

NSX for vSphere Getting Started Guide

VMware NSX for vSphere, release 6.0.x

July 21, 2014

Table of Contents

NSX for vSphere Getting Started Guide.....	1
Introduction.....	3
Installation of NSX for vSphere	4
Infrastructure requirements for NSX-v	4
NSX-v Installation overview	6
Step 1: Install NSX Manager	6
Step 2: Install the NSX Controller Cluster	8
Step 3: Prepare ESXi hosts for NSX	10
L2 Logical Switching	14
Goal of the L2 logical switching lab	14
Create four Logical Switches	14
Add VMs on Web/App/DB Logical Switches	15
Validate that VMs on the same Logical Switch can communicate	16
Distributed Logical Routing.....	18
Goal of the logical routing lab.....	18
Create a single Distributed Logical Router.....	18
Validate that VMs in the different Logical Switches can communicate	21
Distributed Firewalling	23
Goal of the Distributed Firewalling lab	23
Create the Distributed Firewall rules	23
Validate the Distributed Firewall rules	25
Logical Centralized Routing.....	26



Goal of the Logical Centralized Routing lab	26
Create a single Logical Centralized Router (Edge)	27
Configure Dynamic Routing on Logical Distributed and Centralized Routers	29
Validate that dynamic routes are being learned	32
Validate communication from internal to Centralized Router external interface.....	33
Create many-to-one NAT (for traffic initiated from Web-Tier01 to external).....	33
Validate communication from Web-Tier-01 to Internet.....	34
Logical Load Balancing.....	35
Goal of the Logical Load Balancing lab.....	35
Create one new Load Balancer	36
Configure the Load Balancer.....	37
Update the Distributed Firewall rules to allow Load Balancer-to-Web server communication	40
Validate that the Server Pool is UP	40
Create a one-to-one NAT rule on the External Edge Router (for traffic initiated from external to load balancer)	41
Check that external network hosts can communicate to VIP	42
Getting Help and More Information.....	43
NSX-v Documentation	43
Contacting the NSX Technical Services Team	43

NOTE: To obtain the latest information about NSX for vSphere, please visit
<http://www.vmware.com/products/nsx>

Introduction

This document provides step-by-step examples that demonstrate how to set up the following network services in NSX for vSphere:

- Logical Switches
- Logical Distributed Routers
- Distributed Firewalls
- Logical Centralized Routers (Edge)
 - with Dynamic Routing
 - with many-to-one NAT
- Logical Load Balancers (Edge)

At the end, you'll have the following logical network deployed in your lab:

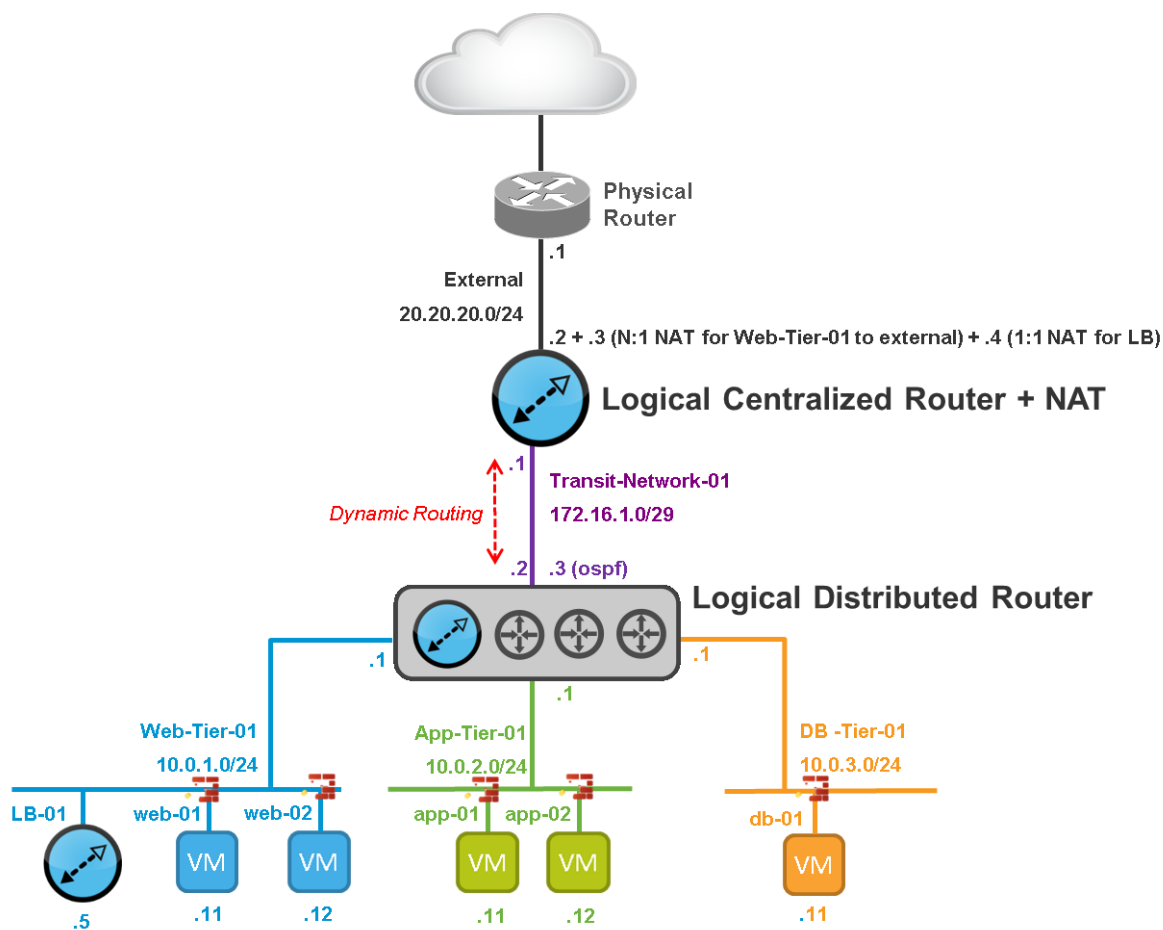


Figure 1 – Logical View of lab

L2 bridging, VPN, and service composer are not covered in this document. Likewise, integrations with third party vendors, such as Palo Alto Networks, Symantec and F5, are not covered here.

Installation of NSX for vSphere

This section guides you through the step-by-step installation, configuration and validation of a new NSX for vSphere (“NSX-v”) deployment.

Infrastructure requirements for NSX-v

VMware elements:

Prior to installing NSX for vSphere, you must deploy:

- vCenter 5.5 with:
 - one or more Compute clusters
 - Management and Edge cluster
- two or more ESXi 5.5 in each cluster

Each ESXi host has the following characteristics:

- Server hardware is listed on the VMware HCL for vSphere 5.5
- 2x Quad Core x86_64 compatible CPUs with a speed of 2Ghz or greater, plus hardware-assisted virtualization support (total of 8 physical cores)
- 32GB of RAM or greater
- 2x Physical NICs
- Either 5GB of Local Disk/Dedicated boot from SAN LUN or supported ESXi embedded device (USB/SD). Local Disk is not required if vSphere Auto Deploy is used.

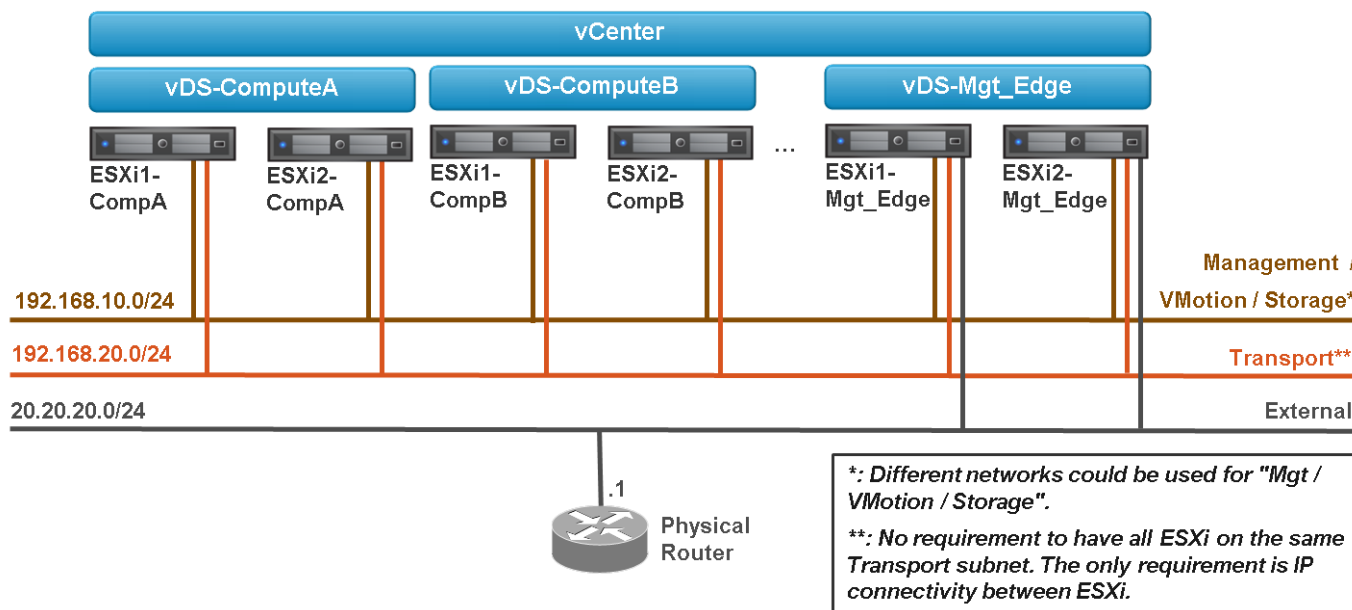


Figure 2 – Infrastructure for NSX

For resource constraints, this lab uses only one Compute Cluster, as shown in the following screenshots.

