

Table of Contents

Lab Overview - HOL-2126-01-NET - VMware NSX-T - Getting Started	2
Lab Guidance	3
Module 1 - Introduction to VMware NSX (15 minutes)	9
Introduction to NSX Data Center	10
Module 1 Conclusion	26
Module 2 - Host and Edge Transport Node Preparation (30 minutes).....	27
Host and Edge Transport Node Preparation - Module Overview	28
Host and Edge Transport Node Preparation.....	29
Module 2 Conclusion	65
Module 3 - Logical Switching (15 minutes)	66
Logical Switching - Module Overview	67
Logical Switching in NSX	68
Module 3 Conclusion	85
Module 4 - Logical Routing (30 minutes).....	86
Logical Routing - Module Overview	87
Logical Routing of East-West Traffic	88
Logical Routing of North-South Traffic	106
High Availability (HA)	119
Equal Cost Multi-Pathing (ECMP)	132
Module 4 Conclusion	143
Module 5 - Distributed Firewall and Tools (30 minutes).....	144
Distributed Firewall and Tools - Module Overview	145
Distributed Firewall in NSX	147
Tools	178
Module 5 Conclusion	187
Module 6 - NSX Services (15 minutes)	188
NSX Edge Services Overview	189
NSX Edge Transport Nodes and Clusters	192
NSX Edge Services - Load Balancing	197
NSX Edge Services - DHCP	222
Module 6 Conclusion	234
Module 7: NSX Operations Overview (15 minutes)	235
NSX operations overview - Module Overview	236
NSX operations overview	237
Module 7 Conclusion	256

Lab Overview - HOL-2126-01-NET - VMware NSX-T - Getting Started

Lab Guidance

Note: It will take more than 90 minutes to complete this lab. The modules are independent of each other so you can start at the beginning of any module and proceed from there. You can use the Table of Contents to access any module of your choosing.

The Table of Contents can be accessed in the upper right-hand corner of the Lab Manual.

Lab Module List:

- [Module 1 - Introduction to NSX](#) (15 minutes) (Basic) In this module you will be introduced to the NSX Data Center platform and its capabilities. You will also explore the new NSX-T 3.0 features.
- [Module 2 - Host and Edge Transport Node Preparation](#) (30 minutes) (Advanced) In this module, you will review the prepared components ready for NSX workloads.
- [Module 3 - Logical Switching](#) (15 minutes) (Advanced) In this module, you will review the NSX lab fabric, configure Logical Switches, and attach them to virtual machines. You will be introduced to the use of the GENEVE protocol for overlay networking.
- [Module 4 - Logical Routing](#) (30 minutes) (Advanced) In this module you will learn how to configure Logical Routing in NSX for North-South and East-West traffic, configure High Availability (HA), and configure Equal Cost Multipathing (ECMP).
- [Module 5 - Distributed Firewall and Tools](#) (30 minutes) (Advanced) In this module we will demonstrate how the distributed firewall (DFW) and operational tools within NSX function and are configured.
- [Module 6 - NSX Services](#) (15 minutes) (Intermediate) In this module you will explore NSX service as VPN, Load balancing and DHCP / IPAM.
- [Module 7 - NSX operations overview](#) (15 minutes) (Intermediate) In this module you will explore basic topology management, flow tracing and operational functions.

Lab Captains:

- **Joe Collon, Staff NSX Systems Engineer, Americas**
- **Kevin Moats, Staff Technical Account Manager, Americas**
- **Victor Monga, Senior Technical Account Manager, Americas**
- **Pearline Vijayakumar, Solution Engineer, APJ**

This lab manual can be downloaded from the Hands-on Labs Document site found [here](#):

<http://docs.hol.vmware.com>


This lab may be available in other languages. To set your language preference and have a localized manual deployed with your lab, you may utilize this document to help guide you through the process:

<http://docs.hol.vmware.com/announcements/nee-default-language.pdf>

Disclaimer: For over a decade, we have collaborated with Intel® to deliver innovative solutions that enable IT to continually transform their data centers. We have incorporated Intel® product and technology information within this lab to help users understand the benefits of how both hardware and software technology matter when trying to deploy in VMware's ecosystem. We believe that this collaboration will have tremendous benefits for our customers.

Disclaimer: Due to the nature off the Hands on Labs environment you may see vSphere or NSX related alarms. These alarms are due to the lack of resources and can be safely ignored.

Location of the Main Console



1. The area in the RED box contains the Main Console. The Lab Manual is on the tab to the Right of the Main Console.
2. A particular lab may have additional consoles found on separate tabs in the upper left. You will be directed to open another specific console if needed.

3. Your lab starts with 90 minutes on the timer. The lab can not be saved. All your work must be done during the lab session. But you can click the **EXTEND** to increase your time. If you are at a VMware event, you can extend your lab time twice, for up to 30 minutes. Each click gives you an additional 15 minutes. Outside of VMware events, you can extend your lab time up to 9 hours and 30 minutes. Each click gives you an additional hour.

Alternate Methods of Keyboard Data Entry

During this module, you will input text into the Main Console. Besides directly typing it in, there are two very helpful methods of entering data which make it easier to enter complex data.

Click and Drag Lab Manual Content Into Console Active Window

You can also click and drag text and Command Line Interface (CLI) commands directly from the Lab Manual into the active window in the Main Console.


Accessing the Online International Keyboard



You can also use the Online International Keyboard found in the Main Console.

1. Click on the Keyboard Icon found on the Windows Quick Launch Task Bar.


Click once in active console window



In this example, you will use the Online Keyboard to enter the "@" sign used in email addresses. The "@" sign is Shift-2 on US keyboard layouts.

1. Click once in the active console window.
2. Click on the **Shift** key.


Click on the @ key



1. Click on the "@ key".

Notice the @ sign entered in the active console window.

Activation Prompt or Watermark




When you first start your lab, you may notice a watermark on the desktop indicating that Windows is not activated.

One of the major benefits of virtualization is that virtual machines can be moved and run on any platform. The Hands-on Labs utilizes this benefit and we are able to run the labs out of multiple datacenters. However, these datacenters may not have identical processors, which triggers a Microsoft activation check through the Internet.

Rest assured, VMware and the Hands-on Labs are in full compliance with Microsoft licensing requirements. The lab that you are using is a self-contained pod and does not have full access to the Internet, which is required for Windows to verify the activation. Without full access to the Internet, this automated process fails and you see this watermark.

This cosmetic issue has no effect on your lab.

Look at the lower right portion of the screen




Please check to see that your lab is finished all the startup routines and is ready for you to start. If you see anything other than "Ready", please wait a few minutes. If after 5 minutes your lab has not changed to "Ready", please ask for assistance.

Module 1 - Introduction to VMware NSX (15 minutes)

Introduction to NSX Data Center

In this module, you will be introduced to NSX Data Center, its capabilities, and the components that make up the solution.

Lab Topology



NOTE: Only focus on given VMs in this topology. Lab has additional VMs for other modules.

This diagram shows the topology that will be used for this lab. The lab includes four (4) ESXi hosts grouped as follows:

- Cluster: **RegionA01-COMP01**
 - Host: **esx-01a.corp.local**
 - Host: **esx-02a.corp.local**
- Cluster: **RegionA01-MGMT**
 - Host: **esx-03a.corp.local**
 - Host: **esx-04a.corp.local**

In addition, this lab includes a standalone KVM host to demonstrate the multi-hypervisor capabilities provided by NSX:


- KVM host: **kvm-01a.corp.local**

Finally, this lab includes four (4) VM form factor NSX Edge Nodes, configured as follows:

- Edge Cluster: **edge-cluster-01**
 - NSX Edge: **nsx-edge-01**

- NSX Edge: **nsx-edge-02**
- Edge Cluster: **edge-cluster-02**
 - NSX Edge: **nsx-edge-03**
 - NSX Edge: **nsx-edge-04**

NSX Logical Lab Topology



This diagram shows the logical network topology that will be used in this lab. The lab includes a single Tier-0 Gateway that has four (4) total connected interfaces, configured as follows:

- Tier-0 Gateway: **Tier-0-gateway-01**
 - Uplink Interface: **Uplink (192.168.100.3/24)**
 - Interface: **LS-web (172.16.10.1/24)**
 - Interface: **LS-app (172.16.20.1/24)**
 - Interface: **LS-db (172.16.30.1/24)**

What is NSX Data Center?

For additional information: [HERE](#)

VMware NSX Data Center provides an agile software-defined infrastructure to build cloud-native application environments. NSX Data Center is focused on providing networking, security, automation, and operational simplicity for emerging application

frameworks and architectures that have heterogeneous endpoint environments and technology stacks. NSX Data Center supports cloud-native applications, bare metal workloads, multi-hypervisor environments, public clouds, and multiple clouds.


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 perfect complement to NSX are Intel® Xeon® Scalable processors. The 2nd generation Intel® Xeon® Scalable processor incorporates advanced compute cores, a new memory hierarchy, connectivity, and acceleration designed to provide high performance and infrastructure efficiency across a wide range of network-intensive workloads. This processor platform delivers up to 1.58X performance improvement over the previous generation of Intel® Xeon® Scalable processors for network workloads, and supports up to five times more virtual network function (VNF) capacity when complemented with Intel® Quick Assist Technology and the Intel® Ethernet 800 Series Ethernet controllers.

New features introduced in NSX 3.0

- NSX Intelligence 1.1: Enhanced workflows: Lab HOL-2126-02 Module 1 - Security planning with NSX
- AAA: Direct integration with AD and OpenLDAP: This lab HOL-2126-01 Module 1 - AAA - LDAP configuration
- Virtual Routing and Forwarding (VRF): Multi-tenant data plane isolation: Lab HOL-2126-03 Module 1 - Advanced NSX networking concepts
- NSX for vSphere with K8s: Pod networking for K8s: Lab HOL-2126-03 Module 4 - Introduction to Kubernetes - NSX integration
- Distributed IDS: Context-based IDS: Lab HOL-2126-02 Module 5 - Distributed Intrusion Detection (IDS)
- NSX Federation: Centralized configuration of multiple NSX Managers + simplified DR - Not featured in HOL

NSX Components (Part 1)



NSX works by implementing three separate but integrated planes: management, control, and data. The three planes are implemented as a set of processes, modules, and agents residing on three types of nodes: manager, controller, and transport.

- Each NSX installation supports a clustered group of three NSX Manager nodes (virtual machines) that also runs management and control plane processes.
- The NSX Manager cluster hosts API services.

Data Plane

For additional information: [HERE](#)

The data plane performs stateless forwarding/transformation of packets based on tables populated by the control plane, reports topology information to the control plane and maintains packet-level statistics.

This is different from the control plane managed state such as MAC:IP tunnel mappings, because the state managed by the control plane is about how to forward the packets, whereas state managed by the data plane is limited to how to manipulate payload.

 The Data Plane Development Kit (DPDK) is a set of data plane libraries and network interface controller drivers for fast packet processing. Starting with

NSX-T Data Center 2.2, DPDK uses a set of optimizations with the Intel® Xeon® Scalable processor family to help improve the packet processing speed.

Compared to the standard way of packet processing, DPDK helps decrease the processor cost and increase the number of packets processed per second. In NSX-T, it is used in two ways, first as a dedicated general-purpose network appliance called Bare Metal Edge Node and secondly at the host level Enhance Datapath Mode, optimizing host throughput for specific use cases.

Control Plane

Computes all ephemeral runtime state based on the configuration from the management plane, disseminates topology information reported by the data plane elements, and pushes the stateless configuration to forwarding engines.

The set of objects that the control plane deals with includes VIFs, logical networks, logical ports, logical routers, IP addresses, and so on. The control plane is split into two parts in NSX:


- Central Control Plane (CCP) runs on the NSX Manager/Controller cluster. It computes an ephemeral runtime state based on the configuration from the management plane and disseminates information reported by the data plane elements via the local control plane.
- Local Control Plane (LCP) runs on the transport nodes, adjacent to the data plane it controls. It monitors local link status, computes ephemeral runtime state based on updates from data plane and CCP, and pushes the stateless configuration to forwarding engines. The LCP shares fate with the data plane element which hosts it.

Management Plane

The management plane provides a single entry point using either a REST API or the NSX user interface to the system, maintains user configuration, and performs operational tasks on all management, control, and data plane nodes in the system. The management plane is the sole source-of-truth for the configured (logical) system, as managed by the user via configuration.

- Management plane agent (MPA) runs on transport nodes, control nodes, and management nodes. It runs independently of the control and data planes and can be restarted independently. The MPA may also perform data plane related tasks on transport nodes as well.

NSX Components (Part 2)



NSX Manager

For additional information: [HERE](#)

NSX Manager provides the graphical user interface (GUI) and the REST APIs for creating, configuring, and monitoring NSX components, such as controllers, segments, and edge nodes.

NSX Manager is the management plane for the NSX ecosystem. NSX Manager provides an aggregated system view and is the centralized network management component of NSX. It provides a method for monitoring and troubleshooting workloads attached to virtual networks created by NSX. NSX Manager has been integrated into the NSX Controllers in a fully active, clustered configuration. It provides configuration and orchestration of:

- Logical networking components – logical switching and routing
- Networking and Edge services

- Security services and distributed firewall - Edge services and security services can be provided by either built-in components of NSX Manager or by integrated 3rd party vendors.

NSX Controller


NSX Controller is an advanced distributed state management system that controls virtual networks and overlay transport tunnels. The NSX Controller function operates as a separate process within the NSX Manager cluster.

Traffic doesn't pass through the controller; instead, the controller is responsible for providing configuration to other NSX Controller components such as the logical switches, logical routers, and edge configuration. Stability and reliability of data transport are central concerns in networking.

Launch Google Chrome




Open a browser by double clicking the **Google Chrome** icon on the desktop.



If you see this pop-up, click **Cancel**

Open NSX Web Interface



Unlike NSX-V, which uses a vSphere Web Client plugin for administering NSX, NSX-T leverages a separate web interface for administration.


- Open the NSX web interface by clicking on the **nsxmgr-01a** bookmark in the bookmark toolbar of Google Chrome.

Log in to NSX Web Interface

Log in to the NSX-T web interface using the following steps:

1. In the User name field type **nsx-admin@corp.local**
2. In the Password field type **VMware1!**
3. Click **LOG IN** (Note: You may need to click Log In twice)

View Transport Zones and Transport Node



Transport Zones

A transport zone controls which hosts a logical switch can reach. It can span one or more host clusters also known as transport nodes.

If two transport nodes are in the same transport zone, VMs hosted on those transport nodes can be attached to the NSX logical switch segments that are also in that transport zone. If VMs are attached to switches that are in different transport zones, the VMs cannot communicate with each other.

A Transport Zone defines a collection of hosts that can communicate with each other across a physical network infrastructure. VM communication between different hosts within the same TZ happens over one or more interfaces defined as a Tunnel End Point (TEP). VM communication to a physical network happens using logical routers and not TEP.

Verify Transport Zones

Transport Zone	ID	Traffic Type	Transport Node Members	Status
TZ-Overlay-1	6e7c...8b66	Overlay	7	Up
TZ-Uplink	69e5...0b7c	VLAN	2	Up
TZ-VLAN-1	adb5...99a3	VLAN	2	Up
Default nsx-overlay-transportzone	1b3a...963e	Overlay	0	Unknown
Default nsx-vlan-transportzone	a95c...aeba	VLAN	0	Unknown

1. Click the **System** tab on the top of the NSX system user interface
2. In the navigation pane on the left click the arrow to expand **Fabric** options
3. Click **Transport Zones**

Host Transport Node

A node (ESXi, KVM, Bare Metal etc.) can serve as a transport node if it contains at least one hostswitch (NVDS). When creating a host transport node and adding it to a transport zone, NSX installs a hostswitch on the host. The hostswitch is used for attaching VMs to NSX logical switch segments and for creating NSX logical router uplinks and downlinks. It is possible to configure multiple transport zones using the same hostswitch.

Verify NSX Unified Appliance

The screenshot shows the NSX-T user interface. At the top, the 'System' tab is selected and highlighted with a red box and a circled '1'. On the left navigation pane, the 'Appliances' option is highlighted with a red box and a circled '2'. In the main content area, under 'NSX Appliances', a cluster is shown with IP 192.168.110.42, status 'Available', and version 3.0.0.0.0.15946739. Below this, a 'System Load' graph shows a value of 2.33. To the right, 'Memory' is shown as 12 GB out of 15 GB allocated. At the bottom of the cluster details, the 'VIEW DETAILS' link is highlighted with a red box and a circled '3'. A large box on the right contains a cube icon and the text 'ADD NSX APPLIANCE'.

1. Click the **System** tab on the top of the NSX system user interface
2. In the navigation pane on the left click **Appliances**
3. Click on **View Details** to verify operational status of controller and manager

Verify Appliance Operational Status

The screenshot shows the NSX-T System Overview page for the appliance at 192.168.110.42. The left sidebar contains navigation options: Home, Networking, Security, Inventory, Plan & Troubleshoot, and System. The main content area is divided into three sections: Activity Summary, Operational Status, and Appliance Details.

Activity Summary:

- System Load: 5.71 (1 min ago)
- Memory: 10 GB allocated, 15 GB available
- Storage: 21 GB allocated, 308 GB available
- Uptime: 0 days, 0 hours

Operational Status:

Component	Status
CLUSTER_BOOT_MANAGER	UP
DATASTORE	UP
CONTROLLER	UP
MANAGER	UP
POLICY	UP
HTTPS	UP
MONITORING	UP
IDPS_REPORTING	UP
CORFU_NONCONFIG	UP
REPO_SYNC	SUCCESS

Appliance Details:

- Version: 3.0.0.0
- Deployment: 1
- Transport Node: 1
- UUID: [redacted]
- CERT THUMBPRINT: [redacted]

At the bottom, there is a "DELETE APPLIANCE" button.

Verify that both controller and manager operational status is **UP**

Verify Host Transport Nodes

The screenshot shows the NSX-T Host Transport Nodes page. The left sidebar has a "Fabric" section with a "Nodes" link highlighted. The main content area shows a table of Host Transport Nodes managed by vcsa-01a.

Host Transport Nodes Table:

Node	ID	IP Addresses	OS Type	NSX Configuration	NSX Version	Host Switch	Tunnels	TEP IP Address	Node Status	Alarms
RegionA01-M...	MoRef L...								2 Hosts Not C...	
RegionA01-C...	MoRef L...			Applied Pro...					2 Hosts Up	
esx-02a.corp...	50e3...f...	192.168.110...	ESXi 6.7.0	Success	3.0.0.0...	1	2	192.168.130...	Up	0
esx-01a.corp...	8d52...e...	192.168.110...	ESXi 6.7.0	Success	3.0.0.0...	1	2	192.168.130...	Up	0

1. In the navigation pane on the left click the arrow to expand **Fabric** options
2. Click on **Nodes**
3. Click on the Managed by drop down to select **vcsa-01a**
4. Click on the arrow beside **RegionA01-Compute** to expand the node selection

Verify that the following two ESXi hosts are listed as Host Transport Nodes.

- esx-01a.corp.local
- esxi-02a.corp.local

Edge Transport Node

An NSX Edge Transport Node can be a physical or virtual form factor. NSX Edge Node provides routing services and connectivity to networks that are external to the NSX deployment.

When virtual machine workloads residing on different NSX segments communicate with one another through a T1, the distributed router (DR) function is used to route the traffic in a distributed, optimized fashion.

However, when virtual machine workloads need to communicate with devices outside of the NSX environment, the service router (SR), which is hosted on an NSX Edge Node, is used.

Verify Edge Transport Nodes

The screenshot shows the NSX-T management console interface. The 'System' tab is selected, and the 'Edge Transport Nodes' sub-tab is highlighted with a red box and a circled '1'. Below the sub-tab, there is a table with columns: Edge, ID, Deployment, Management, Host, Configuration, NSX Version, N-VD, Tunnels, TEP IP Addr, Edge Cluster, Logical Route, Node Status, and Alarms. Four edge nodes are listed, each with a green status dot and 'Up' in the Node Status column. Red arrows point to the 'Node Status' column for each row.


Edge	ID	Deployment	Management	Host	Configuration	NSX Version	N-VD	Tunnels	TEP IP Addr	Edge Cluster	Logical Route	Node Status	Alarms
nsx-edg...	6537...1...	Virtual ...	192.168...		Succ...	3.0.0.0...	2	↑ 3	192.168...	edge-cl...	1	Up	
nsx-edg...	9b0f...7...	Virtual ...	192.168...		Succ...	3.0.0.0...	2	↑ 1	192.168...	edge-cl...	0	Up	
nsx-edg...	1815...8e...	Virtual ...	192.168...		Succ...	3.0.0.0...	1	Not Ave	192.168...		0	Up	
nsx-edg...	5070...c...	Virtual ...	192.168...		Succ...	3.0.0.0...	1	Not Ave	192.168...		0	Up	

1. Click on **Edge Transport Nodes**

Verify that the following four edge nodes are listed as Edge Transport Nodes.

- nsx-edge-01
- nsx-edge-02
- nsx-edge-03
- nsx-edge-04

NSX Components (Part 3)



NSX Gateways

- **Tier-0 (T0) gateway:** This provides both North-South and East-West connectivity. The north edge of a T0 interfaces with the physical network and is where dynamic routing protocols can be configured. The south edge of T0 connects to one or more T1 routing layer(s) and receives routing information from them.
- **Tier-1 (T1) gateway:** This provides East-West connectivity only and hosts logical switch segment interfaces. For Tier-1 attached subnets to be reachable from the physical network, route redistribution towards the Tier-0 layer must be enabled.


Note that a two-tiered routing topology is not mandatory. A single Tier-0 topology can be implemented. In that case, Layer 2 segments are connected directly to the T0 layer, and a Tier-1 router is not configured.

NSX gateway (T0 or T1) is comprised of up to two components. Each gateway can have one Distributed router (DR), and optionally one or more service routers (SR).

- DR: The DR is kernel based and spans hypervisors, providing local routing functions to those VMs that are connected to it, and also exists in any edge nodes the logical router is bound to.
- SR: It is responsible for delivering services that are not currently implemented in a distributed fashion, such as stateful NAT, load balancing, DHCP or VPN services. Service Routers are deployed on the Edge node cluster that is selected when the T0/T1 router is initially configured.
 - A Tier-0 router has an associated DR and SR created always, even if no stateful services are configured.
 - A Tier-1 router has an associated DR always. An SR is created when it is linked to a Tier-0 router and has services such as NAT, LB, DHCP, or VPN configured.

The Management Plane (MP) allocates a VNI and creates a transit segment, then configures a port on both the SR and DR, connecting them to the transit segment. The MP then automatically allocates unique IP addresses for both the SR and DR.

Verify NSX Gateways



1. Click on **Networking**
2. In the navigation pane on the left click on **Tier-0 Gateways**


Verify the following Tier-0 Gateway exists with a Green Success Status

- Tier-0-gateway-01

AAA - LDAP


One of the new features in NSX 3.0 is LDAP configuration. Now, you can configure LDAP users for NSX unified appliance management.

Verify LDAP Configuration



1. Click the **System** tab on the top of the NSX system user interface
2. In the navigation pane on the left click **Users and Roles**
3. Click on **LDAP**
4. Click on **1** to select LDAP Server

Verify LDAP Connectivity



1. Click on the **arrow** to expand the server configuration
2. Click on **Check Status** to verify LDAP server connectivity

LDAP Servers

Identity Source corp.local #Ldap Servers 1

	Hostname/IP	LDAP Protocol	Port	Connection Status
▼	controlcenter.corp.local	LDAP	389	● Successful 
Use StartTLS ● Enabled				Certificate -----BEGIN CERTIFICATE----- MIIFIDCCBHCgAwIBAgITAAAAAAeF5 cqZgiQs9gAAAAAABzANBgkqhkiG9

- Verify that the connection status is **Successful**

Module 1 Conclusion

Congratulations on completing Module 1! Next we will enabling routing between the different logical switches.

Please proceed to any module below which interests you the most.

- [Module 2 - Host and Edge Transport Node Preparation](#) (30 minutes) (Advanced) In this module, you will review the prepared components ready for NSX workloads.
- [Module 3 - Logical Switching](#) (15 minutes) (Advanced) In this module, you will review the NSX lab fabric, configure Logical Switches, and attach them to virtual machines. You will be introduced to the use of the GENEVE protocol for overlay networking.
- [Module 4 - Logical Routing](#) (30 minutes) (Advanced) In this module you will learn how to configure Logical Routing in NSX for North-South and East-West traffic, configure High Availability (HA), and configure Equal Cost Multipathing (ECMP).
- [Module 5 - Distributed Firewall and Tools](#) (30 minutes) (Advanced) In this module we will demonstrate how the distributed firewall (DFW) and operational tools within NSX function and are configured.
- [Module 6 - NSX Services](#) (15 minutes) (Intermediate) In this module you will explore NSX service as VPN, Load balancing and DHCP / IPAM.
- [Module 7 - NSX operations overview](#) (15 minutes) (Intermediate) In this module you will explore basic topology management, flow tracing and operational functions.

How to End Lab



To end your lab click on the **END** button.

Module 2 - Host and Edge Transport Node Preparation (30 minutes)

Host and Edge Transport Node Preparation - Module Overview

The goal of this lab is to explore host and edge transport node preparation in NSX

- We will review host transport node preparation
 - VMware vCenter
 - KVM
- We will review edge transport node preparation
- We will create a new uplink profile, transport zone, host transport node profile and edge cluster

Host and Edge Transport Node Preparation


In this section we will review the NSX Manager and explore the various components that comprise the fabric. We will then validate that the host and edge transport nodes are connected to the NSX fabric.

Launch Google Chrome



- Open a browser by double clicking the **Google Chrome** icon on the desktop.

