NSX Administration Guide

Update 3

Modified on 20 NOV 2017

VMware NSX Data Center for vSphere 6.2

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

https://docs.vmware.com/

If you have comments about this documentation, submit your feedback to

docfeedback@vmware.com

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

Contents

NSX Administration Guide

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

Intended Audience

This manual is intended for anyone who wants to install or use NSX in a VMware vCenter 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 VMware Infrastructure 5.x, including VMware ESX, vCenter Server, and the vSphere Web Client.

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

1

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

Hardware

Table 1-1. Hardware Requirements

Appliance	Memory	vCPU	Disk Space
NSX Manager	16 GB (24 GB with certain NSX deployment sizes*)	4 (8 with certain NSX deployment sizes*)	60 GB
NSX Controller	4 GB	4	20 GB
NSX Edge	■ Compact: 512 MB ■ Large: 1 GB ■ Quad Large: 1 GB ■ X-Large: 8 GB	■ Compact: 1 ■ Large: 2 ■ Quad Large: 4 ■ X-Large: 6	■ Compact: 1 disk 500MB ■ Large: 1 disk 500MB + 1 disk 512MB ■ Quad-Large: 1 disk 500MB + 1 disk 512MB ■ X-Large: 1 disk 500MB + 1 disk 2GB
Guest Introspection	1 GB	2	4 GB
NSX Data Security	512 MB	1	6 GB per ESXi host

As a general guideline, you should increase NSX Manager resources to 8 vCPU and 24 GB of RAM if your NSX managed environment contains more than 256 hypervisors or more than 2000 VMs.

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*.

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, vCenter Server, and ESXi, see the release notes at https://docs.vmware.com/en/VMware-NSX-for-vSphere/index.html.

Note that 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 Servers in Enhanced Linked Mode. See Using Enhanced Linked Mode in *vCenter Server and Host Management* .

To check 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

- If you added ESXi hosts by name to the vSphere inventory, ensure that forward and reverse name resolution is working. 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 enabled on your Web browser, to access the NSX Manager user interface
- From NSX Manager, ensure port 443 is accessible from the ESXi host, the vCenter Server, and the NSX 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 Web Client you are using. See Using the vSphere Web Client in the *vCenter Server and Host Management* documentation for details.

Ports and Protocols Required by NSX 2

The following ports must be open for NSX to operate properly.

Table 2-1. Ports and Protocols required by NSX

Source	Target	Port	Protocol	Purpose	Sensitive	TLS	Authentication
Client PC	NSX Manager	443	TCP	NSX Manager Administrative Interface	No	Yes	PAM Authentication
Client PC	NSX Manager	80	TCP	NSX Manager VIB Access	No	No	PAM Authentication
ESXi Host	vCenter Server	443	TCP	ESXi Host Preparation	No	No	
vCenter Server	ESXi Host	443	TCP	ESXi Host Preparation	No	No	
ESXi Host	NSX Manager	5671	TCP	RabbitMQ	No	Yes	RabbitMQ User/Password
ESXi Host	NSX Controller	1234	TCP	User World Agent Connection	No	Yes	
NSX Controller	NSX Controller	2878, 2888, 3888	TCP	Controller Cluster - State Sync	No	Yes	IPsec
NSX Controller	NSX Controller	7777	TCP	Inter-Controller RPC Port	No	Yes	IPsec
NSX Controller	NSX Controller	30865	TCP	Controller Cluster - State Sync	No	Yes	IPsec
NSX Manager	NSX Controller	443	TCP	Controller to Manager Communication	No	Yes	User/Password
NSX Manager	vCenter Server	443	TCP	vSphere Web Access	No	Yes	
NSX Manager	vCenter Server	902	TCP	vSphere Web Access	No	Yes	
NSX Manager	ESXi Host	443	TCP	Management and provisioning connection	No	Yes	

Table 2-1. Ports and Protocols required by NSX (Continued)

Source	Target	Port	Protocol	Purpose	Sensitive	TLS	Authentication
NSX Manager	ESXi Host	902	TCP	Management and provisioning connection	No	Yes	
NSX Manager	DNS Server	53	TCP	DNS client connection	No	No	
NSX Manager	DNS Server	53	UDP	DNS client connection	No	No	
NSX Manager	Syslog Server	514	TCP	Syslog connection	No	No	
NSX Manager	Syslog Server	514	UDP	Syslog connection	No	No	
NSX Manager	NTP Time Server	123	TCP	NTP client connection	No	Yes	
NSX Manager	NTP Time Server	123	UDP	NTP client connection	No	Yes	
vCenter Server	NSX Manager	80	TCP	Host Preparation	No	Yes	
REST Client	NSX Manager	443	TCP	NSX Manager REST API	No	Yes	User/Password
VXLAN Tunnel End Point (VTEP)	VXLAN Tunnel End Point (VTEP)	8472 (default before NSX 6.2.3) or 4789 (default in new installs of NSX 6.2.3 and later)	UDP	Transport network encapsulation between VTEPs	No	Yes	
ESXi Host	ESXi Host	6999	UDP	ARP on VLAN LIFs	No	Yes	
ESXi Host	NSX Manager	8301, 8302	UDP	DVS Sync	No	Yes	
NSX Manager	ESXi Host	8301, 8302	UDP	DVS Sync	No	Yes	
Guest Introspection VM	NSX Manager	5671	TCP	RabbitMQ	No	Yes	RabbitMQ User/Password
Primary NSX Manager	Secondary NSX Manager	443	TCP	Cross-vCenter NSX Universal Sync Service	No	Yes	

Table 2-1. Ports and Protocols required by NSX (Continued)

Source	Target	Port	Protocol	Purpose	Sensitive	TLS	Authentication
Primary NSX Manager	vCenter Server	443	TCP	vSphere API	No	Yes	
Secondary NSX Manager	vCenter Server	443	TCP	vSphere API	No	Yes	
Primary NSX Manager	NSX Universal Controller Cluster	443	TCP	NSX Controller REST API	No	Yes	User/Password
Secondary NSX Manager	NSX Universal Controller Cluster	443	TCP	NSX Controller REST API	No	Yes	User/Password
ESXi Host	NSX Universal Controller Cluster	1234	TCP	NSX Control Plane Protocol	No	Yes	
ESXi Host	Primary NSX Manager	5671	TCP	RabbitMQ	No	Yes	RabbitMQ User/Password
ESXi Host	Secondary NSX Manager	5671	TCP	RabbitMQ	No	Yes	RabbitMQ User/Password

Ports for Cross-vCenter NSX and Enhanced Linked Mode

If you have a cross-vCenter NSX environment and your vCenter Server systems are in Enhanced Linked Mode, in order to manage any NSX Manager from any vCenter Server system each NSX Manager appliance must have the required connectivity to each vCenter Server system in the environment.

Overview of NSX

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 re-purpose 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. VMware NSX®, the network virtualization platform, is a key product in the SDDC architecture. With NSX, 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 network virtualization programmatically creates, snapshots, deletes, and restores software-based virtual networks. The result is a completely 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 is a completely non-disruptive solution. In fact, with NSX, the physical network infrastructure you already have is all you need to deploy a software-defined data center.

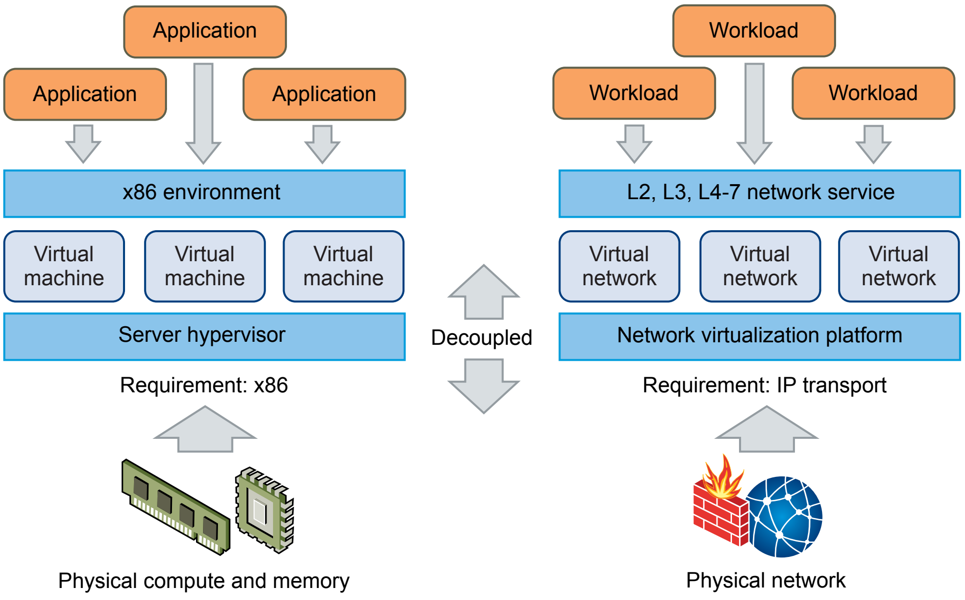

The figure above draws an analogy between compute and network virtualization. With server virtualization, a software abstraction layer (server hypervisor) reproduces the familiar attributes of an x86 physical server (for example, CPU, RAM, Disk, NIC) in software, allowing them to be programmatically 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 can be configured through the vSphere Web Client, a command-line interface (CLI), and a REST API.

This chapter includes the following topics:

- NSX Components
- NSX Edge
- NSX Services

NSX Components

This section describes the components of the NSX solution.

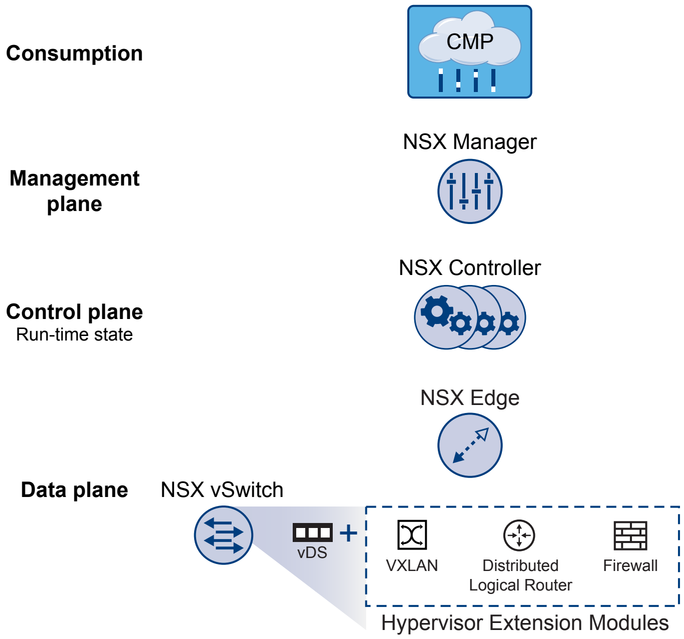

Note that a cloud management platform (CMP) is not a component of NSX, but NSX provides integration into virtually any CMP via the REST API and out-of-the-box integration with VMware CMPs.

Data Plane

The NSX data plane consists of the NSX vSwitch, which is based on the vSphere Distributed Switch (VDS) with additional components to enable services. NSX 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 vSwitch (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 NSX control plane runs in the NSX Controller cluster. NSX Controller is an advanced distributed state management system that provides control plane functions for NSX 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 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 Controllers must be deployed in a three-node cluster. 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 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 NSX management plane is built by the NSX Manager, the centralized network management component of NSX. It provides the single point of configuration and REST API entry-points.

The NSX Manager is installed as a virtual appliance on any ESX™ 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 Managers. 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 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 provides rich integration into virtually any CMP through REST APIs. Out-of-the-box integration is also available through VMware vCloud Automation Center, vCloud Director, and OpenStack with the Neutron plug-in for NSX.

NSX Edge

You can install NSX Edge as an edge services gateway (ESG) or as a distributed logical router (DLR). The number of edge appliances including ESGs and DLRs is limited to 250 on a host.

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 datacenter. 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 ESX 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

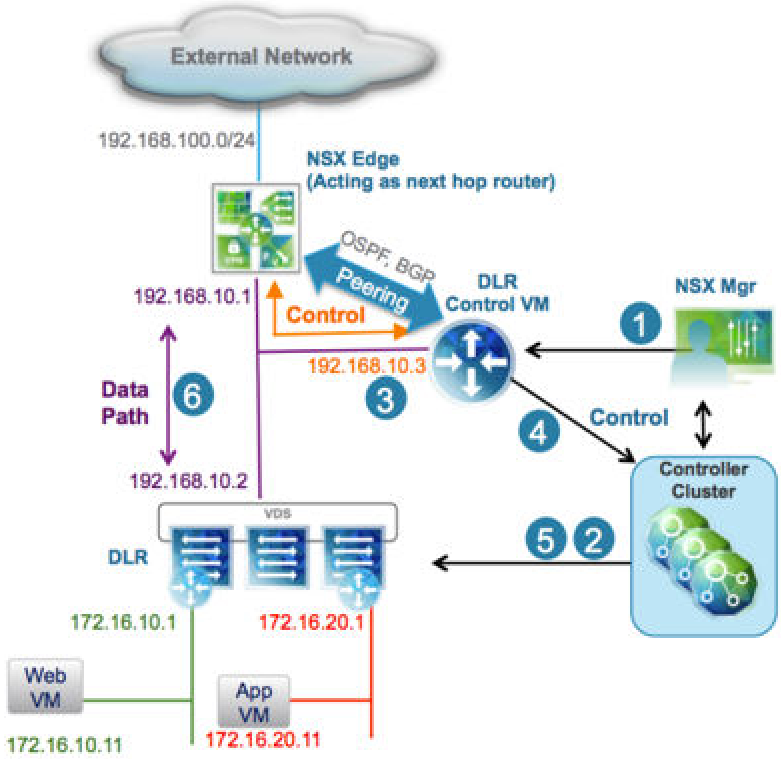

1 A DLR instance is created from the NSX Manager UI (or with API calls), and routing is enabled, leveraging either OSPF or BGP.

2 The NSX Controller leverages 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 components work together to provide the following functional 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 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 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, NSX 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 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.

Data Security provides visibility into sensitive data stored within your organization's virtualized and cloud environments and reports any data security violations.

NSX Extensibility

3rd-party solution providers can integrate their solutions with the NSX 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 6.2 allows you to manage multiple vCenter NSX environments from a single primary NSX Manager.

This chapter includes the following topics:

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 6.1 and earlier, if multiple vCenter NSX environments are deployed, they must be managed separately. In NSX 6.2 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 NSX logical networks. The same logical networks are available across the vCenter NSX 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 new 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 vMotioned across vCenter Servers without having to reconfigure the VM or change firewall rules.
- Enhanced NSX multi-site and disaster recovery capabilities.

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

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. The primary NSX Manager can be used to configure any of the secondary NSX Managers in the environment.

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

NSX Managers can be assigned the standalone role. This is equivalent to pre-NSX 6.2 environments 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 will be 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. The transit role should only be used temporarily, for example, when changing which NSX Manager is the primary.

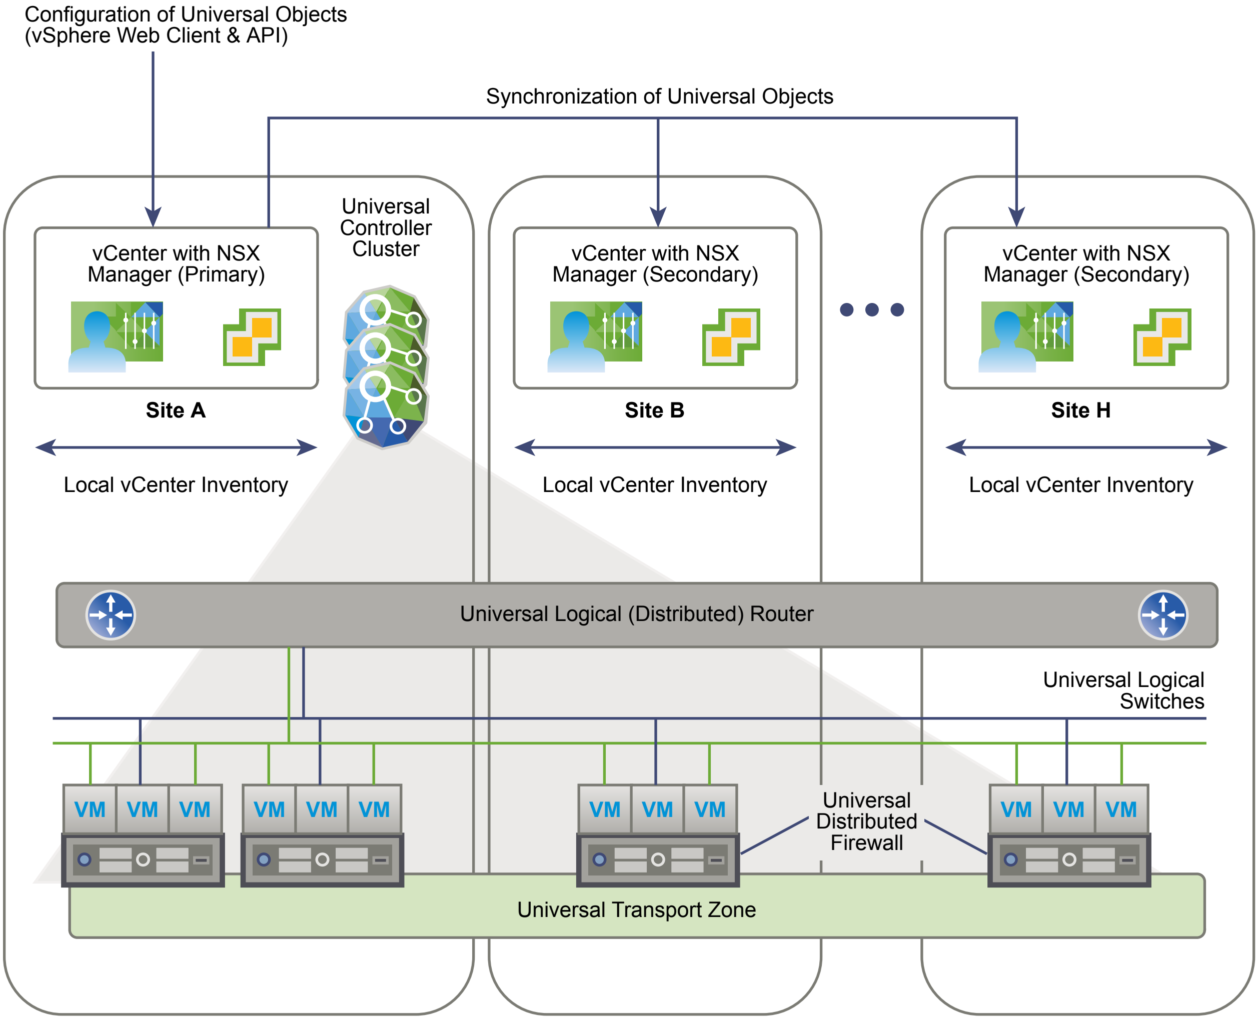
Configuration of Universal Objects (vSphere Web Client & API)
Synchronization of Universal Objects
Universal Controller Cluster
vCenter with NSX Manager (Primary)
Site A
Local vCenter Inventory
vCenter with NSX Manager (Secondary)
Site B
Local vCenter Inventory
vCenter with NSX Manager (Secondary)
Site H
Local vCenter Inventory
Universal Logical (Distributed) Router
Universal Logical Switches
VM VM VM
VM VM VM
VM VM VM
VM VM VM
Universal Distributed Firewall
Universal Transport Zone

Support Matrix for NSX Services in Cross-vCenter NSX

A subset of NSX Services are available for universal synchronization in cross-vCenter NSX.

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

NSX Service	Details	Supports cross-vCenter NSX synchronization in NSX 6.2?
Logical switch	Transport zone	Yes
	Logical switch	Yes

Table 4-1. Support matrix for NSX Services in cross-vCenter NSX (Continued)

NSX Service	Details	Supports cross-vCenter NSX synchronization in NSX 6.2?
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
Data security		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 universal security groups can contain only included objects, no dynamic membership or excluded objects
	Services	Yes
	Service groups	Yes

Universal 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 a vCenter NSX pair.

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 a distributed firewall rule section that is marked for universal synchronization. You can create one universal L2 rule section and one universal L3 rule section. 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 IP Sets
- Universal MAC Sets
- Universal Security Groups
- Universal Services
- Universal Service Groups

Universal network and security objects can be created only from the primary NSX Manager.

Universal security groups can contain only universal IP sets, universal MAC sets, and universal security groups. Membership is defined by included objects only, you cannot use dynamic membership or 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 vCenter NSX setups. The vCenters 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.

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

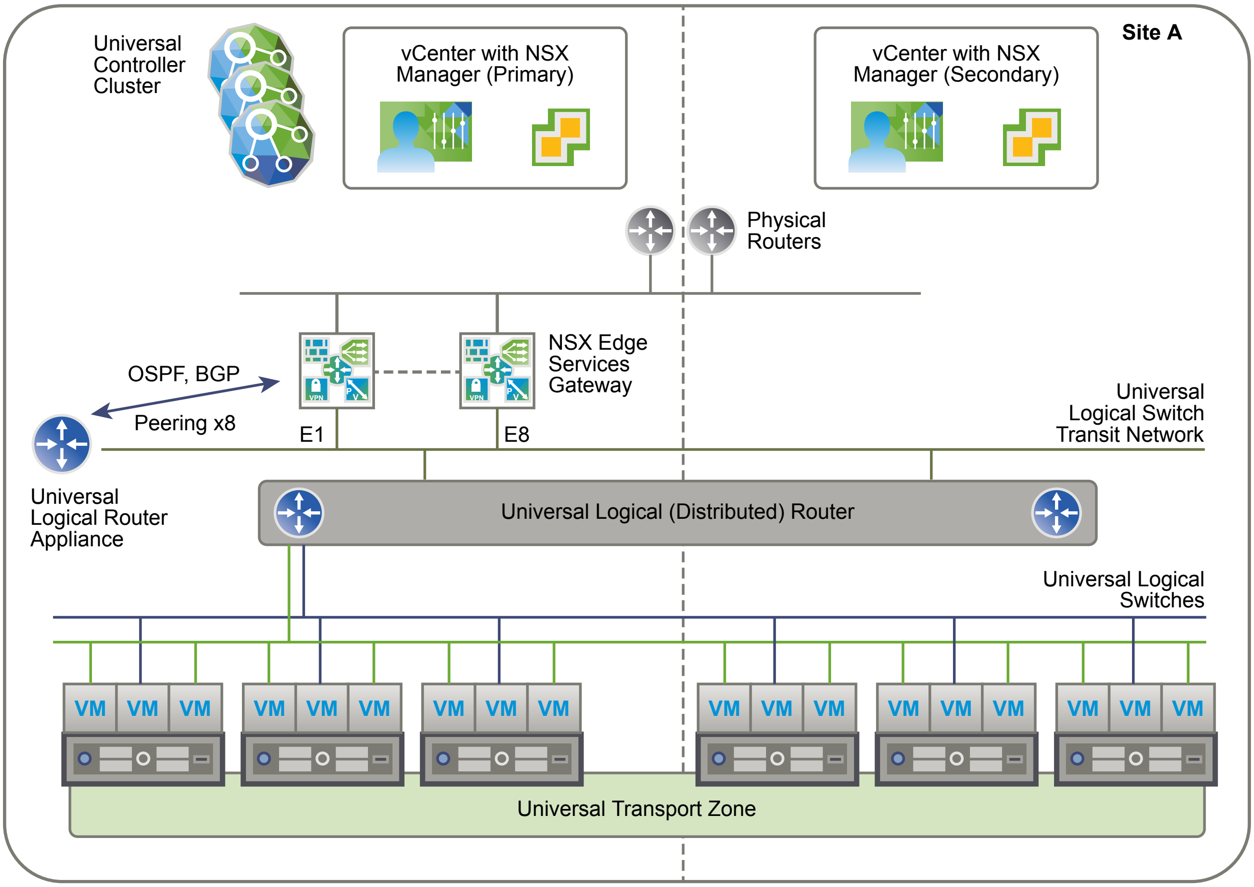
Site A
Universal Controller Cluster
vCenter with NSX Manager (Primary)
vCenter with NSX Manager (Secondary)
Physical Routers
NSX Edge Services Gateway
OSPF, BGP
Peering x8
E1
E8
Universal Logical Switch Transit Network
Universal Logical Router Appliance
Universal Logical (Distributed) Router
Universal Logical Switches
VM
Universal Transport Zone

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

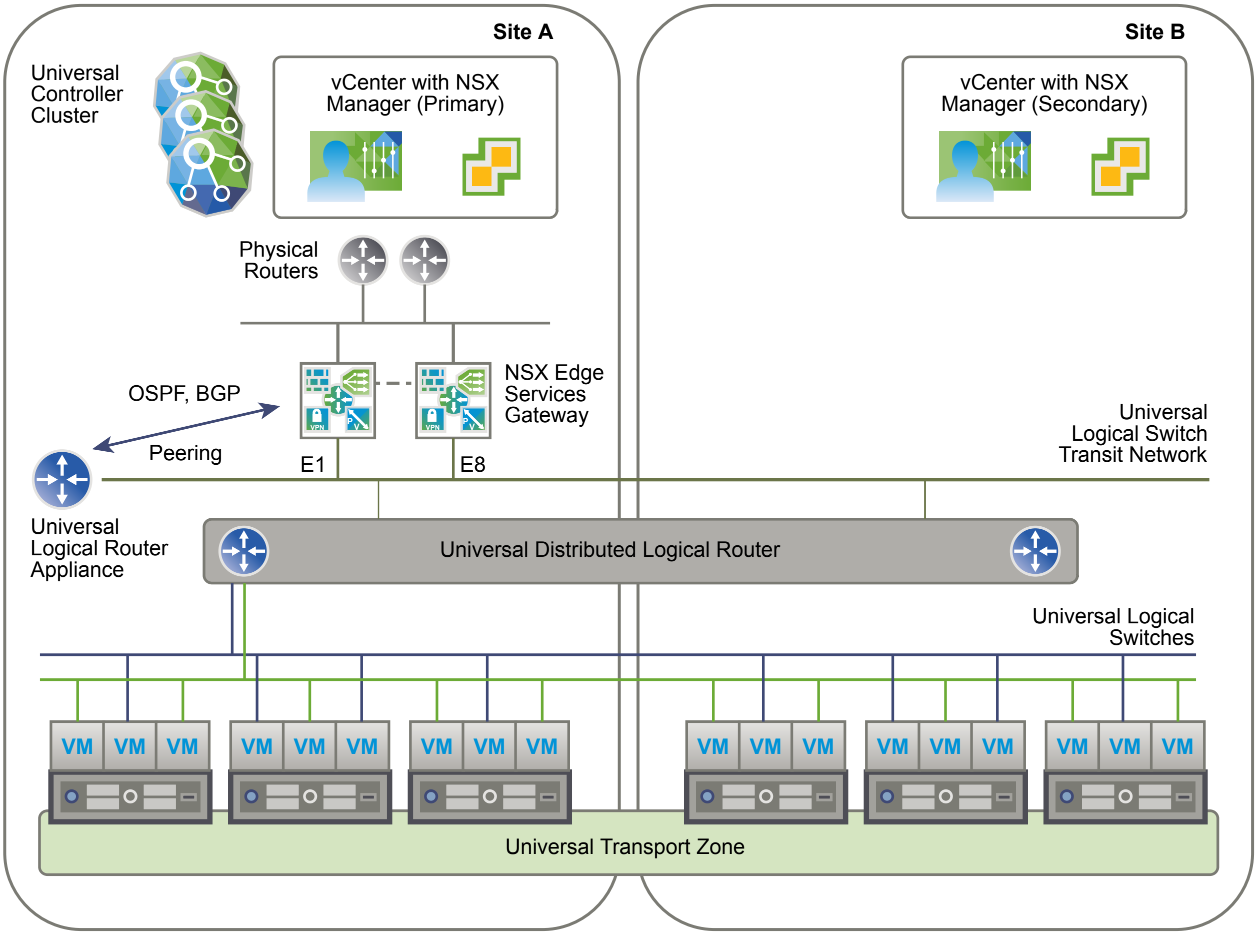
Site A
Site B
Universal Controller Cluster
vCenter with NSX Manager (Primary)
vCenter with NSX Manager (Secondary)
Physical Routers
OSPF, BGP
Peering
NSX Edge Services Gateway
E1
E8
Universal Logical Switch Transit Network
Universal Logical Router Appliance
Universal Distributed Logical Router
Universal Logical Switches
VM
Universal Transport Zone

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. This 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.

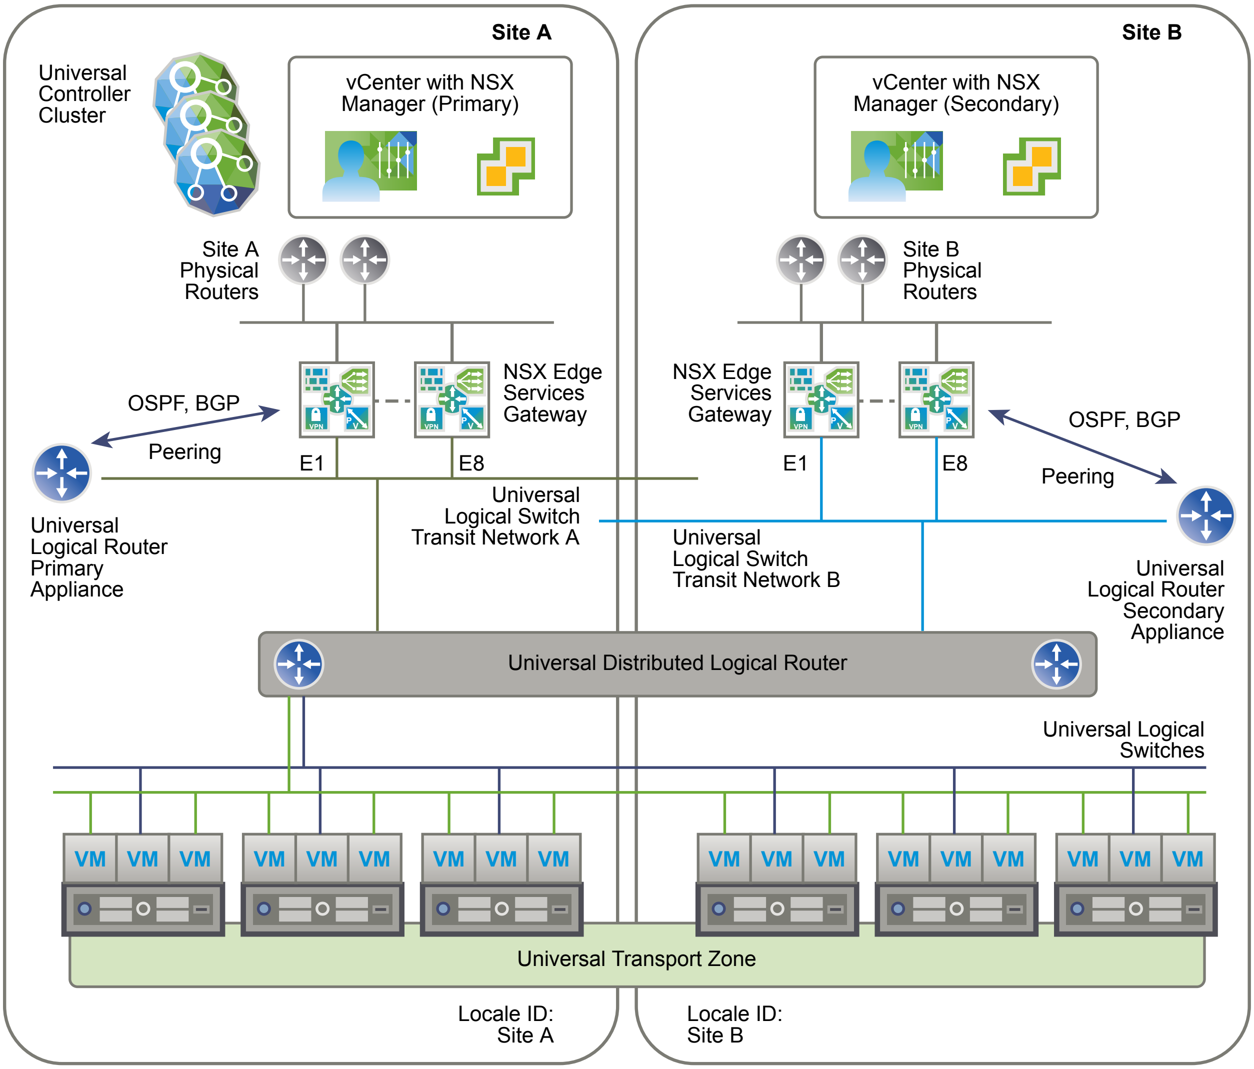
Site A
Universal Controller Cluster
vCenter with NSX Manager (Primary)
Site A Physical Routers
NSX Edge Services Gateway
OSPF, BGP
Peering
E1
E8
Universal Logical Router Primary Appliance
Universal Logical Switch Transit Network A
Site B
vCenter with NSX Manager (Secondary)
Site B Physical Routers
NSX Edge Services Gateway
E1
E8
OSPF, BGP
Peering
Universal Logical Switch Transit Network B
Universal Logical Router Secondary Appliance
Universal Distributed Logical Router
Universal Logical Switches
VM
Universal Transport Zone
Locale ID: Site A
Locale ID: Site B

Modifying NSX Manager Roles

An NSX Manager can have the primary role, the secondary role, or the standalone role. 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 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 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. Similarly, an edge services gateway (ESG) has access to logical switches from only one transport zone.

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

NSX 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. Similarly, an edge services gateway (ESG) has access to logical switches from only one 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.

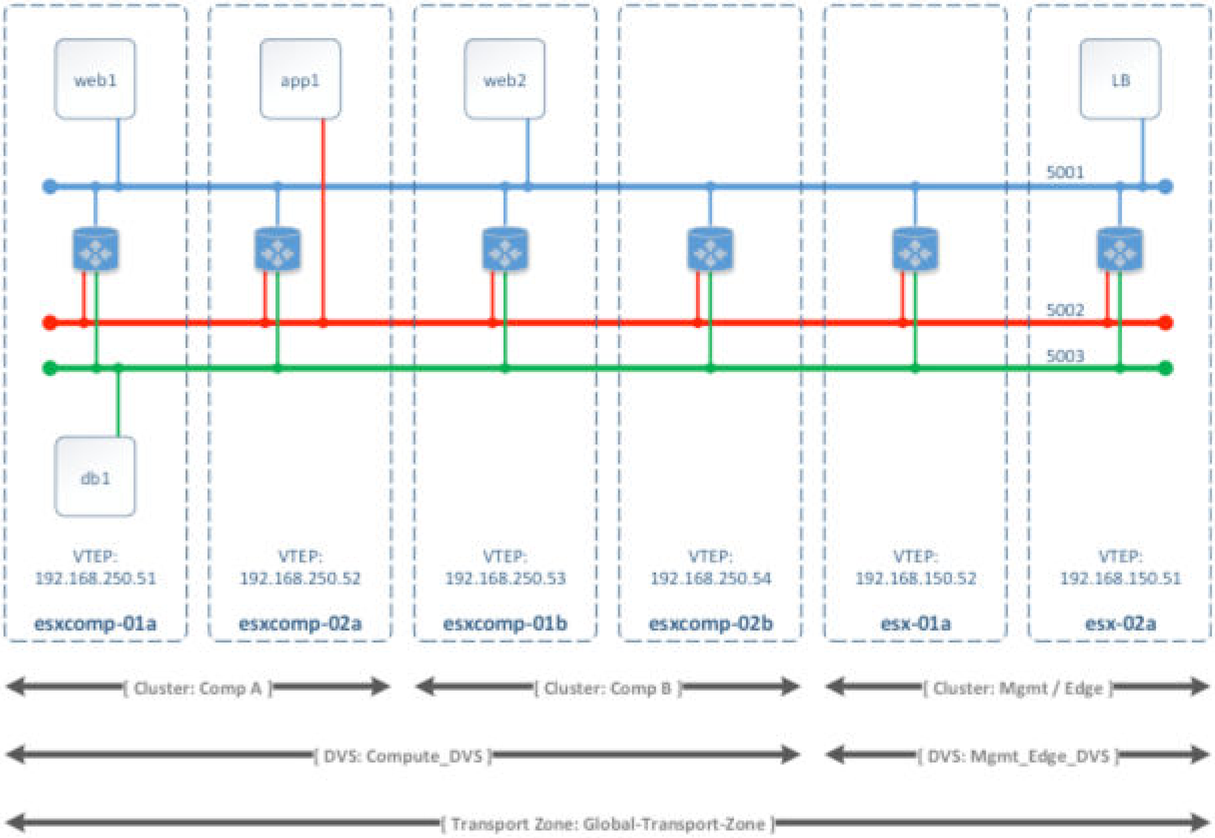

If you do not follow this best practice, keep in mind that if a VDS spans more than one host cluster and the transport zone includes only one (or a subset) of these clusters, any logical switch included within this transport zone can access VMs within all clusters spanned by the VDS. In other words, 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).

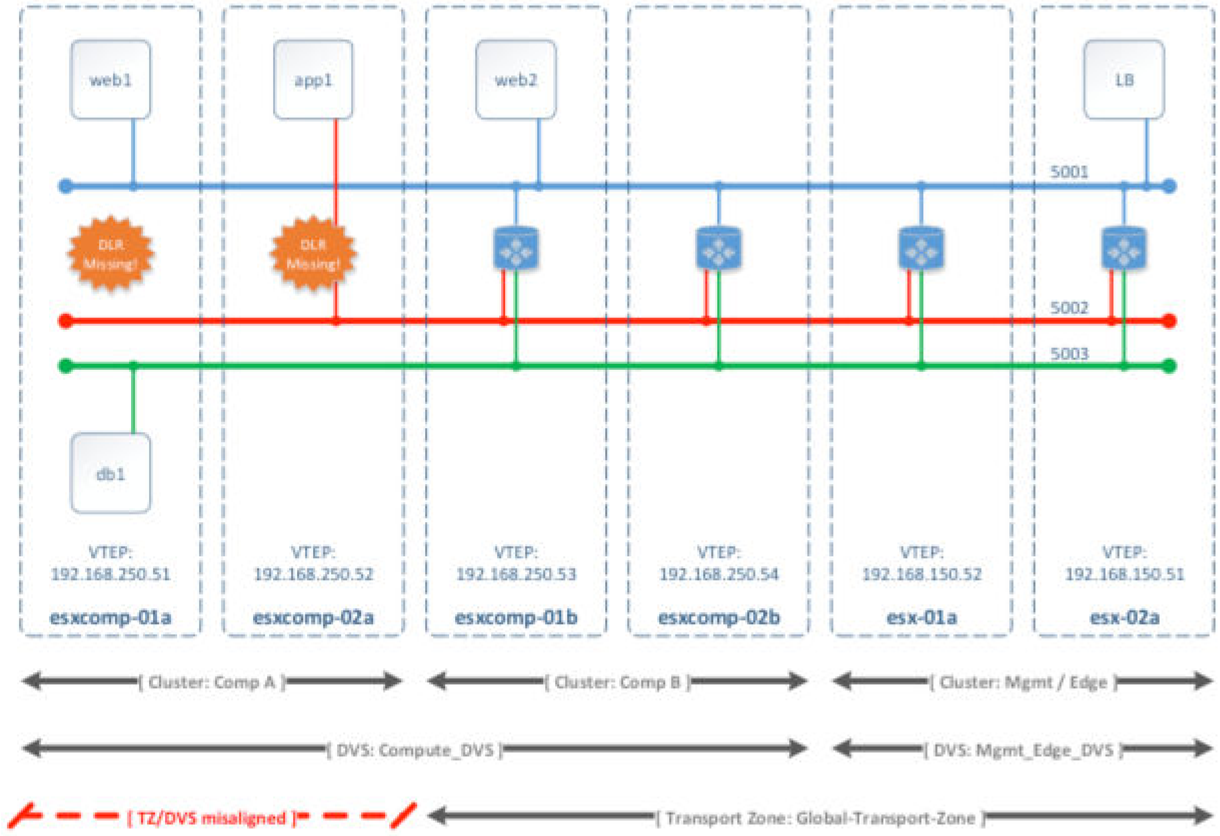

This chapter includes the following topics:

- Add a Transport Zone
- View and Edit a Transport Zone
- Expand a Transport Zone
- Contract a Transport Zone

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. In vCenter, navigate to **Home > Networking & Security > Installation** and select the **Logical Network Preparation** tab.
2. Click **Transport Zones** and click the **New Transport Zone (✚)** icon.
3. (Optional) To add a universal transport zone, 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 headend 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

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.

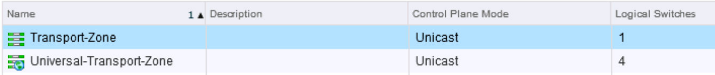

Name 1 ▲	Description	Control Plane Mode	Logical Switches
Transport-Zone		Unicast	1
Universal-Transport-Zone		Unicast	4

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.

View and Edit a Transport Zone

You can view the logical networks in a selected transport zone, the clusters in, and the control plane mode for that transport zone.

Procedure

1 In Transport Zones, double-click a transport zone.

  The Summary tab displays the name and description of the transport zone as well as the number of logical switches associated with it. Transport Zone Details displays the clusters in the transport zone.

2 Click the **Edit Settings** icon in the **Transport Zone Details** section to edit the name, description, or control plane mode of the transport zone.

  If you change the transport zone control plane mode, select **Migrate existing Logical Switches to the new control plane mode** to change the control plane more 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 control plane mode.

3 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

1 In Transport Zones, click a transport zone.

2 Click the **Add Cluster** (✚) icon.

3 Select the clusters you want to add to the transport zone and click **OK**.

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 In **Transport Zones**, double-click a transport zone.

2 In **Transport Zones Details**, click the **Remove Clusters** () icon.

3 Select the clusters that you want to remove.

4 Click **OK**.

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.

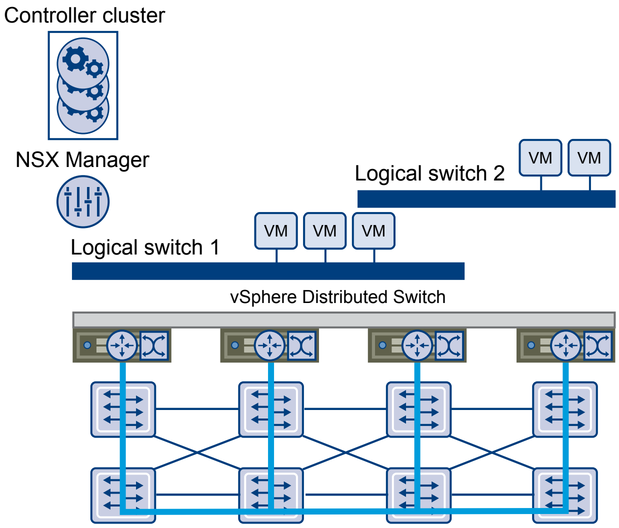

The NSX controller is the central control point for all logical switches within a network and maintains information of all virtual machines, hosts, logical switches, and VXLANs. The controller supports two new logical switch control plane modes, Unicast and Hybrid, These modes decouple NSX from the physical 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

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).
- 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 ESX 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 completely decoupled from 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.

Prerequisites

Table 6-1. Prerequisites for creating a logical switch or universal logical switch

Logical Switch	Universal Logical Switch
■ 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.	■ 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 In the vSphere Web Client, 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 the **New Logical Switch (➕)** icon.

For example:

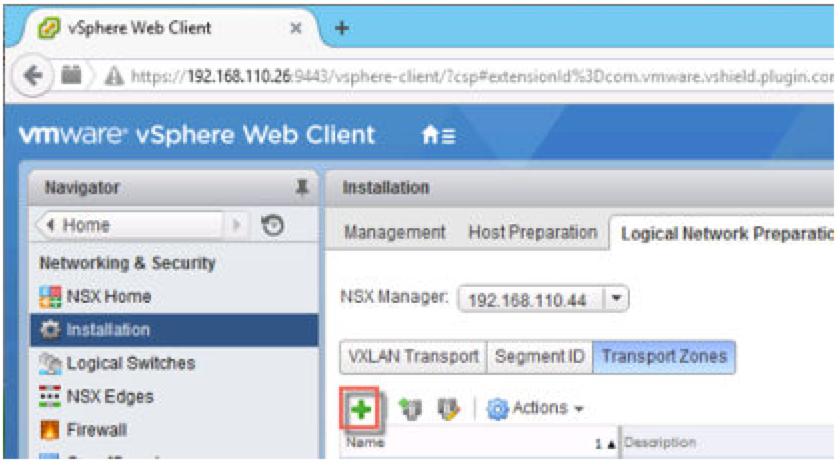

4 Type a name and optional description for the logical switch.

5 Select the transport zone in which you want to create the logical switch. Selecting a universal transport zone will create a universal logical switch.

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 any other multicast addresses assigned on any NSX Manager in the environment.

6 (Optional) Click **Enable IP Discovery** to enable ARP suppression.

This setting 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.

7 (Optional) Click **Enable MAC learning** if your VMs have multiple MAC addresses or are using virtual NICs that are trunking VLANs.

Enabling MAC learning builds a VLAN/MAC pair learning table on each vNIC. This table is stored as part of the dvfilter data. During vMotion, dvfilter saves and restores the table at the new location. The switch then issues RARPs for all the VLAN/MAC entries in the table.

This example shows the app logical switch with default settings.

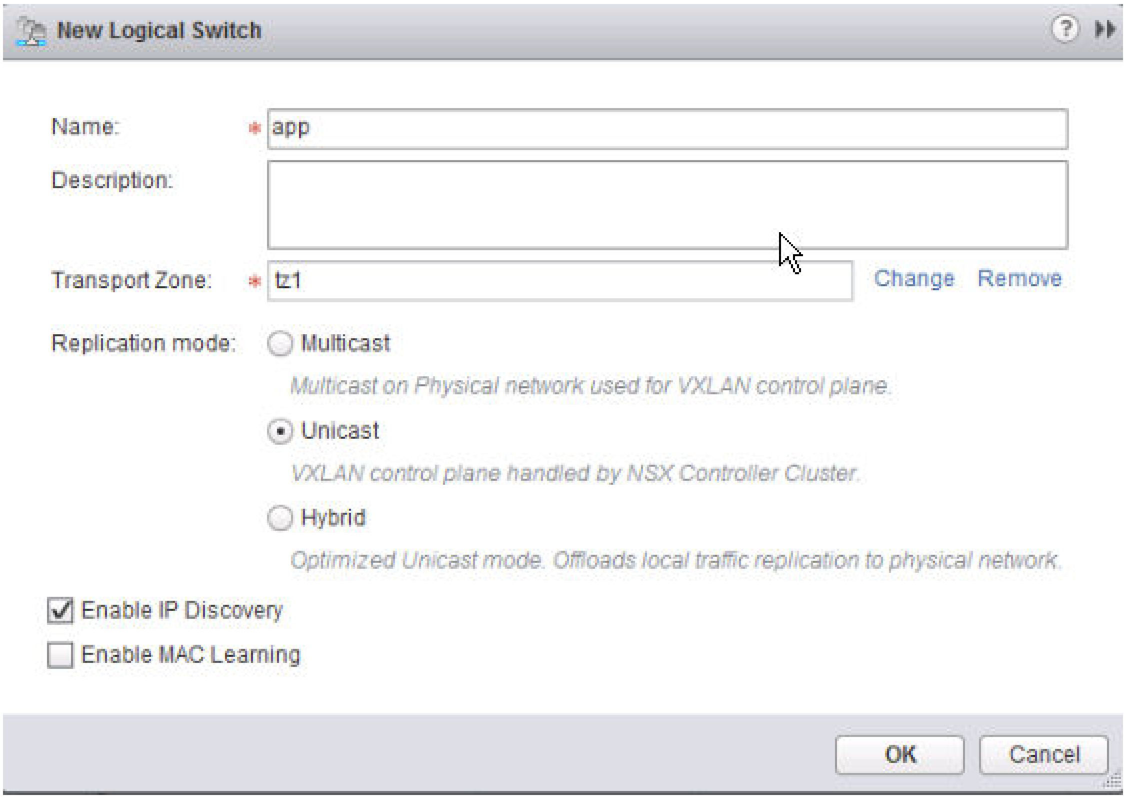

Name: * app

Description:

Transport Zone: * tz1

Replication mode:

○ Multicast

Multicast on Physical network used for VXLAN control plane.

● Unicast

VXLAN control plane handled by NSX Controller Cluster.

○ Hybrid

Optimized Unicast mode. Offloads local traffic replication to physical network.

☑ Enable IP Discovery

□ Enable MAC Learning

DB-Tier-00 is logical switch connected to a transport zone. It is available only on the NSX Manager on which it was created.

DB-Tier-01 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.

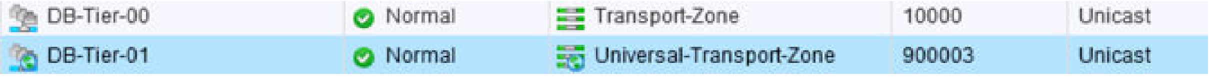

DB-Tier-00	✓ Normal	Transport-Zone	10000	Unicast
DB-Tier-01	✓ Normal	Universal-Transport-Zone	900003	Unicast

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

1. In Logical Switches, select the logical switch to which you want to connect an NSX Edge.
2. Click the **Connect an Edge** () icon.
3. Select the NSX Edge to which you want to connect the logical switch and click **Next**.
4. Select the interface that you want to connect to the logical switch and click **Next**.

   A logical network is typically connected to an internal interface.
5. On the Edit NSX Edge interface page, type a name for the NSX Edge interface.
6. Click **Internal** or **Uplink** to indicate whether this is an internal or uplink interface.
7. Select the connectivity status of the interface.
8. 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.
9. Edit the default MTU if required.
10. Click **Next**.
11. Review the NSX Edge connection details and click **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 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** (➕) icon.
3. Select the virtual machines you want to add to the logical switch.
4. Select the vNICs that you want to connect.
5. Click **Next**.
6. Review the vNICs you selected.
7. Click **Finish**.

Test Logical Switch Connectivity

A ping test checks if 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 in the **Name** column.
2. Click the **Monitor** tab.
3. Click the **Hosts** tab.
4. Click **Browse** in the Source Host section. In the Select Host dialog box, select the destination host.
5. Select the size of the test packet.

   VXLAN standard size is 1550 bytes (should match the physical infrastructure MTU) without fragmentation. This allows NSX to check connectivity and verify that the infrastructure is prepared for VXLAN traffic.

   Minimum packet size allows fragmentation. Hence, with packet size minimized, NSX can check connectivity but not whether the infrastructure is ready for the larger frame size.
6. Click **Browse** in the Destination Host section. In the Select Host dialog box, select the destination host.
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 13 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 **OK**.

Logical Switch Scenario

This scenario presents a situation where company ACME Enterprise has several ESX 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 5.5.

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

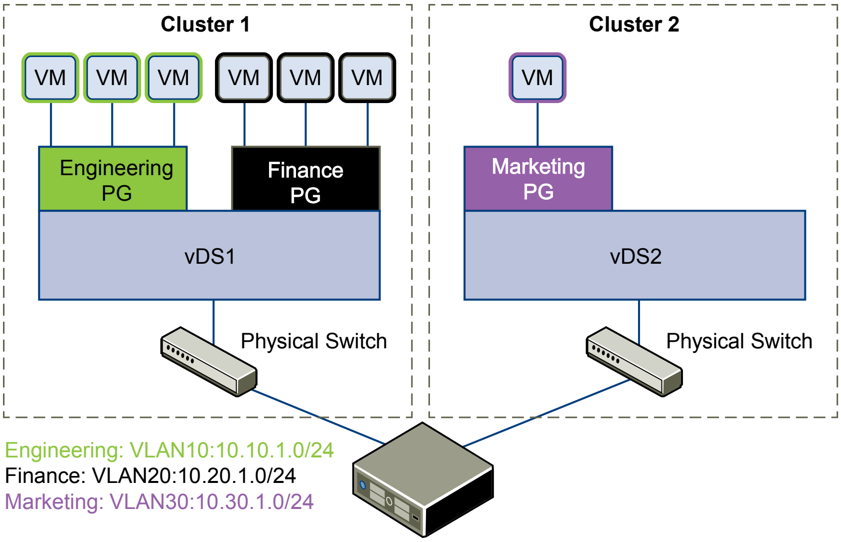

ACME is running out of compute space on Cluster1 while Cluster2 is under-utilized. The ACME network supervisor asks John Admin (ACME's virtualization administrator) to figure out a way to extend the Engineering department to Cluster2 in a way that virtual machines belonging to Engineering on both clusters 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

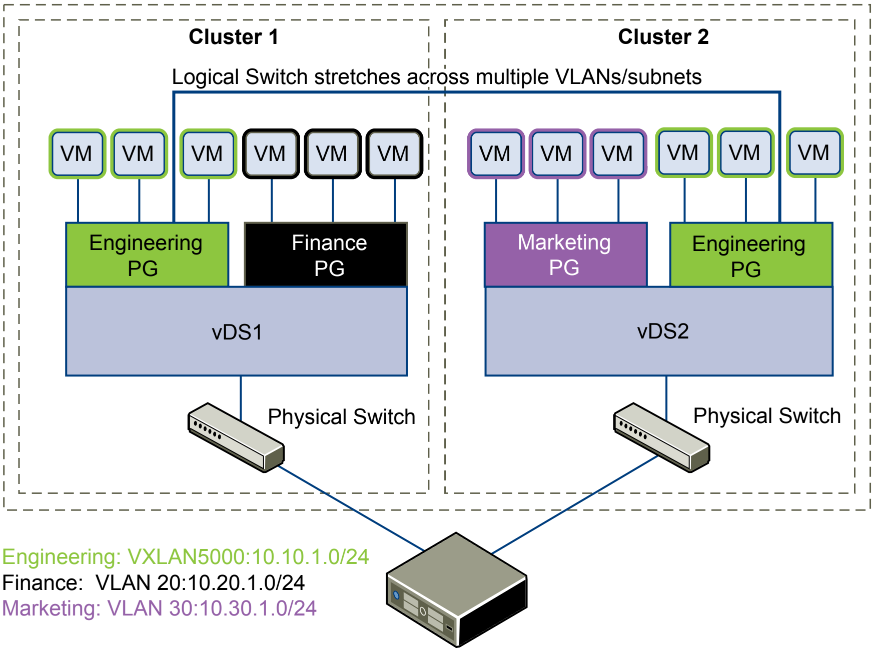
Cluster 1
Cluster 2
Logical Switch stretches across multiple VLANs/subnets
VM
Engineering PG
Finance PG
Marketing PG
Engineering PG
vDS1
vDS2
Physical Switch
Physical Switch
Engineering: VXLAN5000:10.10.1.0/24
Finance: VLAN 20:10.20.1.0/24
Marketing: VLAN 30:10.30.1.0/24

Once John Admin builds a logical switch across the two clusters, he can vMotion virtual machines within the vDS.

Figure 6‑3. vMotion on a logical network

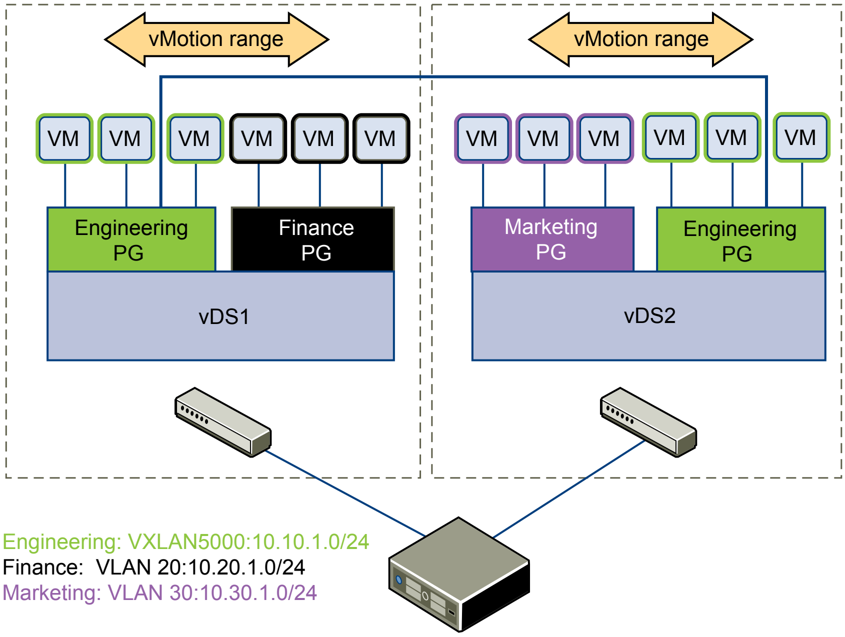

Let us walk through the steps that John Admin follows to build a logical network at ACME Enterprise.

John Admin Assigns Segment ID Pool and Multicast Address Range 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 VMware distributed switches version 5.5.

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**.

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 vDS. 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 virtual 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.

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.

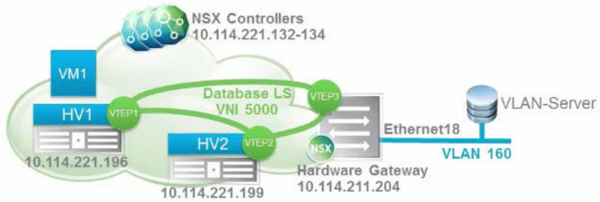

The hardware gateway underlying configuration can have any one of the following components:

- 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.

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.

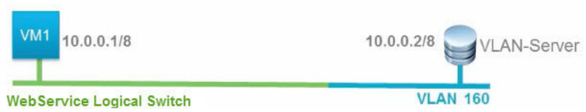

Prerequisites

- 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.
- 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 the *NSX Upgrade* guide.

  You can use the REST API command, `PUT /2.0/vdn/config/vxlan/udp/port/4789` to modify the port number. This API returns no response.

Procedure

1. Set Up the Replication Cluster

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

2. Connect the Hardware Gateway to the NSX Controllers

   You must configure the the OVSDB manager table on the ToR 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 broadcast traffic sent from the hardware gateway. The broadcast traffic can be unknown unicast and multicast traffic,

Note The replication nodes and the hardware gateway switches cannot be on the same IP subnet.

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.

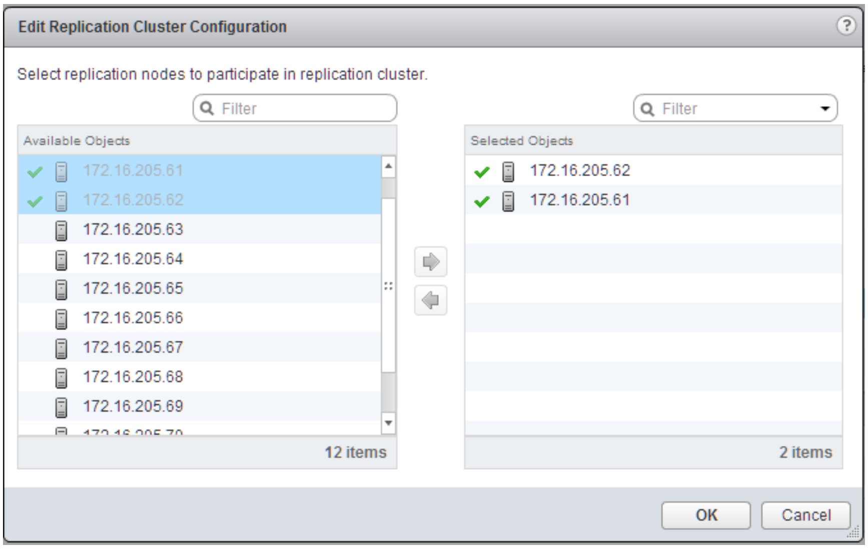

Edit Replication Cluster Configuration

Select replication nodes to participate in replication cluster.

Available Objects
✓ 172.16.205.61
✓ 172.16.205.62
172.16.205.63
172.16.205.64
172.16.205.65
172.16.205.66
172.16.205.67
172.16.205.68
172.16.205.69
172.16.205.70
12 items

Selected Objects
✓ 172.16.205.62
✓ 172.16.205.61
2 items

The selected hypervisors move to the selected objects column.

6 Click **OK**.

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 ToR physical switch to connect the hardware gateway to the NSX Controller.

The Controller passively listens to the connection attempt from ToR. Therefore, the hardware gateway must use the OVSDB manager table to initiate connection.

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 > 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.

4 Click the Add (✚) icon to create the hardware gateway profile details.

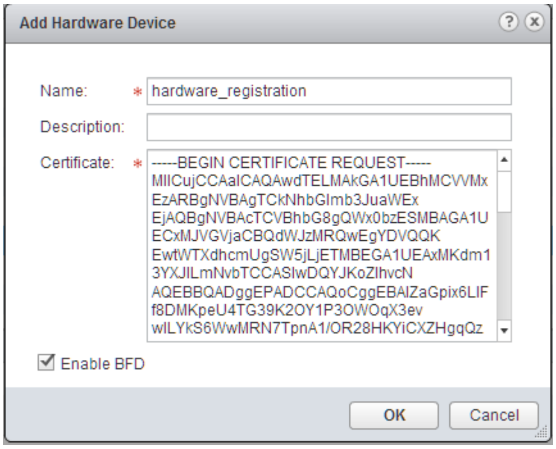

Add Hardware Device

Name: * hardware_registration

Description:

Certificate: *

```
-----BEGIN CERTIFICATE REQUEST-----
MIICujCCAaICAQAwdTELMAkGA1UEBhMCVVMx
EzARBgNVBAgTCkNhbGlmb3JuaWEx
EjAQBgNVBAcTCVBhbG8gQWx0bzESMBAGA1U
ECxMJVGVjaCBQdWJzMRQwEgYDVQQK
EwtWTXdhcmUgSW5jLjETMBEGA1UEAxMKdm1
3YXJILmNvbTCCASIwDQYJKoZIhvcN
AQEBBQADggEPADCCAQoCggEBAIZaGpix6LIF
f8DMKpeU4TG39K2OY1P3OWOqX3ev
wILYkS6WwMRN7TpnA1/OR28HKYiCXZHgqQz
```

☑ Enable BFD

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.

5 Click **OK**.

A profile that represents the hardware gateway is created.

6 Refresh the screen to verify that the hardware gateway is available and running.

The connectivity should be UP.

7 (Optional) Click the hardware gateway profile and right-click to select **View the BFD Tunnel Status** from the drop-down menu.

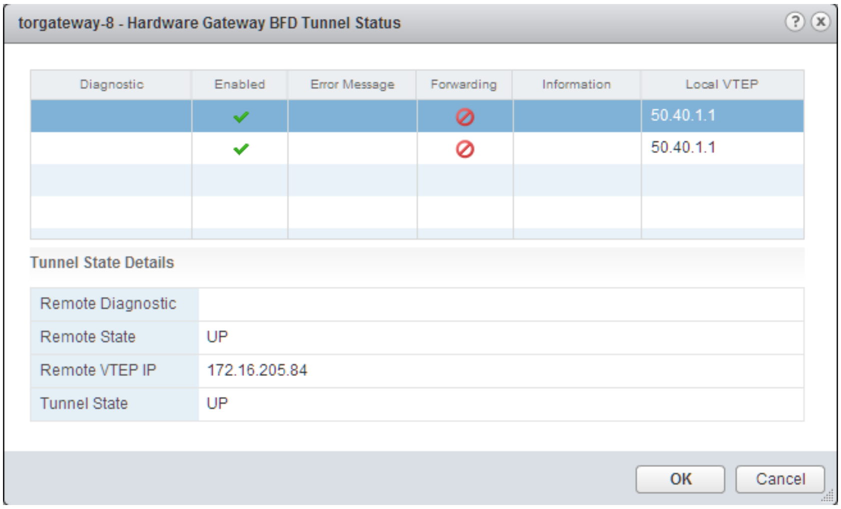

torgateway-8 - Hardware Gateway BFD Tunnel Status

Diagnostic	Enabled	Error Message	Forwarding	Information	Local VTEP
	✓		⊘		50.40.1.1
	✓		⊘		50.40.1.1

Tunnel State Details

Remote Diagnostic	
Remote State	UP
Remote VTEP IP	172.16.205.84
Tunnel State	UP

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 Harware 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.

   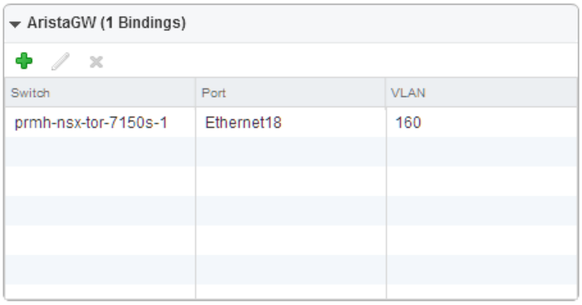

   AristaGW (1 Bindings)

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

   For example, 160.
9. Click **OK**.

The binding is complete.

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

8 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 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 Edge logical router virtual machine. An L2 bridge instance maps to a single VLAN, but there can be multiple bridge instances. The logical router cannot be used as a gateway for devices connected to a bridge.

If High Availability is enabled on the Logical Router and the primary NSX Edge virtual machine goes down, the bridge is automatically moved over to the host with the secondary virtual machine. For this seamless migration to happen, a VLAN must have been configured on the host that has the secondary NSX Edge virtual machine.

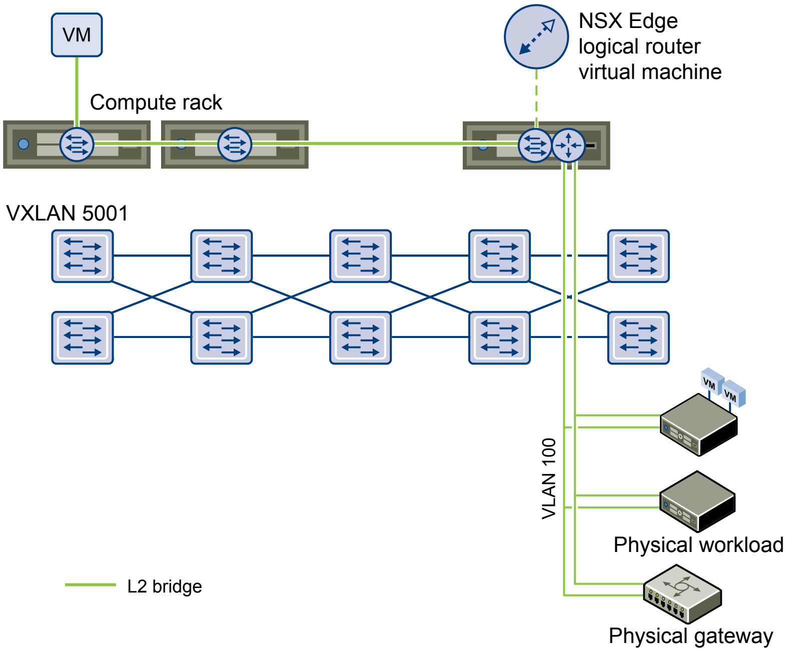

Note that you should not use an L2 bridge to connect a logical switch to another logical switch, a VLAN network to another VLAN network, or to interconnect datacenters. Also, you cannot use a universal logical router to configure bridging and you cannot add a bridge to a universal logical switch.

This chapter includes the following topics:

- Add L2 Bridge
- Add L2 Bridge to a Logically Routed Environment

Add L2 Bridge

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

Prerequisites

An NSX 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** and then click **NSX Edges**.
3. Double click a logical router.
4. Click **Manage** and then click **Bridging**.
5. Click the **Add (+)** icon.
6. Type a name for the bridge.
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 **OK**.

Add L2 Bridge to a Logically Routed Environment

One logical router can have multiple bridging instances, however, the routing and bridging instances cannot share the same vxlan/vlan network. Traffic to and from the bridged vlan and bridged vxlan cannot be routed to the bridged network and vice versa.

Prerequisites

- An NSX 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** and then click **NSX Edges**.

3. Double click the logical router that will be used for bridging.

   Note The bridge instance has to be created in the same routing instance to which the vxlan is connected. One bridge instance can have one vxlan and one vlan, the vxlan and vlan cannot overlap. The same vxlan and vlan cannot connect to more than one bridge instances.

4. Click **Manage** and then click **Bridging**.

   The logical switch that is being used as a router will state Routing Enabled.

5. Click the **Add (✚)** icon.
6. Type a name for the bridge.
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 **OK**.
10. Click **OK** again in the Add Bridge window.
11. Click Publish for changes to the bridging configuration to take effect.