Network fabric:

Configure at least 1600 byte of MTU frame sizes on all the physical switches/routers between ESXi.

vCenter:

Clusters:

- One Compute Cluster “Cluster-CompA” with two ESXi.
- One Management + Edge Cluster “Cluster-Mgt_Edge” with two ESXi.




Figure 3 – vCenter Host View

Networking:

- Virtual Standard Switch (vSS) for Cluster-CompA and Cluster-Mgt_Edge:
 - **Management:** This vSS is used for the ESXi-Compute and ESXi-Mgt_Edge management. *Interface to use:* The interface of the ESXi in Cluster-CompA + Cluster-Mgt_Edge on the Management network is used.
- Virtual Distributed Switch (vDS) for Cluster-CompA:
 - **vDS-CompA:** This vDS will be used for the VM production traffic. *Interface to use:* The interface of the ESXi in Cluster-CompA on the Transport network is used. *Note: No ESXi IP@ is configured yet.*
- Virtual Distributed Switch (vDS) for Cluster-Mgt_Edge:
 - **vDS-Mgt_Edge:** This vDS will be used for the VM production traffic. *Interface to use:* The interface of the ESXi in Cluster-Mgt_Edge on the Transport network is used. *Note: No ESXi IP@ is configured yet. Note2: Create a Management Network for the future logical routers “LogicalRouter_Mgt”*
 - **vDS-External:** This vDS will be used to talk to the physical external network. *Interface to use:* The interface of the ESXi in Cluster-Mgt_Edge on the External network is used. *Note: No ESXi IP@ is configured.*

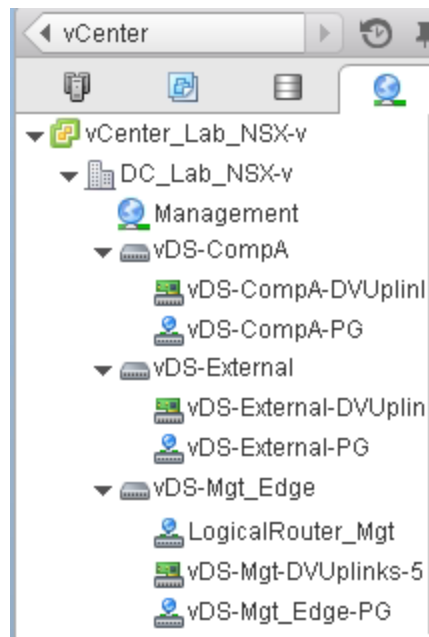


Figure 4 – vCenter Networking View

NSX-v Installation overview

In this step, you'll deploy the NSX Manager and NSX Controller Nodes:

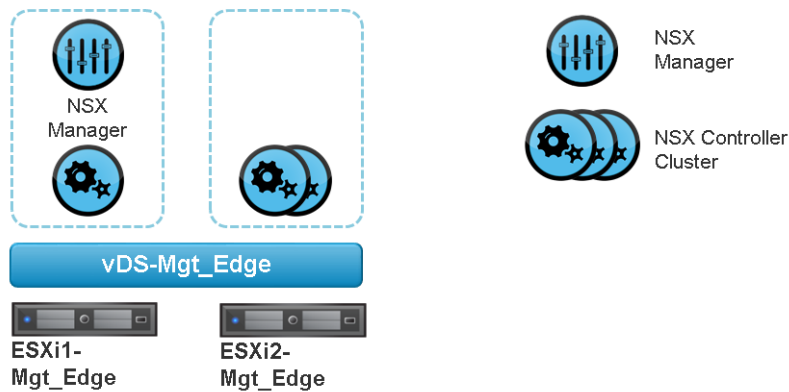


Figure 5 – NSX elements

Step 1: Install NSX Manager

The NSX Manager is the centralized management component of NSX, and runs as a virtual appliance on an ESX host.

1. Install NSX Manager: From **vCenter Home -> Hosts and Clusters**, select Cluster-Mgt_Edge and Deploy OVF Template

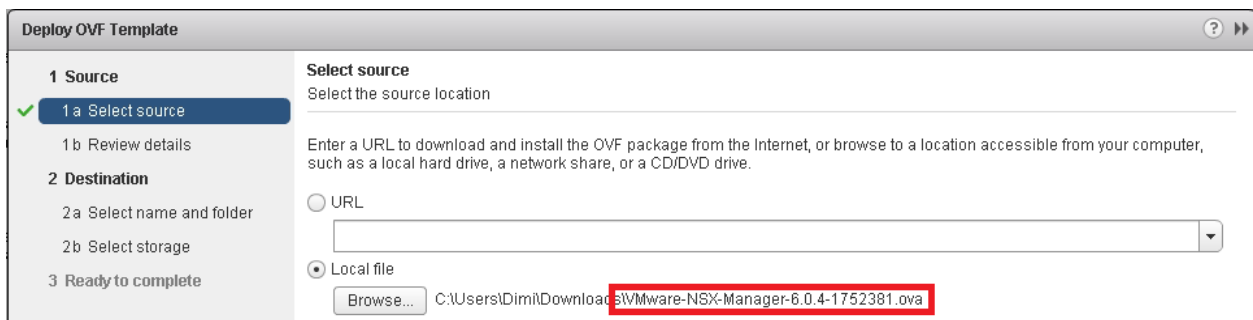


Figure 6 – Installation NSX Manager

2. Register NSX Manager with vCenter: Log in NSX Manager and from [NSX Manager Manage -> NSX Management Services](#), register to vCenter

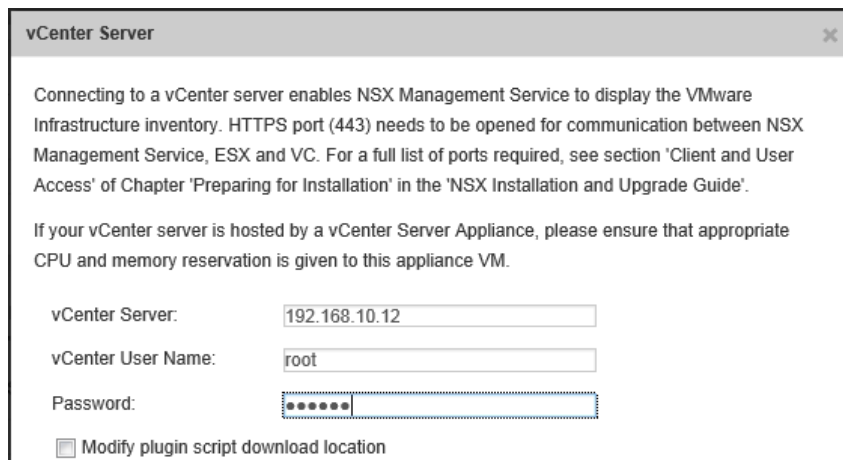


Figure 7 – NSX Manager registration to vCenter

3. Validate registration: Log out of vCenter if already logged in. And re-log in with root (required to get the NSX plugin installed in vCenter). *Note: The first login can take a few minutes.* After registration, you will see the Network & Security plugin in the Inventory:




Figure 8 – NSX plugin in vCenter

Step 2: Install the NSX Controller Cluster

The NSX Controller Cluster is a distributed state management system that controls virtual networks and overlay transport tunnels