Open NSX Web Interface




Unlike NSX-V, which uses a vSphere Web Client plugin for administering NSX, NSX-T leverages a separate web interface for administration.

1. Open the NSX web interface by clicking on the **nsxmgr-01a** bookmark under the RegionA folder in the toolbar of Google Chrome.

***We recommend that you attempt module 1 Introduction to NSX Data Center to navigate through NSX UI.**

Login to the NSX Web Interface




The screenshot shows the VMware® NSX-T™ login page. It features three main components highlighted with red boxes and numbered circles: 1. The username field containing 'nsx-admin@corp.local'. 2. The password field, which is currently empty and masked with dots. 3. The 'LOG IN' button. The interface is clean with a white background and blue accents.

Login to the NSX web interface using the following steps:

1. In the User name field type **nsx-admin@corp.local**
2. In the Password field type **VMware1!**
3. Click **LOG IN**


Navigate to Nodes in the NSX UI



We will now review the state of the NSX fabric.

1. Click **System** tab at the top of the NSX UI
2. Click the arrow to the right of **Fabric** to expand the options
3. Click **Nodes** to view the nodes currently defined in NSX

View Host Transport Node Configuration




The **Nodes** section, under **Fabric**, is where Host and Edge Transport Nodes are defined in the NSX fabric.

Host Transport Nodes are any hypervisor hosts that will participate in NSX overlay functions. NSX includes support for both vSphere/ESXi and KVM hypervisors. This lab is preconfigured with a single, standalone KVM host that is participating in the NSX fabric, as well as a single vCenter server.

Hypervisors can be added individually by selecting **None: Standalone Hosts** from the **Managed by** list, while vCenter servers and their associated inventory can be added by selecting the **vCenter** option.

1. Click the dropdown next to **Managed by**
2. Click **vcsa-01a** from the list of available options

Verify Host Transport Nodes (vCenter)



When viewing **Host Transport Nodes** that are **Managed by vCenter**, the interface displays a list of all clusters in vCenter inventory. In this lab, there are two vCenter clusters: **RegionA01-MGMT** and **RegionA01-COMP01**.

1. Click the arrow to the left of **RegionA01-MGMT (2)** to expand the list of hosts in this cluster
2. Click the arrow to the left of **RegionA01-COMP01 (2)** to expand the list of hosts in this cluster

Verify Host Transport Nodes (vCenter)

Node	ID	IP Addresses	OS Type	NSX Config	NSX Vers	Host Swit	Tunnels	TEP IP Addr	Node Status	Alarms
RegionA01-MGMT (2)										
esx-03a.corp.local	80b4...	192.168.11...	ESXi 6.7.0	Not Conf...		0	Not ...		Not Available	0
esx-04a.corp.local	80b4...	192.168.11...	ESXi 6.7.0	Not Conf...		0	Not ...		Not Available	0
RegionA01-COMP01...										
esx-02a.corp.local	50e3...	192.168.11...	ESXi 6.7.0	Success	3.0.0...	1	↑ 3	192.168.1...	Up	2
esx-01a.corp.local	8d52...	192.168.11...	ESXi 6.7.0	Success	3.0.0...	1	↑ 3	192.168.1...	Up	0

Expanding the list of hosts in both clusters reveals that there are two hosts in each cluster.

- **RegionA01-COMP01** cluster: Hosts **esx-01a.corp.local** and **esx-02a.corp.local**
- **RegionA01-MGMT** cluster: Hosts **esx-03a.corp.local** and **esx-04a.corp.local**

Observe that the hosts in the **RegionA01-COMP01** cluster show an **NSX Configuration** status of **Success**, while the hosts in **RegionA01-MGMT** are **Not Configured**. This means that hosts in the **RegionA01-COMP01** cluster can participate in NSX overlay networking and security, while hosts in **RegionA01-MGMT** cannot.

Verify Edge Transport Nodes

Edge	ID	Deployment	Management	Host	Configuration	NSX Version	N-VD	Tunnels	TEP IP Addr	Edge Cluster	Log	Node Status	Alarms
nsx-edge-01	6537...12...	Virtual M...	192.168.1...		Success	3.0.0.0...	2	↑ 3	192.168.1...	edge-clu...	1	Up	0
nsx-edge-02	9b0f...7...	Virtual M...	192.168.1...		Success	3.0.0.0...	2	↑ 1	192.168.1...	edge-clu...	0	Up	0
nsx-edge-03	1815...8e...	Virtual M...	192.168.1...		Success	3.0.0.0...	1	Not Ava	192.168.1...		0	Up	0
nsx-edge-04	5070...c...	Virtual M...	192.168.1...		Success	3.0.0.0...	1	Not Ava	192.168.1...		0	Up	0

1. Click on the **Edge Transport Nodes** tab

2. You can see that there are 4 edge nodes. **Verify** that the node status for these nodes are all **Up**.

In NSX the Edge Transport Node contains its own TEP within the Edge, and it does not require the hypervisor to perform encapsulation and decapsulation functions on its behalf. When an encapsulated packet is destined for an Edge, it is delivered in its encapsulated form directly to the Edge Node via its TEP address. This allows for greater portability of the Edge Node since it no longer has dependencies on the underlying kernel services of the host.

Do you know why the hosts in the **RegionA01-MGMT** cluster are not configured with NSX kernel modules?

The only VM workloads in the **RegionA01-MGMT** cluster are Edge Transport Nodes. Since these Edge Nodes natively perform their own encapsulation/decapsulation and each has its own TEP, there is no need to configure the hosts in the **RegionA01-MGMT** cluster with NSX kernel modules.

View Uplink Profile Configuration

Uplink Profile	ID	Teaming Policy	Active Uplinks	Standby Uplinks	Transport VLAN	MTU
ESXi-Region01a-COMP01-loadbala...	a2b3...45d7	Load Balance So...	Uplink1,Uplink2		0	1600 (Global MT...
nsx-default-loadbalance-uplink-ho...	0326...dc9f	Load Balance So...	uplink-1,uplink-2,...		0	1600 (Global MT...
nsx-default-uplink-hostswitch-profi...	c352...5f3f	Failover Order	uplink-1	uplink-2	0	1600 (Global MT...
nsx-edge-lag-uplink-profile	ce82...1107	Load Balance So...	uplink-1,uplink-2		0	1600 (Global MT...
nsx-edge-multiple-vteps-uplink-pr...	cf32...e3bc	Failover Order	uplink-1		0	1600 (Global MT...

Uplink Profiles are assigned to Transport Nodes (Host and/or Edge) in the NSX environment and define the configuration of the physical NICs that will be used.

1. On the left side of the NSX user interface, click **Profiles**
2. Click **Uplinks Profiles**
3. Click **nsx-default-uplink-hostswitch-profile**
 - **NOTE:** Click on the name of the Uplink Profile, not the checkbox to its left

Enhanced Network Stack (ENS - also appears as Enhanced Data Path) is a networking stack mode which provides superior network performance when

configured and enabled for optimization of Intel ® NICs. It is primarily utilized in NFV workloads, which require the performance benefits this mode provides.

ENS utilizes the DPDK Poll Mode driver model to significantly improve packet rate and latency for small message sizes.

Verify Uplink Profile Configuration

The screenshot displays the NSX Manager interface for configuring the **nsx-default-uplink-hostswitch-profile**. The left sidebar lists various Uplink Profiles, with **nsx-default-uplink-hostswitch-profile** selected. The main content area shows the configuration details for this profile, including the Summary, LAGs, and Teamings sections. The Summary section is highlighted with a red box, showing the following details:

- Name: nsx-default-uplink-hostswitch-profile
- Description: [blank]
- Transport VLAN: 0
- MTU: 1600 (Global MTU)

The LAGs section shows no LAGs found. The Teamings section is also highlighted with a red box, showing the following details:

Name	ID	LACP Mode	LACP Load Balancing	Num Uplinks	Uplinks	LACP Time
[Default Teaming]		FAILOVER_ORDER		uplink-1	uplink-2	


Observe the following configuration of the **nsx-default-uplink-hostswitch-profile** Uplink Profile:

- Name: **nsx-default-uplink-hostswitch-profile**
- Description: **[blank]**
- Transport VLAN: **0**
- MTU: **1600 (Global MTU)**
- Teaming Policy: **FAILOVER_ORDER**
- Active Uplinks: **uplink-1**
- Standby Uplinks: **uplink-2**

This profile states that two uplinks will be configured in a failover configuration. Traffic will normally utilize uplink-1, and will traverse uplink-2 in the event of a failure of uplink-1.

- 1. Click on 'x' at the top right to close the Uplink Profile summary window

Create a New Uplink Profile



- 1. Click **Add** to create a new Uplink Profile

Create a New Uplink Profile

New Uplink Profile

LAGs

[+ ADD](#) [DELETE](#)

☐ Name*	LACP Mode	LACP Load Balancing*	Uplinks*	LACP Time Out

Teamings

[+ ADD](#) [CLONE](#) [DELETE](#)

☐ Name*	Teaming Policy*	Active Uplinks*	Standby Uplinks
☐ [Default Teaming]	Fallover Order		

Active uplinks and Standby uplinks are user defined labels. These labels will be used to associate with the Physical NICs while adding Transport Nodes.

Transport VLAN

MTU ⓘ

Input the following details in the new Uplink Profile

- Name: **test-uplink-profile**
- Description: **[blank]**
- Transport VLAN: **0**
- MTU: **1600**

Create a New Uplink Profile

New Uplink Profile

No LAGs found

Teamings
[+ ADD](#) [CLONE](#) [DELETE](#)

☒ Name*	Teaming Policy*	Active Uplinks*	Standby Uplinks
☒ [Default Teaming]	Failover Order	uplink-1	 uplink-2

Active uplinks and Standby uplinks are user defined labels. These labels will be used to associate with the Physical NICs while adding Transport Nodes.

Transport VLAN

[CANCEL](#) [ADD](#)

1. Hover over the **Active Uplinks** field under the Teaming configuration and click on the **pencil icon** to edit the Teaming Policy Uplinks with the following parameters

- Active Uplinks: **uplink-1**

New Uplink Profile? ×

Teamings

+ ADD

CLONE

DELETE

☒ Name *	Teaming Policy *	Active Uplinks *	Standby Uplinks
☒ [Default Teaming]	Failover Order	uplink-1	uplink-2

Active uplinks and Standby uplinks are user defined labels. These labels will be used to associate with the Physical NICs while adding Transport Nodes.

Transport VLAN

0

MTU ⓘ

1600

CANCEL

ADD

1. Hover over the **Standby Uplinks** field under the Teaming configuration and click on the **pencil icon** to edit the Teaming Policy Uplinks with the following parameters:

- Standby Uplinks: **uplink-2**

2. Click **Add** to complete the uplink profile creation

View Overlay Transport Zone Configuration

A **Transport Zone** defines the scope of where an NSX segment can exist within the fabric. For example, a dedicated DMZ cluster may contain a DMZ transport zone. Any segments created in this DMZ transport zone could then only be used by VM workloads in the DMZ cluster.

HOL-2126-01-NET

Page 40

There are two types of Transport Zone in NSX, Overlay, and VLAN:

- **Overlay** transport zones are used for NSX Logical Switch segments. Network segments created in an Overlay transport zone will utilize TEPs and Geneve encapsulation, that you will explore in Module 3: Logical Switching.
- **VLAN** transport zones are used for traditional VLAN-backed segments. Network segments created in a VLAN transport zone function similar to a VLAN port group in vSphere.

Second generation Intel® Xeon® Scalable processors (Intel® C620 Series Chipsets) enhance NSX and reduce overhead for near-native I/O performance with SR-IOV. 10/40Gb Intel® Ethernet Network Adapters enable logical networks that allow VMs to communicate across subnets while reducing configuration and management requirements and increasing network responsiveness and flexibility.

The screenshot shows the NSX user interface. On the left sidebar, under the 'Fabric' section, 'Transport Zones' is highlighted with a red box and a circled '1'. The main content area shows the 'Transport Zones' configuration page. At the top, there are tabs for 'Transport Zones' and 'Health Configuration'. Below the tabs, there are buttons for '+ ADD', 'EDIT', 'DELETE', and 'ACTIONS'. A table lists the transport zones:

☐	Transport Zone ↑	ID	Traffic Type	Transport Node Members	Status
☐	TZ-Overlay-1	6e7c...8b66	Overlay	6	Up
☐	TZ-Uplink	69e5...0b7c	VLAN	2	Up
☐	TZ-VLAN-1	adb5...99a3	VLAN	2	Up
☐	Default nsx-overlay-transportzone	1b3a...963e	Overlay	0	Unknown
☐	Default nsx-vlan-transportzone	a95c...aeba	VLAN	0	Unknown

1. On the left side of the NSX user interface, click **Transport Zones**
2. Click **TZ-Overlay-1**
 - **NOTE:** Click on the name of the Transport Zone, not the checkbox to its left

Verify Overlay Transport Zone Configuration

This information is useful for seeing where a given Transport Zone is being used.

Transport Zones Health Configuration

TZ-Overlay-1

Overview Monitor

Total 7 Transport Nodes

Total 5 Switches

Total 16 Switch Ports

Summary EDIT

Name	TZ-Overlay-1
ID	6e7c0d36-b55d-4b38-98ff-c61dfe1d8b66
Location	
Description	
Traffic Type	Overlay
Switch Name	N-VDS-1
Uplink Teaming Policy Names	
Default Transport Zone	No
Logical Ports	16
Logical Switches	5

Observe the following configuration of the **TZ-Overlay** Transport Zone:

- Name: **TZ-Overlay-1**
- Description: **[blank]**
- Traffic Type: **Overlay**
- N-VDS Name: **N-VDS-1**
- Uplink Teaming Policy Names: **[blank]**
- Logical Ports: **16**
- Logical Switches: **5**

1. Click on 'x' at the top right to close the Transport Zone summary window

Create a New Transport Zone

The screenshot shows the VMware vSphere Client interface. The top navigation bar includes tabs for Home, Networking, Security, Inventory, Plan & Troubleshoot, and System. The left sidebar shows a tree view with the following items: System Overview, Configuration, Appliances, Get Started, Fabric (expanded), Nodes, Profiles, and Transport Zones. The main content area is titled 'Transport Zones' and 'Health Configuration'. Below the title, there is a table of transport zones. The table has columns for a checkbox, the name of the transport zone, and its ID. The table contains the following rows:

	Transport Zone ↑	ID
☐	TZ-Overlay-1	6e7c...8b66
☐	TZ-Uplink	69e5...0b7c
☐	TZ-VLAN-1	adb5...99a3
☐	Default nsx-overlay-transportzone	1b3a...963e
☐	Default nsx-vlan-transportzone	a95c...aeba

The 'ADD' button is highlighted with a red box and a circled '1'.

1. Click **Add** to create a new Transport Zone

Continue New Transport Zone Creation

New Transport Zone ⓘ ×

1

Name* TZ-Overlay-Test

Description

Switch Name N-VDS-1

Traffic Type ☒ Overlay ☐ VLAN

Uplink Teaming Policy Names

2

CANCEL ADD


1. Configure the following parameters
 - Name: **TZ-Overlay-Test**
 - Description: **[blank]**
 - Traffic Type: **Overlay**
 - Switch Name: **N-VDS-1**
 - Uplink Teaming Policy Names: **[blank]**
2. Click **Add** to complete the Transport Zone creation

Revisiting Host Transport Node Configuration

Now it's time to review how uplink profiles and transport zones are combined to configure **Host Transport Nodes** in NSX. There are two ways that this can be done:

- **Individual:** An individual host can be configured by choosing the transport zones and uplink profiles that should be associated with the NSX fabric on that host. This configuration is applied on each standalone host in the fabric and is used for KVM and standalone ESXi hosts (not managed by vCenter).
- **Transport Node Profile:** Transport Node Profiles allow you to define the configuration for hosts at a vCenter cluster level. This allows you to apply the configuration to a cluster once, and have all hosts in that cluster automatically receive the appropriate configuration.

Now, lets explore Host Transport Node configuration



1. On the left side of the NSX user interface, click **Profiles**
2. Click **Transport Node Profiles**
3. Click the checkbox to the left of **ESXi-transport-node-profile** to select it
4. Click **Edit** to review the configuration

Verify Host Transport Node Profile

Edit Transport Node Profile - ESXi-transport-n... ? x

Name* ESXi-transport-node-profile

Description

+ ADD SWITCH

▼ N-VDS-1

Type* ☒ N-VDS ☐ VDS

Mode* ☒ Standard (All hosts) ☐ Enhanced Datapath

Name* N-VDS-1

Transport Zone* TZ-Overlay-1 x TZ-VLAN-1 x

OR Create New Transport Zone

NIOC Profile nsx-default-nioc-hostswitch-profile x

OR Create New NIOC Profile

CANCEL SAVE

Observe the following details in the Edit Transport Node Profile dialog:

- Name: **ESXi-transport-node-profile**
- Description: **[blank]**
- Transport Zones (Selected): **TZ-Overlay-1, TZ-VLAN-1**

Verify Host Transport Node Profile N-VDS

OR Create New Transport Zone

NIOC Profile	nsx-default-nioc-hostswitch-profile	OR Create New NIOC Profile
Uplink Profile*	ESXi-Region01a-COMP01-loadbalanced-active-active-profile	OR Create New Uplink Profile
LLDP Profile	LLDP [Send Packet Disabled]	
IP Assignment*	Use IP Pool	
IP Pool*	region-a-tep-pool	

Teaming Policy Switch Mapping

Uplinks	Physical NICs
Uplink1 (active)	vmnic2
Uplink2 (active)	

2 CANCEL SAVE


This profile states that Transport Zones, **TZ-Overlay-1** & **TZ-VLAN-1** will be associated with hosts in this profile. Their connectivity to the physical network will use the **ESXi-Region01a-COMP01-loadbalanced-active-active-profile**. Finally, when a TEP is provisioned on each host, it will assign an IP address from the **region-a-tep-pool** range of IP addresses.

1. **Scroll down** to view the N-VDS-1 configuration
 - Observe the following details in the N-VDS-1 section of the Edit Transport Node Profile dialog:
 - **N-VDS-1**
 - N-VDS Name: **N-VDS-1**
 - NIOC Profile: **nsx-default-nioc-hostswitch-profile**
 - Uplink Profile: **ESXi-Region01a-COMP01-loadbalanced-active-active-profile**
 - LLDP Profile: **LLDP [Send Packet Disabled]**
 - IP Assignment: **Use IP Pool**

- IP Pool: **region-a-tep-pool**
- Physical NICs: **vmnic2** to **uplink-1**

2. Click **CANCEL** to return to the list of Transport Node Profiles

Create a New Host Transport Node Profile



We will just input the values to learn how to create a Transport Node Profile, but we will not save and create it

1. Click **Add** to create a new Transport Node Profile

The screenshot shows the 'Add Transport Node Profile' dialog box. The 'Name' field is filled with 'Test-transport-node-profile'. The 'Description' field is blank. The 'Type' is set to 'N-VDS', 'Mode' is 'Standard (All hosts)', and 'Name' is 'N-VDS-1'. A red box highlights the 'Name', 'Description', and 'New Node Switch' section, with a '1' in a black circle.

- Input the following parameters.
 - Name: **Test-transport-node-profile**
 - Description: **[blank]**
 - N-VDS Name: **N-VDS-1**

Continue New Host Transport Node Profile Creation

Add Transport Node Profile

+ ADD SWITCH

▼ New Node Switch

Type* ☒ N-VDS ☐ VDS

Mode* ☒ Standard (All hosts) ☐ Enhanced Datapath

Name N-VDS-1

Transport Zone* TZ-Overlay-Test  1

OR Create New Transport Zone

NIOC Profile nsx-default-nioc-hostswitch-profile  2

OR Create New NIOC Profile

Uplink Profile* test-uplink-profile  3

OR Create New Uplink Profile

LLDP Profile LLDP [Send Packet Disabled]  4

1. Click the **arrow** beside Transport Zone and select **TZ-Overlay-Test** that you had created earlier
2. Click the **arrow** beside NIOC Profile and select **nsx-default-nioc-hostswitch-profile**
3. Click the **arrow** beside Uplink Profile and select **test-uplink-profile** that you had created earlier
4. Click the **arrow** beside LLDP Profile and select **LLDP [Send Packet Disabled]**

Continue New Host Transport Node Profile Creation

The screenshot shows the 'Add Transport Node Profile' dialog box. It contains the following fields and options:

- Uplink Profile***: test-uplink-profile (with a dropdown arrow and a link 'OR Create New Uplink Profile')
- LLDP Profile**: LLDP [Send Packet Disabled] (with a dropdown arrow)
- IP Assignment***: Use IP Pool (with a dropdown arrow and a checkmark icon, labeled 2)
- IP Pool***: region-a-tep-pool (with a dropdown arrow and a checkmark icon, labeled 3)
- Teaming Policy Switch Mapping**:

Uplinks	Physical NICs
uplink-1 (active)	vmnic2
uplink-2 (standby)	

 (A pencil icon is next to 'uplink-1 (active)', labeled 4)
- PNIC only Migration**: No (toggle switch)
- Network Mappings for Install**: Add Mapping
- Network Mappings for Uninstall**: Add Mapping

At the bottom, there are three buttons: a 'CANCEL' button (labeled 5 and circled in red), an 'ADD' button, and a '5' in a circle next to the 'CANCEL' button.

1. Scroll down to continue the configuration of the new Transport Node Profile
2. Click on the arrow beside IP Assignment and select **Use IP Pool**
3. Click on the arrow beside IP Pool and select **region-a-tep-pool**
4. Click the pencil icon under Physical NICs and key in **vmnic2** for uplink-1 (active)
5. Click **Cancel** to exit the transport node profile creation without saving

Verify Host Transport Nodes

The screenshot shows the NSX-T user interface. On the left sidebar, the 'Nodes' option is highlighted with a red circle and a '1' in a black circle. In the main content area, the 'Managed by' dropdown is set to 'vcsa-01a' with a red circle and a '3' in a black circle. The 'RegionA01-COMP01' cluster is selected with a red circle and a '4' in a black circle. A red arrow points to the 'NSX Configui' column header. The table shows two ESXi hosts: 'esx-02a.corp.local' and 'esx-01a.corp.local', both with a status of 'Up'.

1. On the left side of the NSX user interface, click **Nodes**
2. Click the dropdown next to **Managed by**
3. Click **vcsa-01** from the list of available options
4. Click the **arrow** to the left of **RegionA01-COMP01 (2)** cluster to display the list of nodes

Observe that the RegionA01-COMP01 cluster is configured to use the **ESXi-transport-node-profile** Transport Node Profile. All hosts in this cluster will inherit the configuration that was defined in the profile.

Verify Standalone KVM Transport Node

The screenshot shows the NSX-T user interface. On the left sidebar, the 'Nodes' option is highlighted with a red circle and a '1' in a black circle. In the main content area, the 'Managed by' dropdown is set to 'None: Standalone Hosts' with a red circle and a '2' in a black circle. The 'None: Standalone Hosts' option is highlighted with a red circle and a '1' in a black circle. The table shows one KVM host: 'kvm-01a' with a status of 'Success'.

1. Click the dropdown next to **Managed By**
2. Click **None: Standalone Hosts** from the list of available options

A single, standalone KVM host has been provisioned as part of this lab and has been configured to participate in the NSX fabric.

View Standalone KVM Transport Node Configuration

The screenshot shows the NSX-T vCenter interface. The top navigation bar includes 'Home', 'Networking', 'Security', 'Inventory', 'Plan & Troubleshoot', and 'System'. The 'System' tab is selected, and the 'Host Transport Nodes' sub-tab is active. A table lists the transport nodes, with one node named 'kvm-01a' highlighted. A red circle labeled '1' points to the checkbox next to the node name. Another red circle labeled '2' points to the 'EDIT' button in the actions column. The table has columns for 'Node', 'ID', 'IP Addresses', 'OS Type', and 'Configuration Status'.

Node	ID	IP Addresses	OS Type	Configuration Status
kvm-01a	1966...f465	192.168.110.61	Ubuntu ...	Success

1. Click the checkbox to the left of **kvm-01a.corp.local** to select it
2. Click **Edit** to review the configuration

Verify Standalone KVM Transport Node Configuration

Edit Transport Node - kvm-01a

1 Host Details

2 Configure NSX

Host Details

Name * kvm-01a

Description

IP Addresses * 192.168.110.61

CANCEL NEXT 1

Observe the following details in the Host Details tab of the Edit Transport Node dialog:

- Name: **kvm-01a.corp.local**
- Description: **[blank]**
- IP Addresses: **192.168.110.61**

1. Click **NEXT** to view the Configure NSX settings

Verify Configure NSX Settings for KVM Host Transport Node

Edit Transport Node - kvm-01a

1 Host Details

2 **Configure NSX**

Configure NSX

Name* N-VDS-1

Transport Zone* TZ-Overlay-1

OR Create New Transport Zone

Uplink Profile* KVM-Region01a-single-nic-profile

OR Create New Uplink Profile

LLDP Profile LLDP [Send Packet Disabled]

IP Assignment* Use IP Pool

IP Pool* region-a-tep-pool

Teaming Policy Switch Mapping

Uplinks Physical NICs

Uplink1 (active) eth1

CANCEL 1 **REVIEW** **FINISH**

Observe the following details in the Configure NSX tab of the Edit Transport Node dialog:

- N-VDS Name: **N-VDS-1**
- Associated Transport Zones: **TZ-Overlay-1**
- Uplink Profile: **KVM-Region01a-single-nic-profile**
- LLDP Profile: **LLDP [Send Packet Disabled]**
- IP Assignment: **Use IP Pool**
- IP Pool: **region-a-tep-pool**
- Physical NICs: **eth1** to **uplink-1**

This profile states that a single Transport Zone, **TZ-Overlay-1**, will be associated with the KVM Transport Node host. Its connectivity to the physical network will use the **KVM-Region01a-single-nic-profile**. Finally, when a TEP is provisioned on this host, it will assign an IP address from the **region-a-tep-pool** range of IP addresses.