    The logical switch that is used for bridging will now appear with **Routing Enabled** specified. For more information, see Add a Logical Switch and Connect Virtual Machines to a Logical Switch.

Routing

9

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 Logical (Distributed) Router

Logical router 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. Logical routers can be created on both primary and secondary NSX Managers in a cross-vCenter NSX environment, but universal logical routers can be created only on the primary NSX Manager.

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, Considerations for Management Interface of Distributed Logical Router Control VM, http://kb.vmware.com/kb/2122060.

Prerequisites

- You must have been assigned the Enterprise Administrator or NSX Administrator role.
- You must have an operational controller cluster in your environment before installing a logical router.
- You must create a local segment ID pool, even if you have no plans to create NSX logical switches.
- 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. Make sure the controller cluster is up and available before creating or changing a logical router configuration.
- 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 for logical router VLAN-based ARP proxy to work.
- 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 portgroups 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 should not be created on two different distributed portgroups (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 the same 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.
- Unlike NSX version 6.0 and 6.1, NSX version 6.2 allows logical router-routed logical interfaces (LIFs) to be connected to a VXLAN that is bridged to a VLAN.

- When selecting placement of a logical router virtual appliance, 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, thus 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 *VMware NSX for vSphere Network Virtualization Design Guide* at https://communities.vmware.com/docs/DOC-27683.

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 need to add:
  - If you need to connect a logical switch, you must add a logical router
  - If you need to connect a universal logical switch, you must 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 the **Add** (✚) icon.

4 Select the type of logical router you wish 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 the cross-vCenter NSX environment. This option will be available only if you have assigned a primary NSX Manager, and are making changes from the primary NSX Manager.

  a If you select **Universal Logical (Distributed) Router**, you must also select whether to enable local egress.

5 Type a name for the device.

This name appears in your vCenter inventory. The name should be unique across all logical routers within a single tenant.

Optionally, you can also enter a hostname. This name appears in the CLI. If you do not specify the host name, the Edge ID, which gets created automatically, is displayed in the CLI.

Optionally, you can enter a description and tenant.

6 Deploy an edge appliance.

Deploy Edge Appliance is selected by default. 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.

You can deselect the edge appliance option if you require only static routes, and do not want to deploy an Edge appliance. You cannot add an Edge appliance to the logical router after the logical router has been created.

7 (Optional) Enable High Availability.

Enable High Availability is not selected by default. Select the Enable High Availability check box to enable and configure high availability. High availability is required if you are planning to do dynamic routing.

8 Type and re-type a password for the logical router.

The password must be 12-255 characters and must contain the following:

- At least one upper case letter
- At least one lower case letter
- At last one number
- At least one special character

9 (Optional) Enable SSH and set the log level.

By default, SSH is disabled. If you do not enable SSH, you can still access the logical router by opening the virtual appliance console. Enabling SSH here causes the SSH process to run on the logical router virtual appliance, but you will also need to 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.

By default, the log level is emergency.

For example:

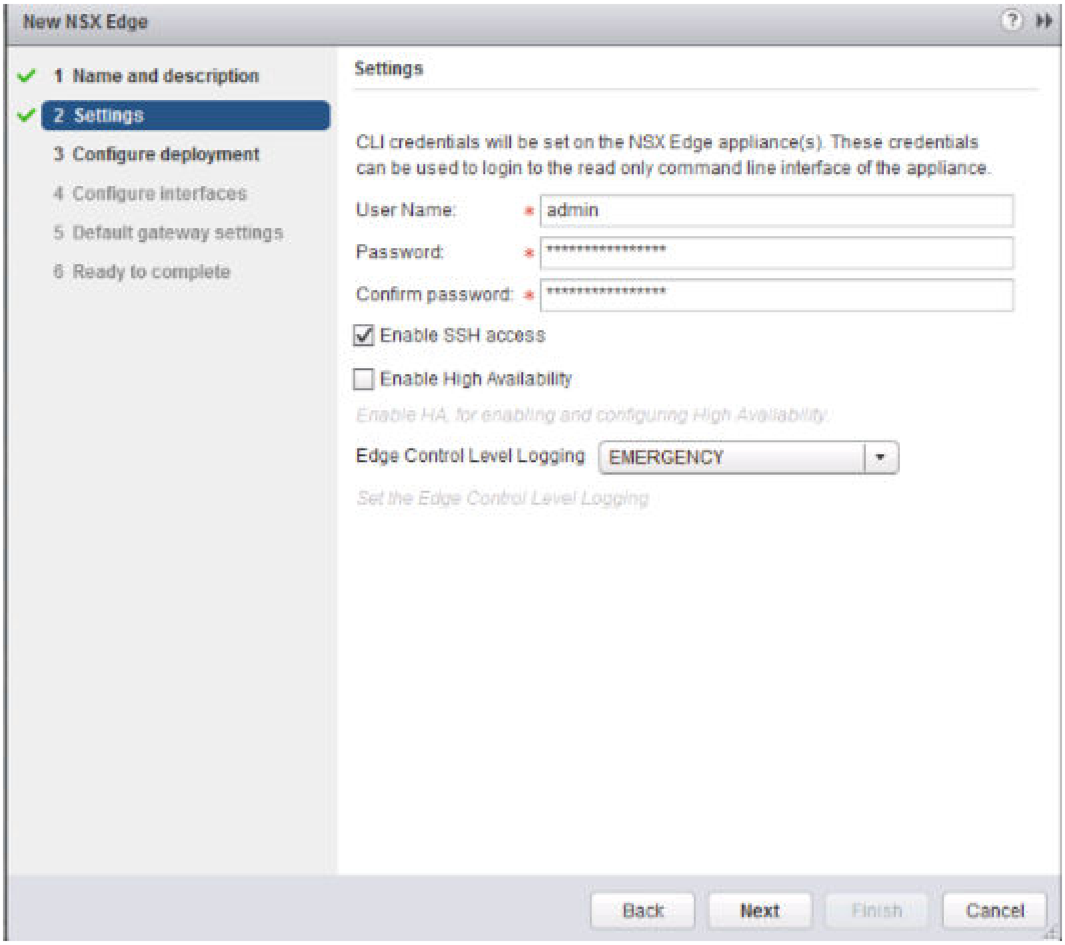

Settings

CLI credentials will be set on the NSX Edge appliance(s). These credentials can be used to login to the read only command line interface of the appliance.

User Name: * admin

Password: * ****************

Confirm password: * ****************

☑ Enable SSH access

□ Enable High Availability

Enable HA, for enabling and configuring High Availability

Edge Control Level Logging EMERGENCY

Set the Edge Control Level Logging

10 Configure deployment.

- If you did not select **Deploy NSX Edge**, the **Add (✚)** icon is grayed out. Click **Next** to continue with configuration.
- If you selected **Deploy NSX Edge**, enter the settings for the logical router virtual appliance that will be added to your vCenter inventory.

For example:

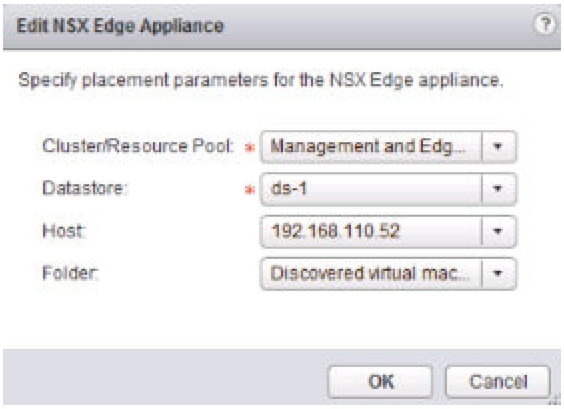

Edit NSX Edge Appliance

Specify placement parameters for the NSX Edge appliance.

Cluster/Resource Pool: * Management and Edg...

Datastore: * ds-1

Host: 192.168.110.52

Folder: Discovered virtual mac...

11 Configure interfaces.

On logical routers, only IPv4 addressing is supported.

In the HA Interface Configuration, if you selected **Deploy NSX Edge** you must connect the interface to a distributed port group. It is recommended to use a VXLAN logical switch for the HA interface. An IP address for each of the two NSX Edge appliances is chosen from the link local address space, 169.250.0.0/16. No further configuration is necessary to configure the HA service.

Note In prior releases of NSX, the HA interface was called the management interface. The HA interface is not supported for remote access to the logical router. You cannot SSH into the HA interface from anywhere that isn't 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. You could, in theory, disable RPF, but this would be counterproductive to high availability. For SSH, use the logical router's protocol address, which is configured later when you configure dynamic routing.

In NSX 6.2, the HA interface of a logical router is automatically excluded from route redistribution.

In **Configure interfaces of this NSX Edge** 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. A logical router uplink interface might connect to an NSX edge services gateway, a third-party router VM for that, or a VLAN-backed dvPortgroup to make the logical router connect to a physical router directly. You must have at least one uplink interface for dynamic routing to work. Uplink interfaces are created as vNICs on the logical router virtual appliance.

The interface configuration that you enter here is modifiable later. You can add, remove, and modify interfaces after a logical router is deployed.

The following example shows an HA interface connected to the management distributed port group. The example also shows two internal interfaces (app and web) and an uplink interface (to-ESG).

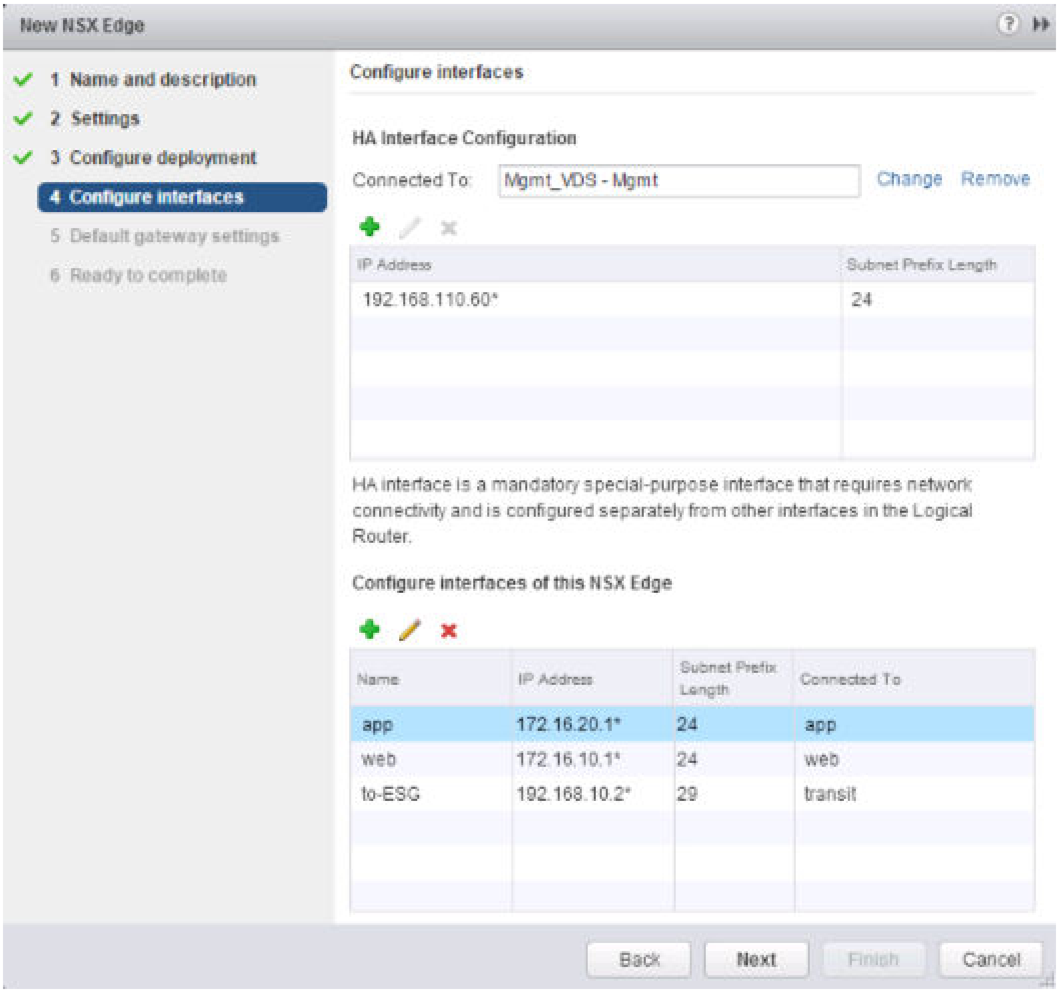

Configure interfaces

HA Interface Configuration

Connected To: Mgmt_VDS - Mgmt

IP Address	Subnet Prefix Length
192.168.110.60*	24

HA interface is a mandatory special-purpose interface that requires network connectivity and is configured separately from other interfaces in the Logical Router.

Configure interfaces of this NSX Edge

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

12 Configure a default gateway.

For example:

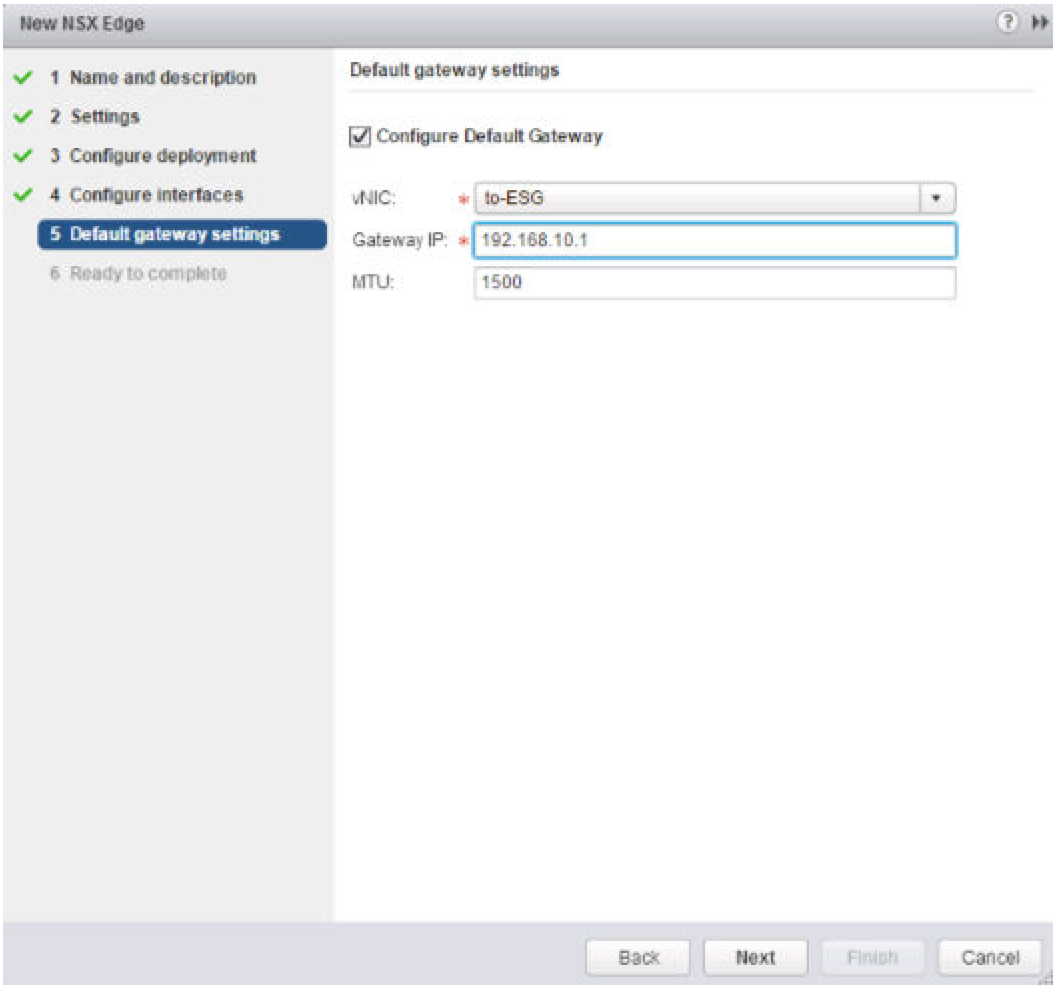

Default gateway settings

☑ Configure Default Gateway

vNIC: * to-ESG

Gateway IP: * 192.168.10.1

MTU: 1500

13 Make sure any VMs attached to the logical switches have their default gateways set properly to the logical router interface IP addresses.

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

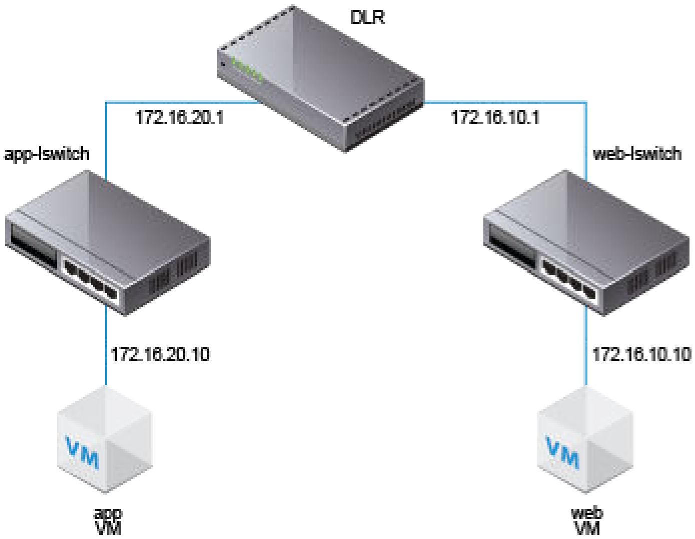

Log in via SSH to the NSX Manager, and run the following commands:

- List all logical router instance information.

```
nsxmgr-l-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-l-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 of the hosts.

```
nsx-mgr-l-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 is helpful to learn which controller is communicating with the host.

  ```
  nsx-mgr-l-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 would 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, start pinging from one VM to the other VM.

  From app VM to web VM:

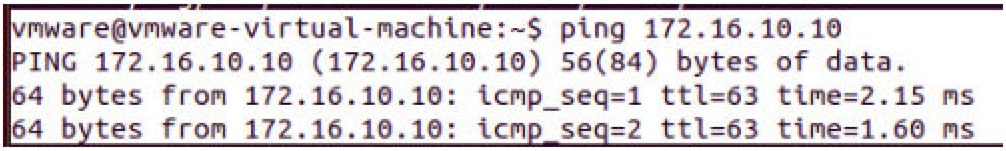

```
vmware@vmware-virtual-machine:~$ ping 172.16.10.10
PING 172.16.10.10 (172.16.10.10) 56(84) bytes of data.
64 bytes from 172.16.10.10: icmp_seq=1 ttl=63 time=2.15 ms
64 bytes from 172.16.10.10: icmp_seq=2 ttl=63 time=1.60 ms
```

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

The `interface-summary` sub-command 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` sub-command shows the routing table that is sent to this controller by the logical router's virtual appliance (also known as the control VM). Note that 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.

-
  ```
  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:~] #esxcli 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:~] # esxcli 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           0
    5001  N/A (headend replication)  Enabled (multicast proxy,ARP proxy)  192.168.110.202
(up)            1                0                0           0
```

Note To enable the vxlan namespace in vSphere 6 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 should contain 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 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

On the hosts where NSX edge appliances are first deployed, NSX enables automatic VM startup/shutdown. If the appliance VMs are later migrated to other hosts, the new hosts might not have automatic VM startup/shutdown enabled. For this reason, VMware recommends that you check all hosts in the cluster to make sure that automatic VM startup/shutdown is enabled. See http://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.vm_admin.doc%2FGUID-5FE08AC7-4486-438E-AF88-80D6C7928810.html.

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.

For example:

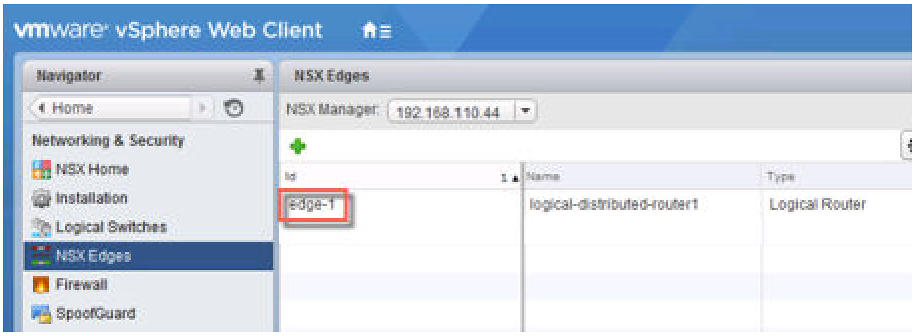

NSX Manager: 192.168.110.44

Id	Name	Type
edge-1	logical-distributed-router1	Logical Router

Add an Edge Services Gateway

You can install multiple NSX Edge services gateway virtual appliances in a data center. Each NSX Edge virtual 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 interface.
- DNS Forwarder: Not supported on uplink interface.
- HA: Not supported on uplink interface, requires at least one internal interface.
- SSL VPN: Listener IP must belong to uplink interface.
- IPsec VPN: Local site IP must belong to uplink interface.
- L2 VPN: Only internal networks can be stretched.

The following figure shows a sample topology with an ESG's uplink interface connected to physical infrastructure through the vSphere distributed switch and the ESG's internal interface connect to an NSX logical router through an NSX logical transit switch.

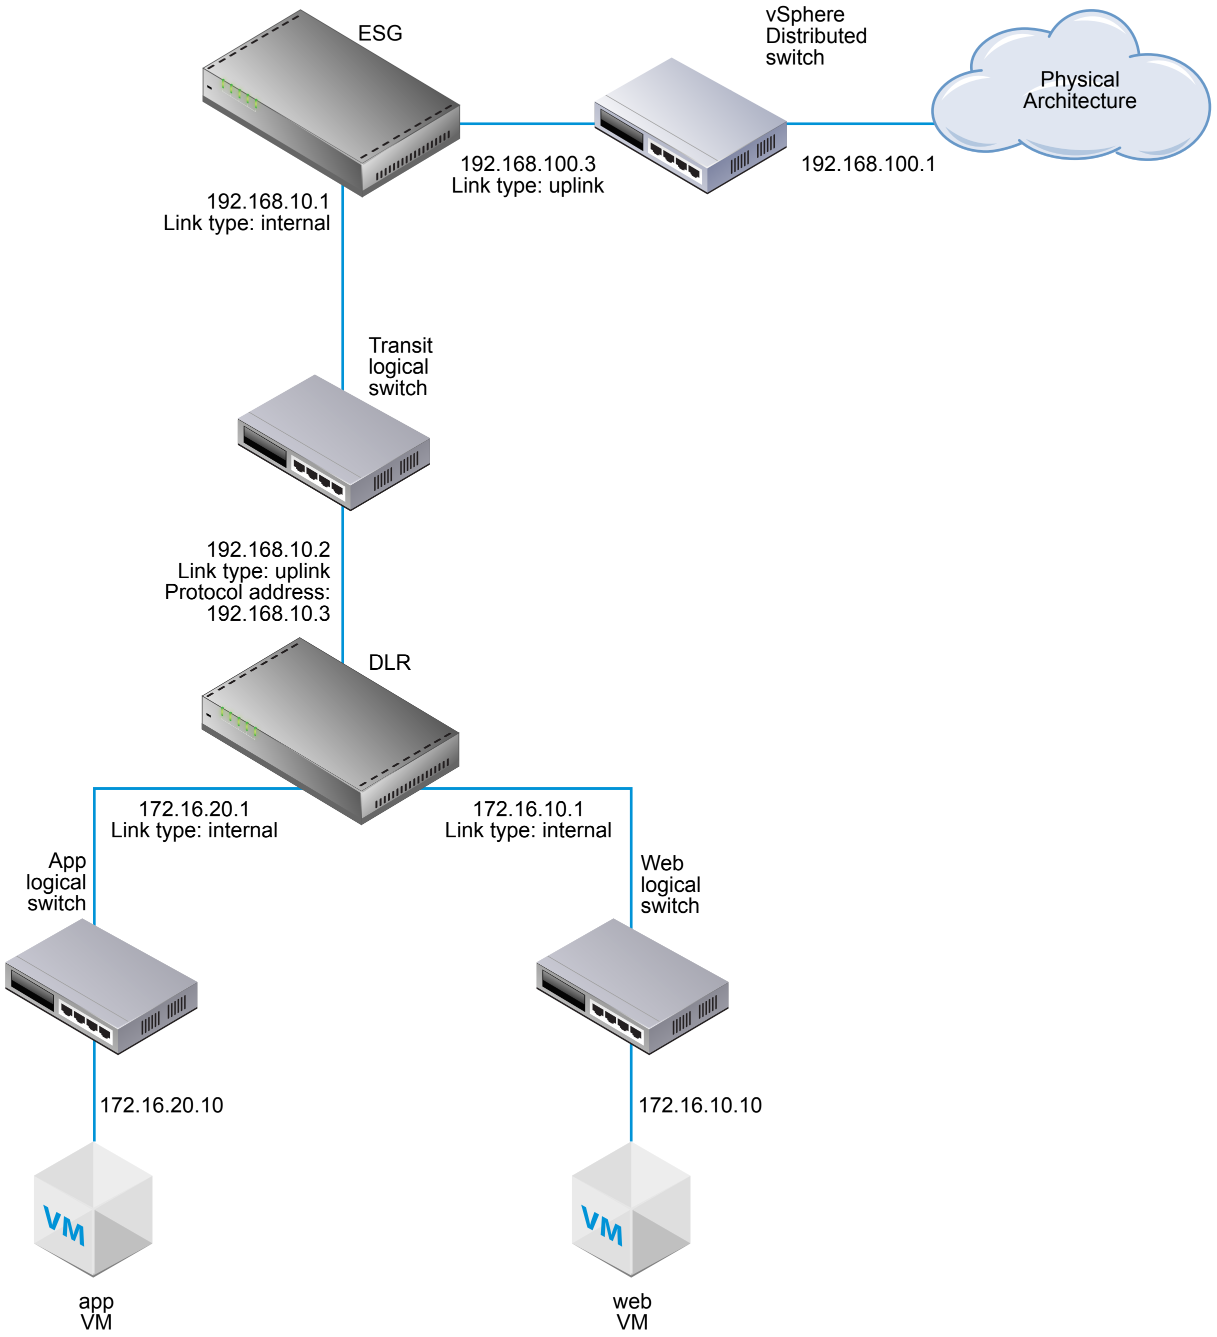

Multiple external IP addresses can be configured for load balancing, site-to-site VPN, and NAT services.

Prerequisites

You must have been 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.

Procedure

1 In vCenter, navigate to **Home > Networking & Security > NSX Edges** and click the **Add** (✚) icon.

2 Select **Edge Services Gateway** and type a name for the device.

This name appears in your vCenter inventory. The name should be unique across all ESGs within a single tenant.

Optionally, you can also enter a hostname. This name appears in the CLI. If you do not specify the host name, the Edge ID, which gets created automatically, is displayed in the CLI.

Optionally, you can enter a description and tenant and enable high availability.

For example:

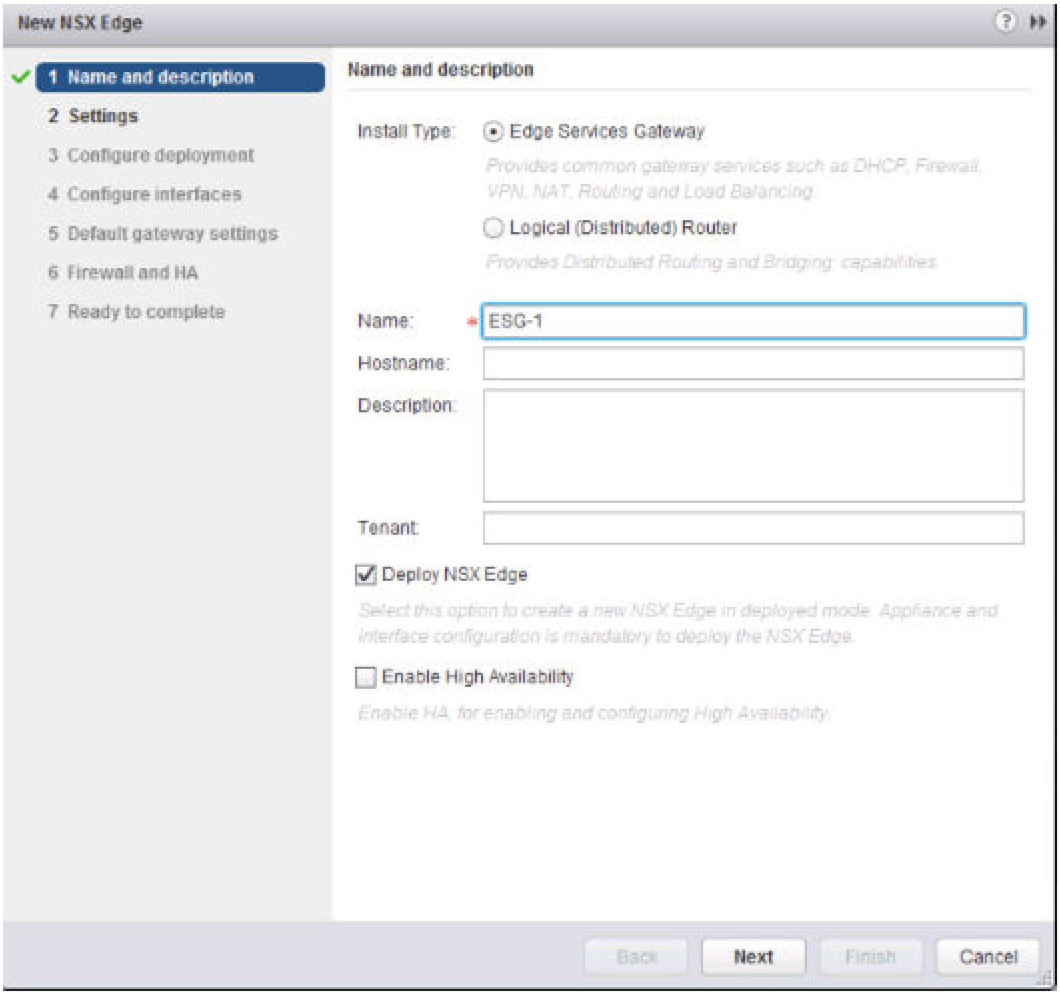

Name and description

Install Type: ● Edge Services Gateway

Provides common gateway services such as DHCP, Firewall, VPN, NAT, Routing and Load Balancing.

○ Logical (Distributed) Router

Provides Distributed Routing and Bridging capabilities.

Name: * ESG-1

Hostname:

Description:

Tenant:

☑ Deploy NSX Edge

Select this option to create a new NSX Edge in deployed mode. Appliance and interface configuration is mandatory to deploy the NSX Edge.

□ Enable High Availability

Enable HA, for enabling and configuring High Availability.

3 Type and re-type a password for the ESG.

The password must be at least 12 characters and must follow 3 of the following 4 rules:

- At least one upper case letter
- At least one lower case letter
- At last one number

- At least one special character

4 (Optional) Enable SSH, high availability, and automatic rule generation, and set the log level.

If you do not enable automatic rule generation, you must manually add firewall, NAT, and routing configuration to allow control traffic for certain, NSX Edge services, including as load balancing and VPN. Auto rule generation does not create rules for data-channel traffic.

By default, SSH and high availability are disabled, and automatic rule generation is enabled. By default, the log level is emergency.

By default, logging is enabled on all new NSX Edge appliances. The default logging level is NOTICE.

For example:

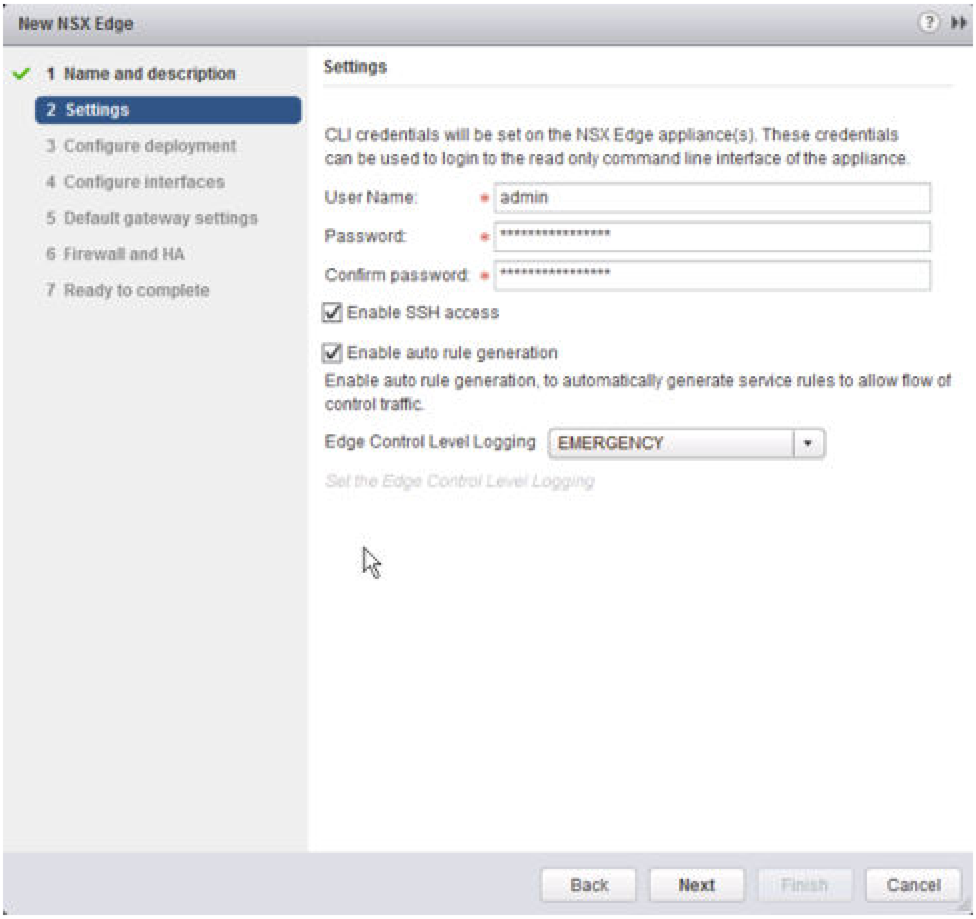

Settings

CLI credentials will be set on the NSX Edge appliance(s). These credentials can be used to login to the read only command line interface of the appliance.

User Name: * admin

Password: * ****************

Confirm password: * ****************

☑ Enable SSH access

☑ Enable auto rule generation

Enable auto rule generation, to automatically generate service rules to allow flow of control traffic.

Edge Control Level Logging EMERGENCY

Set the Edge Control Level Logging

5 Select the size of the NSX Edge instance based on your system resources.

The **Large** NSX Edge has more CPU, memory, and disk space than the **Compact** NSX Edge, and supports a larger number of concurrent SSL VPN-Plus users. The **X-Large** NSX Edge is suited for environments that have a load balancer with millions of concurrent sessions. The Quad Large NSX Edge is recommended for high throughput and requires a high connection rate.

See Chapter 1 System Requirements for NSX.

6 Create an edge appliance.

Enter the settings for the ESG virtual appliance that will be added to your vCenter inventory. If you do not add an appliance when you install NSX Edge, NSX Edge remains in an offline mode until you add an appliance.

If you enabled HA you can add two appliances. If you add a single appliance, NSX Edge replicates its configuration for the standby appliance and ensures that the two HA NSX Edge virtual machines are not on the same ESX host even after you use DRS and vMotion (unless you manually vMotion them to the same host). For HA to work correctly, you must deploy both appliances on a shared datastore.

For example:

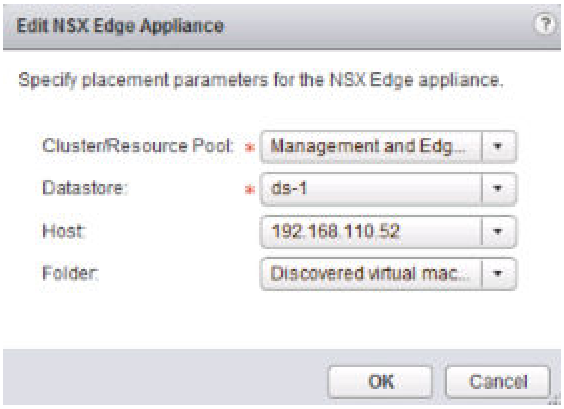
Edit NSX Edge Appliance

Specify placement parameters for the NSX Edge appliance.

Cluster/Resource Pool: * Management and Edg...

Datastore: * ds-1

Host: 192.168.110.52

Folder: Discovered virtual mac...

7 Select **Deploy NSX Edge** to add the Edge in a deployed mode. You must configure appliances and interfaces for the Edge before it can be deployed.

8 Configure interfaces.

On ESGs, both IPv4 and IPv6 addresses are supported.

You must add at least one internal interface for HA to work.

An interface can have multiple non-overlapping subnets.

If you enter more than one IP address for an interface, you can select the primary IP address. An interface can have one primary and multiple secondary IP addresses. NSX Edge considers the primary IP address as the source address for locally generated traffic, for example remote syslog and operator-initiated pings.

You must add an IP address to an interface before using it on any feature configuration.

Optionally, you can enter the MAC address for the interface.

If HA is enabled, you can optionally enter two management IP addresses in CIDR format. Heartbeats of the two NSX Edge HA virtual machines are communicated through these management IP addresses. The management IP addresses must be in the same L2/subnet and be able to communicate with each other.

Optionally, you can modify the MTU.

Enable proxy ARP if you want to allow the ESG to answer ARP requests intended for other machines. This is useful, for example, when you have the same subnet on both sides of a WAN connection.

Enable ICMP redirect to convey routing information to hosts.

Enable reverse path filtering to verify the reachability of the source address in packets being forwarded. In enabled mode, the packet must be received on the interface that the router would use to forward the return packet. In loose mode, the source address must appear in the routing table.

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 allow you to run several cloud instances simultaneous with the same IP and MAC addresses completely isolated or “fenced.”

For example:

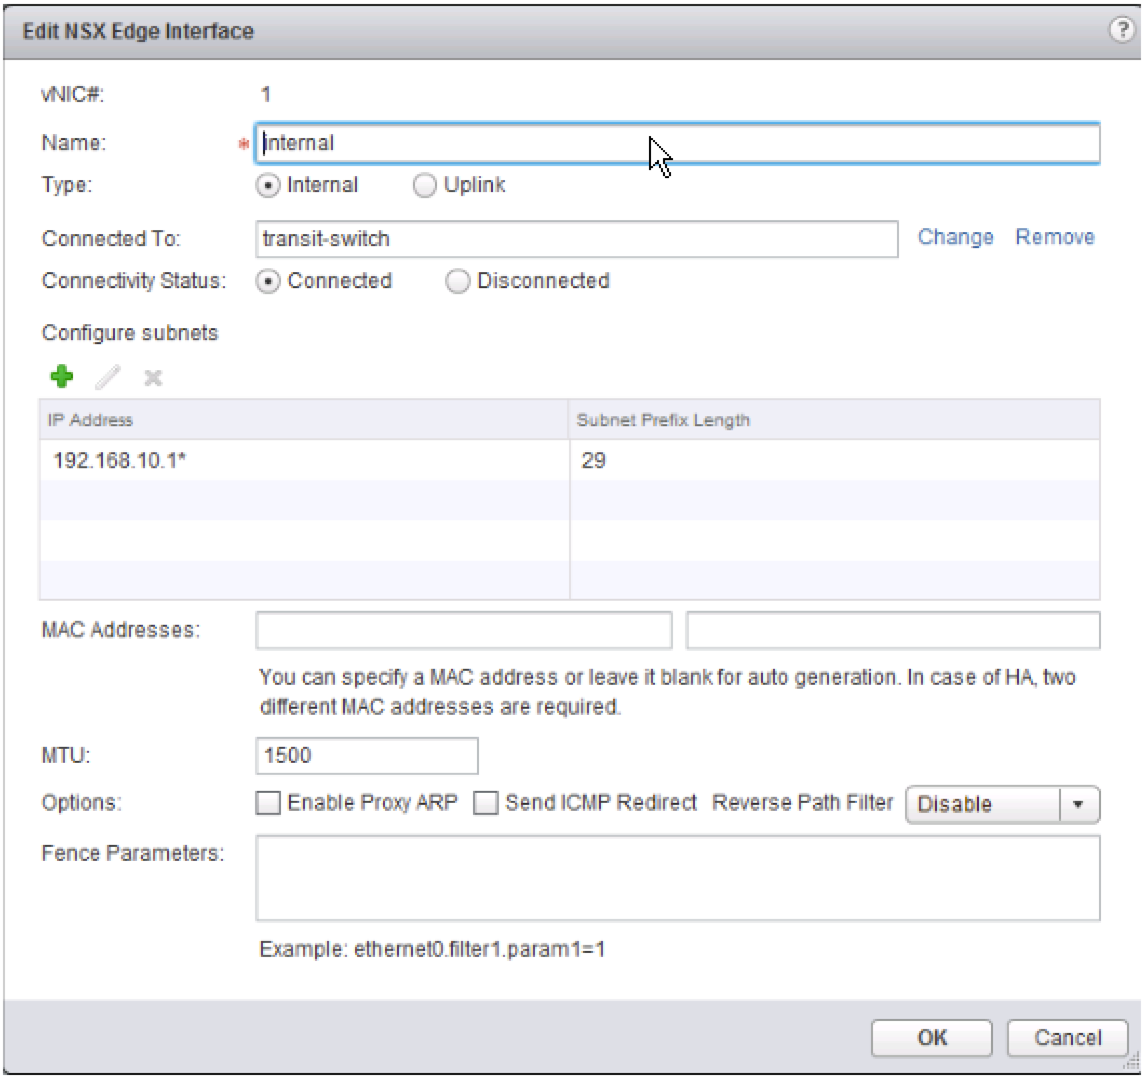

Edit NSX Edge Interface

vNIC#: 1

Name: * Internal

Type: ● Internal ○ Uplink

Connected To: transit-switch

Connectivity Status: ● Connected ○ Disconnected

Configure subnets

IP Address	Subnet Prefix Length
192.168.10.1*	29

MAC Addresses:

You can specify a MAC address or leave it blank for auto generation. In case of HA, two different MAC addresses are required.

MTU: 1500

Options: □ Enable Proxy ARP □ Send ICMP Redirect Reverse Path Filter: Disable

Fence Parameters:

Example: ethernet0.filter1.param1=1

The following example shows two interfaces, one attaching the ESG to the outside world through an uplink portgroup on a vSphere distributed switch and the other attaching the ESG to a logical transit switch to which a distributed logical router is also attached.

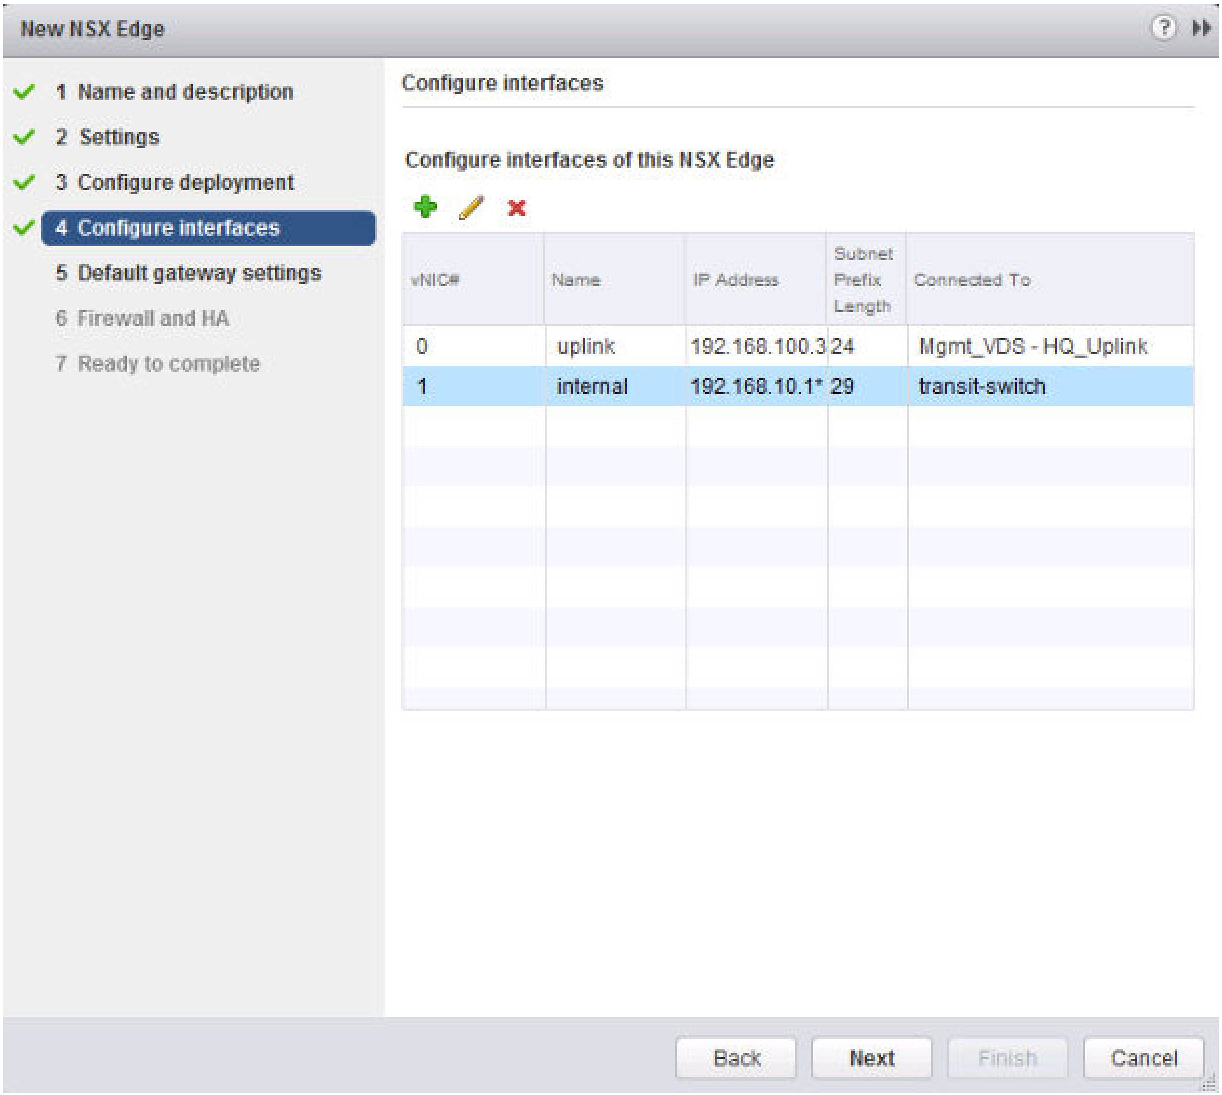

Configure interfaces

Configure interfaces of this NSX Edge

vNIC#	Name	IP Address	Subnet Prefix Length	Connected To
0	uplink	192.168.100.3	24	Mgmt_VDS - HQ_Uplink
1	internal	192.168.10.1*	29	transit-switch

9 Configure a default gateway.

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

For example:

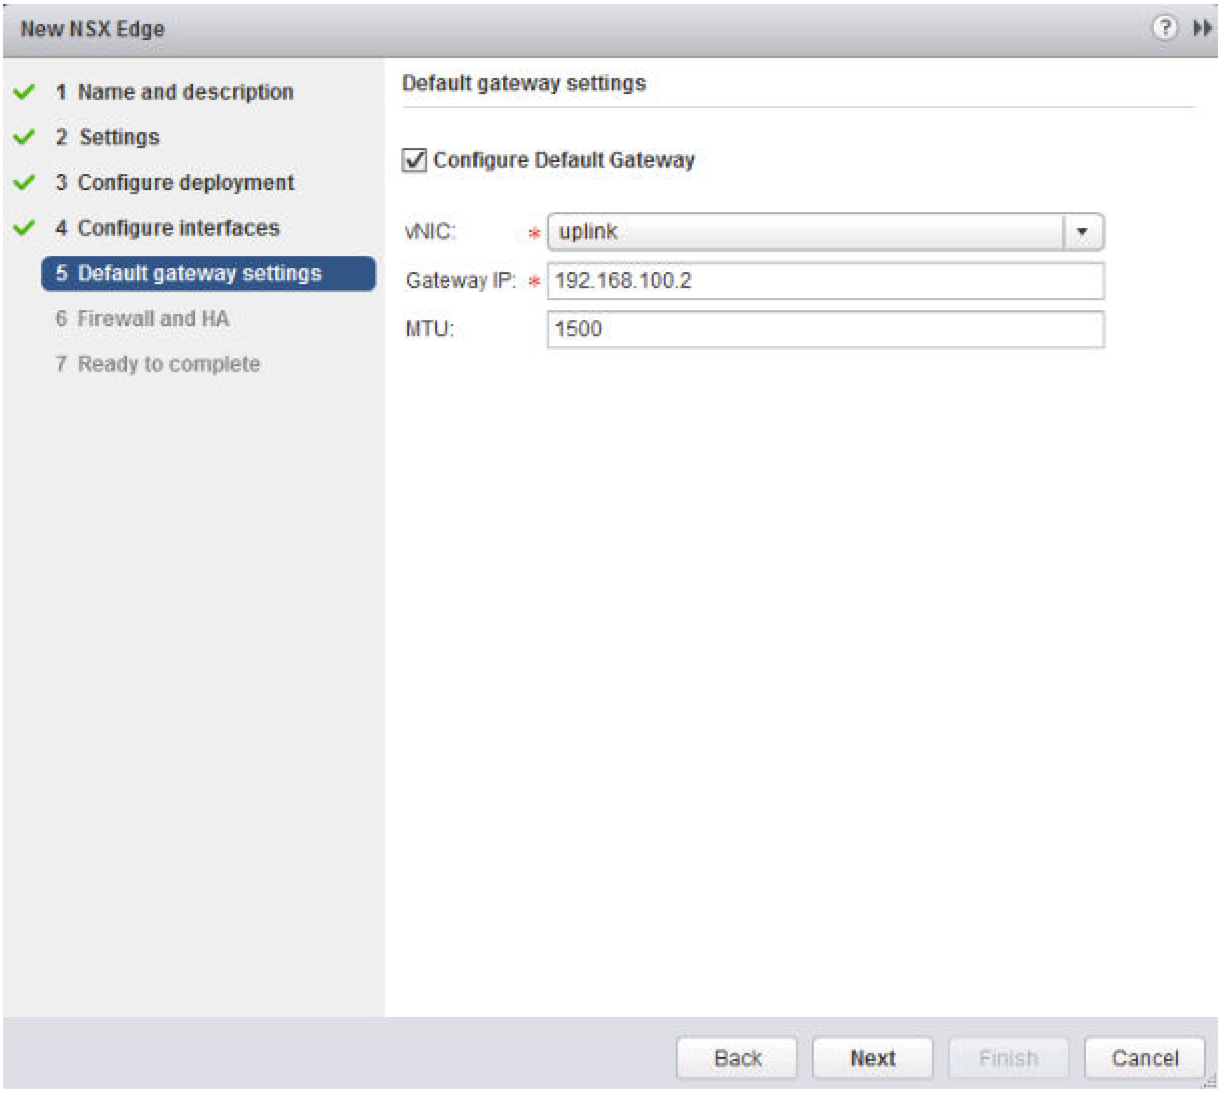

Default gateway settings

☑ Configure Default Gateway

vNIC: * uplink

Gateway IP: * 192.168.100.2

MTU: 1500

10 Configure the firewall policy, logging, and HA parameters.

Caution If you do not configure the firewall policy, the default policy is set to deny all traffic.

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

If you enabled high availability, complete the HA section. By default, HA automatically chooses an internal interface and automatically assigns link-local IP addresses. NSX Edge supports two virtual machines for high availability, both of which are kept up to date with user configurations. If a heartbeat failure occurs on the primary virtual machine, the secondary virtual machine state is changed to active. Thus, one NSX Edge virtual machine is always active on the network. NSX Edge replicates the configuration of the primary appliance for the standby appliance and ensures that the two HA NSX Edge virtual machines are not on the same ESX host even after you use DRS and vMotion. Two virtual machines are deployed on vCenter in the same resource pool and datastore as the appliance you configured. Local link IP addresses are assigned to HA virtual machines in the NSX Edge HA so that they can communicate with each other. Select the internal interface for which to configure HA parameters. If you select ANY for interface but there are no internal interfaces configured, the UI does not display an error. Two Edge appliances are created but since there is no internal interface configured, the new Edge remains in standby and HA is disabled. Once an internal interface is configured, HA will get enabled on the Edge appliance. Type 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. Optionally, 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. You should not use an IP address that exists somewhere else on your network, even if that network is not directly attached to the NSX Edge.

For example:

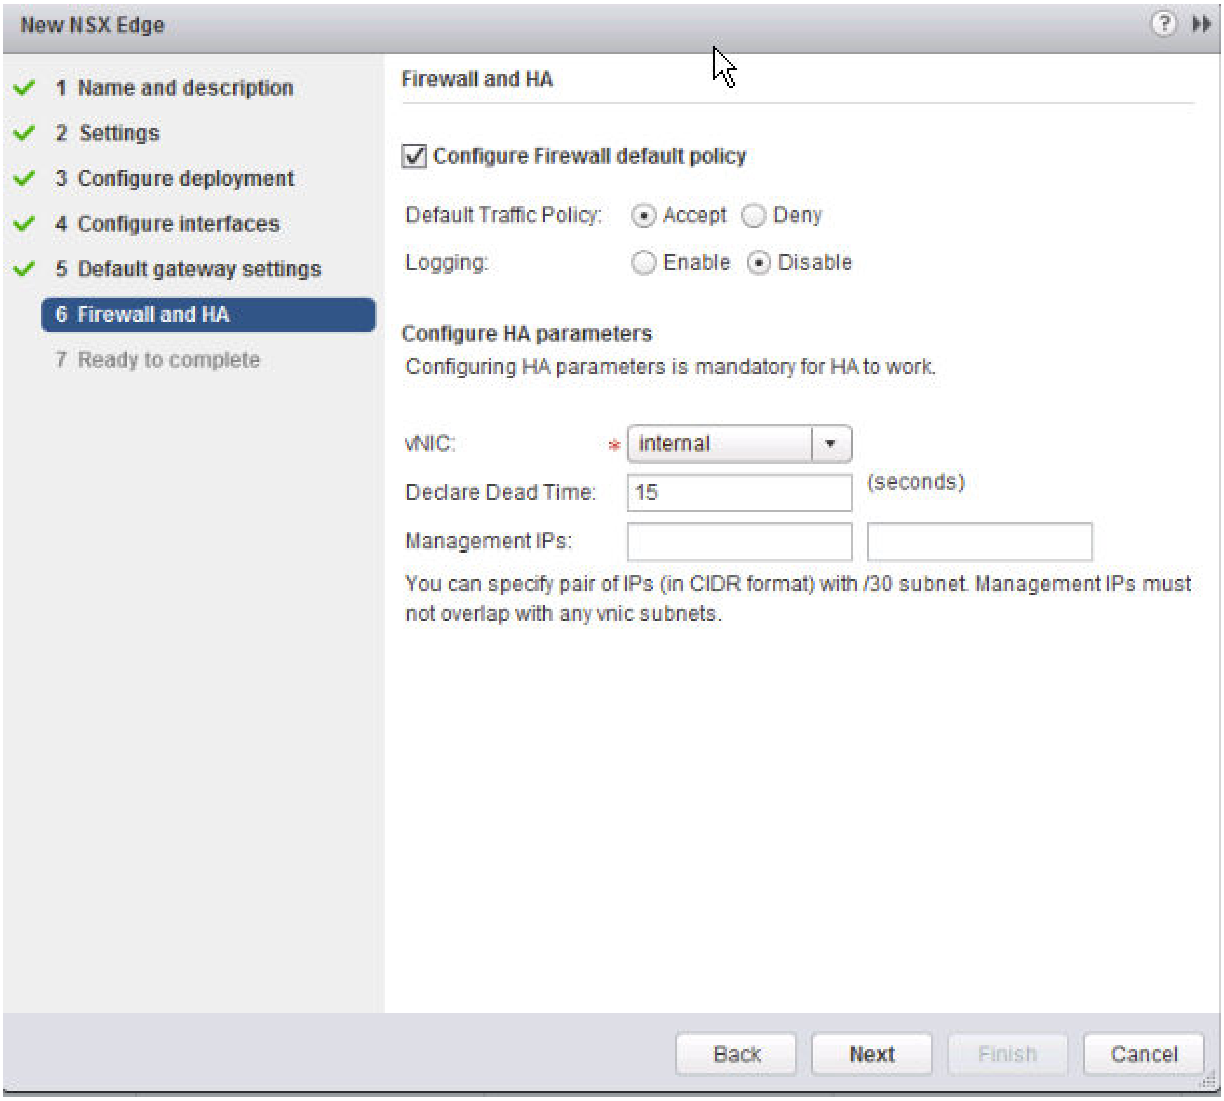

Firewall and HA

☑ Configure Firewall default policy

Default Traffic Policy: ● Accept ○ Deny

Logging: ○ Enable ● Disable

Configure HA parameters

Configuring HA parameters is mandatory for HA to work.

vNIC: * internal

Declare Dead Time: 15 (seconds)

Management IPs:

You can specify pair of IPs (in CIDR format) with /30 subnet. Management IPs must not overlap with any vnic subnets.

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

What to do next

On the hosts where NSX edge appliances are first deployed, NSX enables automatic VM startup/shutdown. If the appliance VMs are later migrated to other hosts, the new hosts might not have automatic VM startup/shutdown enabled. For this reason, VMware recommends that you check all hosts in the cluster to make sure that automatic VM startup/shutdown is enabled. See http://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.vm_admin.doc%2FGUID-5FE08AC7-4486-438E-AF88-80D6C7928810.html.

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 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** and then click **NSX Edges**.

3 Double-click an NSX Edge.

4 Click **Routing** and then click **Global Configuration**.

5 To change Equal-cost multi-path routing (ECMP) configuration click **Edit** next to **Routing Configuration**, then do the following .

Option	Description
For an Edge Services Gateway	To edit ECMP, click **Enable** or **Disable** next to ECMP.
For a Logical Router	a Select ECMP to enable or deselect to disable. b Click OK.

ECMP is a routing strategy that allows next-hop packet forwarding to a single destination can occur over multiple best paths. These best paths can be added statically 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 utilizes 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 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.

Until version 6.1.2, enabling ECMP disabled Distributed Firewall on the Edge Services Gateway virtual machine. Stateful services such as NAT did not work with ECMP. From NSX vSphere version 6.1.3 onwards, ECMP and Distributed Firewall can work together.

6 To change the **Locale ID** on a logical router, click **Edit** next to **Routing Configuration**. Enter a locale ID and click OK.

By default, the locale ID is set to the NSX Manager UUID, but you can override it if local egress was enabled when the universal logical router was created. 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 an option 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 9-1. 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 CSR and get it signed by a 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.

Option	Description
To generate a global certificate	a Log in to the NSX Manager Virtual Appliance. b Click the Manage tab and then click SSL Certificates. c Click **Generate CSR**.
To generate a certificate for an NSX Edge	a Log in to the vSphere Web Client. b Click **Networking & Security** and then click **Edge Services**. c Double-click an NSX Edge. d Click the **Manage** tab and then click **Settings**. e Click the **Certificates** link. f Click **Actions** and select **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.

   Note that SSL VPN-Plus only supports RSA certificates.
5. Edit the default key size if required.
6. For a global certificate, 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. Import the signed certificate.

   a Copy the contents of the signed certificate.

   b Do one of the following.

   - To import a signed certificate at the global level, click **Import** in the NSX Manager Virtual Appliance.
   - To import a signed certificate for an NSX Edge, click **Actions** and select **Import Certificate** in the **Certificates** tab.

   c In the Import CSR dialog box, paste the contents of the signed certificate.

   d 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** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click the **Manage** tab and then ensure that you are in the **Settings** tab.
5. Click **Certificates**.
6. Click the **Add** (✚) icon and select **CA Certificate.**
7. Copy and paste the certificate contents in the Certificate contents text box.
8. Type a description for the CA certificate.
9. Click **OK**.

   You can now sign your own certificates.

Configure a Self-Signed Certificate

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

Prerequisites

Verify that you have a CA certificate so that you can sign your own certificates.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click the **Manage** tab and then ensure that you are in the **Settings** tab.
5. Click **Certificates**.
6. Follow the steps below to generate a CSR.

    a. Click **Actions** and select **Generate CSR**.

    b. In Common name, type the IP address or fully qualified domain name (FQDN) of the NSX Manager.

    c. Type your organization name and unit.

    d. Type the locality, street, state, and country of your organization.

    e. Select the encryption algorithm for communication between the hosts.

       Note that SSL VPN-Plus only supports RSA certificates. VMware recommends RSA for backward compatibility.

    f. Edit the default key size if required.

    g. Type a description for the certificate.

    h. Click **OK**.

    The CSR is generated and displayed in the Certificates list.
7. Verify that the certificate you generated is selected.
8. Click **Actions** and select **Self Sign Certificate**.
9. Type the number of days the self sign certificate is valid for.
10. Click **OK**.

Using Client Certificates

You can create a client certificate through a CAI command or REST call. You can then 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 a reference client certificate for each remote user can be stored and checked against the client certificate presented by the remote user. To prevent future connections from a certain user, you can delete the reference certificate from the security server's list of client certificates. Deleting the certificate denies connections from that user.

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** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click the **Manage** tab and then ensure that you are in the **Settings** tab.
5. Click **Certificates**.
6. Click the **Add** (✚) icon and select **CRL**.
7. In **Certificate contents**, paste the list.
8. (Optional) Type a description.
9. Click **OK**.

Managing Appliances

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

Add an 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** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click the **Manage** tab and then click the **Settings** tab.
5. In **Edge Gateway Appliances**, click the **Add** (✚) icon.
6. Select the cluster or resource pool and datastore for the appliance.
7. (Optional) Select the host on which the appliance is to be added.
8. (Optional) Select the vCenter folder within which the appliance is to be added.
9. Click **Add**.

Edit an Appliance

You can edit a NSX Edge appliance.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click the **Manage** tab and then click the **Settings** tab.
5. In **Edge Gateway Appliances**, select the appliance to change.
6. Click the **Edit** (✎) icon.
7. In the Edit Edge Appliance dialog box, make the appropriate changes.
8. Click **Save**.

Delete an Appliance

You can delete an NSX Edge appliance.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click the **Manage** tab and then click the **Settings** tab.
5. In **Edge Gateway Appliances**, select the appliance to delete.
6. Click the **Delete** (✖) icon.

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. 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 edge services gateway (ESG) can have up to ten internal, uplink, or trunk interfaces. These limits are enforced by the NSX Manager.

An NSX deployment can have up to 1,000 distributed logical router (DLR) instances on a single ESXi host. On a single logical router, you can configure up to 8 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 deployment, see the *VMware® NSX for vSphere Network Virtualization Design Guide* at https://communities.vmware.com/docs/DOC-27683.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click the **Manage** tab and then click the **Interfaces** tab.
5. Select an interface and click the **Edit** (✏) icon.
6. In the Edit Edge Interface dialog box, type a name for the interface.
7. Select **Internal** or **Uplink** to indicate whether this is an internal or external interface.

   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 this interface should be connected.

   a. Click **Select** next to the **Connected To** field.

   b. Depending on what you want to connect to the interface, click the **Logical Switch**, **Standard Portgroup**, or **Distributed Portgroup** tab.

   c. Select the appropriate logical switch or portgroup.

   d. Click **Select**.
9. Select the connectivity status for the interface.

10. In **Configure Subnets**, click the **Add** (✚) icon to add a subnet for the interface.

    An interface can have multiple non-overlapping subnets.

11. In **Add Subnet**, click the **Add** (✚) icon to add an IP address.

    If you enter more than one IP address, you can select the Primary IP address. An interface can have one primary and 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.

12. Type the subnet mask for the interface and click **Save**.
13. Change the default MTU if required.
14. In **Options**, select the required options.

Option	Description
Enable 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 would use to forward the return packet. In loose mode, the source address must appear in the routing table.

15. Type the fence parameters and click **Add**.
16. Click **OK**.

Delete an Interface

You can delete an NSX Edge interface.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click the **Manage** tab and then click the **Interfaces** tab.
5. Select the interface to delete.
6. Click the **Delete** (✖) icon

Enable an Interface

An interface must be enabled for 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** and then click **NSX Edges**.

3 Double-click an NSX Edge.

4 Click the **Manage** tab and then click the **Interfaces** tab.

5 Select the interface to enable.

6 Click the **Enable** (✓) icon.

Disable an Interface

You can disable an interface on an NSX Edge.

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **NSX Edges**.

3 Double-click an NSX Edge.

4 Click the **Manage** tab and then click the **Interfaces** tab.

5 Select the interface to disable.

6 Click the **Disable** icon.

Change Traffic Shaping Policy

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

Procedure

1 Double-click an NSX Edge and navigate to **Manage > Settings > Interfaces**.

2 Select an interface.

3 Click **Actions > Configure Traffic Shaping Policy**.

4 Make appropriate changes.

For more information on the options, see Traffic Shaping Policy.

5 Click **OK**.

Add a Sub Interface

You can add a sub interface on a trunk vNIC, which can then be used by NSX Edge services.

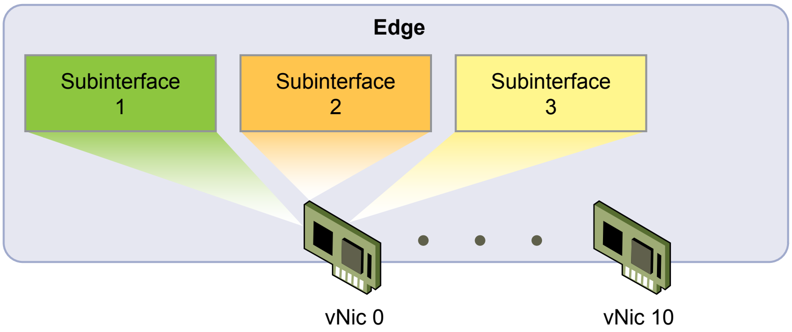

Trunk interfaces can be of the following types:

- VLAN trunk is standard and work with any version of ESXi. This is used to bring tagged VLAN traffic into Edge.
- VXLAN trunk work only with NSX version 6.1. This is used to bring VXLAN traffic into Edge.

A sub interface can be used by the following Edge services:

- DHCP
- Routing (BGP only)
- Load Balancer
- IPSEC VPN
- L2 VPN

A sub interface cannot be used for HA or Logical Firewall. You can, however, use the IP address of the sub interface in a firewall rule.

Procedure

1. In the **Manage > Settings** tab for an NSX Edge, click **Interfaces**.
2. Select an interface and click the **Edit** (✎) icon.
3. In the Edit Edge Interface dialog box, type a name for the interface.
4. In Type, select **Trunk**.
5. Select the standard portgroup or distributed portgroup to which this interface should be connected.
   a. Click **Change** next to the **Connected To** field.
   b. Depending on what you want to connect to the interface, click the **Standard Portgroup** or **Distributed Portgroup** tab.
   c. Select the appropriate portgroup and click **OK**.
   d. Click **Select**.
6. In Sub Interfaces, click the **Add** icon.
7. Click **Enable Sub interface** and type a name for the sub interface.

8. In **Tunnel Id**, type a number between 1 and 4094.

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

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

   - **VLAN** for a VLAN network.

     Type the VLAN ID of the virtual LAN that your sub interface should use. VLAN IDs can range from 0 to 4094.

   - **Network** for a VLAN or VXLAN network.

     Click **Select** and select the distributed portgroup or logical switch. NSX Manager extracts the VLAN ID and uses it in trunk configuration.

   - **None** to create a sub interface without specifying a network or VLAN ID. This sub interface is internal to NSX Edge, and is used to route packets between a stretched network and an unstretched (untagged) network

10. To add subnets to the sub interface, click the **Add** icon in the Configure Subnets area.

11. In Add Subnets, click the **Add** icon to add an IP address. Type the IP address and click **OK**.

    If you enter more than one IP address, you can select the Primary IP address. An interface can have one primary and multiple secondary IP addresses. NSX Edge considers the Primary IP address as the source address for locally generated traffic.

12. Type the subnet prefix length and click **OK**.

13. Edit the default **MTU** value for the sub interface if required.

    The default MTU for a trunk interface is 1600 and 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.

14. Select **Enable Send Redirect** to convey routing information to hosts.

15. Type the MAC address for the interface.

    Since sub interfaces do not support HA, only one MAC address is required.

16. Edit the default MTU of the trunk interface, if required.

17. Click **OK**.

You can now use the sub-interface on Edge services.

What to do next

Configure VLAN trunk if the sub interface added to a trunk vNic is backed by standard portgroup. See Configure VLAN Trunk.

Configure VLAN Trunk

When the sub interface is added to a trunk vNic backed by distributed portgroup, VLAN or VXLAN trunk is automatically configured on the trunk port. When the sub interface is added to a trunk vNic backed by 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** and then click **NSX Edges**.

3 Double-click an NSX Edge.

4 Click the **Monitor** tab and then click the **Settings** tab.

5 Click the **More Actions** (⚙) icon and select **Change Auto Rule configuration**.

6 Make the appropriate changes and click **OK**.

Change CLI Credentials

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

Procedure

1 Log in to the vSphere Web Client.

2. Click **Networking & Security** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click the **Monitor** tab and then click the **Settings** tab.
5. Click the **More Actions** (⚙) icon and select **Change CLI Credentials**.
6. Make the appropriate edits.
7. 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 statefull firewall, or the statefull 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 Edgee 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 Edge replicates the configuration of the primary appliance for the standby appliance or you can manually add two appliances. VMware recommends that you 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 ESX 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, so there is no service disruption during switch 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 ESX 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.

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 thereby ensuring the NSX Edge HA pair is still available to take another failover.

If vSphere HA is not leveraged, 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.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click the **Manage** tab and then click the **Settings** tab.
5. In the **HA Configuration** panel, click **Change**.
6. In the Change HA Configuration dialog box, make changes as appropriate.

   Note If L2 VPN is configured on this Edge appliance before HA is enabled, there must be 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.

7. Click **OK**.

Force Sync NSX Edge with NSX Manager

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

Force sync is used when you need to synchronize the edge configuration as known to the NSX manager to all of the components.

Note For 6.2 and later, force sync avoids data loss for east-west routing traffic, however, north-south routing and bridging may experience an interruption.

Force sync results in the following actions:

- Edge appliances are rebooted, and the latest configuration is applied
- The connection to the host is closed
- If the NSX manager is primary or stand alone, and the edge is a logical distributed router, then the controller cluster is synced
- A message is sent to all relevant hosts to sync the distributed router instance

Important In a cross-vCenter NSX environment, it is required that theNSX Edge instance be force synced first on the primary NSX manager and after that is complete, force sync theNSX Edge instance on secondary NSX managers.

Procedure

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

Configure Remote Syslog Servers

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** and then click **NSX Edges**.
3. Double-click a NSX Edge.
4. Click the **Monitor** tab and then click the **Settings** tab.
5. In the **Details** panel, click **Change** next to Syslog servers.
6. Type the IP address of both remote syslog servers and select the protocol.
7. Click **OK** to save the configuration.

View the Status of an NSX Edge

The status page displays graphs for the traffic flowing through the interfaces of the selected NSX Edge and connection statistics for the firewall and load balancer services.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click the **Monitor** tab.
5. Select the period for which you want to view the statistics.

What to do next

To view more details about NSX Edge, click **Manage** and then click **Settings**.

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. It is recommended that you first apply a force sync and if the issue is 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.

Important In a cross-vCenter environment it is required that theNSX Edge instance be redeployed first on the primary NSX manager and after that is complete, then redeploy the NSX Edge instance on secondary NSX managers. It is required that both the primary and the secondary NSX managers are redeployed.

Prerequisites

Verify the hosts have enough resources to deploy additional NSX Edge Services Gateway appliances during the redeploy. See the Chapter 1 System Requirements for NSX for the resources required for each NSX Edge size.

- For a single NSX Edge instance, there will be two NSX Edge appliances of the appropriate size in the poweredOn state during redeploy.
- Starting in NSX 6.2.3, when redeploying an NSX Edge instance with HA, both replacement appliances are deployed before replacing the old appliances. This means there will be four NSX Edge appliances of the appropriate size in the poweredOn state during upgrade of a given NSX Edge. Once the NSX Edge instance is redeployed, either of the HA appliances could become active.
- Prior to NSX 6.2.3, when redeploying an NSX Edge instance with HA , only one replacement appliance is deployed at time while replacing the old appliances. This means there will be three NSX Edge appliances of the appropriate size in the poweredOn state during the redeploy of a given NSX Edge. Once the NSX Edge instance is redeployed, usually the NSX Edge appliance with HA index 0 becomes active.

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **NSX Edges**.

3 Select an NSX Edge instance.

4 Click the **Actions** (⚙) icon and select **Redeploy Edge**.

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 may not work in the following cases.

- Resource pool on which the NSX Edge was installed is no longer in the vCenter inventory or its Managed Object ID (MoId) has changed.
- Datastore on which the NSX Edge was installed is corrupted/unmounted or in-accessible.
- 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 the above is true, you must update the MoId of the resource pool, datastore, or dvPortGroup using a REST API call. See *NSX API Programming Guide*.

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.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **NSX Edges**.
3. Select an NSX Edge instance.
4. Click the **More Actions** (⚙) icon and select **Download Tech Support Logs**.
5. After the tech support logs are generated, click **Download**.
6. In the Select location for download dialog box, browse to the directory where you want to save the log file.
7. Click **Save**.
8. Click **Close**.

Add a Static Route

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

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click the **Manage** tab and then click the **Routing** tab.
5. Select **Static Routes** from the left panel.

6. Click the **Add** (➕) icon.
7. Type the **Network** in CIDR notation.
8. Type the IP address of the **Next Hop**.

   The router must be able to directly reach the next hop.

   If ECMP is enabled, you can type multiple next hops.
9. Select the **Interface** on which you want to add a static route.
10. For **MTU**, edit the maximum transmission value for the data packets if required.

    The MTU cannot be higher than the MTU set on the NSX Edge interface.
11. 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 9-2. 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.
12. (Optional) Type the **Locale ID**.

    By default, routes have the same locale ID as the NSX Manager. Specifying a locale ID here will associate the route with this locale ID. These routes will be sent only to hosts that have a matching locale ID. See Cross-vCenter NSX Topologies for more information.
13. (Optional) Type a **Description** for the static route.
14. 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 Example: OSPF Configured on the Logical (Distributed) Router.

When you enable a router ID, the field is populated by default with the logical router's uplink interface.

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **NSX Edges**.

3 Double-click a logical router.

4 Click **Routing** and then click **OSPF**.

5 Enable OSPF.

   a Click **Edit** at the top right corner of the window and 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.

6 Configure the OSPF areas.

   a Optionally, delete the not-so-stubby area (NSSA) 51 that is configured by default.

   b In **Area Definitions**, click the **Add** icon.

   c Type an Area ID. NSX Edge supports an area ID in the form of an IP address or decimal number.

   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.

8 Map interfaces to the areas.

   a In **Area to Interface Mapping**, click the **Add** icon 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) If needed, 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 Topology

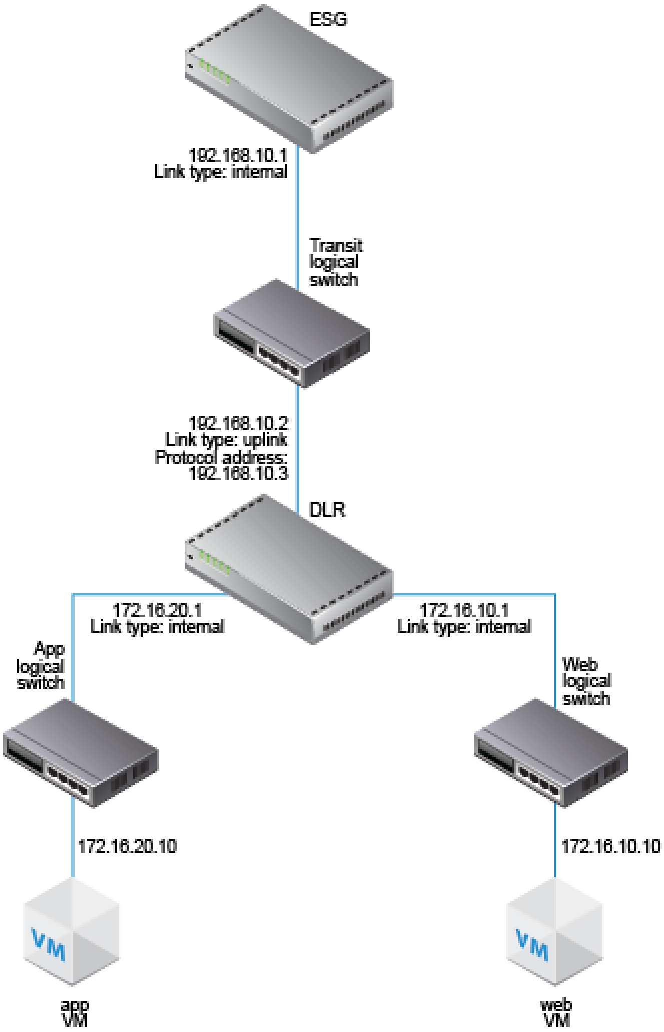

In the following screen, the logical router's default gateway is the ESG's internal interface IP address (192.168.10.1).

The router ID is the logical router's uplink interface---in other words, the IP address that faces the ESG (192.168.10.2).

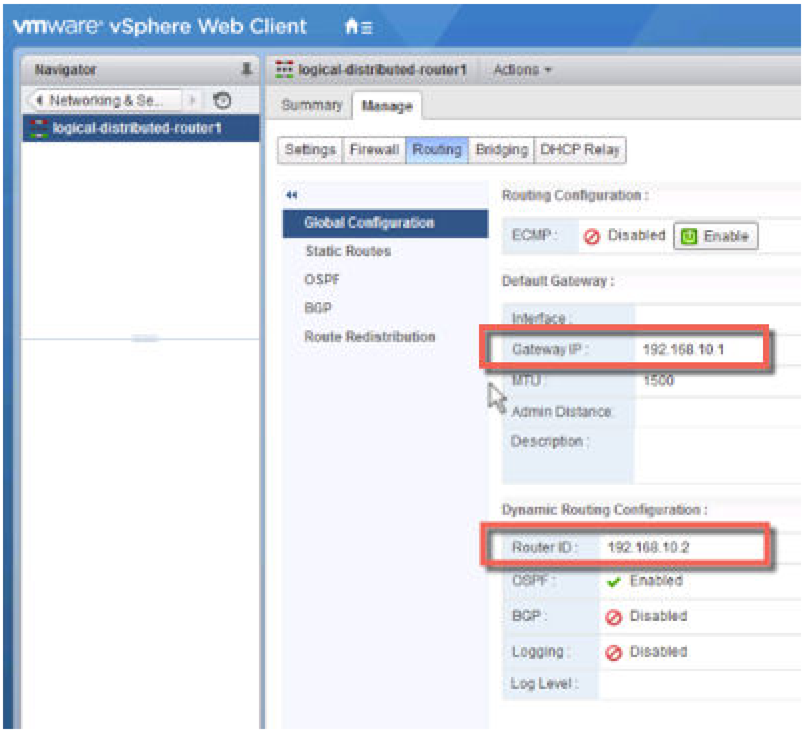

Routing Configuration :

ECMP : Disabled

Default Gateway :

Interface :

Gateway IP : 192.168.10.1

MTU : 1500

Admin Distance :

Description :

Dynamic Routing Configuration :

Router ID : 192.168.10.2

OSPF : ✓ Enabled

BGP : Disabled

Logging : Disabled

Log Level :

The logical router configuration uses 192.168.10.2 as its forwarding address. 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. The area ID configured is 0, and the uplink interface (the interface facing the ESG) is mapped to the area.

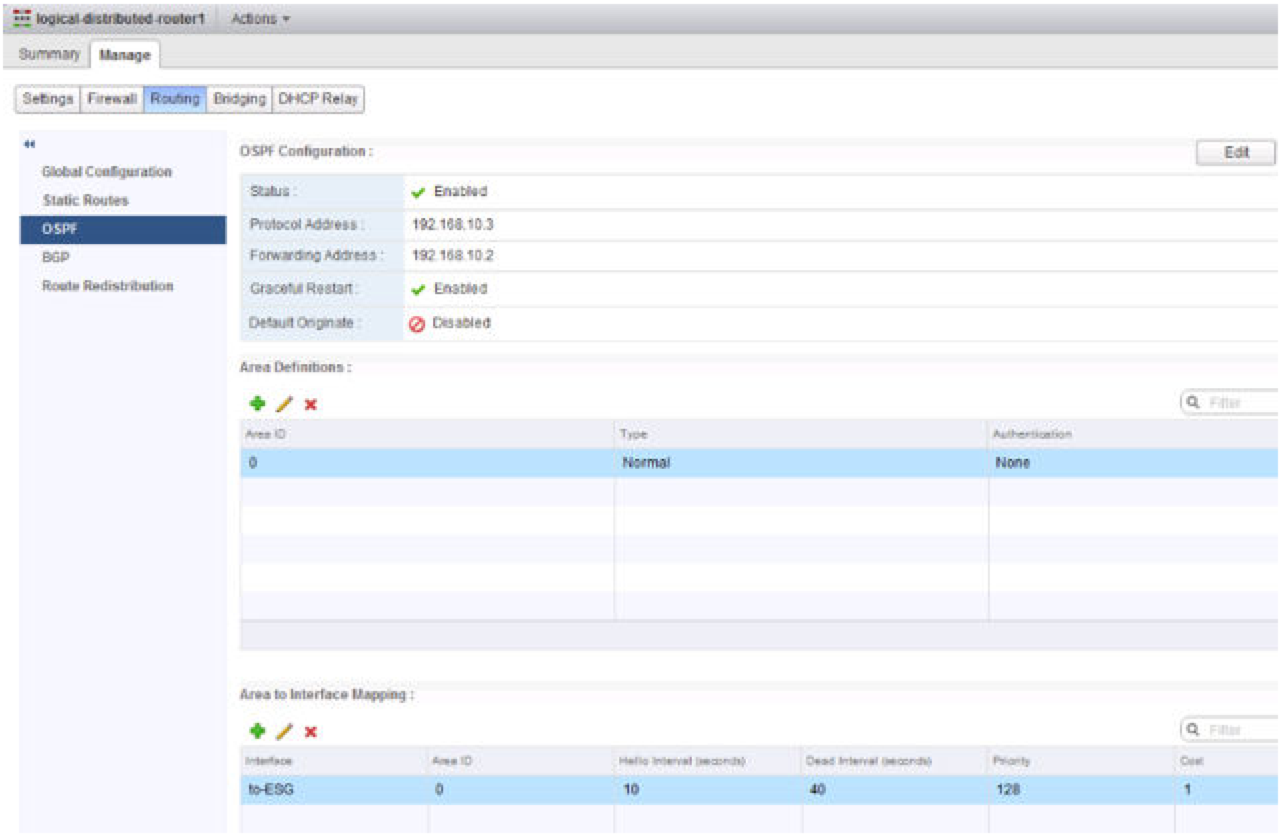

OSPF Configuration :

Status :	✓ Enabled
Protocol Address :	192.168.10.3
Forwarding Address :	192.168.10.2
Graceful Restart :	✓ Enabled
Default Originate :	⊘ Disabled

Area Definitions :

Area ID	Type	Authentication
0	Normal	None

Area to Interface Mapping :

Interface	Area ID	Hello Interval (seconds)	Dead Interval (seconds)	Priority	Cost
to-ESG	0	10	40	128	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.

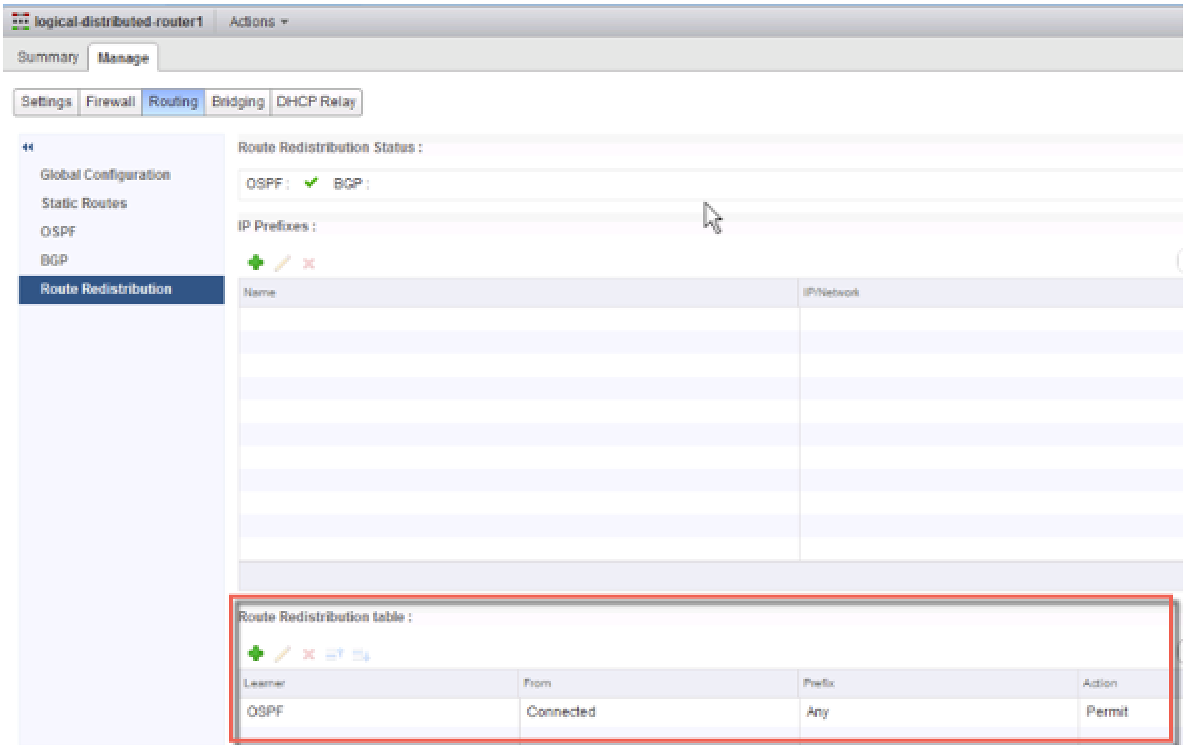

Route Redistribution Status :

OSPF : ✓ BGP :

IP Prefixes :

Name	IP/Network

Route Redistribution table :

Learner	From	Prefix	Action
OSPF	Connected	Any	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:

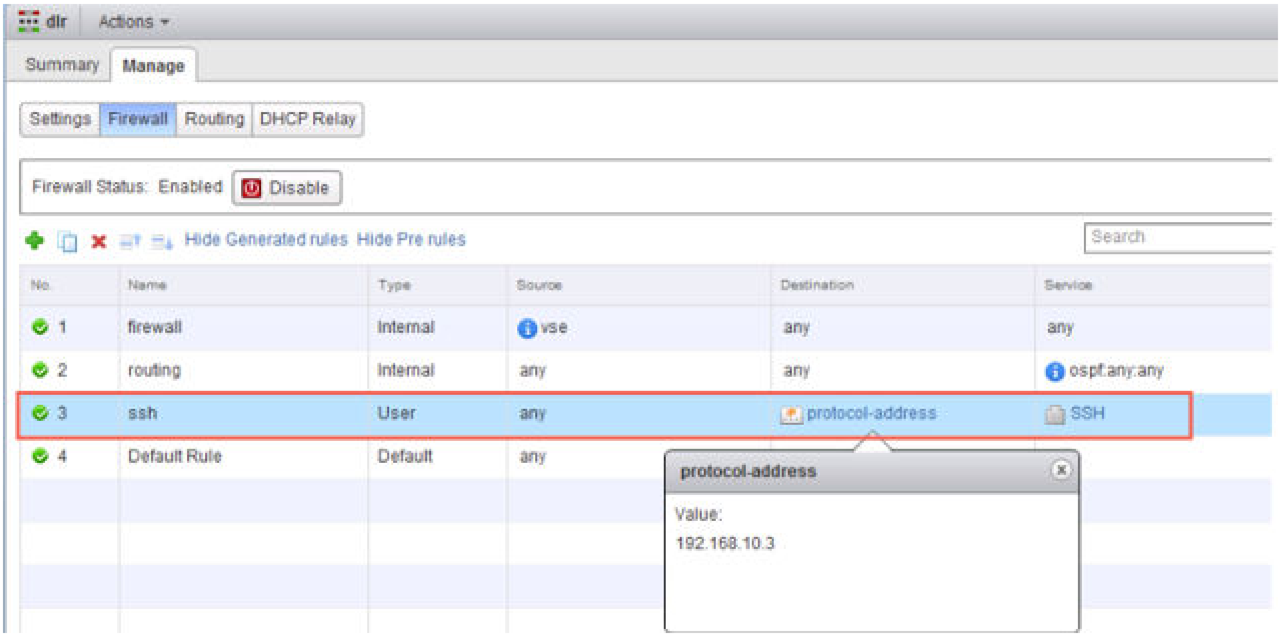

Firewall Status: Enabled

No.	Name	Type	Source	Destination	Service
✓ 1	firewall	Internal	vse	any	any
✓ 2	routing	Internal	any	any	ospf:any:any
✓ 3	ssh	User	any	protocol-address	SSH
✓ 4	Default Rule	Default	any		

protocol-address

Value:
192.168.10.3

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 Example: OSPF Configured on the Edge Services Gateway.

When you enable a router ID, the field 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** and then click **NSX Edges**.

3 Double-click an ESG.

4 Click **Routing** and then click **OSPF**.

5 Enable OSPF.

   a Click **Edit** at the top right corner of the window and click **Enable OSPF**

   b (Optional) Click **Enable Graceful Restart** for packet forwarding to be un-interrupted 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 the **Add** icon.

   c Type an Area ID. NSX Edge supports an area ID in the form of an IP address or decimal number.

   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.

8 Map interfaces to the areas.

a In **Area to Interface Mapping**, click the **Add** icon 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**.

11 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 as shown here) through an uplink portgroup on a vSphere distributed switch.

Figure 9-2. NSX Topology

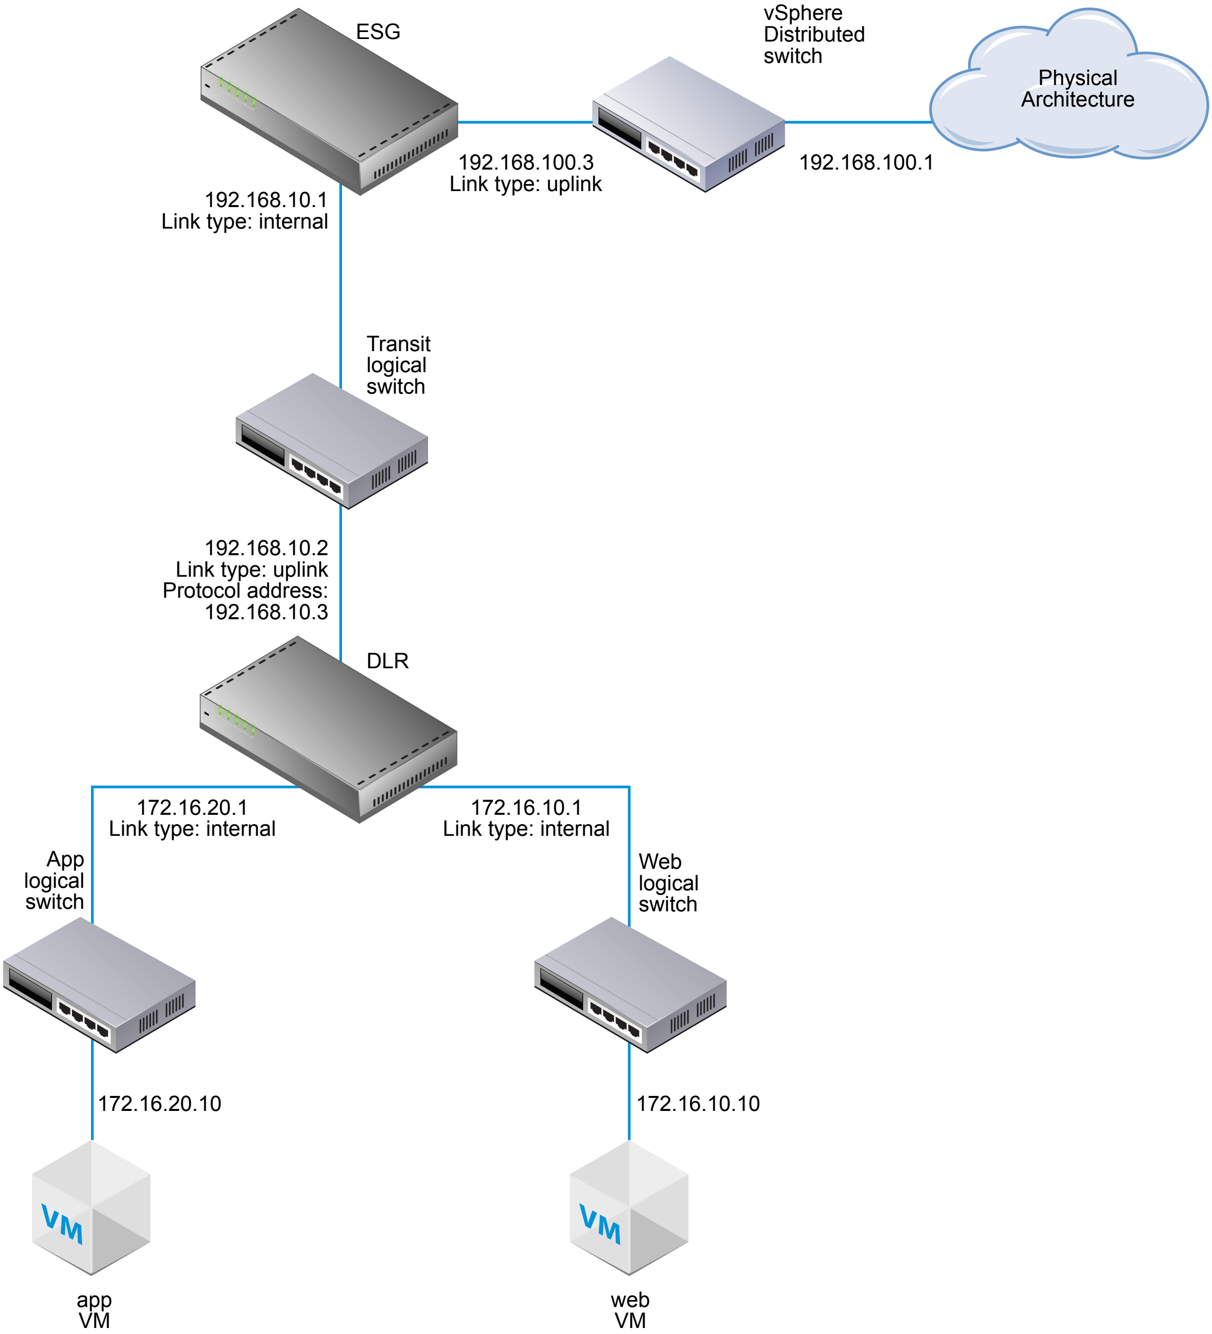
ESG
vSphere Distributed switch
Physical Architecture
192.168.100.3 Link type: uplink
192.168.100.1
192.168.10.1 Link type: internal
Transit logical switch
192.168.10.2 Link type: uplink Protocol address: 192.168.10.3
DLR
172.16.20.1 Link type: internal
172.16.10.1 Link type: internal
App logical switch
Web logical switch
172.16.20.10
172.16.10.10
VM
VM
app VM
web VM

In the following screen, the ESG's default gateway is the ESG's uplink interface to its external peer.

The router ID is the ESG's uplink interface IP address---in other words, the IP address that faces its external peer.

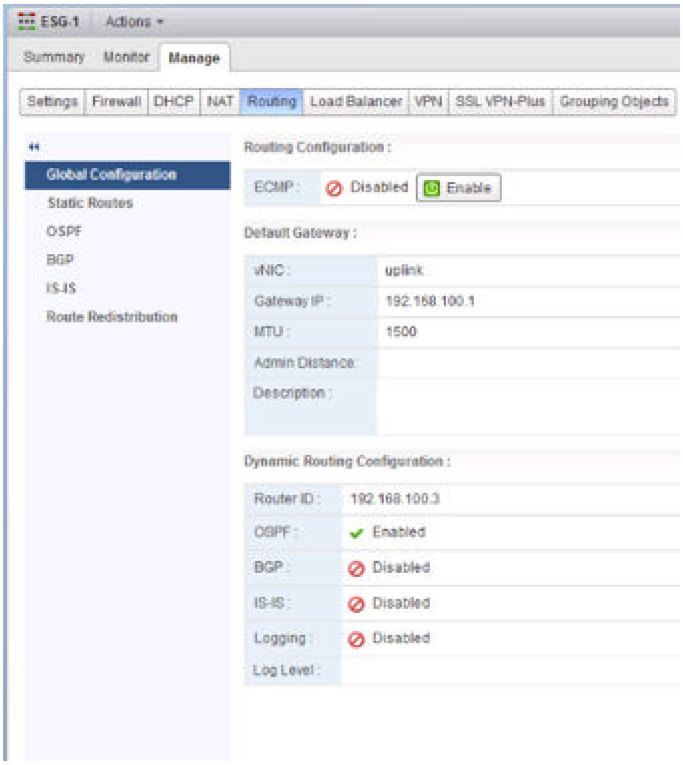

Routing Configuration :

ECMP : Disabled

Default Gateway :

vNIC :	uplink
Gateway IP :	192.168.100.1
MTU :	1500
Admin Distance :	
Description :	

Dynamic Routing Configuration :

Router ID :	192.168.100.3
OSPF :	✓ Enabled
BGP :	Disabled
IS-IS :	Disabled
Logging :	Disabled
Log Level :	

The area ID configured is 0, and the internal interface (the interface facing the logical router) is mapped to the area.

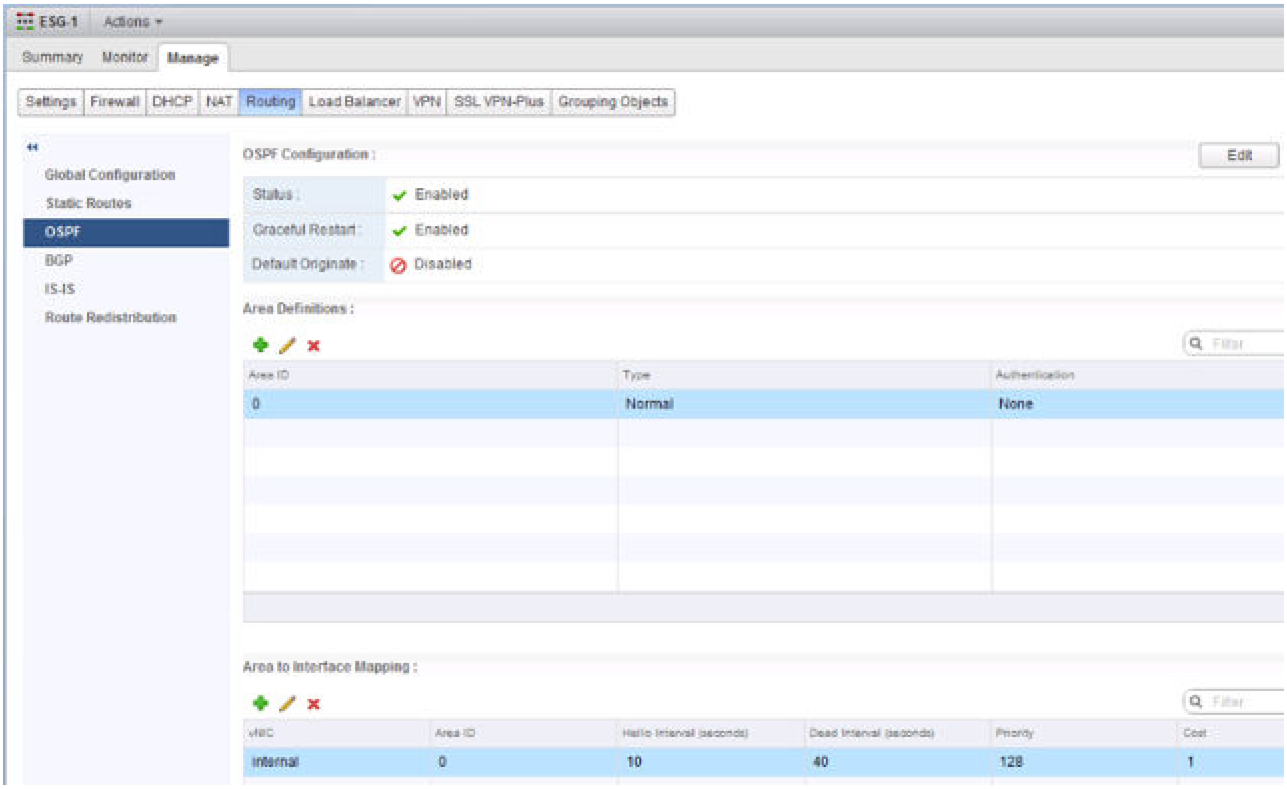

OSPF Configuration :

Status :	✓ Enabled
Graceful Restart :	✓ Enabled
Default Originate :	Disabled

Area Definitions :

Area ID	Type	Authentication
0	Normal	None

Area to Interface Mapping :

vNIC	Area ID	Hello Interval (seconds)	Dead Interval (seconds)	Priority	Cost
internal	0	10	40	128	1

The connected routes are redistributed into OSPF so that the OSPF neighbor (the logical router) can learn about the ESG's uplink network.

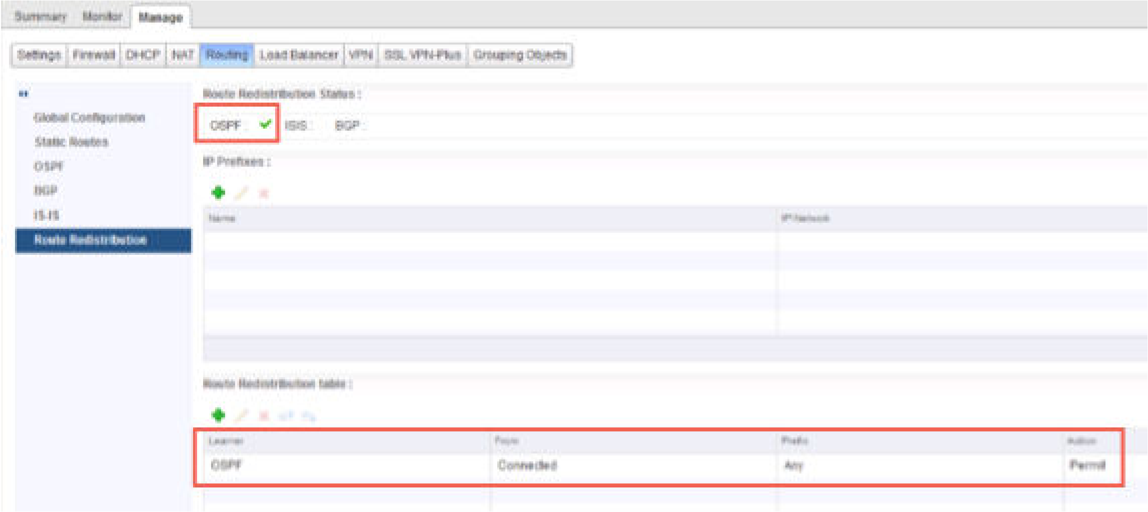

Route Redistribution Status :

OSPF : ✓ ISIS : BGP :

IP Prefixes :

Name	IP Network

Route Redistribution table :

Learner	From	Prefix	Action
OSPF	Connected	Any	Permit

Note Additionally, OSPF can 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.

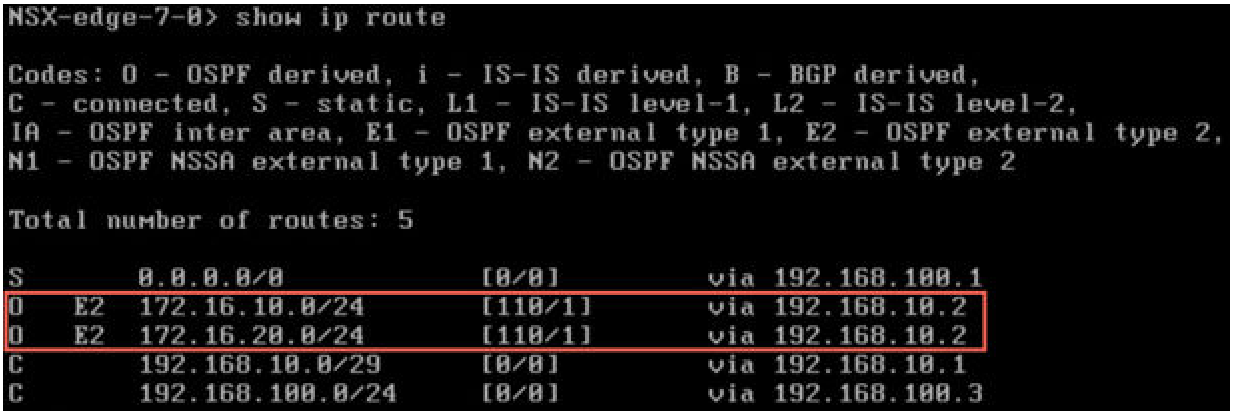