1. Install the first NSX Controller Node: From **NSX Home -> Installation**, add first NSX Controller Node.

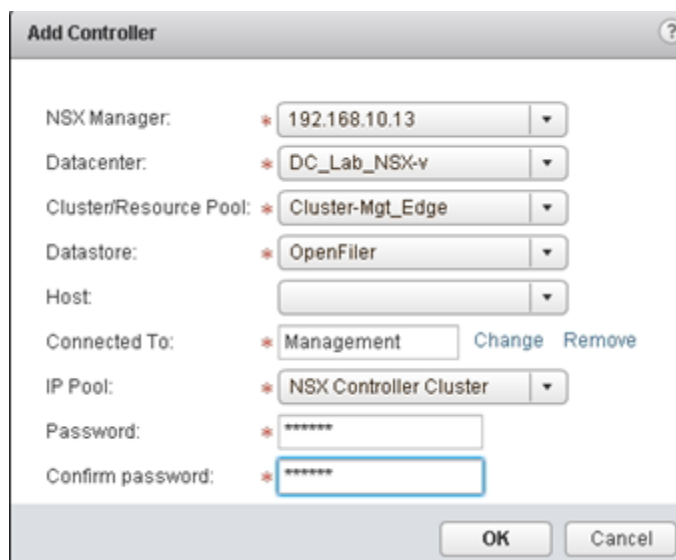


Figure 9 – First NSX Controller Node installation

The IP Pool “NSX Controller Cluster” has been created with the following settings:

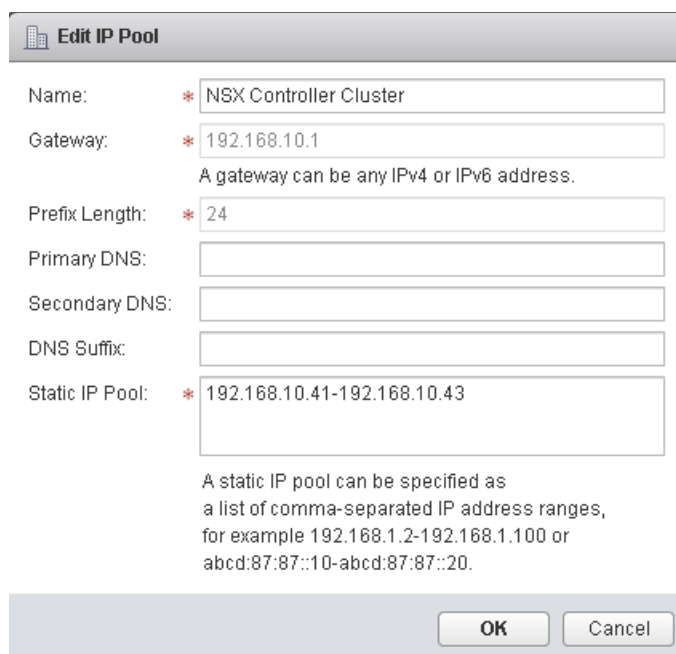


Figure 10 – NSX Controller Cluster IP pool

2. Validate the installation of first NSX Controller Node: The deployment of an NSX Controller Node can take few minutes.

NSX Controller nodes

Name	Node	NSX Manager	Cluster/Resource Pool	Datastore	Host	Software Version	Status
controller-1	192.168.10.41	192.168.10.13	Resources	datastore1 (3)	192.168.10.32	6.0	✓ Normal

Figure 11 – First NSX Controller Node deployed

Note: In rare cases, the Controller takes too long install and is automatically deleted. In such cases, you can install a DHCP server in the Controller's subnet to speed up its installation. That DHCP server can be configured with fake IP addresses since the Controller will still get its IP address from the NSX IP Pool.

3. Install the second and third NSX Controller Nodes:

Note: You can run with only one NSX Controller in a lab (not supported in a production setting), but this will render you unable to test Controller Node high-availability. For a production deployment or to test high-availability, you must install a total of three Controller Nodes.

From [NSX Home](#) -> [Installation](#), add second and third NSX Controller Nodes

Add Controller

NSX Manager: 192.168.10.13
Datacenter: DC_Lab_NSX-v
Cluster/Resource Pool: Cluster-Mgt_Edge
Datastore: OpenFiler
Host:
Connected To: Management Change Remove
IP Pool: NSX Controller Cluster

OK Cancel

Figure 12 – Second and third NSX Controller Nodes installation

4. Validate installation of all three NSX Controller Nodes

Name	Node	NSX Manager	Cluster/Resource Pool	Datastore	Host	Software Version	Status
controller-1	192.168.10.41	192.168.10.13	Resources	OpenFiler	192.168.10.31	6.0	✓ Normal
controller-2	192.168.10.42	192.168.10.13	Resources	OpenFiler	192.168.10.32	6.0	✓ Normal
controller-3	192.168.10.43	192.168.10.13	Resources	OpenFiler	192.168.10.31	6.0	✓ Normal

Figure 13 –NSX Controller Cluster deployed

Step 3: Prepare ESXi hosts for NSX

To provide all the NSX services, special kernel modules and user space tools have to be installed on the ESXi hosts.

1. Install NSX elements on cluster hosts: From **NSX Home -> Installation -> Host Preparation**, click **Install** for all the clusters:

Installation of network virtualization components on vSphere hosts

Clusters & Hosts	Installation Status	Firewall
Cluster-Mgt_Edge	Install	Not Enabled
Cluster-CompA	Install	Not Enabled

Figure 14 –Installation of NSX elements on cluster hosts

2. Check the installation of NSX elements on cluster hosts

Installation of network virtualization components on vSphere hosts

Clusters & Hosts	Installation Status	Firewall
Cluster-Mgt_Edge	✓ 6.0.4 Uninstall	✓ Enabled
Cluster-CompA	✓ 6.0.4 Uninstall	✓ Enabled

Figure 15 – Validation of installation of NSX elements on clusters hosts

3. Configure the VXLAN VTEP interface for Cluster-CompA hosts: From **NSX Home -> Installation -> Host Preparation**, click **Configure** for the Cluster-CompA:

Configure VXLAN networking

Configuring all hosts in cluster "Cluster-CompA" for VXLAN networking.

Switch: * VDS-CompA

VLAN: * 0

MTU: * 1600

VMKNic IP Addressing: * ☐ Use DHCP ☒ Use IP Pool

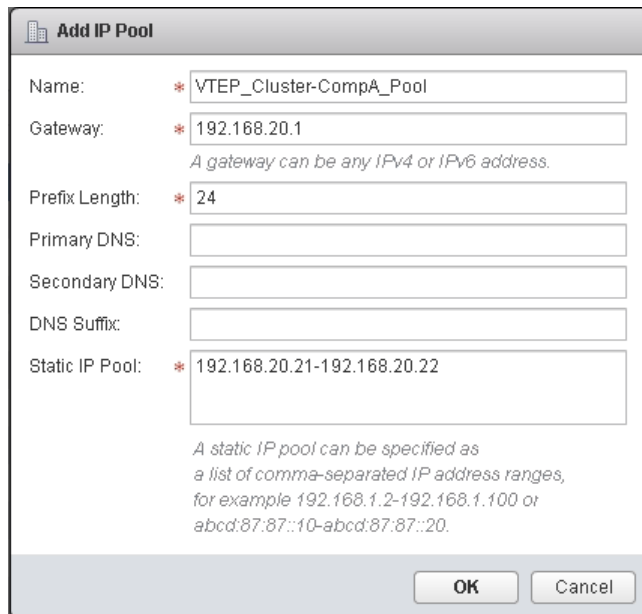
IP Pool: VTEP_Cluster-Co...

VMKNic Teaming Policy: * Fail Over

VTEP: * 1

OK Cancel

Figure 16 – Configuration of VTEP interface for Cluster-CompA hosts



Add IP Pool

Name: * VTEP_Cluster-CompA_Pool

Gateway: * 192.168.20.1
A gateway can be any IPv4 or IPv6 address.

Prefix Length: * 24

Primary DNS:

Secondary DNS:

DNS Suffix:

Static IP Pool: * 192.168.20.21-192.168.20.22
A static IP pool can be specified as a list of comma-separated IP address ranges, for example 192.168.1.2-192.168.1.100 or abcd:87:87::10-abcd:87:87::20.

OK Cancel

Figure 17 – Configuration of VTEP IP@ pool for the Cluster-CompA hosts

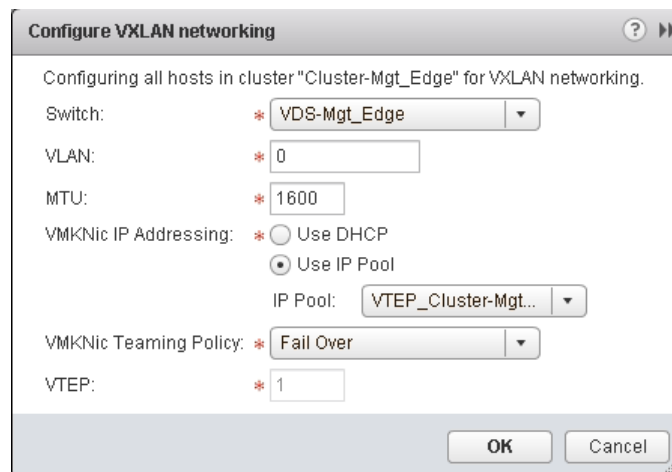
4. Validate the VTEP configuration on the Cluster-CompA hosts. *Note: You may see an “Error Unconfigure” message. This is a known display issue. Refresh the window to see the correct status.*

Installation of network virtualization components on vSphere hosts

Clusters & Hosts	Installation Status	Firewall	VXLAN
▶ Cluster-Mgt_Edge	✓ 6.0.4 Uninstall	✓ Enabled	Configure
▶ Cluster-CompA	✓ 6.0.4 Uninstall	✓ Enabled	✓ Enabled

Figure 18 – Validation VTEP IP@ configuration for the Cluster-CompA hosts

5. Configure the VXLAN VTEP interface for Cluster-Mgt_Edge hosts:



Configure VXLAN networking

Configuring all hosts in cluster "Cluster-Mgt_Edge" for VXLAN networking.