1. Click **CANCEL** to return to the list of Transport Nodes


Launch Putty Session to SSH into KVM Host



We will now login to host `kvm-01a` and verify that the KVM hypervisor is running the `web-03a.corp.local` virtual machine. This workload has already been added to the NSX inventory, and will be used later in this lab.

1. Click the **PuTTY** icon in the taskbar. This will launch the PuTTY terminal client

SSH Into KVM Host




1. Scroll through the list of Saved Sessions until **kvm-01a** is visible
2. Click **kvm-01a** to highlight it
3. Click **Load** to load the saved session

4. Click **Open** to launch the SSH session

- If prompted, click **Yes** to accept the server's host key
- If not automatically logged in, use username **vmware** and password **VMware1!** to log into kvm-01a

Verify VMs Running on KVM Host



```
kvm-01a.corp.local - PuTTY

System information as of Mon Jul 13 12:49:40 EDT 2020

System load:  0.07          Processes:           139
Usage of /:   22.9% of 20.54GB Users logged in:    0
Memory usage: 20%          IP address for eth0: 192.168.110.61
Swap usage:   0%           IP address for virbr0: 192.168.122.1

* Canonical Livepatch is available for installation.
- Reduce system reboots and improve kernel security. Activate at:
  https://ubuntu.com/livepatch

89 packages can be updated.
25 updates are security updates.

Last login: Wed Jul  1 16:48:40 2020 from 192.168.110.10
root@kvm-01a:~# virsh list
Id      Name              State
-----
1       web-03a           running

root@kvm-01a:~#
```

1. Enter **virsh list** to view the virtual machine workloads currently running on this KVM host and confirm that VM **web-03a** is running

```
virsh list
```

Verify TEP Interface on KVM Host

```

root@kvm-01a:~# ifconfig nsx-vtep0.0
nsx-vtep0.0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1600
    inet 192.168.130.51 netmask 255.255.255.0 broadcast 192.168.130.255
    inet6 fe80::4447:92ff:fed5:b3f3 prefixlen 64 scopeid 0x20<link>
    ether 46:47:92:d5:b3:f3 txqueuelen 1000 (Ethernet)
    RX packets 40400 bytes 3873144 (3.8 MB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 17874 bytes 2036414 (2.0 MB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

root@kvm-01a:~#

```

1. Enter **ifconfig nsx-vtep0.0** into the command-line on kvm-01a to see that the TEP interface has been created with an IP address of 192.168.130.51 and an MTU of 1600.

```
ifconfig nsx-vtep0.0
```

View Edge Transport Node Configuration

The screenshot shows the NSX-T web interface with the 'System' tab selected. Under 'System', the 'Edge Transport Nodes' sub-tab is active. A table lists the configured edge transport nodes. The first node, 'nsx-edge-01', is selected with a checkbox. The table columns include Edge, ID, Deployment, Manager, Host, Configuration, NSX Version, N-VC, Tunnels, TEP IP Addr, Edge Cluster, and Log.

Edge	ID	Deployment	Manager	Host	Configuration	NSX Version	N-VC	Tunnels	TEP IP Addr	Edge Cluster	Log
☒	nsx-edge-01	6537...1...	Virtual ...	192.168...	Success	3.0.0.0...	2	↑ 3	192.168...	edge-cl...	
☐	nsx-edge-02	9b0f...7...	Virtual ...	192.168...	Success	3.0.0.0...	2	↑ 1	192.168...	edge-cl...	
☐	nsx-edge-03	1815...8...	Virtual ...	192.168...	Success	3.0.0.0...	1	Not Av	192.168...		
☐	nsx-edge-04	5070...c...	Virtual ...	192.168...	Success	3.0.0.0...	1	Not Av	192.168...		

Similar to the way a Host Transport Node is configured, an Uplink Profile and one or more Transport Zones are used to define an **Edge Transport Node** in NSX. Edge Transport Nodes perform an important function in the NSX fabric. They host the Service Routers (SRs) that are used by Tier-0 and Tier-1 Gateways to perform stateful services such as NAT or load balancing. Most importantly, they host the Tier-0 Service Router

that provides route peering between NSX overlay networking and the physical routed environment.

In this lab, there are four total Edge Transport Nodes, configured in two fault-tolerant clusters of two nodes each.

1. Click **Edge Transport Nodes**
2. Click the checkbox to the left of **nsx-edge-01** to select it
3. Click **Edit** to review the configuration

Verify Edge Transport Node Configuration

Edit Edge Transport Node - nsx-edge-01

Name* nsx-edge-01

Description

+ ADD SWITCH

▼ N-VDS-1 DELETE

Edge Switch Name* N-VDS-1

Transport Zone* TZ-Overlay-1

OR Create New Transport Zone

Uplink Profile* nsx-edge-single-nic-uplink-profile

OR Create New Uplink Profile

IP Assignment* Use IP Pool

IP Pool* region-a-tep-pool

CANCEL SAVE

Observe the following details in the **General** tab of the Edit Edge Transport Node Profile dialog:

- Name: **nsx-edge-01**
- Description: **[blank]**
- Transport Zones (Selected): **TZ-Overlay-1**

Verify Edge Transport Node N-VDS-1 Configuration

Edit Edge Transport Node - nsx-edge-01

Edge Switch Name* N-VDS-1

Transport Zone* TZ-Overlay-1

Uplink Profile* nsx-edge-single-nic-uplink-profile

IP Assignment* Use IP Pool

IP Pool* region-a-tep-pool

OR Create New Transport Zone

OR Create New Uplink Profile

Teaming Policy Switch Mapping

Uplinks	DPDK Fastpath Interfaces
uplink-1 (active)	Edge-Tep-1-RegionA01 (DI...

N-VDS-2 DELETE

CANCEL SAVE

1. **Scroll down** to view the N-VDS-1 configuration

Observe the following details in the N-VDS-1 tab of the Edit Edge Transport Node dialog:

- **N-VDS-1**
 - N-VDS Name: **N-VDS-1**
 - Associated Transport Zones: **TZ-Overlay-1**
 - Uplink Profile: **nsx-edge-single-nic-uplink-profile**
 - IP Assignment: **Use IP Pool**
 - IP Pool: **region-a-tep-pool**
 - Physical NICs: Uplink-1 to **Edge-TEP-1-RegionA01**

Verify Edge Transport Node N-VDS-2 Configuration

Edit Edge Transport Node - nsx-edge-01

▼ N-VDS-2 DELETE

Edge Switch Name* N-VDS-2

Transport Zone* TZ-Uplink × ▼

Uplink Profile* nsx-edge-single-nic-uplink-profile OR Create New Transport Zone

Teaming Policy Switch Mapping

Uplinks DPDK Fastpath Interfaces

uplink-1 (active) Edge-Uplink-1-RegionA01 ... ×

2 CANCEL SAVE

1. **Scroll down** to view the N-VDS-2 configuration

Observe the following details in the N-VDS-1 tab of the Edit Edge Transport Node dialog:

- **N-VDS-2**
 - N-VDS Name: **N-VDS-2**
 - Associated Transport Zones: **TZ-Uplink**
 - Uplink Profile: **nsx-edge-single-nic-uplink-profile**
 - Physical NICs: **Uplink-1** to **Edge-Uplink-1-RegionA01**

This profile states that this Edge Node will host a two Transport Zone, **TZ-Overlay-1 & TZ-Uplink**. One transport zone will be used for route peering with the physical network (TZ-Uplink), while the other transport zone will be used for overlay network services (TZ-Overlay-1). Their connectivity to the physical network will use the **nsx-edge-single-nic-uplink-profile**.

Finally, when a TEP is provisioned on the TZ-Overlay-1 transport zone, it will assign an IP address from the **region-a-tep-pool** range of IP addresses. No TEP will be provisioned on the VLAN transport zone, so the option is disabled.

2. Click **CANCEL** to return to the list of Edge Transport Nodes

View Edge Cluster

The screenshot shows the NSX-T interface with the 'System' tab selected. Under 'System', the 'Edge Clusters' tab is highlighted with a red box and a circled '1'. The 'EDIT' button is highlighted with a red box and a circled '3'. The 'edge-cluster-01' entry is selected with a checkbox, highlighted with a red box, and a circled '2'.

Edge Cluster	ID	Member Type	Cluster Profile
☒ edge-cluster-01	a97e...91c8	Edge Node	nsx-default-edge-high-availability-p...

As we reviewed, there are four Edge Transport Nodes defined in the NSX fabric. For fault tolerance, edge clusters should be configured in two clusters of two nodes each. In the next task, you will be creating the second edge cluster.

1. Click **Edge Clusters**
2. Click the checkbox to the left of **edge-cluster-01** to select it
3. Click **Edit** to review the configuration

Verify Edge Cluster Configuration

Edit Edge Cluster - edge-cluster-01 ? x

Name*
edge-cluster-01

Description

Edge Cluster Profile
nsx-default-edge-high-availability-profile

Transport Nodes

Member Type Edge Node v

☐ Available (2)

nsx-edge-03

nsx-edge-04

< BACK NEXT > 1 - 2 of 2 records

☐ Selected (2)

nsx-edge-01

nsx-edge-02

1

CANCEL


SAVE

Observe the following details in the Edit Edge Cluster dialog:

- Name: **edge-cluster-01**
- Edge Cluster Profile: **nsx-default-edge-high-availability-profile**
- Transport Nodes (Selected): **nsx-edge-01, nsx-edge-02**

1. Click **CANCEL** to return to the list of Edge Clusters

Create a Second Edge Cluster



The screenshot shows the NSX-T interface with the 'System' tab selected. The 'Edge Clusters' section is active, displaying a table of existing clusters. A red box highlights the '+ ADD' button, and a black circle with the number '1' is placed next to it, indicating the first step in creating a new cluster.

	Edge Cluster	ID	Member Type
☐	edge-cluster-01	a97e...91c8	Edge Node

1. Click **Add** to begin creating a new Edge Cluster

Continue Creation of Second Edge Cluster

The screenshot shows the 'Add Edge Cluster' form with the following elements and callouts:

- 1**: Points to the 'Name' field containing 'edge-cluster-02'.
- 2**: Points to the 'Available (2)' checkbox in the 'Transport Nodes' section.
- 3**: Points to the right-pointing arrow button between the 'Available' and 'Selected' node lists.
- 4**: Points to the 'ADD' button at the bottom right of the form.

The form includes fields for 'Name', 'Description', and 'Edge Cluster Profile' (set to 'nsx-default-edge-high-availability-profile'). The 'Transport Nodes' section shows a list of available nodes ('nsx-edge-03', 'nsx-edge-04') and an empty 'Selected' list.

1. Input the following parameters
 - Name: **edge-cluster-02**
 - Edge Cluster Profile: **nsx-default-edge-high-availability-profile**
2. Select the **checkbox** beside Available transport nodes
3. Click the **arrow** to move the nsx-edge-03 and nsx-edge-04 nodes to the Selected column on the right
4. Click **Add** to complete the edge cluster creation

Module 2 Conclusion

Congratulations on completing Module 2! Next we will enabling routing between the different logical switches.`

Please proceed to any module below which interests you the most.

- [Module 1 - Introduction to NSX](#) (15 minutes) (Basic) In this module you will be introduced to the NSX Data Center platform and its capabilities. You will also explore the new NSX-T 3.0 features.
- [Module 3 - Logical Switching](#) (15 minutes) (Advanced) In this module, you will review the NSX lab fabric, configure Logical Switches, and attach them to virtual machines. You will be introduced to the use of the GENEVE protocol for overlay networking.
- [Module 4 - Logical Routing](#) (30 minutes) (Advanced) In this module you will learn how to configure Logical Routing in NSX for North-South and East-West traffic, configure High Availability (HA), and configure Equal Cost Multipathing (ECMP).
- [Module 5 - Distributed Firewall and Tools](#) (30 minutes) (Advanced) In this module we will demonstrate how the distributed firewall (DFW) and operational tools within NSX function and are configured.
- [Module 6 -NSX Services](#) (15 minutes) (Intermediate) In this module you will explore NSX service as VPN, Load balancing and DHCP / IPAM.
- [Module 7 - NSX operations overview](#) (15 minutes) (Intermediate) In this module you will explore basic topology management, flow tracing and operational functions.

How to End Lab




To end your lab click on the **END** button.

Module 3 - Logical Switching (15 minutes)

Logical Switching - Module Overview

The goal of this lab is to explore Logical Switching in NSX.

- We will review the NSX management domain status
- We will review overlay networking in NSX
- We will create a logical switch Segment, connect it to a VM and NSX Gateway, then verify connectivity



Logical Switching in NSX

Now that you have reviewed the NSX components and verified that everything is operating correctly, we will create a logical network segment and connect an existing workload to it.

Create a new Logical Switch Segment

Segment Name	Connectivity	Transport Zone	Subnets
k8s-mgmt	Tier-0-gateway-01 Tier0	TZ-Overlay-1 Overlay	10.0.1/24
k8s-transport	None	TZ-Overlay-1 Overlay	Not Set
LS-app	Tier-0-gateway-01 Tier0	TZ-Overlay-1 Overlay	172.16.20.1/24
LS-db	Tier-0-gateway-01 Tier0	TZ-Overlay-1 Overlay	172.16.30.1/24
LS-web	Tier-0-gateway-01 Tier0	TZ-Overlay-1 Overlay	172.16.10.1/24
Uplink	None	TZ-Uplink VLAN	Not Set

A Segment is a Layer 2 overlay network that provides an isolated broadcast domain.

From the NSX Manager web interface:

1. Click **Networking** in the top level navigation menu
2. Click **Segments** in the menu on the left side of the NSX Networking user interface
3. Click **ADD SEGMENT**

Note that there are a number of preexisting Segments, including LS-web; LS-app; and LS-db. These are used in other lab modules to host a sample three-tiered application.

Add Segment

SEGMENTS SEGMENT PROFILES EDGE BRIDGE PROFILES METADATA PROXIES

ADD SEGMENT EXPAND ALL Filter by Name, Path and more

Segment Name	Connectivity	Transport Zone	Subnets	Pc
LS-new	Tier-0-gateway-01...	TZ-Overlay-1	172.16.60.1/24 *	

SET DHCP CONFIG

Segment needs to have either Subnets or VPN defined, or both.

L2 VPN You have no L2 VPN sessions for this Gateway. For that, go to [VPN](#)

VLAN Enter List of VLANs

Domain Name Enter Fully Qualifier

VPN Tunnel ID

Uplink Teaming Policy Select Uplink Te

IP Address Pool Select IP Pool

REFRESH 1 - 6 of 6 Segments

Enter the following details in the Add Segment dialog:

1. Segment Name: **LS-new**
2. Connectivity: **Tier-0-gateway-01**
3. Transport Zone: **TZ-OVERLAY-1**
4. Subnets: **172.16.60.1/24**

Save Segment Creation

The screenshot shows the 'SEGMENTS' configuration page. At the top, there are tabs for 'SEGMENTS', 'SEGMENT PROFILES', 'EDGE BRIDGE PROFILES', and 'METADATA PROXIES'. Below the tabs, there is a navigation bar with 'ADD SEGMENT', 'EXPAND ALL', and a filter bar. The main form area contains the following fields:

Segment Name	Connectivity	Transport Zone	Subnets	Pc
Metadata Proxy	Set			
Multicast Routing	☒ Enabled			
Address Bindings	Set	Replication Mode	Hierarchical Two-Ti	
Description	Description			

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


> SEGMENT PROFILES

> DHCP STATISTICS

SAVE CANCEL

1. Scroll down to the bottom
2. Click **Save** to save the Add-Segment configuration


Complete and Verify Segment Creation



1. Click **NO** to complete and exit the ADD SEGMENT dialog

Observe that our new Segment, **LS-new**, is now visible in the list of Segments. During the Segment's creation, we connected it to the T0 router **Tier-0-gateway-01** and assigned IP address **172.16.60.1/24** to this T0 interface.

Attach Application VM to the new Segment




We will now connect to vCenter to attach a sample VM workload to our new Segment.

1. Click on the **new tab** icon
2. Click on the **RegionA** folder

3. Click on **vcsa-01a Web Client**

Login to vCenter



VMware® vSphere

administrator@corp.local


☒ Use Windows session authentication

LOGIN

The image shows the VMware vSphere login interface. It includes a username field with 'administrator@corp.local', a password field with masked characters, a checkbox for 'Use Windows session authentication' which is checked and marked with a '1' in a circle, and a blue 'LOGIN' button marked with a '2' in a circle. Red rectangles highlight the checkbox and the login button.

1. Click the **checkbox** for Use Windows session authentication
2. Click **Login**


Find VM web-04a



Once logged in, navigate to Hosts and Clusters:

1. Click on the **Menu** button
2. Click on **Hosts and Clusters**

Edit Settings for web-04a



1. If the tree view on the left is not already expanded, Click the arrow to the left of **RegionA01-COMP01** to expand it
2. Click **web-04a** to select it, then right click to display its Actions menu
3. Click **Edit Settings**.


Change Network Settings

The screenshot shows the 'Edit Settings' window for a virtual machine named 'web-04a'. The 'Virtual Hardware' tab is active. The 'Network adapter 1' is selected, and the 'mgmt-A01-vDS' is chosen from the dropdown menu. The 'Browse...' button is highlighted with a red box and a blue circle, indicating the next step in the process.

Device	Value	Unit	Connect...
CPU	1		
Memory	2	GB	
Hard disk 1	16	GB	
SCSI controller 0	VMware Paravirtual		
Network adapter 1 *	mgmt-A01-vDS		☒ Connect...
CD/DVD drive 1	mgmt-A01-vDS		☐ Connect...
Video card	Specify custom settings		
VMCI device	Device on the virtual machine PCI bus that provides support for the virtual machine communication interface		
Other	Additional Hardware		

1. Click the selection box to the right of **Network adapter 1**
2. Click **Browse...** to bring up the Select Network dialog

Select Network



1. Click **LS-new** to select it
2. Click **OK**

Complete Edit Settings for web-04a

Edit Settings | web-04a

Virtual Hardware | VM Options

ADD NEW DEVICE

> CPU	1	▼	
> Memory	2	GB	▼
> Hard disk 1	16	GB	▼
> SCSI controller 0	VMware Paravirtual		
> Network adapter 1	LS-new	▼	☒ Connected
> CD/DVD drive 1	Host Device	▼	☐ Connected
> Video card	Specify custom settings ▼		
VMCI device	Device on the virtual machine PCI bus that provides support for the virtual machine communication interface		
> Other	Additional Hardware		


CANCEL OK

1

Confirm that **LS-new** is now selected.


1. Click **OK**

Power on VM web-04a



1. Click the **Power on** icon


View IP Address for VM web-04a



Note the IP address for server web-04a. This is the same IP subnet we used when creating our new Segment earlier in this module.

1. Periodically click the **refresh** button until the IP address for web-04a is visible. This may take one to two minutes

Return to NSX Manager



1. Click to return to the **NSX Manager** tab

Note: If you previously closed NSX Manager or it has timed out, click the **nsxmgr-01a** shortcut under RegionA folder in the toolbar and enter the following to login:

- Username: **nsx-admin@corp.local**
- Password: **VMware1!**

View Segment Configuration: LS-new

Segment Name	Connectivity	Transport Zone	Subnets	Ports	Admin State	Status	Alarms
LS-db	Tier-0-gateway-01 Tier0	TZ-Overlay-1 Overlay	172.16.30.1/24	1	Up	Success	0
LS-new	Tier-0-gateway-01 Tier0	TZ-Overlay-1 Overlay	172.16.60.1/24	1	Up	Success	0

Below the table, the configuration details for the selected segment are shown:

- L2 VPN: Not Set
- VLAN: Not Set
- Domain Name: Not Set
- Edge Bridges: 0
- Multicast Routing: Enabled
- Address Bindings: Not Set
- Description: Not Set
- VPN Tunnel ID: Not Set
- Uplink Teaming Policy: Not Set
- IP Address Pool: Not Set
- Metadata Proxy: 0
- Replication Mode: Hierarchical Two-Tier replication
- Tags: 0

Buttons: ADD SEGMENT, EXPAND ALL, REFRESH, VIEW STATISTICS, VIEW RELATED GROUPS.

We will now view the Segment Ports that are configured for our example VM, allowing it to use network overlay services. In addition to creating the LS-new Segment and configuring web-04a to utilize it, a sample three-tiered app has been configured and included with this lab. Segment LS-web has been preconfigured as its web tier, and servers web-01a; web-02a; and web-03a have been connected to it. We will test network connectivity to VMs on these Segments later in this module.

1. Click **Networking** in the top level navigation menu
2. Click **Segments** in the menu on the left side of the NSX Networking user interface
3. If necessary, scroll through the list of Segments until **LS-new** is visible
4. Click the arrow to the left of **LS-new** to expand it
5. It may take one to two minutes before the correct number of segment ports are displayed. You may need to periodically click the REFRESH button until the interface shows 1 Segment Port
6. Click the **1** link under the Ports column to display the Segment Ports dialog

View Segments Ports: LS-new

Segment Ports

Segment LS-new #Segment Ports 1

X

EXPAND ALL

 Search

	Segment Port Name	Connected To	Admin State	Status
⋮ >	web-04a/web-04a.vmx@50e3704f-950e-49ac-9e94-46d4776cfbd4	ID: 8e839e50-88dd-45a1-9b43-d72323e1b536 Type: View More	● Up	● Success
REFRESH				
1 - 1 of 1 Segment Ports				

1

CLOSE

In the list of displayed ports, verify that web-04a is configured to use Segment LS-new and that its status is Up.

1. Click **CLOSE** to return to the Segment list


Verify Connectivity to Web VMs



Now that we have created LS-new and configured server web-04a to utilize it, we will test connectivity.

1. Click the **Command Prompt** icon in the taskbar to launch a Windows Command Prompt

Ping Test of VM web-04a



```
Administrator: Command Prompt
Microsoft Windows [Version 6.3.9600]
(c) 2013 Microsoft Corporation. All rights reserved.

C:\Users\Administrator>ping 172.16.60.14

Pinging 172.16.60.14 with 32 bytes of data:
Reply from 172.16.60.14: bytes=32 time=13ms TTL=62
Reply from 172.16.60.14: bytes=32 time=2ms TTL=62
Reply from 172.16.60.14: bytes=32 time=1ms TTL=62
Reply from 172.16.60.14: bytes=32 time=3ms TTL=62

Ping statistics for 172.16.60.14:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 1ms, Maximum = 13ms, Average = 4ms

C:\Users\Administrator>
```

In the Command Prompt window, perform a ping test of web-04a by entering the following:


```
ping 172.16.60.14
```

The ping test should return successful replies from web-04a. You can also test pinging the web servers listed below on Segment LS-web.

- web-01a - **172.16.10.11** (ESXi)
- web-02a - **172.16.10.12** (ESXi)
- web-03a - **172.16.10.13** (KVM)

Explore Segment Profiles

A segment profile is a configuration template that can be applied to a segment port or a segment as a whole. When it's applied to a segment, it is applied to all the segment ports of the segment but can still be overridden by a segment port specific configuration.



1. Click **Networking** on the top menu bar
2. Click **Segments** on the side bar
3. Click **Segment Profiles**

Verify default-segment-security-profile

Several default profiles (read-only) for different use-cases are preconfigured.

The following default Segment Profiles are available:

- IP Discovery: ARP, DHCP snooping, VM Tools
- QoS: DSCP (trusted or untrusted), CoS, Bandwidth limitations
- MAC Discovery: MAC Change and MAC Learning
- Spoof Guard: Enable/disable Port Bindings based on IP/MAC
- Segment Security: Will be discussed in Security session.

We will now take a look at the default-segment-security-profile

SEGMENTS

SEGMENT PROFILES

BRIDGE PROFILES

METADATA PROXIES

ADD SEGMENT PROFILE

EXPAND ALL

Filter by Name, Path and more

Segment Profile	Type	Assigned To	Tags	Status
default-ip-discovery-profile	IP Discovery Profile		0	Success
default-mac-discovery-profile	MAC Discovery Profile		0	Success
default-segment-security-profile	Segment Security Profile		0	Success
BPDUI BPDUI Filter BPDUI Filter Allow List Enabled DHCP Server Block Server Block - IPv6 Non-IP Traffic Block Rate Limits Receive Broadcast Receive Multicast Enabled Enabled Disabled Disabled 0 0 Client Block Client Block - IPv6 RA Guard Transmit Broadcast Transmit Multicast Disabled Disabled Disabled 0 0				
default-spoofguard-profile	Spoof Guard Profile		0	Success

1. Click the **SEGMENT PROFILES**.
2. Expand the arrow to view configuration of profile

Module 3 Conclusion

Congratulations on completing Module 3! Next we will enabling routing between the different logical switches.`

Please proceed to any module below which interests you the most.

- [Module 1 - Introduction to NSX](#) (15 minutes) (Basic) In this module you will be introduced to the NSX Data Center platform and its capabilities. You will also explore the new NSX-T 3.0 features.
- [Module 2 - Host and Edge Transport Node Preparation](#) (30 minutes) (Advanced) In this module, you will review the prepared components ready for NSX workloads.
- [Module 4 - Logical Routing](#) (30 minutes) (Advanced) In this module you will learn how to configure Logical Routing in NSX for North-South and East-West traffic, configure High Availability (HA), and configure Equal Cost Multipathing (ECMP).
- [Module 5 - Distributed Firewall and Tools](#) (30 minutes) (Advanced) In this module we will demonstrate how the distributed firewall (DFW) and operational tools within NSX function and are configured.
- [Module 6 - NSX Services](#) (15 minutes) (Intermediate) In this module you will explore NSX service as VPN, Load balancing and DHCP / IPAM.
- [Module 7 - NSX operations overview](#) (15 minutes) (Intermediate) In this module you will explore basic topology management, flow tracing and operational functions.

