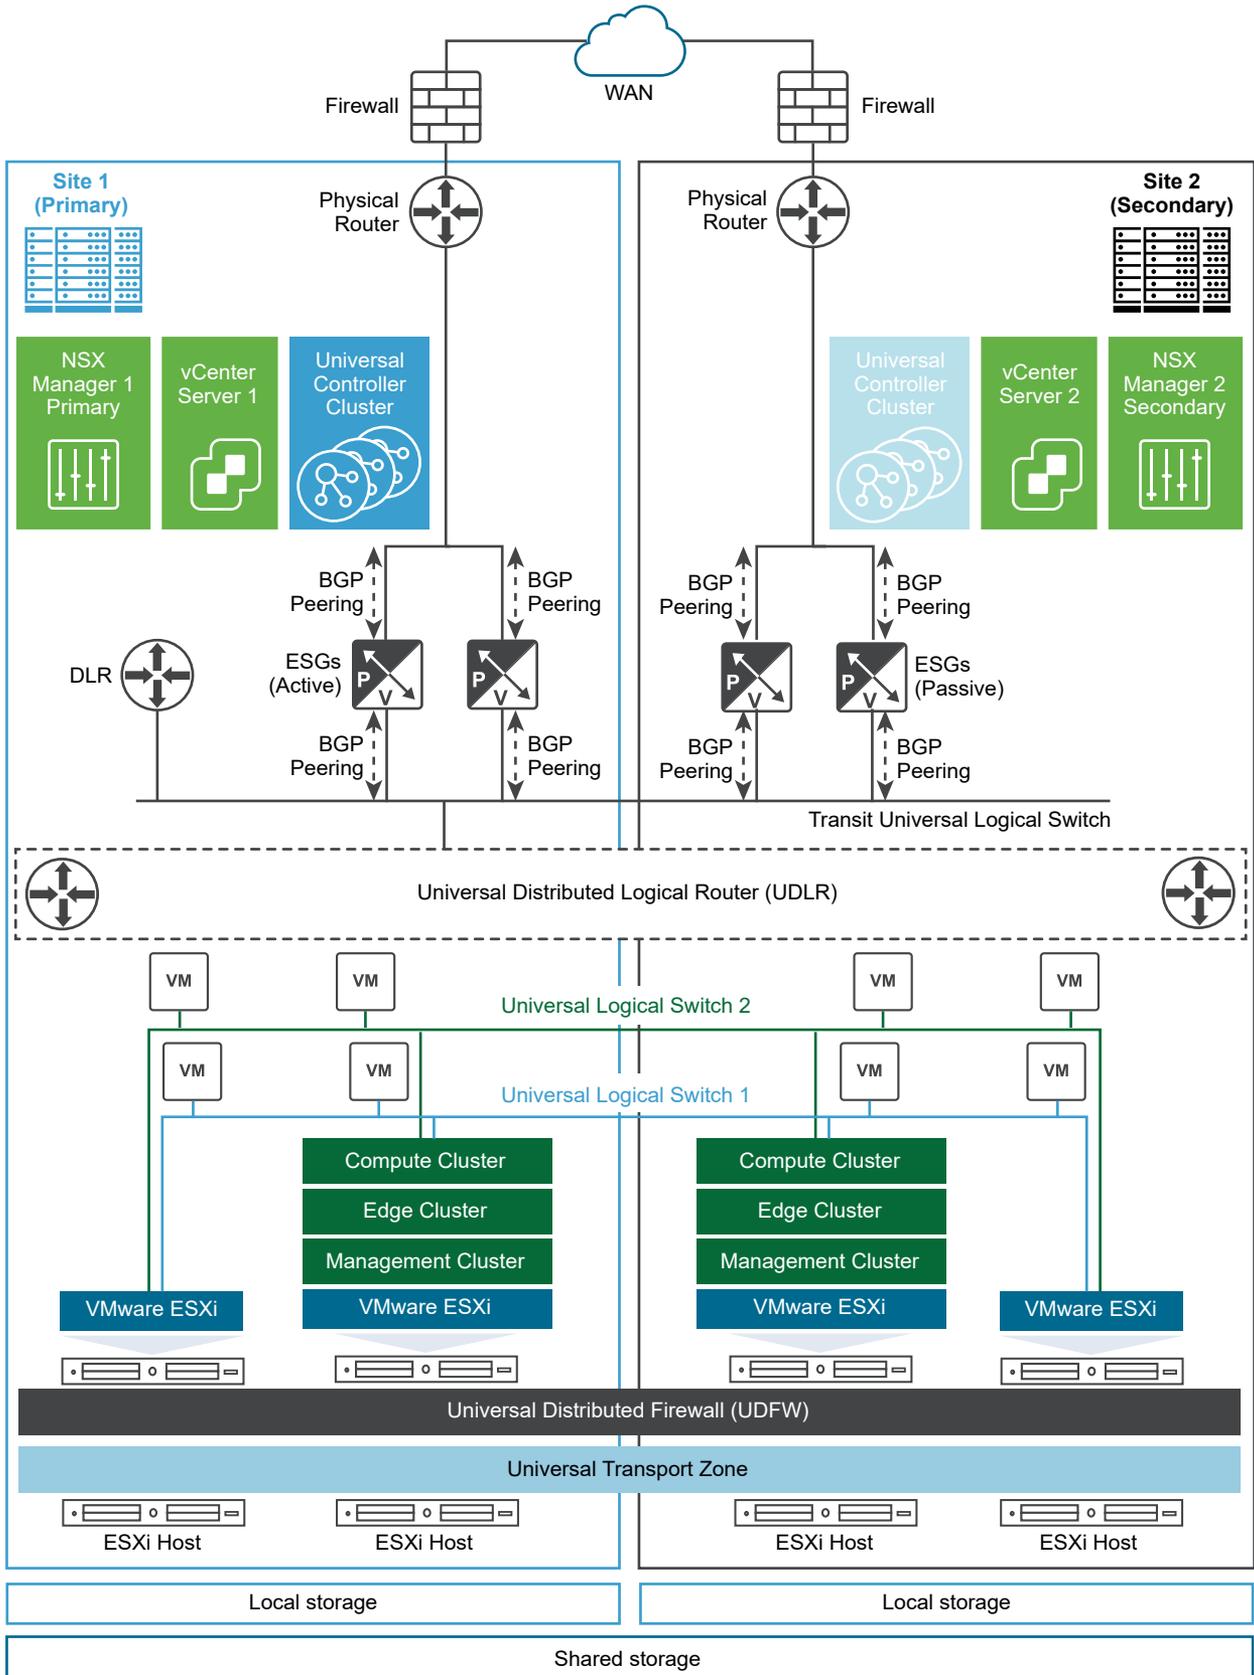
Company ACME Enterprise has two private data center sites in the US, one at Palo Alto, and the other at Austin. During a scheduled maintenance or an unforeseen failure at the Palo Alto site, the company recovers all the applications at its Austin site.