Switch: * VDS-Mgt_Edge

VLAN: * 0

MTU: * 1600

VMKNic IP Addressing: * ☐ Use DHCP ☒ Use IP Pool

IP Pool: VTEP_Cluster-Mgt...

VMKNic Teaming Policy: * Fail Over

VTEP: * 1

OK Cancel

Figure 19 – Configuration of VTEP interface for Cluster-Mgt_Edge hosts

Figure 20 – Configuration of VTEP IP@ pool for the Cluster-Mgt_Edge hosts

6. Validate the VTEP configuration on the Cluster-Mgt_Edge hosts. *Note: You may see an “Error Unconfigure” message. This is a known display issue. Refresh the window to see the correct status.*

Installation of network virtualization components on vSphere hosts				
Clusters & Hosts	Installation Status		Firewall	VXLAN
Cluster-CompA	✓	6.0.4 Uninstall	✓ Enabled	✓ Enabled
Cluster-Mgt_Edge	✓	6.0.4 Uninstall	✓ Enabled	✓ Enabled

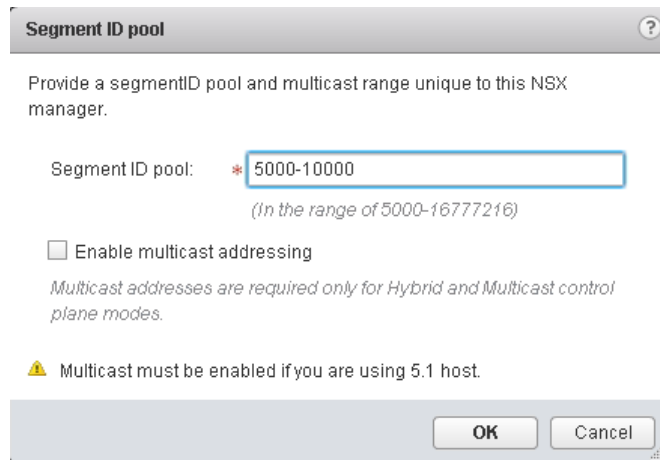
Figure 21 – Validation VTEP IP@ configuration for the Cluster-Mgt_Edge hosts

7. View of the VTEP IP@ allocated to each Cluster hosts. From NSX Home -> Installation -> Logical Network Preparation -> VXLAN Transport:

Clusters & Hosts	Configuration Status	Switch	VLAN	MTU	VMKNic IP Addressing	Teaming Policy	VTEP
Cluster-CompA	✓ Unconfigure	VDS-CompA	0	1600	IP Pool	Fail Over	1
192.168.10.21	✓ Ready				vmk1 : 192.168.20.21		
192.168.10.22	✓				vmk1 : 192.168.20.22		
Cluster-Mgt_Edge	✓ Unconfigure	VDS-Mgt_Edge	0	1600	IP Pool	Fail Over	1
192.168.10.31	✓ Ready				vmk1 : 192.168.20.31		
192.168.10.32	✓ Ready				vmk1 : 192.168.20.32		

Figure 22 – View of the VTEP IP@ allocated to each Cluster hosts

8. Configure VXLAN Segment ID (VXLAN Network Identifier – VNI): From NSX Home -> Installation -> Logical Network Preparation -> Segment ID, click **Edit**. *Note: Since NSX 6.0 with ESXi 5.5, multicast support is no longer required on the physical fabric.*



Segment ID pool

Provide a segmentID pool and multicast range unique to this NSX manager.

Segment ID pool: * 5000-10000
(In the range of 5000-16777216)

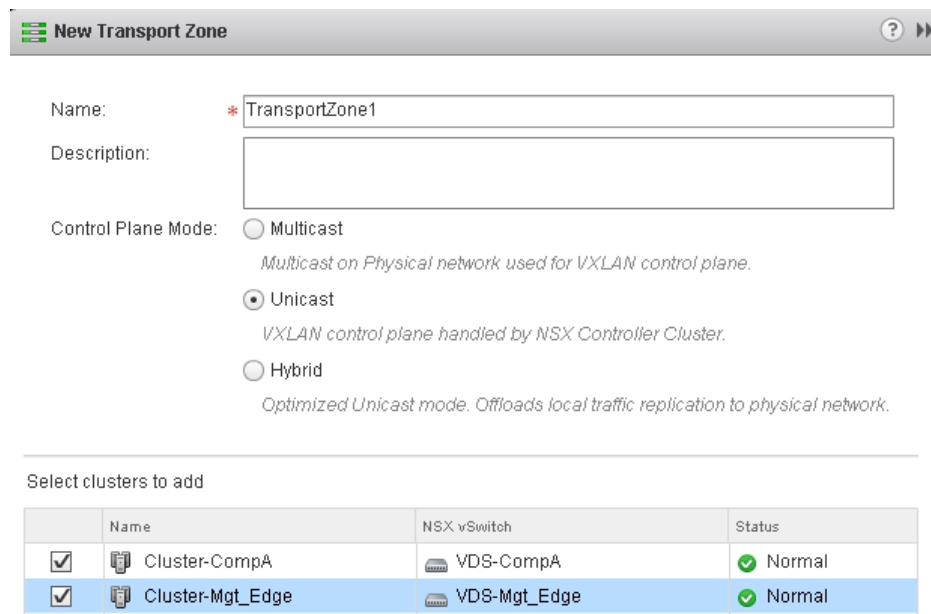
☐ Enable multicast addressing
Multicast addresses are required only for Hybrid and Multicast control plane modes.

⚠ Multicast must be enabled if you are using 5.1 host.

OK Cancel

Figure 23 – View of the VTEP IP@ allocated to each Cluster hosts

- Configure a Transport Zone: The transport zone is the compute diameter of your cloud. You want all your ESXi hosts to participate to your cloud. From [NSX Home -> Installation -> Logical Network Preparation -> Transport Zone](#), click +. *Note: Since NSX 6.0 with ESXi 5.5, multicast support is no longer required on the physical fabric.*



New Transport Zone

Name: * TransportZone1

Description:

Control Plane Mode: ☐ Multicast
Multicast on Physical network used for VXLAN control plane.
☒ Unicast
VXLAN control plane handled by NSX Controller Cluster.
☐ Hybrid
Optimized Unicast mode. Offloads local traffic replication to physical network.

Select clusters to add

	Name	NSX vSwitch	Status
<input checked="" type="checkbox"/>	Cluster-CompA	VDS-CompA	✓ Normal
<input checked="" type="checkbox"/>	Cluster-Mgt_Edge	VDS-Mgt_Edge	✓ Normal

Figure 24 – Creation of the Transport Zone that spans among all Clusters

This completes the installation of the NSX-v elements of your deployment. Proceed to the logical switch set-up steps in the next section.

L2 Logical Switching

Goal of the L2 logical switching lab

In this section, you will create Logical Switches.

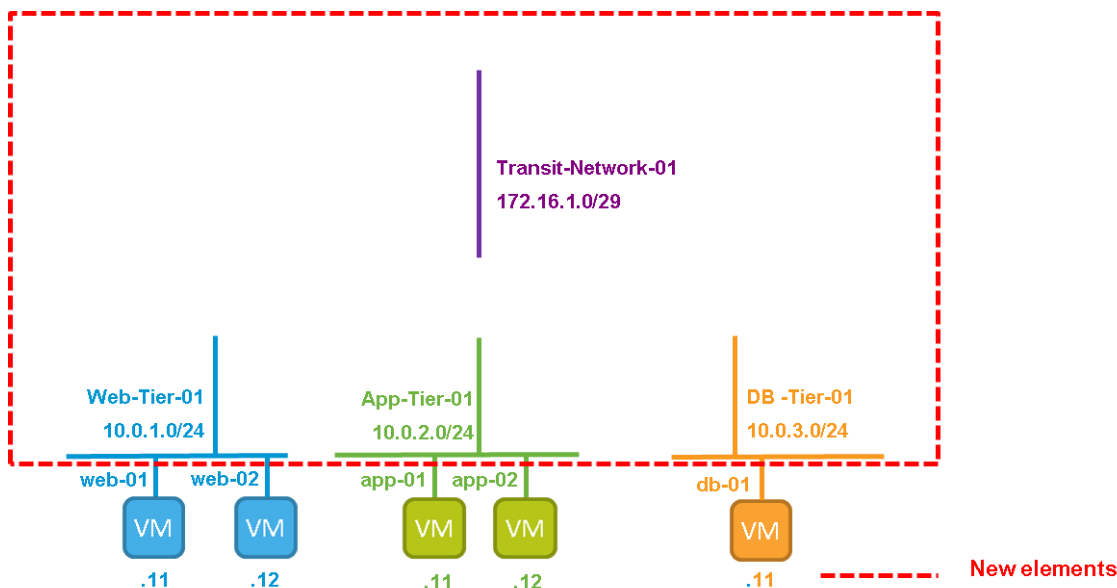


Figure 25 – Logical View Logical Switches

Create four Logical Switches

From **NSX Home** -> **Logical Switches**, create four Logical Switches called:

- Transit-Network-01
- Web-Tier-01
- App-Tier-01
- DB-Tier-01

The screenshot shows the 'New Logical Switch' configuration window. The **Name** field is set to **Transit-Network-01**. The **Description** field is empty. The **Transport Zone** is set to **TransportZone1**. The **Control Plane Mode** is set to **Unicast**, with the description: *VXLAN control plane handled by NSX Controller Cluster.* The **Multicast** option is described as: *Multicast on Physical network used for VXLAN control plane.* The **Hybrid** option is described as: *Optimized Unicast mode. Offloads local traffic replication to physical network.*

Figure 26 –Logical Switch creation

Note: You will notice that one vDS Port Group is automatically created for each Logical Switch.