```
NSX-edge-7-0> show ip route

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** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click **Routing** and then click **BGP**.
5. Click **Edit**.
6. In the Edit BGP Configuration dialog box, click **Enable BGP**.
7. Click **Enable Graceful Restart** for packet forwarding to be un-interrupted during restart of BGP services.
8. Click **Enable Default Originate** to allow NSX Edge to advertise itself as a default gateway to its peers.
9. Type the router ID in **Local AS**. Type the Local AS. This is advertised when BGP peers with routers in other autonomous systems (AS). The path of ASs that a route traverses is used as one metric when selecting the best path to a destination.
10. Click **OK**.
11. In **Neighbors**, click the **Add** icon.
12. Type the IP address of the neighbor.

    When you configure BGP peering between an edge services gateway (ESG) and a logical router, use the logical router's protocol IP address as the ESG's BGP neighbor address.
13. (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).
14. (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 (as long as it is not used anywhere else). When you configure BGP peering between an edge services gateway (ESG) and a logical router, use the logical router's protocol IP address as the ESG neighbor's IP address.
15. Type the remote AS.
16. Edit the default weight for the neighbor connection if required.
17. **Hold Down Timer** displays interval (180 seconds) after not receiving a keep alive message that the software declares a peer dead. Edit if required.

18 **Keep Alive Timer** displays the default frequency (60 seconds) with which the software sends keepalive messages to its peer. Edit if required.

19 If authentication is required, type the authentication password. 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 will not be made.

20 To specify route filtering from a neighbor, click the **Add** icon in the **BGP Filters** area.

Caution A "block all" rule is enforced at the end of the filters.

21 Select the direction to indicate whether you are filtering traffic to or from the neighbor.

22 Select the action to indicate whether you are allowing or denying traffic.

23 Type the network in CIDR format that you want to filter to or from the neighbor.

24 Type the IP prefixes that are to be filtered and click **OK**.

25 Click **Publish Changes**.

Example: Configure BGP Between an ESG and a Logical Router

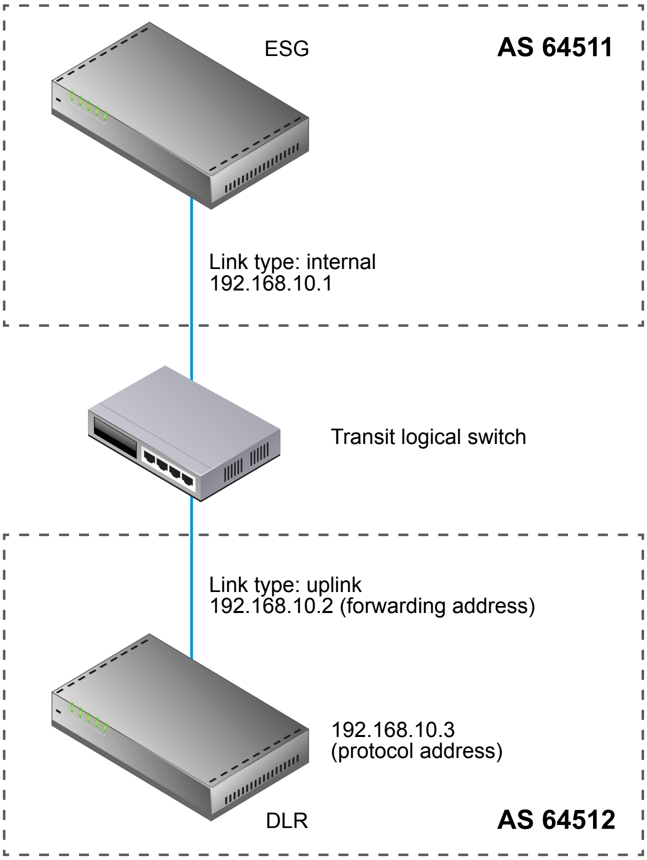

In this topology, the ESG is in AS 64511. The logical router (DLR) is in AS 64512.

The logical router's forwarding address is 192.168.10.2. This is the address configured on the logical router's uplink interface. The logical router's protocol address is 192.168.10.3. This is the address that the ESG will use to form its BGP peering relationship with the logical router.

On the logical router, configure BGP as shown:

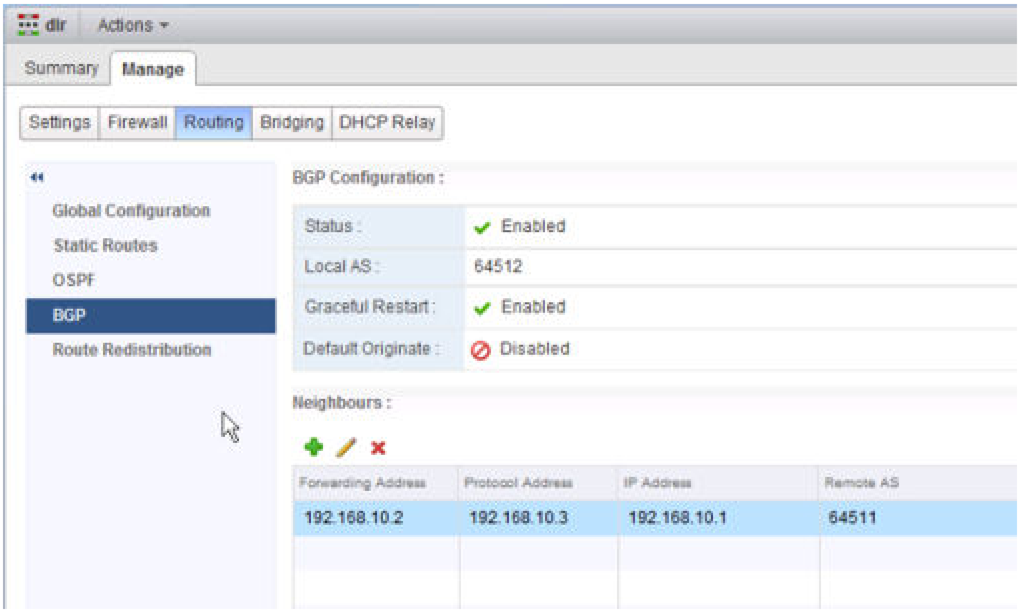

BGP Configuration :

Status :	✓ Enabled
Local AS :	64512
Graceful Restart :	✓ Enabled
Default Originate :	⊘ Disabled

Neighbours :

Forwarding Address	Protocol Address	IP Address	Remote AS
192.168.10.2	192.168.10.3	192.168.10.1	64511

On the ESG, configure BGP as shown:

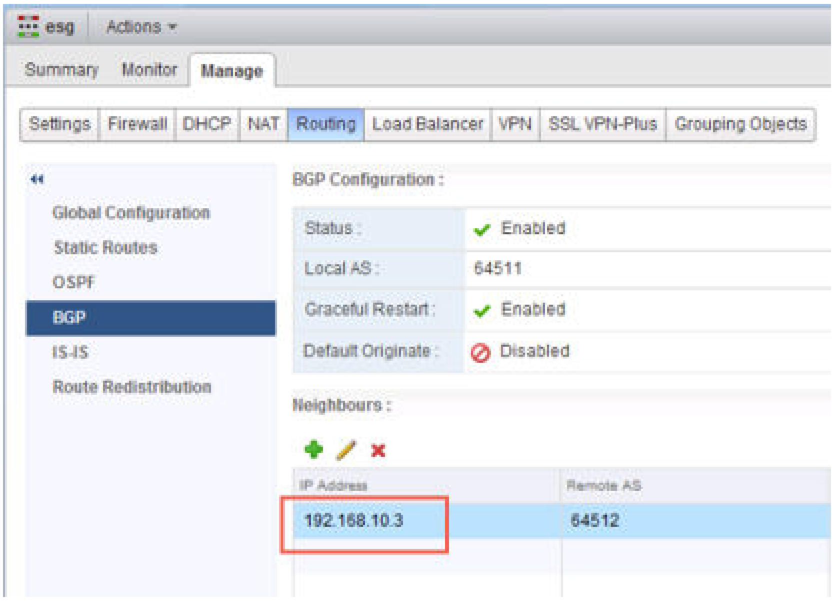

BGP Configuration :

Status :	✓ Enabled
Local AS :	64511
Graceful Restart :	✓ Enabled
Default Originate :	⊘ Disabled

Neighbours :

IP Address	Remote AS
192.168.10.3	64512

The ESG's neighbor address is 192.168.10.3, which is the logical router's protocol address.

Run the `show ip bgp neighbors` command on the logical router, and make sure the BGP state is Established.

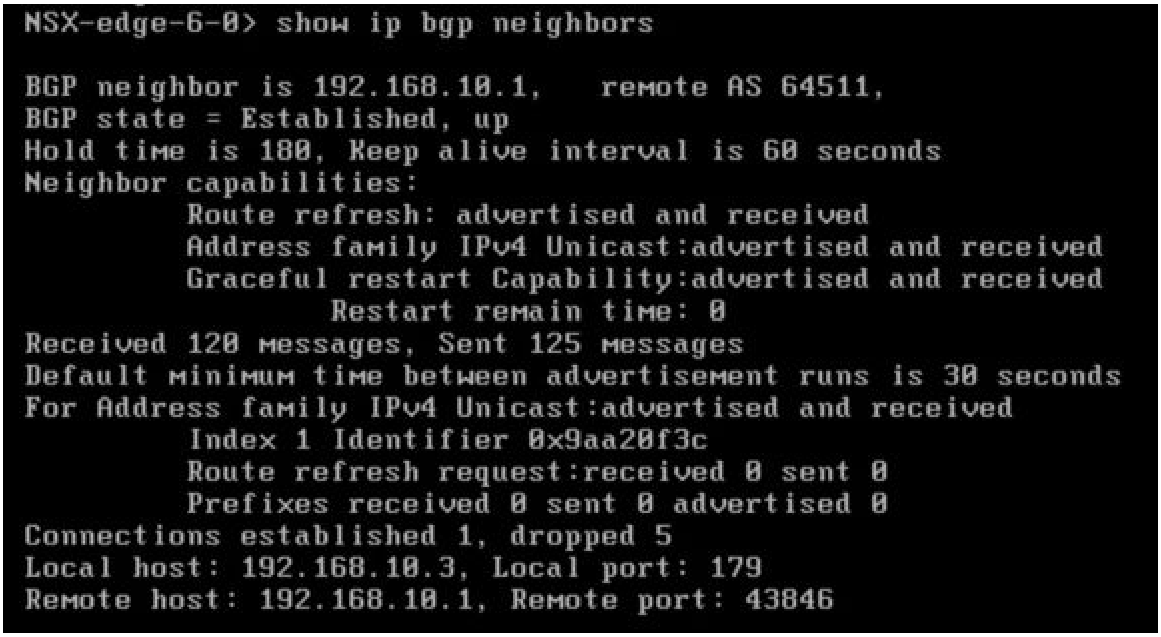

```
NSX-edge-6-0> show ip bgp neighbors

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
```

Run the `show ip bgp neighbors` command on the ESG, and make sure the BGP state is Established.

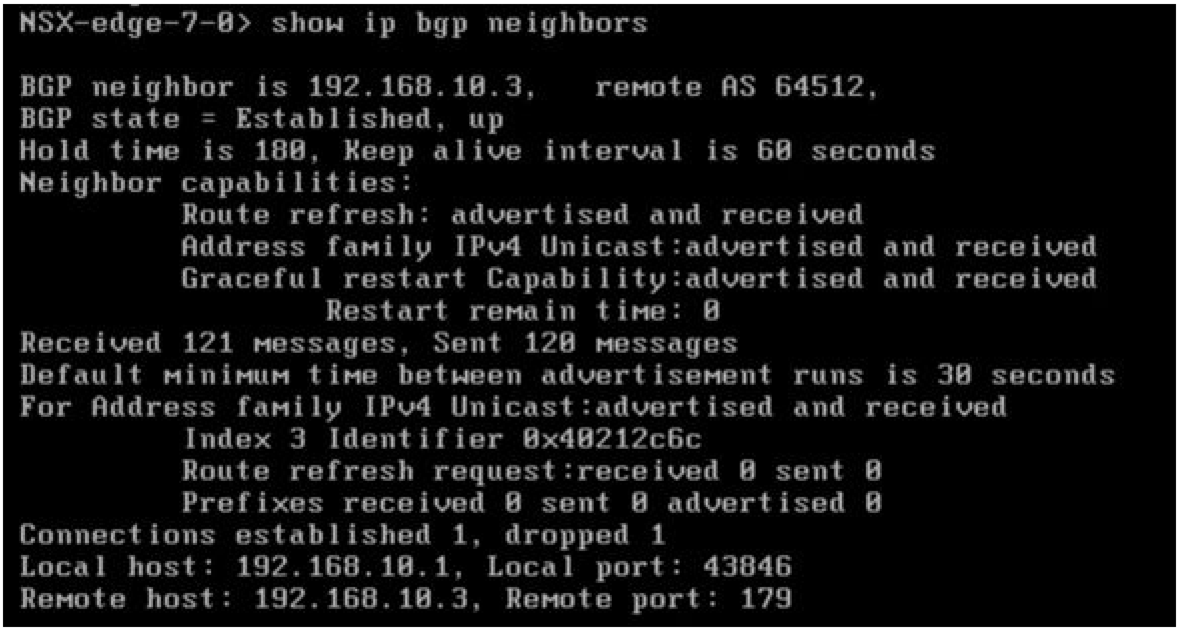

```
NSX-edge-7-0> show ip bgp neighbors

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 IS-IS Protocol

Intermediate System to Intermediate System (IS-IS) is a routing protocol designed to move information by determining the best route for datagrams through a packet-switched network.

A two-level hierarchy is used to support large routing domains. A large domain may be divided into areas. Routing within an area is referred to as Level 1 routing. Routing between areas is referred to as Level 2 routing. A Level 2 Intermediate System (IS) keeps track of the paths to destination areas. A Level 1 IS keeps track of the routing within its own area. For a packet going to another area, a Level 1 IS sends the packet to the nearest Level 2 IS in its own area, regardless of what the destination area is. Then the packet travels via Level 2 routing to the destination area, where it may travel via Level 1 routing to the destination. An IS in both Level 1 and Level 2 is referred to as Level-1-2.

Note NSX support for the IS-IS protocol is currently experimental.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click **Routing** and then click **IS-IS**.
5. Click **Edit** and then click **Enable IS-IS**.
6. Type the System ID and select the IS-IS type.

   Level 1 is intra-area, Level 2 is inter-area, and Level 1-2 is both. Level 2 routers are inter-area routers that can only form relationships with other Level 2 routers. Routing information is exchanged between Level 1 routers and other Level 1 routers. Likewise Level 2 routers only exchange information with other Level 2 routers. Level 1-2 routers exchange information with both levels and are used to connect the inter-area routers with the intra-area routers.
7. Type the **Domain Password** and **Area Password**. The area password is inserted and checked for Level 1 link state packets, and the domain password for Level 2 link state packets.
8. Define the IS-IS areas.
   a. Click the **Add** icon in **Areas**.
   b. Type up to three area IP addresses.
   c. Click **Save**.
9. Configure interface mapping.
   a. Click the **Add** icon in **Interface Mapping**.
   b. Choose the Circuit Type to indicate whether you are configuring the interface for Level-1, Level-2, or Level-1-2 adjacency.
   c. **Hello Interval** displays the default interval in milliseconds between hello packets that are sent on the interface. Edit the default value if required.
   d. **Hello Multiplier** displays the default number of IS-IS hello packets a neighbor must miss before it is declared down. Edit the default value if required.

e **LSP Interval** displays the time delay in milliseconds between successive IS-IS link-state packet (LSP) transmissions. Edit the default value if required.

f **Metric** displays the default metric for the interface. This is used to calculate the cost from each interface via the links in the network to other destinations. Edit the default value if required.

g **Priority** displays the priority of the interface. The interface with the highest priority becomes the designated router. Edit the default value if required.

h In Mesh Group, type the number identifying the mesh group to which this interface belongs. Edit the default value if required.

i Type the authentication password for the interface and click **OK**. Edit the default value if required.

10 Click **Publish Changes**.

Configure Route Redistribution

By default, routers share routes with other routers 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. In 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** and then click **NSX Edges**.

3 Double-click an NSX Edge.

4 Click **Routing** and then click **Route Redistribution**.

5 Click **Edit** next to **Route Redistribution Status**.

6 Select the protocols for which you enable route redistribution and click **OK**.

7 Add an IP prefix.

Entries in the IP Prefix list are processed sequentially.

a Click the **Add** icon in **IP Prefixes**.

b Type a name and IP address of the network.

The IP prefix entered will be exactly matched, except if you include less-than-or-equal-to (LE) or greater-than-or-equal-to (GE) modifiers.

c Click **OK**.

8 Specify redistribution criteria for the IP prefix.

a Click the **Add** icon in **Route Redistribution table**.

b In **Learner Protocol**, select the protocol that is to learn routes from other protocols.

   c In **Allow Learning from**, select the protocols from which routes should be learned.

   d Click **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 In the vSphere Web Client, navigate to **Networking & Security**, then under **Networking & Security Inventory** click an NSX Manager.

2 Click the **Summary** tab. The **ID** field contains 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 will be associated with all routes on this router (static and dynamic). The routes will be 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** and then click **NSX Edges**.

3 Double-click a universal logical (distributed) router.

4 Click the **Routing** tab, then click **Global Configuration**.

5 Click **Edit** next to **Routing Configuration**.

6 Type 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. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **Installation**.
3. Click the **Host Preparation** tab
4. Select the NSX Manager that manages the hosts or clusters you need to configure.
5. Select the host or cluster you want to modify, expanding clusters to display hosts if needed.
6. Click the **Settings** icon (⚙) and click **Change Locale ID**.
7. Type a new locale ID and click **OK.**

   Note 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).

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.

Logical Firewall

10

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

Distributed firewall is a hypervisor kernel-embedded firewall that provides visibility and control for virtualized workloads and networks. You can create access control policies based on VMware vCenter objects like datacenters and clusters and virtual machine names; network constructs like IP or IPSet addresses, VLAN (DVS port-groups), VXLAN (logical switches), security groups, as well as user group identity from Active Directory. Firewall rules are enforced at the vNIC level of each virtual machine to provide consistent access control even when the virtual machine gets vMotioned. The hypervisor-embedded nature of the firewall delivers close to line rate throughput to enable higher workload consolidation on physical servers. The distributed nature of the firewall provides a scale-out architecture that automatically extends firewall capacity when additional hosts are added to a datacenter.

For L2 packets, distributed firewall creates a cache for performance boost. L3 packets are processed in the following sequence:

1. All packets are checked for an existing state. This is done for SYNs too so that bogus or retransmitted SYNs for existing sessions can be detected.
2. If a state match is found, the packets are processed.
3. If a state match is not found, the packet is processed through the rules until a match is found.
   - For TCP packets, a state is set only for packets with a SYN flag. However, rules that do not specify a protocol (service ANY), can match TCP packets with any combination of flags.
   - For UDP packets, 5-tuple details are extracted from the packet. If a state does not exist in the state table, a new state is created using the extracted 5-tuple details. Subsequently received packets are matched against the state that was just created.
   - For ICMP packets, ICMP type, code, and packet direction are used to create a state.

Distributed firewall can help in creating identity-based rules as well. Administrators can enforce access control based on the user's group membership as defined in the enterprise Active Directory. Here are some scenarios where identity-based firewall rules can be used:

- User accessing virtual applications using a laptop or mobile device where AD is used for user authentication
- User accessing virtual applications using VDI infrastructure where the virtual machines are Microsoft Windows based

If you have a third-party vendor firewall solution deployed in your environment, see Redirecting Traffic to a Vendor Solution through Logical Firewall.

Running open VMware Tools on guest or workload virtual machines has not been validated with distributed firewall.

ESXi Threshold Parameters for Distributed Firewall Resource Utilization

Each ESXi host is configured with three threshold parameters for DFW resource utilization: CPU, RAM, and connections per second (CPS). 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.

100 percent of CPU corresponds to the total CPU available on the host.

100 percent of RAM corresponds to the memory allocated for distributed firewall ("total max size"), which is dependent on the total amount of RAM installed in the host.

Table 10-1. Total Max Size

Physical Memory	Total Max Size (MB)
0 - 8GB	160
8GB - 32GB	608
32GB - 64GB	992

Table 10-1. Total Max Size (Continued)

Physical Memory	Total Max Size (MB)
64GB - 96GB	1920
96GB - 128GB	2944
128GB	4222

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:

```
https://NSX-MGR-IP/api/4.0/firewall/stats/eventthresholds

Request body:

<eventThresholds>
  <cpu>
    <percentValue>100</percentValue>
  </cpu>
  <memory>
    <percentValue>100</percentValue>
  </memory>
  <connectionsPerSecond>
    <value>100000</value>
  </connectionsPerSecond>
</eventThresholds>
```

Edge Firewall

Edge Firewall monitors North-South traffic to provide perimeter security functionality including firewall, Network Address Translation (NAT) as well as 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 and/or uplink interfaces work, however, the rules on internal interfaces do not work.

Note NSX-V Edge 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. Edge must defend from Syn-Flood attacks by implementing logic to detect bogus TCP connections and terminate them without consuming Firewall state tracking resources. This feature is disabled by default. To enable this feature in a high risk environment, set the REST API `enableSynFloodProtection` value to 'true' as part of the Firewall Global Configuration.

Working with Firewall Rule Sections

You can add a section to segregate firewall rules. For example, you might want to have the rules for sales and engineering departments in separate sections.

You can create multiple sections for L2 and L3 rules.

Cross-vCenter NSX environments can have one universal L2 rule section and one universal L3 rule section. You must manage universal rules on the primary NSX Manager, and you must create the universal section there before you can add universal rules.

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 a new section 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 > 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 **General** tab to add a section for L3 rules. Click the **Ethernet** tab to add a section for L2 rules.
4. Click the **Add Section** (📁) icon.
5. Type a name for the section and specify the position for the new section. Section names must be unique within NSX Manager.
6. (Optional) To create a universal section, select **Mark this section for Universal Synchronization**.
7. Click **OK** and then click **Publish Changes**.

What to do next

Add rules to the section. You can edit the name of a section by clicking the **Edit section (✏)** icon for that section.

Merge Firewall Rule Sections

You can merge sections and consolidate the rules within those sections. Note that 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 > Firewall**.

2 For the section you want to merge, click the **Merge** (↰) icon and specify 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.

3 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 > Firewall**.

2 Ensure that you are in the **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**.

The section as well as all rules in that section are deleted.

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 17 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*.

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 > 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.

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.

The following vCenter objects can be specified as the source or destination for a firewall rule:

Table 10-2. Objects supported for firewall rules

Source or Destination	Applied To
■ cluster ■ datacenter ■ distributed port group ■ IP set ■ legacy port group ■ logical switch ■ resource pool ■ security group ■ vApp ■ virtual machine ■ vNIC ■ IP address (IPv4 or IPv6)	■ All clusters on which Distributed Firewall has been installed (in other words, all clusters that have been prepared for network virtualization) ■ All Edge gateways installed on prepared clusters ■ cluster ■ datacenter ■ distributed port group ■ Edge ■ legacy port group ■ logical switch ■ security group ■ virtual machine ■ vNIC

Prerequisites

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 REST 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.

If you are adding universal firewall rules, see Add a Universal Firewall Rule

If you are adding an identity-based firewall rule, 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.
- 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.

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*.

Procedure

1 In the vSphere Web Client, navigate to **Networking & Security > Firewall**.

2 Ensure that you are in the **General** tab to add an L3 rule. Click the **Ethernet** tab to add an L2 rule.

3 In the section in which you add a rule, click **Add rule** (✚) icon.

4 Click **Publish Changes**.

A new any any allow 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.

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**.

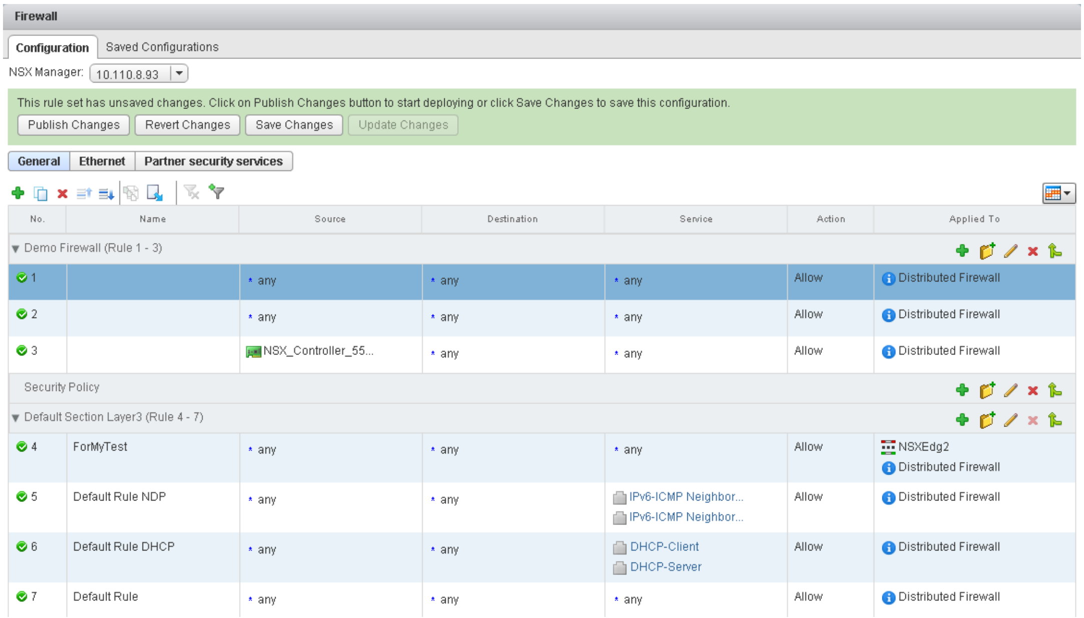

This rule set has unsaved changes. Click on Publish Changes button to start deploying or click Save Changes to save this configuration.

No.	Name	Source	Destination	Service	Action	Applied To
Demo Firewall (Rule 1 - 3)						
1		any	any	any	Allow	Distributed Firewall
2		any	any	any	Allow	Distributed Firewall
3		NSX_Controller_55...	any	any	Allow	Distributed Firewall
Security Policy						
Default Section Layer3 (Rule 4 - 7)						
4	ForMyTest	any	any	any	Allow	NSXEdg2 Distributed Firewall
5	Default Rule NDP	any	any	IPv6-ICMP Neighbor... IPv6-ICMP Neighbor...	Allow	Distributed Firewall
6	Default Rule DHCP	any	any	DHCP-Client DHCP-Server	Allow	Distributed Firewall
7	Default Rule	any	any	any	Allow	Distributed Firewall

5 Point to the **Name** cell of the new rule and click ⊞.

6 Type a name for the new rule.

7 Point to the **Source** cell of the new rule. Additional icons are displayed as described in the table below.

Option	Description
Click IP	To specify source as an IP address. a Select the IP address format. Firewall supports both IPv4 and IPv6 formats. b Type the IP address. You can enter multiple IP addresses in a comma-separated list. The list can contain up to 255 characters.
Click ✎	To specify source as an object other than a specific IP address. a In **View**, select a container from which the communication originated. Objects for the selected container are displayed. b Select one or more objects and click ➡. You can create a new security group or IPSet. Once you create the new object, it is added to the source column by default. For information on creating a new security group or IPSet, see Chapter 22 Network and Security Objects. c To exclude a source from the rule, click **Advanced options**. d Select **Negate Source** to exclude this source from the rule. If **Negate Source** is selected, the rule is applied to traffic coming from all sources except for the source you specified in the previous step. If **Negate Source** is not selected, the rule applies to traffic coming from the source you specified in the previous step. e Click **OK**.

8 Point to the **Destination** cell of the new rule. Additional icons are displayed as described in the table below.

Option	Description
Click IP	To specify destination as an IP address. a Select the IP address format. Firewall supports both IPv4 and IPv6 formats. b Type the IP address. You can enter multiple IP addresses in a comma-separated list. The list can contain up to 255 characters.
Click ✏	To specify destination as an object other than a specific IP address. a In **View**, select a container which the communication is targeting. Objects for the selected container are displayed. b Select one or more objects and click ➡. You can create a new security group or IPSet. Once you create the new object, it is added to the Destination column by default. For information on creating a new security group or IPSet, see Chapter 22 Network and Security Objects. c To exclude a destination port, click **Advanced options**. d Select **Negate Destination** to exclude this destination from the rule. If **Negate Destination** is selected, the rule is applied to traffic going to all destinations except for the destination you specified in the previous step. If **Negate Destination** is not selected, the rule applies to traffic going to the destination you specified in the previous step. e Click **OK**.

9 Point to the **Service** cell of the new rule. Additional icons are displayed as described in the table below.

Option	Description
Click □	To specify service as a port protocol combination. a Select the service protocol. 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. b Type the port number and click **OK**.
Click ✎	To select a pre-defined service/service group or define a new one. a Select one or more objects and click ➡. 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. b Click **OK**.

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.

10 Point to the **Action** cell of the new rule and click ⊞. Make appropriate selections as described in the table below and click **OK**.

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.
Log	Logs all sessions matching this rule. Enabling logging can affect performance.
Do not log	Does not log sessions.

11 In **Applied To**, define the scope at which this rule is applicable. Make appropriate selections as described in the table below and click **OK**.

To apply a rule to	Do this
All prepared clusters in your environment	Select **Apply this rule on all clusters on which Distributed Firewall is enabled**. After you click OK, the Applied To column for this rule displays **Distributed Firewall**.
All NSX Edge gateways in your environment	Select **Apply this rule on all Edge gateways**. After you click OK, the Applied To column for this rule displays **All Edges**. If both the above options are selected, the Applied To column displays **Any**.
One or more cluster, datacenter, distributed virtual port group, NSX Edge, network, virtual machine, vNIC, or logical switch	1 In **Container type**, select the appropriate object.. 2 In the Available list, select one or more objects and click ➡.

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.

12 Click **Publish Changes**.

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**.

When you click **Publish Changes**, the firewall configuration is automatically saved. For information on reverting to an earlier configuration, see Load Firewall Configuration.

What to do next

- Disable 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**.

Load 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 > Firewall**.
2. Ensure that you are in the **General** tab to load an L3 firewall configuration. Click the **Ethernet** tab to load an L2 firewall configuration.
3. Click the **Load configuration** ( ) icon.
4. Select the configuration to load and click **OK**.

   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.

Add a Universal Firewall Rule

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. Edge firewall rules are not supported for vMotion between multiple vCenter Servers.

The primary NSX Manager can contain one universal section for universal L2 rules and one universal section for universal L3 rules. 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.

Table 10-3. Objects supported for universal firewall rules

Source and Destination	Applied To	Service
■ universal MAC set ■ universal IP set ■ universal security group, which can contain an IP set, MAC set, or universal security group ■ universal logical switch	■ universal logical switch ■ Distributed Firewall - applies rules on all clusters on which Distributed Firewall is installed	■ pre-created universal services and service groups ■ user created universal services and services groups

Note that other vCenter objects are not supported for universal rules.

Prerequisites

You must create a universal rule section before you can create universal rules. See Add a Firewall Rule Section.

Procedure

1 In the vSphere Web Client, navigate to **Networking & Security > Firewall**.

2 In NSX Manager, ensure that the primary NSX Manager is selected.

Universal rules can only be added on the primary NSX Manager.

3 Ensure that you are in the **General** tab to add an L3 universal rule. Click the **Ethernet** tab to add an L2 universal rule.

4 In the universal section click the **Add rule** (✚) icon and then click **Publish Changes**.

A new any any allow rule is added at the top of the universal section.

5 Point to the **Name** cell of the new rule and click ✎. Type a name for the rule.

6 Point to the **Source** cell of the new rule. Additional icons are displayed as described in the table below.

Option	Description
Click IP	To specify source as an IP address. a Select the IP address format. Firewall supports both IPv4 and IPv6 formats. b Type the IP address.
Click ✎	To specify a universal IPSet, MACSet, or security group as the source. a In **Object Type**, select a container from which the communication originated. Objects for the selected container are displayed. b Select one or more objects and click ➡. You can create a new security group or IPSet. Once you create the new object, it is added to the source column by default. For information on creating a new security group or IPSet, see Chapter 22 Network and Security Objects. c To exclude a source from the rule, click **Advanced options**. d Select **Negate Source** to exclude this source from the rule. If **Negate Source** is selected, the rule is applied to traffic coming from all sources except for the source you specified in the previous step. If **Negate Source** is not selected, the rule applies to traffic coming from the source you specified in the previous step. e Click **OK**.

7 Point to the **Destination** cell of the new rule. Additional icons are displayed as described in the table below.

Option	Description
Click IP	To specify destination as an IP address. a Select the IP address format. Firewall supports both IPv4 and IPv6 formats. b Type the IP address.
Click ✏	To specify a universal IPSet, MACSet, or security group as the destination. a In **Object Type**, select a container which the communication is targeting. Objects for the selected container are displayed. b Select one or more objects and click ➡. You can create a new security group or IPSet. Once you create the new object, it is added to the Destination column by default. For information on creating a new security group or IPSet, see Chapter 22 Network and Security Objects. c To exclude a destination from the rule, click **Advanced options**. d Select **Negate Destination** to exclude this destination from the rule. If **Negate Destination** is selected, the rule is applied to traffic going to all destinations except for the destination you specified in the previous step. If **Negate Destination** is not selected, the rule applies to traffic going to the destination you specified in the previous step. e Click **OK**.

8 Point to the **Service** cell of the new rule. Additional icons are displayed as described in the table below.

Option	Description
Click □	To specify a service as a port protocol combination. a Select the service protocol. Distributed Firewall supports ALG (Application Level Gateway) for the following protocols: FTP, CIFS, ORACLE TNS, MS-RPC, and SUN-RPC. b Type the port number and click **OK**.
Click ✏	To select a pre-defined universal service/universal service group or define a new one. a Select one or more objects and click ➡. 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. b Click **OK**.

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.

9 Point to the **Action** cell of the new rule and click ✎. Make appropriate selections as described in the table below and click **OK**.

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.
Log	Logs all sessions matching this rule. Enabling logging can affect performance.
Do not log	Does not log sessions.

10 In **Applied To** cell, either accept the default setting, Distributed Firewall, to apply the rule on all clusters with Distributed Firewall enabled, or click the edit icon ✎ to select the universal logical switches on which the rule is to be applied to.

11 Click **Publish Changes**.

The universal rule is replicated on all secondary NSX Managers. The Rule ID stays the same across all NSX instances. To display the Rule ID, click ▦▾ and then click Rule ID.

Universal rules can be edited on the primary NSX Manager and are read only on secondary NSX Managers.

Firewall rules with Universal Section Layer3 and Default Section Layer3:

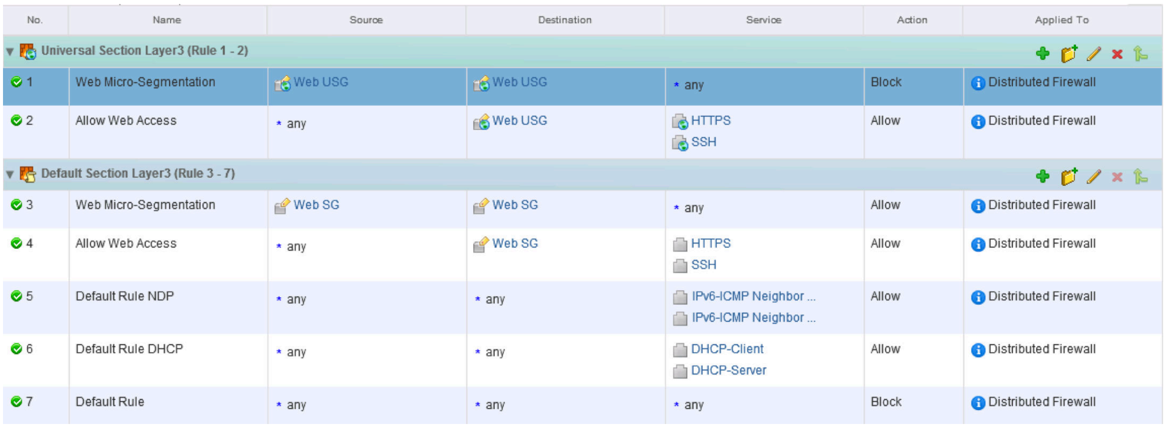

No.	Name	Source	Destination	Service	Action	Applied To
▾ Universal Section Layer3 (Rule 1 - 2)						
✔ 1	Web Micro-Segmentation	Web USG	Web USG	* any	Block	Distributed Firewall
✔ 2	Allow Web Access	* any	Web USG	HTTPS SSH	Allow	Distributed Firewall
▾ Default Section Layer3 (Rule 3 - 7)						
✔ 3	Web Micro-Segmentation	Web SG	Web SG	* any	Allow	Distributed Firewall
✔ 4	Allow Web Access	* any	Web SG	HTTPS SSH	Allow	Distributed Firewall
✔ 5	Default Rule NDP	* any	* any	IPv6-ICMP Neighbor ... IPv6-ICMP Neighbor ...	Allow	Distributed Firewall
✔ 6	Default Rule DHCP	* any	* any	DHCP-Client DHCP-Server	Allow	Distributed Firewall
✔ 7	Default Rule	* any	* any	* any	Block	Distributed Firewall

What to do next

- Disable a rule by clicking ✔ in the No. column, 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.

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.

Procedure

1 In the Firewall tab, click the **Apply Filter** ( ) icon.

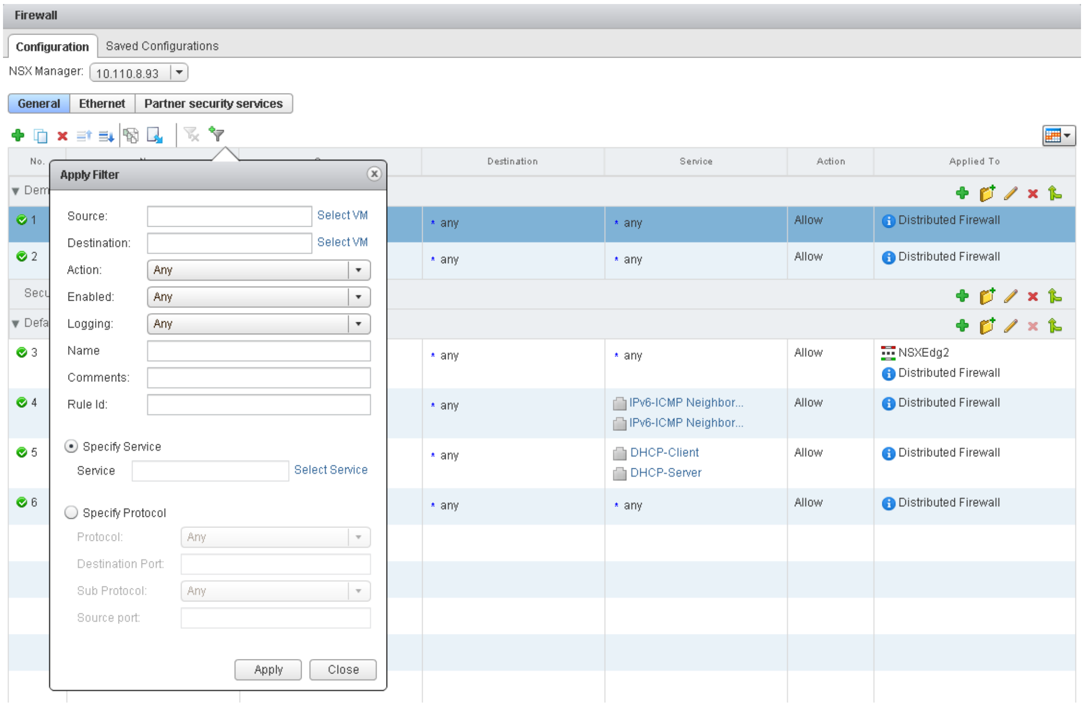

2 Type or select the filtering criteria as appropriate.

3 Click **Apply**.

Rules matching your filtering criteria are displayed.

What to do next

To display all rules again, click the **Remove applied filter** ( ) icon.

Add a Rule and Publish It at a Later Time

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**.

   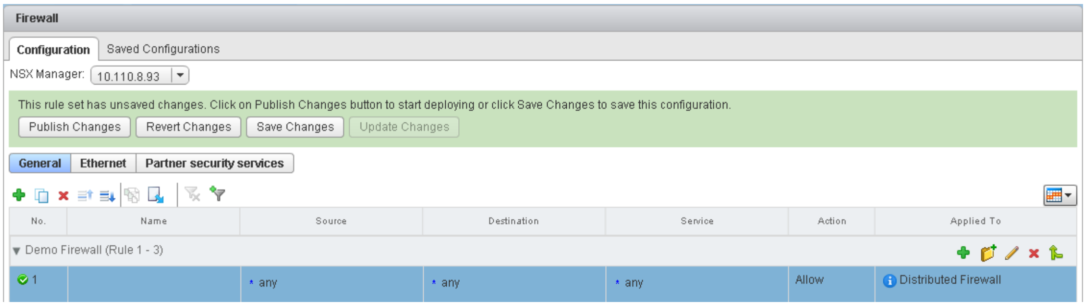

   This rule set has unsaved changes. Click on Publish Changes button to start deploying or click Save Changes to save this configuration.

No.	Name	Source	Destination	Service	Action	Applied To
▼ Demo Firewall (Rule 1 - 3)						
✓ 1		* any	* any	* any	Allow	Distributed Firewall

3. Type a name and description for the configuration and click **OK**.

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. Do one of the following.

   - 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.

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:

1. User-defined pre rules have the highest priority and are enforced in top-to-bottom ordering with a per-virtual NIC level precedence.

2. Auto-plumbed rules.

3. 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 17 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**.

Delete a Firewall Rule

You can delete firewall rules that you created. You cannot delete the default rule or rules managed by Service Composer.

Procedure

1. In the **Firewall** tab, select a rule.
2. Click **Delete selected rule** (✖) icon above the Firewall table.
3. Click **Publish Changes**.

Exclude Virtual Machines from Firewall Protection

You can exclude a set of virtual machines from NSX distributed firewall protection.

NSX Manager, NSX Controllers, and NSX Edge virtual machines are automatically excluded from NSX distributed firewall protection. In addition, VMware recommends that you 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.
- Partner service virtual machines.
- Virtual machines that require promiscuous mode. If these virtual machines are protected by NSX 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. In the vSphere Web Client, click **Networking & Security**.

2 In **Networking & Security Inventory**, click **NSX Managers**.

3 In the **Name** column, click an NSX Manager.

4 Click the **Manage** tab and then click the **Exclusion List** tab.

5 Click the **Add** (✚) icon.

6 Type the name of the virtual machine that you want to exclude and click **Add**.

For example:

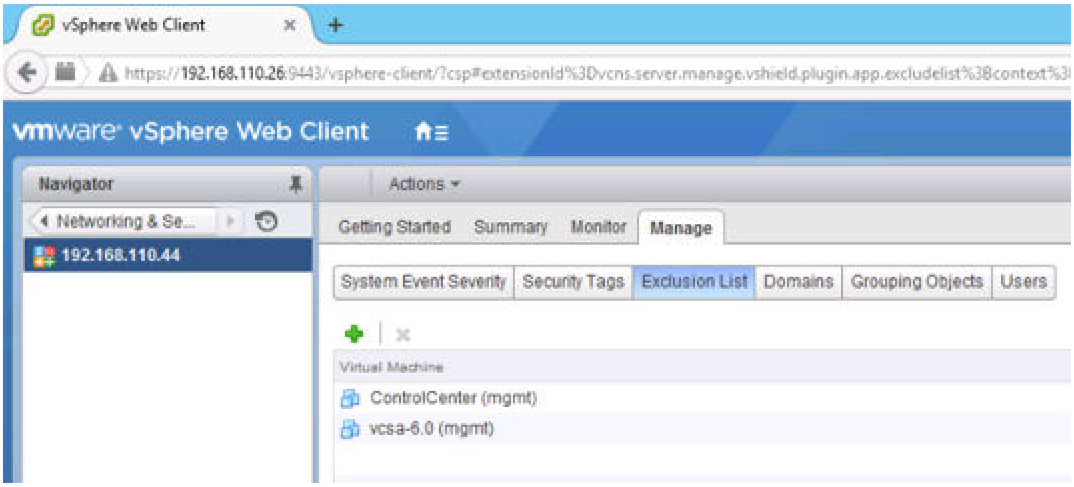

Virtual Machine
ControlCenter (mgmt)
vcsa-6.0 (mgmt)

7 Click **OK**.

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. In order to exclude these 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.

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 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 many other functions:

- allowing copy and paste between VM and host or client desktop
- synchronizing time with the host operating system
- allowing shutdown or restart of the VM from vCenter

- collecting network, disk, and memory usage from the VM and sending it to the host
- determining VM availability by sending and collecting heartbeat

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.

Change IP Detection Type

The IP address of a virtual machine can be detected by VMware Tools, which are installed on the VM, or by DHCP snooping and ARP snooping, which are enabled on the host cluster. These IP discovery methods can be used together in the same NSX installation.

Procedure

1. In the vSphere Web client, navigate to **Networking & Security > Installation > Host Preparation**.
2. Click the cluster you want to change, then click **Actions (⚙) > Change IP Detection Type**.
3. Select the desired detection types and click **OK**.

What to do next

Configure SpoofGuard.

Configure the default firewall rule.

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, click **Networking & Security** and then click **NSX Managers**.
2. In the **Name** column, click the IP address of the appropriate NSX Manager.
3. Click the **Monitor** tab and then click **System Events**.

Firewall Logs

Firewall generates and stores log files, such as audit log, rules message log, and system event log.

Firewall generates three types of 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 are stored on each host in `/var/log/dfwpktlogs.log`.

  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.

  ```
  ~ # 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
  ```

  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
  ```

  To enable rules message logging in vSphere Web Client 6.0 (the UI might differ slightly in vSphere 5.5, but the steps are the same):

  a Enable the **Log** column on the **Networking & Security > Firewall** page.

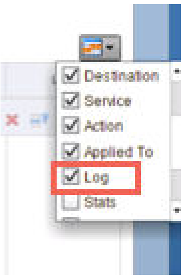

b Enable logging for a rule by hovering over the Log table cell and clicking the pencil icon.

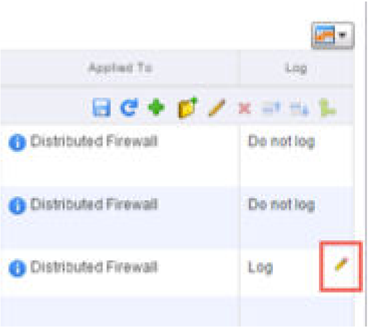

Applied To	Log
Distributed Firewall	Do not log
Distributed Firewall	Do not log
Distributed Firewall	Log

- 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, etc. These are stored in `/home/secureall/secureall/logs/vsm.log`.

To view audit and system event logs in the UI, navigate to **Networking & Security > Installation > Management** and double-click the IP address of the NSX Manager. Then select the **Monitor** tab.

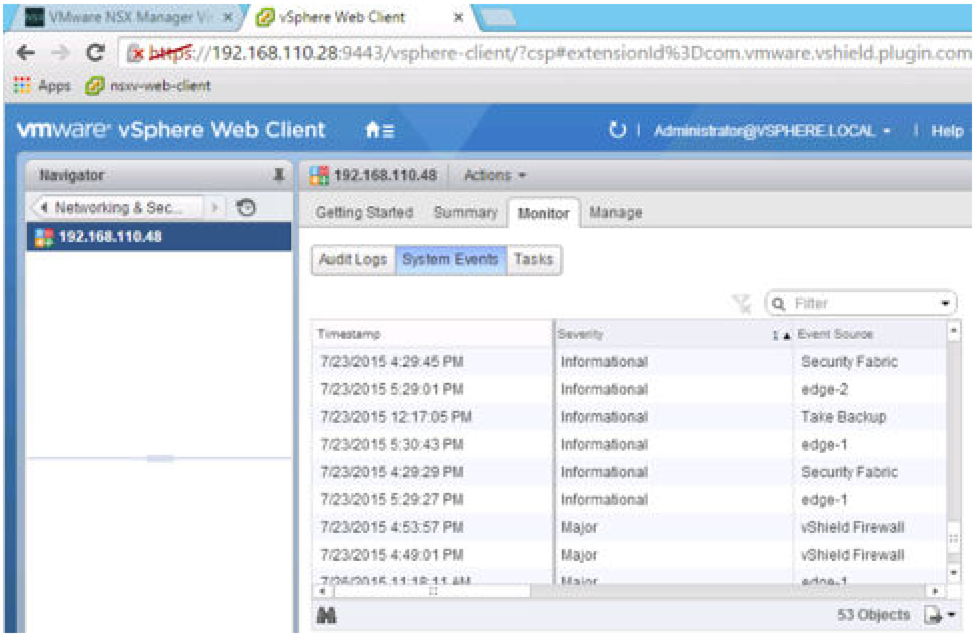

Timestamp	Severity	Event Source
7/23/2015 4:29:45 PM	Informational	Security Fabric
7/23/2015 5:29:01 PM	Informational	edge-2
7/23/2015 12:17:05 PM	Informational	Take Backup
7/23/2015 5:30:43 PM	Informational	edge-1
7/23/2015 4:29:29 PM	Informational	Security Fabric
7/23/2015 5:29:27 PM	Informational	edge-1
7/23/2015 4:53:57 PM	Major	vShield Firewall
7/23/2015 4:49:01 PM	Major	vShield Firewall
7/26/2015 11:18:11 AM	Major	edge-1

For more information, see Chapter 23 Operations and Management.

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 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 created on the Firewall user interface applicable to this NSX Edge are displayed in a read-only mode. Rules are displayed and enforced in the following order:

1 User-defined rules from the Firewall user interface (Read Only).

2 Auto-plumbed rules (rules that enable control traffic to flow for Edge services).

3 User-defined rules on 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 the **Manage** tab and then click **Firewall**.

4 Select the **Default Rule**, which is the last rule in the firewall table.

5 Point to the **Action** cell of the new rule and click ✏.

  a Click **Accept** to allow traffic from or to the specified source and destination.

  b Click **Log** to log all sessions matching this rule.

  Enabling logging can affect performance.

  c Type comments if required.

  d Click **OK**.

6 Click **Publish Changes**.

Add an NSX Edge Firewall Rule

The Edge Firewall tab displays rules created on the centralized Firewall tab in a read-only mode. Any rules that you add here are not displayed on the centralized Firewall tab.

You can add multiple NSX Edge interfaces and/or IP address groups as the source and destination for firewall rules.

Figure 10-1. Firewall rule for traffic to flow from an NSX Edge interface to an HTTP server

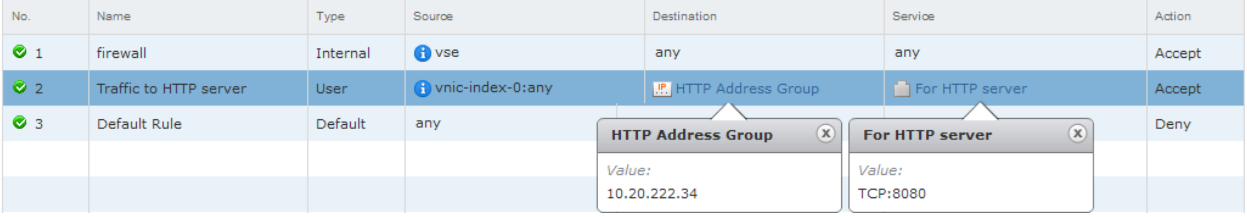

No.	Name	Type	Source	Destination	Service	Action
✓ 1	firewall	Internal	vse	any	any	Accept
✓ 2	Traffic to HTTP server	User	vnic-index-0:any	HTTP Address Group	For HTTP server	Accept
✓ 3	Default Rule	Default	any			Deny

HTTP Address Group ⊗
Value:
10.20.222.34

For HTTP server ⊗
Value:
TCP:8080

Figure 10-2. 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

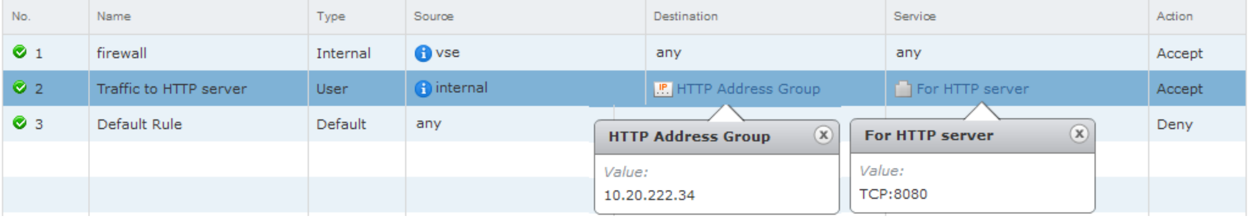

No.	Name	Type	Source	Destination	Service	Action
✓ 1	firewall	Internal	vse	any	any	Accept
✓ 2	Traffic to HTTP server	User	internal	HTTP Address Group	For HTTP server	Accept
✓ 3	Default Rule	Default	any			Deny

HTTP Address Group ⊗
Value:
10.20.222.34

For HTTP server ⊗
Value:
TCP:8080

Note If you select **internal** as the source, the rule is automatically updated when you configure additional internal interfaces.

Figure 10-3. Firewall rule for traffic to allow SSH into a m/c in internal network

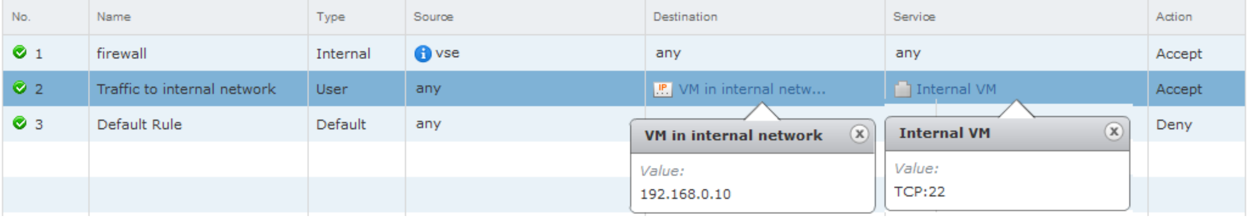

No.	Name	Type	Source	Destination	Service	Action
✓ 1	firewall	Internal	vse	any	any	Accept
✓ 2	Traffic to internal network	User	any	VM in internal netw...	Internal VM	Accept
✓ 3	Default Rule	Default	any			Deny

VM in internal network ⊗
Value:
192.168.0.10

Internal VM ⊗
Value:
TCP:22

Procedure

1 In the vSphere Web Client, navigate to **Networking & Security > NSX Edges**.

2 Double-click an NSX Edge.

3 Click the **Manage** tab and then click the **Firewall** tab.

4 Do one of the following.

Option	Description
To add a rule at a specific place in the firewall table	a Select a rule. b In the No. column, click ✎ and select **Add Above** or **Add Below**. A new any any allow rule is added below the selected rule. If the system defined rule is the only rule in the firewall table, the new rule is added above the default rule.
To add a rule by copying a rule	a Select a rule. b Click the Copy (📋) icon. c Select a rule. d In the No. column, click ✎ and select **Paste Above** or **Paste Below**.
To add a rule anywhere in the firewall table	a Click the **Add** (✚) icon. A new any any allow rule is added below the selected rule. If the system defined 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.

5 Point to the **Name** cell of the new rule and click ⊞.

6 Type a name for the new rule.

7 Point to the **Source** cell of the new rule and click ⊞ or IP.

If you clicked IP, type an IP address.

a Select an object from the drop-down and then make the appropriate selections.

If you select **vNIC Group** and then select **vse**, the rule applies to traffic generated by the NSX Edge. If you select **internal** or **external**, the rule applies to traffic coming from any internal or uplink interface of the selected NSX Edge instance. The rule is automatically updated when you configure additional interfaces. Note that firewall rules on internal interfaces do not work for a Logical Router.

If you select **IP Sets**, you can create a new IP address group. After you create the new group, it is automatically added to the source column. For information on creating an IP Set, see Create an IP Address Group.

b Click **OK**.

8 Point to the **Destination** cell of the new rule and click ⊞ or IP.

  a Select an object from the drop-down and then make the appropriate selections.

  If you select **vNIC Group** and then select **vse**, the rule applies to traffic generated by the NSX Edge. If you select **internal** or **external**, the rule applies to traffic going to any internal or uplink interface of the selected NSX Edge instance. The rule is automatically updated when you configure additional interfaces. Note that firewall rules on internal interfaces do not work for a Logical Router.

  If you select **IP Sets**, you can create a new IP address group. After you create the new group, it is automatically added to the source column. For information on creating an IP Set, see Create an IP Address Group.

  b Click **OK**.

9 Point to the **Service** cell of the new rule and click ⊞ or ▢.

- If you clicked ⊞, select a service. To create a new service or service group, click **New**. After you create the new service, it is automatically added to the Service column. For more information on creating a new service, see Create a Service.

- If you clicked ▢, select a protocol. You can specify the source port by clicking the arrow next to Advanced options. VMware recommends that you avoid specifying the source port from release 5.1 and later. Instead, you can create a service for a protocol-port combination.

Note NSX Edge only supports services defined with L3 protocols.

10 Point to the **Action** cell of the new rule and click ⊞. Make appropriate selections as described in the table below and click **OK**.

Action selected	Results in
Allow	Allows traffic from or to the specified source and destination.
Block	Blocks traffic from or to the specified source and destination.
Reject	Sends reject message for unaccepted packets. RST packets are sent for TCP packets. ICMP unreachable (administratively restricted) packets are sent for other packets.
Log	Logs all sessions matching this rule. Enabling logging can affect performance.
Do not log	Does not log sessions.
Comments	Type comments if required.
Advanced options > Match on Translated	Applies the rule to the translated IP address and services for a NAT rule
Enable Rule Direction	Indicates whether the rule is incoming or outgoing. VMware does not recommend specifying the direction for firewall rules.

11 Click **Publish Changes** to push the new rule to the NSX Edge instance.

What to do next

- Disable a rule by clicking ✓ next to the rule number in the **No.** column.
- Hide generated rules or pre rules (rules added on the centralized Firewall tab) by clicking **Hide Generated rules** or **Hide Pre rules**.
- Display additional columns in the rule table by clicking ▦ and selecting the appropriate columns.

Column Name	Information Displayed
Rule Tag	Unique system generated ID for each rule
Log	Traffic for this rule is being logged or not
Stats	Clicking 📊 shows the traffic affected by this rule (number of sessions, traffic packets, and size)
Comments	Comments for the rule

- Search for rules by typing text in the Search field.

Edit an NSX Edge Firewall Rule

You can edit only the user-defined firewall rules that were added in the Edge Firewall tab. Rules added on the centralized Firewall tab are not editable on the Edge Firewall tab.

Procedure

1. In the vSphere Web Client, navigate to **Networking & Security > NSX Edges**.
2. Double-click an NSX Edge.
3. Click the **Monitor** tab and then click the **Firewall** tab.
4. Select the rule to edit

   Note You cannot change an auto-generated rule or the default rule.

5. Make the desired changes and click **OK**.
6. Click **Publish Changes**.

Change the Priority 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 load balancer traffic from a specific IP address group, and position this rule above the LB allow traffic rule.

Procedure

1. In the vSphere Web Client, navigate to **Networking & Security > NSX Edges**.
2. Double-click an NSX Edge.

3 Click the **Monitor** tab and then click the **Firewall** tab.

4 Select the rule for which you want to change the priority.

> **Note** You cannot change the priority of auto-generated rules or the default rule.

5 Click the **Move Up** (↑) or **Move Down** (↓) icon.

6 Click **OK**.

7 Click **Publish Changes**.

Delete an NSX Edge Firewall Rule

You can delete a user-defined firewall rule that was added in the NSX Edge Firewall tab. Rules added on the centralized Firewall tab cannot be deleted here.

Procedure

1 In the vSphere Web Client, navigate to **Networking & Security > NSX Edges**.

2 Double-click an NSX Edge.

3 Click the **Monitor** tab and then click the **Firewall** tab.

4 Select the rule to delete.

> **Note** You cannot delete an auto-generated rule or the default rule.

5 Click the **Delete** (✖) icon.

Managing NAT Rules

NSX Edge provides network address translation (NAT) service to assign a public address to a computer or group of computers in a private network. Using this 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 to provide access to services running on privately addressed virtual machines.

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 private IP address or the reverse.

Prerequisites

- The translated (public) IP address must have been added to the NSX Edge interface on which you want to add the rule.
- SNAT rules are not supported on sub-interfaces.

Procedure

1 In the vSphere Web Client, navigate to **Networking & Security > NSX Edges**.

2 Double-click an NSX Edge.

3 Click the **Manage** tab and then click the **NAT** tab.

4 Click the **Add** (✚) icon and select **Add SNAT Rule**.

5 Select the interface on which to add the rule.

SNAT rules are not supported on sub-interfaces.

6 Type the original 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	

7 Type the translated (public) 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 Select **Enabled** to enable the rule.

9 Click **Enable logging** to log the address translation.

10 Click **OK** to add the rule.

11 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.

Prerequisites

The original (public) IP address must have been added to the NSX Edge interface on which you want to add the rule.

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **NSX Edges**.

3. Double-click an NSX Edge.
4. Click the **Manage** tab and then click the **NAT** tab.
5. Click the **Add** (✚) icon and select **Add DNAT Rule**.
6. Select the interface on which to apply the DNAT rule.
7. 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`	

8. Type the protocol.
9. Type the original port or port range.

Format	Example
Port number	80
Port range	80-85
`any`	

10. 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
IP address/subnet	192.0.2.0 /24
`any`	

11. Type the translated port or port range.

Format	Example
Port number	80
Port range	80-85
`any`	

12. Select **Enabled** to enable the rule.
13. Select **Enable logging** to log the address translation.
14. Click **Add** to save the rule.

Identity Firewall Overview

11

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 user logs onto in order to apply DFW rules. There are two methods IDFW uses for logon detection: Guest Introspection 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. 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.

OS Supported With IDFW

AD Supported Servers

- Windows 2012
- Windows 2008
- Windows 2008 R2

Guest OS Supported

- Windows 2012
- Windows 2008

- Windows 2008 R2
- Windows 10
- Windows 8 32/64
- Windows 7 32/64

Identity Firewall Workflow

Identity Firewall (IDFW) allows user-based distributed firewall rules (DFW).

User-based distributed firewall rules (DFW) are determined by membership in an Active Directory (AD) group membership. IDFW monitors where Active Directory users are logged into and maps the login to an IP Address, which is used by DFW to apply firewall rules. Identity Firewall requires either guest introspection framework and/or active directory event log scraping.

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. Note that you must configure one or both of these options:
   - 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. For troubleshooting Guest Introspection, see Collecting Guest Introspection Troubleshooting Data.

Working with Active Directory Domains 12

You can a 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.

This chapter includes the following topics:

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 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **NSX Managers**.

3 Click an NSX Manager in the **Name** column and then click the **Manage** tab.

4 Click the **Domain** tab and then click the **Add domain** (✚) icon.

5 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.

6. During sync, to filter out users that no longer have active accounts click **Ignore disabled users** .

7. Click **Next**.

8. In the LDAP Options page, specify the domain controller that the domain is to be synchronized with and select the protocol.

9. Edit the port number if required.

10. Enter the user credentials for the domain account. This user must be able to access the directory tree structure.

11. Click **Next**.

12. (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.

13. 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.

14. Click **Next**.

15. In the Ready to Complete page, review the settings you entered.

16. Click **Finish**.

Attention If an error message occurs stating that the Adding Domain operation failed for the entity because of a domain conflict, the workaround is to is to select Auto Merge.

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.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **NSX Managers**.
3. Click an NSX Manager in the **Name** column and then click the **Manage** tab.
4. Select the domain to be synchronized.
5. Click 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

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **NSX Managers**.
3. Click an NSX Manager in the **Name** column and then click the **Manage** tab.
4. Select a domain and then click the **Edit domain** icon.
5. 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

1. Navigate to **Start > Administrative Tools > Active Directory Users and Computers**.

2 In the navigation tree, expand the node that corresponds to the domain for which you want to enable security log access.

3 Under the node that you just expanded, select the **Builtin** node.

4 Double-click on **Event Log Readers** in the list of groups.

5 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.

Using SpoofGuard 13

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.

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.

SpoofGuard supports both IPv4 and IPv6 addresses. When using IPv4, the SpoofGuard policy supports a single IP address assigned to a vNIC. IPv6 supports multiple IP addresses assigned to a 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 address 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.

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

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 > SpoofGuard**.

2 Click the **Add** icon.

3 Type a name for the policy.

4 Select **Enabled** or **Disabled** to indicate whether the policy is enabled.

5 For **Operation Mode**, select one of the following:

Option	Description
Automatically Trust IP Assignments on Their First Use	Select this option to trust all IP assignments upon initial registration with the NSX Manager.
Manually Inspect and Approve All IP Assignments Before Use	Select this option to require manual approval of all IP addresses. All traffic to and from unapproved IP addresses is blocked.

6 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.

7 Click **Next**.

8 To specify the scope for the policy, click **Add** and select the networks, distributed port groups, or logical switches that this policy should apply to.

A port group or logical switch can belong to only one SpoofGuard policy.

9 Click **OK** and then click **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

1 In the **SpoofGuard** tab, select a policy.

Policy details are displayed below the policy table.

2 In **View**, click one of the option links.

Option	Description
Active Virtual NICs	List of all validated IP addresses
Active Virtual NICs Since Last Published	List of IP addresses that have been validated since the policy was last updated
Virtual NICs IP Required Approval	IP address changes that require approval before traffic can flow to or from these virtual machines
Virtual NICs with Duplicate IP	IP addresses that are duplicates of an existing assigned IP address within the selected datacenter
Inactive Virtual NICs	List of IP addresses where the current IP address does not match the published IP address
Unpublished Virtual NICs IP	List of virtual machines for which you have edited the IP address assignment but have not yet published

3 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 Detected IP(s)**.

Edit an IP Address

You can edit 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 the **SpoofGuard** tab, select a policy.

Policy details are displayed below the policy table.

2 In **View**, click one of the option links.

Option	Description
Active Virtual NICs	List of all validated IP addresses
Active Virtual NICs Since Last Published	List of IP addresses that have been validated since the policy was last updated
Virtual NICs IP Required Approval	IP address changes that require approval before traffic can flow to or from these virtual machines
Virtual NICs with Duplicate IP	IP addresses that are duplicates of an existing assigned IP address within the selected datacenter
Inactive Virtual NICs	List of IP addresses where the current IP address does not match the published IP address
Unpublished Virtual NICs IP	List of virtual machines for which you have edited the IP address assignment but have not yet published

3 For the appropriate vNIC, click the **Edit** icon and make appropriate changes.

4 Click **OK**.

Clear an IP Address

You clear an approved IP address assignment from a SpoofGuard policy.

Procedure

1 In the **SpoofGuard** tab, select a policy.

Policy details are displayed below the policy table.

2 In **View**, click one of the option links.

Option	Description
Active Virtual NICs	List of all validated IP addresses
Active Virtual NICs Since Last Published	List of IP addresses that have been validated since the policy was last updated
Virtual NICs IP Required Approval	IP address changes that require approval before traffic can flow to or from these virtual machines
Virtual NICs with Duplicate IP	IP addresses that are duplicates of an existing assigned IP address within the selected datacenter
Inactive Virtual NICs	List of IP addresses where the current IP address does not match the published IP address
Unpublished Virtual NICs IP	List of virtual machines for which you have edited the IP address assignment but have not yet published

3 Do one of the following.

- To clear a single IP address, click **Clear** next to the IP address.
- To clear multiple IP addresses, select the appropriate vNICs and then click **Clear Approved IP(s)**.

Virtual Private Networks (VPN) 14

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.

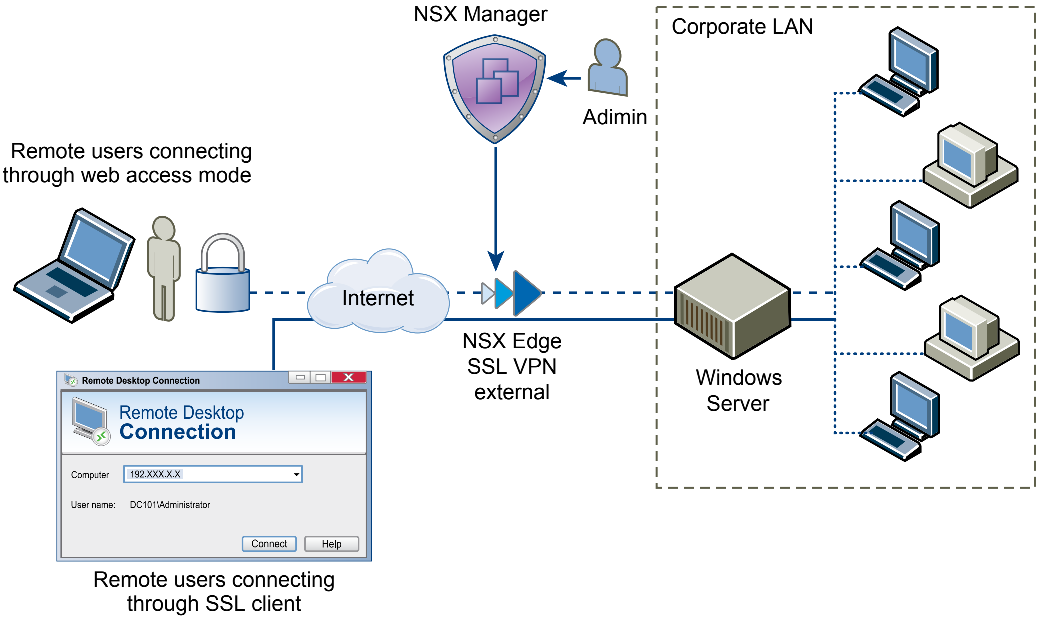
NSX Manager
Adimin
Corporate LAN
Remote users connecting through web access mode
Internet
NSX Edge SSL VPN external
Windows Server
Remote Desktop Connection
Remote Desktop Connection
Computer
192.XXX.X.X
User name:
DC101\Administrator
Connect
Help
Remote users connecting through SSL client

The following client operating systems are supported:

- Windows XP and above (Windows 8 is supported).
- Mac OS X Tiger, Leopard, Snow Leopard, Lion, Mountain Lion, Maverick, and Yosemite. These can be installed either manually or using the Java installer.
- Linux - TCL-TK is required for UI to work. If not present, Linux client can be used using CLI.

For information about troubleshooting SSL VPNs, see https://kb.vmware.com/kb/2126671.

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.

Procedure

1. Add SSL VPN-Plus Server Settings

   You must add SSL VPN server settings to enable SSL on a NSX Edge interface.

2. Add an IP Pool

   The remote user is assigned a virtual IP address from the IP pool that you add.

3. Add a Private Network

   Add the network that you want the remote user to be able to access.

4. Add Authentication

   Instead of a local user, 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.

5. Add Installation Package

   Create an installation package of the SSL VPN-Plus client for the remote user.

6. Add a User

   Add a remote user to the local database.

7. Enable the SSL VPN-Plus Service

   After configuring the SSL VPN-Plus service, enable the service for remote users to begin accessing private networks.

8. 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.

9. Install SSL Client on Remote Site

   This section describes the procedure a remote user can follow on his/her desktop after SSL VPN-Plus is configured. Windows, MAC, and Linux desktops are supported.

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 the encryption method.
6. (Optional) From the Server Certificates table, select the server certificate that you want to add.
7. Click **OK.**

Add an IP Pool

The remote user is assigned a virtual IP address from the IP pool that you add.

Procedure

1. In the **SSL Vpn-Plus** tab, select **IP Pools** from the left panel.
2. Click the **Add** (✚) icon.
3. Type the begin and end IP address for the IP pool.
4. Type the netmask of the IP pool.
5. Type the IP address which is to add the routing interface in the NSX Edge gateway.
6. (Optional) Type a description for the IP pool.
7. Select whether to enable or disable the IP pool.
8. (Optional) In the **Advanced** panel, type the DNS name.
9. (Optional) Type the secondary DNS name.
10. Type the connection-specific DNS suffix for domain based host name resolution.
11. Type the WINS server address.
12. Click **OK**.

Add a Private Network

Add the network that you want the remote user to be able to access.

Procedure

1. In the **SSL Vpn-Plus** tab, select **Private Networks** from the left panel.
2. Click the **Add** (✚) icon
3. Type the private network IP address.
4. Type the netmask of the private network.
5. (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.

   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

Instead of a local user, 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 14-1. AD 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 authentication server.
Port	Displays default port name. Edit if required.
Timeout	Period in seconds within which the AD server must respond.
Status	Select **Enabled** or **Disabled** 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 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.
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 **`sAMAccountName`**.
Search Filter	Filter values by which the search is to be limited. The search filter format is *attribute operator value*.

Table 14-1. 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 14-2. 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 **Enabled** or **Disabled** 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 `sAMAccountName`.
Search Filter	Filter values by which the search is to be limited. The search filter format is *attribute operator value*.
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

Table 14-3. 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 **Enabled** or **Disabled** 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

Table 14-4. RSA-ACE authentication server options

Option	Description
Timeout	Period in seconds within which the AD server must respond.
Configuration File	Click **Browse** to select the `sdconf.rec` file that you downloaded from the RSA Authentication Manager.
Status	Select **Enabled** or **Disabled** to indicate whether the server is enabled.
Source IP Address	IP address of the NSX Edge interface through which the RSA server is accessible.
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 14-5. Local authentication server options

Option	Description
Enable password policy	If selected, defines a password policy. Specify the required values.
Enable password policy	If selected, defines an account lockout policy. Specify the required values. 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. 2 In Retry Duration, type the time period in which the remote user's account gets locked on unsuccessful login attempts. 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. 3 In Lockout Duration, type the time period for which the user account remains locked. After this time, the account is automatically unlocked.
Status	Select **Enabled** or **Disabled** 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.

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
Start client on logon	The SSL VPN client is started when the remote user logs on to his system.
Allow remember password	Enables the option.
Enable silent mode installation	Hides installation commands from remote user.
Hide SSL client network adapter	Hides the VMware SSL VPN-Plus Adapter, which is installed on the remote user's computer along with the SSL VPN installation package.
Hide client system tray icon	Hides the SSL VPN tray icon which indicates whether the VPN connection is active or not.
Create desktop icon	Creates an icon to invoke the SSL client on the user's desktop.
Enable silent mode operation	Hides the pop-up that indicates that installation is complete.
Server security certificate validation	The SSL VPN client validates the SSL VPN server certificate before establishing the secure connection.

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 the Enable icon.

   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 Client on Remote Site

This section describes the procedure a remote user can follow on his/her desktop after SSL VPN-Plus is configured. Windows, MAC, and Linux desktops are supported.

Procedure

1 On the client site, the remote user can type (`https://`*ExternalEdgeInterfaceIP*`/sslvpn-plus/`) in a browser window where *ExternalEdgeInterfaceIP* is the IP address of the Edge external interface where you enabled SSL VPN-Plus.

2 Login to the portal using the user's credentials.

3 Click Full Access tab.

The SSL client is downloaded.

4 Login to the SSL client with the credentials specified in the Users section.

The SSL VPN server certificate is validated depending on the client operating system.

- Windows client

  Windows client is authenticated if the **Server security certificate validation** option was selected when the installation package was created.

- Linux client

  The SSL VPN Linux client validates the server certificate against Firefox's certificate store by default from NSX vSphere version 6.1.3 onwards. If server certificate validation fails, you are prompted to contact your system administrator. If server certificate validation succeeds, a log in prompt is displayed.

  Adding a trusted CA to the trust store i.e Firefox's certificate store is independent of SSL VPN work flow.

- OS X client

The SSL VPN OS X client validates the server certificate against Keychain, a database used to store certificates on OS X, by default from NSX vSphere version 6.1.3 onwards. If server certificate validation fails, you are prompted to contact your system administrator. If server certificate validation succeeds, a log in prompt is displayed.

Adding a trusted CA to the trust store i.e Keychain is independent of SSL VPN work flow.

The remote user can now access the private network.

Configure Web Access SSL VPN-Plus

In web access mode, a remote user can access private networks without a hardware or software SSL client.

Procedure

1 Create a Web Resource

Add a server that the remote user can connect to via a web browser.

2 Add a User

Add a remote user to the local database.

3 Add Authentication

Instead of a local user, 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.

4 Add SSL VPN-Plus Server Settings

You must add SSL VPN server settings to enable SSL on a NSX Edge interface.

5 Enable the SSL VPN-Plus Service

After configuring the SSL VPN-Plus service, enable the service for remote users to begin accessing private networks.

6 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.

Create a Web Resource

Add a server that the remote user can connect to via a web browser.

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **NSX Edges**.

3 Double-click an NSX Edge.

4 Click the **Manage** tab and then click the **SSL VPN-Plus** tab.

5. Select **Web Resource** from the left panel.
6. Click the **Add** (✚) icon.
7. Type a name for the web resource.
8. Type the URL of the web resource that you want the remote user to access.
9. Depending on whether the remote user wants to read from or write to the web resource, select the **HTTPMethod** and type the GET or POST call.
10. Type the description for the web resource. This description is displayed on the web portal when the remote user accesses the web resource.
11. Select **Enable** to enable the web resource. The web resource must be enabled for the remote user to access it.

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**.

Add Authentication

Instead of a local user, 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 14-6. AD 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 authentication server.
Port	Displays default port name. Edit if required.
Timeout	Period in seconds within which the AD server must respond.
Status	Select **Enabled** or **Disabled** 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 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.
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 `sAMAccountName`.
Search Filter	Filter values by which the search is to be limited. The search filter format is *attribute operator value*.

Table 14-6. 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 14-7. 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 **Enabled** or **Disabled** 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 `sAMAccountName`.
Search Filter	Filter values by which the search is to be limited. The search filter format is *attribute operator value*.
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

Table 14-8. 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 **Enabled** or **Disabled** 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

Table 14-9. RSA-ACE authentication server options

Option	Description
Timeout	Period in seconds within which the AD server must respond.
Configuration File	Click **Browse** to select the `sdconf.rec` file that you downloaded from the RSA Authentication Manager.
Status	Select **Enabled** or **Disabled** to indicate whether the server is enabled.
Source IP Address	IP address of the NSX Edge interface through which the RSA server is accessible.
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 14-10. Local authentication server options

Option	Description
Enable password policy	If selected, defines a password policy. Specify the required values.
Enable password policy	If selected, defines an account lockout policy. Specify the required values. 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. 2 In Retry Duration, type the time period in which the remote user's account gets locked on unsuccessful login attempts. 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. 3 In Lockout Duration, type the time period for which the user account remains locked. After this time, the account is automatically unlocked.
Status	Select **Enabled** or **Disabled** 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.

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 the encryption method.

6 (Optional) From the Server Certificates table, select the server certificate that you want to add.

7 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 the ⏻ Enable icon.

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**.

SSL VPN-Plus Logs

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/SSL VPN Client/`.

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.
Force virtual keyboard	Allow remote users to enter web or client login information only via the virtual keyboard.
Randomize keys of virtual keyboard	Make the virtual keyboard keys random.

Select	To
Enable forced timeout	Disconnect the remote user after the specified timeout period is over. Type the timeout period in minutes.
Session idle timeout	If there is no activity on the user session for the specified period, end the user session after that period is over.
User notification	Type a message to be displayed to the remote user after he logs in.
Enable public URL access	Allow remote user to access any site which is not configured (and not listed on web portal) by administrator.

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 **Monitor** 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 select the image file for the remote user's logo.

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 If desired, change the client banner.

9 Click **OK**.

Working with IP Pools

You can edit or delete an IP pool.

For information on adding an IP pool, see Configure Network Access SSL VPN-Plus or Configure Web 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 or Configure Web 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**.

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 or Configure Web 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 or Configure Web 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, IP unicast traffic, and no dynamic routing protocol are supported between the NSX Edge instance and remote VPN routers.

Behind each remote VPN router, you can configure multiple subnets to connect to the internal network behind an NSX Edge through IPSec tunnels.

Note Subnets and the internal network behind a NSX Edge must have address ranges that do not overlap.

If the local and remote peer across an IPsec VPN have overlapping IP addresses, traffic forwarding across the tunnel might be not consistent depending on whether local connected routes and auto-plumbed routes exist.

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 routers use this public address to access the NSX Edge instance.

You can place remote VPN routers behind a NAT device as well. You must provide the VPN native address and the VPN Gateway ID to set up the tunnel. On both ends, static one-to-one NAT is required for the VPN address.

The number of tunnels needed is defined by the number of local subnets multiplied by the number of peer subnets. For example, if there are 10 local subnets and 10 peer subnets you need 100 tunnels. The maximum number of tunnels supported is determined by the ESG size, as shown below.

Table 14-11. Number of IPSec Tunnels per ESG

ESG	Number of IPSec Tunnels
Compact	512
Large	1600
Quad-Large	4096
X-Large	6000

The following IPSec VPN algorithms are supported:

- AES (AES128-CBC)

- AES256 (AES256-CBC)
- Triple DES (3DES192-CBC)
- AES-GCM (AES128-GCM)
- DH-2 (Diffie–Hellman group 2)
- DH-5 (Diffie–Hellman group 5)

For IPSec VPN configuration examples, see Chapter 24 NSX Edge VPN Configuration Examples.

For IPSec VPN troubleshooting, see https://kb.vmware.com/kb/2123580.

Configuring IPSec VPN Service

You can set up an NSX Edge tunnel between a local subnet and a peer subnet.

Note If you connect to a remote site via IPSec VPN, the IP address of that site cannot be learnt by Dynamic Routing on the Edge uplink.

1 Enable IPSec VPN Service

   You must enable the IPSec VPN service for traffic to flow from the local subnet to the peer subnet.

2 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.

3 Specify Global IPSec VPN Configuration

   This enables IPSec VPN on the NSX Edge instance.

4 Enable Logging for IPSec VPN

   You can enable logging of all IPSec VPN traffic.

5 Configure IPSec VPN Parameters

   You must configure at least one external IP address on the NSX Edge to provide IPSec VPN service.

Enable IPSec VPN Service

You must enable the IPSec VPN service for traffic to flow from the local subnet to the peer subnet.

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **NSX Edges**.

3 Double-click an NSX Edge.

4 Click the **Manage** tab and then click the **VPN** tab.

5 Click **IPSec VPN**.

6 Click **Enable**.

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, generate a CSR, copy the privacy-enhanced mail (PEM) file content, and save it in a file in `req/edge1.req`.

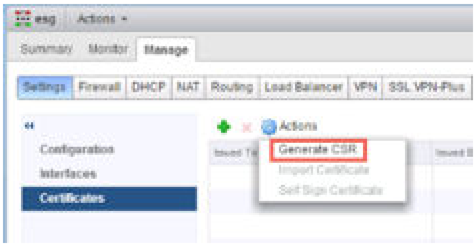

See Configure a CA Signed Certificate.

6 Run the 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 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 Upload the CA certificate in the file `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 uploaded CA certificate and save the configuration.

13 On the **Certifcate** tab, click the uploaded certificate 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.

Check the status by clicking **Show IPsec Statistics**. Click the channel to see the tunnel status. Both the channel and tunnel status should be green.

Specify Global IPSec VPN Configuration

This enables 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 VPNs. They can only be used in load balancing and SSL VPNs.

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **NSX Edges**.

3 Double-click an NSX Edge.

4 Click the **Manage** tab and then click the **VPN** tab.

5 Click **IPSec VPN**.

6 Click **Change** next to Global configuration status.

7 Type a global pre-shared key for those sites whose peer endpoint is set to any and select **Display shared key** to display the key.

8. Select Enable certificate authentication and select the appropriate certificate.
9. Click **OK**.

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** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click the **Manage** tab and then click the **VPN** tab.
5. Click **IPSec VPN**.
6. Click ▸ next to **Logging Policy** and click **Enable logging** to log the traffic flow between the local subnet and peer subnet and select the logging level.
7. Select the log level and 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** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click the **Monitor** tab and then click the **VPN** tab.
5. Click **IPSec VPN**.
6. Click the **Add** (✚) icon.
7. Type a name for the IPSec VPN.
8. Type the IP address of the NSX Edge instance in **Local Id**. This will be the peer Id on the remote site.
9. Type the IP address 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.
10. Type the subnets to share between the sites in CIDR format. Use a comma separator to type multiple subnets.

11 Type the Peer Id to uniquely identify the peer site. For peers using certificate authentication, this ID must be the common name in the peer's certificate. For PSK peers, this ID can be any string. VMware recommends that you use the public IP address of the VPN or a FQDN for the VPN service as the peer ID.

12 Type the IP address of the peer site in Peer Endpoint. If you leave this blank, NSX Edge waits for the peer device to request a connection.

13 Type the internal IP address of the peer subnet in CIDR format. Use a comma separator to type multiple subnets.

14 Select the Encryption Algorithm.

15 In Authentication Method, select one of the following:

Option	Description
PSK (Pre Shared Key)	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.
Certificate	Indicates that the certificate defined at the global level is to be used for authentication.

16 Type the shared key in if anonymous sites are to connect to the VPN service.

17 Click **Display Shared Key** to display the key on the peer site.

18 In Diffie-Hellman (DH) Group, select the cryptography scheme that will allow the peer site and the NSX Edge to establish a shared secret over an insecure communications channel.

19 In Extension, type one of the following:

- securelocaltrafficbyip=*IPAddress* to re-direct Edge's local traffic over the IPSec VPN tunnel. This is the default value.
- passthroughSubnets=*PeerSubnetIPAddress* to support overlapping subnets .

20 Click **OK**.

NSX Edge creates a tunnel from the local subnet to the peer subnet.

What to do next

Enable the IPSec VPN service.

Edit IPSec VPN Service

You can edit an IPSec VPN service.

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **NSX Edges**.

3 Double-click an NSX Edge.

4. Click the **Monitor** tab and then click **VPN** tab.
5. Click **IPSec VPN**.
6. Select the IPSec service that you want to edit.
7. Click the **Edit** (✎) icon.
8. Make the appropriate edits.
9. Click **OK**.

Disable IPSec Service

You can disable an IPSec service.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click the **Monitor** tab and then click **VPN** tab.
5. Click **IPSec VPN**.
6. Select the IPSec service that you want to disable.
7. Click the **Disable** (⊘) icon.

Delete IPSec Service

You can delete an IPSec service.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click the **Monitor** tab and then click **VPN** tab.
5. Click **IPSec VPN**.
6. Select the IPSec service that you want to delete.
7. Click the **Delete** (✖) icon.

L2 VPN Overview

L2 VPN allows you to configure a tunnel between two sites. Virtual machines remain on the same subnet in spite of being moved between these sites, which enables you to extend your datacenter. An NSX Edge at one site can provide all services to virtual machines on the other site.

In order to create the L2 VPN tunnel, you configure an L2 VPN server and L2 VPN client.

Configuring L2 VPN

To stretch your network using L2 VPN, you configure an L2 VPN server (destination Edge) and an L2 VPN client (source Edge). You must then enable the L2 VPN service on both the server and the client.

Prerequisites

A sub interface must have been added on a trunk interface of the NSX Edge. See Add a Sub Interface.

Procedure

1 L2 VPN Best Practices

   Configuring L2 VPN according to best practices can avoid problems such as looping and duplicate pings and responses.

2 Configure L2 VPN Server

   The L2 VPN server is the destination NSX Edge to which the client is to be connected.

3 Add Peer Sites

   You can connect multiple sites to the L2 VPN server.

4 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 will fail.

5 Configure L2 VPN Client

   The L2 VPN client is the source NSX Edge that initiates communication with the destination Edge (L2 VPN server).

6 Enable L2 VPN Service on Client

   You must enable the L2 VPN service on the L2 VPN client (source NSX Edge).

L2 VPN Best Practices

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
    a Deploy the Edges and the VMs on separate ESXi hosts.
    b Configure the Teaming and Failover Policy for the Distributed Port Group associated with the Edge’s Trunk vNic as follows:
        1 Load balancing as “Route based on originating virtual port.”
        2 Configure only one uplink as Active and the other uplink as Standby.
    c Configure the teaming and failover policy for the distributed port group associated with the VMs as follows:
        1 Any teaming policy is okay.
        2 Multiple active uplinks can be configured.
    d Configure Edges to use sink port mode and disable promiscuous mode on the trunk vNic.
- Option 2: Edges and VMs on the same ESXi 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 NSX 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 you 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 that represents an Edge gateway with the purpose of acting as an L2 VPN client to be deployed on a host that is not managed by NSX.

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, we recommend using sink port mode in the L2 VPN standalone NSX Edge configuration.

Prerequisites

You need the port number of the trunk vNIC of the standalone edge.

Procedure

1 Retrieve the dvsUuid value:

   a Go to vCenter Mob UI at https://<vc-ip>/mob?vmodl=1.

   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.

2 Modify the selectionSet in the vCenter managed object browser (MOB).

   a Log in to the vCenter Mob UI at https://<vc-ip>/mob?vmodl=1.

   b Click **content**.

   c Click **DVSManager**.

   d Click **updateOpaqueDataEx**.

   e In the **selectionSet** value box paste the following XML block:

```
<selectionSet xsi:type="DVPortSelection">
     <dvsUuid>c2 1d 11 50 6a 7c 77 68-e6 ba ce 6a 1d 96 2a 15</dvsUuid> <!--example only-->
     <portKey>393</portKey>  <!--port number of the DVPG where SINK to be set-->
</selectionSet>
```

   Use the dvsUuid value that you retrieved from the vCenter MOB.

f On the opaqueDataSpec value box paste one of the following XML blocks:

To enable SINK port:

```
<opaqueDataSpec>
                <operation>edit</operation>
                <opaqueData>
                     <key>com.vmware.etherswitch.port.extraEthFRP</key>
                     <opaqueData
xsi:type="vmodl.Binary">AAABAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAA=</opaqueData>
                 </opaqueData>
</opaqueDataSpec>
```

To disable SINK port:

```
<opaqueDataSpec>
                <operation>edit</operation>
                <opaqueData>
                     <key>com.vmware.etherswitch.port.extraEthFRP</key>
                     <opaqueData
xsi:type="vmodl.Binary">AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAA=</opaqueData>
                 </opaqueData>
</opaqueDataSpec>
```

g Set the isRuntime boolean to **false**.

h Click **Invoke Method**.

Configure L2 VPN Server

The L2 VPN server is the destination NSX Edge to which the client is to be connected.

Procedure

1. In the **L2 VPN** tab, select **Server** and click **Change**.
2. In **Listener IP**, type the primary or secondary IP address of an external interface of the NSX Edge.
3. The default port for the L2 VPN service is 443. Edit this if required.
4. Select the encryption algorithm for communication between the server and the client.
5. Select the certificate to be bound to SSL VPN server.
6. Click **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 In the L2 VPN tab, ensure that the **L2 VPN Mode** is **Server**.

2 In **Site Configuration Details**, click the **Add** icon.

3 Type a unique name for the peer site.

4 Type the user name and password with which the peer site is to be authenticated. User credentials on the peer site should be the same as those on the client side.

5 In **Stretched Interfaces**, click **Select Sub Interfaces** to select the sub interfaces to be stretched with the client.

   a In Select Object, select the trunk interface for the Edge.

      Sub interfaces configured on the trunk vNIC are displayed.

   b Double-click the sub interfaces to be stretched.

   c Click **OK**.

6 If the default gateway for virtual machines is same across the two sites, type the gateway IP addresses for which the traffic should be locally routed or for which traffic is to be blocked over the tunnel in **Egress Optimization Gateway Address**.

7 Click **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 will fail.

Procedure

1 For the destination NSX Edge, navigate to **Manage > VPN > L2 VPN**.

2 In **L2VPN Service Configuration**, click **Enable**.

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 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. In the L2 VPN tab, set the **L2 VPN Mode** to **Client** and click **Change**.
2. Type the address of the L2 VPN server to which this client is to be connected. The address can be the host name or IP address.
3. If required, edit the default port to which the L2 VPN client should connect to.
4. Select the encryption algorithm for communicating with the server.
5. In **Stretched Interfaces**, click **Select Sub Interfaces** to select the sub interfaces to be stretched to the server.

   a. In **Select Object**, select the trunk interface for the Edge.

      Sub interfaces configured on the trunk vNIC are displayed.

   b. Double-click the sub interfaces to be stretched.

   c. Click **OK**.
6. Type a description.
7. In **Egress Optimization Gateway Address**, type the gateway IP address of the sub interfaces or the IP addresses to which traffic should not flow over the tunnel.
8. In **User Details**, type the user credentials to get authenticated at the server..
9. Click the **Advanced** tab.

   If the client NSX Edge does not have direct access to the internet and needs to reach the source (server) NSX Edge via a proxy server, specify **Proxy Settings**.
10. To enable only secure proxy connections, select **Enable Secure Proxy**.
11. Type the proxy server address, port, user name, and password.
12. To enable server certificate validation, select **Validate Server Certificate** and select the appropriate CA certificate.
13. Click **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. For the source NSX Edge, navigate to **Manage > VPN > L2 VPN**.
2. In **L2VPN Service Configuration**, click **Enable**.

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.
- If a trunk vNic backed by standard portgroup is being stretched, enable L2 VPN traffic manually by the following steps:

  a. Set **Promiscuous mode** to **Accept**.

  b. Set **Forged Transmits** to **Accept**.

  For more information, see *ESXi and vCenter Server 5.5 Documentation*.

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.

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. Enabling a sink port is the recommended best practice.

  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 datacenter 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 should use. Click **Next**.
   - The Public interface is the uplink interface.
   - The Trunk interface is used to create sub-interfaces for the networks that will be stretched. Connect this interface to the trunk port group you created.
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. This should match the cipher used on the L2VPN server.

   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. Type the L2 VPN server address and port.

   h. Type the user name and password with which the peer site is to be authenticated.

   i. In Sub Interfaces VLAN (Tunnel ID), type VLAN ID(s) of the network(s) 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, you can 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 would not be allowed: 10(100)-14(104). You would need to rewrite this as 10(100),11(101),12(102),13(103),14(104).

   j. If the standalone NSX Edge does not have direct access to the internet and needs to reach the source (server) NSX Edge via a proxy server, type the proxy address, port, user name, and password.

   k. If a Root CA is available, you can paste it in to the Certificate section.

   l. 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 virtual machine.

Note the trunk vNIC port number and configure a sink port. See Configure a Sink Port .

Make any further configuration changes with the standalone edge command line interface. See the *NSX Command Line Interface Reference*.

View L2 VPN Statistics

You can view L2 VPN statistics such as tunnel status, bytes sent and received etc. for the source and destination NSX Edge.

Procedure

1 In the L2 VPN tab. ensure that the **L2 VPN Mode** is **Client**.

2 Click **Fetch Status** and expand **Tunnel Status**.

If the L2 VPN server has multiple peer sites, statistics are displayed for all the peer sites.

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.

Logical Load Balancer 15

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 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.

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
- 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.

NSX load balancer supports the layer 4 and layer 7 load balancing engines. The layer 4 load balancer is packet-based and layer 7 load balancer is socket-based.

A packet-based load balancing is implemented on the TCP and UDP layer. Packet-based load balancing does not stop the connection or buffer the whole request, instead 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. You can select Acceleration Enable in both the global configuration and relevant virtual server configuration to enable packet-based load balancing.

A 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 server-facing connection. Socket-based load balancing is the default mode for TCP, HTTP, and HTTPS virtual servers.

The key concepts of the NSX load balancer are, virtual server, server pool, server pool member, and service monitor.

Virtual Server Abstract of an application service, represented by a unique combination of IP, port, and protocol 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.

You begin by setting global options for the load balancer. You now create a server pool consisting of backend server members and associate a service monitor with the pool to manage and share the backend servers efficiently.

You then 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. The application profile 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 to define health check parameters for the load balancer.

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, or source IP hash.

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.

- Configure Load Balancer Service

  You can specify global load balancer configuration parameters.

- 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 backend servers flexibly and efficiently. A pool manages load balancer distribution methods and has a service monitor attached to it for health check parameters.

- Create an Application Profile

  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. Using profiles enhances your control over managing network traffic, and makes traffic-management tasks easier and more efficient.

- Add an Application Rule

  You can write an application rule to directly manipulate and manage application traffic.

- Add Virtual Servers

  Add an NSX Edge internal or uplink interface as a virtual server.

Configure Load Balancer Service

You can specify global load balancer configuration parameters.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click **Manage** and then click the **Load Balancer** tab.
5. Click **Edit**.

6 Select the check boxes next to the options to enable.

Option	Description
Enable Load Balancer	Allows the NSX Edge load balancer to distribute traffic to internal servers for load balancing.
Enable Acceleration	When enabled globally, each virtual IP uses the faster L4 LB engine rather than L7 LB engine. The L4 TCP VIP is processed before the Edge Firewall so no Allow firewall rule is required. L7 HTTP/HTTPS VIPs are processed after the Edge Firewall. If Enable Acceleration is selected, an Edge Firewall rule must allow access to the L7 HTTP/HTTPS VIP. If the Enable Acceleration flag is selected with L4 TCP VIP, and the server pool is in non-transparent mode, an SNAT rule is added. Therefore, make sure that Firewall is enabled on NSX Edge. If the Enable Acceleration flag is deselected with L4 TCP VIP and the firewall is enable, then the Edge Firewall rule must allow access to the L7 HTTP/HTTPS VIP.
Logging	NSX Edge load balancer collects traffic logs. 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 `show log follow` command to list the load balancing logs. Debug and Info options log end-user requests. Warning, Error, and Critical options do not log end-users requests.
Enable Service Insertion	Allows the load balancer to work with third party vendor services. If you have a third party vendor load balancer service deployed in your environment, see Using a Partner Load Balancer.

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

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **NSX Edges**.

3 Double-click an NSX Edge.

4 Click **Manage** and then click the **Load Balancer** tab.

5 In the left navigation panel, click **Service Monitoring**.

6 Click the **Add** (✚) icon.

7 Enter a name for the service monitor.

8 Enter the interval in seconds at which a server is to be pinged.

9 Enter the number of times the server must be pinged before it is declared down.

10 Enter the maximum time in seconds within which a response from the server must be received.

11 Select the way in which to send the health check request to the server from the drop-down menu.

12 For HTTP and HTTPS traffic, perform the steps below.

a Enter the string that the monitor expects to match in the status line of HTTP response in the Expected section.

For example, 200,301,302,401.

b Select the method to detect server status from the drop-down menu.

c Enter the URL to be used in the sample request.

d If you select the POST method, enter the data to be sent.

13 Enter the string to be matched in the response content in the Receive section.

If the string in the Expected section is not matched, the monitor does not try to match the Receive content.

14 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.

```
<extension>eregi="(OK1|OK2)"</extension>
```

Refer to the table for supported protocol extensions.

Table 15-1. 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=*STRING*	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 criti 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=*INTEGER*	Close connection once more than the specified number of bytes are received.

Table 15-1. Extensions for TCP Protocol (Continued)

Monitor Extension	Description
delay=*INTEGER*	Seconds to wait between sending string and polling for response.
certificate=*INTEGER*[,*INTEGER*]	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).
ssl-version=3	Force SSL handshake using sslv3. sslv3 and tlsv1 are disabled in the health check option by default.
ssl-version=10	Force SSL handshake using tls 1.0.
ssl-version=11	Force SSL handshake using tls 1.1.
ssl-version=12	Force SSL handshake using tls 1.2.
ciphers='ECDHE-RSA-AES256-GCM-SHA384'	Display ciphers used in HTTPS health check.
warning=DOUBLE	Response time in seconds to result in warning status.
critical=DOUBLE	Response time in seconds to result in critical status.