Currently, ACME Enterprise achieves this disaster recovery the traditional way by performing the following tasks manually:

- Remapping IP address
- Synchronizing security policies
- Updating other services that use the application IP addresses, such as DNS, security policies, and other services.

This traditional approach to a disaster recovery consumes significant additional time to complete 100% recovery at its site in Austin. To achieve a fast disaster recovery with a minimal downtime, ACME Enterprise decides to deploy NSX Data Center 6.4.5 or later in a Cross-vCenter environment, as shown in the following logical topology diagram.

Figure 24-1. Multi-Site Cross-vCenter NSX Topology in Active - Passive Mode and Local Egress Disabled



In this topology, site 1 at Palo Alto is the primary (protected) data center, and site 2 at Austin is the secondary (recovery) data center. Each site has a single vCenter Server, which is paired with its own NSX Manager. The NSX Manager at site 1 (Palo Alto) is assigned the role of a primary NSX Manager, and the NSX Manager at site 2 (Austin) is assigned the role of a secondary NSX Manager.

ACME Enterprise deploys the Cross-vCenter NSX across both sites in an Active - Passive mode. 100% applications (workloads) run on site 1 at Palo Alto, and 0% applications run on site 2 at Austin. That is, by default, site 2 is in passive or standby mode.

Both sites have their own Compute, Edge, and Management Clusters and ESGs that are local to that site. As local egress is disabled on the UDLR, only a single UDLR Control VM is deployed on the primary site. The UDLR Control VM is connected to the universal transit logical switch.