How to End Lab



To end your lab click on the **END** button.


Module 4 - Logical Routing (30 minutes)

Logical Routing - Module Overview

The goal of this lab is to demonstrate the Logical Routing capabilities of NSX

- We will configure Logical Routing of East-West traffic
- We will configure Logical Routing of North-South traffic
- We will configure High Availability (HA)
- We will configure Equal Cost Multi-Pathing (ECMP)

Logical Routing Topology



The lab environment we are building currently includes a Tier-0 Gateway which connects to outside networks. In this module, we will build a Tier-1 Gateway that will handle routing between a sample three-tiered application's network segments, and move those segments from the existing Tier-0 Gateway to the newly created Tier-1 Gateway.

Logical Routing of East-West Traffic


In this lesson we will explore Logical Routing of East-West traffic in NSX.

Launch Google Chrome



- If not continuing from a previous module, open a browser by double clicking the **Google Chrome** icon on the desktop.

Open NSX Web Interface



- If not already open, launch the NSX web interface by clicking on the **nsxmgr-01a** bookmark in the **Region A** bookmark folder of Google Chrome.

Login to the NSX Web Interface

VMware® NSX-T™

admin 1

..... 2

LOG IN 3

If you do not already have an active session, login to the NSX web interface using the following steps:

1. In the Username field type **admin**
2. In the Password field type **VMware1!VMware1!**
3. Click **LOG IN**

View Existing Tier-0 Gateway

Home Networking 1 Inventory Plan & Troubleshoot System

Tier-0 Gateways

ADD GATEWAY EXPAND ALL Filter by Name, Path and more

Tier-0 Gateway Name	HA Mode	Linked Tier-1 Gateways	Linked Segments	Status	Alarms
Tier-0-gateway-01	Active Active	0	4 3	Success	0

2


NSX includes two tiers of routing: Tier-0 (T0) and Tier-1 (T1). An NSX deployment will typically consist of at least one T0 Gateway that includes Uplink connections to the physical network. This lab has been preconfigured with a three-tiered application spanning three NSX Segments: LS-web, LS-app and LS-db. These Segments have been connected to an existing T0 Gateway.

A Tier-0 Gateway can provide routing to multiple Tier-1 Gateways, allowing multiple isolated tenant environments to exist behind a single Tier-0 Gateway. In this module, we will examine the existing T0 Gateway, then create a new T1 Gateway and migrate the existing three-tiered app Segments over to it.

From the NSX Manager web interface:

1. Click **Networking** in the top level navigation menu
2. Click **Tier-0 Gateways** in the menu on the left side of the NSX Networking user interface
3. Click the **4** link under Linked Segments for T0 Tier-0-gateway-01 to display the Linked Segments dialog. This number may be different, depending on how many previous modules have been completed

View Linked Segments




The list of Linked Segments shows all Segments connected to this T0 Gateway. Observe that LS-web, LS-app and LS-db are currently connected to the existing Tier-0-gateway-01

router. The displayed segments may be different, depending on how many previous modules have been completed.

1. Click **CLOSE** to return to the list of Tier-0 Gateways

Create New Tier-1 Gateway



In this step we will create a new Tier-1 Gateway. We will migrate the existing three-tier app Segments to this Gateway, enabling East-West routing between the app tiers to occur within this new T1 Gateway.

1. Click **Tier-1 Gateways** in the menu on the left side of the NSX Networking user interface
2. Click **ADD TIER-1 GATEWAY**

Add Tier-1 Gateway

Tier-1 Gateways

ADD TIER-1 GATEWAY

COLLAPSE ALL

Filter by Name, Path and more

Tier-1 Gateway Name	Linked Tier-0 Gateway	#Linked Segments	Status	Alarms
1 Tier-1*	Tier-0-gateway-012			
Edge Cluster Select Edge Cluste Edges Set Edges Pool Select Pool Allocation Enable Standby Relocation Tags Tag (R) Scope Max 30 allowed. Click (+) to save. > Route Advertisement NOTE - Before further configurations can be done, fill out mandatory fields above (*), click 'Save' below.				

REFRESH

No Tier-1 Gateways


1. In the Tier-1 Gateway Name field, enter **Tier-1-gateway-01**
2. In the Linked Tier-0 Gateway field, select **Tier-0-gateway-01**

1. Scroll to the bottom of the ADD TIER-1 GATEWAY dialog
2. Click SAVE

Continue Configuring Tier-1 Gateway

1. Click **YES** to continue configuring Tier-1-gateway-01


Configure Route Advertisement



Now that the Tier-1 Gateway has been created, we need NSX to advertise the Segments behind it to the physical network. When LS-web, LS-app and LS-db are directly connected to Tier-0-gateway-01, they are advertised via BGP. Once we move these Segments to our new Tier-1 Gateway, they will no longer be directly connected to the existing T0 and will therefore no longer be advertised. We will configure the new T1 Gateway to advertise its Connected Segments to its T0 gateway, allowing the networks to be advertised via BGP and making them reachable from outside of the NSX environment.

1. Click the arrow to the left of **Route Advertisement** to expand it
2. Click the toggle to the right of **All Connected Segments & Service Ports** to enable it
3. Click **SAVE**

Close Editing Mode



1. Click **CLOSE EDITING** at the bottom of the edit dialog to exit edit mode and return to the list of Tier-1 Gateways

Verify Creation of New Tier-1 Gateway

Tier-1 Gateways					
ADD TIER-1 GATEWAY		EXPAND ALL		Filter by Name, Path and more	
	Tier-1 Gateway Name	Linked Tier-0 Gateway	#Linked Segments	Status	Alarms
>	Tier-1-gateway-01	Tier-0-gateway-01	0	Success	0

Verify that the new T1 Gateway **Tier-1-gateway-01** has been created. Confirm that it is linked to T0 Gateway **Tier-0-gateway-01**, has **0** Linked Segments, and has a Status of **Success**.

Connect Segment LS-web to Tier-1 Gateway

SEGMENTS									
ADD SEGMENT		EXPAND ALL		Filter by Name, Path and more					
	Segment Name	Connectivity	Transport Zone	Subnets	Ports	Admin State	Status	Alarms	
>	k8s-mgmt	Tier-0-gateway-01 Tier0	TZ-Overlay-1 Overlay	10.0.11/24	3	Up	Success	0	
>	k8s-trans...	None	TZ-Overlay-1 Overlay	Not Set	0	Up	Success	0	
>	LS-app	Tier-0-gateway-01 Tier0	TZ-Overlay-1 Overlay	172.16.20.1/24	1	Up	Success	0	
>	LS-db	Tier-0-gateway-01 Tier0	TZ-Overlay-1 Overlay	172.16.30.1/24	1	Up	Success	0	
>	LS-web	Tier-0-gateway-01 Tier0	TZ-Overlay-1 Overlay	172.16.10.1/24	3	Up	Success	0	
>	Uplink	None	TZ-Uplink VLAN	Not Set	0	Up	Success	0	

1. Click **Segments** in the menu on the left side of the NSX Networking user interface
2. Click the **More Options** icon to the left of **LS-web** to display its Options menu, then click **Edit**

Edit Segment LS-web

LS-we **1** Tier-1-... TZ-Overlay 172.16.10.1/24 3

CIDR e.g. 10.22.12.2/23

Gateway CID

CIDR e.g. fc7e:f206:db42::1/48

SET DHCP CONFIG

Segment needs to have either Subnets or VPN defined, or both.

L2 VPN You have no L2 VPN sessions for this Gateway. For that, go to [VPN Services](#). Note that for L2 sessions

VLAN Enter List of VLANs

Domain Name Enter Fully Qualified Domain

Edge Bridges Set

Address Bindings Set

Description Description

VPN Tunnel ID

Uplink Teaming Policy Select Uplink Teaming P

IP Address Pool Select IP Pool

Metadata Proxy Set

Replication Mode Hierarchical Two-Tier replic.

Tags Tag (Re Scope (

Max 30 allowed. Click (+) to save.

2 SAVE CANCEL | Unsaved Changes

1. Select **Tier-1-gateway-01** for Connectivity
2. Click **SAVE**. You may need to scroll to the bottom of the dialog window if the button is not visible

Close Editing Mode

✓ Changes Saved!

> SEGMENT PROFILES

> DHCP STATIC BINDINGS

1 CLOSE EDITING

1. Click **CLOSE EDITING** at the bottom of the edit dialog to exit edit mode and return to the list of Segments

Minimize Segment LS-web View

⋮	>	LS-app	Tier-0-gateway-01 Tier0	TZ-Overlay-1 Overlay	172.16.20.1/24	1	● Up	● Success	0
⋮	>	LS-db	Tier-0-gateway-01 Tier0	TZ-Overlay-1 Overlay	172.16.30.1/24	1	● Up	● Success	0
⋮	▼	LS-web	Tier-1-gateway-01 Tier1	TZ-Overlay-1 Overlay	172.16.10.1/24	3	● Up	● Success	0
		L2 VPN	Not Set		VPN Tunnel ID		Not Set		
		VLAN	Not Set		Uplink Teaming Policy		Not Set		
		Domain Name	Not Set		IP Address Pool		Not Set		
		Edge Bridges	0		Metadata Proxy		0		
								VIEW STATISTICS	
								VIEW RELATED GROUPS	

1. Click the arrow to the left of **LS-web** to minimize it

Connect Segment LS-app to Tier-1 Gateway

<<

Network Overview

Network Topology

Connectivity

Tier-O Gateways

Tier-1 Gateways

Segments

Network Services

VPN

NAT

Load Balancing

Forwarding Policies

SEGMENTS

SEGMENT PROFILES

EDGE BRIDGE PROFILES

METADATA PROXIES

ADD SEGMENT

EXPAND ALL

Filter by Name, Path and more

	Segment Name	Connectivity	Transport Zone	Subnets	Ports	Admin State	Status ⓘ	Alarms
⋮ > ⓘ	k8s-mgmt	Tier-0-gateway-01 Tier0	TZ-Overlay-1 Overlay	10.0.11/24	3	● Up	● Success ⓘ	0 —
⋮ > ⓘ	k8s-transport	None	TZ-Overlay-1 Overlay	Not Set	0	● Up	● Success ⓘ	0 —
1 ⋮ > ⓘ	LS-app	Tier-0-gateway-01 Tier0	TZ-Overlay-1 Overlay	172.16.20.1/24	1	● Up	● Success ⓘ	0 —
⋮ > ⓘ	LS-db	Tier-0-gateway-01 Tier0	TZ-Overlay-1 Overlay	172.16.30.1/24	1	● Up	● Success ⓘ	0 —
⋮ > ⓘ	LS-web	Tier-1-gateway-01 Tier1	TZ-Overlay-1 Overlay	172.16.10.1/24	3	● Up	● Success ⓘ	0 —
⋮ > ⓘ	Uplink	None	TZ-Uplink VLAN	Not Set	0	● Up	● Success ⓘ	0 —

1. Click the **More Options** icon to the left of **LS-app** to display its Options menu, then click **Edit**

Edit Segment LS-app

LS-app **1** Tier-1... TZ-Overlay 172.16.20.1/2 1 ☒

CIDR e.g. 10.22.12.2/23

Gateway CID

CIDR e.g. fc7e:f206:db42::1/48

[SET DHCP CONFIG](#)

Segment needs to have either Subnets or VPN defined, or both.

L2 VPN You have no L2 VPN sessions for this Gateway. For that, go to [VPN Services](#). Note that for L2 sessions

VLAN Enter List of VLANs

Domain Name Enter Fully Qualified Domain

Edge Bridges [Set](#)

Address Bindings [Set](#)

Description Description

VPN Tunnel ID

Uplink Teaming Policy Select Uplink Teaming P

IP Address Pool Select IP Pool

Metadata Proxy [Set](#)

Replication Mode Hierarchical Two-Tier replic.

Tags Tag (Re Scope (

Max 30 allowed. Click (+) to save.

2 [SAVE](#) [CANCEL](#) | Unsaved Changes

1. Select **Tier-1-gateway-01** for Connectivity
2. Click **SAVE**

Close Editing Mode

✓ Changes Saved!

> SEGMENT PROFILES

> DHCP STATIC BINDINGS

1 [CLOSE EDITING](#)

1. Click **CLOSE EDITING** at the bottom of the edit dialog to exit edit mode and return to the list of Segments

Minimize Segment LS-app View

⋮	>	ⓘ	k8s-mgmt	Tier-0-gateway-01 Tier0	TZ-Overlay-1 Overlay	10.0.1/24	3	● Up	● Success	0
⋮	>	ⓘ	k8s-transport	None	TZ-Overlay-1 Overlay	Not Set	0	● Up	● Success	0
⋮	▼	ⓘ	LS-app	Tier-1-gateway-01 Tier1	TZ-Overlay-1 Overlay	172.16.20.1/24	1	● Up	● Success	0
		1	L2 VPN	Not Set		VPN Tunnel ID		Not Set		
			VLAN	Not Set		Uplink Teaming Policy		Not Set		
			Domain Name	Not Set		IP Address Pool		Not Set		
			Edge Bridge	0		Metadata Proxy		0		
								VIEW STATISTICS		
								VIEW RELATED GROUPS		

1. Click the arrow to the left of **LS-app** to minimize it

Connect Segment LS-db to Tier-1 Gateway

Network Overview

Network Topology

Connectivity

Tier-0 Gateways

Tier-1 Gateways

Segments

Network Services

VPN

NAT

Load Balancing

Forwarding Policies

IP Management

SEGMENTS

SEGMENT PROFILES

EDGE BRIDGE PROFILES

METADATA PROXIES

ADD SEGMENT

EXPAND ALL

Filter by Name, Path and more

	Segment Name	Connectivity	Transport Zone	Subnets	Ports	Admin State	Status	Alarms
>	k8s-mgmt	Tier-0-gateway-01 Tier0	TZ-Overlay-1 Overlay	10.0.1/24	3	Up	Success	0
>	k8s-transport	None	TZ-Overlay-1 Overlay	Not Set	0	Up	Success	0
>	LS-app	Tier-1-gateway-01 Tier1	TZ-Overlay-1 Overlay	172.16.20.1/24	1	Up	Success	0
>	LS-db	Tier-0-gateway-01 Tier0	TZ-Overlay-1 Overlay	172.16.30.1/24	1	Up	Success	0
>	LS-web	Tier-1-gateway-01 Tier1	TZ-Overlay-1 Overlay	172.16.10.1/24	3	Up	Success	0
>	Uplink	None	TZ-Uplink VLAN	Not Set	0	Up	Success	0

1. Click the **More Options** icon to the left of **LS-db** to display its Options menu, then click **Edit**

Edit Segment LS-db

1. Select **Tier-1-gateway-01** for Uplink & Type
2. Click **SAVE**

Close Editing Mode

1. Click **CLOSE EDITING** at the bottom of the edit dialog to exit edit mode and return to the list of Segments

Minimize Segment LS-db View


			kubernetes	Tier-0-gateway-01 Tier0	TZ-Overlay-1 Overlay	172.16.1.1/24	0			Success	0
			k8s-transport	None	TZ-Overlay-1 Overlay	Not Set	0			Success	0
			LS-app	Tier-1-gateway-01 Tier1	TZ-Overlay-1 Overlay	172.16.20.1/24	1			Success	0
			LS-db	Tier-1-gateway-01 Tier1	TZ-Overlay-1 Overlay	172.16.30.1/24	1			Success	0
1 L2 VPN Not Set VPN Tunnel ID Not Set VLAN Not Set Uplink Teaming Policy Not Set VIEW STATISTICS VIEW RELATED GROUPS											

1. Click the arrow to the left of **LS-db** to minimize it

Verify Connectivity From Admin Desktop

Now that we have migrated our Segments from the existing Tier-0 Gateway to our new Tier-1 Gateway, we will test connectivity.

Open Command Prompt



1. Click the **Start Menu icon** in the Windows Task Bar
2. Click **Command Prompt**

Ping Test of Segment Gateways

```
Administrator: Command Prompt

C:\Users\Administrator>ping -n 2 172.16.10.1 1

Pinging 172.16.10.1 with 32 bytes of data:
Reply from 172.16.10.1: bytes=32 time=2ms TTL=62
Reply from 172.16.10.1: bytes=32 time=2ms TTL=62

Ping statistics for 172.16.10.1:
    Packets: Sent = 2, Received = 2, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 2ms, Maximum = 2ms, Average = 2ms

C:\Users\Administrator>ping -n 2 172.16.20.1 2

Pinging 172.16.20.1 with 32 bytes of data:
Reply from 172.16.20.1: bytes=32 time=1ms TTL=62
Reply from 172.16.20.1: bytes=32 time=1ms TTL=62

Ping statistics for 172.16.20.1:
    Packets: Sent = 2, Received = 2, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 1ms, Maximum = 1ms, Average = 1ms

C:\Users\Administrator>ping -n 2 172.16.30.1 3

Pinging 172.16.30.1 with 32 bytes of data:
Reply from 172.16.30.1: bytes=32 time=1ms TTL=62
Reply from 172.16.30.1: bytes=32 time=1ms TTL=62

Ping statistics for 172.16.30.1:
    Packets: Sent = 2, Received = 2, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 1ms, Maximum = 1ms, Average = 1ms

C:\Users\Administrator>
```

From the Command Prompt, verify you are able to ping the gateway IP addresses of the Segments we have connected to T1 Gateway **Tier-1-gateway-01**.

1. Ping the gateway of the LS-web Segment:

```
ping -n 2 172.16.10.1
```

2. Ping the gateway of the LS-app Segment:

```
ping -n 2 172.16.20.1
```

3. Ping the gateway of the LS-db Segment:


```
ping -n 2 172.16.30.1
```

Verify East-West Connectivity From web-01a



We will now login to VM web-01a and verify that we can reach the other VMs that comprise our sample three-tiered application.

1. Click the **PuTTY** icon in the taskbar. This will launch the PuTTY terminal client




1. Scroll through the list of Saved Sessions until **web-01a.corp.local** is visible
2. Click **web-01a.corp.local** to highlight it
3. Click **Load** to load the saved session
4. Click **Open** to launch the SSH session



- If prompted, click **Yes** to accept the server's host key
- If not automatically logged in, use username **vmware** and password **VMware1!** to log into web-01a

Ping app-01a and db-01a VMs



```
web-01a.corp.local - PuTTY
Using username "root".
Authenticating with public key "controlcenter" from agent
Last login: Thu May 14 18:35:28 2020 from 192.168.110.10
 01:22:03 up 3 days, 6:58, 0 users, load average: 0.00, 0.00, 0.00
root@web-01a [ ~ ]# ping -c 2 app-01a.corp.local
PING app-01a.corp.local (172.16.20.11) 56(84) bytes of data:
64 bytes from app-01a.corp.local (172.16.20.11): icmp_seq=1 ttl=63 time=1.51 ms
64 bytes from app-01a.corp.local (172.16.20.11): icmp_seq=2 ttl=63 time=0.805 ms

--- app-01a.corp.local ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 3ms
rtt min/avg/max/mdev = 0.805/1.156/1.508/0.353 ms
root@web-01a [ ~ ]# ping -c 2 db-01a.corp.local
PING db-01a.corp.local (172.16.30.11) 56(84) bytes of data:
64 bytes from db-01a.corp.local (172.16.30.11): icmp_seq=1 ttl=63 time=1.49 ms
64 bytes from db-01a.corp.local (172.16.30.11): icmp_seq=2 ttl=63 time=1.05 ms

--- db-01a.corp.local ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 3ms
rtt min/avg/max/mdev = 1.049/1.269/1.490/0.223 ms
root@web-01a [ ~ ]#
```

From the SSH session, verify you are able to ping the IP addresses of the app and db tier VMs that are now routing through **Tier-1-gateway-01**.

1. Ping app-tier VM app-01a:

```
ping -c 2 app-01a.corp.local
```

2. Ping db-tier VM db-01a:

```
ping -c 2 db-01a.corp.local
```

Logical Routing of North-South Traffic

In this lesson we will explore Logical Routing of North-South traffic in NSX.

Review Edge Transport Node Configuration

Edge	ID	Deployment	Managemen	Host	Configuratic	NSX Version	N-VC	Tunnels	TEP IP Addr	Edge Cluste	Logical Rou	Node Status	Alarms
☐	nsx-edg...	6537...1...	Virtual ...	192.168...	Succ...	3.0.0.0...	2	↑4	192.168...	edge-cl...	1	Up	0
☐	nsx-edg...	9b0f...7...	Virtual ...	192.168...	Succ...	3.0.0.0...	2	↑1	192.168...	edge-cl...	0	Up	0
☐	nsx-edg...	1815...8...	Virtual ...	192.168...	Succ...	3.0.0.0...	1	Not Ava	192.168...		0	Up	0
☐	nsx-edg...	5070...c...	Virtual ...	192.168...	Succ...	3.0.0.0...	1	Not Ava	192.168...		0	Up	0

In NSX, the Edge Node provides computational power and North-South routing services to the NSX fabric. Edge Nodes are appliances with pools of capacity that can host routing uplinks as well as non-distributed, stateful services. Edge Nodes can be clustered for scalability and high availability, running in both active-active and active-standby configurations.

We will review the Edge Node and Edge Cluster configurations. We will then review the North-South connectivity provided to the existing Tier-0 Gateway by the Edge Cluster.


1. Click **System** in the top level navigation menu
2. Click **Fabric** in the menu on the left side of the NSX System user interface
3. Click **Nodes**
4. Click **Edge Transport Nodes**

Basic configuration information can be viewed from the list of Edge Nodes, including its TEP address and Edge Cluster configuration. Observe the following:

- Edge Nodes nsx-edge-01 and nsx-edge-02 are members of Edge Cluster **edge-cluster-01**
- Edge Nodes nsx-edge-03 and nsx-edge-04 are not currently members of an Edge Cluster. These Edge Nodes are clustered as part of a separate module
- Edge Node nsx-edge-01 has a single Logical Router instance. The remaining three Edge Nodes do not have any Logical Routers provisioned

Note: You can view any truncated field, such as Edge; Edge Cluster; or TEP IP Addresses by hovering the mouse pointer over the field. A tooltip will appear with the full value.

Transport Zones



Host Transport Nodes Edge Transport Nodes Edge Clusters ESXi Bridge Clusters NCP Clusters													
+ ADD EDGE VM EDIT DELETE ACTIONS													
Edge	ID	Deployment	Manager	Host	Configuratic	NSX Versior	N-VC	Tunnels	TEP IP Addr	Edge Cluste	Logical Rou	Node Statu	Alarms
1 nsx-edg...	6537...1...	Virtual ...	192.168...		Succ...	3.0.0.0...	2	↑ 4	192.168...	edge-cl...	1	Up	0
nsx-edg...	9b0f...7...	Virtual ...	192.168...		Succ...	3.0.0.0...	2	↑ 1	192.168...	edge-cl...	0	Up	0
nsx-edg...	1815...8...	Virtual ...	192.168...		Succ...	3.0.0.0...	1	Not Ave	192.168...		0	Up	0
nsx-edg...	5070...c...	Virtual ...	192.168...		Succ...	3.0.0.0...	1	Not Ave	192.168...		0	Up	0

When used for North-South routing, an NSX Edge Transport Node will be configured with two Transport Zones:

- **VLAN Transport Zone:** Used to connect to the physical network fabric and peer with the physical core router. The Tier-0 Gateway creates an interface in this Transport Zone to pass traffic with the environment outside of NSX
- **Overlay Transport Zone:** Used to connect to the NSX overlay network and communicate with other Edge and Host Transport notes in the NSX fabric. The Tier-0 Gateway creates an interface in this Transport Zone to pass traffic with the overlay environment

We will now review the Transport Zone configuration of the Edge Transport Node.

1. Click nsx-edge-01 to view its configuration

Review Transport Zone Configuration

nsx-edge-01

Overview Monitor Tunnels Related ▾

Summary EDIT

Name	nsx-edge-01
ID	6537759c-6f70-47fa-bc48-a10f6e101232
Location	
Description	
External ID	6537759c-6f70-47fa-bc48-a10f6e101232
Configuration State	Success
Deployment Type	Virtual Machine
Management IP	192.168.110.101
Host	
NSX Version	3.0.0.0.0.15946012
Controller Connectivity	Up
Manager Connectivity	Up
Transport Zones	TZ-Overlay-1 TZ-Uplink
Edge Cluster	edge-cluster-01
Logical Routers	1

Tags MANAGE

Observe that nsx-edge-01 is configured with two Transport Zones: **TZ-Uplink** and **TZ-Overlay-1**. These two Transport Zones are used for the Northbound (VLAN) and Southbound (Overlay) interfaces of the Tier-0 gateway, respectively.

1. Click the Close icon to return to the full list of Edge Transport Nodes

View Tier-0 Gateway

Home Networking Security Inventory Plan & Troubleshoot System

Tier-0 Gateways

ADD GATEWAY ▾

COLLAPSE ALL Filter by Name, Path and more

Tier-0 Gateway Name	HA Mode	Linked Tier-1 Gateways	Linked Segments	Status	Alarms
Tier-0-gateway-01	Active Active	1	1	Success	0

IP Address Management Not Set

Edge Cluster edge-cluster-01

> Additional Settings

> Route Distinguisher for VRF

Gateways

> EVPN Settings

Tags 0

INTERFACES

External and Service Interfaces

VIEW DNS

VIEW NAT

VIEW GATEWAY FIREWALL

We will now review the configuration of the existing Tier-0 (T0) Gateway **Tier-0-gateway-01**. This T0 Gateway router is configured to use the Uplink connections provided by Edge Cluster edge-cluster-01, which is comprised of Edge Nodes nsx-edge-01 and nsx-edge-02.