From **vCenter Home -> Networking**

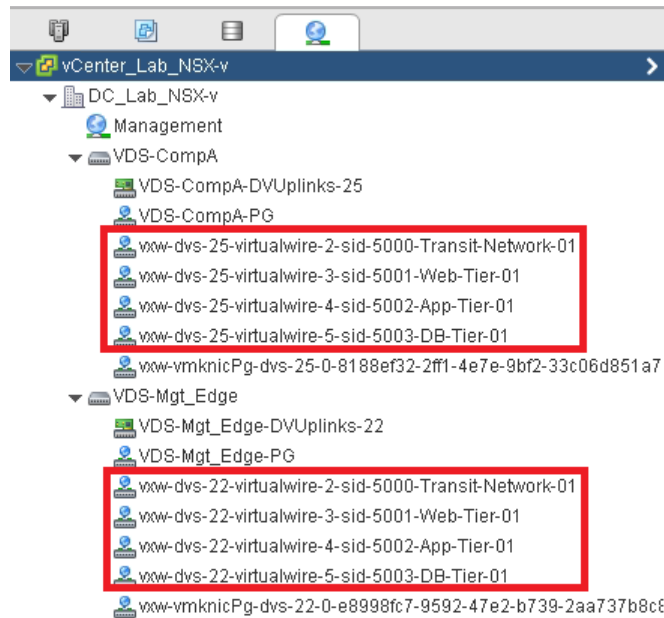


Figure 27 –vDS Port Groups created for each logical switch

Add VMs on Web/App/DB Logical Switches

You have VMs on the different Cluster-CompA hosts:




Figure 28 – VMs in Cluste-CompA

From **NSX Home -> Logical Switches**, add VMs to the appropriate logical switch




Figure 29 – Add VMs onLlogical Switch

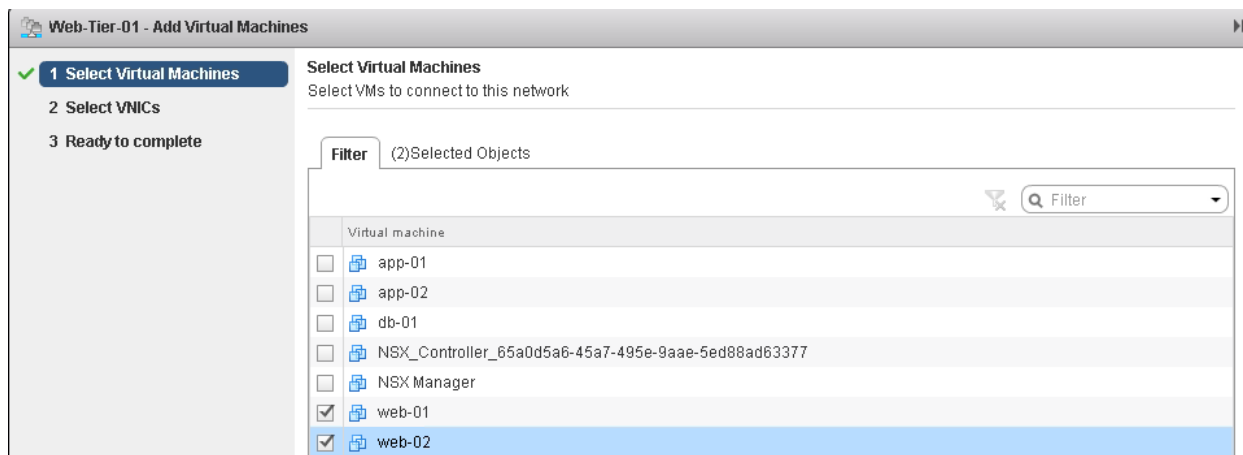


Figure 30 – Select VMs

Note: You can check the VMs are connected to the correct Logical Switch on vCenter too:

From vCenter Home -> Hosts and Clusters, look at the VM Hardware




Figure 31 – Validate VM Network adapter is connected to vDS port group

Validate that VMs on the same Logical Switch can communicate

```
root@web-01:~# ping 10.0.1.12
PING 10.0.1.12 (10.0.1.12) 56(84) bytes of data.
64 bytes from 10.0.1.12: icmp_req=1 ttl=64 time=12.9 ms
64 bytes from 10.0.1.12: icmp_req=2 ttl=64 time=0.711 ms
```

Figure 32 – ping between Web VMs

Note: The VM traffic flow in the fabric is:

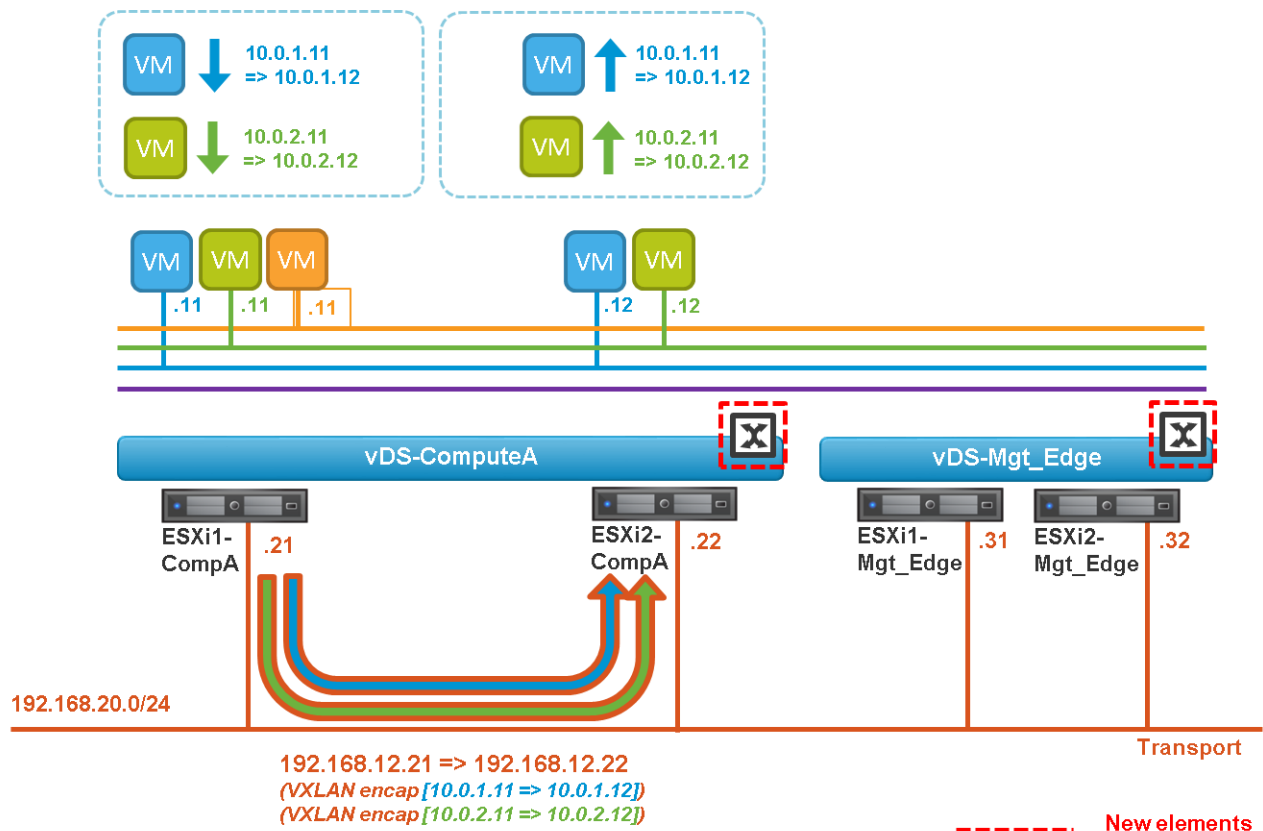


Figure 33 – Logical Switch traffic flow

Distributed Logical Routing

Goal of the logical routing lab

In this step, you'll create a Distributed Logical Router.