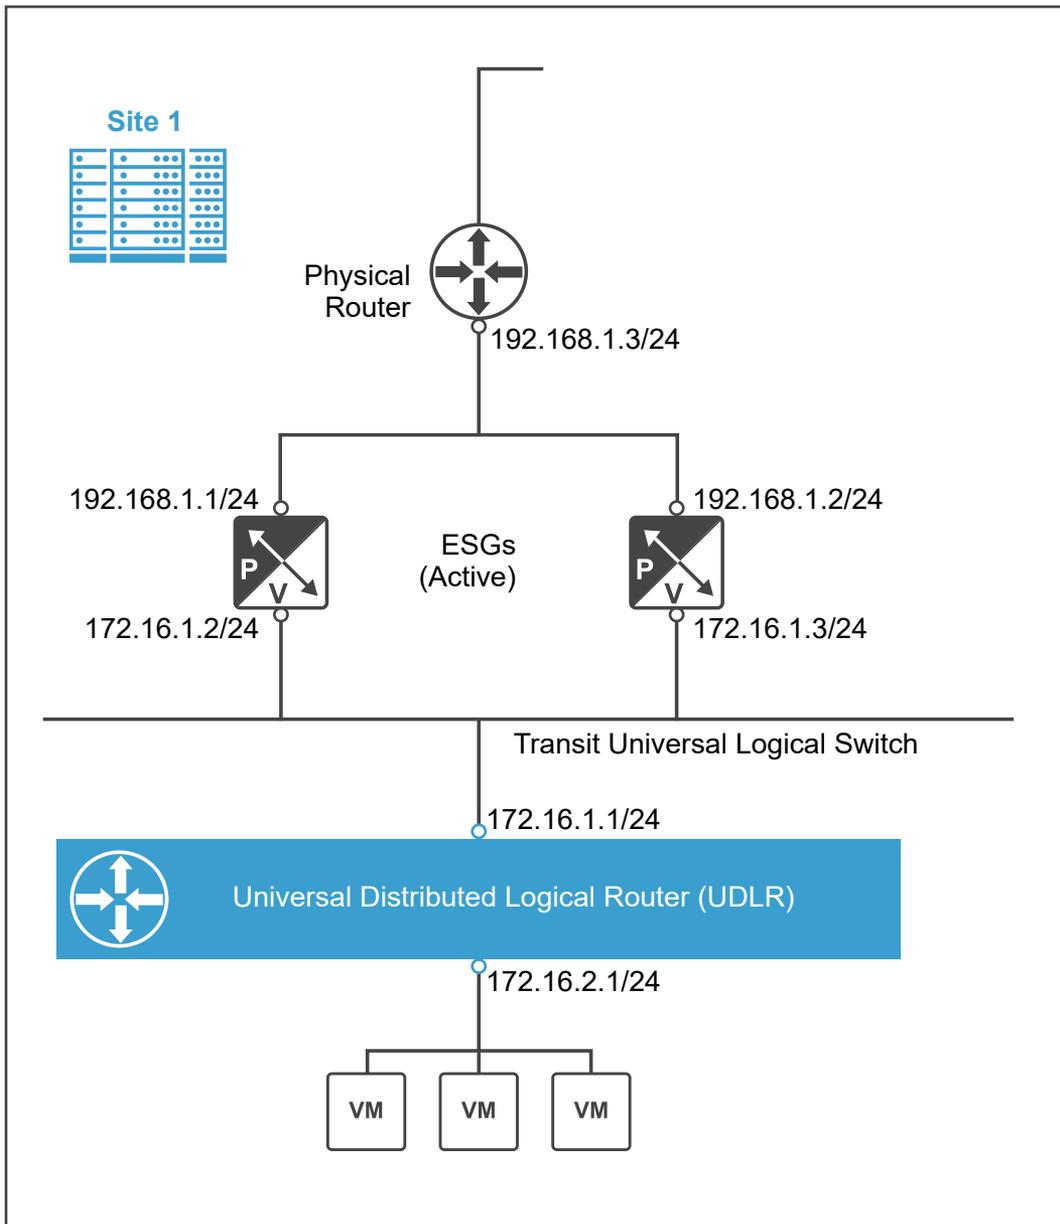
The NSX administrator creates universal objects that span two vCenter domains at site 1 and site 2. The universal logical networks use universal networking and security objects, such as Universal Logical Switches (ULS), Universal Distributed Logical Routers (UDLR), and Universal Distributed Firewall (UDFW).

The administrator does the following configuration tasks at site 1:

- Creates a universal transport zone from the primary NSX Manager.
- Deploys a Universal Controller Cluster with three controller nodes.
- Adds the local Compute, Edge, and Management clusters to the universal transport zone from the primary NSX Manager.
- Disables local egress, enables ECMP, and enables Graceful Restart on the UDLR Control VM (Edge Appliance VM).
- Configures dynamic routing using BGP between the Edge Services Gateways (ESGs) and the UDLR Control VM.
- Disables ECMP and enables Graceful Restart on both the ESGs.
- Disables firewall on both the ESGs because ECMP is enabled on the UDLR Control VM and to ensure that all traffic is allowed.

The following diagram shows a sample configuration of the uplink and downlink interfaces on the ESGs and the UDLRs at site 1.