1. Click **Networking** in the top level navigation menu
2. Click **Tier-0 Gateways** in the menu on the left side of the NSX-T Networking user interface
3. Click the arrow to the left of **Tier-0-gateway-01** to expand it
4. Click the arrow to the left of **INTERFACES** to expand it
5. Click the **1** link to the right of External and Service Interfaces to display the Interfaces dialog

View Tier-0 Gateway Interfaces

Interfaces

Tier-O Gateways Tier-O-gatew... #Interfaces 1

COLLAPSE ALL Search

	Name	Type	IP Address / Mask	Connected To(Segment)	Status
1	Uplink-1	External	192.168.100.3/24	Uplink	Success

Edge Node nsx-edge-01 MTU Not Set

Tags 0 PIM Disabled

ND Profile default URPF Mode Strict

VIEW STATISTICS

2 CLOSE

1. Click the arrow to the left of **Uplink-1** to expand it

Review the **Uplink-1** interface configured on Tier-0-gateway-01. This is the North-South Uplink interface the T0 uses to peer with the external routed environment. Observe the following settings:

- Name: **Uplink-1**
- Type: **External**
- IP Address / Mask: **192.168.100.3/24**
- Connected To(Segment): **Uplink**
- Status: **Success**
- Edge Node: **nsx-edge-01**

From this screen we can determine that the T0 Gateway has a single Uplink interface that uses IP address **192.168.100.3**. This Uplink is hosted on the Edge Node **nsx-edge-01** and is currently **Up**.

2. Click **CLOSE** to return to the list of Tier-0 Gateways

View Tier-0 Gateway BGP Configuration

Tier-0 Gateway Name	HA Mode	Linked Tier-1 Gateways	Linked Segments	Status	Alarms
Tier-0-gateway-01	Active Active	1	1	Success	0

IP Address Management: Not Set

Edge Cluster: edge-cluster-01

> Additional Settings

> Route Distinguisher for VRF Gateways

> EVPN Settings

Tags: 0

INTERFACES

External and Service Interfaces: 1

> ROUTING

> MULTICAST

1 **BGP**

Local AS	65001	Inter SR iBGP	On
BGP	On	ECMP	On
Graceful Restart	Helper Only	Multipath Relax	On
Graceful Restart Timer	180 seconds	Graceful Restart Stop Timer	600 seconds
Route Aggregation	0	BGP Neighbors	2 1

Border Gateway Protocol (BGP) is a communication protocol used by routers to exchange route information. When two or more routers are configured in this way and communicate with one another, they are called neighbors. We will now review Tier-0-gateway-01's BGP configuration.

1. Click the arrow to the left of **BGP** to expand it
2. Click the 1 link to the right of BGP Neighbors to display the BGP Neighbors dialog

Verify Tier-0 Gateway BGP Neighbors

BGP Neighbors

Tier-O Gateways Tier-O-gatew... #Neighbors 1

EXPAND ALL Search

	IP Address	BFD	Remote AS number	Route Filter	Allowas-in	Status
>	192.168.100.1	Disabled	65002	1	Disabled	Success

REFRESH 1 - 1 of 1 BGP Neighbors

1 CLOSE


Tier-0-gateway-01 is configured with one BGP neighbor. Review the settings for the 192.168.100.1 neighbor:

- IP Address: **192.168.100.1**
- BFD: **Disabled**
- Remote AS number: **65002**
- Route Filter: **1**
- Allowas-in: **Disabled**
- Status: **Success**

In this instance, Tier-0-gateway-01 is peering with a router at IP address **192.168.100.1** using BGP AS number **65002**. Its status is currently Up, indicated as **Success**.

1. Click **CLOSE** to return to the list of Tier-0 Gateways

View Network Topology




Introduced with NSX-T 3.0 is a new Network Topology visualization view. This view displays a graphical representation of the NSX environment, including Tier-0 and Tier-1 Gateways, Segments, and their connectivity to one another. As we observed in the previous steps:

- **Tier-0-gateway-01** has an Uplink interface with an IP address of 192.168.100.3
- **Tier-0-gateway-01** is connected to **Tier-1-gateway-01** via an automatically assigned subnet of 100.64.208.0/31
- The Segments **LS-web**, **LS-app** and **LS-db** are connected to, and routed by, **Tier-1-gateway-01**


NOTE: If previous modules in this lab have not been completed, the Network Topology view may differ from the one above.

Open Command Prompt



1. Click the **Start Menu icon** in the Windows Task Bar
2. Click **Command Prompt**

Trace Route Path from Main Console to Virtual Machines




1. Trace route path from your Admin Console to web server web-01a on Segment LS-web

```
tracert -d 172.16.10.11
```

Observe that:

- The first hop is the physical gateway (**192.168.110.1**)
- The physical router then routes the packet to the Uplink interface of Tier-0-gateway-01 (**192.168.100.3**)
- Tier-0-gateway-01 routes the packet to its connected interface of Tier-1-gateway-01. This is the autoconfigured subnet we reviewed above (**100.64.208.1**) and is not advertised to the external network; therefore it will display as "timed out" in the route path
- Finally, the packet is delivered to server web-01a on NSX Segment LS-web (**172.16.10.11**)

Test Connectivity to Sample Web App



We will now test connectivity to our sample three-tiered application.

1. Open a new Tab
2. Click on the **3 Tier App** bookmark
3. Select **Web-01a (Running on esx-01a)**

Web Site is Working

Customer Database Access

Accessed via: web-01a

Name Filter (blank for all records):


Rank	Name	Universe	Revenue
1	CHOAM	Dune	\$1.7 trillion
2	Acme Corp.	Looney Tunes	\$348.7 billion
3	Sirius Cybernetics Corp.	Hitchhiker's Guide	\$327.2 billion
4	Buy n Large	Wall-E	\$291.8 billion
5	Aperture Science, Inc.	Valve	\$163.4 billion
6	SPECTRE	007	\$157.1 billion
7	Very Big Corp. of America	Monty Python	\$146.6 billion
8	Frobozz Magic Co.	Zork	\$112.9 billion
9	Warbucks Industries	Lil' Orphan Annie	\$61.5 billion
10	Tyrell Corp.	Bladerunner	\$59.4 billion
11	Wayne Enterprises	Batman	\$31.3 billion
12	Virtucon	Austin Powers	\$24.9 billion
13	Globex	The Simpsons	\$23.7 billion
14	Umbrella Corp.	Resident Evil	\$22.6 billion
15	Wonka Industries	Charlie and the Chocolate Factory	\$21.0 billion
16	Stark Industries	Iron Man	\$20.3 billion
17	Clampett Oil	Beverly Hillbillies	\$18.1 billion
18	Oceanic Airlines	Lost	\$7.8 billion
19	Brawnido	Idiocracy	\$5.8 billion
20	Cyberdyne Systems Corp.	Terminator	\$5.5 billion
21	Paper Street Soap Company	Fight Club	\$5.0 billion
22	Gringotts	Harry Potter	\$4.4 billion
23	Oscorp	Spider-Man	\$3.1 billion
24	Nakatomi Trading Corp.	Die-Hard	\$2.5 billion
25	Los Pollos Hermanos	Breaking Bad	\$1.3 billion

Verify that access to our three-tiered web app is working.

High Availability (HA)

In this lesson we will configure High Availability (HA) for Tier-0 and Tier-1 Gateways in NSX.

Review Edge Cluster Configuration



1. Click **System** in the top level navigation menu
2. Click **Fabric** in the menu on the left side of the NSX System user interface
3. Click **Nodes**
4. Click **Edge Clusters**
5. Click the **2** link for edge-cluster-01's Edge Transport Nodes

Recall from earlier in this module that Edge Transport Nodes provide computational power and North-South routing services to the NSX fabric. Tier-0 and Tier-1 Gateways can provision stateful services, such as NAT or Load Balancing, that are hosted on an Edge Transport Node.

In the event of a power or hardware failure, the loss of an Edge Node could occur. In this instance, any services hosted on that Edge Node would be lost as well. For this reason, Edge Transport Nodes are grouped into an Edge Cluster in NSX. Edge Clusters provide fault tolerance and resilience that can withstand individual failures within the cluster.

View Edge Cluster edge-cluster-01

Host Transport Nodes Edge Transport Nodes **Edge Clusters** ESXi Bridge Clusters NCP Clusters

edge-cluster-01

Overview **Related**

Transport Nodes

Transport Node ↑	ID
nsx-edge-01	6537...1232
nsx-edge-02	9b0f...7b64

Observe that Edge Cluster edge-cluster-01 consists of two individual Edge Transport Nodes: nsx-edge-01 and nsx-edge-02.

1. Click the **close icon**

Modify Existing Tier-0 Gateway

Home **Networking** 1 Inventory Plan & Troubleshoot System

Tier-0 Gateways

ADD GATEWAY EXPAND ALL Filter by Name, Path and more

Tier-0 Gateway Name	HA Mode	Linked Tier-1 Gateways	Linked Segments	Status	Alarms
Tier-0-gateway-01	Active Active	1	1	Success	0

We will now modify the existing Tier-0 Gateway to leverage NSX's Edge Clustering capabilities.

1. Click **Networking** in the top level navigation menu
2. Click **Tier-0 Gateways** in the menu on the left side of the NSX Networking user interface
3. Click the **More Options** icon to the left of **Tier-0-gateway-01** to display its Options menu, then click **Edit**

Edit Interfaces of Tier-0 Gateway

Tier-0 Gateway Name: Tier-0-gateway-C * HA Mode: Active Active * Linked Tier-1 Gateways: Linked Segments: Status: Alarms:

IP Address Management: No Dynamic IP Allocation

Edge Cluster: edge-cluster-01

> Additional Settings

> Route Distinguisher for VRF

Gateways

> EVPN Settings

Tags: Tag (Re) Scope (t) + Max 30 allowed. Click (+) to save.

1 INTERFACES

External and Service Interfaces 1 2

> ROUTING

> MULTICAST

> BGP

> ROUTE RE-DISTRIBUTION

CLOSE EDITING

Observe that our existing Tier-0 Gateway is configured for an HA Mode of Active Active. This allows the use of multiple Edge Nodes in the Edge Cluster simultaneously. Also note that the Tier-0 Gateway is configured to use Edge Cluster edge-cluster-01.

1. Click the arrow to the left of **INTERFACES** to expand it
2. Click the **1** link to the right of **External and Service Interfaces** to display the Set Interfaces dialog

Add Uplink to Tier-0 Gateway for nsx-edge-02

Set Interfaces

Tier-0 Gateways Tier-O-gatew... #Interfaces 1

ADD INTERFACE 2 COLLAPSE ALL

	Name	Type	IP Address / Mask	Connected To(Segment)	Status
1	Uplink-1	External	192.168.100.3/24	Uplink	Success
	Edge Node	nsx-edge-01	MTU	Not Set	VIEW STATISTICS
	Tags	0	PIM	Disabled	
	ND Profile	default	URPF Mode	Strict	

CLOSE

1. Click the arrow to the left of **Uplink-1** to expand it

Observe that the existing **Uplink-1** interface is running on Edge Node **nsx-edge-01** and is configured for IP address **192.168.100.3**. If a failure were to occur on this Edge Node, North-South connectivity to the NSX environment would be lost. We will now add a second Uplink interface to the Tier-0 Gateway that leverages nsx-edge-02, the second Edge Node in Edge Cluster edge-cluster-01.

2. Click **ADD INTERFACE**

Configure Uplink for Tier-0 Gateway on nsx-edge-02

Name	Type	IP Address / Mask	Connected To(Segment)	Status
Uplink-2	External	192.168.100.4/24	Uplink	
Enter IP Address Masks CIDR e.g. IPv4 10.1/24 or IPv6 fc7e:f206:dbb				
Edge Node *	nsx-edge-02	MTU	Enter MTU Minimum 64	
PIM	☐ Disabled	ND Profile	default	
Tags	Tag (Requ... Scope (Op... Max 30 allowed. Click (+) to save.	URPF Mode	Strict	
SAVE CANCEL				

1. Name: **Uplink-2**
2. Type: **External**
3. IP Address / Mask: **192.168.100.4/24**
 - **Note:** You must press **ENTER** or click the **Add Item(s): 192.168.100.4/24** link once you've entered the IP address
4. Connected To(Segment): **Uplink**
5. Edge Node: **nsx-edge-02**
6. Click **SAVE**

Verify Second Uplink on Tier-0 Gateway

Set Interfaces

Tier-0 Gateways Tier-0-gatew... #Interfaces 2

ADD INTERFACE COLLAPSE ALL Search

Name	Type	IP Address / Mask	Connected To(Segment)	Status
Uplink-1	External	192.168.100.3/24	Uplink	Success
Uplink-2				
Edge Node	nsx-edge-02	MTU	Not Set	
Tags	0	PIM	Disabled	
ND Profile	default	URPF Mode	Strict	

1. Click the arrow to the left of **Uplink-2** to expand it

2. Click **CLOSE**

1. Click the arrow to the left of **Uplink-2** to expand it

Confirm that our Tier-0 Gateway now has two interfaces: **Uplink-1** and **Uplink-2**. Interface **Uplink-2** exists on Edge Node **nsx-edge-02** with IP address **192.168.100.4/24**.

2. Click **CLOSE**

Add Secondary Uplink to BGP Source IPs

The screenshot shows the NSX configuration interface for BGP. The left sidebar lists the configuration sections: INTERFACES, ROUTING, MULTICAST, BGP, and ROUTE RE-DISTRIBUTION. The BGP section is expanded, showing various settings. The 'BGP Neighbors' section at the bottom right has a link '1' circled in red, indicating the next step in the configuration process. The 'CLOSE EDITING' button is visible at the bottom right.

Section	Setting	Value	Status
INTERFACES	External and Service Interfaces	2	
ROUTING			
MULTICAST			
BGP	Local AS	65001	
	BGP	On	On
	Graceful Restart	Helper Only	
	Graceful Restart Timer	180	Range 1 to 3600 seconds
	Route Aggregation	Set	
	Inter SR iBGP	On	On
	ECMP	On	On
	Multipath Relax	On	On
	Graceful Restart Stale Timer	600	Range 1 to 3600 seconds
	BGP Neighbors	1	1
ROUTE RE-DISTRIBUTION			

We will now configure BGP to advertise from the second interface that we defined on nsx-edge-02. This will allow BGP on the Tier-0 Gateway to establish peering from the interfaces on both nsx-edge-01 and nsx-edge-02. During normal operation, both Edges will be considered viable paths into and out of the NSX environment. In the event that an Edge Transport Nodes fail, its BGP neighbor state will be lost and its path information will be removed from the BGP routing table. Traffic will continue to flow through the remaining Edge Transport Node. Upon recovery of the lost Edge Transport Node, its BGP state will be reestablished and its path information will be added back to the BGP routing table automatically.

1. Click the **1** link to the right of **BGP Neighbors** to display the Set BGP Neighbors dialog

Set BGP Neighbors

Set BGP Neighbors

Tier-0 Gateways Tier-0-gatew... #Neighbors 1

ADD BGP NEIGHBOR EXPAND ALL Search

	IP Address	BFD	Remote AS number	Route Filter	Allowas-in	Status
1 > 	192.168.100.1	Disabled	65002	1	Disabled	Success  

REFRESH 1 - 1 of 1 BGP Neighbors

CLOSE

1. Click the **More Options** icon to the left of **192.168.100.1** to display its Options menu, then click **Edit**

Edit BGP Neighbor Settings

Set BGP Neighbors

Tier-0 Gateways Tier-0-gatew... #Neighbors 1

ADD BGP NEIGHBOR COLLAPSE ALL Search

IP Address	BFD	Remote AS number	Route Filter	Allowas-in	Status
192.168.100.1	Disabled	65002 E.g. 65000	1	Disabled	

Source Addresses: 192.168.100.3, 192.168.100.4

Max Hop Limit: 1

Graceful Restart: Helper Only

Description: Description

BFD Interval: 500

BFD Multiplier: 3
E.g. 2

Hold Down Time: 15

Password:

Keep Alive Time: 5

SAVE CANCEL


REFRESH 1 - 1 of 1 BGP Neighbors

CLOSE

1. Source Addresses: Add **192.168.100.4/24** (this should be in addition to the existing 192.168.100.3/24 entry)
 - **Note:** You must press **ENTER** or click the **Add Item(s): 192.168.100.4/24** link once you've entered the IP address
2. Click the arrow to the left of **TIMERS & PASSWORD** to expand it
3. Hold Down Time: **15**
4. Keep Alive Time: **5**
5. Click **SAVE**
6. Click **CLOSE**


In addition to adding the secondary interface's IP address as a BGP Source, we also modified the BGP Hold Down and Keep Alive timers to 15 and 5 seconds, respectively. We lowered these values in order to speed up BGP reconvergence. This will be useful when we test Edge Node failure later in this module.

Open PuTTY




1. Click the **Start Menu icon** in the Windows Task Bar
2. Click **PuTTY**

Connect to nsx-edge-02



1. Scroll through the list of Saved Sessions until **nsx-edge-02** is visible
2. Click **nsx-edge-02** to select and highlight it
3. Click **Load** to load the saved session
4. Click **Open** to launch the SSH session

Log in to nsx-edge-02



Log in to nsx-edge-02 with the following credentials:

1. Password: **VMware1!VMware1!**

Once you are authenticated to the Edge Node, maximize the PuTTY window for better visibility.

- Click the Windows **Maximize icon**

List Logical Routers Connected to nsx-edge-02

```
nsx-edge-02> get logical-routers
Logical Router
UUID                               VRF  LR-ID  Name                                     Type                               Ports
736a80e3-23f6-5a2d-81d6-bbefb2786666 0    0      0                                         TUNNEL                             3
6a7f2777-f857-41a9-ba92-f6162e53bad2 1    2049   DR-Tier-1-gateway-01                   DISTRIBUTED_ROUTER_TIER1           7
275e6800-7f8b-4b0e-a1b4-3e2ff8ea0c1e 2    2050   SR-Tier-0-gateway-01                   SERVICE_ROUTER_TIER0               6
1b7906c1-12e9-4469-a7d8-1a42ed65dea3 4    1      DR-Tier-0-gateway-01                   DISTRIBUTED_ROUTER_TIER0           5

nsx-edge-02>
```

- Get a list of Logical Routers connected to nsx-edge-02 by running the following command:

```
get logical-routers
```

Note the VRF number of the Logical Router **SR-Tier-0-gateway-01**.

NOTE: The VRF number of **SR-Tier-0-gateway-01** may differ from the screenshot.

Verify BGP Neighbor Relationship with Upstream Router

```
nsx-edge-02> vrf 2
nsx-edge-02(tier0_sr)> get bgp neighbor summary
BFD States: NC - Not configured, AC - Activating, D - Disconnected
AD - Admin down, DW - Down, IN - Init, UP - Up
BGP summary information for VRF default for address-family: ipv4Unicast
Router ID: 192.168.100.4 Local AS: 65001

Neighbor          AS           State Up/DownTime  BFD InMsgs  OutMsgs  InPfx  OutPfx
192.168.100.1     65002        Estab 00:03:15  NC 192      192      1      5

nsx-edge-02(tier0_sr)>
```

- Enter the VRF routing context on the Edge Node by entering the following command (**NOTE:** Replace "2" in the command below with the VRF number found in the previous step):

```
vrf 2
```

2. Get the BGP neighbor status by running the following command:

```
get bgp neighbor summary
```

Verify the neighbor relationship with **192.168.100.1** is showing a state of **Estab** (Established).


HA Confirmed

As you can see, we now have two edge nodes that have established connections with our external router, providing redundant North-South routing to the NSX environment.

Equal Cost Multi-Pathing (ECMP)


In this lesson we will test Equal Cost Multi-Pathing (ECMP) by simulating an Edge Node failure.

Open PuTTY



1. Click the **Start Menu icon** in the Windows Task Bar
2. Click **PuTTY**


Open Command Prompt



We will now start a ping session to one of the sample web servers located on NSX Segment LS-web. Once this has been done, we will shut down nsx-edge-01, simulating a failure of an Edge Node. We should then observe BGP detecting the loss of connectivity to the Edge Node, routing all traffic through nsx-edge-02.

1. Click the **Start Menu icon** in the Windows Task Bar
2. Click **Command Prompt**

Ping VM web-01a (Before Failure)



The screenshot shows a Windows Command Prompt window titled "Administrator: Command Prompt". The command prompt displays the following text:

```
Microsoft Windows [Version 6.3.9600]
(c) 2013 Microsoft Corporation. All rights reserved.

C:\Users\Administrator>ping web-01a.corp.local 1

Pinging web-01a.corp.local [172.16.10.11] with 32 bytes of data:
Reply from 172.16.10.11: bytes=32 time=2ms TTL=61
Reply from 172.16.10.11: bytes=32 time=2ms TTL=61
Reply from 172.16.10.11: bytes=32 time=2ms TTL=61
Reply from 172.16.10.11: bytes=32 time=7ms TTL=61

Ping statistics for 172.16.10.11:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 2ms, Maximum = 7ms, Average = 3ms

C:\Users\Administrator>
```

In the screenshot, the command `ping web-01a.corp.local` is highlighted with a red rectangle, and the number `1` is circled in white. A red arrow points to the first line of the ping output: `Reply from 172.16.10.11: bytes=32 time=2ms TTL=61`.

1. Ping web-01a.corp.local on NSX Segment LS-web:

```
ping web-01a.corp.local
```

You should observe ping replies from the web server.

Traceroute Path to VM web-01a (Before Failure)

```
Administrator: Command Prompt
Microsoft Windows [Version 6.3.9600]
(c) 2013 Microsoft Corporation. All rights reserved.
C:\Users\Administrator>tracert -d web-01a.corp.local
Tracing route to web-01a.corp.local [172.16.10.11]
over a maximum of 30 hops:
  0  <1 ms    <1 ms    <1 ms    192.168.110.1
  1  1 ms     1 ms     1 ms     192.168.100.3
  2  *        *        *        Request timed out.
  3  3 ms     2 ms     2 ms     172.16.10.11
Trace complete.
C:\Users\Administrator>
```

By configuring the second interface on nsx-edge-02, we now have two fault tolerant paths into the NSX environment. We will perform a traceroute before simulating an Edge Node failure, observing that traffic will fail over to the secondary interface of the Tier-0 Gateway.


1. Traceroute to web-01a.corp.local to display its path:

```
tracert -d web-01a.corp.local
```

The first hop is the IP address of the vPod router (the gateway of your admin desktop). Observe the second hop, **192.168.100.3**, which is the IP address of the Tier-0 Gateway interface on nsx-edge-01. Traffic is then routed to the Tier-1 gateway via the unadvertised network in hop three, and finally delivered to web-01a at 172.16.10.11 in the fourth and final hop.

NOTE: Because both paths are equally valid, your traceroute may traverse the Tier-0 interface on nsx-edge-02 instead of nsx-edge-01. If this is the case, your second route hop will display **192.168.100.4 instead of 192.168.100.3**. If this occurs, please substitute nsx-edge-02 in the following steps to test fault tolerance.


Power Off Appropriate Edge



We will now connect to vCenter and simulate a failure by powering down the Edge Node your trace route utilized. The loss of this Edge Node will cause all traffic to route through the remaining Edge Node.


1. Click on the **new tab** icon
2. Click on the **Region A** folder
3. Click on **vcsa-01a WEB Client (HTML)**

Login to vCenter



1. Click the **checkbox** for Use Windows session authentication
2. Click **Login**


Find Appropriate Edge Node



Once logged in, navigate to Hosts and Clusters:


1. Click on the **Menu** button
2. Click on **Hosts and Clusters**

Power Off Edge Node VM



NOTE: Because both paths are equally valid, the output of the traceroute you performed earlier may traverse the Tier-0 interface on either nsx-edge-01 or nsx-edge-02. If your second route hop in the earlier traceroute displays **192.168.100.3**, please power off **nsx-edge-01**. Likewise, if the second route hop in the earlier traceroute displays **192.168.100.4**, please power off nsx-edge-02 to test fault tolerance.

1. Right click on **nsx-edge-01** (or **nsx-edge-02**, as mentioned above)
2. Click **Power**
3. Click **Shut Down Guest OS**



1. Click **YES**

Ping VM web-01a (During Failure)

```
Administrator: Command Prompt
Microsoft Windows [Version 6.3.9600]
(c) 2013 Microsoft Corporation. All rights reserved.

C:\Users\Administrator>ping web-01a.corp.local

Pinging web-01a.corp.local [172.16.10.11] with 32 bytes of data:
Reply from 172.16.10.11: bytes=32 time=2ms TTL=61
Reply from 172.16.10.11: bytes=32 time=2ms TTL=61
Reply from 172.16.10.11: bytes=32 time=2ms TTL=61
Reply from 172.16.10.11: bytes=32 time=7ms TTL=61

Ping statistics for 172.16.10.11:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 2ms, Maximum = 7ms, Average = 3ms

C:\Users\Administrator>ping web-01a.corp.local 1
Pinging web-01a.corp.local [172.16.10.11] with 32 bytes of data:
Request timed out.
Request timed out.
Request timed out.
Request timed out.

Ping statistics for 172.16.10.11:
    Packets: Sent = 4, Received = 0, Lost = 4 (100% loss),

C:\Users\Administrator>
```



Return to the Command Prompt window and perform the following:

1. Ping web-01a.corp.local:

```
ping web-01a.corp.local
```

Ping requests should time out for approximately 15 seconds before BGP reconverges and removes the failed path through nsx-edge-01. The amount of time required is determined by the BGP timers, which we changed to specify a Keep Alive time of 5 seconds and a Hold-Down time of 15 seconds.