Table 15-2. Extensions for HTTP/HTTPS Protocol

Monitor Extension	Description
no-body	Do not wait for document body: stop reading after headers. Note that this still does an HTTP GET or POST, not a HEAD.
max-age=*SECONDS*	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.
content-type=*STRING*	Specify Content-Type header media type in POST calls.
linespan	Allow regex to span newlines (must precede -r or -R).
regex=*STRING* or ereg=*STRING*	Search page for regex STRING.
eregi=*STRING*	Search page for case-insensitive regex STRING.
invert-regex	Return CRITICAL if found, OK if not.
proxy-authorization=*AUTH_PAIR*	Username:password on proxy-servers with basic authentication.
useragent=*STRING*	String to be sent in HTTP header as `User Agent`.
header=*STRING*	Any other tags to be sent in HTTP header. Use multiple times for additional headers.
onredirect=ok|warning|critical|follow|sticky|stickyport	How to handle redirected pages. `sticky` is like follow but stick to the specified IP address. `stickyport` also ensures port stays the same.
pagesize=*INTEGER:INTEGER*	Minimum page size required (bytes) : Maximum page size required (bytes).
warning=DOUBLE	Response time in seconds to result in warning status.
critical=DOUBLE	Response time in seconds to result in critical status.

Table 15-3. Extensions for HTTPS Protocol

Monitor Extension	Description
sni	Enable SSL/TLS hostname extension support (SNI).
certificate=*INTEGER*	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=AUTH_PAIR	Username:password on sites with basic authentication.

15 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 backend servers flexibly and efficiently. A 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** and then click **NSX Edges**.

3 Double-click an NSX Edge.

4 Click **Manage** and then click the **Load Balancer** tab.

5 In the left navigation panel, click **Pools**.

6 Click the **Add** (✚) icon.

7 Type a name and description for the load balancer pool.

8 Select the algorithm balancing method for each enabled service.

Option	Description
IP-HASH	Selects a server based on a hash of the source IP address and the total weight of all the running servers. Algorithm parameters are disabled for this option.
LEASTCONN	Distributes client requests to multiple servers based on the number of connections already on the server. New connections are sent to the server with the fewest connections. Algorithm parameters are disabled for this option.
ROUND_ROBIN	Each server is used in turn according to the weight assigned to it. This is the smoothest and fairest algorithm when the server's processing time remains equally distributed. Algorithm parameters are disabled for this option.

Option	Description
URI	The left part of the URI (before the question mark) is hashed and divided by the total weight of the running servers. The result designates which server receives the request. This ensures that a URI is always directed to the same server as long as no server goes up or down. The URI algorithm parameter has two options `uriLength=<len>` and `uriDepth=<dep>`. The length parameter range should be 1<=len<256. The depth parameter range should be 1<=dep<10. 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. 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.
HTTPHEADER	HTTP header name is looked up in each HTTP request. The header name in parenthesis 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. The HTTPHEADER algorithm parameter has one option `headerName=<name>`. For example, you can use **`host`** as the HTTPHEADER algorithm parameter.
URL	URL parameter specified in the argument is looked up in the query string of each HTTP GET request. 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. If no value or parameter is found, then a round robin algorithm is applied. The URL algorithm parameter has one option `urlParam=<url>`.

9 (Optional) Select an existing default or custom monitor from the **Monitors** drop-down menu.

10 Add members to the pool.

a Click the **Add** icon (✚).

b Enter the name and IP address of the server member or click **Select** to assign grouping objects.

The grouping objects can be either vCenter or NSX.

c Enter the port where the member is to receive traffic on 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.

d Enter the proportion of traffic this member is to handle in the Weight section.

e Enter the maximum number of concurrent connections the member can handle.

If the incoming requests goes higher than the maximum, they are queued and wait for a connection be released.

f Enter the minimum number of concurrent connections a member must always accept.

g Click **OK**.

11. Check **Transparent** to make client IP addresses visible to the backend servers.

    If Transparent is not selected (default value), backend servers see the traffic source IP address as a Load balancer internal IP address. If Transparent is selected, 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.

12. Click **OK**.

Create an Application Profile

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. Using profiles enhances your control over managing network traffic, and makes traffic-management tasks easier and more efficient.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click **Manage** and then click the **Load Balancer** tab.
5. In the left navigation panel, click **Application Profiles**.
6. Click the **Add** (✚) icon.
7. Type a name for the profile and select the traffic type for which you are creating the profile from the drop-down menu.

Traffic Type	Persistence Method Supported
TCP	Source IP, MSRDP
HTTP	Cookie, Source IP
HTTPS	Cookie, SSL Session ID (SSL Passthrough enabled) , Source IP
UDP	Source IP

8. Type the URL to which you want to redirect HTTP traffic.

   For example, you can direct traffic from `http://myweb.com` to `https://myweb.com`.

9 Specify persistence type for the profile from the drop-down menu.

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.

- Select **Cookie** persistence to insert a unique cookie to identify the session the first time a client accesses the site.

  The cookie is referred in subsequent requests to persist the connection to the appropriate server.

- Select **Source IP** persistence to track sessions based on the source IP address.

  When a client requests a connection to a virtual server that supports source address affinity persistence, the load balancer checks to see if that client previously connected, and if so, returns the client to the same pool member.

- Select Microsoft Remote Desktop Protocol **MSRDP** persistence to maintain persistent sessions between Windows clients and servers that are running the Microsoft Remote Desktop Protocol (RDP) service.

  The recommended scenario for enabling MSRDP persistence is to create a load balancing pool that consists of members running Windows Server 2003 or Windows Server 2008, where all members belong to a Windows cluster and participate in a Windows session directory.

10 Type a cookie name and select the mode by which the cookie should be inserted.

Option	Description
Insert	NSX Edge sends a cookie. If the server sends one or more cookies, the client receives an extra cookie (the server cookie(s) + the Edge cookie). If the server does not send a cookie, the client receives the Edge cookie.
Prefix	This option is selected if your client does not support more than one cookie. **Note** 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 (as a prefix) its cookie information in the server cookie value. This cookie added information is removed when NSX Edge sends it to the server.
App Session	The server does not send a cookie. Instead, it sends the user session information as a URL. For example, http://mysite.com/admin/UpdateUserServlet;jsessionid=OI24B9ASD7BSSD, where `jsessionid` is the user session information and is used for the persistence. It is not possible to see the App Session persistence table for troubleshooting.

11 Enter the persistence expiration time in seconds.

The default value of persistence is five minutes.

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 of time, even if the existing connections are still alive.

12 (Optional) Create an application profile for HTTPS traffic.

Supported HTTPS traffic pattern.

- SSL Offloading - Client -> HTTPS -> LB (terminate SSL) -> HTTP -> server
- SSL Proxy - Client -> HTTPS -> LB (terminate SSL) -> HTTPS -> server
- SSL Passthrough - Client -> HTTPS-> LB (SSL passthrough) -> HTTPS -> server
- Client -> HTTP-> LB -> HTTP -> servers

a (Optional) Check **Insert X-Forwarded-For HTTP header** to identify the originating IP address of a client connecting to a Web server through the load balancer.

b Check **Configure Service Certificate** to select the applicable service certificate, CA certificates, and CRLs used to terminate the HTTPS traffic from the client on the load balancer in the **Virtual Server Certificates** tab.

This is only required if the Client -> LB connection is HTTPS.

c (Optional) Check **Enable Pool Side SSL** to enable the HTTPS communication between the load balancer and the backend servers.

You can use the pool side SSL to configure End-to-End SSL.

d (Optional) Check **Configure Service Certificate** to select the applicable service certificate, CA certificates, and CRLs used to authenticate the load balancer from the server side in the **Pool Certificates** tab.

This is only required for the pattern Client -> HTTPS -> LB -> HTTPS -> servers.

You can configure service certificate if the NSX Edge load balancer has CA certificate and CRL already configured and needs to verify service certificate from backend servers. This option can also be used to provide load balancer certificate to backend server if the backend server needs to verify the load balancer side service certificate.

13 Enter the cipher algorithms or cipher suite negotiated during the SSL/TLS handshake.

For example, you can allow only `3DES` cipher suites to be used.

14 Specify whether client authentication is to be ignored or required from the drop-down menu.

If you select required, the client must provide a certificate after the request or the handshake is aborted.

15 Click **OK**.

Add an Application Rule

You can write an application rule to directly manipulate and manage application traffic.

For application rule examples, see Application Rule Examples.

Procedure

1 Log in to the vSphere Web Client.

2. Click **Networking & Security** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click **Manage** and then click the **Load Balancer** tab.
5. In the left navigation panel, click **Application Rules**.
6. Click the **Add** (✚) icon.
7. Type the name and script for the rule.

   For information on the application rule syntax, see http://cbonte.github.io/haproxy-dconv/configuration-1.5.html.
8. Click **OK**.

Application Rule Examples

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 a 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(mstshash) 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

# Provde 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_
```

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.

```
acl block_url_list path_beg -i /private /finance
block 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 authent_url url /authent.php
acl cookie_present hdr_sub(cookie) cookie1=
redirect prefix /authent.php if !authent_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 prefix /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://maitenance.xyz.com/maintenance.htm
```

NT LAN Manager (NTLM) Authentication

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
```

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 HTTP.

```
rspirep ^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_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 rests the connection if the 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.

Add Virtual Servers

Add an NSX Edge internal or uplink interface as a virtual server.

Prerequisites

- Verify that an application profile is available. See Create an Application Profile.
- If you are associating an application rule with the virtual server, see Create an Application Profile.
- If you are enabling acceleration to use the faster load balancer, the acceleration should be enabled when configuring the load balancer. See Configure Load Balancer Service.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click **Manage** and then click the **Load Balancer** tab.
5. In the left navigation panel, click **Virtual Servers**.
6. Click the **Add** (✚) icon.
7. Check **Enable Virtual Server** to make this virtual server available for use.
8. (Optional) Check **Enable Acceleration** for the NSX Edge load balancer to use the faster L4 load balancer engine rather than L7 load balancer engine.

   If a virtual server configuration such as application rules, HTTP type, or cookie persistence, is using the L7 load balancer engine, then the L7 load balancer engine is used even if you enabled acceleration or not.

   You can use the `show service load balancer virt` CLI command to confirm the load balancer engine in use.
9. 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.
10. Enter a name and description for the virtual server.
11. Click **Select IP Address** to set the IP address that the load balancer is listening on and type the protocol that the virtual server will handle.

    The Select IP Address dialog box shows only the primary IP address. If you are creating a VIP using a secondary IP address, enter it manually.
12. Select the protocol that the virtual server handles from the drop-down menu.
13. Enter the port number that the load balancer listens on.

    You can also set a range of ports for example, 80,8001-8004,443, to share the virtual server configuration such as, server pool, application profile, and application rule.

    To use FTP, the TCP protocol must have port 21 assigned to it.
14. Select the application rule.
15. Enter the maximum concurrent connections that the virtual server can process in the Connection Limit section.
16. Enter the maximum incoming new connection requests per second in the Connection Rate Limit section.
17. (Optional) Click the **Advanced** tab and add the application rule to associate it with the virtual server.

18 Click **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** and then click **NSX Edges**.

3 Double-click an NSX Edge.

4 Click the **Monitor** tab and then click the **Load Balancer** tab.

5 In the left navigation panel, click **Application Profiles**.

6 Select a profile and click the **Edit** (✎) icon.

7 Make the appropriate changes to traffic, persistence, certificate, or cipher configuration and click **Finish**.

Configure SSL Termination for a Load Balancer

Without SSL termination configured, HTTP request 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** to ensure that a valid certificate is present.

Procedure

1 In the application profile at **Manage > Load Balancer > Application Profiles**.

2 Select **HTTPS** type from the drop-down menu.

3 Verify that **Enable SSL Passthrough** is deselected.

4 Verify that **Configure Service Certificate** is selected.

5 Select the appropriate certificate from the list.

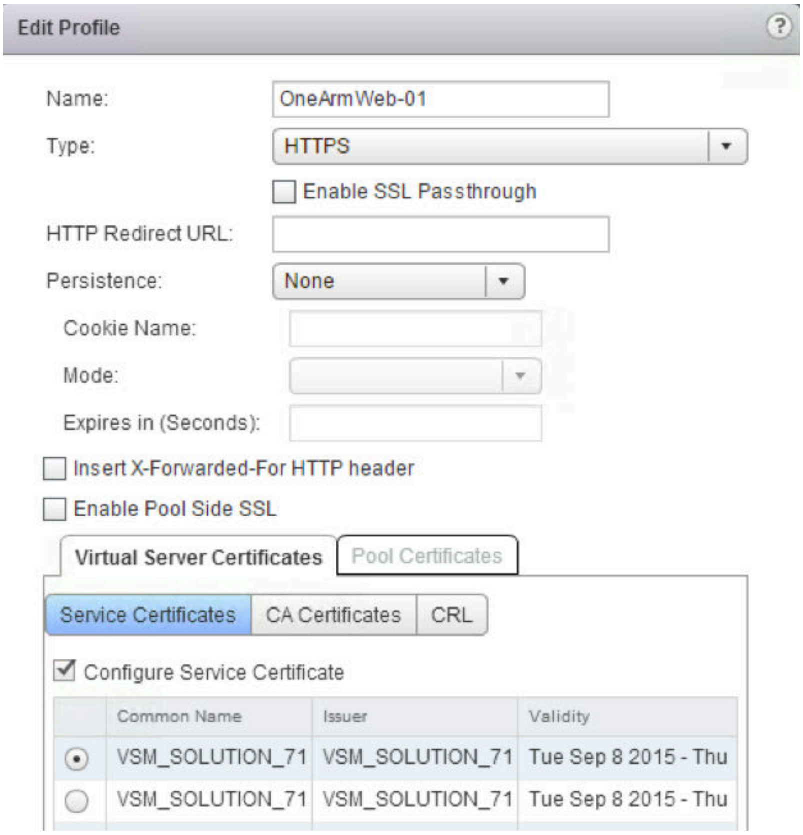

Edit Profile

Name: OneArmWeb-01

Type: HTTPS

□ Enable SSL Passthrough

HTTP Redirect URL:

Persistence: None

Cookie Name:

Mode:

Expires in (Seconds):

□ Insert X-Forwarded-For HTTP header

□ Enable Pool Side SSL

☑ Configure Service Certificate

	Common Name	Issuer	Validity
●	VSM_SOLUTION_71	VSM_SOLUTION_71	Tue Sep 8 2015 - Thu
○	VSM_SOLUTION_71	VSM_SOLUTION_71	Tue Sep 8 2015 - Thu

Delete an Application Profile

You can delete an application profile.

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **NSX Edges**.

3 Double-click an NSX Edge.

4 Click **Manage** and then click the **Load Balancer** tab.

5 In the left navigation panel, click **Application Profiles**.

6 Select a profile and click the **Delete** icon.

Managing Service Monitors

After you create a service monitor to define health check parameters for a network traffic and associate it with a pool, you can update the existing service monitor or delete it to save system resources.

Edit a Service Monitor

You can edit a service monitor.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click **Manage** and then click the **Load Balancer** tab.
5. In the left navigation panel, click **Service Monitoring**.
6. Select a service monitor and click the **Edit** icon.
7. Make the appropriate changes and click **OK**.

Delete a Service Monitor

You can delete a service monitor.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click **Manage** and then click the **Load Balancer** tab.
5. In the left navigation panel, click **Service Monitoring**.
6. Select a service monitor and click the **Delete** 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** and then click **NSX Edges**.
3. Double-click an NSX Edge.

4 Click the **Monitor** tab and then click the **Load Balancer** tab.

5 Ensure that you are in the Pool tab.

6 Select the pool to edit.

7 Click the **Edit** (✎) icon.

8 Make the appropriate changes and click **Finish**.

Configure a Load Balancer to Use Transparent Mode

Transparent indicates whether client IP addresses are visible to the backend servers. If Transparent is not selected (default value), backend servers see the traffic source IP as a Load balancer internal IP. If Transparent is selected, 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.

Procedure

- In the server pool configuration at **Manage > Load Balancer > Pools**, enable transparent mode.

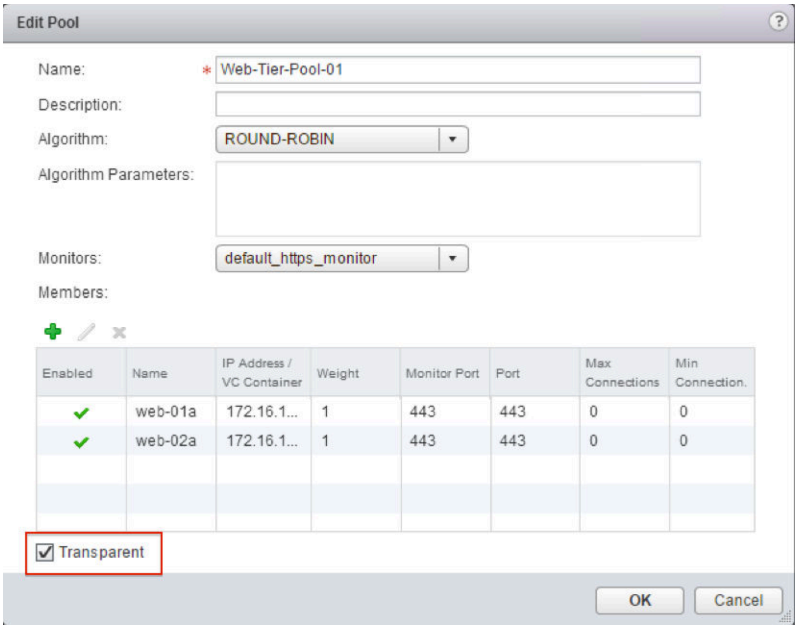

Edit Pool

Name: * Web-Tier-Pool-01

Description:

Algorithm: ROUND-ROBIN

Algorithm Parameters:

Monitors: default_https_monitor

Members:

Enabled	Name	IP Address / VC Container	Weight	Monitor Port	Port	Max Connections	Min Connection.
✓	web-01a	172.16.1...	1	443	443	0	0
✓	web-02a	172.16.1...	1	443	443	0	0

☑ Transparent

Delete a Server Pool

You can delete a server pool.

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **NSX Edges**.

3 Double-click an NSX Edge.

4 Click the **Monitor** tab and then click the **Load Balancer** tab.

5. Ensure that you are in the Pool tab.
6. Select the pool to delete
7. Click the **Delete** (✖) icon.

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** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click the **Monitor** tab and then click the **Load Balancer** tab.
5. Click **Virtual Servers** tab.
6. Select the virtual server to edit.
7. Click the **Edit** (✏) icon.
8. Make the appropriate changes and click **Finish**.

Delete a Virtual Server

You can delete a virtual server.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click the **Monitor** tab and then click the **Load Balancer** tab.
5. Click **Virtual Servers** tab.
6. Select the virtual server to delete.
7. Click the **Delete** (✖) 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

You can edit an application rule.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click **Manage** and then click the **Load Balancer** tab.
5. In the left navigation panel, click **Application Rules**.
6. Select a rule and click the **Edit** icon.
7. Make the appropriate changes and click **OK**.

Delete an Application Rule

You can delete an application rule.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click **Manage** and then click the **Load Balancer** tab.
5. In the left navigation panel, click **Application Profiles**.
6. Select a profile and click the **Delete** icon.

Load Balance Web Servers using NTLM Authentication

By default NSX Load Balancer closes the server TCP connection after each client request. Since NTLM authentication requires multiple HTTP requests in the same TCP session, authentication through NSX Load Balancer is broken.

Prerequisites

To work around this, add the following application rule on the VIP load balancing the Web servers using NTLM authentication:

```
add # NTLM authentication and keep the server connection open between requests
no option http-server-close
```

This application rule keeps the server connection open between requests.

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 incoming client traffic.

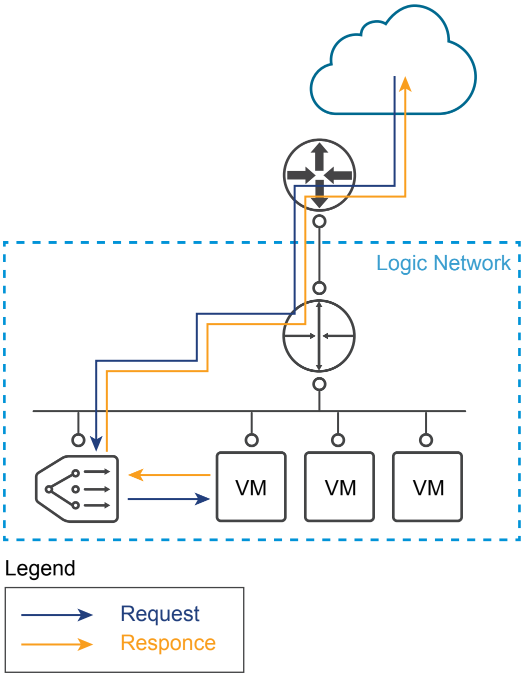

In proxy mode, the load balancer uses its own IP address as the source address to send requests to a backend server. The backend server views all traffic as being sent from the load balancer and responds to the load balancer directly. This mode is also called SNAT mode or non-transparent mode.

A typical NSX one-armed load balancer is deployed on the same subnet with its backend 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 backend servers. For the return traffic, reverse NAT is required to change the source IP address from the backend 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 would break.

After the ESG receives the traffic, it performs two operations: Destination Network Address Translation (DNAT) to change the VIP address to the IP address of one of the load balanced machines, and 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 then back to the client. This option is much easier to configure than the Inline mode, but has two potentials 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/HTTPS applications is to enable Insert X-Forwarded-For in the HTTP application profile so that the client IP address will be carried in the X-Forwarded-For HTTP header in the request sent to the backend server.

If client IP address visibility is required on the backend server for applications other than HTTP/HTTPS, you can configure the IP pool to be transparent. In case clients are not on the same subnet as the backend server, inline mode is recommended. Otherwise, you must use the load balancer IP address as the default gateway of the backend server.

Note Usually, there are two methods to guarantee connection integrity:

- SNAT/proxy/non-transparent mode (discussed above)
- Direct server return (DSR)

In DSR mode, the backend server responds directly to the client. Currently, NSX load balancer does not support DSR.

Procedure

1 Create a certificate by double-clicking the Edge and then selecting **Manage > Settings > Certificate**.

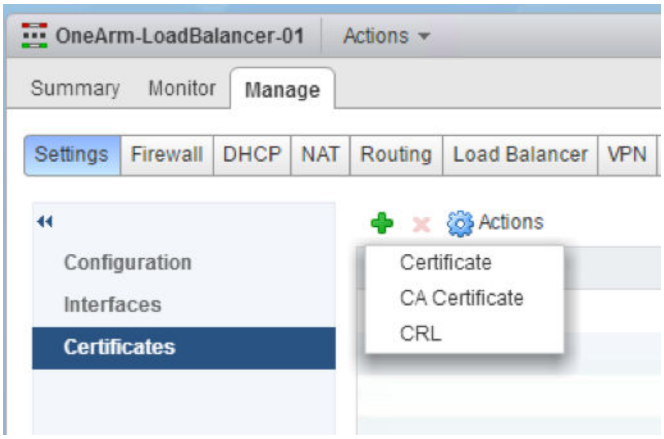

2 Enable the load balancer service by selecting **Manage > Load Balancer > Global Configuration > Edit**.

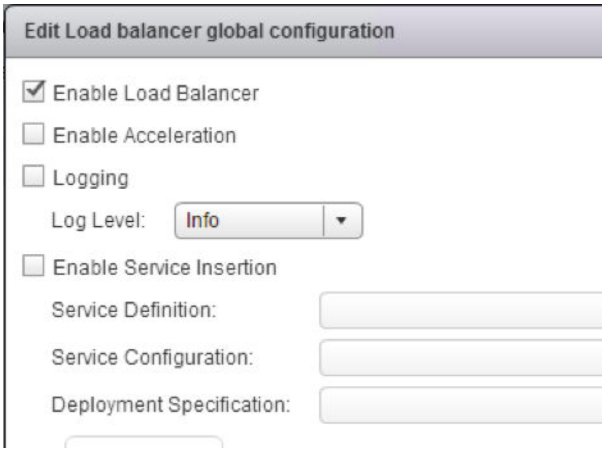

Edit Load balancer global configuration

☑ Enable Load Balancer

□ Enable Acceleration

□ Logging

Log Level: Info

□ Enable Service Insertion

Service Definition:

Service Configuration:

Deployment Specification:

3 Create an HTTPS application profile by selecting **Manage > Load Balancer > Application Profiles**.

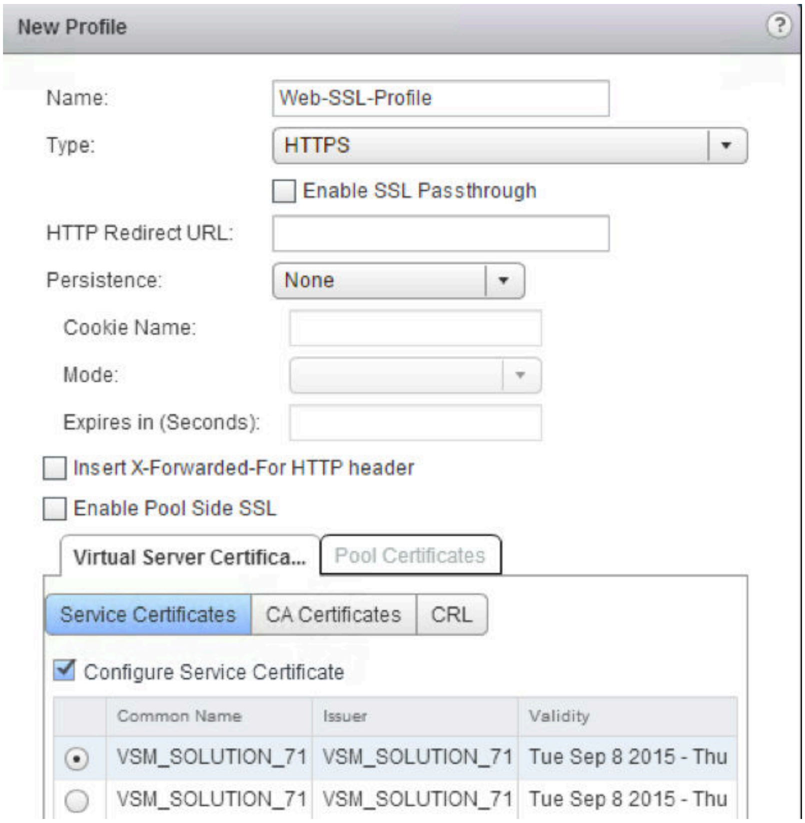

New Profile

Name: Web-SSL-Profile

Type: HTTPS

□ Enable SSL Passthrough

HTTP Redirect URL:

Persistence: None

Cookie Name:

Mode:

Expires in (Seconds):

□ Insert X-Forwarded-For HTTP header

□ Enable Pool Side SSL

Virtual Server Certifica... | Pool Certificates

Service Certificates | CA Certificates | CRL

☑ Configure Service Certificate

	Common Name	Issuer	Validity
●	VSM_SOLUTION_71	VSM_SOLUTION_71	Tue Sep 8 2015 - Thu
○	VSM_SOLUTION_71	VSM_SOLUTION_71	Tue Sep 8 2015 - Thu

Note The screenshot above uses self-signed certificates for documentation-purposes only.

4 Optionally, click **Manage > Load Balancer > Service Monitoring** and edit the default service monitoring to change it from basic HTTP/HTTPS to specific URL/URIs, as required.

5 Create server pools by selecting **Manage > Load Balancer > Pools**.

To use SNAT mode, leave the **Transparent** check box unchecked in the pool configuration.

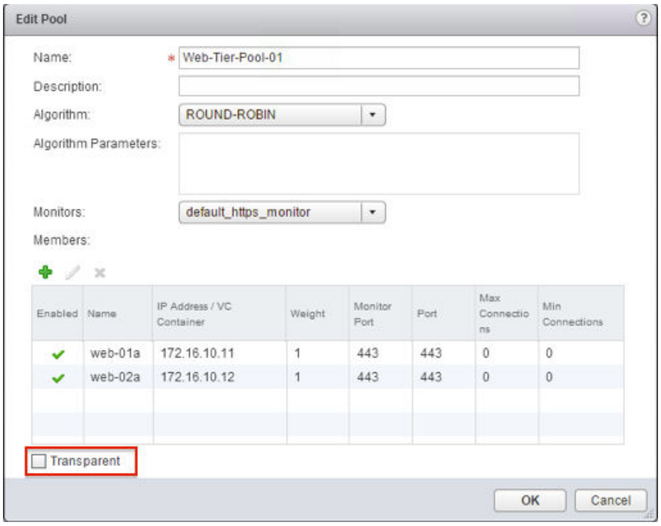

Name:	* Web-Tier-Pool-01
Description:	
Algorithm:	ROUND-ROBIN
Algorithm Parameters:	
Monitors:	default_https_monitor
Members:	

Enabled	Name	IP Address / VC Container	Weight	Monitor Port	Port	Max Connections	Min Connections
✓	web-01a	172.16.10.11	1	443	443	0	0
✓	web-02a	172.16.10.12	1	443	443	0	0

☐ Transparent

Ensure that the VMs are listed and enabled.

6 Optionally, click **Manage > Load Balancer > Pools > Show Pool Statistics** to check the status.

Make sure that the member status is UP.

7 Create a virtual server by selecting **Manage > Load Balancer > Virtual Servers**.

If you would like to use the L4 load balancer for UDP or higher-performance TCP, check **Enable Acceleration**. If you check **Enable Acceleration**, make sure that the firewall status is **Enabled** on the load balancer NSX Edge, because a firewall is required for L4 SNAT.

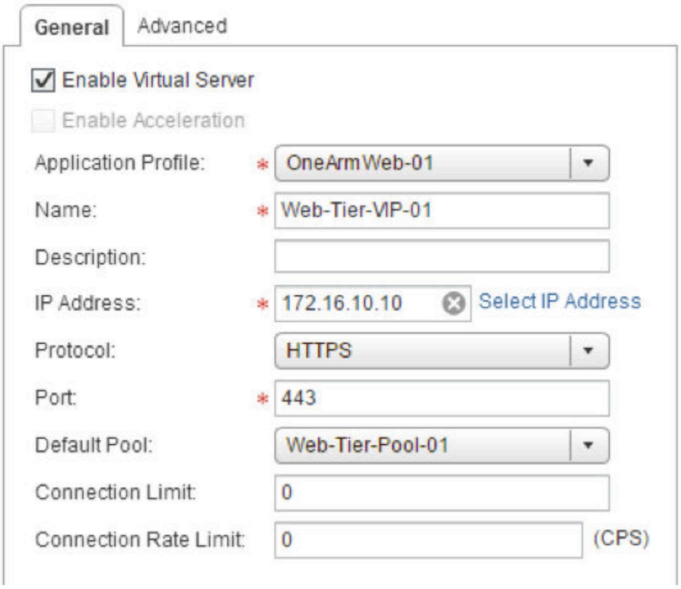

☑ Enable Virtual Server

☐ Enable Acceleration

Application Profile:	* OneArmWeb-01
Name:	* Web-Tier-VIP-01
Description:	
IP Address:	* 172.16.10.10
Protocol:	HTTPS
Port:	* 443
Default Pool:	Web-Tier-Pool-01
Connection Limit:	0
Connection Rate Limit:	0 (CPS)

Ensure that the IP address is tied to the server pool.

8 Optionally, if using an application rule, check the configuration in **Manage > Load Balancer > Application Rules**.

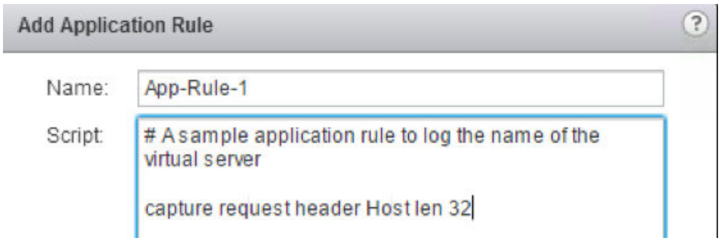

Add Application Rule

Name: App-Rule-1

Script:

```
# A sample application rule to log the name of the
virtual server

capture request header Host len 32
```

9 If using an application rule, ensure that the application rule is associated with the virtual server in **Manage > Load Balancer > Virtual Servers > Advanced**.

For supported examples, see: https://communities.vmware.com/docs/DOC-31772.

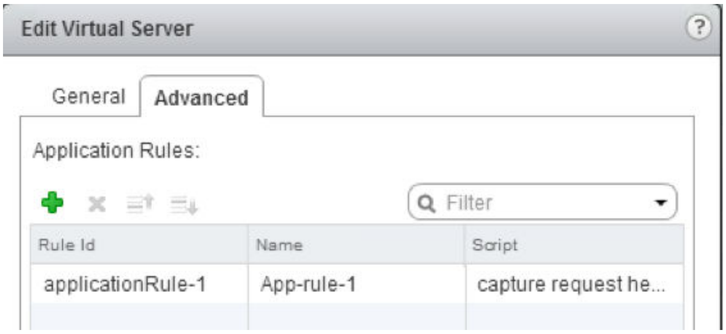

Edit Virtual Server

Application Rules:

Rule Id	Name	Script
applicationRule-1	App-rule-1	capture request he...

In non-transparent mode, the backend server cannot see the client IP, but can see the load balancer internal IP address. As a workaround for HTTP/HTTPS traffic, check **Insert X-Forwarded-For HTTP header**. With this option checked, the Edge load balancer adds the header "X-Forwarded-For" with the value of the client source IP address.

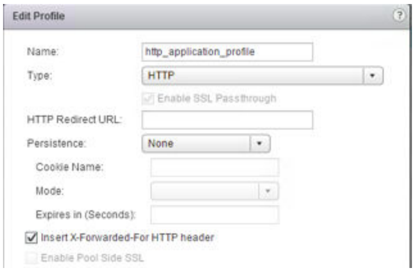

Edit Profile

Name: http_application_profile

Type: HTTP

☑ Enable SSL Passthrough

HTTP Redirect URL:

Persistence: None

Cookie Name:

Mode:

Expires in (Seconds):

☑ Insert X-Forwarded-For HTTP header

□ Enable Pool Side SSL

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 you configure the NSX load balancer, you can provide the NSX Edge device uplink interface IP address for vCenter Single Sign-On.

Prerequisites

- Perform the PSC High Availability preparation tasks listed in the knowledge. See http://kb.vmware.com/kb/2113315.
- Save the /ha/lb.crt and /ha/lb_rsa.key from 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 internal logical switch.

Procedure

1 Add PSC CA certificates to the NSX Edge.

   a Save the PSC root.cer and certificate, RSA and passphrase generated from the OpenSSL command.

   b Double-click the Edge and select**Manage > Settings > Certificate** .

   c Add the saved content root.cer file to the CA certificate contents.

   d Add the saved passphrase to the private key section.

2 Enable the load balancer service.

   a Select **Manage > Load Balancer > Edit**.

   b Check the **Enable Load Balancing** and **Logging** options.

3 Create application profiles with TCP and HTTPS protocols.

a Select **Manage > Load Balancer > Application Profiles**.

b Create a TCP application profile.

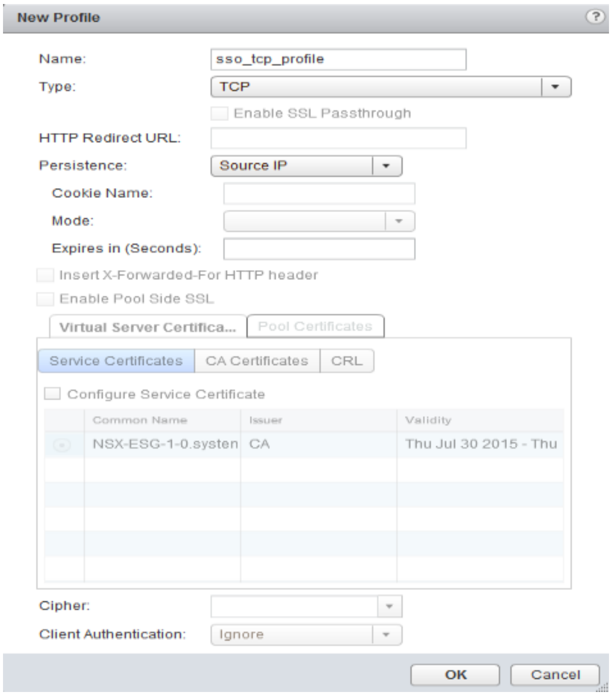

New Profile

Name: sso_tcp_profile

Type: TCP

□ Enable SSL Passthrough

HTTP Redirect URL:

Persistence: Source IP

Cookie Name:

Mode:

Expires in (Seconds):

□ Insert X-Forwarded-For HTTP header

□ Enable Pool Side SSL

Virtual Server Certifica... | Pool Certificates

Service Certificates | CA Certificates | CRL

□ Configure Service Certificate

	Common Name	Issuer	Validity
○	NSX-ESG-1-0.systen	CA	Thu Jul 30 2015 - Thu

Cipher:

Client Authentication: Ignore

c Create an HTTPS application profile.

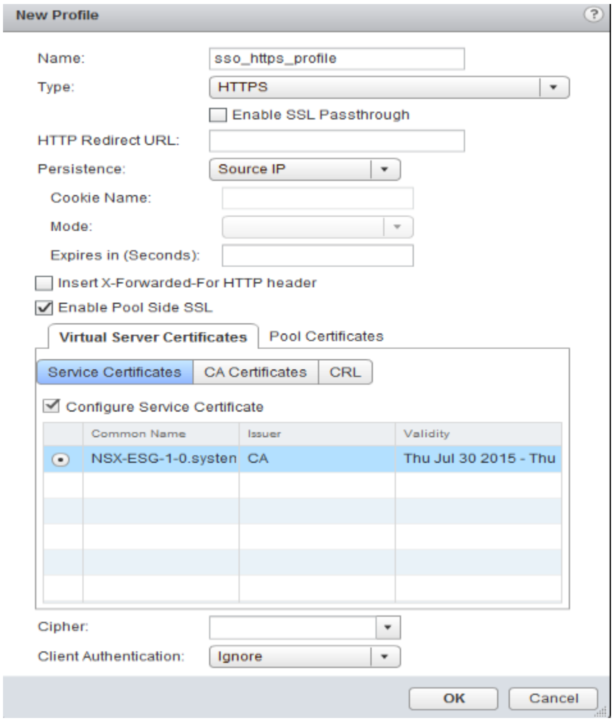

New Profile

Name: sso_https_profile

Type: HTTPS

□ Enable SSL Passthrough

HTTP Redirect URL:

Persistence: Source IP

Cookie Name:

Mode:

Expires in (Seconds):

□ Insert X-Forwarded-For HTTP header

☑ Enable Pool Side SSL

Virtual Server Certificates | Pool Certificates

Service Certificates | CA Certificates | CRL

☑ Configure Service Certificate

	Common Name	Issuer	Validity
●	NSX-ESG-1-0.systen	CA	Thu Jul 30 2015 - Thu

Cipher:

Client Authentication: Ignore

4 Create application pools to add member PSC nodes.

a Select **Manage > Load Balancer > Pools**.

b Create two application pools with monitor port 443.

Use the PSC node IP address.

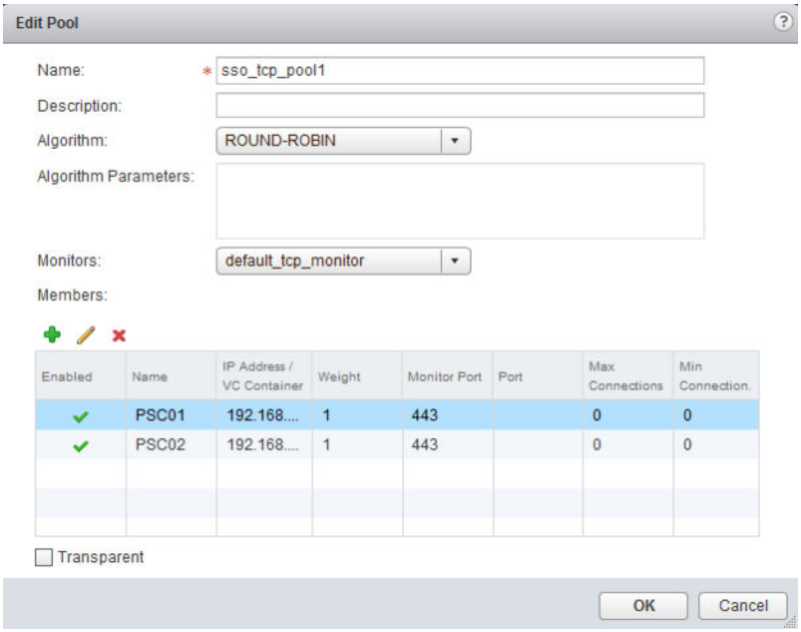

Edit Pool

Name:	* sso_tcp_pool1
Description:	
Algorithm:	ROUND-ROBIN
Algorithm Parameters:	
Monitors:	default_tcp_monitor
Members:	

Enabled	Name	IP Address / VC Container	Weight	Monitor Port	Port	Max Connections	Min Connection.
✓	PSC01	192.168....	1	443		0	0
✓	PSC02	192.168....	1	443		0	0

□ Transparent

c Create two application pools with monitor port 389.

Use the PSC node IP address.

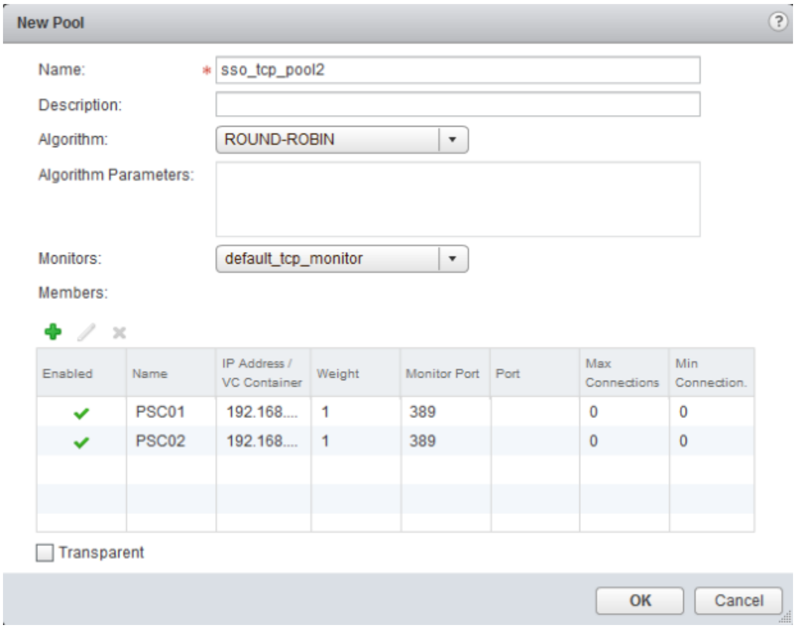

New Pool

Name:	* sso_tcp_pool2
Description:	
Algorithm:	ROUND-ROBIN
Algorithm Parameters:	
Monitors:	default_tcp_monitor
Members:	

Enabled	Name	IP Address / VC Container	Weight	Monitor Port	Port	Max Connections	Min Connection.
✓	PSC01	192.168....	1	389		0	0
✓	PSC02	192.168....	1	389		0	0

□ Transparent

5 Create virtual servers for the TCP and HTTPS protocols.

a Select **Manage > Load Balancer > Virtual Servers** .

b Create a virtual server for TCP VIP.

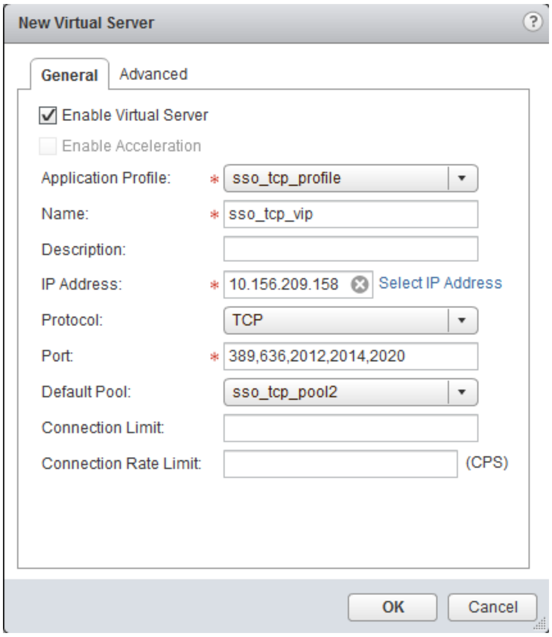

New Virtual Server

General | Advanced

☑ Enable Virtual Server

□ Enable Acceleration

Application Profile: *	sso_tcp_profile
Name: *	sso_tcp_vip
Description:	
IP Address: *	10.156.209.158
Protocol:	TCP
Port: *	389,636,2012,2014,2020
Default Pool:	sso_tcp_pool2
Connection Limit:	
Connection Rate Limit:	(CPS)

c Create a virtual server for HTTPS VIP.

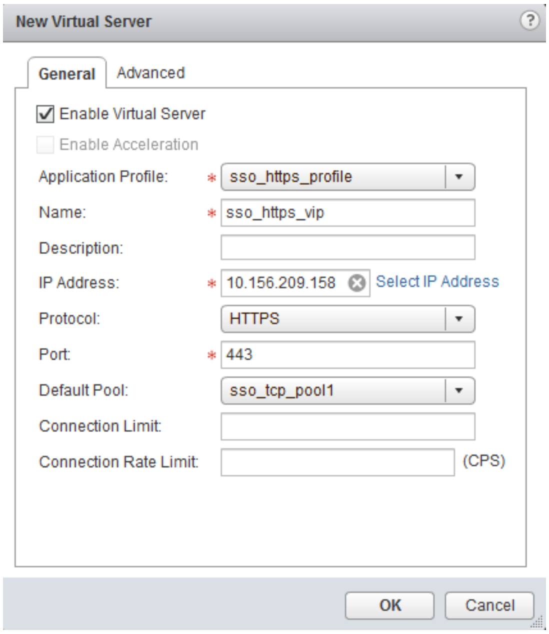

New Virtual Server

General | Advanced

☑ Enable Virtual Server

□ Enable Acceleration

Application Profile: *	sso_https_profile
Name: *	sso_https_vip
Description:	
IP Address: *	10.156.209.158
Protocol:	HTTPS
Port: *	443
Default Pool:	sso_tcp_pool1
Connection Limit:	
Connection Rate Limit:	(CPS)

Other Edge Services

16

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 DNS Servers

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.

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, thus one IP address can belong to only one IP pool.

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **NSX Edges**.

3 Double-click an NSX Edge.

4 Click the **Manage** tab and then click the **DHCP** tab.

5 Click the **Add** (➕) icon.

6 Configure the pool.

Option	Action
Auto Configure DNS	Select to use the DNS service configuration for the DHCP binding.
Lease never expires	Select to bind the address to the MAC address of the virtual machine forever. If you select this, **Lease Time** is disabled.
Start IP	Type the starting IP address for the pool.
End IP	Type the ending IP address for the pool.
Domain Name	Type the domain name of the DNS server. This is optional.
Primary Name Server	If you did not select **Auto Configure DNS**, type the **Primary Nameserver** for the DNS service. You must enter the IP address of a DNS server for hostname-to-IP address resolution. This is optional.
Secondary Name Server	If you did not select **Auto Configure DNS**, type the **Secondary Nameserver** for the DNS service. You must enter the IP address of a DNS server for hostname-to-IP address resolution. This is optional.
Default Gateway	Type 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 is optional.
Subnet Mask	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 distributed router.
Lease Time	Select whether to lease the address to the client for the default time (1 day), or type a value in seconds. You cannot specify the lease time if you selected **Lease never expires**. This is optional.

7 Click **OK**.

Enable the DHCP Service

Enable 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 have been added.

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **NSX Edges**.

3. Double-click an NSX Edge.
4. Click the **Manage** tab and then click the **DHCP** tab.
5. Click **Enable**.
6. Select **Enable logging** if required and select the log level.
7. Click **Publish Changes**.

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** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click the **Manage** tab and then click the **DHCP** tab.
5. Select a DHCP pool and click the **Edit** icon.
6. Make the appropriate changes and click **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** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click the **Manage** tab and then click the **DHCP** tab.
5. Select **Bindings** from the left panel.
6. Click the **Add** (✚) icon.

7 Configure the binding.

Option	Action
Auto Configure DNS	Select to use the DNS service configuration for the DHCP binding.
Lease never expires	Select to bind the address to the MAC address of the virtual machine forever.
Interface	Select the NSX Edge interface to bind.
VM Name	Select the virtual machine to bind.
VM vNIC Index	Select the virtual machine NIC to bind to the IP address.
Host Name	Type the host name of the DHCP client virtual machine.
IP Address	Type 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	Type the domain name of the DNS server.
Primary Name Server	If you did not select **Auto Configure DNS**, type the **Primary Nameserver** 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 **Auto Configure DNS**, type the **Secondary Nameserver** for the DNS service. You must enter the IP address of a DNS server for hostname-to-IP address resolution.
Default Gateway	Type 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 Time	If you did not select **Lease never expires**, select whether to lease the address to the client for the default time (1 day), or type a value in seconds.

8 Click **Add**.

9 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** and then click **NSX Edges**.

3 Double-click an NSX Edge.

4 Click the **Manage** tab and then click the **DHCP** tab.

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 **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.

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.

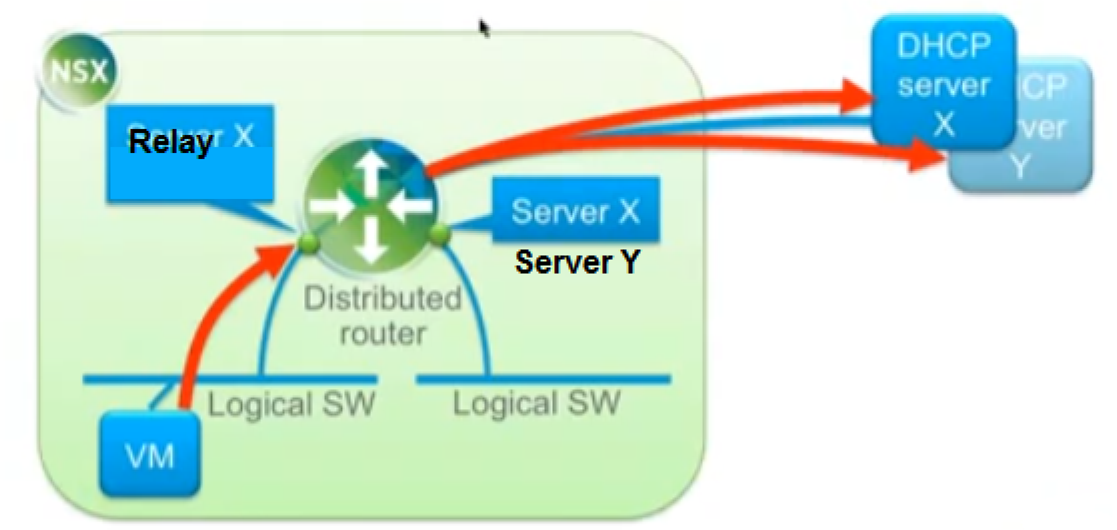

Note

- 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.

Procedure

1 In the vSphere Web Client, navigate to **Networking & Security > NSX Edges**.

2 Double-click the appropriate Edge and ensure that that you are in the **Manage > DHCP** tab.

3 Click **Edit** next to **DHCP Relay Global Configuration**.

4 To add an IP set as the a server:

   a Click the **Add** icon and select the IP set.

   b Move the selected IP set to the Selected Objects list by clicking the ➡ icon.

   c Click **OK**.

5 To add IP addresses or domain names, type the address or name in the appropriate area.

6 Click **OK**.

Add Relay Agents

Add the Edge interfaces from which DHCP messages are to be relayed to the external DHCP relay server(s).

Procedure

1 In the **DHCP Relay Agents** area, click the **Add** icon.

2 In **vNIC**, ensure that an internal vNIC is selected.

   The **Gateway IP Address** displays the primary IP address of the selected vNic.

3 Click **OK**.

Configure DNS Servers

You can configure external DNS servers to which NSX Edge can relay name resolution requests from clients. NSX Edge will relay client application requests to the DNS servers to fully resolve a network name and cache the response from the servers.

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **NSX Edges**.

3 Double-click a NSX Edge.

4. Click the **Manage** tab and then click the **Settings** tab.
5. In the **DNS Configuration** panel, click **Change**.
6. Click **Enable DNS Service** to enable the DNS service.
7. Type IP addresses for both DNS servers.
8. Change the default cache size if required.
9. Click **Enable Logging** to log DNS traffic and select the log level.

   Generated logs are sent to the syslog server.
10. Click **Ok**.

Service Composer

17

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.

Security Policy

A security policy is a collection of the following service configurations.

Table 17-1. 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	Data Security or 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 17-1. Service Composer overview

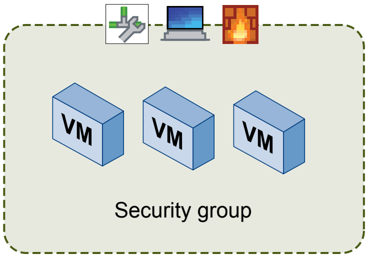

If a virtual machine belongs to more than one security group, the services that are applied to the virtual machine depends on the precedence of the security policy mapped to the security groups.

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
- Graphical View of Service Composer
- Working with Security Tags
- Viewing Effective Services
- Working with Security Policies
- Edit a Security Group
- 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 17-2. Service Composer in action

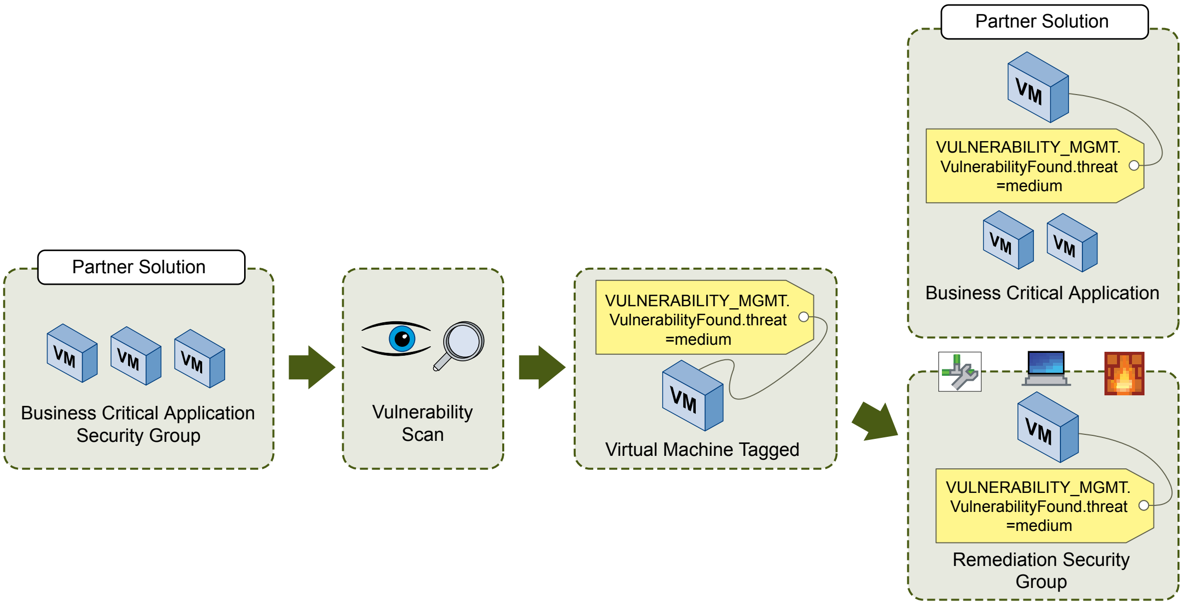
Partner Solution
VM
VM
VM
Business Critical Application Security Group
Vulnerability Scan
VULNERABILITY_MGMT. VulnerabilityFound.threat =medium
VM
Virtual Machine Tagged
Partner Solution
VM
VULNERABILITY_MGMT. VulnerabilityFound.threat =medium
VM
VM
Business Critical Application
VM
VULNERABILITY_MGMT. VulnerabilityFound.threat =medium
Remediation Security Group

This topic will now take you through the steps required to consume the security services offered by Service Composer.

1 Create a Security Group in Service Composer

You create a security group at the NSX Manager level.

2. 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.

3. 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.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **Service Composer**.
3. Click the **Security Groups** tab and then click the **Add Security Group** icon.
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. 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.

   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.

   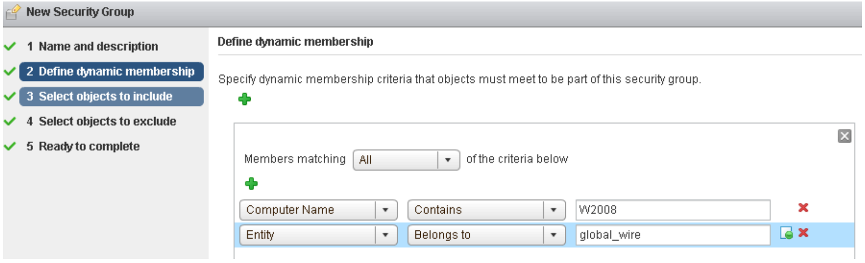

   New Security Group

   ✓ 1 Name and description
   ✓ 2 Define dynamic membership
   ✓ 3 Select objects to include
   ✓ 4 Select objects to exclude
   ✓ 5 Ready to complete

   Define dynamic membership

   Specify dynamic membership criteria that objects must meet to be part of this security group.

   Members matching All of the criteria below

Computer Name	Contains	W2008
Entity	Belongs to	global_wire

6. Click **Next**.
7. On the Select objects to include page, select the object type from the drop-down.

8 Double-click the object 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
- 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**.

Membership of a security group is determined as follows:

{Expression result (derived from Step 5) + Inclusions (specified inStep 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.

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, Data Security, 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.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **Service Composer**.
3. Click the **Security Policies** tab.
4. Click the **Create Security Policy** () icon.
5. In the Add Security Policy dialog box, type a name for the security policy.
6. Type a description for the security policy.

   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.
7. Select **Inherit security policy from specified 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.
8. Click **Next**.

9 In the Guest Introspection Services page, click 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.

  If you select Data Security, you must have a data security policy in place. See Chapter 19 Data Security.

  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 an 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.

  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.

10 Click **Next**.

11 On the Firewall page, click the **Add Firewall Rule** (✚) icon.

Here, you are defining firewall rules for the security groups(s) that this security policy will be applied to.

a Type a name and description for the firewall rule you are adding.

b Select **Allow** or **Block** to indicate whether the rule needs to allow or block traffic to the selected destination.

c 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 **Change** and select the appropriate security groups.

d 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.

e Select the services and/or service groups to which the rule applies to.

f Select **Enabled** or **Disabled** to specify the rule state.

g Select **Log** to log sessions matching this rule.

Enabling logging may affect performance.

h 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.

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 by selecting **Synchronize Firewall Rules** from the **Actions** menu in the Security Policies tab.

12 Click **Next**.

The Network Introspection Services page displays NetX services that you have integrated with your VMware virtual environment.

13 Click the **Add Network Introspection Service** (✚) icon.

a In the Add Network Introspection Service dialog box, type 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 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.

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.

Edit Service Composer Firewall Applied To Setting

You can set the applied to setting for all firewall rules created though 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 Log in to the vSphere Web Client.

2 Click **Networking & Security**, click **Service Composer**, and click the **Security Policies** tab.

3 Click **Actions > Edit Firewall Policy Settings**. Select a default setting for Applied To and click OK.

Option	Description
Distributed Firewall	Firewall rules are applied to all clusters on which Distributed Firewall is installed.
Policy's Security Groups	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*.

Example: Applied To Behavior

In this example scenario, your default firewall rule action 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

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 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **Service Composer**.

3 Click the **Security Policy** tab.

4 Select a security policy and click the **Apply Security Policy** () icon.

5 Select the security group that you want to apply the policy to.

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.

6 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.

7 Click **OK**.

Graphical View of Service Composer

Service Composer offers a canvas 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.

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 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **Service Composer**.

3 Click the **Canvas** tab.

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.

Figure 17-3. Service Composer canvas top level view

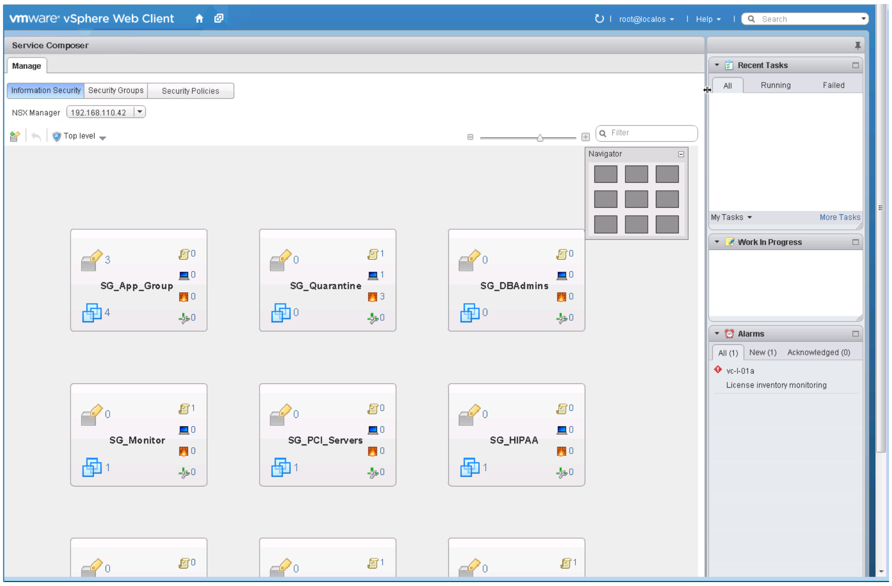

Each rectangular box in the canvas represents a security group and the icons within the box represents security group members and details about the security policy mapped to the security group.

Figure 17-4. Security group

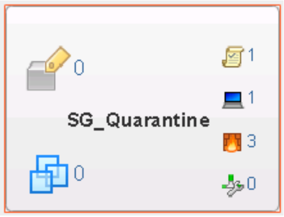

A number next to each icon indicates the number of instances - for example, 📜 1 indicates that 1 security policy is mapped to that security group.

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. ■ You can create a new security policy by clicking the Create **Security Policy** (📜) icon. The newly created security policy object is automatically mapped to the security group. ■ Map additional security policies to the security group by clicking the **Apply Security Policy** (📜) icon.

Icon	Click to display
💻	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 17-5. Details displayed when you click an icon in the security group

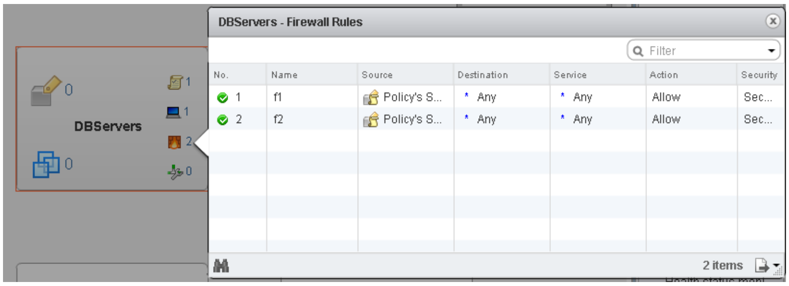

DBServers - Firewall Rules

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...

2 items

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.

View Applied Security Tags

You can view the security tags applied to virtual machines in your environment.

Prerequisites

A data security or 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 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **NSX Managers**.

3 Click an NSX Manager in the **Name** column and then click the **Manage** tab.

4 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.

5 Click the number in the **VM Count** column to view the virtual machines to which that tag in that row has been applied.

Add a Security Tag

You can manually add a security tag and apply it to a virtual machine. This is especially useful when you are using a non-NETX solution in your environment and hence, cannot register the solution tags with NSX Manager.

Prerequisites

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **NSX Managers**.

3 Click an NSX Manager in the **Name** column and then click the **Manage** tab.

4 Click the **Security Tags** tab.

5. Click the **New Security Tag** (✚) icon.
6. Type a name and description for the tag and click **OK**.

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 a virtual machine.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **NSX Managers**.
3. Click an NSX Manager in the **Name** column and then click the **Manage** tab.
4. Click the **Security Tags** tab.
5. Select a security tag and click the **Assign Security Tag** (✚) icon.
6. Select one or more virtual machines and click **OK**.

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. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **NSX Managers**.
3. Click an NSX Manager in the **Name** column and then click the **Manage** tab.
4. Click the **Security Tags** tab.
5. Select a security tag and click the **Edit Security Tag** (✎) icon.
6. Make the appropriate changes and click **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. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **NSX Managers**.
3. Click an NSX Manager in the **Name** column and then click the **Manage** tab.
4. Click the **Security Tags** tab.

5 Select a security tag and click 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. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **Service Composer**.
3. Click the **Security Policies** tab.
4. Click a security policy in the **Name** column.
5. Ensure that you are in the **Manage > Information Security** tab.

Each of the three tabs (**Endpoint Services**, **Firewall**, **Network Introspection Services**) displays the corresponding services for the security policy.

Services that are not effective are greyed out. The **Overridden** column displays the services that are actually applied on the security policy and the **Inherited from** 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. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **Service Composer**.
3. Click the **Security Policies** tab.
4. Click a security policy in the **Name** column.
5. 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 - data security, 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. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **Service Composer**.
3. Click the **Security Policies** tab.
4. Click the **Manage Precedence** (≡↕) icon.

5 In the Manage Precedence dialog box, select the security policy that you want to change the precedence for and click the **Move Up** (↑) or **Move Down** (↓)icon.

6 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 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **Service Composer**.

3 Click the **Security Policies** tab.

4 Select the security policy that you want to edit and click the **Edit Security Policy** (✎) icon.

5 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 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **Service Composer**.

3 Click the **Security Policies** tab.

4 Select the security policy that you want to delete and click the **Delete Security Policy** (✖) icon.

Edit a Security Group

You can edit a security group.

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **Service Composer**.

3 Click the **Security Groups** tab.

4 Select the security group you want to edit and click the **Edit Security Group** icon.

5 Make the appropriate changes and click **OK**.

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 17-6. Configuring Service Composer

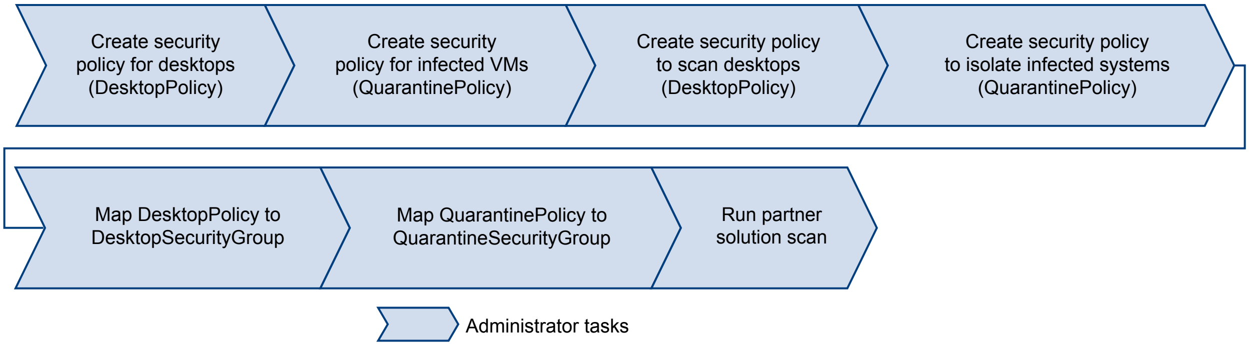
Create security policy for desktops (DesktopPolicy)
Create security policy for infected VMs (QuarantinePolicy)
Create security policy to scan desktops (DesktopPolicy)
Create security policy to isolate infected systems (QuarantinePolicy)
Map DesktopPolicy to DesktopSecurityGroup
Map QuarantinePolicy to QuarantineSecurityGroup
Run partner solution scan
Administrator tasks

Figure 17-7. Service Composer Conditional Workflow

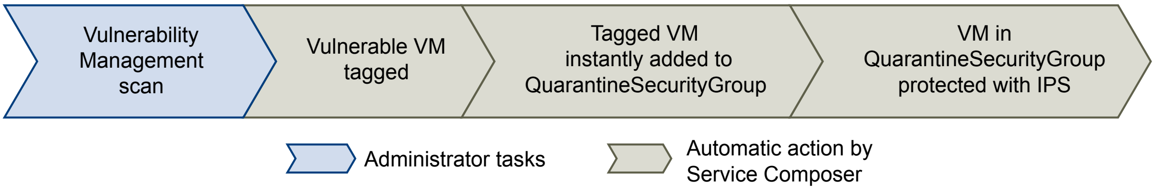
Vulnerability Management scan
Vulnerable VM tagged
Tagged VM instantly added to QuarantineSecurityGroup
VM in QuarantineSecurityGroup protected with IPS
Administrator tasks
Automatic action by Service Composer

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.

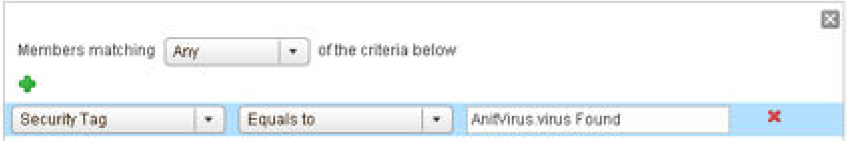

Members matching Any of the criteria below

✚

Security Tag | Equals to | AntiVirus.virus Found ✖

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 (0)

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 very 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.

   a Click the Information Security tab.

   b Click the number next to the icon to see that the SP_App is mapped.

8 Export the SP_App policy.

   a Click the Security Policies tab and then click the **Export Blueprint** ( ) icon.

   b In **Name**, type `Template_ App_` and in **Prefix**, type `FromAppArchitect`.

   c Click Next.

   d Select the SP_App policy and click Next.

   e Review your selections and click Finish.

   f 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.

   a Click **Actions** and then click the **Import Service Configuration** icon.

   b Select the `FromAppArtchitect_Template_App` file from your desktop (you saved it in step 7).

   c 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

18

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 vCenter Server must be correctly configured for Guest Introspection security:

- Not all guest operating systems are supported by Guest Introspection. Virtual machines with non-supported operating systems are not protected by the security solution.
- 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 ESX host to another within the resource pool.

This chapter includes the following topics:

Install Guest Introspection

Installing Guest Introspection automatically installs a new VIB and a service virtual machine on each host in the cluster. Guest Introspection is required for NSX Data Security, Activity Monitoring, and several third-party security solutions.

For autodeploy setup on stateless hosts, you must manually restart VMware NSX for vSphere 6.x Service Virtual Machines (SVM) after an ESXi host reboot. For more information, see the Knowledge Base article http://kb.vmware.com/kb/2120649.

Caution In a VMware NSX for vSphere 6.x environment, when a Service VM (SVM) is migrated (vMotion/SvMotion), you may experience these symptoms:

- An interruption in the service (workload VM) for which the Service VM (SVM) is providing data
- ESXi host fails with a purple diagnostic screen contains backtraces similar to:

```
@BlueScreen: #PF Exception 14 in world wwww:WorldName IP 0xnnnnnnnn addr 0x0
PTEs:0xnnnnnnnn;0xnnnnnnnn;0x0;
0xnnnnnnnn:[0xnnnnnnnn]VmMemPin_DecCount@vmkernel#nover+0x1b
0xnnnnnnnn:[0xnnnnnnnn]VmMemPinUnpinPages@vmkernel#nover+0x65
0xnnnnnnnn:[0xnnnnnnnn]VmMemPin_ReleaseMainMemRange@vmkernel#nover+0x6
0xnnnnnnnn:[0xnnnnnnnn]P2MCache_ReleasePages@vmkernel#nover+0x2a
0xnnnnnnnn:[0xnnnnnnnn]DVFilterVmciUnmapGuestPage@com.vmware.vmkapi#v2_2_0_0+0x34
```

This is a known issue affecting VMware ESXi 5.5.x and 6.x hosts. To work around the issue, do not manually migrate a Service VM (SVM) (vMotion/SvMotion) to another ESXi host in the cluster. To migrate a SVM to another datastore (svMotion), VMware recommends a cold migration by turning the SVM off, and then migrating it to another datastore.

Prerequisites

The installation instructions that follow assume that you have the following system:

- A datacenter with supported versions of vCenter Server and ESXi installed on each host in the cluster.
- If the hosts in your clusters were upgraded from vCenter Server version 5.0 to 5.5, you must open ports 80 and 443 on those hosts.
- 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 using NSX for 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 6.2 installed and running.
- Ensure the NSX Manager and the prepared hosts that will run Guest Introspection services are linked to the same NTP server and that time is synchronized. Failure to do so may 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 you want to assign an IP address to the NSX Guest Introspection service virtual machine from an IP pool, create the IP pool before installing NSX Guest Introspection. See Working with IP Pools in *the NSX Administration Guide*.

vSphere Fault Tolerance does not work with Guest Introspection.

Procedure

1 On the **Installation** tab, click **Service Deployments**.

2 Click the **New Service Deployment** (✚) icon.

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 as soon as it is installed or select a deployment date and time.

5 Click **Next**.

6 Select the datacenter and cluster(s) 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**, follow the steps below for each host in the cluster.

   a On the vSphere Web Client home page, click **vCenter** and then click **Hosts**.

   b Click a host in the **Name** column and then click the **Manage** tab.

   c Click **Agent VM Settings** and click **Edit**.

   d Select the datastore and click **OK**.

8 Select the distributed virtual port group to host the management interface. If the datastore is set to **Specified on host**, the network must also be **Specified on host**.

   The selected port group must be able to reach the NSX Manager’s port group and must be available on all hosts in the selected cluster.

   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 NSX Guest Introspection service virtual machine through Dynamic Host Configuration Protocol (DHCP). Select this option if your hosts are on different subnets.
An IP pool	Assign an IP address to the NSX 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**.

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. In some cases, resolving the errors displays additional errors. Take the required action and click **Resolve** again.

What to do next

Install VMware Tools on guest virtual machines.

Install VMware Tools on the Guest Virtual Machines

VMware Tools include the NSX Thin Agent that must be installed on each guest virtual machine to be protected. Virtual machines with VMware Tools installed are automatically protected whenever they are started up on an ESX host that has the security solution installed. That is, protected virtual machines retain the security protection through shut downs and restarts, and even after a vMotion move to another ESX host with the security solution installed.

Prerequisites

Ensure that the guest virtual machine has ESX 5.1 or later and a supported version of Windows installed. The following Windows operating systems are supported for NSX Guest Introspection:

- Windows Vista (32 bit)
- Windows 7 (32/64 bit)
- Windows XP SP3 and above (32 bit)
- Windows 2003 SP2 and above (32/64 bit)
- Windows 2003 R2 (32/64 bit)
- Windows 2008 (32/64 bit)
- Windows 2008 R2 (64 bit)
- Windows 8 (32/64) -- from vSphere 5.5 and later
- Win2012 (64) -- from vSphere 5.5 and later
- Windows 8.1 (32/64) -- from vSphere 5.5 Patch 2 and later
- Win2012 R2 (64) -- from vSphere 5.5 Patch 2 and later

Procedure

1 Follow the procedure Manually Install or Upgrade VMware Tools in a Windows Virtual Machine.

2 After you select **Custom** setup in step 7, expand the **VMCI Driver** section, select **vShield Drivers**, and select **This feature will be installed on the local hard drive**.

3 Follow the remaining steps in the procedure.

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 ESX host-resident Guest Introspection module, and the protected virtual machine-resident thin agent.

Procedure

1. In the vSphere Web Client, click **vCenter**, and then click **Datacenters**.
2. In the **Name** column, click a datacenter.
3. Click **Monitor** and then click **Endpoint**.

   The Guest Introspection Health and Alarms page displays the health of the objects under the datacenter 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 Alarms

Alarms signal the vCenter Server administrator about Guest Introspection events that require attention. Alarms are automatically cancelled in case the alarm state is no longer present.

vCenter Server alarms can be displayed without a custom vSphere plug-in. See the *vCenter Server Administration Guide* on events and alarms.

Upon registering as a vCenter Server extension, NSX Manager defines the rules that create and remove alarms, based on events coming from the three Guest Introspection components: SVM, Guest Introspection module, and thin agent. Rules can be customized. For instructions on how to create new custom rules for alarms, see the vCenter Server documentation. In some cases, there are multiple possible causes for the alarm. The tables that follow list the possible causes and the corresponding actions you might want to take for remediation.

Host Alarms

Host alarms are generated by events affecting the health status of the Guest Introspection module.

Table 18-1. Errors (Marked Red)

Possible Cause	Action
The Guest Introspection module has been installed on the host, but is no longer reporting status to the NSX Manager.	1 Ensure that Guest Introspection is running by logging in to the host and typing the command `/etc/init.d/vShield-Endpoint-Mux start`. 2 Ensure that the network is configured properly so that Guest Introspection can connect to NSX Manager. 3 Reboot the NSX Manager.

SVM Alarms

SVM alarms are generated by events affecting the health status of the SVM.

Table 18-2. Red SVM Alarms

Problem	Action
There is a protocol version mismatch with the Guest Introspection module	Ensure that the Guest Introspection module and SVM have a protocol that is compatible with each other.
Guest Introspection could not establish a connection to the SVM	Ensure that the SVM is powered on and that the network is configured properly.
The SVM is not reporting its status even though guests are connected.	Internal error. Contact your VMware support representative.

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`.

The following table lists Guest Introspection events reported by the SVM and the NSX Manager.

Table 18-3. Guest Introspection Events

Description	Severity	VC Arguments
Guest Introspection solution *SolutionName* enabled. Supporting version *versionNumber* of the VFile protocol.	info	timestamp
ESX module enabled.	info	timestamp
ESX module uninstalled.	info	timestamp
The NSX Manager has lost connection with the ESX module.	info	timestamp
Guest Introspection solution *SolutionName* was contacted by a non-compatible version of the ESX module.	error	timestamp, solution version, ESX module version
A connection between the ESX module and *SolutionName* failed.	error	timestamp, ESX module version, solution version
Guest Introspection failed to connect to the SVM.	error	timestamp
Guest Introspection lost connection with the SVM.	error	timestamp

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`.

Collecting Guest Introspection Troubleshooting Data

VMware Technical Support routinely requests diagnostic information or a support bundle when a support request is handled. This diagnostic information contains logs and configuration files for your virtual machines.

Identity Firewall Troubleshooting Data

If your identity-based firewall environment uses Guest Introspection, diagnostic information is found in the following the Knowledge Base articles: Troubleshooting vShield Endpoint / NSX Guest Introspection https://kb.vmware.com/kb/2094261 and Collecting logs in VMware NSX for vSphere 6.x Guest Introspection Universal Service Virtual Machine https://kb.vmware.com/kb/2144624.

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 NSX Data Security, Activity 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.

To uninstall Guest Introspection:

1. In vCenter, navigate to **Home > Networking & Security > Installation** and select the **Service Deployments** tab.
2. Select a Guest Introspection instance and click the delete icon.
3. Either delete now or schedule the deletion for a later time.

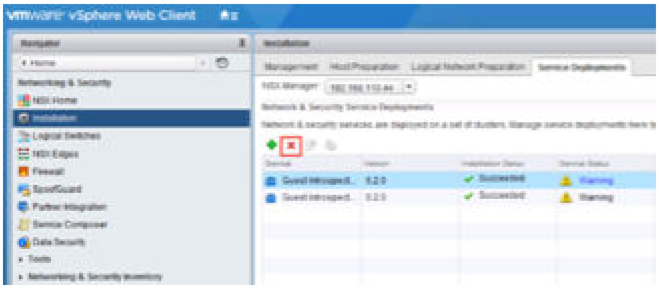

NSX Manager: 192.168.110.44

Network & Security Service Deployments

Network & security services are deployed on a set of clusters. Manage service deployments here

Service	Version	Installation Status	Service Status
Guest Introspecti...	6.2.0	✓ Succeeded	Warning
Guest Introspecti...	6.2.0	✓ Succeeded	Warning

Data Security 19

NSX Data Security provides visibility into sensitive data stored within your organization's virtualized and cloud environments. Based on the violations reported by NSX Data Security, you can ensure that sensitive data is adequately protected and assess compliance with regulations around the world.

Note As of NSX 6.2.3, the NSX Data Security feature has been deprecated. In NSX 6.2.3, you can continue to use this feature at your discretion, but be aware that this feature will be removed from NSX in a future release.

To begin using NSX Data Security, you create a policy that defines the regulations that apply to data security in your organization and specifies the areas of your environment and files to be scanned. A regulation is composed of content blades, which identify the sensitive content to be detected. NSX supports PCI, PHI, and PII related regulations only.

When you start a Data Security scan, NSX analyzes the data on the virtual machines in your vSphere inventory and reports the number of violations detected and the files that violated your policy.

This chapter includes the following topics:

Install NSX Data Security

Note As of NSX 6.2.3, the NSX Data Security feature has been deprecated. In NSX 6.2.3, you can continue to use this feature at your discretion, but be aware that this feature will be removed from NSX in a future release.

Prerequisites

NSX Guest Introspection must be installed on the cluster where you are installing Data Security.

If you want to assign an IP address to the Data Security service virtual machine from an IP pool, create the IP pool before installing Data Security. See Grouping Objects in the *NSX Administration Guide*.

Procedure

1 In the **Installation** tab, click **Service Deployments**.

2 Click the **New Service Deployment** (✚) icon.

3 In the Deploy Network and Security Services dialog box, select **Data Security** and click **Next**.

4 In **Specify schedule** (at the bottom of the dialog box), select **Deploy now** to deploy Data Security as soon as it is installed or select a deployment date and time.

5 Click **Next**.

6 Select the datacenter and cluster(s) where you want to install Data Security 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**.

The selected datastore must be available on all hosts in the selected cluster.

If you selected **Specified on host**, the datastore for the ESX host must be specified in the **AgentVM Settings** of the host before it is added to the cluster. See *vSphere API/SDK Documentation*.

8 Select the distributed virtual port group to host the management interface. This port group must be able to reach the NSX Manager's port group.

If the datastore is set to **Specified on host**, the network to be used must be specified in the **agentVmNetwork** property of each host in the cluster. See *vSphere API/SDK Documentation*.

When you add a host(s) to the cluster, the **agentVmNetwork** property for the host must be set before it is added to the cluster.

The selected port group must be available on all hosts in the selected cluster.

9 In IP assignment, select one of the following:

Select	To
DHCP	Assign an IP address to the Data Security service virtual machine through Dynamic Host Configuration Protocol (DHCP).
An IP pool	Assign an IP address to the Data Security service virtual machine from the selected IP pool.

Note that static IP address are not supported.

10 Click **Next** and then click **Finish** on the Ready to complete page.

11 Monitor the deployment until the **Installation Status** column displays **Succeeded**.

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. In some cases, resolving the errors displays additional errors. Take the required action and click **Resolve** again.

NSX Data Security User Roles

A user's role determines the actions that the user can perform.

Role	Actions Allowed
Security Administrator	Create and publish policies and view violation reports. Cannot start or stop a data security scan.
NSX Administrator	Start and stop data security scans.
Auditor	View configured policies and violation reports.

Defining a Data Security Policy

To detect sensitive data in your environment, you must create a data security policy. You must be a Security Administrator to create policies.

To define a policy, you must specify the following:

Regulations

A regulation is a data privacy law for protecting PCI (Payment Card Industry), PHI (Protected Health Information) and PII (Personally Identifiable Information) information. You can select the regulations that your company needs to comply to. When you run a scan, Data Security identifies data that violates the regulations in your policy and is sensitive for your organization.

File Filters

You can create filters to limit the data being scanned and exclude file types unlikely to contain sensitive data from the scan.

Select Regulations

After you select the regulations that you want your company data to comply with, NSX can identify files that contain information in violation of these regulations.

Prerequisites

You must have the Security Administrator role.

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking and Security** and then click **Data Security**.

3 Click the **Manage** tab.

4 Click **Edit** and click **All** to display all available regulations.

5 Select the regulations for which you want to detect compliance.

Note For information on available regulations, see the *Data Security Reference* guide.

Certain regulations require additional information for NSX Data Security to recognize sensitive data. If you selected a regulation that monitors Group Insurance Numbers, Patient Identification Numbers, Medical Record Numbers, Health Plan Beneficiary Numbers, US Bank Account Numbers, Custom Accounts, or Student identification numbers, specify a regular expression pattern for identifying that data.

6 Check the accuracy of the regular expression.

Specifying incorrect regular expressions can slow down the discovery process. For more information on regular expressions, see Creating Regular Expressions.

7 Click **Next**.

8 Click **Finish.**

9 Click **Publish Changes** to apply the policy.

Specify File Filters

You can restrict the files that you want to monitor based on size, last modified date, or file extensions.

Prerequisites

You must have been assigned the Security Administrator role.

Procedure

1 In the **Manage** tab of the Data Security panel, click **Edit** next to **Files to scan**.

2 You can either monitor all files on the virtual machines in your inventory, or select the restrictions you want to apply.

Option	Description
Monitor all files on the guest virtual machines	NSX Data Security scans all files.
Monitor only the files that match the following conditions	Select the following options as appropriate. ■ **Size** indicates that NSX Data Security should only scan files less than the specified size. ■ **Last Modified Date** indicates that NSX Data Security should scan only files modified between the specified dates. ■ **Types:** Select **Only files with the following extensions** to enter the file types to scan. Select **All files, except those with extensions** to enter the file types to exclude from the scan.

For information on file formats that NSX Data Security can detect, see the *Data Security Reference* guide.

3 Click **Save.**

4 Click **Publish Changes** to apply the policy.

Running a Data Security Scan

Running a data security scan identifies data in your virtual environment that violates your policy.

Prerequisites

You must be a NSX Administrator to start, pause, or stop a data security scan.

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking and Security** and then click **Data Security**.

3 Click the **Manage** tab.

4 Click **Start** next to Scanning.

   Note If a virtual machine is powered off, it will not be scanned until it is powered on.

   If a scan is in progress, the available options are **Pause** and **Stop**.

If Data Security is part of a Service Composer policy, virtual machines in the security group mapped to that Service Composer policy are scanned once during a scan. If the policy is edited and published while a scan is running, the scan restarts. This rescan ensures that all virtual machines comply with the edited policy. A rescan is triggered by publishing an edited policy, not by data updates on your virtual machines.

If a virtual machine is moved to an excluded cluster or resource pool while the data security scan is in progress, the files on that virtual machine are not scanned. In case a virtual machine is moved by vMotion to another host, the scan continues on the second host. Files that were scanned while the virtual machine was on the previous host are not rescanned.

When the Data Security engine starts scanning a virtual machine, it records the scan start time. When the scan ends, it records the end of the scan. You can view the scan start and end time for a cluster, host, or virtual machine on the **Tasks and Events** tab.

NSX Data Security throttles the number of virtual machines concurrently scanned on a host to minimize impact on performance. VMware recommends that you pause the scan during normal business hours to avoid any performance overhead.

Viewing and Downloading Reports

When you start a security scan, NSX displays the start and end time of each scan, the number of virtual machines scanned, and the number of violations detected.

Prerequisites

You have the Security Administrator or Auditor role.

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking and Security** and then click **Data Security**.

3 Click the **Reports** tab.

4 Specify the report for Violation counts or for Violating files.

Creating Regular Expressions

A regular expression is a pattern that describes a certain sequence of text characters, otherwise known as strings. You use regular expressions to search for, or match, specific strings or classes of strings in a body of text.

Using a regular expression is like performing a wildcard search, but regular expressions are far more powerful. Regular expressions can be very simple or very complex. An example of a simple regular expression is *cat.*

This finds the first instance of the letter sequence cat in any body of text that you apply it to. If you want to make sure it only finds the word *cat*, and not other strings like *cats* or *hepcat*, you could use this slightly more complex regular expression: *\bcat\b*.

This expression includes special characters that ensure a match occurs only if there are word breaks on both sides of the *cat* sequence. As another example, to perform a near equivalent to the typical wildcard search string *c+t*, you could use this regular expression: *\bc\w+t\b*.

This means find a word boundary (\b) followed by a *c*, followed by one or more non-whitespace characters, non-punctuation characters (\w+), followed by a *t*, followed by a word boundary (\b). This expression finds *cot*, *cat*, *croat*, but not *crate.*

Expressions can be very complex. The following expression finds any valid email address.

\b[A-Za-z0-9._%-]+@[A-Za-z0-9.-]+\.[A-Za-z]{2,4}\b

For more information on creating regular expressions, see http://userguide.icu-project.org/strings/regexp.

Uninstall NSX Data Security

Uninstall NSX data security either because you are no longer using it or because you are upgrading NSX Manager. NSX data security does not support a direct upgrade. Before upgrading NSX Manager, it is important to first uninstall NSX data security and then reinstall it after the upgrade.

As of NSX 6.2.3, the NSX Data Security feature has been deprecated. In NSX 6.2.3, you can continue to use this feature at your discretion, but be aware that this feature will be removed from NSX in a future release.

Procedure

1 In the **Installation** tab, click **Service Deployments**.

2 Select the NSX Data Security service and click the **Delete Service Deployment** (✖) icon.

3 In the Confirm Delete dialog box, click **Delete now** or select a date and time for the delete to take effect.

4 Click **OK**.

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 ESX Server host or across multiple ESX server hosts for production, testing, or development purposes.

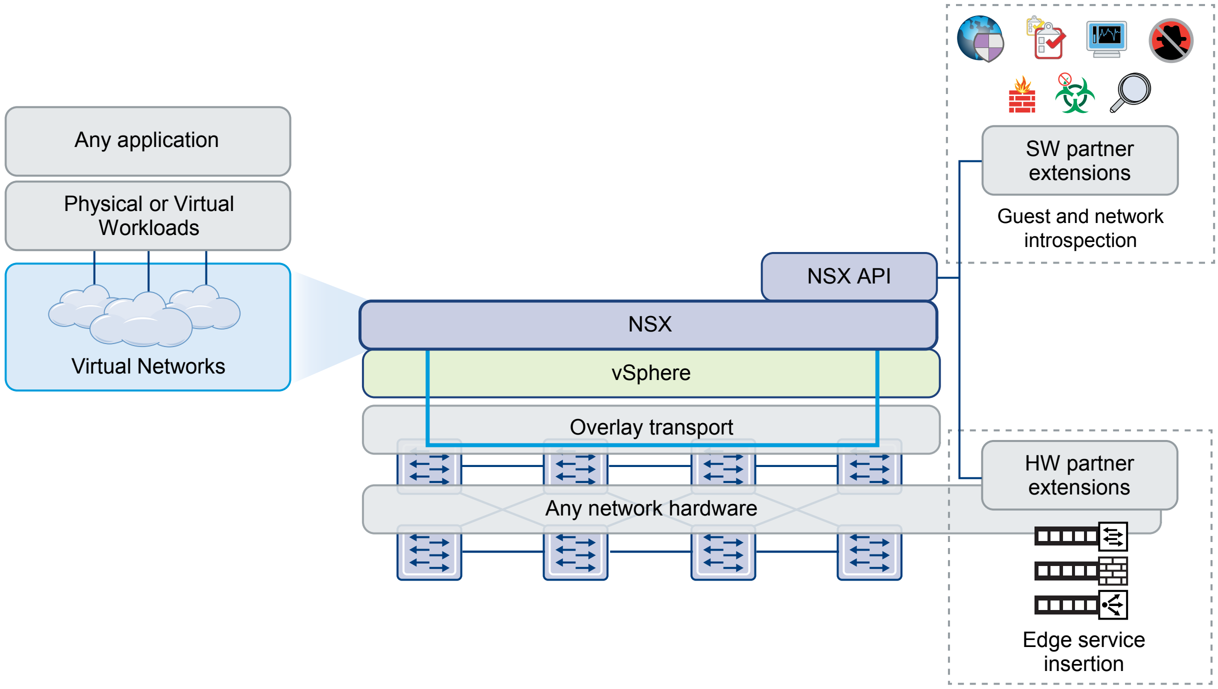

There are various deployment methods for inserting third party services into NSX.

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 3rd-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.

Prerequisites

Ensure that:

- The partner solution is registered with NSX Manager.
- NSX Manager can access the partner solution's management console.

Procedure

1. Click **Networking & Security** and then click **Installation**.
2. Click the **Service Deployments** tab and click the **New Service Deployment** (➕) icon.
3. In the Deploy Network and Security Services dialog box, select the appropriate solution(s).
4. In **Specify schedule** (at the bottom of the dialog box), select **Deploy now** to deploy the solution immediately or select a deployment date and time.
5. Click **Next**.
6. Select the datacenter and cluster(s) where you want to deploy the solution and click **Next**.
7. Select the datastore on which to add the solution service virtual machines storage or select **Specified on host**.

   The selected datastore must be available on all hosts in the selected cluster.

   If you selected **Specified on host**, the datastore for the ESX host must be specified in the **AgentVM Settings** of the host before it is added to the cluster. See *vSphere API/SDK Documentation*.
8. Select the distributed virtual port group to host the management interface. This port group must be able to reach the NSX Manager’s port group.

   If the network is set to **Specified on host**, the network to be used must be specified in the **Agent VM Settings > Network** property of each host in the cluster. See *vSphere API/SDK Documentation*.

   You must set the agent VM network property on a host before you add it to a cluster. Navigate to **Manage > Settings > Agent VM Settings > Network** and click **Edit** to set the agent VM network.

   The selected port group must be available on all hosts in the selected cluster.
9. In IP assignment, select one of the following:

Select	To
DHCP	Assign an IP address to the service virtual machine through Dynamic Host Configuration Protocol (DHCP).
An IP pool	Assign an IP address to the 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** displays Successful. If the status displays Failed, click the icon next to Failed and take action to resolve the error.

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.

Traffic for the specified Edge is load balanced by the third party vendor's management console.

Remove 3rd-Party Integration

This example describes how to remove a 3rd-party integration solution from NSX.

There is a correct order of software when removing any 3rd-party software solution.

Procedure

1. In the Sphere Web Client, navigate to **Networking & Security > Service Composer**, and delete the rules (or security polices) that are redirecting traffic to the 3rd-party solution.
2. Navigate to **Service Definitions** and double-click the name of the 3rd-party solution.
3. Click **Related Objects** and delete the related objects.
4. Navigate to **Installation > Service Deployments** and delete the 3rd-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.

The 3rd-party integration solution is removed from NSX.

What to do next

Make notes of the configuration settings, and then remove NSX from the 3rd-party solution. For example, you may need to delete rules that reference other objects and then delete the objects.

User Management

21

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

To deploy and administer NSX, certain vCenter permissions are required. NSX provides extensive read and read/write permissions for various users and roles.

Roles Definition

The available roles are as follows:

roles = system_write, system_urm, super_user, vshield_admin, security_admin, auditor, dlp_svm, epsec_host, enterprise_admin, component_manager_user, replicator local_user_roles = system_write, system_urm, super_user, security_admin, auditor, dlp_svm, epsec_host, component_manager_user, replicator

system_roles = system_write, system_urm, dlp_svm, epsec_host, replicator

Permission Types

The permission types are read and write.

Roles Access Definition

The role access definitions determine whether a role has read or read/write permission.

super_user.object_permission = read, write

vshield_admin.object_permission = read, write

security_admin.object_permission = read, write

auditor.object_permission = read

system_write.object_permission = read, write

system_urm.object_permission = read

dlp_svm.object_permission = read, write

epsec_host.object_permission = read, write

enterprise_admin.object_permission = read, write

replicator.object_permission = read, write

Root Definition

The root definition describes the superuser roles.

super_user.superuser = true

system_write.superuser = true

Role to Object Access for Global Scope

vshield_admin.object_access_scope.global = true

super_user.object_access_scope.global = true

system_write.object_access_scope.global = true

system_urm.object_access_scope.global = true

dlp_svm.object_access_scope.global = true

epsec_host.object_access_scope.global = true

enterprise_admin.object_access_scope.global = true

Role to Object Access for Universal Scope

replicator.object_access_scope.universal=true

system_write.object_access_scope.universal=true

Services

The following services are available in NSX:

administration, urm, edge, app, namespace, spoofguard, dlp, epsec, library, install, vdn, eam, si, truststore, component_manager, ipam, secfabric, security_policy, messaging, replicator

Feature Definitions

The feature definitions within each service are as follows:

administration.featurelist = administration.configuration, administration.update, administration.system_events, administration.audit_logs, administration.debug

urm.featurelist = urm.user_account_management, urm.object_access_control, urm.feature_access_control

edge.featurelist = edge.system, edge.nat, edge.firewall, edge.dhcp, edge.loadbalancer, edge.vpn, edge.syslog, edge.support, edge.routing, edge.certificate, edge.appliance, edge.highavailability, edge.dns, edge.vnic, edge.ssh, edge.autoplumbing, edge.statistics, edge.bridging, edge.systemcontrol

app.featurelist = app.config, app.firewall, app.flow, app.forcesync, app.syslog, app.techsupport

pgi.featurelist = pgi.switch, pgi.portgroup, pgi.lkm

namespace.featurelist = namespace.config

spoofguard.featurelist = spoofguard.config

dlp.featurelist = dlp.scan_scheduling, dlp.reports, dlp.policy, dlp.svm_interaction

epsec.featurelist = epsec.registration, epsec.health_monitoring, epsec.manager, epsec.policy, epsec.svm_priv, epsec.scan, epsec.reports

library.featurelist = library.grouping, library.host_preparation, library.tagging

install.featurelist = install.app, install.epsec, install.dlp

vdn.featurelist = vdn.config_nsm, vdn.provision

eam.featurelist = eam.install

si.featurelist = si.service, si.serviceprofile

truststore.featurelist = truststore.trustentity_management

component_manager.featurelist = healthstatus

ipam.featurelist = ipam.configuration, ipam.ipallocation secfabric.featurelist = secfabric.deploy, secfabric.alarms

security_policy.featurelist = security_policy.configuration, security_policy.security_group_binding

blueprint_sam.featurelist = blueprint_sam.reports, blueprint_sam.ad_config, blueprint_sam.control_data_collection, blueprint_sam.techsupport, blueprint_sam.db_maintain

messaging.featurelist = messaging.messaging

replicator.featurelist = replicator.configuration

Feature Access Definitions

For each feature and role combination, the feature access definition denotes whether the user has read-only or read/write permissions.

When a feature and role combination is not listed, this means the user with that role has no access to this feature.

For example:

auditor.app.firewall = read

security_admin.app.firewall = read, write

This means the auditor role on the app.firewall feature has read-only access, whereas the security_admin role on the app.firewall feature has read/write access.

Feature Access Definitions - system_urm

system_urm.urm.user_account_management = read

Feature Access Definitions - vshield_admin

vshield_admin.administration.configuration = read, write

vshield_admin.administration.update = read, write

vshield_admin.administration.system_events = read, write

vshield_admin.administration.audit_logs = read

vshield_admin.urm.user_account_management = read, write

vshield_admin.urm.object_access_control = read

vshield_admin.urm.feature_access_control = read

vshield_admin.edge.system = read, write

vshield_admin.edge.appliance = read, write

vshield_admin.edge.highavailability = read, write

vshield_admin.edge.vnic = read, write

vshield_admin.edge.dns = read

```
vshield_admin.edge.ssh = read, write
vshield_admin.edge.autoplumbing = read
vshield_admin.edge.statistics = read
vshield_admin.edge.nat = read
vshield_admin.edge.dhcp = read
vshield_admin.edge.loadbalancer = read
vshield_admin.edge.vpn = read
vshield_admin.edge.syslog = read, write
vshield_admin.edge.support = read, write
vshield_admin.edge.routing = read
vshield_admin.edge.firewall = read
vshield_admin.edge.bridging = read
vshield_admin.edge.certificate = read
vshield_admin.edge.systemcontrol = read, write
vshield_admin.library.grouping = read
vshield_admin.app.config = read, write
vshield_admin.app.forcesync = read, write
vshield_admin.app.syslog = read, write
vshield_admin.app.techsupport = read, write
vshield_admin.namespace.config = read, write
vshield_admin.dlp.scan_scheduling = read, write
vshield_admin.epsec.reports = read, write
vshield_admin.epsec.registration = read, write
vshield_admin.epsec.health_monitoring = read
vshield_admin.epsec.policy = read, write
vshield_admin.epsec.scan_scheduling = read, write
vshield_admin.library.host_preparation = read, write
vshield_admin.library.tagging = read
vshield_admin.install.app = read, write
vshield_admin.install.epsec = read, write
vshield_admin.install.dlp = read, write
vshield_admin.vdn.config_nsm = read, write
```

vshield_admin.vdn.provision = read, write

vshield_admin.eam.install = read, write

vshield_admin.si.service = read, write

vshield_admin.si.serviceprofile = read, write

vshield_admin.truststore.trustentity_management = read, write

vshield_admin.ipam.configuration = read, write

vshield_admin.ipam.ipallocation = read, write

vshield_admin.secfabric.deploy = read, write

vshield_admin.secfabric.alarms = read_write

vshield_admin.blueprint_sam.ad_config = read, write

vshield_admin.blueprint_sam.control_data_collection = read, write

vshield_admin.blueprint_sam.techsupport = read, write

vshield_admin.blueprint_sam.db_maintain = read, write

vshield_admin.messaging.messaging = read, write

vshield_admin.replicator.configuration = read, write

Feature Access Definitions - security_admin

security_admin.administration.system_events = read, write

security_admin.administration.audit_logs = read

security_admin.edge.system = read

security_admin.edge.appliance = read

security_admin.edge.highavailability = read

security_admin.edge.vnic = read, write

security_admin.edge.dns = read, write

security_admin.edge.ssh = read, write

security_admin.edge.autoplumbing = read, write

security_admin.edge.statistics = read

security_admin.edge.nat = read, write

security_admin.edge.dhcp = read, write

security_admin.edge.loadbalancer = read, write

security_admin.edge.vpn = read, write

security_admin.edge.syslog = read, write security_admin.edge.support = read, write

security_admin.edge.routing = read, write

security_admin.edge.firewall = read, write

security_admin.edge.bridging = read, write

security_admin.edge.certificate = read, write

security_admin.edge.systemcontrol = read, write

security_admin.app.firewall = read, write

security_admin.app.flow = read, write

security_admin.app.forcesync = read

security_admin.app.syslog = read

security_admin.namespace.config = read

security_admin.spoofguard.config = read, write

security_admin.dlp.reports = read, write

security_admin.dlp.policy = read, write

security_admin.epsec.policy = read, write

security_admin.epsec.reports = read

security_admin.epsec.health_monitoring = read

security_admin.library.grouping = read, write

security_admin.library.tagging = read, write

security_admin.install.app = read

security_admin.install.epsec = read

security_admin.install.dlp = read

security_admin.vdn.config_nsm = read

security_admin.vdn.provision = read

security_admin.eam.install = read

security_admin.si.service = read, write

security_admin.si.serviceprofile = read

security_admin.truststore.trustentity_management = read, write

security_admin.ipam.configuration = read, write

security_admin.ipam.ipallocation = read, write

security_admin.secfabric.alarms = read

security_admin.secfabric.deploy = read

security_admin.security_policy.configuration = read, write

security_admin.security_policy.security_group_binding = read, write

security_admin.blueprint_sam.reports = read

security_admin.blueprint_sam.ad_config = read

security_admin.blueprint_sam.control_data_collection = read

security_admin.blueprint_sam.db_maintain = read

security_admin.messaging.messaging = read, write

security_admin.replicator.configuration = read

Feature Access Definitions - auditor

auditor.administration.system_events = read

auditor.administration.audit_logs = read

auditor.edge.appliance = read

auditor.edge.highavailability = read

auditor.edge.vnic = read

auditor.edge.dns = read

auditor.edge.ssh = read

auditor.edge.autoplumbing = read

auditor.edge.statistics = read

auditor.edge.nat = read

auditor.edge.dhcp = read

auditor.edge.loadbalancer = read

auditor.edge.vpn = read

auditor.edge.syslog = read

auditor.edge.routing = read

auditor.edge.firewall = read

auditor.edge.bridging = read

auditor.edge.system = read

auditor.edge.certificate = read

auditor.edge.systemcontrol = read

auditor.app.firewall = read

auditor.app.flow = read

```
auditor.app.forcesync = read

auditor.app.syslog = read

auditor.namespace.config = read

auditor.spoofguard.config = read

auditor.dlp.scan_scheduling = read

auditor.dlp.policy = read

auditor.dlp.reports = read

auditor.library.grouping = read

auditor.epsec_host.health_monitoring = read

auditor.epsec.policy = read

auditor.epsec.reports = read

auditor.epsec.registration = read

auditor.vdn.config_nsm = read

auditor.epsec.scan_scheduling = read

auditor.vdn.provision = read

auditor.si.service = read

auditor.si.serviceprofile = read

auditor.truststore.trustentity_management = read

auditor.secfabric.alarms = read

auditor.secfabric.deploy = read

auditor.security_policy.configuration = read

auditor.security_policy.security_group_binding = read

auditor.blueprint_sam.reports = read

auditor.blueprint_sam.ad_config = read

auditor.blueprint_sam.control_data_collection = read

auditor.blueprint_sam.db_maintain = read

auditor.library.tagging = read

auditor.ipam.configuration = read

auditor.ipam.ipallocation = read

auditor.messaging.messaging = read

auditor.replicator.configuration = read
```

Feature Access Definitions - dlp_svm

dlp_svm.dlp.svm_interaction = read, write

dlp_svm.epsec.svm_priv = read, write

dlp_svm.epsec.registration = read

dlp_svm.epsec.policy = read

dlp_svm.epsec.scan_scheduling = read

dlp_svm.library.host_preparation = read, write

dlp_svm.library.tagging = read, write

Feature Access Definitions - epsec_host

epsec_host.epsec.registration = read

epsec_host.epsec.health_monitoring = write

Feature Access Definitions - enterprise_admin

enterprise_admin.administration.configuration = read, write

enterprise_admin.administration.update = read, write

enterprise_admin.administration.system_events = read, write

enterprise_admin.administration.audit_logs = read

enterprise_admin.urm.user_account_management = read, write

enterprise_admin.urm.object_access_control = read

enterprise_admin.urm.feature_access_control = read

enterprise_admin.edge.system = read, write

enterprise_admin.edge.appliance = read, write

enterprise_admin.edge.highavailability = read, write

enterprise_admin.edge.vnic = read, write

enterprise_admin.edge.dns = read, write

enterprise_admin.edge.ssh = read, write

enterprise_admin.edge.autoplumbing = read, write

enterprise_admin.edge.statistics = read, write

enterprise_admin.edge.nat = read, write

enterprise_admin.edge.dhcp = read, write enterprise_admin.edge.loadbalancer = read, write

enterprise_admin.edge.vpn = read, write

enterprise_admin.edge.syslog = read, write

enterprise_admin.edge.support = read, write

enterprise_admin.edge.routing = read, write

enterprise_admin.edge.firewall = read, write

enterprise_admin.edge.bridging = read, write

enterprise_admin.edge.certificate = read, write

enterprise_admin.edge.systemcontrol = read, write

enterprise_admin.library.grouping = read, write

enterprise_admin.library.host_preparation = read, write

enterprise_admin.library.tagging = read, write

enterprise_admin.app.config = read, write

enterprise_admin.app.forcesync = read, write

enterprise_admin.app.syslog = read, write

enterprise_admin.app.techsupport = read, write

enterprise_admin.app.firewall = read, write

enterprise_admin.app.flow = read, write

enterprise_admin.namespace.config = read, write

enterprise_admin.dlp.scan_scheduling = read, write

enterprise_admin.dlp.reports = read, write

enterprise_admin.dlp.policy = read, write

enterprise_admin.epsec.registration = read, write

enterprise_admin.epsec.health_monitoring = read

enterprise_admin.epsec.scan_scheduling = read, write

enterprise_admin.epsec.reports = read, write

enterprise_admin.epsec.policy = read, write

enterprise_admin.install.app = read, write

enterprise_admin.install.epsec = read, write

enterprise_admin.install.dlp = read, write

enterprise_admin.eam.install = read, write

enterprise_admin.spoofguard.config = read, write enterprise_admin.vdn.config_nsm = read, write

enterprise_admin.vdn.provision = read, write

enterprise_admin.si.service = read, write

enterprise_admin.si.serviceprofile = read, write

enterprise_admin.truststore.trustentity_management = read, write

enterprise_admin.ipam.configuration = read, write

enterprise_admin.ipam.ipallocation = read, write

enterprise_admin.secfabric.deploy = read, write

enterprise_admin.secfabric.alarms = read, write

enterprise_admin.security_policy.configuration = read, write

enterprise_admin.security_policy.security_group_binding = read, write

enterprise_admin.blueprint_sam.reports = read

enterprise_admin.blueprint_sam.ad_config = read, write

enterprise_admin.blueprint_sam.control_data_collection = read, write

enterprise_admin.blueprint_sam.techsupport = read, write

enterprise_admin.blueprint_sam.db_maintain = read, write

enterprise_admin.messaging.messaging = read, write

enterprise_admin.replicator.configuration = read, write

Feature Access Definitions - component_manager_user

component_manager_user.component_manager.healthstatus = read

Feature Access Definitions - replicator

replicator.administration.configuration = read, write

replicator.administration.update = read, write

replicator.administration.system_events = read, write

replicator.administration.audit_logs = read

replicator.urm.user_account_management = read, write

replicator.urm.object_access_control = read

replicator.urm.feature_access_control = read

replicator.edge.system = read, write

replicator.edge.appliance = read, write

replicator.edge.highavailability = read

replicator.edge.vnic = read, write

replicator.edge.dns = read

replicator.edge.ssh = read

replicator.edge.autoplumbing = read, write

replicator.edge.statistics = read

replicator.edge.nat = read

replicator.edge.dhcp = read, write

replicator.edge.loadbalancer = read

replicator.edge.vpn = read

replicator.edge.syslog = read

replicator.edge.support = read

replicator.edge.routing = read, write

replicator.edge.firewall = read

replicator.edge.bridging = read

replicator.edge.certificate = read

replicator.edge.systemcontrol = read

replicator.library.grouping = read, write

replicator.library.host_preparation = read, write

replicator.library.tagging = read, write

replicator.app.config = read, write

replicator.app.forcesync = read, write

replicator.app.syslog = read, write

replicator.app.techsupport = read, write

replicator.app.firewall = read, write

replicator.app.flow = read, write

replicator.namespace.config = read, write

replicator.dlp.scan_scheduling = read, write

replicator.dlp.reports = read, write

replicator.dlp.policy = read, write

replicator.epsec.registration = read, write

replicator.epsec.health_monitoring = read replicator.epsec.scan_scheduling = read, write

replicator.epsec.reports = read, write

replicator.epsec.policy = read, write

replicator.install.app = read, write

replicator.install.epsec = read, write

replicator.install.dlp = read, write

replicator.eam.install = read, write

replicator.spoofguard.config = read, write

replicator.vdn.config_nsm = read, write

replicator.vdn.provision = read, write

replicator.si.service = read, write

replicator.si.serviceprofile = read, write

replicator.truststore.trustentity_management = read, write

replicator.ipam.configuration = read, write

replicator.ipam.ipallocation = read, write

replicator.secfabric.deploy = read, write

replicator.secfabric.alarms = read, write

replicator.security_policy.configuration = read, write

replicator.security_policy.security_group_binding = read, write

replicator.blueprint_sam.reports = read

replicator.blueprint_sam.ad_config = read, write

replicator.blueprint_sam.control_data_collection = read, write

replicator.blueprint_sam.techsupport = read, write

replicator.blueprint_sam.db_maintain = read, write

replicator.messaging.messaging = read, write

replicator.replicator.configuration = read, write

Overwrite Role Feature Permissions on Secondary Node on Universal Objects

secondary.super_user.edge.highavailability = read, write

secondary.enterprise_admin.edge.highavailability = read, write

secondary.vshield_admin.edge.highavailability = read, write

```
secondary.super_user.edge.ssh = read, write
secondary.enterprise_admin.edge.ssh = read, write
secondary.security_admin.edge.ssh = read, write
secondary.vshield_admin.edge.ssh = read, write
secondary.super_user.edge.syslog = read, write
secondary.enterprise_admin.edge.syslog = read, write
secondary.security_admin.edge.syslog = read, write
secondary.vshield_admin.edge.syslog = read, write
secondary.super_user.edge.support = read, write
secondary.enterprise_admin.edge.support = read, write
secondary.security_admin.edge.support = read, write
secondary.vshield_admin.edge.support = read, write
secondary.super_user.edge.routing = read, write
secondary.security_admin.edge.routing = read, write
secondary.enterprise_admin.edge.routing = read, write
secondary.super_user.edge.appliance = read, write
secondary.vshield_admin.edge.appliance = read, write
secondary.enterprise_admin.edge.appliance = read, write
secondary.super_user.edge.vnic = read, write
secondary.vshield_admin.edge.vnic = read, write
secondary.enterprise_admin.edge.vnic = read, write
secondary.super_user.edge.firewall = read, write
secondary.vshield_admin.edge.firewall = read, write
secondary.enterprise_admin.edge.firewall = read, write
```

Configure Single Sign On

SSO makes vSphere and NSX 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 improves the security of user authentication for vCenter users and enables NSX to authenticate users from other identity services such as AD, NIS, and LDAP.

With SSO, NSX supports authentication using authenticated Security Assertion Markup Language (SAML) tokens from a trusted source via REST API calls. NSX Manager can also acquire authentication SAML tokens for use with other VMware solutions.

NSX caches group information for SSO users. Changes to group memberships will take up to 60 minutes to propagate from the identity provider (for example, active directory) to NSX.

Prerequisites

- To use SSO on NSX Manager, you must have vCenter Server 5.5 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 http://kb.vmware.com/kb/2072435 and http://kb.vmware.com/kb/2113115.

- NTP server must be specified so that the SSO server time and NSX Manager time is in sync.

  For example:

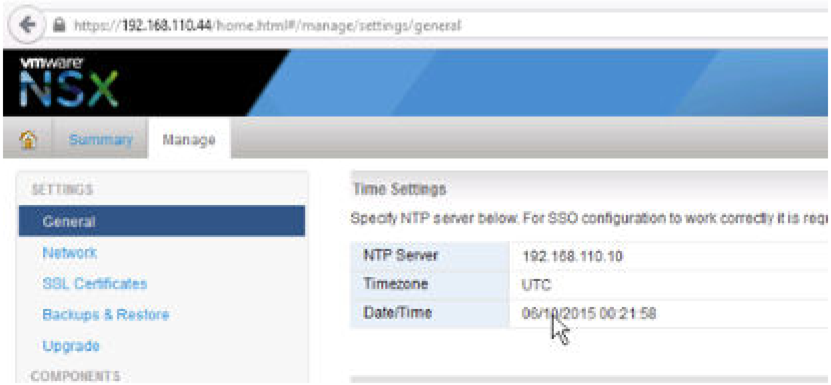

https://192.168.110.44/home.html#/manage/settings/general

Time Settings

Specify NTP server below. For SSO configuration to work correctly it is requ

NTP Server	192.168.110.10
Timezone	UTC
Date/Time	06/10/2015 00:21:58

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 with the password that you configured during NSX Manager installation.

2 Click the **Manage** tab, then click **NSX Management Service**.

3 Type the name or IP address of the host that has the lookup service.

If you are using vCenter to perform the lookup service, enter the vCenter Server's IP address or hostname, and enter the vCenter Server user name and password.

4 Type the port number.

Enter port 443 if you are using vSphere 6.0. For vSphere 5.5, use port number 7444.

The Lookup Service URL is displayed based on the specified host and port.

For example:

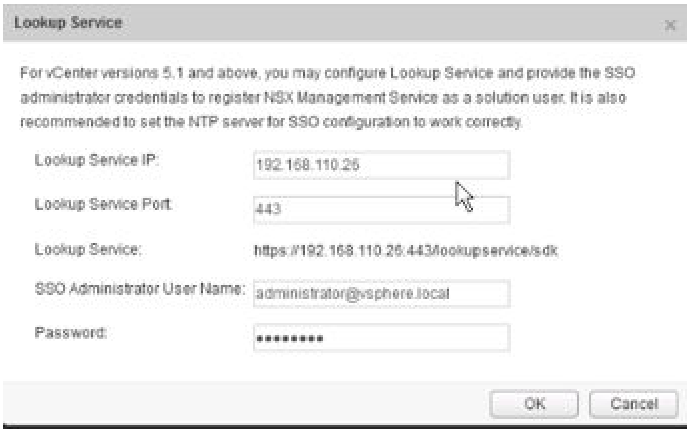

Lookup Service

For vCenter versions 5.1 and above, you may configure Lookup Service and provide the SSO administrator credentials to register NSX Management Service as a solution user. It is also recommended to set the NTP server for SSO configuration to work correctly.

Lookup Service IP: 192.168.110.26

Lookup Service Port: 443

Lookup Service: https://192.168.110.26:443/lookupservice/sdk

SSO Administrator User Name: administrator@vsphere.local

Password: *********

5 Check that the certificate thumb print matches the certificate of the vCenter 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.

6 Confirm that the Lookup Service status is **Connected**.

For example:

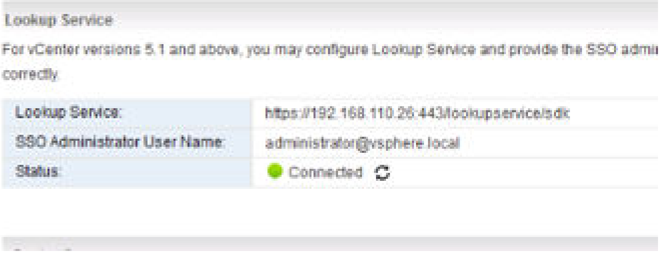

Lookup Service

For vCenter versions 5.1 and above, you may configure Lookup Service and provide the SSO admi
correctly.

Lookup Service:	https://192.168.110.26:443/lookupservice/sdk
SSO Administrator User Name:	administrator@vsphere.local
Status:	● Connected

What to do next

Assign a role to the SSO user.

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.

The following rules are enforced:

- A user can have only one role.
- You cannot add a role to a user or remove an assigned role from a user. You can, however, change the assigned role for a user.

Table 21-1. NSX Manager User Roles

Right	Permissions
Enterprise Administrator	NSX operations and security.
NSX Administrator	NSX operations only: for example, install virtual appliances, configure port groups.
Security Administrator	NSX security only: for example, define data security policies, create port groups, create reports for NSX modules.
Auditor	Read only.

The Enterprise Administrator and NSX Administrator roles can be assigned only to vCenter users.

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 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 configuration.

1 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **NSX Managers**.

3 Click an NSX Manager in the Name column and then click the **Manage** tab.

4 Click **Users**.

5 Click **Add**.

The Assign Role window opens.

6 Click **Specify a vCenter user** or **Specify a vCenter group**.

7 Type the vCenter **User** or **Group** name for the user.

Refer to the example below for more information.

Domain name: corp.vmware.com

Alias: corp

Group name: group1@corp.vmware.com

User name : user1@corp.vmware.com

When a group is assigned a role on the NSX Manager any user from that group can be logged in to the NSX Manager user interface.

When assigning a role to a user, type the user alias. For example, user1@corp.

8 Click **Next**.

9 Select the role for the user and click **Next**. For more information on the available roles, see Managing User Rights.

10 Click **Finish**.

The user account appears in the Users table.

Understanding 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 may 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 has been registered with NSX Manager, and SSO has been configured.

1 Assign a role to Sally.

   a Log in to the vSphere Web Client.

   b Click **Networking & Security** and then click **NSX Managers**.

   c Click an NSX Manager in the Name column and then click the **Manage** tab.

   d Click **Users** and then click **Add**.

   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 datacenter.

   a Click the Home icon and then click **vCenter Home > Datacenters**.

   b Select a datacenter and click **Actions > All vCenter Actions > Add Permission**.

   c Click **Add** and select the domain CORP.

   d In **Users and Groups**, select **Show Groups First**.

e Select NetEng and click **OK**.

f In **Assigned Role**, select **Read-only** and un-select **Propagate to children** and click **OK**.

3 Log out of vSphere Web Client and log back in as smoore@corp.local.

Sally can perform NSX operations only. For example, install virtual appliances, create logical switches, and so on..

Example: Inherit Permissions Through a User-Group Membership Scenario

Group option	Value
Name	G1
Role assigned	Auditor (Read only)
Resources	Global root

User option	Value
Name	John
Belongs to group	G1
Role assigned	None

John belongs to group G1, which has been assigned the auditor role. John inherits the group role and resource permissions.

Example: User Member of Multiple Groups Scenario

Group option	Value
Name	G1
Role assigned	Auditor (Read only)
Resources	Global root

Group option	Value
Name	G2
Role assigned	Security Administrator (Read and Write)
Resources	Datacenter1

User option	Value
Name	Joseph
Belongs to group	G1, G2
Role assigned	None

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, John has the following permissions:

- Read, write (Security Administrator role) for Datacenter1
- Read only (Auditor) for global root

Example: User Member of Multiple Roles Scenario

Group option	Value
Name	G1
Role assigned	Enterprise Administrator
Resources	Global root

User option	Value
Name	Bob
Belongs to group	G1
Role assigned	Security Administrator (Read and Write)
Resources	Datacenter1

Bob has been 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

Edit a User Account

You can edit a user account to change the role or scope. You cannot edit the **admin** account.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then under **Networking & Security Inventory** click **NSX Managers**.
3. Click an NSX Manager in the Name column and then click the **Manage** tab.
4. Click **Users**.
5. Select the user you want to edit.
6. Click **Edit**.
7. Make changes as necessary.
8. 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. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then under **Networking & Security Inventory** click **NSX Managers**.
3. Click an NSX Manager in the Name column and then click the **Manage** tab.
4. Click **Users**.
5. Select the user you want to change the role for.
6. Click **Change Role**.
7. Make changes as necessary.
8. 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. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then under **Networking & Security Inventory** click **NSX Managers**.
3. Click an NSX Manager in the Name column and then click the **Manage** tab.
4. Click **Users**.
5. Select a user account.
6. 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. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then under **Networking & Security Inventory** click **NSX Managers**.

3 Click an NSX Manager in the Name column and then click the **Manage** tab.

4 Click **Users**.

5 Select a user account.

6 Click **Delete**.

7 Click **OK** to confirm deletion.

If you delete a vCenter user account, only the role assignment for NSX Manager is deleted. The user account on vCenter is not deleted.

Network and Security Objects 22

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.

This chapter includes the following topics:

Working with IP Address Groups

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 vice versa.

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then under **Networking & Security Inventory** click **NSX Managers**.

3 Click an NSX Manager in the **Name** column and then click the **Manage** tab.

   ◆ You must select the primary NSX Manager if you need to manage universal IP address groups.

4 Click the **Grouping Objects** tab, then click **IP Sets**.

5 Click the **Add** (➕) icon.

6 Type a name for the address group.

7 (Optional) Type a description for the address group.

8. Type the IP addresses to be included in the group.
9. (Optional) Select **Enable inheritance to allow visibility at underlying scopes.**
10. (Optional) Select **Mark this object for Universal Synchronization** to create a universal IP address group.
11. Click **OK**.

Edit an IP Address Group

Prerequisites

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then under **Networking & Security Inventory** click **NSX Managers**.
3. Click an NSX Manager in the **Name** column and then click the **Manage** tab.
   - You must select the primary NSX Manager if you need to manage universal IP address groups.
4. Click the **Grouping Objects** tab, then click **IP Sets**.
5. Select the group that you want to edit and click the **Edit** (✏) icon.
6. In the Edit IP Sets dialog box, make the appropriate changes.
7. Click **OK**.

Delete an IP Address Group

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then under **Networking & Security Inventory** click **NSX Managers**.
3. Click an NSX Manager in the **Name** column and then click the **Manage** tab.
   - You must select the primary NSX Manager if you need to manage universal IP address groups.
4. Click the **Grouping Objects** tab and then click **IP Sets**.
5. Select the group that you want to delete and click the **Delete** (✖) icon.

Working with MAC Address Groups

Create a MAC Address Group

You can create a MAC address group 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 vice versa.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then under **Networking & Security Inventory** click **NSX Managers**.
3. Click an NSX Manager in the **Name** column and then click the **Manage** tab.
   - You must select the primary NSX Manager if you need to manage universal MAC address groups.
4. Click the **Grouping Objects** tab and then click **MAC Sets**.
5. Click the **Add** (✚) icon.
6. Type a name for the address group.
7. (Optional) Type a description for the address group.
8. Type the MAC addresses to be included in the group.
9. (Optional) Select **Enable inheritance to allow visibility at underlying scopes.**
10. (Optional) Select **Mark this object for Universal Synchronization** to create a universal MAC address group.
11. Click **OK**.

Edit a MAC Address Group

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then under **Networking & Security Inventory** click **NSX Managers**.
3. Click an NSX Manager in the **Name** column and then click the **Manage** tab.
   - You must select the primary NSX Manager if you need to manage universal MAC address groups.
4. Click the **Grouping Objects** tab and then click **MAC Sets**.
5. Select the group that you want to edit and click the **Edit** (✏) icon.

6. In the Edit MAC Set dialog box, make the appropriate changes.
7. Click **OK**.

Delete a MAC Address Group

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then under **Networking & Security Inventory** click **NSX Managers**.
3. Click an NSX Manager in the **Name** column and then click the **Manage** tab.
   - You must select the primary NSX Manager if you need to manage universal MAC address groups.
4. Click the **Grouping Objects** tab and then click **MAC Sets**.
5. Select the group that you want to delete and click the **Delete** (✖) icon.

Working with IP Pools

You can edit or delete an IP pool.

For information on adding an IP pool, see Configure Network Access SSL VPN-Plus or Configure Web Access SSL VPN-Plus.

Create an IP Pool

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then under **Networking & Security Inventory** click **NSX Managers**.
3. Click an NSX Manager in the **Name** column and then click the **Manage** tab.
4. Click the **Grouping Objects** tab and then click **IP Pool**.
5. Click the **Add New IP Pool** icon.
6. Type a name for the IP pool and type the default gateway and prefix length.
7. (Optional) Type the primary and secondary DNS and the DNS suffix.
8. Type the IP address ranges to be included in the pool and click **OK**.

Edit an IP Pool

You can edit an IP pool - you cannot edit the CIDR and gateway.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then under **Networking & Security Inventory** click **NSX Managers**.
3. Click an NSX Manager in the **Name** column and then click the **Manage** tab.
4. Click the **Grouping Objects** tab and then click **IP Pools**.
5. Select an IP pool and click the **Edit** icon.
6. Make the appropriate changes and click **OK**.

Delete IP Pool

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then under **Networking & Security Inventory** click **NSX Managers**.
3. Click an NSX Manager in the **Name** column and then click the **Manage** tab.
4. Click the **Grouping Objects** tab and then click **IP Pool**.
5. Select the IP pool that you want to delete and click the **Delete** icon.

Working with Security Groups

A security group is a collection of assets or grouping objects from your vSphere inventory.

Create a Security Group

You create a security group at the NSX Manager level.

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. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then under **Networking & Security Inventory** click **NSX Managers**.
3. Click an NSX Manager in the **Name** column and then click the **Manage** tab.
   - You must select the primary NSX Manager if you need to manage universal security groups.

4 Click the **Grouping Objects** tab, click **Security Group**, then click the **Add Security Group** icon.

5 Type a name and optionally a description for the security group.

6 (Optional) If you need to create a universal security group, select **Mark this object for universal synchronization**.

7 Click **Next**.

8 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.

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.

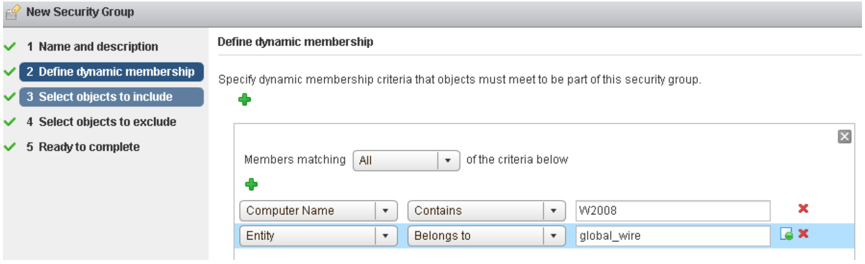

New Security Group

✓ 1 Name and description
✓ 2 Define dynamic membership
✓ 3 Select objects to include
✓ 4 Select objects to exclude
✓ 5 Ready to complete

Define dynamic membership

Specify dynamic membership criteria that objects must meet to be part of this security group.

Members matching All of the criteria below

Computer Name	Contains	W2008
Entity	Belongs to	global_wire

9 Click **Next**.

10 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 22-1. Objects that can be included in security groups and universal security groups.

Security Group	Universal Security Group
■ Other security groups to nest within the security group you are creating. ■ Cluster ■ Logical Switch ■ Network ■ Virtual App ■ Datacenter ■ IP sets ■ Directory Groups **Note** 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 Chapter 11 Identity Firewall Overview. ■ MAC Sets ■ Security tag ■ vNIC ■ Virtual Machine ■ Resource Pool ■ Distributed Virtual Port Group	■ Other universal security groups to nest within the universal security group you are creating. ■ Universal IP sets ■ Universal MAC sets

The objects selected here are always included in the security group regardless of whether or not they match the criteria in Step 8.

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.

11 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.

12 Click **Finish**.

Membership of a security group is determined as follows:

{Expression result (derived from Step 8) + Inclusions (specified in Step 10} - Exclusion (specified in Step 11)

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. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then under **Networking & Security Inventory** click **NSX Managers**.
3. Click an NSX Manager in the **Name** column and then click the **Manage** tab.
   - You must select the primary NSX Manager if you need to manage universal security groups.
4. Click the **Grouping Objects** tab and then click **Security Group**.
5. Select the group that you want to edit and click the **Edit** (✏) icon.
6. In the Edit Security Group dialog box, make the appropriate changes.
7. Click **OK**.

Delete a Security Group

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then under **Networking & Security Inventory** click **NSX Managers**.
3. Click an NSX Manager in the **Name** column and then click the **Manage** tab.
   - You must select the primary NSX Manager if you need to manage universal security groups.
4. Click the **Grouping Objects** tab and then click **Security Group**.
5. Select the group that you want to delete and click the **Delete** (✖) 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 create a service and then define rules for that service.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then under **Networking & Security Inventory** click **NSX Managers**.
3. Click an NSX Manager in the **Name** column and then click the **Manage** tab.
   - You must select the primary NSX Manager if you need to manage universal services.
4. Click the **Grouping Objects** tab and then click **Service**.
5. Click the **Add** (✚) icon.
6. Type a **Name** to identify the service.
7. (Optional) Type a **Description** for the service.
8. Select a **Protocol**.
   a. Depending on the protocol selected, you may be prompted to enter further information, such as destination port.
9. (Optional) Select **Enable inheritance to allow visibility at underlying scopes.**
10. (Optional) Select **Mark this object for Universal Synchronization** to create a universal service.
11. Click **OK**.

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. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then under **Networking & Security Inventory** click **NSX Managers**.
3. Click an NSX Manager in the **Name** column and then click the **Manage** tab.
   - You must select the primary NSX Manager if you need to manage universal service groups.
4. Click the **Grouping Objects** tab and then click **Service Groups**.
5. Click the **Add** icon.
6. Type a **Name** to identify the service group.
7. (Optional) Type a **Description** for the service group.
8. (Optional) Select **Mark this object for Universal Synchronization** to create a universal service group.

9. In Members, select the services or service groups that you want to add to the group.
10. (Optional) Select **Enable inheritance to allow visibility at underlying scopes.**
11. Click **OK**.

Edit a Service or Service Group

You can edit services and service groups.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then under **Networking & Security Inventory** click **NSX Managers**.
3. Click an NSX Manager in the **Name** column and then click the **Manage** tab.
   - You must select the primary NSX Manager if you need to manage universal services or service groups.
4. Click the **Grouping Objects** tab and then click **Service** or **Service Groups**.
5. Select a custom service or service group and click the **Edit** (✏) icon.
6. Make the appropriate changes.
7. Click **OK**.

Delete a Service or Service Group

You can delete services or service group.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then under **Networking & Security Inventory** click **NSX Managers**.
3. Click an NSX Manager in the **Name** column and then click the **Manage** tab.
   - You must select the primary NSX Manager if you need to manage universal services or service groups.
4. Click the **Grouping Objects** tab and then click **Service** or **Service Groups**.
5. Select a custom service or service group and click the **Delete** (✖) icon.
6. Click **Yes**.

   The service or service group is deleted.

Operations and Management 23

This chapter includes the following topics:

Change Controller Password

To ensure data security, you can change passwords for NSX controllers.

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **Installation**.

3 Under Management, select the controller for which you want to change the password.

4 Click **Actions** and then click **Change Controller Cluster Password**.

5 Type a new password and click **OK**.

The controller password has changed.

Recover from an NSX Controller Failure

In case of an NSX Controller failure, you may still have two controllers that are working. The cluster majority is maintained, and the control plane continues to function. Even so, it is important to delete all three controllers and add new ones, so as to maintain a fully functional three-node cluster.

We recommend deleting the controller cluster when one or more of the controllers encounter catastrophic, unrecoverable errors or when one or more of the controller VMs become inaccessible and cannot be fixed.

We recommend deleting all controllers in such a case, even if some of the controllers seem healthy. The recommended process is to create a new controller cluster and use the Update Controller State mechanism on the NSX Manager to synchronize the state to the controllers.

Procedure

1 Login to vSphere Web Client.

2 From **Networking & Security**, click **Installation > Management**.

3 In the NSX Controller nodes section, click each controller and take screen shots/print-screens of the details screens or write down the configuration information for later reference.

For example:

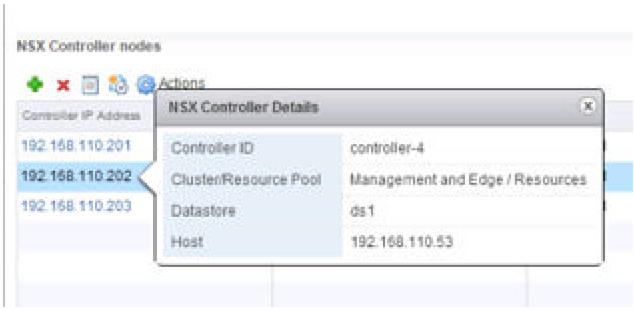

Controller IP Address
192.168.110.201
192.168.110.202
192.168.110.203

Controller ID	controller-4
Cluster/Resource Pool	Management and Edge / Resources
Datastore	ds1
Host	192.168.110.53

4 In the NSX Controller nodes section, delete all three of them by selecting each one and clicking the **Delete Node (x)** icon.

When there are no controllers in the system, the hosts are operating in what is called "headless" mode. New VMs or vMotioned VMs will have networking issues until new controllers are deployed and the synchronization is completed.

5 Deploy three new NSX Controller nodes by clicking the **Add Node (+)** icon.

6 In the Add Controller dialog box, select the datacenter on which you are adding the nodes, and configure the controller settings.

a Select the appropriate cluster.

b Select a Host in the cluster and storage.

c Select the distributed port-group.

d Select the IP pool from which IP addresses are to be assigned to the node.

e Click **OK**, wait for installation to complete, and ensure all nodes have a status of Normal.

7 Resynchronize the controller state by clicking **Actions > Update Controller State**.

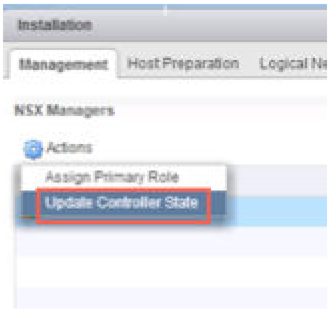

Update Controller State pushes the current VXLAN and Distributed Logical Router configuration (including Universal Objects in a Cross-VC NSX deployment) from NSX Manager to the Controller Cluster.

Change VXLAN Port

You can change the port used for VXLAN traffic.

Starting in NSX 6.2.3, 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 to NSX 6.2.3, 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. In a cross-vCenter NSX environment the change will propagate to all NSX Manager appliances and all hosts in the cross-vCenter NSX environment.

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 Click the **Logical Network Preparation** tab, then click **VXLAN Transport**.

2 Click the **Change** button in the VXLAN Port panel. Enter the port you want to switch to. 4789 is the port assigned by IANA for VXLAN.

It will take 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>
```

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.

Procedure

1 In the vSphere Web Client, navigate to **Networking & Security > Installation > Host Preparation**.

2 Select a cluster or expand clusters and select a host. Click **Actions** (⚙) then **Communication Channel Health**.

The communication channel health information is displayed.