Figure 24-2. Site 1: Sample Interface Configuration



The administrator does the following configuration tasks at site 2:

- Adds the local Compute, Edge, and Management clusters to the universal transport zone from the secondary NSX Manager.
- Specifies similar downlink interfaces on the ESGs as configured at site 1 ESGs.
- Specifies similar BGP configuration on the ESGs as configured at site 1 ESGs.
- Powers down the ESGs on the secondary site when site 1 is active.

Now, let us walk through the steps that the NSX administrator can perform to achieve a disaster recovery in the following scenarios:

- Scenario 1: Scheduled full site failure at site 1
- Scenario 2: Unscheduled full site failure at site 1
- Scenario 3: Full failback to site 1

This chapter includes the following topics:

- [Scenario 1: Scheduled Full Site Failure](#)
- [Scenario 2: Unscheduled Full Site Failure](#)
- [Scenario 3: Full Failback to Primary Site](#)

## Scenario 1: Scheduled Full Site Failure

In this scenario, the NSX administrator performs a scheduled maintenance of the network infrastructure at site 1. During the scheduled maintenance window, the administrator shuts down site 1, and performs a smooth failover to the secondary site 2.

The NSX administrator wants to meet the following key objectives:

- Achieve a full site failover at site 2 with minimal downtime.
- Retain site 1 application IP addresses at site 2 after the failover.
- Automatically recover all Edge interface settings and BGP protocol configuration settings at site 2.

---

### Note

- The administrator can do the failover tasks manually by using either the vSphere Web Client or by running the NSX REST APIs. In addition, the administrator can automate some failover tasks by running a script file that contains the APIs to run during the failover. This scenario explains manual failover steps using the vSphere Web Client. However, if any step requires the use of either the CLI or the NSX REST APIs, adequate instructions are provided.
- In this scenario, the disaster recovery workflow is specific to the topology explained earlier, which has a primary NSX Manager and a single secondary NSX Manager. The workflow with multiple secondary NSX Managers is not in the scope of this scenario.

---

**Important** When the failover to the secondary site 2 is in progress or partly completed, avoid powering on the NSX Manager at site 1 to failback to the primary site 1. Ensure that the failover process is first completed by using the procedure in this scenario. Only after a clean failover is done to the secondary site 2, restore or failback all the workloads to the original primary site 1. For detailed instructions about the failback process, see [Scenario 3: Full Failback to Primary Site](#).

---

### Prerequisites

- NSX Data Center 6.4.5 or later is installed at both sites 1 and 2.

- vCenter Server at sites 1 and 2 are deployed with Enhanced Linked Mode.
- At site 1 and site 2, the following conditions are met:
  - No application-specific security policies are configured on a non-NSX firewall, if any.
  - No application-specific firewall rules are configured on a non-NSX firewall, if any.
  - Firewall is disabled on both the ESGs because ECMP is enabled on the UDLRs and to ensure that all traffic is allowed.
- At site 2, the following conditions are met before the failover:
  - Similar downlink interfaces are configured manually on the ESGs as configured at site 1.
  - Similar BGP configuration is done manually on the ESGs as configured at site 1.
  - ESGs are in powered down state when the primary site 1 is active or running.

### Procedure

- 1 Shut down site 1 on the scheduled date.
  - a Power off the primary NSX Manager and all the three controller nodes that are associated with the primary NSX Manager.
  - b On the **Installation and Upgrade** page, navigate to **Management > NSX Managers**.
    - If you refresh the **NSX Managers** page in the current browser session, the role of the primary NSX Manager changes to `Unknown`.
    - If you log out from the vSphere Web Client and log in again or start a new vSphere Web Client browser session, the primary NSX Manager is no longer displayed on the **NSX Managers** page.
  - c Navigate to **Networking & Security > Dashboard > Overview**.
    - If you refresh the **Dashboard** page in the current browser session, the following error message is displayed: `Could not establish connection with NSX Manager. Please contact administrator..` This error means that the primary NSX Manager is no longer reachable.
    - If you log out from the vSphere Web Client and log in again or start a new vSphere Web Client browser session, the primary NSX Manager is no longer available in the **NSX Manager** drop-down menu.
  - d Navigate to **Networking & Security > Installation and Upgrade > Management > NSX Controller Nodes**. Select the secondary NSX Manager, and ensure that the status of all three controller nodes is `Disconnected`.
  - e Power off all the NSX Edges and the universal distributed logical router (UDLR) control VM.

## 2 Promote the secondary NSX Manager to a primary role.

- a On the **Installation and Upgrade** page, navigate to **Management > NSX Managers**.
- b Select the secondary NSX Manager.
- c Click **Actions > Disconnect from Primary NSX Manager**. When prompted to continue with the disconnect operation, click **Yes**.