Ping VM web-01a (After Reconvergence)



```
Administrator: Command Prompt
Ping statistics for 172.16.10.11:
  Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
  Approximate round trip times in milli-seconds:
    Minimum = 2ms, Maximum = 7ms, Average = 3ms

C:\Users\Administrator>ping web-01a.corp.local

Pinging web-01a.corp.local [172.16.10.11] with 32 bytes of data:
Request timed out.
Request timed out.
Request timed out.
Request timed out.

Ping statistics for 172.16.10.11:
  Packets: Sent = 4, Received = 0, Lost = 4 (100% loss),

C:\Users\Administrator>ping web-01a.corp.local 1

Pinging web-01a.corp.local [172.16.10.11] with 32 bytes of data:
Reply from 172.16.10.11: bytes=32 time=2ms TTL=61
Reply from 172.16.10.11: bytes=32 time=2ms TTL=61
Reply from 172.16.10.11: bytes=32 time=3ms TTL=61
Reply from 172.16.10.11: bytes=32 time=2ms TTL=61

Ping statistics for 172.16.10.11:
  Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
  Approximate round trip times in milli-seconds:
    Minimum = 2ms, Maximum = 3ms, Average = 2ms

C:\Users\Administrator>
```


Once 15 seconds have elapsed, repeat the ping test and verify reconvergence.

1. Ping web-01a.corp.local:

```
ping web-01a.corp.local
```

You should observe that connectivity has been restored and ping replies are being received for server web-01a. If this is not the case, please wait a moment and try again.

Traceroute Path to VM web-01a (After Reconvergence)



```
Administrator: Command Prompt
Microsoft Windows [Version 6.3.9600]
(c) 2013 Microsoft Corporation. All rights reserved.
C:\Users\Administrator>tracert -d web-01a.corp.local
Tracing route to web-01a.corp.local [172.16.10.11]
over a maximum of 30 hops:
  0  <1 ms    <1 ms    <1 ms    192.168.110.1
  1  1 ms     1 ms     1 ms     192.168.100.4
  2  *        *        *        Request timed out.
  3  3 ms     2 ms     2 ms     172.16.10.11
Trace complete.
C:\Users\Administrator>
```


Since BGP has reconverged, its path to the Tier-0 Gateway should now be through its interface on nsx-edge-02 (or nsx-edge-01, as noted earlier).

1. Traceroute to web-01a.corp.local to display its path:

```
tracert -d web-01a.corp.local
```

The first hop is the IP address of the vPod router (the gateway of your admin desktop). The second hop is now **192.168.100.4** (or **192.168.100.3**), the IP address of the Tier-0 Gateway interface on nsx-edge-02 (or nsx-edge-01). Traffic is then routed to the Tier-1 gateway via the unadvertised network in hop three, and finally delivered to web-01a at 172.16.10.11 in the fourth and final hop.

Power On Edge Node VM



Now that we have tested fault tolerance on the Edge Node, we will return the Edge Node VM to a Powered On state. Return to the vSphere Client then perform the following.

1. Right click on **nsx-edge-01** (or **nsx-edge-02**, as mentioned previously)
2. Click **Power**
3. Click **Power On**


Module 4 Conclusion

Congratulations on completing Module 4.

Please proceed to any module below which interests you the most.

- [Module 1 - Introduction to VMware NSX](#) (15 minutes) In this module you will be introduced to NSX, its capabilities, and the components that make up the solution.
- [Module 2 - Host and Edge Transport Node Preparation](#) (30 minutes) In this module you will review the NSX lab fabric.
- [Module 3 - Logical Switching](#) (15 minutes) In this module you will learn how to configure Logical Switches and attach them to virtual machines.
- [Module 5 - Distributed Firewall and Tools](#) (30 minutes) In this module you will learn how to configure the Distributed Firewall, Spoof Guard, and use various tools within NSX-T.
- [Module 6 - NSX Edge Services](#) (15 minutes) Learn how to configure and enable the NSX Load balancer for our web servers.

How to End Lab



To end your lab click on the **END** button.

Module 5 - Distributed Firewall and Tools (30 minutes)

Distributed Firewall and Tools - Module Overview

The goal of this module is to demonstrate how the distributed firewall (DFW) and operational tools within NSX function and are configured. The Distributed Firewall in NSX-T 3.0 is installed by default with a final Allow all rule. This means that all traffic is permitted and Micro Segmentation is "off". To enable Micro Segmentation you need to change the last rule from Allow to Deny. In this module we will execute the following operations:

DFW Section:


- Verify the functionality of a 3 Tier App while the NSX DFW final rule is "Allow Any"
- Change the final rule to "Deny all"
- Verify that our 3 Tier Web App no longer functions and is being blocked by the default deny all rule
- Enable preconfigured 3 Tier App DFW rules
- Verify the preconfigured rules are being utilized with Log Insight
- Verify our 3 Tier Web App functions again
- Explore NSX-T grouping and Recreate DFW allow rules
- Verify our 3 Tier Web App functions again

Tools Section:

- Use trace flow to verify your new rules are utilized


HOL-2126 Logical Diagram

This diagram illustrates what virtual machines make up our 3 Tier Web App for testing



3 Tier Web App ports

This diagram illustrates the port requirements for our 3 Tier Web App



Distributed Firewall in NSX

In this Chapter we will review and configure the Distributed Firewall of NSX-T

By default the NSX DFW is configured with a final Allow All rule. This means that all traffic is allowed and in order to "enable Micro segmentation" or block all traffic we will need to change the final rule to Deny. Before we change the final rule lets verify our precreated 3 Tier App works as expected.


NOTE: Only focus on given VMs for the 3-tier app. Lab has additional VMs for other modules.

Launch Chrome



1. Double click on the **Chrome icon** on the desktop

Connect to 3 Tier Web App



1. Click on the **3 Tier App** bookmark bar folder
2. Click on the **Web-01a** shortcut

Verify connectivity to Web-01a

Customer Database Access

Accessed via: web-01a

Name Filter (blank for all records):

Rank	Name	Universe	Revenue
1	CHOAM	Dune	\$1.7 trillion
2	Acme Corp.	Looney Tunes	\$348.7 billion

1. Verify you have successfully connected to Web-01a and it has retrieved data from the App server

Note: Feel free to test Web-02a and Web-03a shortcuts to verify they work as well.

Now that we have verified our 3 Tier App works let's change our final rule to Deny All

Launch NSX-T Manager

Customer Database Access

Accessed via: web-01a

Name Filter (blank for all records):

Rank	Name	Universe	Revenue
1	CHOAM	Dune	\$1.7 trillion
2	Acme Corp.	Looney Tunes	\$348.7 billion

Login to NSX-T Manager

1. **Click** to open a new tab
2. Click **Region A** to open the bookmark folder
3. Click **NSX-T shortcut** to launch log in page

VMware® NSX-T™

nsx-admin@corp.local

.....

LOG IN

1. Input **nsx-admin@corp.local** for the user name and **VMware1!** for the password
2. Click **Log in**

Note: This log in is utilizing the new Active Directory integration in version 3.0

Navigate to the DFW management page

The screenshot shows the NSX-T Security page. The 'Security' tab is selected in the top navigation bar (1). In the left sidebar, 'Distributed Firewall' is highlighted under the 'East West Security' section (2). In the main content area, the 'DISTRIBUTED FIREWALL' section is visible. A table lists rules, and the 'Default Layer3' rule is highlighted, with its expand icon (a right-pointing chevron) circled (3).

	☐	Name	ID	Sources
>	☐	3-tier-app	(3)	Applied To
>	☐	K8s-Top	(0)	Applied To
>	☐	K8s-Bottom	(0)	Applied To
>	☐	Catch all rul...	(2)	Applied To
>	☐	Default Lay...	(1)	Applied To

1. Click on **Security**
2. Click on **Distributed Firewall**
3. Click to **expand** the Default Layer3 Section


Set the Default Layer 2 rule to Drop



1. Click the **drop down arrow** to change the default rule action
2. Choose **Drop** as the rule action


Note: Once we publish this change explicit allow rules must be made for communication to be allowed in the environment. Let's verify our 3 Tier App is being blocked.

Publish the change



- 1. Click **Publish** to change the default rule action


Verify 3 Tier App traffic is blocked



1. Switch back to the **first Chrome tab**
2. Click the **Refresh** icon
3. Verify the App can no longer be accessed


Note: It may take up to 20 seconds for the page to timeout, you can also verify web-02a and web-03a cannot be accessed. Now that we know the app can not be reached lets enable the preconfigured rules and test again.

Switch back to the NSX-T management tab




1. Click the **NSX-T management tab**

Explore the preconfigured 3 Tier App rules



1. Verify you are in the **application section** of the DFW
2. Expand the **3 Tier App section**
3. Review the **preconfigured rules** required for the 3 Tier App to function

Enable the preconfigured DFW rules



1. Enable each preconfigured rule by **clicking the enable / disable slider** to the right of the rule
2. Click **Publish** to save your changes

Now that the allow rules are enabled lets test the 3 Tier App connectivity and verify the rules are applied.

Test 3 Tier App connectivity

The screenshot shows a web browser window with the following elements:


- Tab 1 (NSX):** Labeled with a red box and a circled '1'.
- Address Bar:** Shows the URL `https://web-01a.corp.local/cgi-bin/app.py`. The refresh icon is circled with a '2'.
- Breadcrumbs:** HOL Admin > Region A > 3-tier-app > phoenix app.
- Page Title:** Customer Database Access.
- Status:** A box says 'Accessed via: web-01a', with a red arrow pointing to it from a circled '3'.
- Form:** 'Name Filter (blank for all records):' followed by an input field and an 'Apply' button. A red arrow points to the input field from the circled '3'.
- Table:** A table with 4 columns: Rank, Name, Universe, and Revenue. It contains 3 rows of data.

Rank	Name	Universe	Revenue
1	CHOAM	Dune	\$1.7 trillion
2	Acme Corp.	Looney Tunes	\$348.7 billion
3	Sirius Cybernetics Corp.	Hitchhiker's Guide	\$327.2 billion

1. Switch back to the **first Chrome tab**
2. Click the **refresh icon**
3. Verify the App can be **accessed** again

We have now enabled the DFW within NSX-T and enabled the preconfigured 3 Tier App rules. Next we will verify through logging the rules work as expected.

Verify 3-tier rules are being logged



1. Change focus back to the **NSX** tab
2. Click the **gear icon** next to the 3-tier-client-access rule

Note log label

Settings

Rule > 3-tier client access

Logging ☒ Enable

Direction In-Out ▾

IP Protocol ☐ IPV4 ☐ IPV6 ☒ IPV4-IPV6

Log Label 3-tier client access


Notes

CANCEL APPLY

1. **Note** that the log label is set to "3-tier client access" and **logging is enabled**
2. Click **Cancel** to exit

Note: You can also explore each individual rule to see the log label set for each.

Launch Log Insight




1. **Click** to open a new tab
2. Click to **open the Region A bookmark folder**
3. Click **vRealize Log Insight shortcut** to launch log in page

Log into Log Insight

A screenshot of the VMware vRealize Log Insight login page. The page title is 'VMware vRealize® Log Insight™'. There is a text input field for the username, containing 'admin'. Below it is a password input field, represented by a yellow bar with dots. A red box labeled '1' highlights the username and password fields. Below the password field is a blue button labeled 'LOGIN'. A red box labeled '2' highlights the 'LOGIN' button.

1. Input **admin** for the user name and **VMware1!** for the password
2. Click **Log in**

Review hit count for the 3-tier app rules



To ensure you are viewing the page shown you may need to select **Dashboard 1** under **My Dashboards** in the left side panel.


1. Note that we have a **recent hit** on all three 3-tier app rules
2. Click to **launch the interactive analytics** for the 3-tier client access rule
3. Optional: If you do not see any hits / logs for the rules you can change the time frame to the **Last hour of data**

Review the log detail



1. Note that we were able to log the successful rule hit by using the log label "**3-tier client access**"
2. Note that the allowed traffic was from 192.168.110.10 (your local machine) to TCP 443 on 172.16.10.11 (web-01a)


Delete the preconfigured 3 Tier App policy



We will now delete and reconfigure the rules and groups to take a more detailed look at how they are configured. If you would like to skip this configuration you can jump ahead to the tools section or the next module.


1. Switch back to the **NSX-T Chrome tab**
2. Click the **3 dots** next to the 3 Tier App policy
3. Click **Delete policy**

Publish the changes




1. Click **Publish** to publish the changes

Navigate to the Groups screen in the Inventory




1. Click on **Inventory**
2. Click **Groups**

Delete the preconfigured 3-tier-db servers and 3-tier-app-servers groups



1. Click the the **three dots next to the 3-tier-app-servers** group
2. Click **Delete**


Confirm delete



1. Click **Delete** to confirm

Do this for two of the three 3-tier groups (**app_servers** and **db_servers**)

Delete 3-tier-db-servers



1. Delete the **3-tier-db-servers** by following the same steps

Note: the 3-tier-web-servers group cannot be deleted because it is in use for another rule, we will reuse this group.

Refresh the screen

	Name	Compute Members
 > 	3-tier-web-servers	View Members
 > 	AD-group-app-users	View Members
 > 	AD-group-it-admins	View Members
 > 	AD-servers	View Members
 > 	all -phoenix-servers	View Members
 > 	all-3-tier-servers	View Members
 > 	phoenix-app-servers	View Members
 > 	phoenix-db-servers	View Members
 > 	phoenix-web-servers	View Members
 > 	user-vm-01	View Members
 REFRESH		

1. Click **refresh**
2. **Verify** the groups were deleted

Recreate App Servers group

The screenshot shows a 'Add Group' dialog box. At the top left, there is a button labeled 'ADD GROUP' circled in red, with a black circle containing the number '1' next to it. Below this, there is a table with two columns: 'Name' and 'Compute Me'. The 'Name' column has a text input field containing '3-tier-app-servers', which is also circled in red, with a black circle containing the number '2' next to it. The 'Compute Me' column has a button labeled 'Set Members' circled in red, with a black circle containing the number '3' next to it. Below the table, there is a 'Description' label and a text input field. At the bottom of the dialog, there are two buttons: 'SAVE' and 'CANCEL'.

1. Click **Add Group**
2. Input **3-tier-app-servers**
3. Click **Set Members**

Select App Server group members

Select Members | app_servers

Add Compute Members either by creating or by directly adding them. You can also add Identity members separately. Identity members intersect with the Compute members to form the effective membership of the group.

Membership Criteria (0) **Members (1)** IP/MAC Addresses (0) AD Groups (0)

2 app-01a X

Select Category VirtualMachine (selected: 1) Maximum: 500 Members with IP and MAC Addresses

	Name
☒	app-01a
☐	core-A
☐	db-01a
☐	k8s-db-01a
☐	k8s-master

1 REFRESH 1 - 11 of 11


Show Only Selected ☐

4 CANCEL APPLY

1. Click on **Members**
2. Select Category **Virtual Machine** from the drop down
3. **Check box** app-01a
4. Click **Apply**

Note: This is just one way of creating a group of virtual machines, the groups you previously deleted utilized tags instead of static members.


Save the group



The screenshot shows a configuration form for a group named '3-tier-app-servers'. The form has two tabs: 'Name' and 'Compute Members'. The 'Name' tab is active, showing a text input field with '3-tier-app-servers' and a red asterisk. Below the input field is a 'Description' label and a text area. At the bottom, there are two buttons: 'SAVE' (highlighted with a red box and a circled '1') and 'CANCEL'.

1. Click **Save** to save the group

Create DB Servers group



The screenshot shows a configuration form for a group named '3-tier-db-servers'. The form has two tabs: 'Name' and 'Compute Members'. The 'Name' tab is active, showing a text input field with '3-tier-db-servers' and a red asterisk. Below the input field is a 'Description' label and a text area. At the bottom, there are two buttons: 'SAVE' and 'CANCEL'. The 'Compute Members' tab is also visible, showing a 'Set Members' button. The 'ADD GROUP' button is highlighted with a red box and a circled '1'. The 'Name' tab is highlighted with a circled '2'. The 'Set Members' button is highlighted with a red box and a circled '3'.

1. Click **Add Group**
2. Input **3-tier-db-servers**
3. Click **Set Members**

Select DB Server group members

Select Members | db_servers

Add Compute Member by creating or by directly adding them. You can also add Identity members separately. Identity members intersect with the Compute members to form an effective membership of the group.

Membership Criteria (0) **Members (1)** IP/MAC Addresses (0) AD Groups (0)

db-01a X

Select Category VirtualMachine (selected: 1) Maximum: 500 Members with IP and MAC Addresses

	Name
☐	app-01a
☐	core-A
☒	db-01a
☐	k8s-db-01a
☐	k8s-master

1 - 11 of 11

Show Only Selected ☐

CANCEL APPLY

1. Click on **Members**
2. Select Category **Virtual Machine** from the drop down
3. **Check box** db-01a
4. Click **Apply**

Note: This is just one way of creating a group of virtual machines, the groups you previously deleted utilized tags instead of static members.

Save the group


1. Click **Save** to save the group

Verify your new groups have been created

	Name	Compute Members	
⋮ >	3-tier-app-servers	View Members	2
⋮ >	3-tier-db-servers	View Members	1
⋮ >	3-tier-web-servers	View Members	


1. **Verify** your two groups were created and the 3-tier-web-servers group still remains.
2. Optional: You can click each group's **View Members link** to verify the correct vm's are included.

3 Tier App port requirements




As a reminder here are the port requirements for the 3 Tier App to function. Next let's go to the Distributed Firewall section to create the rules.

Navigate to the DFW section




1. Click **Security**
2. Click **Distributed Firewall**
3. Click **Application**

Create a new Security Policy




1. Click **Add Policy**
2. Type **3 Tier App** in the text field

Add Client Access rule




1. Click the **three dots** next to the 3 Tier App Policy
2. Click **Add Rule**

Configure the rule



1. Click on the name field and name the rule **Client Access**
2. Leave the Source as **Any**
3. **Click** on the pencil icon under Destinations

Set destination group



1. **Check** the check box next to 3-tier-web-servers group you created earlier
2. Click **Apply** to save the Destination

Edit the service

	☐	Name	ID	Sources	Destinations	Services	Profiles	Applied To	Action
⋮	▼	☐ 3 Tier App (1)	Applied To	D...			1		
⋮		☐ Client access		Any	3-tier-w...	Any		None	DFW

1. Click the **pencil icon** under Services

Set the service

Set Service

Rule > Client Access

1

HTTPS X

ADD NEW SERVICE

	Name	Service Entries	Tags	Status
☐	CIM-HTTPS	TCP (Source: Destination: 5989)	0	● Up
☒	HTTPS	TCP (Source: Destination: 443)	0	● Up
☐	HTTPS, net.tcp binding	TCP (Source: Destination: 32843,32844,32845)	0	● Up
☐	SAP HTTPS	TCP (Source: Destination: 50001)	0	● Up
☐	VMware-VCO-VCO-HTTPS	TCP (Source: Destination: 8283)	0	● Up
☐	VMware-VCO-WebHTTPS	TCP (Source: Destination: 8281)	0	● Up

☒ 1  REFRESH


1 - 6 of 6 Services

CANCEL

APPLY

1. Type **HTTPS** in the search box to find the HTTPS service
2. **Check** the check box next to the HTTPS service
3. Click **Apply**

Verify and publish the Client Access Rule



1. Verify the **Client Access rule** is configured as follows
2. Click **Publish**

Client Access Rule settings:

Name: Client Access


Sources: Any

Destinations: 3-tier-web-servers

Services: HTTPS

Action: Allow

Create Web to App access rule



1. Use the same process to **create the Web to App Access rule** as you did for the Client Access rule
2. Click **Publish**

Web to App Access Rule settings:

Name: Web to App Access

Sources: 3-tier-web-servers

Destinations: 3-tier-app-servers

Services: TCP_8443

Action: Allow

Create App to DB access rule

The screenshot shows the Palo Alto Networks Distributed Firewall (DFW) interface. The 'APPLICATION' category is selected, showing 12 rules. A red box highlights the '3 Tier App (3)' rule, which is 'Uninitialized'. A red circle with the number '1' is next to the rule's expand icon, and a red circle with the number '2' is next to the 'PUBLISH' button in the top right corner.

Name	ID	Sources	Destinations	Services	Profiles	Applied To	Action
3 Tier App (3)			D...				Uninitialized
App to DB ac...		3-tier-a...	3-tier-d...	MySQL	None	DFW	Allow
Web to App ...	5096	3-tier-w...	3-tier-a...	TCP_84...	None	DFW	Allow
Client access	5097	Any	3-tier-w...	Any	None	DFW	Allow

1. Use the same process to **create the App to DB Access rule** as you did for the Client Access rule
2. Click **Publish**

App to DB Access Rule settings:

Name: App to DB Access

Sources: 3-tier-app-servers

Destinations: 3-tier-db-servers

Services: MySQL TCP 3306

Action: Allow

Test the 3 Tier App

Customer Database Access

Accessed via: web-01a

Name Filter (blank for all records):

Rank	Name	Universe	Revenue
1	CHOAM	Dune	\$1.7 trillion
2	Acme Corp.	Looney Tunes	\$348.7 billion
3	Sirius Cybernetics Corp.	Hitchhiker's Guide	\$327.2 billion
4	Buy n Large	Wall-E	\$291.8 billion

1. Switch back to the 3 Tier App **Chrome tab**
2. Click the **refresh** button
3. Verify the 3 Tier App functions properly

Congratulations you have successfully configured micro segmentation rules for a 3 Tier App!!!

Tools


NSX provides a variety of tools to make operations and troubleshooting easier.

Traceflow, IPFIX, Port Mirroring and Consolidated Capacity can be found under the Plan & Troubleshoot section. In this chapter, we will explore Traceflow as it relates to the 3-tier app DFW rules.

Here is a list of the tools and their function.

- **Traceflow:** test data plane between two vms by injecting traffic
- **Port Mirroring:** packet copy function applied between NICs on the same Transport Node
- **IPFIX:** generate and send flow records of traffic to a remote collector
- **Consolidated Capacity:** View current object inventory and maximum capacity


Verify Security policy



Before we explore Traceflow let's verify our 3-tier app rules are set as expected and the final rule in the DFW is set to deny.


1. Within the NSX Manager, click on **Security**
2. Click on **Distributed Firewall**
3. Click the arrow to **expand the 3-tier-app security policy**
4. Click the arrow to **expand the Default Layer3 security policy**

Verify DFW rules



1. **Make note** of the web to app allow rule that permits TCP 8443 traffic from web to app servers
2. **Verify** all three 3-tier app rules are **enabled**
3. **Verify** the final default rule is set to **block**

Traceflow



Traceflow traces the transport node-level path of a packet injected at a logical port. Traceflow supports L2 and L3 and is supported across ESXi, KVM Edge (including NAT). Using details provided by the user, a Traceflow packet will be created and inserted into the source logical port. When this packet travels from source to destination, all the logical network entities will report the observations. These observations are consolidated and shown in UI.

1. Click on **Plan & Troubleshoot** on the top bar
2. Click on **Traceflow**

Select web-01a as the Source

The screenshot shows a 'Source' configuration window with a 'RESET' button. The 'Type' is set to 'Virtual Machine' (indicated by a circled '1'). The 'VM Name' field is open, showing a list of VMs: phoenix-db-01, phoenix-web-01, phoenix-web-02, user-vm-01, web-01a, web-02a, web-03a, and web-04a. The 'web-01a' option is highlighted with a red box (indicated by a circled '2').

1. Click the **drop down arrow** next to VM name
2. Select **web-01a**

Select App 01a as the Destination

The screenshot shows a 'Destination' configuration window with a 'RESET' button. The 'Type' is set to 'Virtual Machine'. The 'VM Name' field is open, showing a list of VMs: app-01a, core-A, db-01a, and k8smaster-01a. The 'app-01a' option is highlighted with a red box. A red box also highlights the drop-down arrow next to the VM Name field.

1. Click the **drop down arrow** next to VM name
2. Select **app-01a**

Set the Protocol and port

Traceflow

1

2

Protocol Type
TCP

Source Port 0

Destination Port 8443

TCP Flags SYN X

Source RESET

Type Virtual Machine

VM Name web-01a
Host:8d528a95-b198-40b8-9e62-71305c82e0e3

Identify a virtual interface for packet injection

Virtual Interface Network adapter 1
Segment Port: default: 610c8284-3f60-4d58-8b8c-7f520a77b954

IP Address * 172.16.10.11

MAC Address * 00:50:56:9c:a5:b3

Destination RESET

Type Virtual Machine

VM Name app-01a
Host:8d528a95-b198-40b8-9e62-71305c82e0e3

Virtual Interface Network adapter 1
Segment Port: default: 56456de3-041d-4190-8c9a-b7803bafbd1c

IP Address * 172.16.20.11

MAC Address * 00:50:56:9c:d1:1e


3

ADVANCED SETTINGS

TRACE

1. Set the Protocol type to **TCP**
2. Set the **Destination port to 8443**
3. Click **Trace**

Traceflow Allow Results



Once the trace has completed, NSX-T Traceflow will provide a map and detailed information detailing the packets path.

1. As you can see in the diagram, our packet was **injected, forwarded, and delivered** as expected.
2. Note that the flow was allowed by the **web to app DFW** rule created earlier in the module. (Your rule ID may be different if you continued on in the previous section and recreated the security policy.)
3. Feel free to **scroll down** to view more details of the flow.

Edit the trace



1. Click the **edit** button

Warning

Click Proceed to start a new traceflow request. Doing so would discard the current trace results.