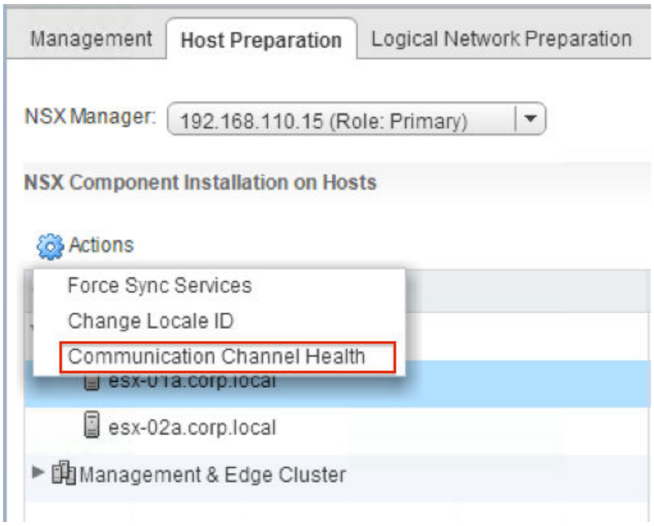
Management
Host Preparation
Logical Network Preparation
NSX Manager:
192.168.110.15 (Role: Primary)
NSX Component Installation on Hosts
Actions
Force Sync Services
Change Locale ID
Communication Channel Health
esx-01a.corp.local
esx-02a.corp.local
Management & Edge Cluster

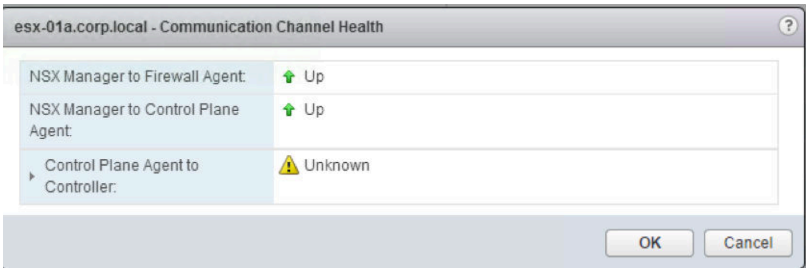
esx-01a.corp.local - Communication Channel Health

NSX Manager to Firewall Agent:	Up
NSX Manager to Control Plane Agent:	Up
Control Plane Agent to Controller:	Unknown

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 http://www.vmware.com/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 Log in to the vSphere Web Client.

2 Select **Networking & Security**.

3 Under Networking & Security Inventory, select **NSX Managers**.

4 Double-click the NSX Manager you want to modify.

5 Click the **Summary** tab.

6 Click **Edit** in the Customer Experience Improvement Program dialog box.

7 Select or deselect the **Join the VMware Customer Experience Improvement Program** option.

8 (Optional) Configure the recurrence settings.

9 Click **OK**.

System Events and Audit Logs

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).

With the NSX ticket logger feature, you can track the changes you make with a ticket ID. Audit logs for operations tracked by a ticket will include the ticket ID.

About NSX Logs

This section describes how 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 syslog for hosts managed by a vCenter Server, see VMware vSphere ESXi and vCenter Server 5.5 Documentation.

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.

NSX Manager

To specify a syslog server, see Specify a Syslog Server.

To download technical support logs, see Download Technical Support Logs for NSX.

NSX Edge

To specify a syslog server, see Configure Remote Syslog Servers.

To download technical support logs, see Download Tech Support Logs for NSX Edge.

Firewall

You must configure the remote syslog server for each cluster that has firewall enabled. The remote syslog server is specified in the `Syslog.global.logHost` attribute. See *ESXi and vCenter Server 5.5 Documentation*.

Here is a sample line from a host log file.

```
2013-10-02T05:41:12.670Z cpu11:1000046503)vsip_pkt: INET, match, PASS, Rule 0/3, Ruleset domain-c7,
Rule ID 100, OUT, Len 60, SRC 10.24.106.96, DST 10.24.106.52, TCP SPORT 59692, DPORT 22 S
```

which consists of three parts:

Table 23-1. Components of log file entry

	Value in example
VMKernel common log portion consists of date, time, CPU, and WorldID	2013-10-02T05:41:12.670Z cpu11:1000046503)
Identifier	vsip_pkt
Firewall specific portion	INET, match, PASS, Rule 0/3, Ruleset domain-c7, Rule ID 100, OUT, Len 60, SRC 10.24.106.96, DST 10.24.106.52, TCP SPORT 59692, DPORT 22 S

Table 23-2. Firewall specific portion of log file entry

Entity	Possible Values
AF Value	INET, INET6
Reason	Possible values: match, bad-offset, fragment, short, normalize, memory, bad-timestamp, congestion, ip-option, proto-cksum, state-mismatch, state-insert, state-limit, src-limit, synproxy, spoofguard
Action	PASS, DROP, SCRUB, NOSCRUB, NAT, NONAT, BINAT, NOBINAT, RDR, NORDR, SYNPROXY_DROP, PUNT, REDIRECT, COPY
Rule identifier	*Identifier*
Rule value	Ruleset ID and Rule position (Internal details)
Rule set identifier	*Identifier*
Rule set value	Ruleset name
Rule ID identifier	*Identifier*
Rule ID	ID matched
Direction	ROUT, IN
Length identifier	Len followed by variable
Length value	Packet length
Source identifier	SRC
Source IP address	*IP address*
Destination identifier	*IP address*

Table 23-2. Firewall specific portion of log file entry (Continued)

Entity	Possible Values
Protocol	TCP, UDP, PROTO
Source port identifier	SPORT
Source port	Source port number for TDP and UDP
Source port identifier	Destination port identifier
Destination port	Destination port number for TDP and UDP
Flag	Flag for TCP

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 ticked ID.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click the **Manage** tab.
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

   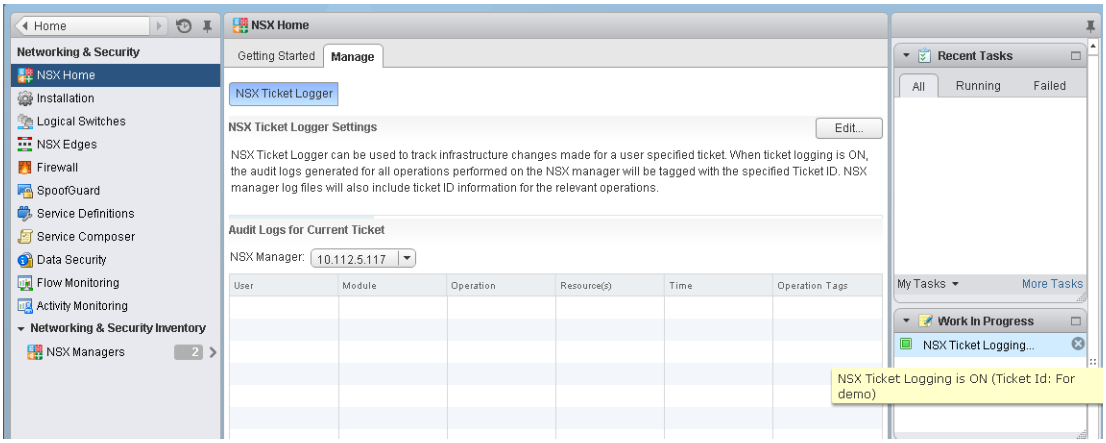

   NSX Ticket Logger can be used to track infrastructure changes made for a user specified ticket. When ticket logging is ON, the audit logs generated for all operations performed on the NSX manager will be tagged with the specified Ticket ID. NSX manager log files will also include ticket ID information for the relevant operations.

   NSX Manager: 10.112.5.117

User	Module	Operation	Resource(s)	Time	Operation Tags

   NSX Ticket Logging is ON (Ticket Id: For demo)

   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 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 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then under **Networking & Security Inventory** click **NSX Managers**.

3 Click an NSX Manager in the **Name** column and then click the **Monitor** tab.

4 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.

NSX Manager Virtual Appliance Events

The following events are specific to the NSX Manager virtual appliance.

Table 23-3. NSX Manager Virtual Appliance Events

	Power Off	Power On	Interface Down	Interface Up
Local CLI	Run `show log follow` command.	Run `show log follow` command.	Run `show log follow` command.	Run `show log follow` command.
GUI	NA	NA	NA	NA

Table 23-4. NSX Manager Virtual Appliance Events

	CPU	Memory	Storage
Local CLI	Run `show process monitor` command.	Run `show system memory` command.	Run `show filesystem` command.
GUI	NA	NA	NA

About the Syslog Format

The system event message logged in the syslog has the following structure.

```
syslog header (timestamp + hostname + sysmgr/)
Timestamp (from the service)
Name/value pairs
Name and value separated by delimiter '::' (double colons)
Each name/value pair separated by delimiter ';;' (double semi-colons)
```

The fields and types of the system event contain the following information.

```
Event ID :: 32 bit unsigned integer
Timestamp :: 32 bit unsigned integer
Application Name :: string
Application Submodule :: string
Application Profile :: string
Event Code :: integer (possible values: 10007 10016 10043 20019)
Severity :: string (possible values: INFORMATION LOW MEDIUM HIGH CRITICAL)
Message ::
```

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 1,000, 000 audit logs.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then under **Networking & Security Inventory** click **NSX Managers**.
3. In the **Name** column, click an NSX server and then click the **Monitor** tab.
4. Click the **Audit Logs** tab.
5. 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.
6. In the **Audit Log Change Details**, select **Changed Rows** to display only those properties whose values have changed for this audit log operation.

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**.

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.

Specify a Syslog Server

If you specify a syslog server, NSX Manager sends all audit logs and system events to the syslog server.

Syslog data is useful for troubleshooting and reviewing data logged during installation and configuration.

NSX Edge supports two syslog servers. NSX Manager and NSX Controllers support one syslog server.

Procedure

1 In a Web browser, navigate to the NSX Manager appliance GUI at https://<nsx-manager-ip> or https://<nsx-manager-hostname>.

2 Log in as admin with the password that you configured during NSX Manager installation.

3 Click **Manage Appliance Settings**.

For example:

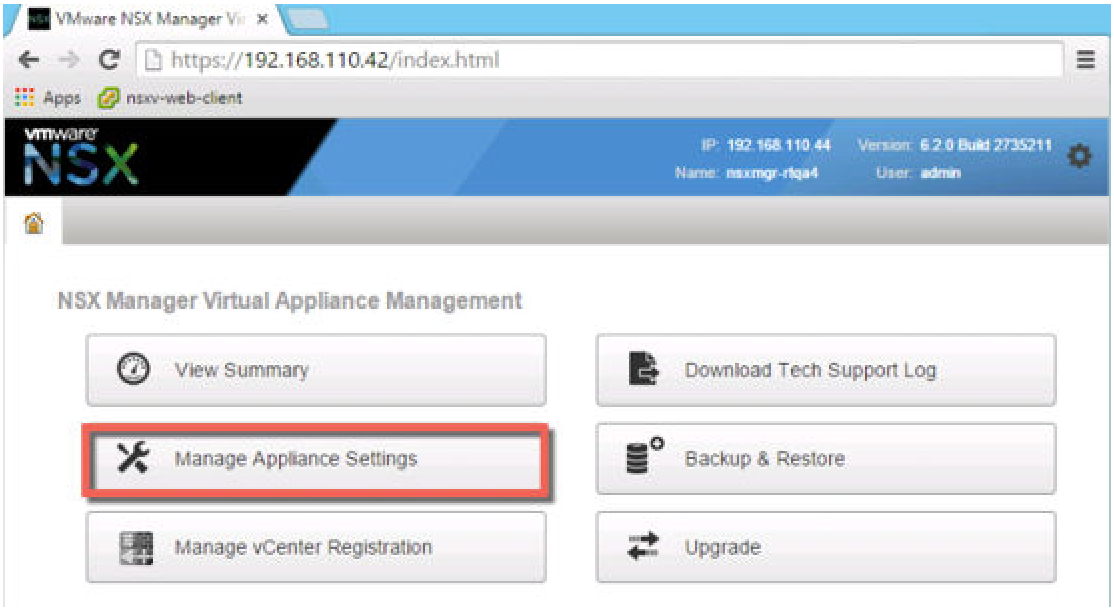
VMware NSX Manager Vi
https://192.168.110.42/index.html
Apps
nsxv-web-client
vmware NSX
IP: 192.168.110.44
Name: nsxmgr-rtqa4
Version: 6.2.0 Build 2735211
User: admin
NSX Manager Virtual Appliance Management
View Summary
Download Tech Support Log
Manage Appliance Settings
Backup & Restore
Manage vCenter Registration
Upgrade

4 From the Settings panel, click **General**.

5 Click **Edit** next to **Syslog Server**.

6 Type the IP address or hostname, port, and protocol of the syslog server.

If you do not specify a port, the default UDP port for the IP address/host name of the syslog server is used.

For example:

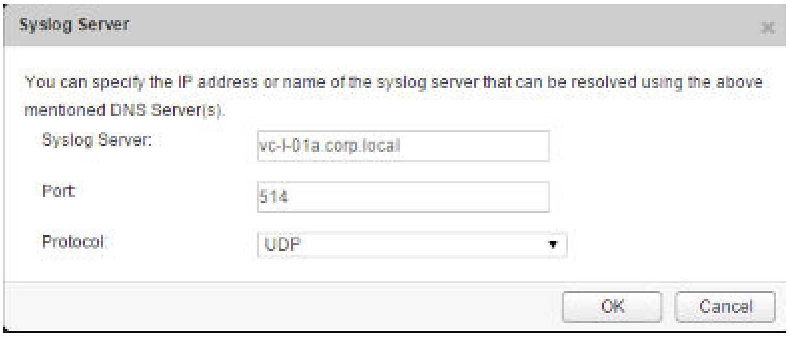
Syslog Server

You can specify the IP address or name of the syslog server that can be resolved using the above mentioned DNS Server(s).

Syslog Server:	vc-l-01a.corp.local
Port:	514
Protocol:	UDP

7 Click **OK**.

vCenter Server remote logging is enabled, and logs are stored in your standalone syslog server.

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.

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 you 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**.

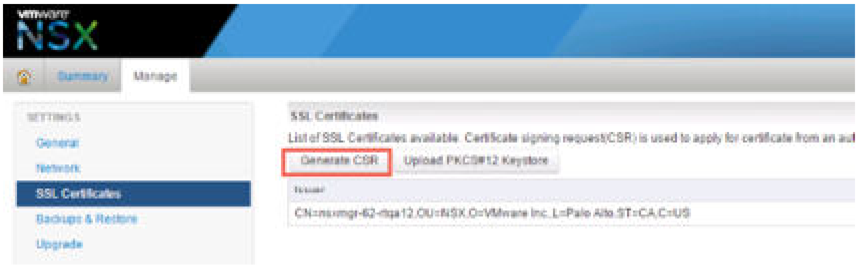

List of SSL Certificates available. Certificate signing request(CSR) is used to apply for certificate from an aut

CN=nsxmgr-62-rtqa12,OU=NSX,O=VMware Inc.,L=Palo Alto,ST=CA,C=US

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 **US**.

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 obtain the NSX Manager certificate, make sure the certificate and private key are in the PKCS#12 format. If your NSX Manager certificate and private key are not in PKCS#12 format, you must convert them before you can import them into NSX Manager.

Prerequisites

Verify that OpenSSL is installed on the system. You can download openssl from http://www.openssl.org.

Procedure

◆ After receiving the signed certificate from the authorized signer, use OpenSSL to generate a PKCS#12 (.pfx or .p12) keystore file from the certificate file and your private key.

For example:

```
openssl pkcs12 –export –out server.p12 –inkey server.key –in server.crt –certfile CACert.crt
```

In this example, CACert.crt is the name of the root certificate that was returned by the certificate authority.

What to do next

Import the signed SSL certificate 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**.

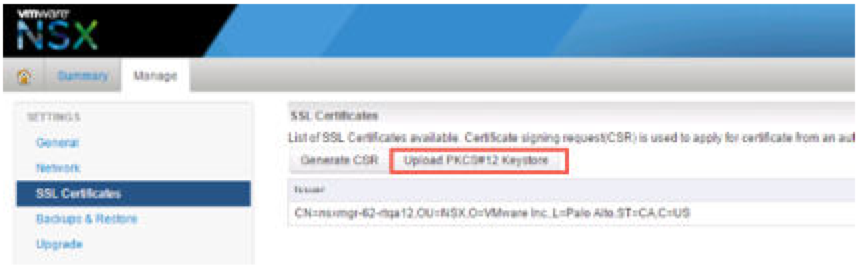

List of SSL Certificates available. Certificate signing request(CSR) is used to apply for certificate from an aut

Issuer

CN=nsxmgr-62-rtqa12,OU=NSX,O=VMware Inc.,L=Palo Alto,ST=CA,C=US

5 Click **Choose File** to locate the file.

6 Click **Import**.

7 To apply the certificate, reboot the NSX Manager appliance.

The certificate is stored in NSX Manager.

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 vmwNsxMSnmpDisabled 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 vmwNsxMSnmpManagerConfigUpdated 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 SNM 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 Log in to the vSphere Web Client.

2 Select **Networking & Security > Networking & Security Inventory > NSX Managers**.

3 Select an NSX Manager IP address.

4 Select the **Manage > System Events** tabs.

5 Click **Edit** to configure the SNMP settings.

Option	Description
Service	Sends out SNMP trap. By default, this option is disabled.
Group Notification	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.
Receivers	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.

6 Click **OK**.

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 Select the **Manage > System 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.

Option	Description
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. Log in to the vSphere Web Client.
2. Select **Networking & Security > Networking & Security Inventory > NSX Managers**.
3. Select an NSX Manager IP address.
4. Select the **Manage > System Events** tabs.
5. Select a system event under the System Traps section.
6. 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.
7. Change the severity of the system event from the drop-down menu.
8. If you change the severity from Informational to critical, check the **Enable as SNMP Trap** checkbox.
9. Click **OK**.
10. (Optional) Click the **Enable** (✓) icon or **Disable** (⊘) icon in the header to enable or disable sending a system trap.
11. (Optional) Click the **Copy** (📋) icon to copy one or more event rows to your clipboard.

NSX Backup and Restore

Proper backup of all NSX 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 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.

We recommend taking backups in the following scenarios:

- Before an NSX or vCenter upgrade.
- After an NSX or vCenter upgrade.
- After Day Zero deployment and initial configuration of NSX components, such as after the creation of NSX Controllers, 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 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 NSX Manager Data

You can back up NSX Manager data by performing an on-demand backup or a scheduled backup.

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 new backup file before and after performing an NSX upgrade, one backup for the old version and another backup for the new version.

Procedure

1. Log in to the NSX Manager Virtual Appliance.
2. Under Appliance Management, click **Backups & Restore**.
3. To specify the backup location, click **Change** next to FTP Server Settings.
   a. Type 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 required.
   d. Type the user name and password required to login to the backup system.

e In the **Backup Directory** field, type the absolute path where backups will be stored.

To determine the absolute path, you can 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 Type a text string in **Filename Prefix**.

This text is prepended to each backup filename for easy recognition on the backup system. For example, if you type **ppdb**, the resulting backup is named as ppdbHH_MM_SS_DayDDMonYYYY.

g Type the pass phrase to secure the backup.

You will need this pass phrase to restore the backup.

h Click **OK**.

For example:

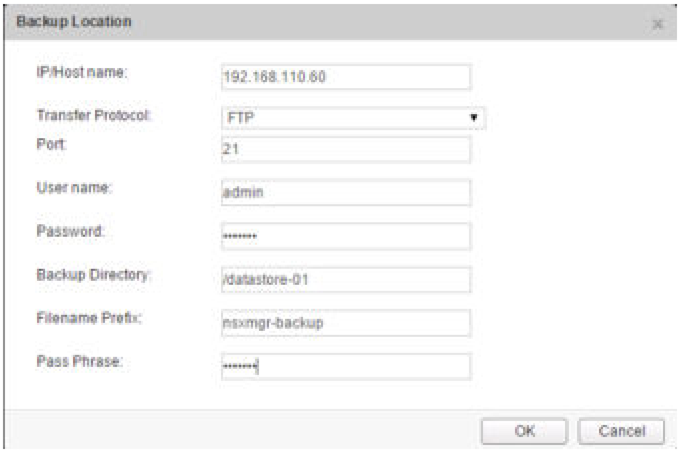

Backup Location

IP/Host name:	192.168.110.60
Transfer Protocol:	FTP
Port:	21
User name:	admin
Password:	•••••••
Backup Directory:	/datastore-01
Filename Prefix:	nsxmgr-backup
Pass Phrase:	•••••••

4 For an on-demand backup, click **Backup**.

A new file is added under **Backup History**.

5 For scheduled backups, click **Change** next to Scheduling.

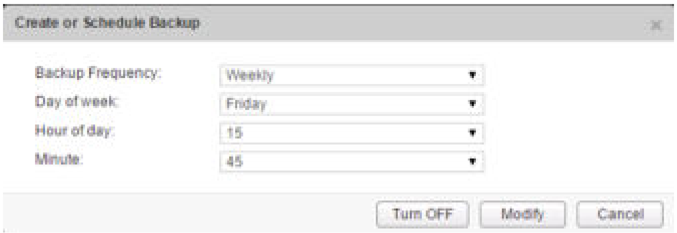

Create or Schedule Backup

Backup Frequency: Weekly

Day of week: Friday

Hour of day: 15

Minute: 45

a 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 as this field is not applicable to a daily frequency.

b For a weekly backup, select the day of the week the data should be backed up.

c For a weekly or daily backup, select the hour at which the backup should begin.

d Select the minute to begin and click **Schedule**.

6 To exclude logs and flow data from being backed up, click **Change** next to Exclude.

a Select the items you want to exclude from the backup.

b Click **OK**.

7 Save your FTP server IP/hostname, credentials, directory details, and pass phrase. This information is needed 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 officially 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 screen shots of the current settings for the old NSX Manager appliance or note them so that they can be used to specify IP information and backup location information for the newly deployed NSX Manager appliance.

Procedure

1 Take screen shots or note all settings on the existing NSX Manager appliance.

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 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.

6. In the Backups History section, select the check box for the backup to restore and click **Restore**.

Back Up NSX Edges

All NSX Edge configurations (logical routers and edge services gateways) are backed up as part of NSX Manager data backup.

If you have an intact NSX Manager configuration, you can recreate an inaccessible or failed Edge appliance VM by redeploying the NSX Edge (click the **Redeploy NSX Edge** icon in the vSphere Web Client).

Taking individual NSX Edge backups is not supported.

Back Up vSphere Distributed Switches

You can export vSphere distributed switch and distributed port group configurations to a file.

The file preserves valid network configurations, enabling distribution of these configurations to other deployments.

This functionality is available only with the vSphere Web Client 5.1 or later. VDS settings and port-group settings are imported as part of the import.

As a best practice, export the VDS configuration before preparing the cluster for VXLAN. For detailed instructions, 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 VM snapshots, see http://kb.vmware.com/kb/1015180.

Useful links for vCenter 5.5:

- http://kb.vmware.com/kb/2057353
- http://kb.vmware.com/kb/2034505

- http://www.vmware.com/files/pdf/techpaper/vmware-vcenter-server-availability-guide.pdf

Useful links for vCenter 6.0:

- https://pubs.vmware.com/vsphere-60/topic/com.vmware.vsphere.install.doc/GUID-539B47B4-114B-49BC-9736-F14058127ECA.html
- http://kb.vmware.com/kb/2110294

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. You can live monitor TCP and UDP connections to and from a selected vNIC. You can also exclude flows by specifying filters.

Flow monitoring can thus be used as a forensic tool to detect rogue services and examine outbound sessions.

View Flow Monitoring Data

You can view traffic sessions on virtual machines within the specified time span. The last 24 hours of data are displayed by default, the minimum time span is one hour and the maximum is two weeks.

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 Installation Guide*.

Procedure

1. Log in to the vSphere Web Client.
2. Select **Networking & Security** from the left navigation pane and then select **Flow Monitoring**.
3. Ensure that you are in the **Dashboard** tab.

4 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.

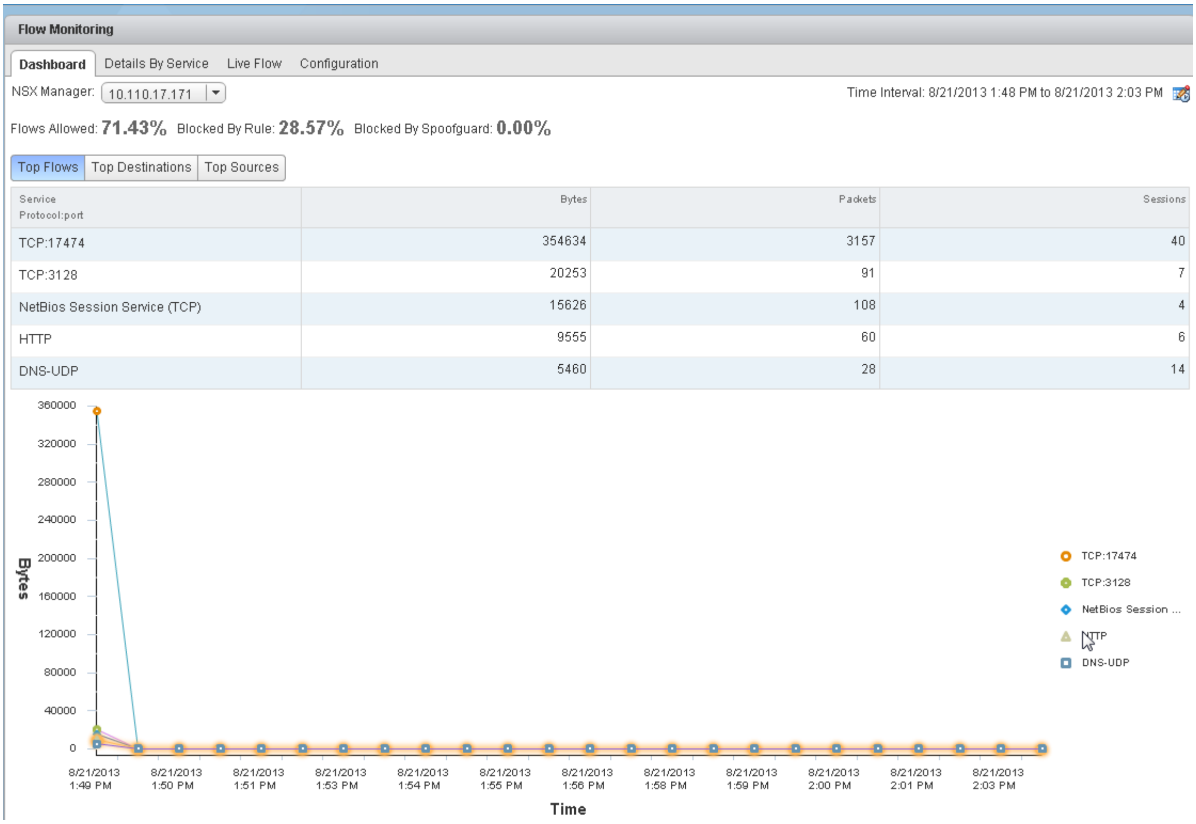

NSX Manager: 10.110.17.171

Time Interval: 8/21/2013 1:48 PM to 8/21/2013 2:03 PM

Flows Allowed: 71.43% Blocked By Rule: 28.57% Blocked By Spoofguard: 0.00%

Service Protocol:port	Bytes	Packets	Sessions
TCP:17474	354634	3157	40
TCP:3128	20253	91	7
NetBios Session Service (TCP)	15626	108	4
HTTP	9555	60	6
DNS-UDP	5460	28	14

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.

5 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.

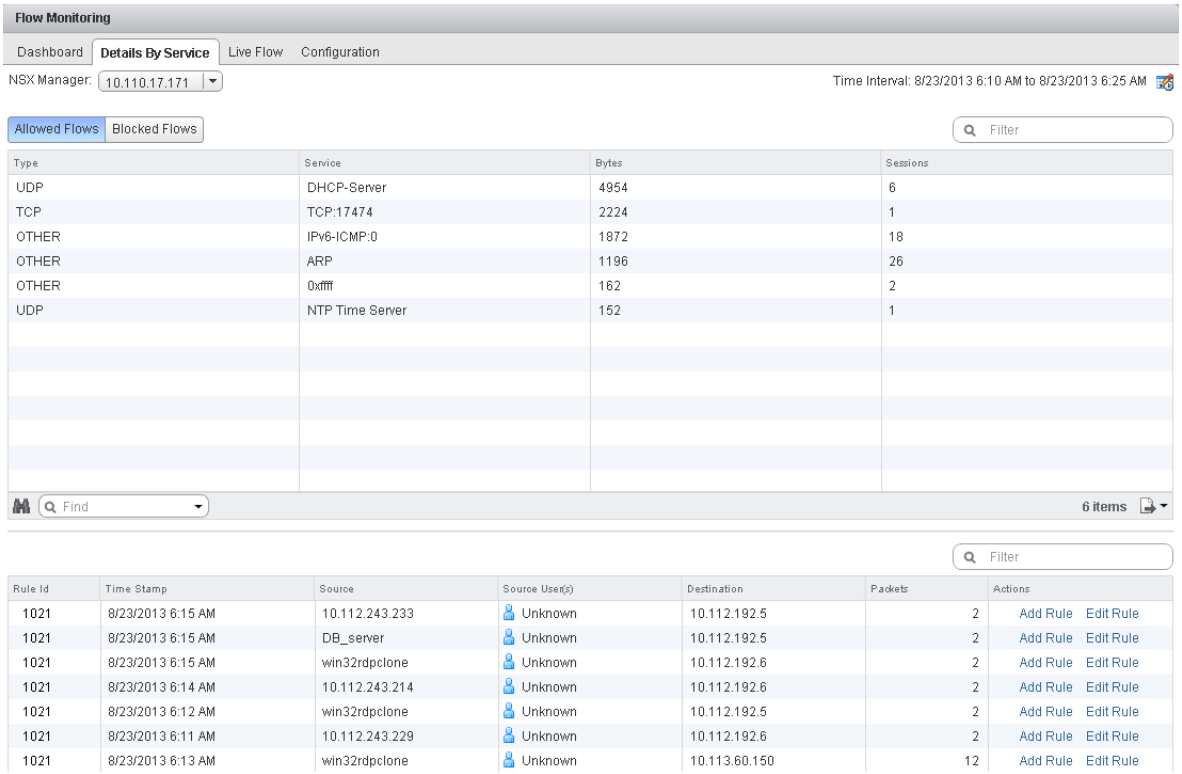

NSX Manager: 10.110.17.171

Time Interval: 8/23/2013 6:10 AM to 8/23/2013 6:25 AM

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

6 items

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

6 Click an item in the table to display the rules that allowed or blocked that traffic flow.

7 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 Log in to the vSphere Web Client.

2 Select **Networking & Security** from the left navigation pane and then select **Flow Monitoring**.

3 Click next to **Time interval**.

4 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.

5 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. Log in to the vSphere Web Client.
2. Select **Networking & Security** from the left navigation pane and then select **Flow Monitoring**.
3. Click the **Details by Service** tab.
4. 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.
5. Click a rule ID to view rule details.
6. 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. Log in to the vSphere Web Client.
2. Select **Networking & Security** from the left navigation pane and then select **Flow Monitoring**.
3. Click the **Live Flow** tab.

4 Click **Browse** and select a vNIC.

5 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.

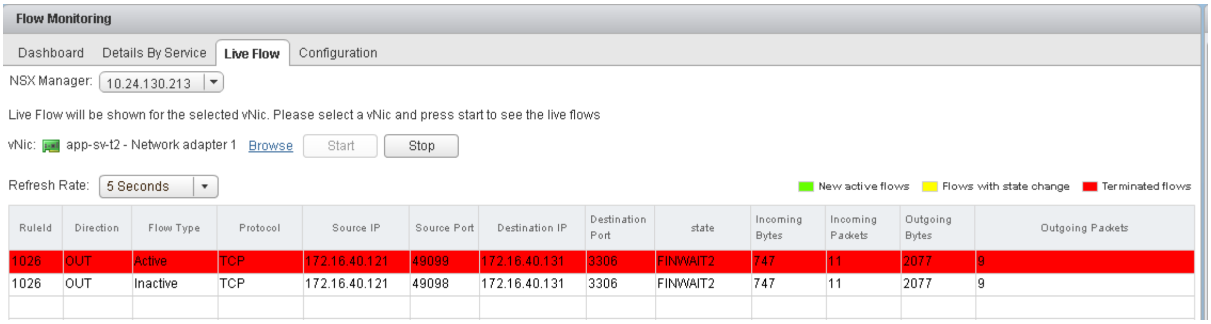

NSX Manager: 10.24.130.213

Live Flow will be shown for the selected vNic. Please select a vNic and press start to see the live flows

vNic: app-sv-t2 - Network adapter 1

Refresh Rate: 5 Seconds

RuleId	Direction	Flow Type	Protocol	Source IP	Source Port	Destination IP	Destination Port	state	Incoming Bytes	Incoming Packets	Outgoing Bytes	Outgoing Packets
1026	OUT	Active	TCP	172.16.40.121	49099	172.16.40.131	3306	FINWAIT2	747	11	2077	9
1026	OUT	Inactive	TCP	172.16.40.121	49098	172.16.40.131	3306	FINWAIT2	747	11	2077	9

6 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 Log in to the vSphere Web Client.

2 Select **Networking & Security** from the left navigation pane and then select **Flow Monitoring**.

3 Select the **Configuration** tab.

4 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**.

5 To specify filtering criteria, click **Flow Exclusion** and follow the steps below.

a Click the tab corresponding to the flows you want to exclude.

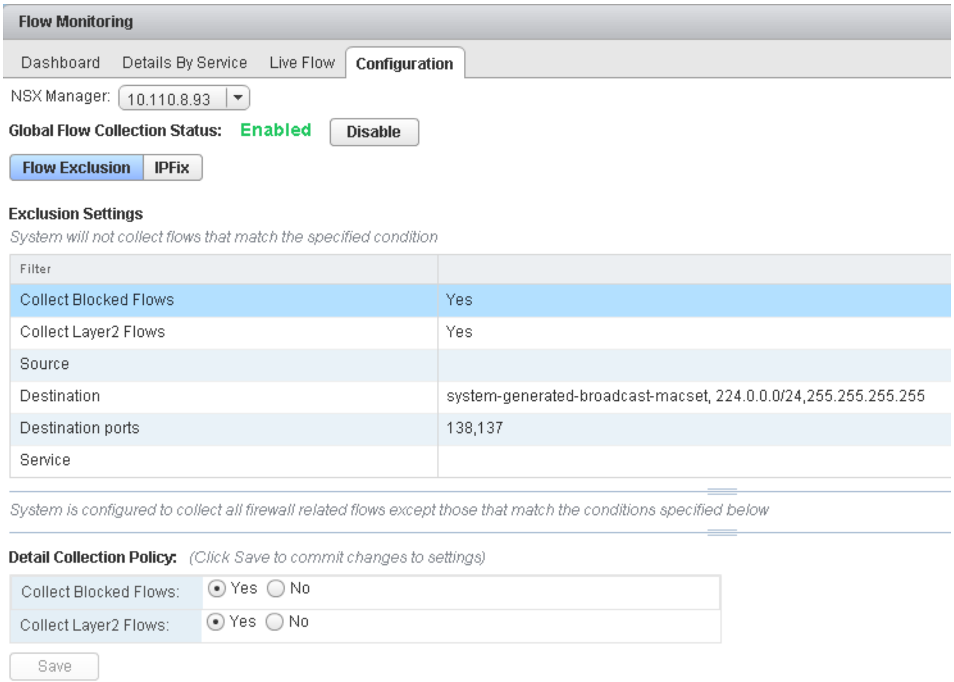

NSX Manager: 10.110.8.93

Global Flow Collection Status: Enabled

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

b Specify the required information.

If you selected	Specify the following information
Collect Blocked Flows	Select No to exclude blocked flows.
Collect Layer2 Flows	Select No to exclude Layer2 flows.
Source	Flows are not collected for the specified sources. 1 Click the **Add** icon. 2 In View, select the appropriate container. 3 Select the objects to exclude.
Destination	Flows are not collected for the specified destinations. 1 Click the **Add** icon. 2 In View, select the appropriate container. 3 Select the objects to exclude.
Destination ports	Excludes flows to the specified ports. Type the port numbers to exclude.
Service	Excludes flows for the specified services and service groups. 1 Click the **Add** icon. 2 Select the appropriate services and/or service groups.

c Click **Save**.

6 To configure flow collection, click **IPFix** and follow the steps below.

a Click **Edit** next to IPFix Configuration and click **Enable IPFix Configuration**.

b In **Observation DomainID**, type a 32-bit identifier that identifies the firewall exporter to the flow collector.

c 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 5. For example, if the flow is active for 30 minutes and the export timeout is 5 minutes, then the flow will be exported 7 times during its lifetime. Once each for creation and deletion, and 5 times during the active period.

d In **Collector IPs**, click the Add (✚) icon and type the IP address and UDP port of the flow collector.

e Click **OK**.

Activity Monitoring

Activity monitoring provides visibility into the applications that are in use on the Windows desktop virtual machines that are managed by vCenter. This visibility helps ensure that security policies at your organization are being enforced correctly.

A security policy may mandate who is allowed access to what applications. The cloud administrator can generate Activity Monitoring reports to see if the IP based firewall rule that they set is doing the intended work. By providing user and application level detail, activity monitoring translates high level security policies to low level IP address and network based implementation.

Figure 23-2. Your virtual environment today

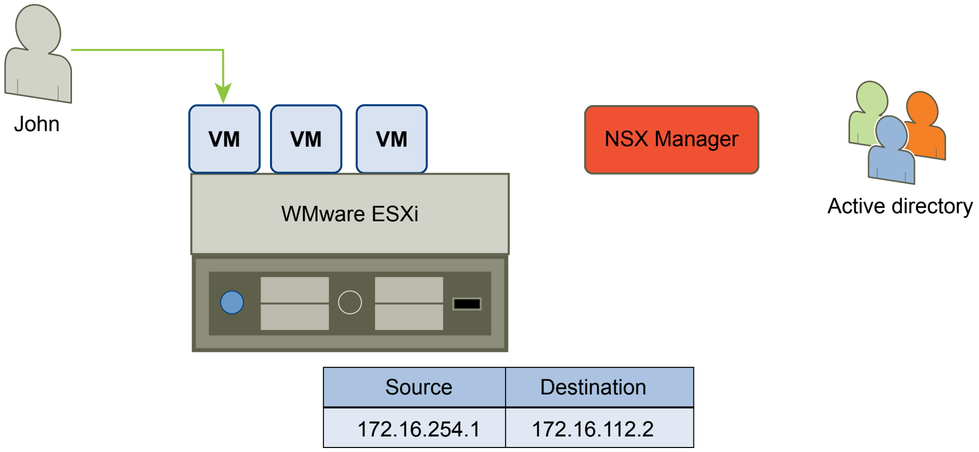

Source	Destination
172.16.254.1	172.16.112.2

Once you enable data collection for Activity Monitoring, you can run reports to view inbound traffic (such as virtual machines being accessed by users) as well as outbound traffic (resource utilization, interaction between inventory containers, and AD groups that accessed a server).

Figure 23-3. Your virtual environment with Activity Monitoring

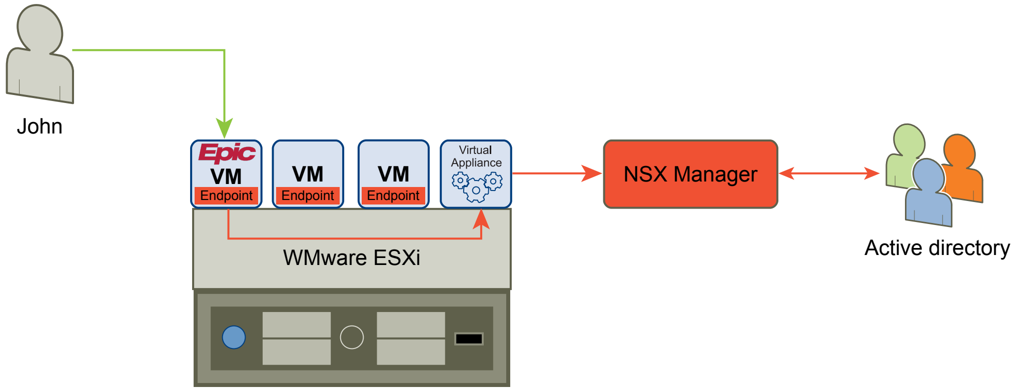

User	AD group	App name	Originating VM name	Destination VM name	Source IP	Destination IP
John	Doctors	Epic.exe	DoctorsWS13	EpicSVR3	172.16.254.1	172.16.112.2

Important Activity monitoring is not supported on Linux VMs.

Set Up Activity Monitoring

For activity monitoring to work, there are several required procedures that must be performed, including installation of the guest introspection driver, installation of guest introspection VMs, and enabling NSX activity monitoring. Optionally, you can also use service composer to control which VMs are monitored.

Prerequisites

- NSX must be installed and operational.
- NSX Manager must be linked with the AD server where it will get groups to which to match Windows VMs users.
- The vCenter inventory must contain one or more Windows desktop VMs.
- VMware Tools must be running and current on your Windows desktop VMs.

Procedure

1 On the Windows VMs in your vCenter inventory, install the Guest Introspection driver if it is not already installed.

   a Navigate to **Control Panel\Programs\Programs and Features**, right-click **VMware Tools** and select **Change**.

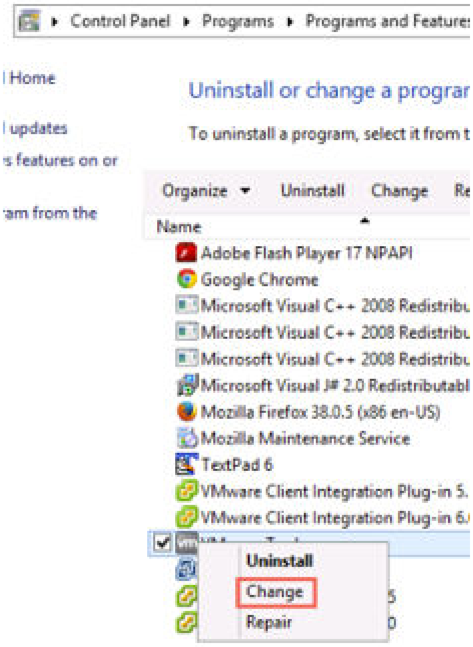
Control Panel ▸ Programs ▸ Programs and Feature
Home
updates
s features on or
am from the
Uninstall or change a prograr
To uninstall a program, select it from t
Organize ▾
Uninstall
Change
Re
Name
Adobe Flash Player 17 NPAPI
Google Chrome
Microsoft Visual C++ 2008 Redistribu
Microsoft Visual C++ 2008 Redistribu
Microsoft Visual C++ 2008 Redistribu
Microsoft Visual J# 2.0 Redistributabl
Mozilla Firefox 38.0.5 (x86 en-US)
Mozilla Maintenance Service
TextPad 6
VMware Client Integration Plug-in 5.
VMware Client Integration Plug-in 6.
Uninstall
Change
Repair

   b Select **Modify**.

   c Under **VMCI Driver**, click **Guest Introspection Drivers > Will be installed on local hard drive**.

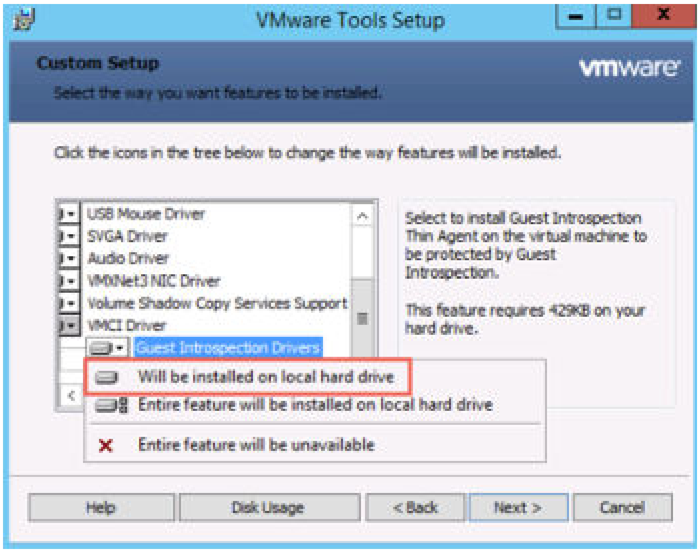
VMware Tools Setup
Custom Setup
Select the way you want features to be installed.
vmware
Click the icons in the tree below to change the way features will be installed.
USB Mouse Driver
SVGA Driver
Audio Driver
VMXNet3 NIC Driver
Volume Shadow Copy Services Support
VMCI Driver
Guest Introspection Drivers
Will be installed on local hard drive
Entire feature will be installed on local hard drive
Entire feature will be unavailable
Help
Disk Usage
< Back
Next >
Cancel

Select to install Guest Introspection Thin Agent on the virtual machine to be protected by Guest Introspection.

This feature requires 429KB on your hard drive.

The guest introspection driver detects what applications are running on each Windows VM and sends this information to the guest introspection VM.

2 Install the guest introspection VMs.

When first launching the VMware Tools install, choose the **Custom** option. In the VMCI folder, select **Guest Introspection Driver**. The driver is not selected by default.

To add the driver after VMware Tools is already installed:

a In the vCenter Web Client, navigate to **Networking & Security > Installation > Service Deployments**.

b Add a new service deployment.

c Select **Guest Introspection**.

d Select the host clusters that contain Windows VMs.

e Select the appropriate datastores, networks, and IP addressing mechanism. If you are not using DHCP for your guest introspection VMs, create and assign an IP pool.

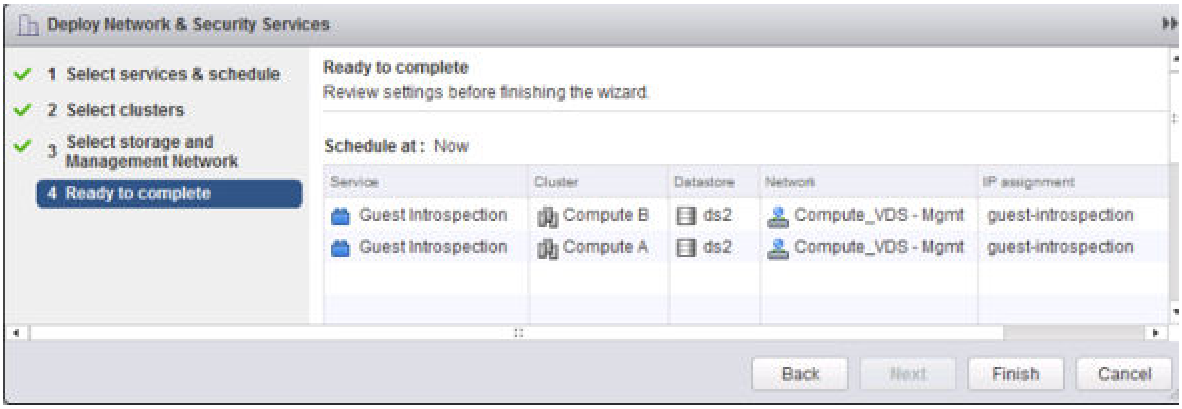
Deploy Network & Security Services
✓ 1 Select services & schedule
✓ 2 Select clusters
✓ 3 Select storage and Management Network
4 Ready to complete
Back
Next
Finish
Cancel

Ready to complete
Review settings before finishing the wizard.

Schedule at : Now

Service	Cluster	Datastore	Network	IP assignment
Guest Introspection	Compute B	ds2	Compute_VDS - Mgmt	guest-introspection
Guest Introspection	Compute A	ds2	Compute_VDS - Mgmt	guest-introspection

Two guest introspection VMs are installed, one on each host within each cluster.

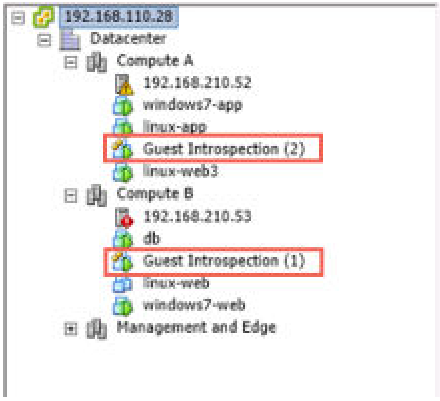

192.168.110.28
Datacenter
Compute A
192.168.210.52
windows7-app
linux-app
Guest Introspection (2)
linux-web3
Compute B
192.168.210.53
db
Guest Introspection (1)
linux-web
windows7-web
Management and Edge

3 Enable activity monitoring on the Windows VMs.

   a In the **Hosts and Clusters** view, select the Windows VM, and select the **Summary** tab.

   b In NSX Activity Monitoring, click **Edit** and click **Yes**.

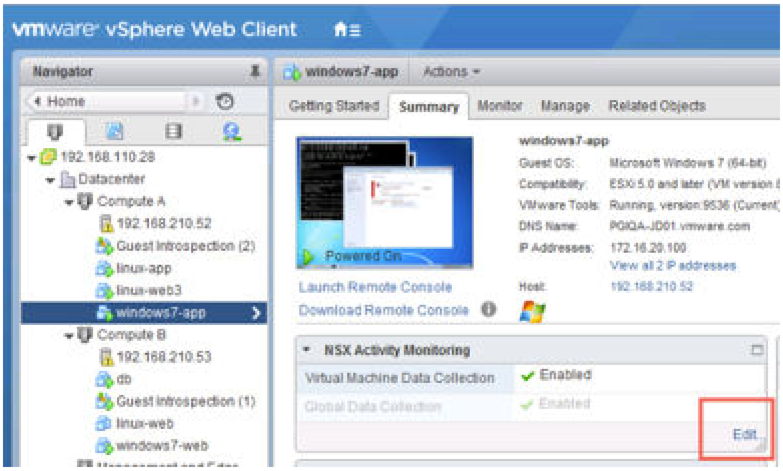

windows7-app

Guest OS: Microsoft Windows 7 (64-bit)

Compatibility: ESXi 5.0 and later (VM version

VMware Tools: Running, version:9536 (Current)

DNS Name: PGIQA-JD01.vmware.com

IP Addresses: 172.16.20.100

View all 2 IP addresses

Host: 192.168.210.52

NSX Activity Monitoring

Virtual Machine Data Collection	✓ Enabled
Global Data Collection	✓ Enabled

   Repeat this step for all Windows VMs that you want to monitor.

4 (Optional) Modify the list of vCenter objects that are monitored, or define a dynamic membership rule.

   a In the vCenter Web Client, navigate to **Networking & Security > Service Composer**.

   b Edit the **Activity Monitoring Data Collection** security group.

c Define a dynamic membership rule so that as new Windows VMs are added to the cluster, the VM will automatically be monitored.

d Select vCenter objects to include or exclude in the activity monitoring security group.

The VMs on which you enabled activity monitoring are automatically included in the activity monitoring security group.

In this example, all VMs with names starting with "win" are automatically added to the activity monitoring security group. This means that activity monitoring will be automatically enabled on them.

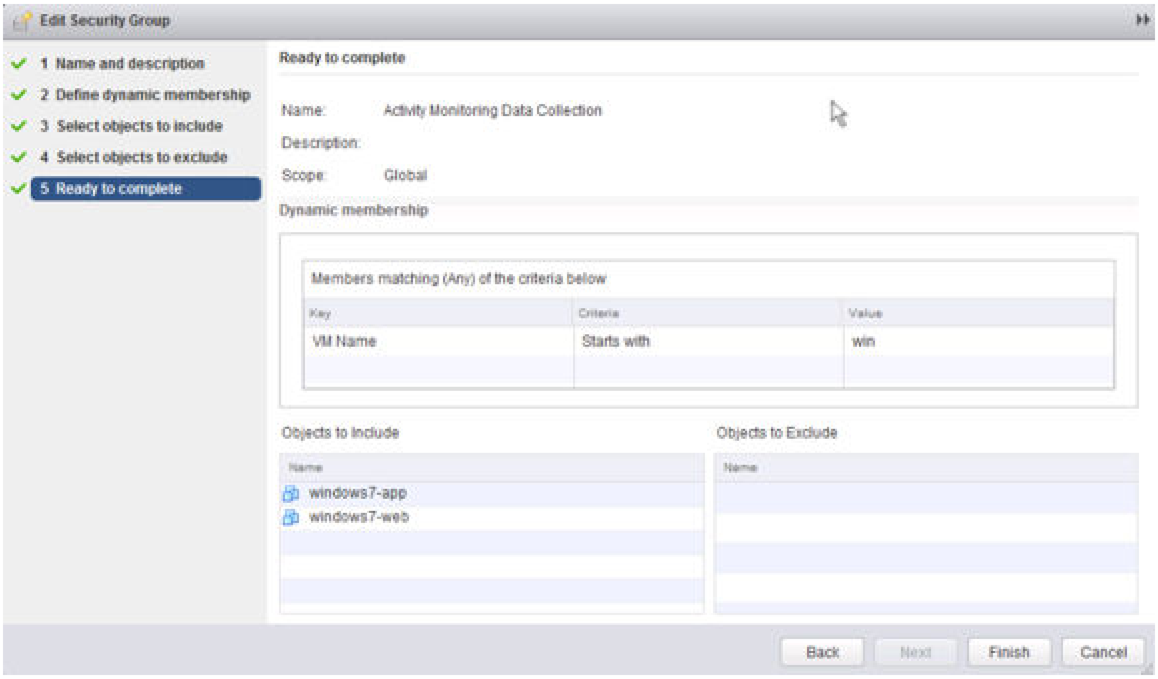
Edit Security Group
✓ 1 Name and description
✓ 2 Define dynamic membership
✓ 3 Select objects to include
✓ 4 Select objects to exclude
✓ 5 Ready to complete
Back
Next
Finish
Cancel

Ready to complete

Name: Activity Monitoring Data Collection

Description:

Scope: Global

Dynamic membership

Members matching (Any) of the criteria below

Key	Criteria	Value
VM Name	Starts with	win

Objects to Include

Name
windows7-app
windows7-web

Objects to Exclude

Name

Activity Monitoring Scenarios

This section describes some hypothetical scenarios for Activity Monitoring.

User Access to Applications

Our hypothetical company, ACME Enterprise, only permits approved users to access specific applications on corporate assets.

Their security policy mandates are:

- Allow only authorized users to access critical business applications
- Allow only authorized applications on corporate servers
- Allow access to only required ports from specific networks

Based on the above, they need controlled access for employees based on user identity to safeguard corporate assets. As a starting point, the security operator at ACME Enterprise needs to be able to verify that only administrative access is allowed to the MS SQL servers.

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then **Activity Monitoring**.

3 Click the **Inbound Activity** tab.

4 In **Outbound from** leave value as **All Observed AD Groups** to see access from any and all employees.

5 In **Where destination virtual machine**, select **includes**, and leave **all observed destination virtual machines** selected.

6 In **And where destination application**, select **includes**, click **all observed destination applications** and select the MS SQL servers.

7 Click **Search**.

The search results show that only administrative users are accessing the MS SQL servers. Notice that are no groups (such as Finance or HR) accessing these servers.

8 We can now invert this query by setting the **Outbound from** value to HR and Finance AD groups.

9 Click **Search**.

No records are displayed, confirming that no users from either of these groups can access MS SQL servers.

Applications on Datacenter

As part of their security policies, ACME Enterprise needs Visibility into all data center applications. This can help Identify rogue applications that either capture confidential information or siphon sensitive data to external sources.

John, Cloud Administrator at ACME Enterprise, wants to confirm that access to the SharePoint server is only through Internet Explorer and no rogue application (such as FTP or RDP) can access this server.

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then **Activity Monitoring**.

3 Click the **VM Activity** tab.

4 In **Where source VM**, select **includes**, and leave **All observed virtual machines** selected to capture traffic originating from all virtual machines in the datacenter.

5 In **Where destination VM**, select **includes**, click **All observed virtual machines**, and select the SharePoint server.

6 Click **Search**.

The **Outbound App Product Name** column in the search results show that all access to the SharePoint server was only through Internet Explorer. The relatively homogenous search results indicate that there is a firewall rule applied to this SharePoint server preventing all other access methods.

Also note that the search results display the source user of the observed traffic rather than the source group. Clicking the arrow in the search result displays details about the source user such as the AD group to which the user belongs.

Verify Open Ports

Once John Admin knows that the ACME Enterprise share point server is being accessed only by authorized applications, he can ensure that the company allows only required ports to be open based on expected use.

Prerequisites

In the Applications on Datacenter scenario, John Admin had observed traffic to the ACME Enterprise share point server. He now wants to ensure that all access from the share point server to the MSSQL server is through expected protocols and applications.

Procedure

1. Click the **Go Home** icon.
2. Click **vCenter Home** and then click **Virtual Machines**.
3. Select **win_sharepoint** and then click the **Monitor** tab.
4. Click **Activity Monitoring**.
5. In **Where destination**, select **win2K-MSSQL**.
6. Click **Search**.

Search results show traffic from the share point server to the MSSQL server. The **User** and **Outbound App** columns show that only systems processes are connecting to the MSSQL server, which is what John expected to see.

The **Inbound Port** and **App** columns show that all access is to the MSSQL server running on the destination server.

Since there are too many records in the search results for John to analyze in a web browser, he can export all the entire result set and save the file in a CSV format by clicking the icon on the bottom right side of the page.

Enable Data Collection

You must enable data collection for one or more virtual machines on a vCenter Server before running an Activity Monitoring report. Before running a report, ensure that the enabled virtual machines are active and are generating network traffic.

You should also register NSX Manager with the AD Domain Controller. See Register a Windows Domain with NSX Manager.

Note that only active connections are tracked by Activity Monitoring. Virtual machine traffic blocked by firewall rules at the vNIC level is not reflected in reports.

Enable Data Collection on a Single Virtual Machine

You must enable data collection at least five minutes before running an Activity Monitoring report.

Prerequisites

Procedure

1 Log in to the vSphere Web Client.

2 Click **vCenter** and then click **VMs and Templates**.

3 Select a virtual machine from the left inventory panel.

4 Click the **Manage** tab and then click the **Settings** tab.

5 Click **NSX Activity Monitoring** from the left panel.

6 Click **Edit**.

7 In the Edit NSX Activity Monitoring Data Collection Settings dialog box, click **Yes**.

Enable Data Collection for Multiple Virtual Machines

The Activity Monitoring Data Collection security group is a pre-defined security group. You can add multiple virtual machines to this security group at a time, and data collection is enabled on all of these virtual machines.

You must enable data collection at least five minutes before running an Activity Monitoring report.

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **Service Composer**.

3 Click the **Security Groups** tab.

4 Select the Activity Monitoring Data Collection security group and click the **Edit** (✏) icon.

5 Follow the wizard to add virtual machines to the security group.

Data collection is enabled on all virtual machines you added to this security group, and disabled on any virtual machines you excluded from the security group.

View Virtual Machine Activity Report

You can view traffic to or from a virtual machine or a set of virtual machines in your environment.

You can either do a quick query using the default search criteria by clicking **Search**, or tailor the query according to your requirements.

Prerequisites

- Guest Introspection must be installed in your environment.
- A domain must be registered with NSX Manager. For information on domain registration, see Register a Windows Domain with NSX Manager.
- Data collection must be enabled on one or more virtual machines.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then **Activity Monitoring**.
3. Click the **VM Activity** tab.
4. Click the link next to **Where source**. Select the virtual machines for which you want to view outbound traffic. Indicate whether you want to include or exclude the selected virtual machine(s) from the report.
5. Click the link next to **Where destination**. Select the virtual machines for which you want to view inbound traffic. Indicate whether you want to include or exclude the selected virtual machine(s) from the report.
6. Click the **During period** (📝) icon and select the time period for the search.
7. Click **Search**.

Search results filtered by the specified criterion are displayed. Click a row to view detailed information about the user for that row.

You can export a specific record or all records on this page and save them to a directory in a .csv format by clicking the ⇨▾ icon on the bottom right side of the page.

View Inbound Activity

You can view all inbound activity to a server by desktop pool, security group, or AD group.

Figure 23-4. View inbound activity

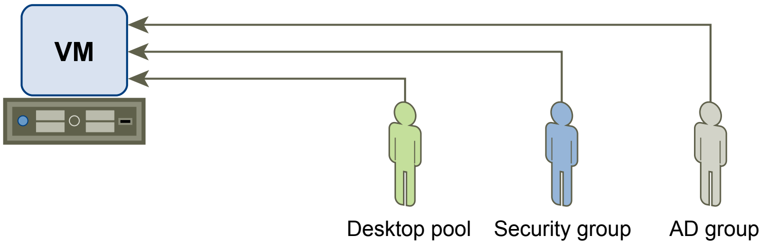

You can either do a quick query using the default search criteria by clicking **Search**, or tailor the query according to your requirements.

Prerequisites

- Guest Introspection must be installed in your environment.
- A domain must be registered with NSX Manager. For information on domain registration, see Register a Windows Domain with NSX Manager.
- Data collection must be enabled on one or more virtual machines.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then **Activity Monitoring**.
3. Click the **Inbound Activity** tab.
4. Click the link next to **Originating from**.
5. Select the type of user group that you want to view activity for.
6. In **Filter type**, select one or more group and click OK.
7. In **Where destination virtual machine**, select **includes** or **excludes** to indicate whether the selected virtual machines should be included in or excluded from the search.
8. Click the link next to **And where destination virtual machine**.
9. Select one or more virtual machine and click **OK**.
10. In **And where destination application**, select **includes** or **excludes** to indicate whether the selected applications should be included in or excluded from the search.
11. Click the link next to **And where destination application**.
12. Select one or more application and click **OK**.
13. Click the **During period** (✎) icon and select the time period for the search.
14. Click **Search**.

Search results filtered by the specified criterion are displayed. Click anywhere in the results table to view information about the users that accessed the specified virtual machines and applications.

You can export a specific record or all records on this page and save them to a directory in a .csv format by clicking the ⇥ icon on the bottom right side of the page.

View Outbound Activity

You can view what applications are being run by a security group or desktop pool and then drill down into the report to find out which client applications are making outbound connections by a particular group of users. You can also discover all user groups and users who are accessing a particular application, which can help you determine if you need to adjust identity firewall in your environment.

Figure 23-5. VIew Outbound activity

Prerequisites

- Guest Introspection must be installed in your environment.
- A domain must be registered with NSX Manager. For information on domain registration, see Register a Windows Domain with NSX Manager.
- Data collection must be enabled on one or more virtual machines.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then **Activity Monitoring**.
3. Ensure that the **Outbound Activity** tab is selected in the left pane.
4. Click the link next to **Originating from**.

   All groups discovered through guest introspection are displayed.
5. Select the type of user group that you want to view resource utilization for.
6. In **Filter**, select one or more group and click **OK**.
7. In **Where application**, select **includes** or **excludes** to indicate whether the selected application should be included in or excluded from the search.
8. Click the link next to **Where application**.
9. Select one or more application and click **OK**.
10. In **And where destination**, select **includes** or **excludes** to indicate whether the selected virtual machines should be included in or excluded from the search.
11. Click the link next to **And where destination**.
12. Select one or more virtual machine and click **OK**.
13. Click the **During period** (📅) icon and select the time period for the search.
14. Click **Search**.

    Scroll to the right to see all the information displayed.

Search results filtered by the specified criterion are displayed. Click a row to view information about users within that AD group that used the specified application to access the specified virtual machines.

You can export a specific record or all records on this page and save them to a directory in a .csv format by clicking the 📤▾ icon on the bottom right side of the page.

View Interaction between Inventory Containers

You can view the traffic passing between defined containers such as AD groups, security groups and/or desktop pools. This can help you identify and configure access to shared services and to resolve mis-configured relationships between Inventory container definitions, desktop pools and AD groups.

Figure 23-6. Interaction between containers

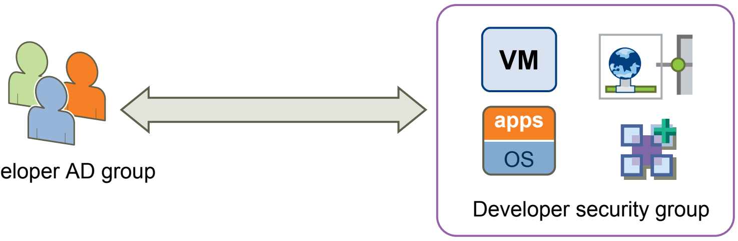

You can either do a quick query using the default search criteria by clicking **Search**, or tailor the query according to your requirements.

Prerequisites

- Guest Introspection must be installed in your environment.
- A domain must be registered with NSX Manager. For information on domain registration, see Register a Windows Domain with NSX Manager.
- Data collection must be enabled on one or more virtual machines.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then **Activity Monitoring**.
3. Select the **Inter Container Interaction** tab in the left pane.
4. Click the link next to **Originating from**.

   All groups discovered through guest introspection are displayed.
5. Select the type of user group that you want to view resource utilization for.
6. In **Filter**, select one or more group and click **OK**.
7. In **Where the destination is**, select **is** or **is not** to indicate whether the selected group should be included in or excluded from the search.
8. Click the link next to **Where the destination is**.
9. Select the group type.
10. In **Filter**, select one or more group and click **OK**.
11. Click the **During period** (📝) icon and select the time period for the search.
12. Click **Search**.

Search results filtered by the specified criterion are displayed. Click in a row to view information about the users that accessed the specified containers.

You can export a specific record or all records on this page and save them to a directory in a .csv format by clicking the icon on the bottom right side of the page.

Example: Interaction between Inventory Containers Query

- Verify allowed communication

  If you have defined containers in your vCenter inventory and then added a rule to allow communication between these containers, you can verify that the rule is working by running this query with the two containers specified in the **Originating from** and **Where the destination is** fields.

- Verify denied communication

  If you have defined containers in your vCenter inventory and then added a rule to deny communication between these containers, you can verify that the rule is working by running this query with the two containers specified in the **Originating from** and **Where the destination is** fields.

- Verify denied intra-container communication

  If you have implemented a policy that does not allow members of a container communicating with other members of the same container, you can run this query to verify that the policy works. Select the container in both **Originating from** and **Where the destination is** fields.

- Eliminate unnecessary access

  Suppose you have defined containers in your vCenter inventory and then added a rule to allow communication between these containers. There may be members in either container that do not interact with the other container at all. You may then choose to remove these members from the appropriate container to optimize security control. To retrieve such a list, select the appropriate containers in both **Originating from** and **Where the destination is** fields. Select **is not** next to the **Where the destination is** field.

View Outbound AD Group Activity

You can view the traffic between members of defined Active Directory groups and can use this data to fine-tune your firewall rules.

You can either do a quick query using the default search criteria by clicking **Search**, or tailor the query according to your requirements.

Prerequisites

- Guest Introspection must be installed in your environment.
- A domain must be registered with NSX Manager. For information on domain registration, see Register a Windows Domain with NSX Manager.
- Data collection must be enabled on one or more virtual machines.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then **Activity Monitoring**.
3. Select the **AD Groups & Containers** tab in the left pane.
4. Click the link next to **Originating from**.

   All groups discovered through guest introspection are displayed.
5. Select the type of user group that you want to include in the search.
6. In **Filter**, select one or more group and click **OK**.
7. In **Where AD Group**, select **includes** or **excludes** to indicate whether the selected AD group should be included in or excluded from the search.
8. Click the link next to **Where AD Group**.
9. Select one or more AD groups and click **OK**.
10. Click the **During period** (📝) icon and select the time period for the search.
11. Click **Search**.

Search results filtered by the specified criterion are displayed. Click in a row to view information about the members of the specified AD group that are accessing network resources from within the specified security group or desktop pool.

You can export a specific record or all records on this page and save them to a directory in a .csv format by clicking the 📄▾ icon on the bottom right side of the page.

Override Data Collection

In case of an emergency such as a network overload, you can turn off data collection at a global level. This overrides all other data collection settings.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then **Activity Monitoring**.
3. Click the **Settings** tab.
4. Select the vCenter Server for which you want to overwrite data collection.
5. Click **Edit**.
6. De-select **Collect reporting data**.
7. Click **OK**.

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 vCenter Web Client, navigate to **Home > Networking & Security > Traceflow**.

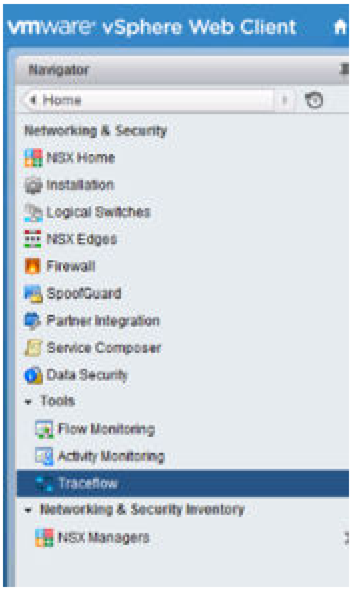

2 Select the traffic type: Unicast, broadcast, or multicast.

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.

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 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.

6 For a Layer 2 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.

7 Configure other required and optional settings.

8 Click **Trace**.

Example: Scenarios

The following example shows a Layer 2 traceflow involving two VMs that are running on a single ESXi host. The two VMs are connected to a single logical switch.