The secondary NSX Manager is disconnected from the primary NSX Manager, and enters into a `Transit` role.

- d Click **Actions > Assign Primary Role**.

The secondary NSX Manager at site 2 is promoted to a primary role.

---

**Caution** As local egress is disabled on the UDLR, the UDLR Control VM (Edge Appliance VM) is deployed only at the original primary site (site 1). Before site 1 fails, the UDLR Control VM is not available at the secondary site (site 2), which is now promoted to primary. Therefore, redeploy the UDLR Control VM at the promoted primary site (site 2) before redeploying the NSX Controller Cluster.

If the controller nodes are deployed before deploying the UDLR Control VM, the forwarding tables on the UDLR are flushed out. This results in a downtime immediately after the first controller node is deployed at site 2. This situation might result in communication outages. To avoid this situation, deploy the UDLR Control VM before deploying the NSX Controller nodes.

## 3 Power on the NSX Edges that are in the powered down state, and deploy the UDLR Control VM (Edge Appliance VM) at the secondary site 2 (promoted primary).

For instructions about deploying the UDLR Control VM, see the *NSX Cross-vCenter Installation Guide*.

While deploying the UDLR Control VM, configure the following resource settings:

- Select the data center as **site 2**.
- Select the cluster/resource pool.
- Select the datastore.

---

**Note** After deploying the UDLR Control VM, the following configuration settings are automatically recovered at site 2:

- BGP protocol routing configuration
- BGP password configuration
- Uplink and internal interface settings

## 4 Deploy the three NSX Controller Cluster nodes at site 2 (promoted primary).

For detailed instructions about deploying NSX Controllers, see the *NSX Cross-vCenter Installation Guide*.

- 5 Update the NSX Controller Cluster state.
  - a On the **Installation and Upgrade** page, click **NSX Managers**.
  - b Select the promoted primary NSX Manager.
  - c Click **Actions > Update Controller State**.
- 6 Force sync routing service on each cluster at site 2.
  - a On the **Installation and Upgrade** page, click **Host Preparation**.
  - b Select the promoted primary NSX Manager.
  - c Select one cluster at a time, and then click **Actions > Force Sync Services**.
  - d Select **Routing**, and click **OK**.
- 7 Migrate the workload VMs from site 1 to site 2.

---

**Note** The workload VMs continue to exist at site 1. Therefore, you must manually migrate the workload VMs to site 2.

---

## Results

The manual recovery of NSX components and the failover from the primary site (site 1) to the secondary site (site 2) is complete.

## What to do next

Verify whether the failover to site 2 is 100% complete by doing these steps on site 2 (promoted primary site):

- 1 Check whether the NSX Manager has the primary role.
- 2 Check whether the Control VM (Edge Appliance VM) is deployed on the UDLR.
- 3 Check whether the status of all controller cluster nodes is *Connected*.
- 4 Check whether the host preparation status is *Green*.
- 5 Log in to the CLI console of the UDLR Control VM (Edge Appliance VM), and do these steps:
  - a Check whether all BGP neighbors are established and the status is UP by running the `show ip bgp neighbors` command.
  - b Check whether all BGP routes are being learned from all BGP neighbors by running the `show ip route bgp` command.

After a complete failover to site 2, all workloads run on the secondary site (promoted primary) and traffic is routed through the UDLR and the NSX Edges at site 2.

After the scheduled maintenance is done, the administrator powers on the NSX Manager and controller cluster nodes at the primary site 1, and restores all the workloads to the original primary site 1. For detailed instructions about doing a manual failback to the primary site, see [Scenario 3: Full Failback to Primary Site](#).

## Scenario 2: Unscheduled Full Site Failure

In this scenario, a natural disaster strikes at the primary site 1 in Palo Alto, and site 1 goes down completely. The NSX administrator performs a manual failover to the secondary site 2 in Austin.

As the primary site has gone down due to unforeseen circumstances, the administrator cannot do any failover preparation before the actual failure occurs.

The NSX administrator wants to meet the following key objectives:

- Achieve a full site failover at site 2 with minimal downtime.
- Retain site 1 application IP addresses at site 2 after the failover.
- Automatically recover all Edge interface settings and BGP protocol configuration settings at site 2.

---

### Note

- The administrator can do the failover tasks manually by using either the vSphere Web Client or by running the NSX REST APIs. In addition, the administrator can automate some failover tasks by running a script file that contains the APIs to run during the failover. This scenario explains manual failover steps using the vSphere Web Client. However, if any step requires the use of either the CLI or the NSX REST APIs, adequate instructions are provided.
- In this scenario, the disaster recovery workflow is specific to the topology explained earlier, which has a primary NSX Manager and a single secondary NSX Manager. The workflow with multiple secondary NSX Managers is not in the scope of this scenario.