1

CANCELPROCEED

1. Click **proceed**

Change to port 22

IP Address
IPv4

Traffic Type
Unicast

Protocol Type
TCP

Source Port
0

1

Destination Port
22

TCP Flags
SYN x

Source RESET

Type
Virtual Machine

VM Name
web-01a

Identify a virtual interface for packet injection

Virtual Interface
Network adapter 1

IP Address *
172.16.10.11

MAC Address *
00:50:56:9c:a5:b3

Destination RESET

Type
Virtual Machine

VM Name
app-01a

Virtual Interface
Network adapter 1

IP Address *
172.16.20.11

MAC Address *
00:50:56:9c:d1:1e


2

ADVANCED SETTINGS

TRACE

1. Change the **destination port to 22**
2. Click **trace**

Traceflow block results



- 1. Note that the flow was **dropped** this time by the default Drop firewall rule

This concludes the tools section of this module, feel free to view flows between different vm's and ports within the environment.

Module 5 Conclusion

Congratulations on completing Module 5! Next we will explore NSX services.

Please proceed to any module below which interests you the most.

- [Module 1 - Introduction to NSX](#) (15 minutes) (Basic) In this module you will be introduced to the NSX Data Center platform and its capabilities. You will also explore the new NSX-T 3.0 features.
- [Module 2 - Host and Edge Transport Node Preparation](#) (30 minutes) (Advanced) In this module, you will review the prepared components ready for NSX workloads.
- [Module 3 - Logical Switching](#) (15 minutes) (Advanced) In this module, you will review the NSX lab fabric, configure Logical Switches, and attach them to virtual machines. You will be introduced to the use of the GENEVE protocol for overlay networking.
- [Module 4 - Logical Routing](#) (30 minutes) (Advanced) In this module you will learn how to configure Logical Routing in NSX for North-South and East-West traffic, configure High Availability (HA), and configure Equal Cost Multipathing (ECMP).
- [Module 6 -NSX Services](#) (15 minutes) (Intermediate) In this module you will explore NSX service as VPN, Load balancing and DHCP / IPAM.
- [Module 7 - NSX operations overview](#) (15 minutes) (Intermediate) In this module you will explore basic topology management, flow tracing and operational functions.

How to End Lab



To end your lab click on the **END** button.

Module 6 - NSX Services (15 minutes)

NSX Edge Services Overview

The goal of this lab is to explore some of the various services available in NSX. In this module you will complete the following tasks:

- Create an NSX Edge Cluster to provide fault tolerance to your services
- Create and configure Load Balancing in NSX
- Create and configure IPAM in NSX via DHCP Server

NSX Edge Transport Node Overview

NSX Edge Nodes are service appliances with pools of capacity, dedicated to running network and security services in the NSX fabric that cannot be distributed to the hypervisors. Edge Nodes are used to provide routed connectivity between the overlay and the physical infrastructure via the Service Router (SR) component of the Tier-0 Gateway, and can also provide additional centralized, non-distributed services such as load balancing, NAT and VPN. Services provided by the NSX Edge Transport Node include:

- Connectivity to Physical Infrastructure
- NAT
- DHCP Server / Relay
- Gateway (Layer 3) Firewall
- Load Balancer
- L2 Bridging (Overlay to Physical/VLAN)
- VPN

As soon as one of these services is configured or an external interface is defined on a Tier-0 or Tier-1 gateway, a Service Router (SR) is instantiated on the selected Edge node. The Edge node is also a transport node in NSX, hosting its own TEP address. This allows it to communicate with other nodes in the overlay network. NSX Edge Transport Nodes are typically configured for one or more Overlay Transport Zones, and will also be connected to one or more VLAN transport zones when used for North-South (Uplink) connectivity.

Beginning with the NSX-T Data Center 3.0 release, support for the Intel® QuickAssist Technology (QAT) is provided on bare metal servers. Intel® QAT provides the hardware acceleration for various cryptography operations, such as IPsec VPN bulk cryptography, offloading the function from the Intel® Xeon® Scalable processor.

The QAT feature is enabled by default if the NSX Edge is deployed on a bare metal server with an Intel® QuickAssist PCIe card that is based on the installed C62x chipset (Intel® QuickAssist Adapter 8960 or 8970). The single root I/O virtualization (SR-IOV) interface must be enabled in the BIOS firmware

NSX Edge Transport Node Form Factors

NSX Edge Node is available for deployment in either a virtual machine (VM) or bare metal form factor. When deployed as a VM, the Edge Node benefits from native vSphere features such as Distributed Resource Scheduler (DRS) and vMotion. Deploying the Edge Node on bare metal allows direct access to the device's hardware resources, providing increased performance and lower latency than the VM form factor.

- 2nd Gen Intel® Xeon® Scalable processors, with Intel® Virtualization Technology (Intel® VT), built into and enhanced in five successive generations of Intel® Xeon® processors, enables live migration of VMs across Intel Xeon processor generations.


Consider the network bandwidth requirements within your data center when planning vMotion. A 10 GbE NIC can vMotion up to 8 VMs simultaneously.

Launch Google Chrome



- If not continuing from a previous module, open a browser by double clicking the **Google Chrome** icon on the desktop.

Launch NSX Manager in a New Tab



1. Click on the **new tab** icon

2. Click on the **Region A** folder
3. Click on **nsxmgr-01a**

NSX Manager Login

VMware® NSX-T™

admin 1


..... 2

LOG IN 3

Enter the following login credentials for the NSX Manager.

1. In the Username field type **admin**
2. In the Password field type **VMware1!VMware1!**
3. Click **LOG IN**

NSX Edge Transport Nodes and Clusters



When deploying centralized services in the NSX fabric, the instance of that service is provisioned and realized on an NSX Edge Node. If the Edge Node hosting this service were to experience a failure, any services running on the Edge Node would also fail as a result. To prevent a failure from impacting these services, Edge Nodes are grouped into logical objects called Edge Clusters.

An Edge Cluster is a group of one or more Edge Nodes that specifies the fault domain for services and how they should be recovered. Your lab is provisioned with four Edge Transport Nodes: **nsx-edge-01**, **nsx-edge-02**, **nsx-edge-03** and **nsx-edge-04**. We will now review the existing Edge Cluster configuration and create one additional cluster.

1. Click **System** in the top level navigation menu
2. Click **Fabric** in the menu on the left side of the NSX System user interface
3. Click **Nodes** in the Fabric submenu
4. Click **Edge Clusters**

Observe that one Edge Cluster is currently configured: **edge-cluster-01**. We will review its configuration before creating a second cluster.

5. Click the **checkbox** to the left of edge-cluster-01 to select it
6. Click **EDIT**

View edge-cluster-01 Configuration

Edit Edge Cluster - edge-cluster-01 ⓘ ×

Name* edge-cluster-01

Description

Edge Cluster Profile nsx-default-edge-high-availability-profile × ▾

Transport Nodes

Member Type Edge Node ▾

☐ Available (2)		☐ Selected (2)
☐ nsx-edge-03	➤	☐ nsx-edge-01 ←
☐ nsx-edge-04	➤	☐ nsx-edge-02 ←


< BACK NEXT > 1 - 2 of 2 records

1 CANCEL SAVE

Observe that as stated above, there are four Edge Nodes. Edge Cluster **edge-cluster-01** is configured to use Edge Nodes **nsx-edge-01** and **nsx-edge-02**, indicated in the Selected column. nsx-edge-03 and nsx-edge-04 are displayed as Available and are not part of this Edge Cluster.

1. Click **CANCEL** to return to the list of Edge Clusters

Create New Edge Cluster



The screenshot shows the VMware NSX Manager interface. The top navigation bar includes Home, Networking, Security, Inventory, Plan & Troubleshoot, and System. The left sidebar shows System Overview, Configuration, Appliances, Get Started, Fabric, Nodes, Profiles, and Transport Zones. The main content area is titled 'Edge Clusters' and includes a table with columns: Edge Cluster, ID, Member Type, Cluster Profile, and Edge Transport Nodes. A red circle highlights the '+ ADD' button, and a black circle with the number '1' is next to it. The table contains one entry: 'edge-cluster-01' with ID 'a97e...91c8', Member Type 'Edge Node', Cluster Profile 'nsx-default-edge-high-avalla...', and 2 Edge Transport Nodes.

Edge Cluster	ID	Member Type	Cluster Profile	Edge Transport Nodes
edge-cluster-01	a97e...91c8	Edge Node	nsx-default-edge-high-avalla...	2

We will now define a new Edge Cluster that will be used for the services we configure in this module.

- Click **ADD**

Add Edge Cluster

Add Edge Cluster ? ×

Name* **1**

Description

Edge Cluster Profile nsx-default-edge-high-availability-profile × ▼

Transport Nodes

2 Other Type Edge Node ▼

3 ☒ Available (2)	☐ Selected (0)
☒ nsx-edge-03	No records found
☒ nsx-edge-04	

4 CANCEL ADD

1. In the Name field, enter **edge-cluster-02**
2. Click the **checkbox** to the left of Available to select all available Edge Nodes (nsx-edge-03 and nsx-edge-04)
3. Click the **right arrow icon** to move the Edge Nodes from Available to the Selected column
4. Click **ADD**

Confirm New Edge Cluster

Host Transport Nodes Edge Transport Nodes Edge Clusters ESXi Bridge Clusters NCP Clusters					
+ ADD EDIT DELETE ACTIONS ▾					Search
☐	Edge Cluster	ID	Member Type	Cluster Profile	Edge Transport Nodes
☐	edge-cluster-01	a97e...91c8	Edge Node	nsx-default-edge-high-availa...	2
☐	edge-cluster-02	6b75...a84b	Edge Node	nsx-default-edge-high-availa...	2


Confirm that Edge Cluster **edge-cluster-02** was created successfully, and that it is comprised of **2** Edge Transport Nodes.

NSX Edge Services - Load Balancing

We will now create a load balancer in NSX. In this section of the module you will execute the following tasks:

- Create a new load balancer
- Create health checks for HTTPS services on web-01a and web-02a web servers
- Create a virtual IP (VIP) to load balance web server traffic to 2 separate web servers
- Test the load balancing services

Create New Tier-1 Gateway



In order to create a Load balancer we need a Tier-1 Gateway deployed to at least one edge gateway.

1. Click **Networking** in the top level navigation menu
2. Click **Tier-1 Gateways** in the menu on the left side of the NSX Networking user interface
3. Click **ADD TIER-1 GATEWAY**

Add Tier-1 Gateway

The screenshot shows the 'Tier-1 Gateways' configuration interface. It includes a table with columns for 'Tier-1 Gateway Name', 'Linked Tier-0 Gateway', '#Linked Segments', 'Status', and 'Alarms'. Below the table, there are configuration fields for 'Edge Cluster', 'Fail Over', 'Edges Pool Allocation Size', 'Tags', 'Route Advertisement', and various toggle switches for 'All Static Routes', 'All DNS Forwarder Routes', 'All Connected Segments & Service Ports', 'All IPsec Local Endpoints', 'All NAT IP's', 'All LB VIP Routes', and 'All LB SNAT IP Routes'. A 'SAVE' button is at the bottom left.

Numbered callouts in the image indicate the following steps:

1. Enter **Tier-1-LB-01** for Tier-1 Gateway Name
2. Select **Tier-0-gateway-01** for Linked Tier-0 Gateway
3. Select **edge-cluster-02** for Edge Cluster
4. Click the arrow to the left of **Route Advertisement** to expand it
5. Click the toggle to enable **All LB VIP Routes**
6. Click **SAVE**

1. Enter **Tier-1-LB-01** for Tier-1 Gateway Name
2. Select **Tier-0-gateway-01** for Linked Tier-0 Gateway
3. Select **edge-cluster-02** for Edge Cluster

Note: Tier-1 gateways used for Load balancing services must be placed on edge nodes of medium or large size

4. Click the arrow to the left of **Route Advertisement** to expand it
5. Click the toggle to enable **All LB VIP Routes**
6. Click **SAVE**

End Configuring Tier-1 Gateway

Tier-1 Gateway Tier-1-LB-01 is successfully created.
Want to continue configuring this Tier-1 Gateway?

YES

NO

1

1. Click **NO** to end editing the Tier-1 Gateway

Edit Tier-0 Gateway

HomeNetworkingSecurityInventoryPlan & TroubleshootSystem

Network Overview

Network Topology

Connectivity1

Tier-0 Gateways

Tier-1 Gateways

Segments

Network Services

Tier-0 Gateways

ADD GATEWAYEXPAND ALLFilter by Name, Path and more

	Tier-0 Gateway Name	HA Mode	Linked Tier-1 Gateways	Linked Segments	Status ⓘ	Alarms
⋮ 2	Tier-0-gateway-01	Active Active	0	4	Success ⓘ	0

1. Click **Tier-0 Gateways** in the menu on the left side of the NSX Networking user interface
2. Click the **More Options** icon to the left of **Tier-0-Gateway-01** to display its Options menu, then click **Edit**

Edit Tier-0 Route Re-distribution

Tier-0 Gateway Name	HA Mode	Linked Tier-1 Gateways	Linked Segments	Status ⓘ	Alarms
Tier-0-gateway-C *	Active Active ▾ *				
IP Address Management No Dynamic IP Allocation					
Edge Cluster edge-cluster-01 ⓧ ▾					
Additional Settings					
Route Distinguisher for VRF Gateways					
EVPN Settings					
Tags					
Tag (Re) ▾ Scope (I) ▾ +					
Max 30 allowed. Click (+) to save.					
INTERFACES					
ROUTING					
MULTICAST					
BGP					
1 ▾ ROUTE RE-DISTRIBUTION					
Route Re-distribution 3 2					
Route Re-distribution Status ☒ On					
CLOSE EDITING					

1. Click the arrow to the left of **ROUTE RE-DISTRIBUTION** to expand it
2. Click the **3** link to the right of Route Re-distribution to display the Set Route Re-distribution dialog

Set Route Re-distribution

Set Route Re-distribution

Tier-0 Gateways Tier-O-gatew... #Route Re-distribution 3

ADD ROUTE RE-DISTRIBUTION

1	Name	Route Re-distribution	Route Map
	T1-LB	1	Not Set
	T-1 NAT	1	Not Set
	T-0 and T-1 Connected segments	2	Not Set

CANCEL APPLY

Since we rely on the Tier-0 Gateway to re-distribute the routes from the Tier-1 to the physical fabric, we also need to allow the Load Balancer and SNAT routes to be re-distributed at the Tier-0 level as well. This has already been configured in your lab, but would typically need to be enabled. We will now review the route redistribution configuration.

1. Click the **More Options** icon to the left of **T1-LB** to display its Options menu, then click **Edit**

Select Route Re-distribution

Set Route Re-distribution

Tier-0 Gateways Tier-0-gatew... #Route Re-distribution 3

ADD ROUTE RE-DISTRIBUTION

Search

Name	Route Re-distribution	Route Map
T1-LB	1*	Select Route Map
ADD CANCEL		
T-1 NAT	1	Not Set
T-0 and T-1 Connected segments	2	Not Set

CANCEL APPLY

1. Click the **1** link under Route Re-distribution to display the Set Route Re-distribution dialog

Set Route Re-distribution Options

Set Route Re-distribution

Tier-O Gateways

Tier-O-gatew...

#Selected Sources 1

Select sources below

Tier-O Subnets

☐ Static Routes

☐ IPsec Local IP

☐ EVPN TEP IP

☐ Connected Interfaces & Segments

☐ Service Interface Subnet

☐ Loopback Interface Subnet

☐ NAT IP

☐ DNS Forwarder IP

☐ External Interface Subnet

☐ Connected Segment

Advertised Tier-1 Subnets

☐ DNS Forwarder IP

☒ LB VIP

☐ LB SNAT IP

☐ Connected Interfaces & Segments

☐ Service Interface Subnet

☐ Static Routes

☐ NAT IP

☐ IPsec Local Endpoint

☐ Connected Segment

2

CANCEL

APPLY

1. Observe that **LB VIP** is selected
2. Click **CANCEL**

Cancel Route Re-distribution Changes

Set Route Re-distribution

Tier-0 Gateways Tier-0-gatew... #Route Re-distribution 3

ADD ROUTE RE-DISTRIBUTION

Q Search

Name	Route Re-distribution	Route Map
T1-LB ADD CANCEL 1	1*	Select Route Map
T-1 NAT	1	Not Set
T-0 and T-1 Connected segments	2	Not Set

2 CANCEL APPLY

- 1. Click **CANCEL** to cancel editing of the T1-LB Re-distribution
- 2. Click **CANCEL** to close the Set Route Re-distribution dialog without making any changes

Close Editing

Route Re-distribution

Status

1 CLOSE EDITING

- 1. Click **CLOSE EDITING**

Create New Load Balancer

The screenshot shows the NSX Networking interface for creating a new Load Balancer. The left sidebar has 'Load Balancing' selected. The top navigation bar shows 'LOAD BALANCERS' as the active tab. The main form area contains the following fields and controls:

- Name:** LB-01
- Type:** Server Load Balancer
- Size:** Small
- Attachment:** Tier-1-LB-01
- Description:** Enter Description
- Error Log Level:** Info
- Tags:** Tag (Requ) Scope (Op) Max 30 allowed. Click (+) to save.
- Admin State:** Enabled
- NOTE:** Before further configurations can be done, fill out mandatory fields above (*), click 'Save' below.
- VIRTUAL SERVERS:** (Link to view virtual servers)
- Buttons:** SAVE, CANCEL

Now that we have the T-1 and routing requirements set up let's create our Load Balancer.

1. Click **Load Balancing** in the menu on the left side of the NSX Networking user interface
2. Click on **ADD LOAD BALANCER**
3. Enter **LB-01** for Name
4. Select **Small** for Size
5. Select our recently created **Tier-1-LB-01** gateway for Attachment
6. Click **SAVE**

Click no to continue


The screenshot shows a confirmation dialog box with the following text:

✓ Load Balancer LB-01 is successfully created.
Want to continue configuring this Load Balancer?

Buttons: YES | NO (highlighted with a red circle and a '1' callout)


1. Click **NO** to continue

Wait for Successful Configuration



1. Click the **refresh** icon periodically until the Status shows Success
 - **NOTE:** It may take up to 3 - 4 minutes for the Status to display Success. During this time you may see the status transition through other states, including Failed and Unknown. This is normal, and occurs while Policy Manager attempts to realize the desired configuration on the NSX Manager.

Create Health Monitor



Lets create a health monitor for our 3 Tier App

1. Click **MONITORS**
2. Click **ADD ACTIVE MONITOR**
3. Click **HTTPS** in the drop down menu

Add Health Monitor

Name	Protocol	Monitoring Port	Monitoring Interval	Timeout Period (seconds)	Server Pools
1 Webapp-monitor	HTTPS	2 443	5	5	
Description		Enter Description		Fall Count	3
Tags		Tag (ReqL) Scope (Of) +		Rise Count	3
Max 30 allowed. Click (+) to save.					
Additional Properties					
HTTP Request		Configure 3		HTTP Response Configure	
SSL Configuration		Configure			
SAVE CANCEL					

1. Enter **Webapp-monitor** for Name
2. Enter **443** for Monitoring Port
3. Click the **Configure** link to the right of HTTP Request

Set HTTP Request Configuration

HTTP Request and Response Configuration

Active Health Monitor - Webapp-monitor

HTTP Request Configuration **HTTP Response Configuration** 3

HTTP Method Get

HTTP Request URL 1 /cgi-bin/app.py

HTTP Request Version 2 1.1

[ADD](#)

Header Name	Header Value
 No found.	

HTTP Request Body

[CANCEL](#) [APPLY](#)

1. Enter **/cgi-bin/app.py** for HTTP Request URL
2. Select **1.1** for the HTTP Request Version
3. Click the **HTTP Response Configuration** tab

Set HTTP Response Configuration

HTTP Request and Response Configuration

Active Health Monitor -

HTTP Request Configuration **HTTP Response Configuration**

HTTP Response Code 1 or more response codes

HTTP Response Body **1** Customer Database

2 CANCEL APPLY

1. Enter **Customer Database** for HTTP Response Body
2. Click **APPLY**

Save Health Monitor

Name	Protocol	Monitoring Port	Monitoring Interval	Timeout Period (seconds)	Server Pools
Webapp-monitor *	HTTPS	443	5	5	
Description		Enter Description		Fall Count	3
Tags		Tag (ReqL) Scope (O) +		Rise Count	3
Max 30 allowed. Click (+) to save.					
Additional Properties HTTP Request Configure HTTP Response Configure SSL Configuration Configure					
1 SAVE CANCEL					

1. Click **SAVE**

Create New Server Pool

LOAD BALANCERS	VIRTUAL SERVERS	SERVER POOLS	PROFILES	MONITORS
ADD SERVER POOL 2 1 COLLAPSE ALL Filter by Name, Path and more				
Name	Algorithm	Members/Group	Virtual Servers	Status
3 Web-servers-01 *	Round Robin	Select Members 4		
Description		Enter Description		
SNAT Translation Mode		Automap		
Additional Properties Active Monitor Set Monitors				
SAVE CANCEL				

In this step we will create a new Server Pool. A Server Pool is a list of the systems the load balancer will monitor and deliver traffic to for a given Virtual Server.

1. Click **SERVER POOLS**
2. Click **ADD SERVER POOL**
3. Enter **Web-servers-01** for pool Name
4. Click the **Select Members** link to display the Configure Server Pool Members dialog

Add Pool Member

Configure Server Pool Members

Server Pool -

☒ Enter individual members ☐ Select a group

ADD MEMBER

Name	IP	Port	Weight	State	Backup Member	Max Concurrent Connections
web-01a	172.16.10.11 e.g. 10.10.10.10 or fc7e:f206:db42::	443 e.g. 80, 443	1	Enabled	☒ Disabled	

SAVE CANCEL

1 - 1 of 1 Members

CANCEL **APPLY**

1. Click **ADD MEMBER**
2. Enter **web-01a** for Name
3. Enter **172.16.10.11** for IP
4. Enter **443** for Port
5. Click **SAVE**

Add Second Pool Member

Configure Server Pool Members

Server Pool -

☒ Enter individual members
 ☐ Select a group

ADD MEMBER Search

Name	IP	Port	Weight	State	Backup Member	Max Concurrent Connections
web-02a *	172.16.10.12 * e.g. 10.10.10.10 or fc7e:f206:db42::	443 e.g. 80, 443	1	Enabled	☐ Disabled	
SAVE CANCEL						
web-01a	172.16.10.11	443	1	Enabled	☒ Disabled	

1 - 2 of 2 Members

CANCEL **APPLY**

1. Click **ADD MEMBER**
2. Enter **web-02a** for Name
3. Enter **172.16.10.12** for IP
4. Enter **443** for Port
5. Click **SAVE**
6. Click **APPLY**

Select Health Monitor

Name	Algorithm	Members/Group	Virtual Servers	Status	Alarms
Web-servers-01 *	Round Robin	2			
Description Enter Description		Active Monitor Set Monitors 1			
SNAT Translation Mode Automap					
> Additional Properties					
SAVE CANCEL					

We will now select a Health Monitor for this Pool. Health Monitors define how the load balancer will check the pool members to determine their ability to accept incoming connections.

- 1. Click the **Set Monitors** link to display the Select Active Monitors dialog

Select Active Monitors

Select Active Monitors

Server Pool - Your New Server Pool

Webapp-monitor X

ADD ACTIVE MONITOR EXPAND ALL Filter by Name, Path and more

	Name	Protocol	Monitoring Port	Monitoring Interval	Timeout Period (seconds)	Server Pools
☐	default-http-lb-monitor	HTTP	80	5	5	0
☐	default-https-lb-monitor	HTTPS	443	5	5	0
☐	default-icmp-lb-monitor	ICMP		5	5	0
☐	default-tcp-lb-monitor	TCP		5	5	0
☒	Webapp-monitor	HTTPS	443	5	5	0

☒ 1 REFRESH 1 - 5 of 5 Active Monitors

CANCEL

2

APPLY

- 1. Click the **checkbox** to the left of Webapp-monitor to select it
- 2. Click **APPLY**

Save Server Pool

Name	Algorithm	Members/Group	Virtual Servers	Status	Alarms
Web-servers-01 *	Round Robin	2			

Description

Enter Description

Active Monitor

1 Monitor

SNAT Translation Mode

Automap

> Additional Properties


1

SAVE

CANCEL

1. Click **SAVE**


Create virtual server



The last step is to define a Virtual Server. The Virtual Server has an IP address that accepts incoming connections and routes them to a pool member. How a pool member is chosen gets specified during configuration, and can be based on a number of factors including availability; load; and number of connections.

1. Click **VIRTUAL SERVERS**
2. Click **ADD VIRTUAL SERVER**
3. Click **L4 TCP** in the drop down menu

Configure the Virtual Server




Do the following to configure the virtual server settings

1. Enter **Webapp-VS** for Name
2. Enter **172.16.10.10** for IP Address (This IP address has a DNS record of webapp.corp.local)
3. Enter **443** for Ports
4. Click the **Add Item(s): 443** link that appears, to add port 443 to the list of desired ports

5. Select the **LB-01** Load Balancer instance that was created earlier in this module
6. Select the **Web-servers-01** Server Pool that was created earlier in this module
7. Click **SAVE**

Confirm Successful Virtual Server Configuration




	Name	IP Address	Ports	Type	Load Balancer	Server Pool	Status	Alarms
	Webapp-VS	172.16.10.10	443	L4 TCP	LB-01	Web-servers-01	Success	0

Our new load balancer configuration is now complete. We have configured a Virtual Server on port **443** with an IP address of **172.16.10.10**, sending traffic to the web servers in Server Pool **Web-servers-01**. This server pool has two IP addresses, **172.16.10.11** and **172.16.10.12**, corresponding to web-01a and web-02a respectively. The Load Balancer is monitoring the health and availability of the web application by connecting to the pool members every five seconds with the URL **/cgi-bin/app.py** and expecting a response that contains the text **"Customer Database"**.