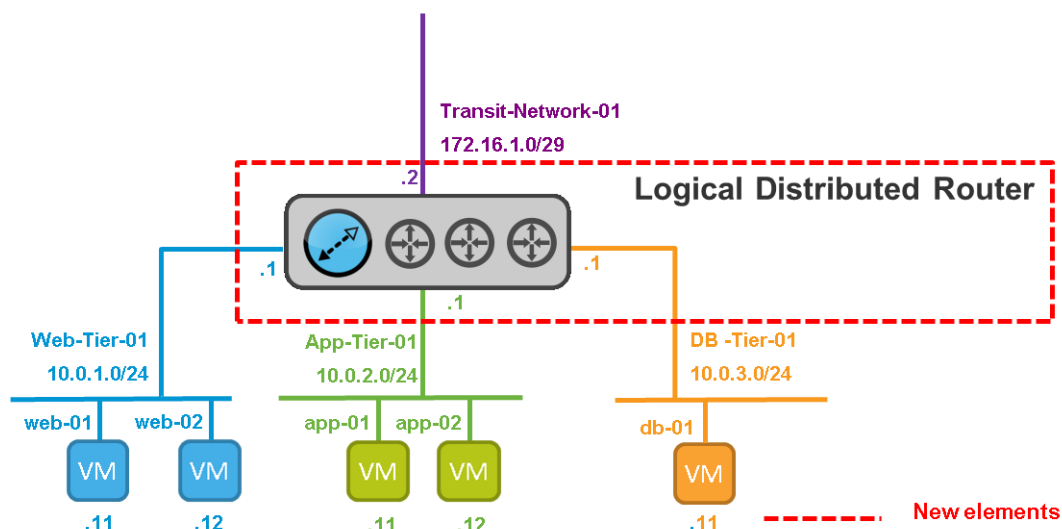


Figure 34 – Logical View Distributed Logical Router

Create a single Distributed Logical Router

From **NSX Home** -> **NSX Edges**, create a Distributed Logical Router with four interfaces (LIFS)

- Uplink to Transit-Network-01 with an IP of 172.16.1.2/29
- Internal connected to Web-Tier-01 Logical Switch with IP 10.0.1.1/24
- Internal connected to App-Tier-01 Logical Switch with IP 10.0.2.1/24
- Internal connected to DB-Tier-01 Logical Switch with IP 10.0.3.1/24

New NSX Edge

1 Name and description

2 CLI credentials

3 Configure deployment

4 Configure interfaces

5 Configure HA

6 Ready to complete

Name and description

Install Type:

Edge Services Gateway

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

Logical (Distributed) Router

Provides Distributed Routing and Bridging capabilities.

Enable High Availability

Enable HA, for enabling and configuring High Availability.

Name:

* LogicalRouter-01

Hostname:

Description:

Tenant:

Figure 35 – Logical Distributed Router creation, first pane

New NSX Edge

1 Name and description

2 CLI credentials

3 Configure deployment

4 Configure interfaces

5 Configure HA

6 Ready to complete

Configure deployment

Datacenter:

* DC_Lab_NSX-v

NSX Edge Appliances

+

x

Resource Pool	Host	Datastore	Folder
Cluster-Mgt_Ed...	192.168.10.32	OpenFiler	

Figure 36 – Logical Distributed Router creation, second pane

NSX for vSphere Getting Started Guide

Copyright © 2014, VMware. All rights reserved.

19

New NSX Edge

1 Name and description

2 CLI credentials

3 Configure deployment

4 Configure interfaces

5 Configure HA

6 Ready to complete

Configure interfaces

Management Interface Configuration

Connected To: *

LogicalRouter_Mgt

Change

Remove

+

x

IP Address	Subnet Prefix Length

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

Configure interfaces of this NSX Edge

+

x

Name	IP Address	Subnet Prefix Length	Connected To
Transit-01	172.16.1.2*	24	Transit-Network-01
Web-01	10.0.1.1*	24	Web-Tier-01
App-01	10.0.2.1*	24	App-Tier-01
DB-01	10.0.3.1*	24	DB-Tier-01

Figure 37 – Logical Distributed Router creation, third pane

Note: One Management Interface must be configured. This interface is to access the Logical Router Control VM via SSH for management/troubleshooting (the VM production traffic doesn't reach the Logical Router Control VM - see Figure 39 and Figure 40). For SSH access, configure a management IP address (not shown in the screenshot above).

NSX for vSphere Getting Started Guide

Copyright © 2014, VMware. All rights reserved.

20

Validate that VMs in the different Logical Switches can communicate

```
root@web-01:~# ping 10.0.2.11
PING 10.0.2.11 (10.0.2.11) 56(84) bytes of data.
64 bytes from 10.0.2.11: icmp_req=1 ttl=63 time=0.693 ms
64 bytes from 10.0.2.11: icmp_req=2 ttl=63 time=0.474 ms
64 bytes from 10.0.2.11: icmp_req=3 ttl=63 time=0.454 ms
```

Figure 38 – ping between Web and App VM

Note: The Logical Router Control VM (in the Mgt_Edge Cluster) is not involved in the L3 VM traffic flow.

The VM traffic flow in the fabric is shown below.