---

**Important** If the primary site 1 powers on while the failover to the secondary site 2 is in progress, first ensure that the failover process is completed by using the procedure in this scenario. Only after a clean failover is done to the secondary site 2, restore or failback all the workloads to the original primary site 1. For detailed instructions about the failback process, see [Scenario 3: Full Failback to Primary Site](#).

---

### Prerequisites

- NSX Data Center 6.4.5 or later is installed at both sites 1 and 2.
- vCenter Server at sites 1 and 2 are deployed with Enhanced Linked Mode.
- At site 1 and site 2, the following conditions are met:
  - No application-specific security policies are configured on a non-NSX firewall, if any.
  - No application-specific firewall rules are configured on a non-NSX firewall, if any.
  - Firewall is disabled on both the ESGs because ECMP is enabled on the UDLRs and to ensure that all traffic is allowed.
- At site 2, the following conditions are met before the failover:
  - Similar downlink interfaces are configured manually on the ESGs as configured at site 1.

- Similar BGP configuration is done manually on the ESGs as configured at site 1.
- ESGs are in powered down state when the primary site 1 is active or running.

## Procedure

- 1 Verify that the primary NSX Manager at site 1 is down.
  - a On the **Installation and Upgrade** page, navigate to **Management > NSX Managers**.
    - If you refresh the **NSX Managers** page in the current browser session, the role of the primary NSX Manager changes to *Unknown*.
    - If you log out from the vSphere Web Client and log in again or start a new vSphere Web Client browser session, the primary NSX Manager is no longer displayed on the **NSX Managers** page.
  - b Navigate to **Networking & Security > Dashboard > Overview**.
    - If you refresh the **Dashboard** page in the current browser session, the following error message is displayed: *Could not establish connection with NSX Manager. Please contact administrator..* This error means that the primary NSX Manager is no longer reachable.
    - If you log out from the vSphere Web Client and log in again or start a new vSphere Web Client browser session, the primary NSX Manager is no longer available in the **NSX Manager** drop-down menu.
- 2 Promote the secondary NSX Manager to a primary role.
  - a On the **Installation and Upgrade** page, navigate to **Management > NSX Managers**.
  - b Select the secondary NSX Manager.
  - c Click **Actions > Disconnect from Primary NSX Manager**. When prompted to continue with the disconnect operation, click **Yes**.
 

The secondary NSX Manager is disconnected from the primary NSX Manager, and enters into a *Transit* role.
  - d Click **Actions > Assign Primary Role**.
 

The secondary NSX Manager at site 2 is promoted to a primary role.

---

**Caution** As local egress is disabled on the UDLR, the UDLR Control VM (Edge Appliance VM) is deployed only at the original primary site (site 1). Before site 1 fails, the UDLR Control VM is not available at the secondary site (site 2), which is now promoted to primary. Therefore, redeploy the UDLR Control VM at the promoted primary site (site 2) before redeploying the NSX Controller Cluster.

If the controller nodes are deployed before deploying the UDLR Control VM, the forwarding tables on the UDLR are flushed out. This results in a downtime immediately after the first controller node is deployed at site 2. This situation might result in communication outages. To avoid this situation, deploy the UDLR Control VM before deploying the NSX Controller nodes.

---

- 3 Power on the NSX Edges that are in the powered down state, and deploy the UDLR Control VM (Edge Appliance VM) at the secondary site 2 (promoted primary).

For instructions about deploying the UDLR Control VM, see the *NSX Cross-vCenter Installation Guide*.

While deploying the UDLR Control VM, configure the following resource settings:

- Select the data center as **site 2**.
- Select the cluster/resource pool.
- Select the datastore.

---

**Note** After deploying the UDLR Control VM, the following configuration settings are automatically recovered at site 2:

- BGP protocol routing configuration
  - BGP password configuration
  - Uplink and internal interface settings
- 

- 4 Deploy the three NSX Controller Cluster nodes at site 2 (promoted primary).

For detailed instructions about deploying NSX Controllers, see the *NSX Cross-vCenter Installation Guide*.

- 5 Update the NSX Controller Cluster state.

- a On the **Installation and Upgrade** page, click **NSX Managers**.
- b Select the promoted primary NSX Manager.
- c Click **Actions > Update Controller State**.

- 6 Force sync routing service on each cluster at site 2.