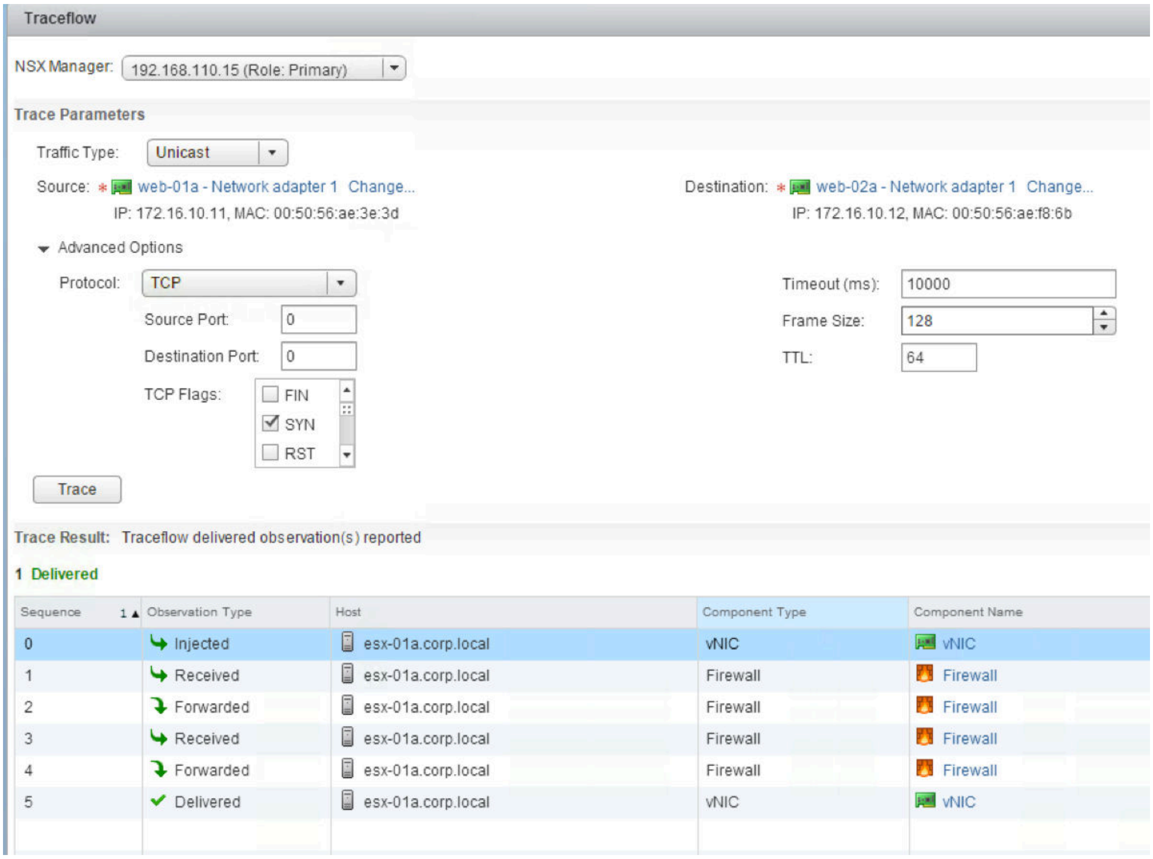

Traceflow

NSX Manager: 192.168.110.15 (Role: Primary)

Trace Parameters

Traffic Type: Unicast

Source: * web-01a - Network adapter 1 Change...
IP: 172.16.10.11, MAC: 00:50:56:ae:3e:3d

Destination: * web-02a - Network adapter 1 Change...
IP: 172.16.10.12, MAC: 00:50:56:ae:f8:6b

Advanced Options

Protocol: TCP

Source Port: 0

Destination Port: 0

TCP Flags: □ FIN ☑ SYN □ RST

Timeout (ms): 10000

Frame Size: 128

TTL: 64

Trace

Trace Result: Traceflow delivered observation(s) reported

1 Delivered

Sequence	Observation Type	Host	Component Type	Component Name
0	Injected	esx-01a.corp.local	vNIC	vNIC
1	Received	esx-01a.corp.local	Firewall	Firewall
2	Forwarded	esx-01a.corp.local	Firewall	Firewall
3	Received	esx-01a.corp.local	Firewall	Firewall
4	Forwarded	esx-01a.corp.local	Firewall	Firewall
5	Delivered	esx-01a.corp.local	vNIC	vNIC

The following example shows a Layer 2 traceflow involving two VMs that are running on two different ESXi hosts. The two VMs are connected to a single logical switch.

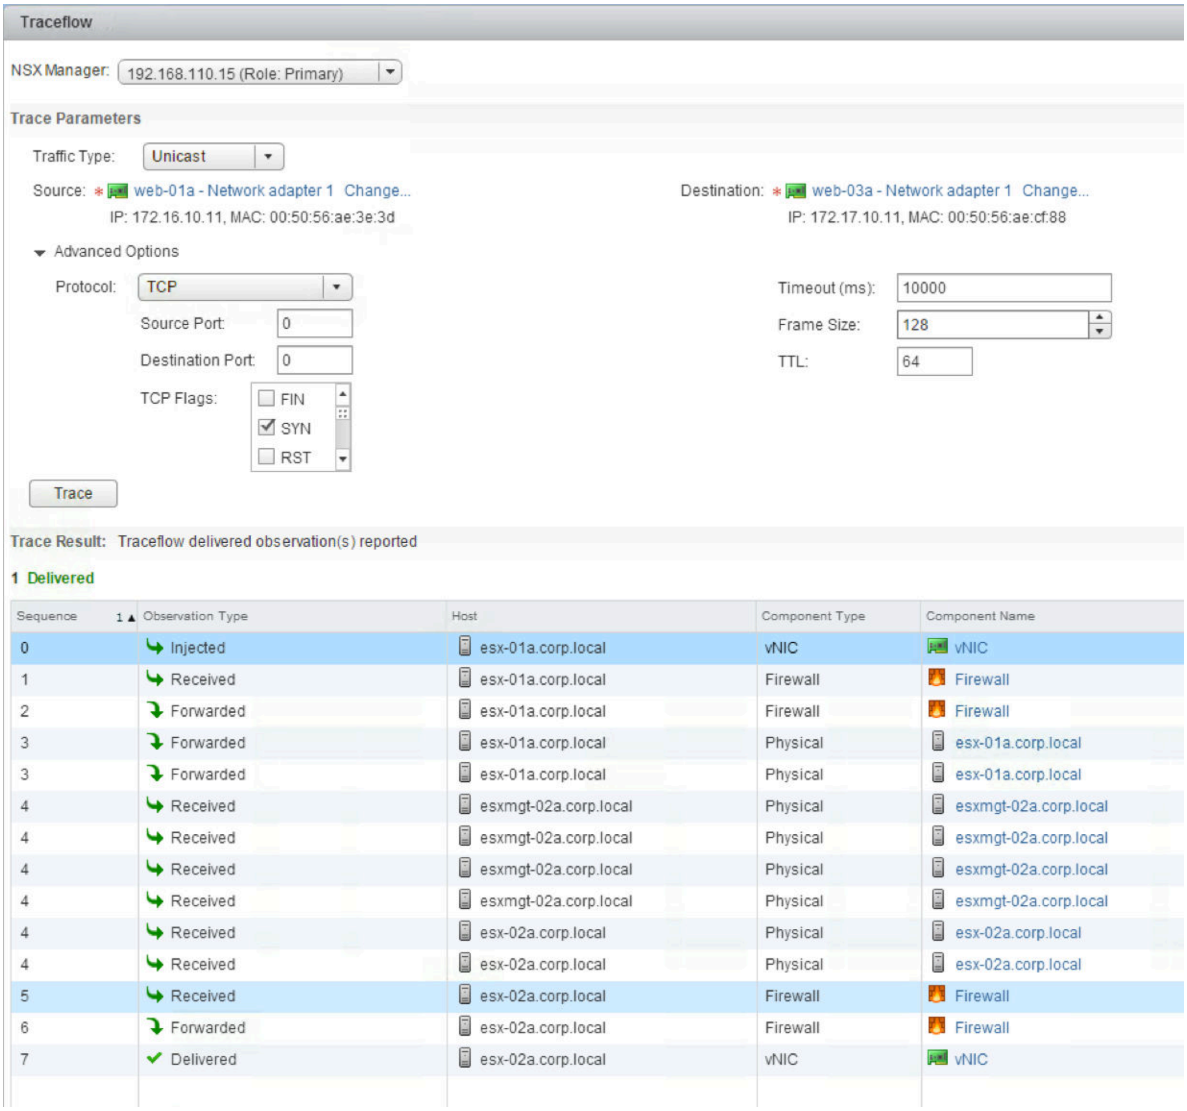

Traceflow

NSX Manager: 192.168.110.15 (Role: Primary)

Trace Parameters

Traffic Type: Unicast

Source: * web-01a - Network adapter 1 Change...
IP: 172.16.10.11, MAC: 00:50:56:ae:3e:3d

Destination: * web-03a - Network adapter 1 Change...
IP: 172.17.10.11, MAC: 00:50:56:ae:cf:88

Advanced Options

Protocol: TCP

Source Port: 0

Destination Port: 0

TCP Flags: ☐ FIN ☑ SYN ☐ RST

Timeout (ms): 10000

Frame Size: 128

TTL: 64

Trace

Trace Result: Traceflow delivered observation(s) reported

1 Delivered

Sequence 1 ▲	Observation Type	Host	Component Type	Component Name
0	Injected	esx-01a.corp.local	vNIC	vNIC
1	Received	esx-01a.corp.local	Firewall	Firewall
2	Forwarded	esx-01a.corp.local	Firewall	Firewall
3	Forwarded	esx-01a.corp.local	Physical	esx-01a.corp.local
3	Forwarded	esx-01a.corp.local	Physical	esx-01a.corp.local
4	Received	esxmgt-02a.corp.local	Physical	esxmgt-02a.corp.local
4	Received	esxmgt-02a.corp.local	Physical	esxmgt-02a.corp.local
4	Received	esxmgt-02a.corp.local	Physical	esxmgt-02a.corp.local
4	Received	esxmgt-02a.corp.local	Physical	esxmgt-02a.corp.local
4	Received	esx-02a.corp.local	Physical	esx-02a.corp.local
4	Received	esx-02a.corp.local	Physical	esx-02a.corp.local
5	Received	esx-02a.corp.local	Firewall	Firewall
6	Forwarded	esx-02a.corp.local	Firewall	Firewall
7	Delivered	esx-02a.corp.local	vNIC	vNIC

The following example shows a Layer 3 traceflow. The two VMs are connected to two different logical switches that are separated by a logical router.

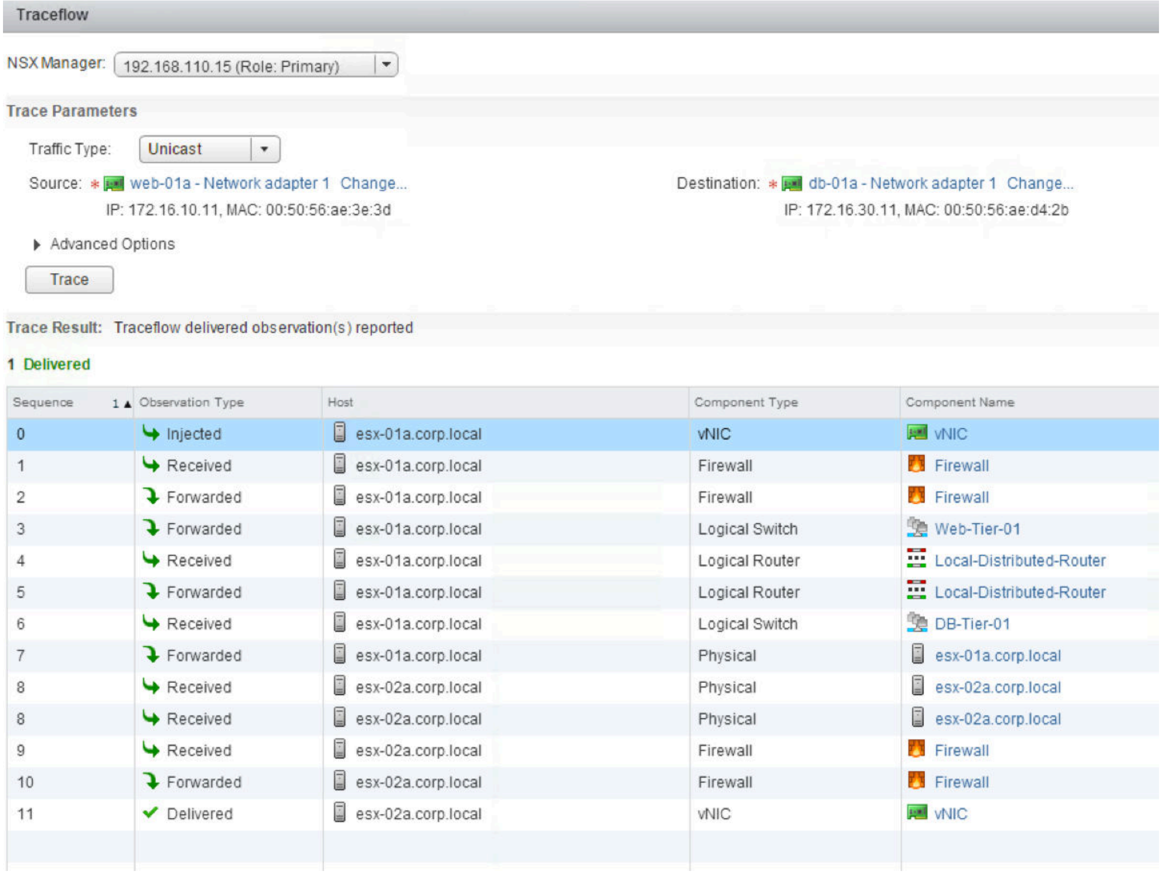

Traceflow

NSX Manager: 192.168.110.15 (Role: Primary)

Trace Parameters

Traffic Type: Unicast

Source: * web-01a - Network adapter 1 Change...
IP: 172.16.10.11, MAC: 00:50:56:ae:3e:3d

Destination: * db-01a - Network adapter 1 Change...
IP: 172.16.30.11, MAC: 00:50:56:ae:d4:2b

Advanced Options

Trace

Trace Result: Traceflow delivered observation(s) reported

1 Delivered

Sequence 1 ▲	Observation Type	Host	Component Type	Component Name
0	Injected	esx-01a.corp.local	vNIC	vNIC
1	Received	esx-01a.corp.local	Firewall	Firewall
2	Forwarded	esx-01a.corp.local	Firewall	Firewall
3	Forwarded	esx-01a.corp.local	Logical Switch	Web-Tier-01
4	Received	esx-01a.corp.local	Logical Router	Local-Distributed-Router
5	Forwarded	esx-01a.corp.local	Logical Router	Local-Distributed-Router
6	Received	esx-01a.corp.local	Logical Switch	DB-Tier-01
7	Forwarded	esx-01a.corp.local	Physical	esx-01a.corp.local
8	Received	esx-02a.corp.local	Physical	esx-02a.corp.local
8	Received	esx-02a.corp.local	Physical	esx-02a.corp.local
9	Received	esx-02a.corp.local	Firewall	Firewall
10	Forwarded	esx-02a.corp.local	Firewall	Firewall
11	✓ Delivered	esx-02a.corp.local	vNIC	vNIC

The following example shows a broadcast traceflow in a deployment that has three VMs connected to a single logical switch. Two of the VMs are on one host (esx-01a), and the third is on another host (esx-02a). The broadcast is sent from one of the VMs on host 192.168.210.53.

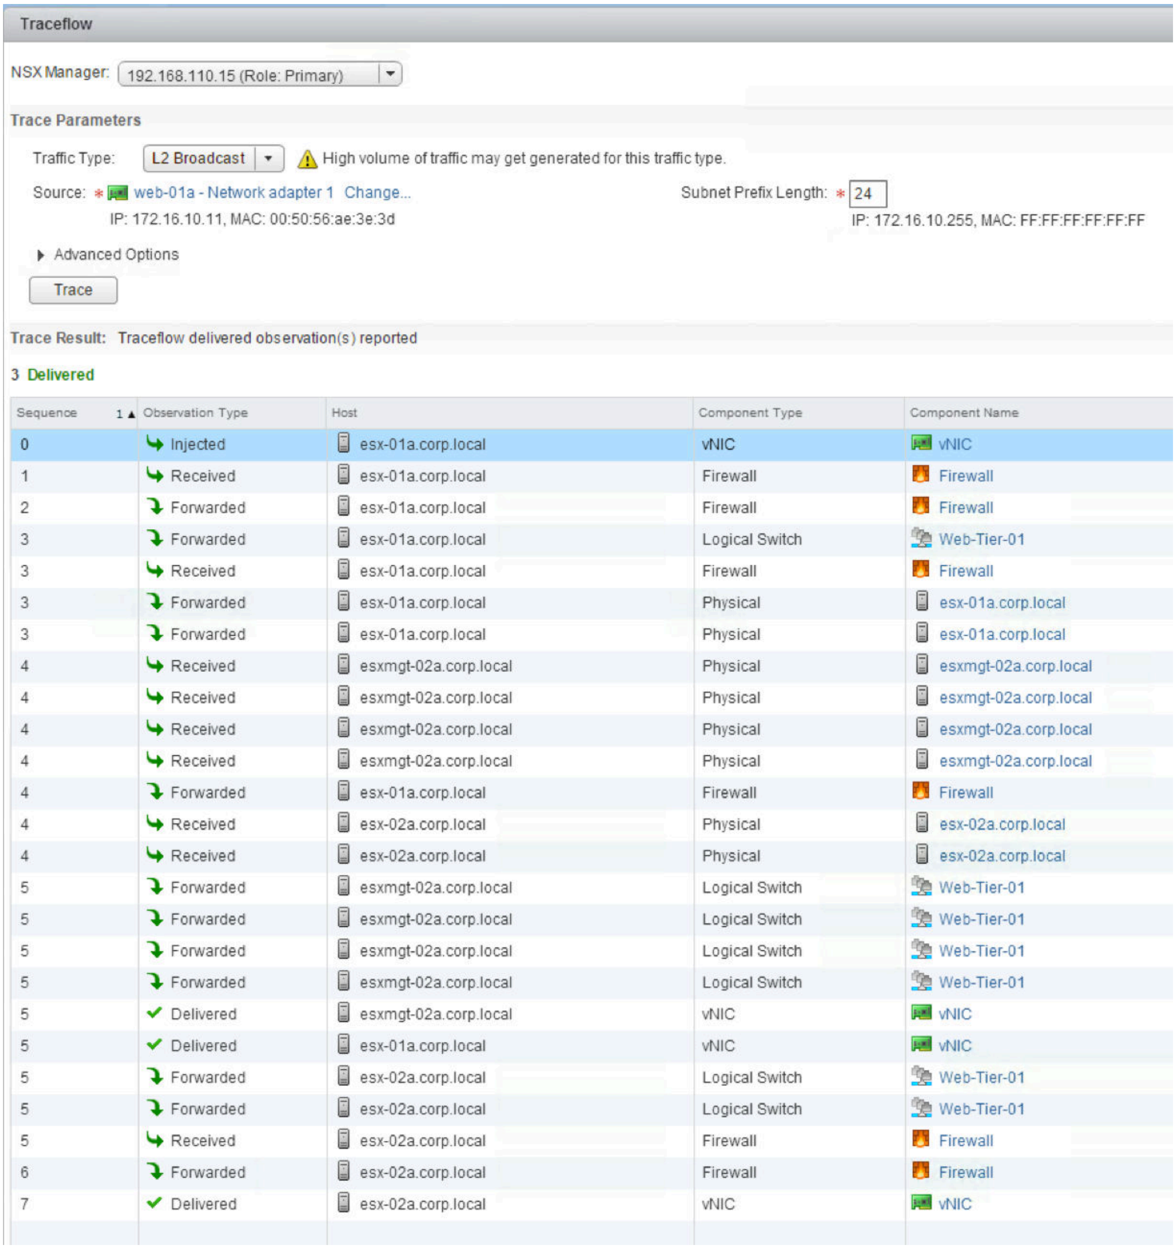

Traceflow

NSX Manager: 192.168.110.15 (Role: Primary)

Trace Parameters

Traffic Type: L2 Broadcast ⚠ High volume of traffic may get generated for this traffic type.

Source: * web-01a - Network adapter 1 Change...
IP: 172.16.10.11, MAC: 00:50:56:ae:3e:3d

Subnet Prefix Length: * 24
IP: 172.16.10.255, MAC: FF:FF:FF:FF:FF:FF

Advanced Options

Trace

Trace Result: Traceflow delivered observation(s) reported

3 Delivered

Sequence 1 ▲	Observation Type	Host	Component Type	Component Name
0	Injected	esx-01a.corp.local	vNIC	vNIC
1	Received	esx-01a.corp.local	Firewall	Firewall
2	Forwarded	esx-01a.corp.local	Firewall	Firewall
3	Forwarded	esx-01a.corp.local	Logical Switch	Web-Tier-01
3	Received	esx-01a.corp.local	Firewall	Firewall
3	Forwarded	esx-01a.corp.local	Physical	esx-01a.corp.local
3	Forwarded	esx-01a.corp.local	Physical	esx-01a.corp.local
4	Received	esxmgt-02a.corp.local	Physical	esxmgt-02a.corp.local
4	Received	esxmgt-02a.corp.local	Physical	esxmgt-02a.corp.local
4	Received	esxmgt-02a.corp.local	Physical	esxmgt-02a.corp.local
4	Received	esxmgt-02a.corp.local	Physical	esxmgt-02a.corp.local
4	Forwarded	esx-01a.corp.local	Firewall	Firewall
4	Received	esx-02a.corp.local	Physical	esx-02a.corp.local
4	Received	esx-02a.corp.local	Physical	esx-02a.corp.local
5	Forwarded	esxmgt-02a.corp.local	Logical Switch	Web-Tier-01
5	Forwarded	esxmgt-02a.corp.local	Logical Switch	Web-Tier-01
5	Forwarded	esxmgt-02a.corp.local	Logical Switch	Web-Tier-01
5	Forwarded	esxmgt-02a.corp.local	Logical Switch	Web-Tier-01
5	Delivered	esxmgt-02a.corp.local	vNIC	vNIC
5	Delivered	esx-01a.corp.local	vNIC	vNIC
5	Forwarded	esx-02a.corp.local	Logical Switch	Web-Tier-01
5	Forwarded	esx-02a.corp.local	Logical Switch	Web-Tier-01
5	Received	esx-02a.corp.local	Firewall	Firewall
6	Forwarded	esx-02a.corp.local	Firewall	Firewall
7	Delivered	esx-02a.corp.local	vNIC	vNIC

The following example shows what happens when multicast traffic is sent in a deployment that has multicast configured.

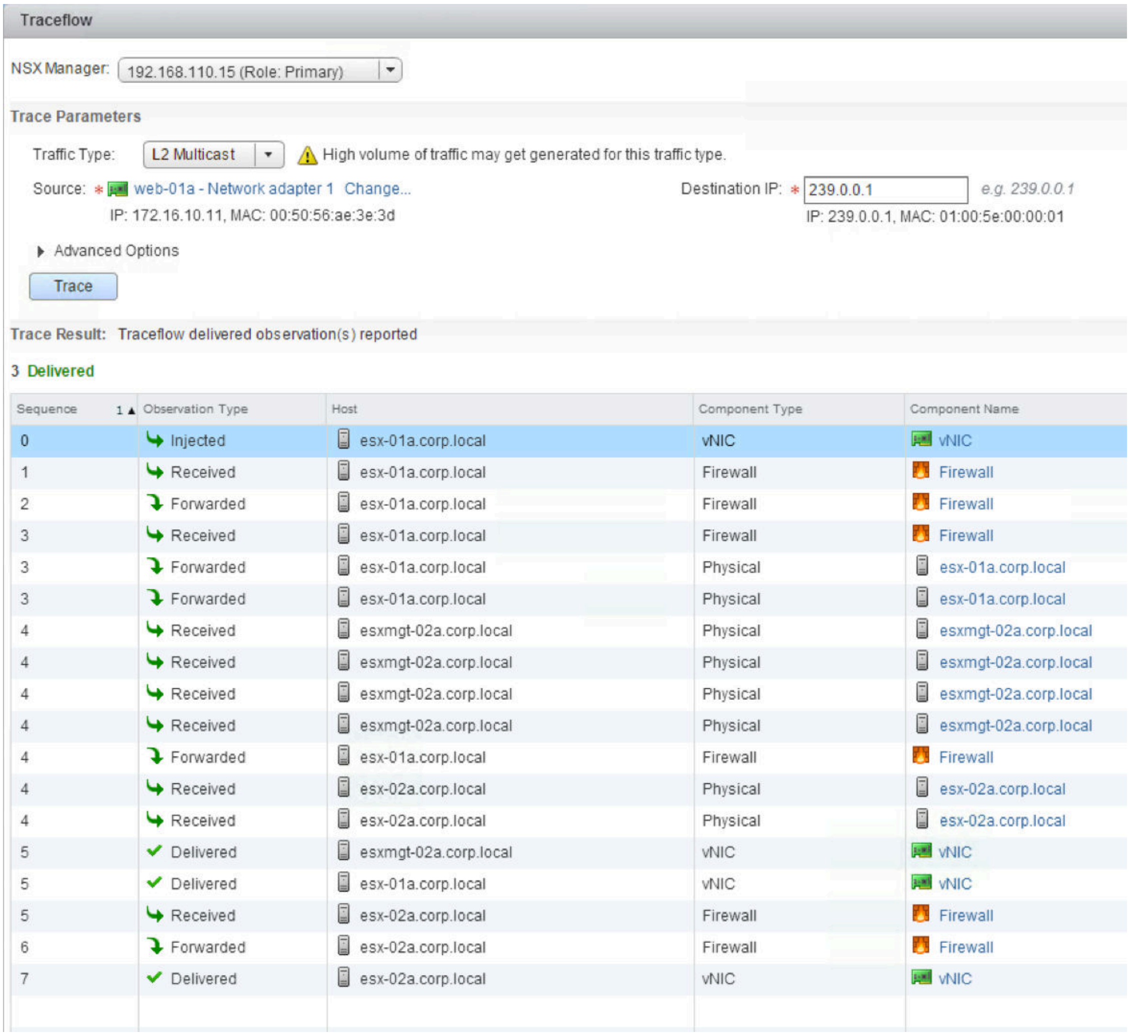

Traceflow

NSX Manager: 192.168.110.15 (Role: Primary)

Trace Parameters

Traffic Type: L2 Multicast ⚠ High volume of traffic may get generated for this traffic type.

Source: * web-01a - Network adapter 1 Change...
IP: 172.16.10.11, MAC: 00:50:56:ae:3e:3d

Destination IP: * 239.0.0.1 e.g. 239.0.0.1
IP: 239.0.0.1, MAC: 01:00:5e:00:00:01

Advanced Options

Trace

Trace Result: Traceflow delivered observation(s) reported

3 Delivered

Sequence	Observation Type	Host	Component Type	Component Name
0	Injected	esx-01a.corp.local	vNIC	vNIC
1	Received	esx-01a.corp.local	Firewall	Firewall
2	Forwarded	esx-01a.corp.local	Firewall	Firewall
3	Received	esx-01a.corp.local	Firewall	Firewall
3	Forwarded	esx-01a.corp.local	Physical	esx-01a.corp.local
3	Forwarded	esx-01a.corp.local	Physical	esx-01a.corp.local
4	Received	esxmgt-02a.corp.local	Physical	esxmgt-02a.corp.local
4	Received	esxmgt-02a.corp.local	Physical	esxmgt-02a.corp.local
4	Received	esxmgt-02a.corp.local	Physical	esxmgt-02a.corp.local
4	Received	esxmgt-02a.corp.local	Physical	esxmgt-02a.corp.local
4	Forwarded	esx-01a.corp.local	Firewall	Firewall
4	Received	esx-02a.corp.local	Physical	esx-02a.corp.local
4	Received	esx-02a.corp.local	Physical	esx-02a.corp.local
5	✓ Delivered	esxmgt-02a.corp.local	vNIC	vNIC
5	✓ Delivered	esx-01a.corp.local	vNIC	vNIC
5	Received	esx-02a.corp.local	Firewall	Firewall
6	Forwarded	esx-02a.corp.local	Firewall	Firewall
7	✓ Delivered	esx-02a.corp.local	vNIC	vNIC

The following example shows what happens when a traceflow is dropped because of a distributed firewall rule that blocks ICMP traffic sent to the destination address. Notice that the traffic never leaves the original host, even though the destination VM is on another host.

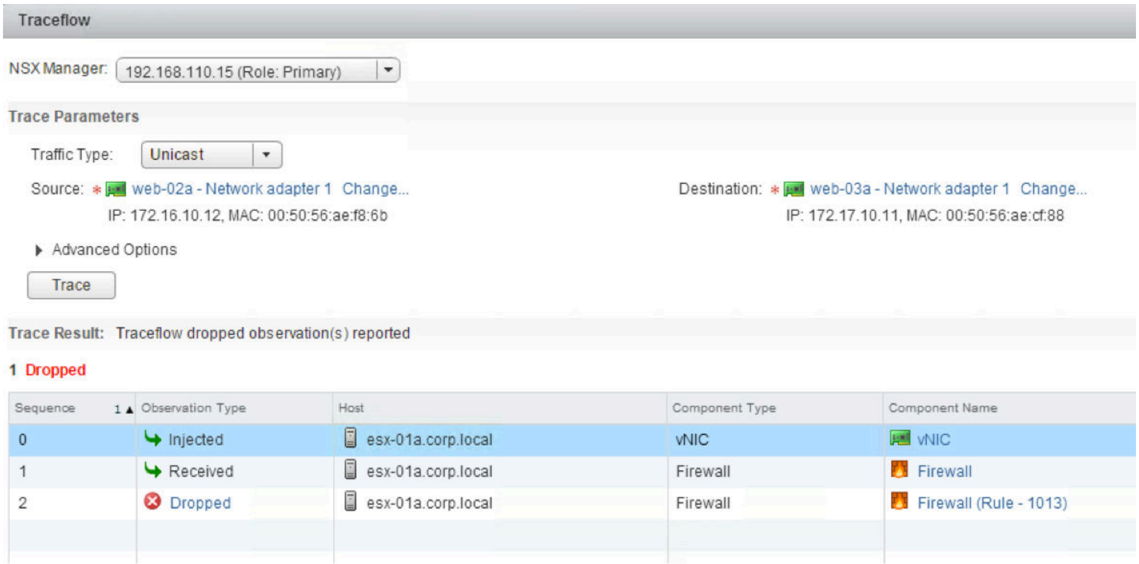

Traceflow

NSX Manager: 192.168.110.15 (Role: Primary)

Trace Parameters

Traffic Type: Unicast

Source: * web-02a - Network adapter 1 Change...
IP: 172.16.10.12, MAC: 00:50:56:ae:f8:6b

Destination: * web-03a - Network adapter 1 Change...
IP: 172.17.10.11, MAC: 00:50:56:ae:cf:88

Advanced Options

Trace

Trace Result: Traceflow dropped observation(s) reported

1 Dropped

Sequence	Observation Type	Host	Component Type	Component Name
0	Injected	esx-01a.corp.local	vNIC	vNIC
1	Received	esx-01a.corp.local	Firewall	Firewall
2	✖ Dropped	esx-01a.corp.local	Firewall	Firewall (Rule - 1013)

The following example shows what happens when a traceflow destination is on the other side of an edge services gateway, such as an IP address on the Internet or any internal destination that must be routed through the edge services gateway. The traceflow is not allowed, by design, because traceflow is supported for destinations that are either on the same subnet or are reachable through distributed logical routers (DLRs).

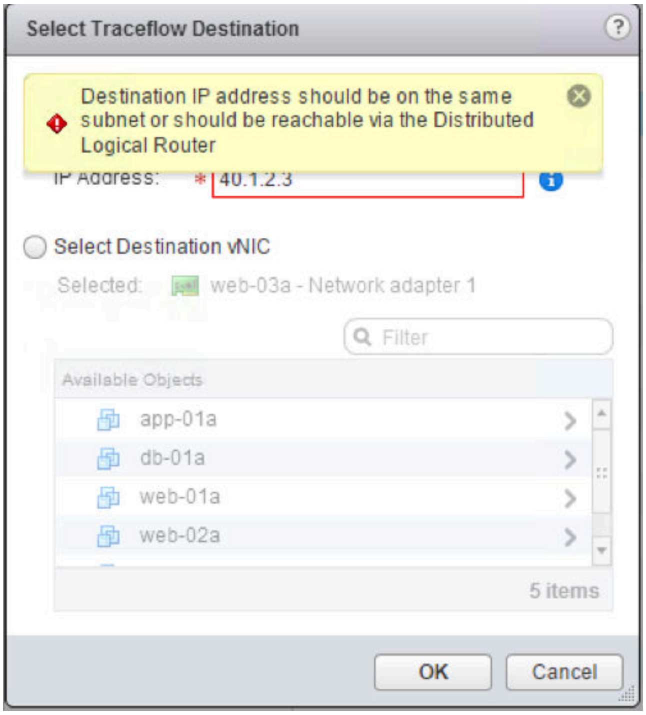

Select Traceflow Destination

Destination IP address should be on the same subnet or should be reachable via the Distributed Logical Router

IP Address: * 40.1.2.3

○ Select Destination vNIC

Selected: web-03a - Network adapter 1

The following example shows what happens when the traceflow destination is a VM that is on another subnet and is powered off.

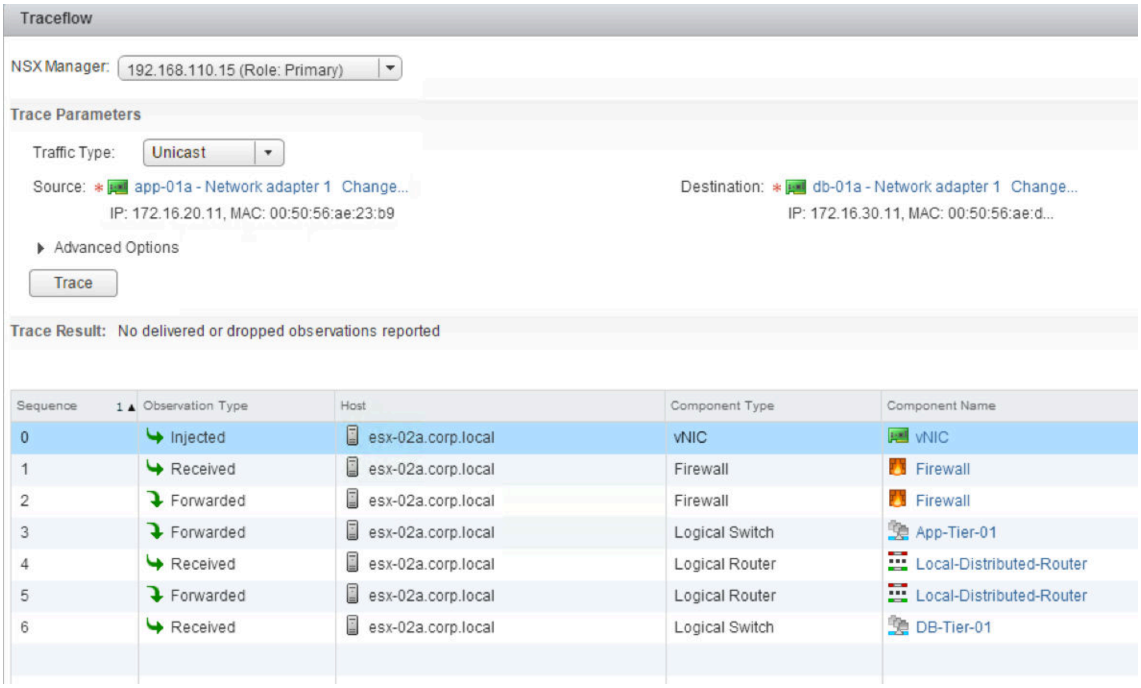

Traceflow

NSX Manager: 192.168.110.15 (Role: Primary)

Trace Parameters

Traffic Type: Unicast

Source: * app-01a - Network adapter 1 Change...
IP: 172.16.20.11, MAC: 00:50:56:ae:23:b9

Destination: * db-01a - Network adapter 1 Change...
IP: 172.16.30.11, MAC: 00:50:56:ae:d...

Trace Result: No delivered or dropped observations reported

Sequence 1 ▲	Observation Type	Host	Component Type	Component Name
0	Injected	esx-02a.corp.local	vNIC	vNIC
1	Received	esx-02a.corp.local	Firewall	Firewall
2	Forwarded	esx-02a.corp.local	Firewall	Firewall
3	Forwarded	esx-02a.corp.local	Logical Switch	App-Tier-01
4	Received	esx-02a.corp.local	Logical Router	Local-Distributed-Router
5	Forwarded	esx-02a.corp.local	Logical Router	Local-Distributed-Router
6	Received	esx-02a.corp.local	Logical Switch	DB-Tier-01

NSX Edge VPN Configuration Examples

24

This scenario contains configuration examples for a basic point-to-point 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 192.0.2.0/24 to the internet. NSX Edge interfaces are configured as follows:

- Uplink interface: 198.51.100.1
- Internal interface: 192.0.2.1

The remote gateway connects the 172.16.0.0/16 internal network to the internet. The remote gateway interfaces are configured as follows:

- Uplink interface: 10.24.120.90/24
- Internal interface: 172.16.0.1/16

Figure 24-1. NSX Edge connecting to a remote VPN gateway

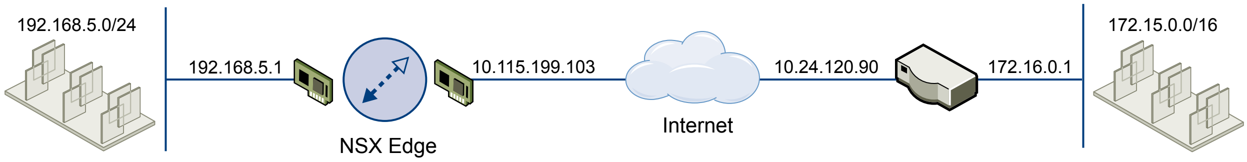

Note For NSX Edge to NSX Edge IPSEC tunnels, you can use the same scenario by setting up the second NSX Edge as the remote gateway.

This chapter includes the following topics:

- Terminology
- IKE Phase 1 and Phase 2
- Configuring IPSec VPN Service Example
- Using a Cisco 2821 Integrated Services Router
- Using a Cisco ASA 5510
- Configuring a WatchGuard Firebox X500
- Troubleshooting NSX Edge Configuration Example

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.

These are some of the standards you may 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 is designed to be 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.
- Diffie-Hellman (DH) key exchange is a cryptographic protocol that allows two parties that have no prior knowledge of each other to jointly establish a shared secret key over an insecure communications channel. VSE supports DH group 2 (1024 bits) and group 5 (1536 bits).

IKE Phase 1 and Phase 2

IKE is a standard method used to arrange secure, 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
- TripleDES / AES [Configurable]
- SHA-1
- MODP group 2 (1024 bits)
- pre-shared secret [Configurable]
- SA lifetime of 28800 seconds (eight hours) with no kbytes rekeying
- ISAKMP aggressive mode disabled

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 one keys as a base or by performing a new key exchange). The IKE Phase 2 parameters supported by NSX Edge are:

- TripleDES / AES [Will match the Phase 1 setting]

- SHA-1
- ESP tunnel mode
- MODP group 2 (1024 bits)
- Perfect forward secrecy for rekeying
- SA lifetime of 3600 seconds (one hour) with no kbytes rekeying
- Selectors for all IP protocols, all ports, between the two networks, using IPv4 subnets

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, 3DES/AES128, sha1, and DH Group 2/5. 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 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 one, the Cisco device sends the message with flag NO_PROPOSAL_CHOSEN and terminates 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 doesn't match, the Cisco device sends a message with flag INVALID_ID_INFORMATION; Phase 1 fails.

Phase 2: Quick Mode Transactions

The following transactions occur in 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>).

Configuring IPSec VPN Service Example

You must configure VPN parameters and then enable the IPSEC service.

Procedure

1. Configure NSX Edge VPN Parameters Example

   You must configure at least one external IP address on NSX Edge to provide IPSec VPN service.

2. Enable IPSec VPN Service Example

   You must enable the IPSec VPN service for traffic to flow from the local subnet to the peer subnet.

Configure NSX Edge VPN Parameters Example

You must configure at least one external IP address on NSX Edge to provide IPSec VPN service.

Procedure

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **NSX Edges**.
3. Double-click an NSX Edge.
4. Click the **Monitor** tab and then click the **VPN** tab.

5. Click **IPSec VPN**.
6. Click the **Add** (✚) icon.
7. Type a name for the IPSec VPN.
8. Type the IP address of the NSX Edge instance in **Local Id**. This will be the peer Id on the remote site.
9. Type the IP address 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.
10. Type the subnets to share between the sites in CIDR format. Use a comma separator to type multiple subnets.
11. Type the Peer Id to uniquely identify the peer site. For peers using certificate authentication, this ID must be the common name in the peer's certificate. For PSK peers, this ID can be any string. VMware recommends that you use the public IP address of the VPN or a FQDN for the VPN service as the peer ID
12. Type the IP address of the peer site in Peer Endpoint. If you leave this blank, NSX Edge waits for the peer device to request a connection.
13. Type the internal IP address of the peer subnet in CIDR format. Use a comma separator to type multiple subnets.
14. Select the Encryption Algorithm.
15. In Authentication Method, select one of the following:

Option	Description
PSK (Pre Shared Key)	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.
Certificate	Indicates that the certificate defined at the global level is to be used for authentication.

16. Type the shared key in if anonymous sites are to connect to the VPN service.
17. Click **Display Shared Key** to display the key on the peer site.
18. In Diffie-Hellman (DH) Group, select the cryptography scheme that will allow the peer site and the NSX Edge to establish a shared secret over an insecure communications channel.
19. Change the MTU threshold if required.
20. Select whether to enable or disable the Perfect Forward Secrecy (PFS) threshold. In IPsec negotiations, Perfect Forward Secrecy (PFS) ensures that each new cryptographic key is unrelated to any previous key.
21. Click **OK**.

    NSX Edge creates a tunnel from the local subnet to the peer subnet.

What to do next

Enable the IPSec VPN service.

Enable IPSec VPN Service Example

You must enable the IPSec VPN service for traffic to flow from the local subnet to the peer subnet.

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **NSX Edges**.

3 Double-click an NSX Edge.

4 Click the **Monitor** tab and then click the **VPN** tab.

5 Click **IPSec VPN**.

6 Click **Enable**.

What to do next

Click **Enable Logging** to log the traffic flow between the local subnet and peer subnet.

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
```

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.

Troubleshooting NSX Edge Configuration Example

Use this information to help you troubleshoot negotiation problems with your setup.

Successful Negotiation (both Phase 1 and Phase 2)

The following examples display a successful negotiating result between NSX Edge and a Cisco device.

NSX Edge

From the NSX Edge command line interface (ipsec auto -status, part of show service ipsec command):

```
000 #2: "s1–c1":500 STATE_QUICK_I2 (sent QI2, IPsec SA established);
      EVENT_SA_REPLACE in 2430s; newest IPSEC; eroute owner; isakmp#1; idle;
      import:admin initiate
000 #2: "s1–c1" esp.f5f6877d@10.20.131.62 esp.7aaf335f@10.20.129.80
      tun.0@10.20.131.62 tun.0@10.20.129.80 ref=0 refhim=4294901761
000 #1: "s1–c1":500 STATE_MAIN_I4 (ISAKMP SA established); EVENT_SA_REPLACE in
      27623s; newest ISAKMP; lastdpd=0s(seq in:0 out:0); idle;
      import:admin initiate
```

Cisco

```
ciscoasa# show crypto isakmp sa detail

    Active SA: 1
    Rekey SA: 0 (A tunnel will report 1 Active and 1 Rekey SA during rekey)
Total IKE SA: 1

    IKE Peer: 10.20.129.80
    Type : L2L        Role    : responder
    Rekey : no        State   : MM_ACTIVE
    Encrypt : 3des    Hash    : SHA
    Auth : preshared  Lifetime: 28800
    Lifetime Remaining: 28379
```

Phase 1 Policy Not Matching

The following lists Phase 1 Policy Not Matching Error logs.

NSX Edge

NSX Edge hangs in STATE_MAIN_I1 state. Look in /var/log/messages for information showing that the peer sent back an IKE message with "NO_PROPOSAL_CHOSEN" set.

```
000 #1: "s1-c1":500 STATE_MAIN_I1 (sent MI1,
      expecting MR1); EVENT_RETRANSMIT in 7s; nodpd; idle;
      import:admin initiate
000 #1: pending Phase 2 for "s1-c1" replacing #0
Aug 26 12:31:25 weiqing-desktop ipsec[6569]:
      | got payload 0x800(ISAKMP_NEXT_N) needed: 0x0 opt: 0x0
Aug 26 12:31:25 weiqing-desktop ipsec[6569]:
      | ***parse ISAKMP Notification Payload:
Aug 26 12:31:25 weiqing-desktop ipsec[6569]:
      |    next payload type: ISAKMP_NEXT_NONE
Aug 26 12:31:25 weiqing-desktop ipsec[6569]: |    length: 96
Aug 26 12:31:25 weiqing-desktop ipsec[6569]:
      |    DOI: ISAKMP_DOI_IPSEC
Aug 26 12:31:25 weiqing-desktop ipsec[6569]: |    protocol ID: 0
Aug 26 12:31:25 weiqing-desktop ipsec[6569]: |    SPI size: 0
Aug 26 12:31:25 weiqing-desktop ipsec[6569]:
      |    Notify Message Type: NO_PROPOSAL_CHOSEN
Aug 26 12:31:25 weiqing-desktop ipsec[6569]:
      "s1-c1" #1: ignoring informational payload,
       type NO_PROPOSAL_CHOSEN msgid=00000000
```

Cisco

If debug crypto is enabled, an error message is printed to show that no proposals were accepted.

```
ciscoasa# Aug 26 18:17:27 [IKEv1]:
      IP = 10.20.129.80, IKE_DECODE RECEIVED
      Message (msgid=0) with payloads : HDR + SA (1)
      + VENDOR (13) + VENDOR (13) + NONE (0) total length : 148
Aug 26 18:17:27 [IKEv1 DEBUG]: IP = 10.20.129.80,
      processing SA payload
Aug 26 18:17:27 [IKEv1]: Phase 1 failure:  Mismatched attribute
      types for class Group Description:  Rcv'd: Group 5
      Cfg'd: Group 2
Aug 26 18:17:27 [IKEv1]: Phase 1 failure:  Mismatched attribute
      types for class Group Description:  Rcv'd: Group 5
      Cfg'd: Group 2
Aug 26 18:17:27 [IKEv1]: IP = 10.20.129.80, IKE_DECODE SENDING
      Message (msgid=0) with payloads : HDR + NOTIFY (11)
      + NONE (0) total length : 124
Aug 26 18:17:27 [IKEv1 DEBUG]: IP = 10.20.129.80,
      All SA proposals found unacceptable
Aug 26 18:17:27 [IKEv1]: IP = 10.20.129.80, Error processing
      payload: Payload ID: 1
Aug 26 18:17:27 [IKEv1 DEBUG]: IP = 10.20.129.80, IKE MM Responder
      FSM error history (struct &0xd8355a60)  <state>, <event>:
      MM_DONE, EV_ERROR-->MM_START, EV_RCV_MSG-->MM_START,
      EV_START_MM-->MM_START, EV_START_MM-->MM_START,
      EV_START_MM-->MM_START, EV_START_MM-->MM_START,
      EV_START_MM-->MM_START, EV_START_MM
Aug 26 18:17:27 [IKEv1 DEBUG]: IP = 10.20.129.80, IKE SA
      MM:9e0e4511 terminating:  flags 0x01000002, refcnt 0,
      tuncnt 0
Aug 26 18:17:27 [IKEv1 DEBUG]: IP = 10.20.129.80, sending
      delete/delete with reason message
```

Phase 2 Not Matching

The following lists Phase 2 Policy Not Matching Error logs.

NSX Edge

NSX Edge hangs at STATE_QUICK_I1. A log message shows that the peer sent a NO_PROPOSAL_CHOSEN message.

```
000 #2: "s1-c1":500 STATE_QUICK_I1 (sent QI1, expecting
      QR1); EVENT_RETRANSMIT in 11s; lastdpd=-1s(seq in:0 out:0);
      idle; import:admin initiate
Aug 26 12:33:54 weiqing-desktop ipsec[6933]: | got payload
      0x800(ISAKMP_NEXT_N) needed: 0x0 opt: 0x0
Aug 26 12:33:54 weiqing-desktop ipsec[6933]: | ***parse
      ISAKMP Notification Payload:
Aug 26 12:33:54 weiqing-desktop ipsec[6933]: |    next payload
      type: ISAKMP_NEXT_NONE
Aug 26 12:33:54 weiqing-desktop ipsec[6933]: |    length: 32
Aug 26 12:33:54 weiqing-desktop ipsec[6933]:
      |    DOI: ISAKMP_DOI_IPSEC
Aug 26 12:33:54 weiqing-desktop ipsec[6933]: |    protocol ID: 3
Aug 26 12:33:54 weiqing-desktop ipsec[6933]: |    SPI size: 16
Aug 26 12:33:54 weiqing-desktop ipsec[6933]: |    Notify Message
      Type: NO_PROPOSAL_CHOSEN
Aug 26 12:33:54 weiqing-desktop ipsec[6933]: "s1-c1" #3:
      ignoring informational payload, type NO_PROPOSAL_CHOSEN
      msgid=00000000
```

Cisco

Debug message show that Phase 1 is completed, but Phase 2 failed because of policy negotiation failure.

```
Aug 26 16:03:49 [IKEv1]: Group = 10.20.129.80,
      IP = 10.20.129.80, PHASE 1 COMPLETED
Aug 26 16:03:49 [IKEv1]: IP = 10.20.129.80, Keep-alive type
      for this connection: DPD
Aug 26 16:03:49 [IKEv1 DEBUG]: Group = 10.20.129.80,
      IP = 10.20.129.80, Starting P1 rekey timer: 21600 seconds
Aug 26 16:03:49 [IKEv1]: IP = 10.20.129.80, IKE_DECODE RECEIVED
      Message (msgid=b2cdcb13) with payloads : HDR + HASH (8)
      + SA (1) + NONCE (10) + KE (4) + ID (5) + ID (5) + NONE (0)
      total length : 288
.
```

```
.
.
Aug 26 16:03:49 [IKEv1]: Group = 10.20.129.80, IP = 10.20.129.80,
      Session is being torn down. Reason: Phase 2 Mismatch
```

PFS Mismatch

The following lists PFS Mismatch Error logs.

NSX Edge

PFS is negotiated as part of Phase 2. If PFS does not match, the behavior is similar to the failure case described in Phase 2 Not Matching.

```
000 #4: "s1-c1":500 STATE_QUICK_I1 (sent QI1, expecting
      QR1); EVENT_RETRANSMIT in 8s; lastdpd=-1s(seq in:0 out:0);
      idle; import:admin initiate
Aug 26 12:35:52 weiqing-desktop ipsec[7312]: | got payload 0x800
      (ISAKMP_NEXT_N) needed: 0x0 opt: 0x0
Aug 26 12:35:52 weiqing-desktop ipsec[7312]:
      | ***parse ISAKMP Notification Payload:
Aug 26 12:35:52 weiqing-desktop ipsec[7312]: |    next payload
      type: ISAKMP_NEXT_NONE
Aug 26 12:35:52 weiqing-desktop ipsec[7312]: |    length: 32
Aug 26 12:35:52 weiqing-desktop ipsec[7312]:
      |    DOI: ISAKMP_DOI_IPSEC
Aug 26 12:35:52 weiqing-desktop ipsec[7312]: |    protocol ID: 3
Aug 26 12:35:52 weiqing-desktop ipsec[7312]: |    SPI size: 16
Aug 26 12:35:52 weiqing-desktop ipsec[7312]: |    Notify Message
      Type: NO_PROPOSAL_CHOSEN
Aug 26 12:35:52 weiqing-desktop ipsec[7312]: "s1-c1" #1: ignoring
      informational payload, type NO_PROPOSAL_CHOSEN
      msgid=00000000
Aug 26 12:35:52 weiqing-desktop ipsec[7312]: | info:  fa 16 b3 e5
      91 a9 b0 02  a3 30 e1 d9  6e 5a 13 d4
Aug 26 12:35:52 weiqing-desktop ipsec[7312]: | info:  93 e5 e4 d7
Aug 26 12:35:52 weiqing-desktop ipsec[7312]:
      | processing informational NO_PROPOSAL_CHOSEN (14)
```

Cisco

```
<BS>Aug 26 19:00:26 [IKEv1 DEBUG]: Group = 10.20.129.80,
      IP = 10.20.129.80, sending delete/delete with
      reason message
Aug 26 19:00:26 [IKEv1 DEBUG]: Group = 10.20.129.80,
      IP = 10.20.129.80, constructing blank hash payload
Aug 26 19:00:26 [IKEv1 DEBUG]: Group = 10.20.129.80,
      IP = 10.20.129.80, constructing blank hash payload
Aug 26 19:00:26 [IKEv1 DEBUG]: Group = 10.20.129.80,
      IP = 10.20.129.80, constructing IKE delete payload
Aug 26 19:00:26 [IKEv1 DEBUG]: Group = 10.20.129.80,
      IP = 10.20.129.80, constructing qm hash payload
Aug 26 19:00:26 [IKEv1]: IP = 10.20.129.80, IKE_DECODE SENDING
      Message (msgid=19eb1e59) with payloads : HDR + HASH (8)
      + DELETE (12) + NONE (0) total length : 80
Aug 26 19:00:26 [IKEv1]: Group = 10.20.129.80, IP = 10.20.129.80,
      Session is being torn down. Reason: Phase 2 Mismatch
```

PSK not Matching

The following lists PSK Not Matching Error logs

NSX Edge

PSK is negotiated in the last round of Phase 1. If PSK negotiation fails, NSX Edge state is STATE_MAIN_I4. The peer sends a message containing INVALID_ID_INFORMATION.

```
Aug 26 11:55:55 weiqing-desktop ipsec[3855]:
      "s1-c1" #1: transition from state STATE_MAIN_I3 to
      state STATE_MAIN_I4
Aug 26 11:55:55 weiqing-desktop ipsec[3855]: "s1-c1" #1:
      STATE_MAIN_I4: ISAKMP SA established
      {auth=OAKLEY_PRESHARED_KEY
      cipher=oakley_3des_cbc_192 prf=oakley_sha group=modp1024}
Aug 26 11:55:55 weiqing-desktop ipsec[3855]: "s1-c1" #1: Dead Peer
      Detection (RFC 3706): enabled
Aug 26 11:55:55 weiqing-desktop ipsec[3855]: "s1-c1" #2:
      initiating Quick Mode PSK+ENCRYPT+TUNNEL+PFS+UP+SAREFTRACK
      {using isakmp#1 msgid:e8add10e proposal=3DES(3)_192-SHA1(2)_160
      pfsgroup=OAKLEY_GROUP_MODP1024}
Aug 26 11:55:55 weiqing-desktop ipsec[3855]: "s1-c1" #1:
      ignoring informational payload, type INVALID_ID_INFORMATION
      msgid=00000000
```

Cisco

```
Aug 26 15:27:07 [IKEv1]: IP = 10.115.199.191,
      IKE_DECODE SENDING Message (msgid=0) with payloads : HDR
      + KE (4) + NONCE (10) + VENDOR (13) + VENDOR (13)
      + VENDOR (13) + VENDOR (13) + NAT-D (130) + NAT-D (130)
      + NONE (0) total length : 304
Aug 26 15:27:07 [IKEv1]: Group = 10.115.199.191,
      IP = 10.115.199.191, Received encrypted Oakley Main Mode
      packet with invalid payloads, MessID = 0
Aug 26 15:27:07 [IKEv1]: IP = 10.115.199.191, IKE_DECODE SENDING
      Message (msgid=0) with payloads : HDR + NOTIFY (11)
      + NONE (0) total length : 80
```

```
Aug 26 15:27:07 [IKEv1]: Group = 10.115.199.191,
      IP = 10.115.199.191, ERROR, had problems decrypting
      packet, probably due to mismatched pre-shared key.
      Aborting
```

Packet Capture for a Successful Negotiation

The following lists a packet capture session for a successful negotiation between NSX Edge and a Cisco device.

```
No.     Time        Source        Destination   Protocol Info
9203    768.394800  10.20.129.80  10.20.131.62  ISAKMP  Identity Protection
                                                           (Main Mode)
Frame 9203 (190 bytes on wire, 190 bytes captured)
Ethernet II, Src: Vmware_9d:2c:dd (00:50:56:9d:2c:dd),
      Dst: Cisco_80:70:f5 (00:13:c4:80:70:f5)
Internet Protocol, Src: 10.20.129.80 (10.20.129.80),
      Dst: 10.20.131.62 (10.20.131.62)
User Datagram Protocol, Src Port: isakmp (500), Dst Port: isakmp (500)
Internet Security Association and Key Management Protocol
  Initiator cookie: 92585D2D797E9C52
  Responder cookie: 0000000000000000
  Next payload: Security Association (1)
  Version: 1.0
  Exchange type: Identity Protection (Main Mode) (2)
  Flags: 0x00
  Message ID: 0x00000000
  Length: 148
  Security Association payload
    Next payload: Vendor ID (13)
    Payload length: 84
    Domain of interpretation: IPSEC (1)
    Situation: IDENTITY (1)
    Proposal payload # 0
      Next payload: NONE (0)
      Payload length: 72
      Proposal number: 0
      Protocol ID: ISAKMP (1)
      SPI Size: 0
      Proposal transforms: 2
      Transform payload # 0
         Next payload: Transform (3)
         Payload length: 32
         Transform number: 0
         Transform ID: KEY_IKE (1)
         Life-Type (11): Seconds (1)
         Life-Duration (12): Duration-Value (28800)
         Encryption-Algorithm (1): 3DES-CBC (5)
         Hash-Algorithm (2): SHA (2)
         Authentication-Method (3): PSK (1)
         Group-Description (4): 1536 bit MODP group (5)
      Transform payload # 1
         Next payload: NONE (0)
         Payload length: 32
         Transform number: 1
         Transform ID: KEY_IKE (1)
         Life-Type (11): Seconds (1)
         Life-Duration (12): Duration-Value (28800)
         Encryption-Algorithm (1): 3DES-CBC (5)
         Hash-Algorithm (2): SHA (2)
         Authentication-Method (3): PSK (1)
         Group-Description (4): Alternate 1024-bit MODP group (2)
  Vendor ID: 4F456C6A405D72544D42754D
    Next payload: Vendor ID (13)
    Payload length: 16
    Vendor ID: 4F456C6A405D72544D42754D
  Vendor ID: RFC 3706 Detecting Dead IKE Peers (DPD)
    Next payload: NONE (0)
    Payload length: 20
    Vendor ID: RFC 3706 Detecting Dead IKE Peers (DPD)

No.     Time        Source        Destination   Protocol Info
9204    768.395550  10.20.131.62  10.20.129.80  ISAKMP Identity Protection
                                                        (Main Mode)

Frame 9204 (146 bytes on wire, 146 bytes captured)
Ethernet II, Src: Cisco_80:70:f5 (00:13:c4:80:70:f5),
      Dst: Vmware_9d:2c:dd (00:50:56:9d:2c:dd)
Internet Protocol, Src: 10.20.131.62 (10.20.131.62),
      Dst: 10.20.129.80 (10.20.129.80)
User Datagram Protocol, Src Port: isakmp (500), Dst Port: isakmp (500)
Internet Security Association and Key Management Protocol
  Initiator cookie: 92585D2D797E9C52
  Responder cookie: 34704CFC8C8DBD09
  Next payload: Security Association (1)
  Version: 1.0
  Exchange type: Identity Protection (Main Mode) (2)
  Flags: 0x00
  Message ID: 0x00000000
  Length: 104
  Security Association payload
    Next payload: Vendor ID (13)
    Payload length: 52
    Domain of interpretation: IPSEC (1)
    Situation: IDENTITY (1)
    Proposal payload # 1
      Next payload: NONE (0)
      Payload length: 40
      Proposal number: 1
      Protocol ID: ISAKMP (1)
      SPI Size: 0
      Proposal transforms: 1
      Transform payload # 1
         Next payload: NONE (0)
         Payload length: 32
         Transform number: 1
         Transform ID: KEY_IKE (1)
         Encryption-Algorithm (1): 3DES-CBC (5)
         Hash-Algorithm (2): SHA (2)
         Group-Description (4): Alternate 1024-bit MODP group (2)
         Authentication-Method (3): PSK (1)
         Life-Type (11): Seconds (1)
         Life-Duration (12): Duration-Value (28800)
  Vendor ID: Microsoft L2TP/IPSec VPN Client
    Next payload: NONE (0)
    Payload length: 24
    Vendor ID: Microsoft L2TP/IPSec VPN Client

No.     Time        Source        Destination   Protocol Info
9205    768.399599  10.20.129.80  10.20.131.62  ISAKMP Identity Protection
                                                        (Main Mode)

Frame 9205 (222 bytes on wire, 222 bytes captured)
Ethernet II, Src: Vmware_9d:2c:dd (00:50:56:9d:2c:dd),
      Dst: Cisco_80:70:f5 (00:13:c4:80:70:f5)
Internet Protocol, Src: 10.20.129.80 (10.20.129.80),
      Dst: 10.20.131.62 (10.20.131.62)
User Datagram Protocol, Src Port: isakmp (500), Dst Port: isakmp (500)
Internet Security Association and Key Management Protocol
  Initiator cookie: 92585D2D797E9C52
  Responder cookie: 34704CFC8C8DBD09
  Next payload: Key Exchange (4)
  Version: 1.0
  Exchange type: Identity Protection (Main Mode) (2)
  Flags: 0x00
  Message ID: 0x00000000
  Length: 180
  Key Exchange payload
    Next payload: Nonce (10)
    Payload length: 132
    Key Exchange Data (128 bytes / 1024 bits)
  Nonce payload
    Next payload: NONE (0)
    Payload length: 20
    Nonce Data


No.     Time        Source        Destination   Protocol Info
9206    768.401192  10.20.131.62  10.20.129.80  ISAKMP Identity Protection
                                                        (Main Mode)
Frame 9206 (298 bytes on wire, 298 bytes captured)
Ethernet II, Src: Cisco_80:70:f5 (00:13:c4:80:70:f5),
      Dst: Vmware_9d:2c:dd (00:50:56:9d:2c:dd)
Internet Protocol, Src: 10.20.131.62 (10.20.131.62),
      Dst: 10.20.129.80 (10.20.129.80)
User Datagram Protocol, Src Port: isakmp (500), Dst Port: isakmp (500)
Internet Security Association and Key Management Protocol
  Initiator cookie: 92585D2D797E9C52
  Responder cookie: 34704CFC8C8DBD09
  Next payload: Key Exchange (4)
  Version: 1.0
  Exchange type: Identity Protection (Main Mode) (2)
  Flags: 0x00
  Message ID: 0x00000000
  Length: 256
  Key Exchange payload
    Next payload: Nonce (10)
    Payload length: 132
    Key Exchange Data (128 bytes / 1024 bits)
  Nonce payload
    Next payload: Vendor ID (13)
    Payload length: 24
    Nonce Data
  Vendor ID: CISCO-UNITY-1.0
    Next payload: Vendor ID (13)
    Payload length: 20
    Vendor ID: CISCO-UNITY-1.0
  Vendor ID: draft-beaulieu-ike-xauth-02.txt
    Next payload: Vendor ID (13)
    Payload length: 12
    Vendor ID: draft-beaulieu-ike-xauth-02.txt
  Vendor ID: C1B7EBE18C8CBD099E89695E2CB16A4A
    Next payload: Vendor ID (13)
    Payload length: 20
    Vendor ID: C1B7EBE18C8CBD099E89695E2CB16A4A
  Vendor ID: CISCO-CONCENTRATOR
    Next payload: NONE (0)
    Payload length: 20
    Vendor ID: CISCO-CONCENTRATOR

No.     Time        Source        Destination   Protocol Info
9207    768.404990  10.20.129.80  10.20.131.62  ISAKMP Identity Protection
                                                            (Main Mode)

Frame 9207 (110 bytes on wire, 110 bytes captured)
Ethernet II, Src: Vmware_9d:2c:dd (00:50:56:9d:2c:dd),
      Dst: Cisco_80:70:f5 (00:13:c4:80:70:f5)
Internet Protocol, Src: 10.20.129.80 (10.20.129.80),
      Dst: 10.20.131.62 (10.20.131.62)
User Datagram Protocol, Src Port: isakmp (500), Dst Port: isakmp (500)
Internet Security Association and Key Management Protocol
  Initiator cookie: 92585D2D797E9C52
  Responder cookie: 34704CFC8C8DBD09
  Next payload: Identification (5)
  Version: 1.0
  Exchange type: Identity Protection (Main Mode) (2)
  Flags: 0x01
  Message ID: 0x00000000
  Length: 68
  Encrypted payload (40 bytes)

No.     Time        Source        Destination   Protocol Info
9208    768.405921  10.20.131.62  10.20.129.80  ISAKMP   Identity Protection
                                                            (Main Mode)
Frame 9208 (126 bytes on wire, 126 bytes captured)
Ethernet II, Src: Cisco_80:70:f5 (00:13:c4:80:70:f5),
      Dst: Vmware_9d:2c:dd (00:50:56:9d:2c:dd)
Internet Protocol, Src: 10.20.131.62 (10.20.131.62),
      Dst: 10.20.129.80 (10.20.129.80)
User Datagram Protocol, Src Port: isakmp (500), Dst Port: isakmp (500)
Internet Security Association and Key Management Protocol
  Initiator cookie: 92585D2D797E9C52
  Responder cookie: 34704CFC8C8DBD09
  Next payload: Identification (5)
  Version: 1.0
  Exchange type: Identity Protection (Main Mode) (2)
  Flags: 0x01
  Message ID: 0x00000000
  Length: 84
  Encrypted payload (56 bytes)

No.     Time        Source        Destination   Protocol Info
9209    768.409799  10.20.129.80  10.20.131.62  ISAKMP   Quick Mode

Frame 9209 (334 bytes on wire, 334 bytes captured)
Ethernet II, Src: Vmware_9d:2c:dd (00:50:56:9d:2c:dd),
      Dst: Cisco_80:70:f5 (00:13:c4:80:70:f5)
Internet Protocol, Src: 10.20.129.80 (10.20.129.80),
      Dst: 10.20.131.62 (10.20.131.62)
User Datagram Protocol, Src Port: isakmp (500), Dst Port: isakmp (500)
Internet Security Association and Key Management Protocol
  Initiator cookie: 92585D2D797E9C52
  Responder cookie: 34704CFC8C8DBD09
  Next payload: Hash (8)
  Version: 1.0
  Exchange type: Quick Mode (32)
  Flags: 0x01
  Message ID: 0x79a63fb1
  Length: 292
  Encrypted payload (264 bytes)

No.     Time        Source        Destination   Protocol Info
9210    768.411797  10.20.131.62  10.20.129.80  ISAKMP   Quick Mode

Frame 9210 (334 bytes on wire, 334 bytes captured)
Ethernet II, Src: Cisco_80:70:f5 (00:13:c4:80:70:f5),
      Dst: Vmware_9d:2c:dd (00:50:56:9d:2c:dd)
Internet Protocol, Src: 10.20.131.62 (10.20.131.62),
      Dst: 10.20.129.80 (10.20.129.80)
User Datagram Protocol, Src Port: isakmp (500), Dst Port: isakmp (500)
Internet Security Association and Key Management Protocol
  Initiator cookie: 92585D2D797E9C52
  Responder cookie: 34704CFC8C8DBD09
  Next payload: Hash (8)
  Version: 1.0
  Exchange type: Quick Mode (32)
  Flags: 0x01
  Message ID: 0x79a63fb1
  Length: 292
  Encrypted payload (264 bytes)

No.     Time        Source        Destination   Protocol Info
9211    768.437057  10.20.129.80  10.20.131.62  ISAKMP   Quick Mode
```

```
Frame 9211 (94 bytes on wire, 94 bytes captured)
Ethernet II, Src: Vmware_9d:2c:dd (00:50:56:9d:2c:dd),
      Dst: Cisco_80:70:f5 (00:13:c4:80:70:f5)
Internet Protocol, Src: 10.20.129.80 (10.20.129.80),
      Dst: 10.20.131.62 (10.20.131.62)
User Datagram Protocol, Src Port: isakmp (500), Dst Port: isakmp (500)
Internet Security Association and Key Management Protocol
  Initiator cookie: 92585D2D797E9C52
  Responder cookie: 34704CFC8C8DBD09
  Next payload: Hash (8)
  Version: 1.0
  Exchange type: Quick Mode (32)
  Flags: 0x01
  Message ID: 0x79a63fb1
  Length: 52
  Encrypted payload (24 bytes)
```