1. Click the **refresh** icon periodically until the Status shows Success
 - **NOTE:** It may take up to 60 seconds for the Status to display Success

Test the New Virtual Server



1. Click to open a **new tab** in Chrome
2. Click the **3-tier-app** bookmark folder to expand it

- Click the **Webapp VIP (Will not work until load balancer config)** shortcut from the drop down


Verify the Web Application

Customer Database Access			
Accessed via: web-01a		1	
Name Filter (blank for all records):			Apply
Rank	Name	Universe	Revenue
1	CHOAM	Dune	\$1.7 trillion
2	Acme Corp.	Looney Tunes	\$348.7 billion
3	Sirius Cybernetics Corp.	Hitchhiker's Guide	\$327.2 billion
4	Buy n Large	Wall-E	\$291.8 billion
5	Aperture Science, Inc.	Valve	\$163.4 billion
6	SPECTRE	007	\$157.1 billion
7	Very Big Corp. of America	Monty Python	\$146.6 billion
8	Frobozz Magic Co.	Zork	\$112.9 billion
9	Warbucks Industries	Lil' Orphan Annie	\$61.5 billion
10	Tyrell Corp.	Bladerunner	\$59.4 billion
11	Wayne Enterprises	Batman	\$31.3 billion

You should now see the test web application and the Customer Database information.

- Note the web server** your browser connected to


Connect to vCenter



We will now log into the vCenter web client so we can manually fail one of the web servers and test fault tolerance. If you have an existing tab with the vCenter client, click to select it. Otherwise perform the following to open a new tab and connect to vCenter.


1. Click to open a **new tab** in Chrome
2. Click the **Region A** bookmark folder to expand it
3. Click the **vcsa-01a Web Client** shortcut from the drop down

Login to vCenter




1. Click the **checkbox** to the left of Use Windows session authentication
2. Click **LOGIN**

Select Hosts and Clusters View



1. Click the **Menu** link to display the vCenter menu
2. Click **Hosts and Clusters**


Power Off the Web Server You Connected To



Recall the web server that served your request when we recently connected to the Webapp Virtual Server. We will now power that server off and verify that traffic gets directed to the remaining web server.

1. Click to select the **web server** you noted down in the previous step
2. Click **Actions**
3. Click **Power**
4. Click **Shut Down Guest OS**

Confirm Guest Shut Down



1. Click **YES** to confirm Shut Down of Guest OS

Verify the Web Application

Customer Database Access


Accessed via: web-02a

Name Filter (blank for all records):

Rank	Name	Universe	Revenue
1	CHOAM	Dune	\$1.7 trillion
2	Acme Corp.	Looney Tunes	\$348.7 billion
3	Sirius Cybernetics Corp.	Hitchhiker's Guide	\$327.2 billion
4	Buy n Large	Wall-E	\$291.8 billion
5	Aperture Science, Inc.	Valve	\$163.4 billion
6	SPECTRE	007	\$157.1 billion
7	Very Big Corp. of America	Monty Python	\$146.6 billion

1. Click the **Customer Database tab** to change focus back to the web application
2. Click the **Refresh** icon
3. Verify that you are now connecting to the **remaining web server pool member**
 - **Note:** This should be the opposite of the server you initially connected to

Power On the Web Server



Before completing this section of the module, we will return the web server we had powered off to an operational state.

1. Click the **web server** that was recently powered off to ensure it is selected
2. Click the **Power On** icon

NSX Edge Services - DHCP

We will now explore IP Address Management (IPAM) via the DHCP Server component of NSX. In this section of the module you will execute the following tasks:

- Create a new DHCP Server
- Attach the DHCP Server to an NSX Segment
- Configure a virtual machine to leverage DHCP services
- Test DHCP, then restore the server's static IP address configuration

Create New DHCP Server

The screenshot shows the NSX Policy Manager interface. The top navigation bar has 'Networking' selected. The left sidebar has 'DHCP' selected. The main content area shows the 'DHCP' configuration page. The 'ADD DHCP PROFILE' button is highlighted. The form below it has the following fields:


Profile Name	Profile Type	Server IP Address	Lease Time (seconds)	Where Used
DHCP-01	DHCP Server	Enter IP Addresses CIDR e.g. 10.22.12.2/23 or IPv6 fc7e:1206:db42::1/48	86400 Should be between 60 and 4294967295	0

Below the table, there are fields for 'Edge Cluster' (Select Edge Cluster), 'Edges' (Set), and 'Tags' (Tag (Re) Scope (t) Max 30 allowed. Click (+) to save.). At the bottom of the form, there are 'SAVE' and 'CANCEL' buttons.

The first step in configuring our DHCP Server is to define it in the NSX Policy Manager interface.

1. Click **Networking** in the top level navigation menu
2. Click **DHCP** in the menu on the left side of the NSX Networking user interface
3. Click **ADD DHCP PROFILE**
4. Enter **DHCP-01** for Profile Name
5. Click **SAVE**


Configure Tier-0 Gateway for DHCP Service



The next step in configuring our DHCP Server is to attach it to a Gateway. In this exercise we will attach it to the existing T0 Gateway Tier-0-Gateway-01.

1. Click Tier-0 Gateways in the menu on the left side of the NSX Networking user interface
2. Click the **More Options** icon to the left of **Tier-0-Gateway-01** to display its Options menu, then click Edit

Configure IP Address Management



1. Click the **No Dynamic IP Allocation** link to the right of IP Address Management to display the Set IP Address Management dialog

Set IP Address Management

Set IP Address Management

×

Choose either local or remote DHCP or No Dynamic IP Allocation.

Type

DHCP Server

DHCP Server

DHCP-01

Server Details

Server Name	DHCP-01
Lease Time	86400 seconds
Server Address	100.96.0.1/30

CANCEL

SAVE

1. Select **DHCP Server** for Type
2. Select **DHCP-01** for DHCP Server
3. Click **SAVE**

Save Tier-0 Gateway Configuration

Tier-0 Gateway Name	HA Mode	Linked Tier-1 Gateways	Linked Segments	Status ⓘ	Alarms
Tier-0-gateway-C *	Active Active *				
IP Address Management Local 1 Servers					
Edge Cluster edge-cluster-01 ⓧ					
> Additional Settings > Route Distinguisher for VRF Gateways > EVPN Settings					
Tags Tag (Re) ⌵ Scope (t) ⌵ + Max 30 allowed. Click (+) to save.					
1 SAVE CANCEL Unsaved Changes					
> INTERFACES > ROUTING > MULTICAST > BGP > ROUTE RE-DISTRIBUTION					
2 CLOSE EDITING					

1. Click **SAVE**
2. Click **CLOSE EDITING**

Configure Segment LS-web for DHCP

Home

Networking

Security

Inventory

Plan & Troubleshoot

System

<<

Network Overview

Network Topology

Connectivity

Tier-0 Gateways

Tier-1 Gateways

Segments

Network Services

VPN

NAT

Load Balancing

SEGMENTS

SEGMENT PROFILES

EDGE BRIDGE PROFILES

METADATA PROXIES

ADD SEGMENT

EXPAND ALL

Filter by Name, Path and more

Segment Name	Connectivity	Transport Zone	Subnets	Ports	Admin State	Status	Alarms
> k8s-mgmt	Tier-0-gateway-01 Tier0	TZ-Overlay-1 Overlay	10.0.11/24	3	Up	Success	0
> k8s-transport	None	TZ-Overlay-1 Overlay	Not Set	0	Up	Success	0
> LS-app	Tier-1-gateway-01 Tier1	TZ-Overlay-1 Overlay	172.16.20.1/24	1	Up	Success	0
> LS-db	Tier-1-gateway-01 Tier1	TZ-Overlay-1 Overlay	172.16.30.1/24	1	Up	Success	0
> LS-web	Tier-1-gateway-01 Tier1	TZ-Overlay-1 Overlay	172.16.10.1/24	3	Up	Success	0
> Uplink	None	TZ-Uplink VLAN	Not Set	0	Up	Success	0

The final step in configuring DHCP services is to attach our new DHCP Server to the desired Segment. In this example we will use the existing Segment LS-web.

1. Click **Segments** in the menu on the left side of the NSX Networking user interface
2. Click the **More Options** icon to the left of **LS-web** to display its Options menu, then click Edit

Edit Segment LS-web

Segment Name	Connectivity	Transport Zone	Subnets	Ports	Admin State	Status	Alarms
LS-web	Tier-0-gat...	TZ-Overlay-1	172.16.10.1/24 CIDR e.g. 10.22.12.2/23 Gateway CID CIDR e.g. fc7e:f206:db42::1/48	3	☒	i	

Segment needs to have either Subnets or VPN defined, or both.

L2 VPN: You have no L2 VPN sessions for this Gateway. For that, go to [VPN Services](#) . Note that for L2 sessions to work, you

VPN Tunnel ID:

VLAN: Enter 1 list of VLA Ne

Unlink Teaming Policy: Select 1 Unlink Teaming Policy

1. Select **Tier-0-Gateway-01** for Connectivity
2. Click **SET DHCP CONFIG** to display the DHCP Options dialog

Set DHCP Config

The screenshot shows the 'Set DHCP Config' dialog box for the 'LS-web' segment. The configuration includes:

- Segment:** LS-web
- IPv4 Gateway:** 172.16.10.1/24
- IPv6 Gateway:** Not Set
- DHCP Type:** Gateway DHCP
- DHCP Profile:** DHCP-01
- Message:** IPv6 server settings are not supported for Gateway DHCP
- IPv4 Server:** Selected tab
- DHCP Config:** Enabled (Callout 1 points to the toggle switch)
- DHCP Server Address:** 100.96.0.1/30
- DHCP Ranges:** 99 Maximum | Format 172.16.14.10-172.16.14.100 or 172.16.14.0/24. The range **172.16.10.100-172.16.10.105** is entered (Callout 2 points to the input field).
- Add Item(s):** 172.16.10.100-172.16.10.105 (Callout 3 points to the 'Add Item(s)' link)
- Lease Time (seconds):** Default value is 86400
- DNS Servers:** Enter IP Addresses (e.g. 10.10.10.10)
- Buttons:** CANCEL and APPLY (Callout 4 points to the APPLY button)


In this step, we will enable DHCP for the Segment and provide a range of available IP addresses for the DHCP Server to use.

1. Click the toggle to the right of DHCP Config to set it to **Enabled**
2. Enter **172.16.10.100-172.16.10.105** for DHCP Ranges
3. Click the **Add Item(s): 172.16.10.100-172.16.10.105** link that appears, to add the provided DHCP range
4. Click **APPLY**

Save Segment Configuration

1. Click **SAVE**
2. Click **CLOSE EDITING**


Connect to vCenter



We will now log into the vCenter web client so we can manually fail one of the web servers and test fault tolerance. If you have an existing tab with the vCenter client, click to select it. Otherwise perform the following to open a new tab and connect to vCenter.

1. Click to open a **new tab** in Chrome
2. Click the **Region A** bookmark folder to expand it
3. Click the **vcsa-01a Web Client** shortcut from the drop down

Login to vCenter



VMware® vSphere

administrator@corp.local

1


☒ Use Windows session authentication

2

LOGIN


1. Click the **checkbox** to the left of Use Windows session authentication
2. Click **LOGIN**

Select Hosts and Clusters View



1. Click the **Menu** link to display the vCenter menu
2. Click **Hosts and Clusters**


Establish Console Connection to web-01a



We will now reconfigure server web-01a to obtain a DHCP address from the NSX DHCP Server we just configured.


1. Click **web-01a** to select it
2. Click the **Launch Console** icon

Launch Console



1. Click **OK** to launch a Web Console for server web-01a

Configure DHCP on web-01a



1. Enter **root** at the login: prompt for the username and press Enter
2. Enter **VMware1!** at the Password: prompt and press Enter. The characters you enter will not be displayed
3. Enter **ifconfig -a** and press Enter
4. Observe that server web-01 has a static IP address of 172.16.10.11

We will now reconfigure the server to request a DHCP address.

5. Enter the following text and press Enter:

```
/root/2126-01/dhcp.sh
```

The web server will reboot and obtain an IP address via DHCP.

Confirm DHCP Configuration on web-01a

```

web-01a
Welcome to Photon 3.0 (x86_64) - Kernel 4.19.79-1.ph3-esx (tty1)
Hint: Num Lock on
1 web-01a login: root
2 Password:
3 root@web-01a:~# ifconfig -a
eth0: Link encap:Ethernet HWaddr 00:50:56:00:a5:b3
      inet addr:172.16.10.100 Bcast:172.16.10.255 Mask:255.255.255.0
      inet6 addr: fe80::250:56ff:fe9c:a5b3 Scope:Link
      UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
      RX packets:66 errors:0 dropped:0 overruns:0 frame:0
      TX packets:26 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:1000
      RX bytes:5298 (5.2 KB) TX bytes:2704 (2.7 KB)

lo: Link encap:Local Loopback
      inet addr:127.0.0.1 Mask:255.0.0.0
      inet6 addr: ::1/128 Scope:Host
      UP LOOPBACK RUNNING MTU:65536 Metric:1
      RX packets:90 errors:0 dropped:0 overruns:0 frame:0
      TX packets:90 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:1000
      RX bytes:6792 (6.7 KB) TX bytes:6792 (6.7 KB)

root@web-01a [ ~ ]# /root/2126-01/static.sh

```

1. Enter **root** at the login: prompt for the username and press Enter
2. Enter **VMware1!** at the Password: prompt and press Enter. The characters you enter will not be displayed
3. Enter **ifconfig -a** and press Enter
4. Observe that server web-01 has now obtained an IP address from our NSX DHCP Server, and has been assigned an IP address of 172.16.10.100.

We will now return the web server to using its previously configured static IP address.

5. Enter the following text and press Enter:

```
/root/2126-01/static.sh
```

The web server will reboot with its original IP address of 172.16.10.11.


Module 6 Conclusion

Congratulations on completing Module 6.

Please proceed to any module below which interests you the most.

- [Module 1 - Introduction to NSX-T](#) (15 minutes) In this module you will be introduced to NSX-T, its capabilities, and the components that make up the solution.
- [Module 2 - Host and Edge Transport Node Preparation](#) (30 minutes) In this module, you will review the prepared components ready for NSX workloads.
- [Module 3 - Logical Switching](#) (15 minutes) In this module, you will review the NSX lab fabric, configure Logical Switches, and attach them to virtual machines. You will be introduced to the use of the GENEVE protocol for overlay networking.
- [Module 4 - Logical Routing](#) (30 minutes) In this module you will learn how to configure Logical Routing in NSX for North-South and East-West traffic, configure High Availability (HA), and configure Equal Cost Multipathing (ECMP).
- [Module 5 - Distributed Firewall and Tools](#) (30 minutes) In this module we will demonstrate how the distributed firewall (DFW) and operational tools within NSX function and are configured.
- [Module 7 - NSX Operations Overview](#) (30 Minutes) In this module you will explore basic topology management, flow tracing and operational functions.

How to End Lab



To end your lab click on the **END** button.

Module 7: NSX Operations Overview (15 minutes)

NSX operations overview - Module Overview


The Goal of this lab is to explore basic topology management, flow tracing and operational functions.

NSX operations:

- Network Topology feature review
- Security overview configuration and capacity
- Inventory overview configuration and capacity
- System overview Configuration and Capacity
- Flow tracing

HOL-2126 Logical Diagram


This diagram illustrates the virtual machines that make up our 3 Tier Web App we will be exploring.



NSX operations overview


NSX provides a variety of tools to make operations and troubleshooting easier. In this module we will explore some of these tools.

Launch NSX-T Manager



1. **Click** to open a new tab
2. Click **Region A** to open the bookmark folder
3. Click **NSX-T shortcut** to launch log in page

Login to NSX-T Manager



VMware® NSX-T™

nsx-admin@corp.local


.....

LOG IN


1. Input **nsx-admin@corp.local** for the user name and **VMware1!** for the password
2. Click **Log in**

Note: This log in is utilizing the new Active Directory integration in version 3.0

Navigate to Network topology view




T-0, T-1 and connected segments



In this view we can see the relationship between configured T-0, T-1 gateways, logical segments and connected virtual machines.

1. **Note** the relationship between logical segments and connected virtual machines.
2. Click on the **T-0 gateway**

Tier-0-gateway-01



In this view we can see the name and details of the T-0 gateway

1. **Note** the Tier-0-gateway-01 name and configuration details
2. Click on the **BGP NEIGHBOR** link for more details

Review BGP Neighbor screen

BGP Neighbors

Tier-O Gateways Tier-O-gatew... #Neighbors

EXPAND ALL Search

	IP Address	BFD	Remote AS number	Route Filter	Allowas-in	Status
⋮ >	192.168.100.1	Disabled	65002	1	Disabled	Success 🔄 📄

1


REFRESH

2 BGP Neighbors

CLOSE


1. **Note** the Neighbor details
2. Click **Close** to return to the topology view

Explore the view



Take a moment to explore the topology, hover over items to see more detail, click on groups of VM's to expand and collapse the view.

Security Overview configuration



1. Click **Security**
2. Click **Security overview**
3. Click **Configuration**
4. Note all the detail on the page and the ability to click on links for further inspection / configuration. Feel free to click on links and explore.

Once you are done exploring Click **Security -> Security overview -> Configuration** to return to this view.

Security Overview Capacity

Limit	Current Inventory	Maximum Capacity
System-wide Firewall Sections	9	10,000
Active Directory Groups (Identity Firewall)	58	100,000
Distributed Firewall Sections	8	10,000
Active Directory Domains (Identity Firewall)	1	4
System-wide Firewall Rules	14	100,000
Distributed Firewall Rules	13	100,000
System-wide Endpoint Protection Enabled Hosts	0	256
System-wide Endpoint Protection Enabled Virtual Machines	0	7,500

1. Click **Security**
2. Click **Security overview**
3. Click **Capacity**
4. Note the current Security Inventory configuration and the maximum capacity, this is an easy way to keep an eye on sizing in the environment.

Inventory Overview Configuration

Inventory Overview

- Groups: 12
- Services: 411
- Virtual Machines: 15
- Context Profiles: 63
- Physical Servers: 0

GROUPS

- Used in Policies: 10
- Unused in Policies: 2

Identity Groups

- With Members: 12
- Without Members: -2
- Used in Policies: 2
- Unused in Policies: 0

1. Click **Inventory**
2. Click **Inventory overview**
3. Click **Configuration**
4. Note all the detail on the page and the ability to click on links for further inspection / configuration. Feel free to click on links and explore.

Once you are done exploring Click **Inventory -> Inventory overview - > Configuration** to return to this view.

Inventory Overview Capacity

Limit	Current Inventory	Maximum Capacity
Groups	12	10,000
Hypervisor Hosts	3	1,024
vSphere Clusters	2	64

Maximum capacity displayed above is for large appliance only.

1. Click **Inventory**
2. Click **Inventory overview**
3. Click **Capacity**
4. Note the current Inventory configuration and the maximum capacity, this is an easy way to keep an eye on sizing in the environment.

System overview Configuration

NSX Management Nodes: 1

Host Transport Nodes: 5

Edge Transport Nodes: 4

Transport Zones: 5

NSX Intelligence Nodes: 1

1. Click **System**
2. Click **System overview**

3. Click **Configuration**
4. Note the detail on the page and the ability to click on links for further inspection / configuration. Feel free to click on links and explore.

Once you are done exploring Click **System-> System overview ->Configuration** to return to this view.

System overview Capacity

Limit	Current Inventory	Maximum Capacity
System-wide Edge Nodes	2	0.62% 320
Edge Clusters	1	0.62% 160


Maximum capacity displayed above is for large appliance only.

1. Click **System**
2. Click **System overview**
3. Click **Capacity**
4. Note the current System configuration and the maximum capacity, this is an easy way to keep an eye on sizing in the environment.

Flowtracing

Flowtracing is the last tool we will look at in this module. This is a repeat of the tools section in Module 5 - Distributed Firewall and tools. If you have already completed Module 5 you can skip this section and end your lab. The flowtracing tool tests data plane connectivity between two vms by injecting traffic.


Verify Security policy



Before we explore Traceflow lets verify our 3-tier app rules are set as expected and the final rule in the DFW is set to deny.


1. Within the NSX Manager, click on **Security**
2. Click on **Distributed Firewall**
3. Click the arrow to **expand the 3-tier-app security policy**
4. Click the arrow to **expand the Default Layer3 security policy**

Verify DFW rules



1. **Make note** of the web to app allow rule that permits TCP 8443 traffic from web to app servers
2. **Verify** all three 3-tier app rules are **enabled**
3. **Verify** the final default rule is set to **block**

Traceflow



Traceflow traces the transport node-level path of a packet injected at a logical port. Traceflow supports L2 and L3 and is supported across ESXi, KVM Edge (including NAT). Using details provided by the user, a Traceflow packet will be created and inserted into the source logical port. When this packet travel from source to destination, all the logical network entities will report the observations. These observations are consolidated and shown in UI.

1. Click on **Plan & Troubleshoot** on the top bar
2. Click on **Traceflow**

Select web-01a as the Source

The screenshot shows a 'Source' configuration window with a 'RESET' button. The 'Type' is set to 'Virtual Machine'. The 'VM Name' field is open, showing a list of VMs: phoenix-db-01, phoenix-web-01, phoenix-web-02, user-vm-01, web-01a, web-02a, web-03a, and web-04a. A red box highlights the dropdown arrow next to the VM Name field, and another red box highlights 'web-01a' in the list. A circled '1' points to the 'Type' dropdown, and a circled '2' points to the 'VM Name' dropdown.

1. Click the **drop down arrow** next to VM name
2. Select **web-01a**

Select App 01a as the Destination

The screenshot shows a 'Destination' configuration window with a 'RESET' button. The 'Type' is set to 'Virtual Machine'. The 'VM Name' field is open, showing a list of VMs: app-01a, core-A, db-01a, and k8smaster-01a. A red box highlights the dropdown arrow next to the VM Name field, and another red box highlights 'app-01a' in the list.

1. Click the **drop down arrow** next to VM name
2. Select **app-01a**

Set the Protocol and port

Traceflow

1

2

Protocol Type
TCP

Source Port 0

Destination Port 8443

TCP Flags SYN X

Source RESET

Type Virtual Machine

VM Name web-01a

Host:8d528a95-b198-40b8-9e62-71305c82e0e3

Identify a virtual interface for packet injection

Virtual Interface Network adapter 1

Segment Port: default: 610c8284-3f60-4d58-8b8c-7f520a77b954

IP Address * 172.16.10.11

MAC Address * 00:50:56:9c:a5:b3

Destination RESET

Type Virtual Machine

VM Name app-01a

Host:8d528a95-b198-40b8-9e62-71305c82e0e3

Virtual Interface Network adapter 1

Segment Port: default: 56456de3-041d-4190-8c9a-b7803bafbd1c

IP Address * 172.16.20.11

MAC Address * 00:50:56:9c:d1:1e


3

ADVANCED SETTINGS

TRACE

1. Set the Protocol type to **TCP**
2. Set the Destination port to **8443**

Traceflow Allow Results




Once the trace has started, NSX-T Traceflow will provide back a map along with information that the packet took, from one end to another.

1. As you can see in the diagram, our packet was **injected, forwarded, and delivered** as expected.
2. Note that the flow was allowed by the **web to app DFW** rule created earlier in the module. (your rule ID may be different if you continued on in the previous section)
3. Feel free to **scroll down** to view more details of the flow.

Edit the trace



1. Click the **edit** button




1. Click **proceed**

Change to port 22

1. Change the **destination port to 22**
2. Click **trace**

Traceflow block results



1. Note that the flow was **dropped** this time by the default Drop firewall rule

This concludes the tools section of this module, feel free to view flows between different vm's and ports within the environment.


Module 7 Conclusion

Congratulations on completing Module 7.

Please proceed to any module below which interests you the most.


- [Module 1 - Introduction to NSX](#) (15 minutes) (Basic) In this module you will be introduced to the NSX Data Center platform and its capabilities. You will also explore the new NSX-T 3.0 features.
- [Module 2 - Host and Edge Transport Node Preparation](#) (30 minutes) (Advanced) In this module, you will review the prepared components ready for NSX workloads.
- [Module 3 - Logical Switching](#) (15 minutes) (Advanced) In this module, you will review the NSX lab fabric, configure Logical Switches, and attach them to virtual machines. You will be introduced to the use of the GENEVE protocol for overlay networking.
- [Module 4 - Logical Routing](#) (30 minutes) (Advanced) In this module you will learn how to configure Logical Routing in NSX for North-South and East-West traffic, configure High Availability (HA), and configure Equal Cost Multipathing (ECMP).
- [Module 5 - Distributed Firewall and Tools](#) (30 minutes) (Advanced) In this module we will demonstrate how the distributed firewall (DFW) and operational tools within NSX function and are configured.
- [Module 6 - NSX Services](#) (15 minutes) (Intermediate) In this module you will explore NSX service as VPN, Load balancing and DHCP / IPAM.

Test Your Skills!



Now that you've completed this lab, try testing your skills with VMware Odyssey, our newest Hands-on Labs gamification program. We have taken Hands-on Labs to the next level by adding gamification elements to the labs you know and love. Experience the fully automated VMware Odyssey as you race against the clock to complete tasks and reach the highest ranking on the leaderboard. Try the Network Security Odyssey lab.

How to End Lab



To end your lab click on the **END** button.

Conclusion

Thank you for participating in the VMware Hands-on Labs. Be sure to visit <http://hol.vmware.com/> to continue your lab experience online.

Lab SKU: HOL-2126-01-NET

Version: 20200929-013628