- a On the **Installation and Upgrade** page, click **Host Preparation**.
- b Select the promoted primary NSX Manager.
- c Select one cluster at a time, and then click **Actions > Force Sync Services**.
- d Select **Routing**, and click **OK**.

- 7 Migrate the workload VMs from site 1 to site 2.

---

**Note** The workload VMs continue to exist at site 1. Therefore, you must manually migrate the workload VMs to site 2.

---

## Results

The manual recovery of NSX components and the failover from the primary site (site 1) to the secondary site (site 2) is complete.

**What to do next**

Verify whether the failover to site 2 is 100% complete by doing these steps on site 2 (promoted primary site):

- 1 Check whether the NSX Manager has the primary role.
- 2 Check whether the Control VM (Edge Appliance VM) is deployed on the UDLR.
- 3 Check whether the status of all controller cluster nodes is `Connected`.
- 4 Check whether the host preparation status is `Green`.
- 5 Log in to the CLI console of the UDLR Control VM (Edge Appliance VM), and do these steps:
  - a Check whether all BGP neighbors are established and the status is `UP` by running the `show ip bgp neighbors` command.
  - b Check whether all BGP routes are being learned from all BGP neighbors by running the `show ip route bgp` command.

After a complete failover to site 2, all workloads run on the secondary site (promoted primary) and traffic is routed through the UDLR and the NSX Edges at site 2.

## Scenario 3: Full Failback to Primary Site

In this scenario, primary site 1 is down either due to a scheduled maintenance or an unplanned power failure. All workloads are running on the secondary site 2 (promoted primary site), and traffic is being routed through the UDLR and the NSX Edges at site 2. Now, the original primary site 1 is up again and the NSX administrator wants to recover NSX components and restore all the workloads at the original primary site 1.

The NSX administrator wants to meet the following key objectives:

- Achieve a full failback of all workloads from site 2 to original primary site 1 with minimal downtime.
- Retain the application IP addresses after failback to site 1.
- Automatically recover all Edge interface settings and BGP protocol configuration settings at site 1.

---

### Note

- The administrator can do the failback tasks manually by using either the vSphere Web Client or by running the NSX REST APIs. In addition, the administrator can automate some failback tasks by running a script file that contains the APIs to run during the failback. This scenario explains manual failback steps using the vSphere Web Client. However, if any step requires the use of either the CLI or the NSX REST APIs, adequate instructions are provided.
  - In this scenario, the disaster recovery workflow is specific to the topology explained earlier, which has a primary NSX Manager and a single secondary NSX Manager. The workflow with multiple secondary NSX Managers is not in the scope of this scenario.
-

## Prerequisites

- NSX Data Center 6.4.5 or later is installed at both sites 1 and 2.
- vCenter Server at sites 1 and 2 are deployed with Enhanced Linked Mode.
- At site 1 and site 2, the following conditions are met:
  - No application-specific security policies are configured on a non-NSX firewall, if any.
  - No application-specific firewall rules are configured on a non-NSX firewall, if any.
  - Firewall is disabled on both the ESGs because ECMP is enabled on the UDLRs and to ensure that all traffic is allowed.
- At site 2 (promoted primary), no changes are made in the universal logical components before initiating the failback process.

## Procedure

- 1 When the primary site 1 is up again, make sure that the NSX Manager and the controller cluster nodes are powered on and running.
  - a Navigate to **Networking & Security > Dashboard > Overview**.
  - b Select the primary NSX Manager from the drop-down menu.
  - c In the **System Overview** pane, check the status of the NSX Manager and the controller cluster nodes.  
  
A filled green dot next to NSX Manager and the controller nodes mean that both the NSX components are powered on and running.
- 2 Before initiating the failback process, verify the following:
  - a On the **Installation and Upgrade** page, navigate to **Management > NSX Managers**, and observe that NSX Managers at both sites have a primary role.
  - b On the **NSX Controller Nodes** page, ensure that the Universal Controller Cluster (UCC) nodes exist at both sites.
- 3 Shut down all the three UCC nodes that are associated with site 2 (promoted primary).
- 4 On the **NSX Controller Nodes** page, delete all the three UCC nodes that are associated with site 2 (promoted primary).

---

**Tip** You can use the NSX REST APIs to remove one controller node at a time by running the following API call: `https://NSX_Manager_IP/api/2.0/vdn/controller/{controllerID}`. However, delete the last controller node forcefully by running the following API call: `https://NSX_Manager_IP/api/2.0/vdn/controller/{controllerID}?forceRemoval=true`.

---

- 5 Ensure that there are no changes in the universal components at site 2 (promoted primary) before proceeding to the next step.