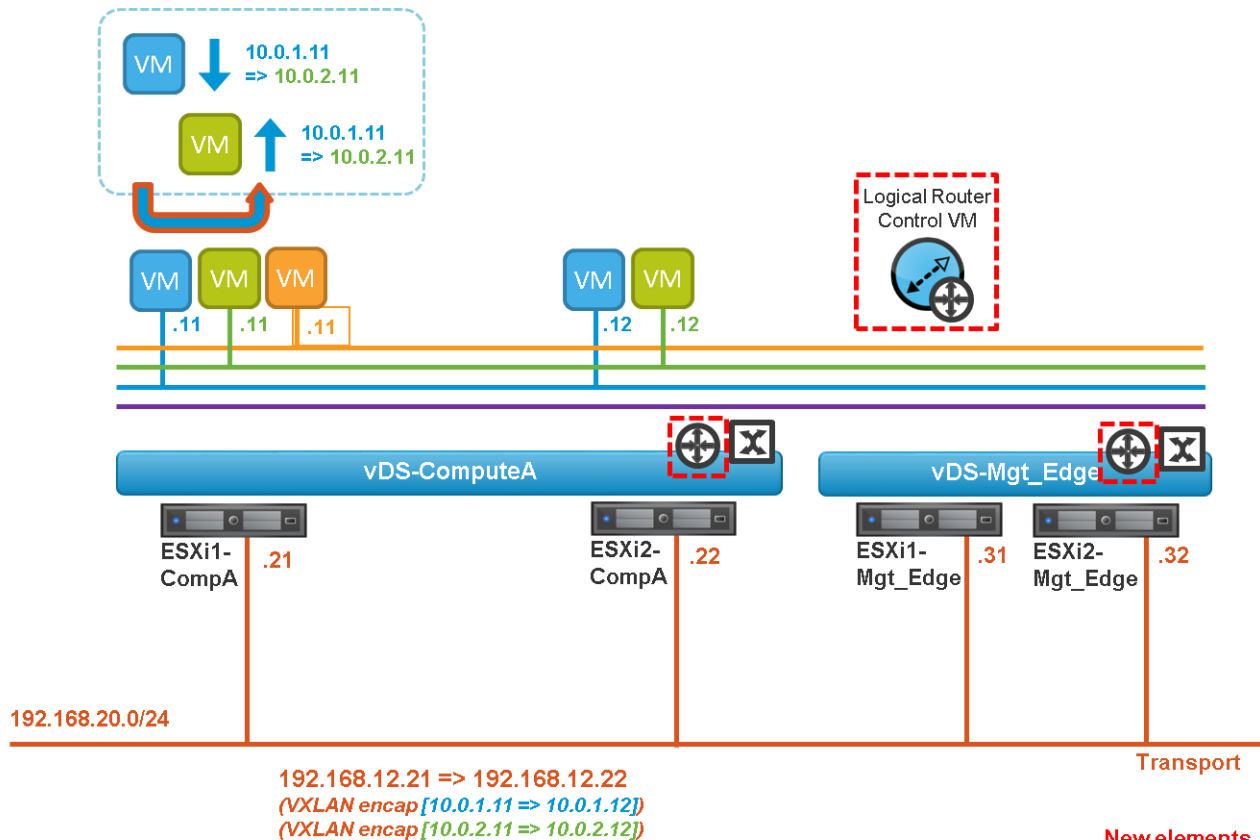


Figure 39 – L3 traffic flow – case both VMs are in the same ESXi host

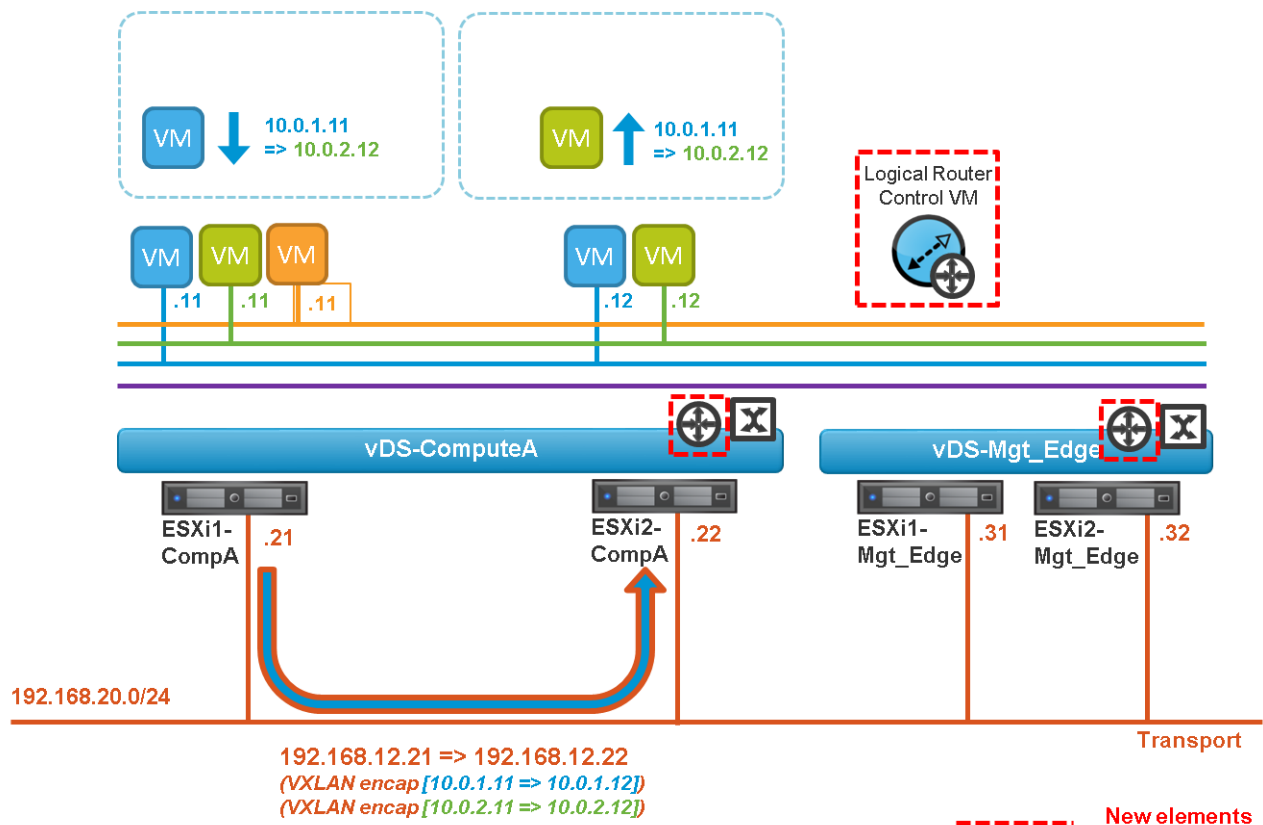


Figure 40 – L3 traffic flow – case both VMs are in different ESXi hosts

Distributed Firewalling

Goal of the Distributed Firewalling lab

In this step, you'll create the Distributed Firewall rules.

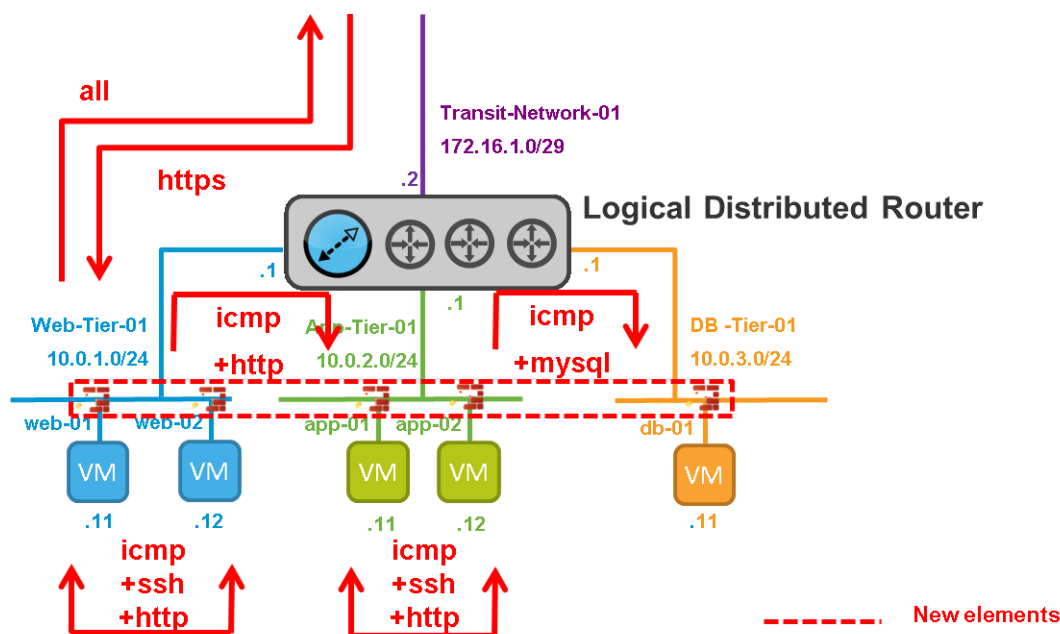


Figure 41 – Logical View Distributed Firewall

Create the Distributed Firewall rules

For ease of use, the example below is using Logical Switch Names for the “Source” and “Destination” instead of subnets.

This option works only if you have the VM Tools installed on the VMs.

If you do not have the VM Tools on your VMs, use subnet.

From **NSX Home** -> **Firewall**, create the rules:

- 1) External access: Source any, Destination Web-Tier-01, Allow https, Apply To Web-Tier-01
- 2) Inter Web-Tier-01: Source Web-Tier-01, Destination Web-Tier-01, Allow icmp + ssh + http, Apply To Web-Tier-01
- 3) Inter Web-Tier-01_block: Source Web-Tier-01, Destination Web-Tier-01, Block any, Apply To Web-Tier-01
- 4) Web-Tier-01-App-Tier-01: Source Web-Tier-01, Destination App-Tier-01, Allow icmp + http, Apply To Web-Tier-01 + App-Tier-01
- 5) Inter App-Tier-01: Source App-Tier-01, Destination App-Tier-01, Allow icmp + ssh + http, Apply To App-Tier-01
- 6) App-Tier-01-DB-Tier-01: Source App-Tier-01, Destination DB-Tier-01, Allow icmp + mysql, Apply To App-Tier-01 + DB-Tier-01
- 7) Web-Tier-01-External: Source Web-Tier-01, Destination any, Allow all, Apply To Web-Tier-01

8) Everything else: Source any, Destination any, Block any, Apply To Web-Tier-01 + App-Tier-01 + DB-Tier-01

Note: To display the field “Apply To”, click on the grid:




Figure 42 –Distributed Firewall fields selection

No.	Name	Source	Destination	Service	Action	Applied To
▼ Application-01 (Rule 1 - 8)						
1	External access	any	Web-Tier-01	HTTPS	Allow	Web-Tier-01
2	Inter Web-Tier-01	Web-Tier-01	Web-Tier-01	SSH HTTP ICMP Echo	Allow	Web-Tier-01
3	Inter Web-Tier-01_block	Web-Tier-01	Web-Tier-01	any	Block	Web-Tier-01
4	Web-Tier-01-App-Tier-01	Web-Tier-01	App-Tier-01	HTTP ICMP Echo	Allow	Web-Tier-01 App-Tier-01
5	Inter App-Tier-01	App-Tier-01	App-Tier-01	SSH HTTP ICMP Echo	Allow	App-Tier-01
6	App-Tier-01-DB-Tier-01	App-Tier-01	DB-Tier-01	ICMP Echo MySQL	Allow	App-Tier-01 DB-Tier-01
7	Web-Tier-01-External	Web-Tier-01	any	any	Allow	Web-Tier-01
8	Everything else	any	any	any	Block	Web-Tier-01 App-Tier-01 DB-Tier-01

Figure 43 –Distributed Firewall rules

Validate the Distributed Firewall rules

```
root@web-01:~# ping 10.0.2.11
PING 10.0.2.11 (10.0.2.11) 56(84) bytes of data.
64 bytes from 10.0.2.11: icmp_req=1 ttl=63 time=0.693 ms
64 bytes from 10.0.2.11: icmp_req=2 ttl=63 time=0.474 ms
64 bytes from 10.0.2.11: icmp_req=3 ttl=63 time=0.454 ms
^C
--- 10.0.2.11 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2003ms
rtt min/avg/max/mdev = 0.454/0.540/0.693/0.109 ms
root@web-01:~# ssh 10.0.2.11

ssh: connect to host 10.0.2.11 port 22: Connection timed out
```