- 6 Remove the primary role on the NSX Manager at site 2 (promoted primary).
  - a On the **Installation and Upgrade** page, navigate to **Management > NSX Managers**.
  - b Select the NSX Manager at site 2, and click **Actions > Remove Primary Role**.  
A message prompts you to ensure that the controllers owned by the NSX Manager at site 2 are deleted before removing the primary role.
  - c Click **Yes**.  
The NSX Manager at site 2 enters into a Transit role.
- 7 On the primary NSX Manager at site 1, remove the associated secondary NSX Manager.
  - a On the **NSX Managers** page, select the NSX Manager that is associated with site 1.
  - b Click **Actions > Remove Secondary Manager**.
  - c Select the **Perform Operation even if NSX Manager is inaccessible** check box.
  - d Click **Remove**.
- 8 Register the NSX Manager at site 2, which is in Transit, as the secondary of primary NSX Manager at site 1.

---

**Caution** Because local egress is disabled on the UDLR Control VM (Edge Appliance VM), the Control VM is automatically deleted. Therefore, before registering the NSX Manager at site 2 (currently in Transit role) with a secondary role, make sure that the controller cluster nodes at site 2 are deleted. If the controller cluster nodes are not deleted, network traffic disruption can occur.

---

- a On the **Installation and Upgrade** page, navigate to **Management > NSX Managers**.
- b Select the NSX Manager that is associated with site 1.
- c Click **Actions > Add Secondary Manager**.
- d Select the NSX Manager that is associated with site 2.
- e Enter the user name and password of the NSX Manager at site 2, and accept the security certificate.
- f Click **Add**.

After completing all these substeps, observe the following results:

- NSX Manager at site 1 has a primary role, and NSX Manager at site 2 has a secondary role.
- On the NSX Manager at site 2, three shadow controller nodes appear with status as `Disconnected`. The following message is displayed: `Can read or update controller cluster properties only on Primary or Standalone Manager.`

This message means that the secondary NSX Manager at site 2 is unable to establish connectivity with the Universal Controller Cluster nodes on the primary NSX Manager at site 1. However, after a few seconds, the connection gets reestablished and the status changes to Connected.

- 9 Power on the Control VM (Edge Appliance VM) on the UDLR and the NSX Edges at site 1.
  - a Navigate to **Networking > VMs > Virtual Machines**.
  - b Right-click the VM Name (VM ID) of the UDLR Control VM and click **Power on**.
  - c Repeat step (b) for the Edge VMs that you want to power on.
  - d Wait until the UDLR Control VM and Edge VMs are up and running before proceeding to the next step.
- 10 Make sure that the UDLR Control VM (Edge Appliance VM) that is associated with the secondary NSX Manager at site 2 is automatically deleted.
  - a Navigate to **Networking & Security > NSX Edges**.
  - b Select the secondary NSX Manager, and then click a UDLR.
  - c On the **Status** page, observe that no Edge Appliance VM is deployed on the UDLR.
- 11 Update the NSX Controller state on the primary site 1 so that the controller services are synced with the secondary site 2.
  - a On the **Installation and Upgrade** page, click **NSX Managers**.
  - b Select the primary NSX Manager at site 1.
  - c Click **Actions > Update Controller State**.
- 12 Migrate the workload VMs from site 2 to site 1.

---

**Note** The workload VMs continue to exist at site 2. Therefore, you must manually migrate the workload VMs to site 1.

---

## Results

The manual failback of all NSX components and workloads from the secondary site (site 2) to the primary site (site 1) is complete.

## What to do next

Verify whether the failback to primary site 1 is 100% complete by doing these steps on site 1:

- 1 Check whether the NSX Manager has the primary role.
- 2 Check whether the Control VM (Edge Appliance VM) is deployed on the UDLR.
- 3 Check whether the status of all controller cluster nodes is *Connected*.
- 4 Perform a Communication Health Check on each host cluster that is prepared for NSX.
  - a Navigate to **Installation and Upgrade > Host Preparation**.

- b Select the NSX Manager at site 1.
  - c Select one cluster at a time, and check whether the Communication Channel Health status of the cluster is UP.
  - d For each host in the cluster, check whether the Communication Channel Health status of the host is UP.
  - e Check whether the host preparation status is *Green*.
- 5 Log in to the CLI console of the UDLR Control VM (Edge Appliance VM), and do these steps:
- a Check whether all BGP neighbors are established and the status is UP by running the `show ip bgp neighbors` command.
  - b Check whether all BGP routes are being learned from all BGP neighbors by running the `show ip route bgp` command.

After a complete failback to site 1, all workloads run on the primary site 1 and traffic is routed through the UDLR and the NSX Edges at site 1.