Figure 44 – ping and ssh between Web and App VM

Note: The non-authorized traffic is dropped at the beginning:

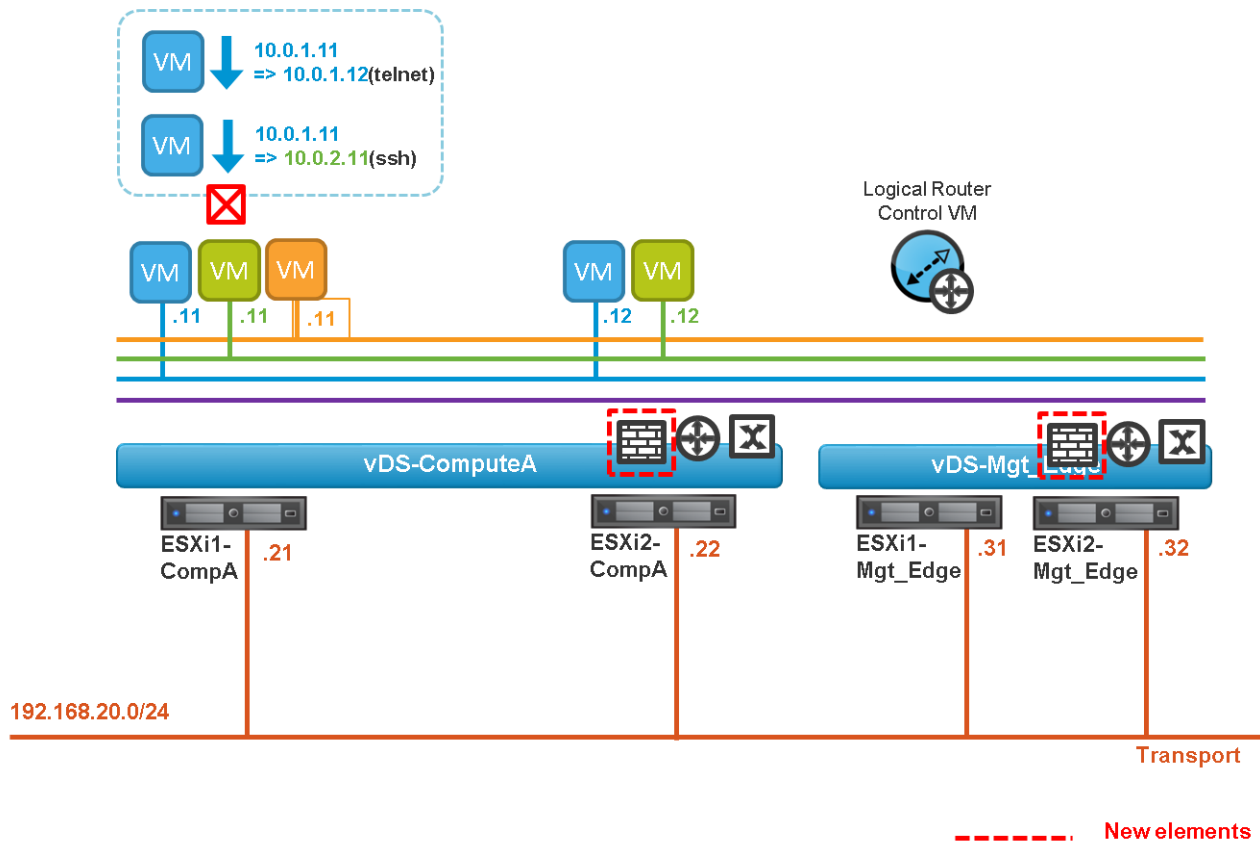


Figure 45 – Distributed Firewall traffic flow

Logical Centralized Routing

Goal of the Logical Centralized Routing lab

In this step, you'll create a Logical Centralized Router (Edge) with:

- dynamic routing
- many-to-one NAT

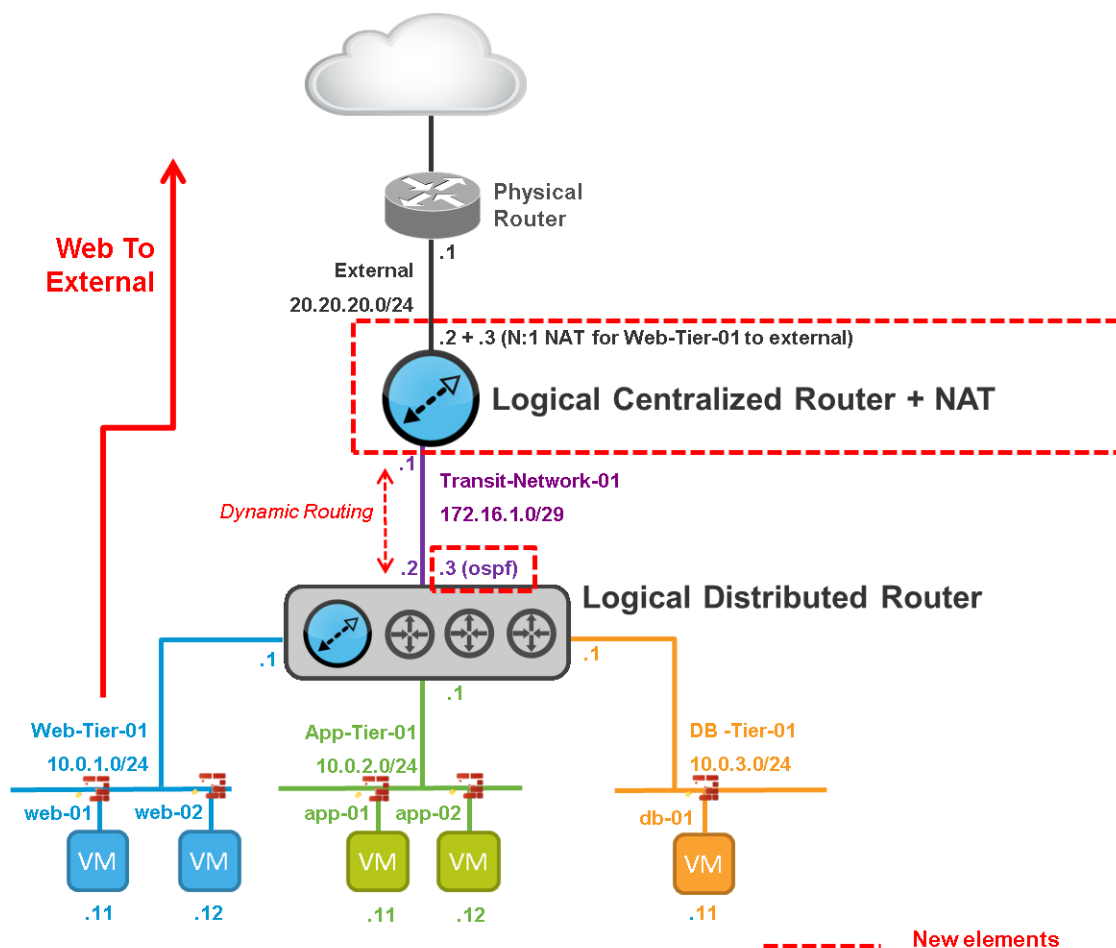


Figure 46 – Logical View Logical Centralized Router

Create a single Logical Centralized Router (Edge)

From **NSX Home -> NSX Edges**, create an Edge Service Gateway with two interfaces (LIFS)

- Uplink to External with an IP of 20.20.20.2/24
- Internal to Transit-Network-01 with an IP of 172.16.1.1/29

The screenshot shows the 'New NSX Edge' configuration window. On the left is a sidebar with a list of steps: 1 Name and description (highlighted with a green checkmark), 2 CLI credentials, 3 Configure deployment, 4 Configure interfaces, 5 Default gateway settings, 6 Firewall and HA, and 7 Ready to complete. The main area is titled 'Name and description'. It contains the following fields and options:

- Install Type:** Two radio buttons. 'Edge Services Gateway' is selected. Below it is a description: 'Provides common gateway services such as DHCP, Firewall, VPN, NAT, Routing and Load Balancing.' The other option is 'Logical (Distributed) Router' with the description: 'Provides Distributed Routing and Bridging capabilities.'
- Enable High Availability:** An unchecked checkbox. Below it is the text: 'Enable HA, for enabling and configuring High Availability.'
- Name:** A text field containing 'CentralizedRouter-01'.
- Hostname:** An empty text field.
- Description:** A large empty text area.
- Tenant:** An empty text field.

Figure 47 – Logical Centralized Router creation, first pane

The screenshot shows the 'New NSX Edge' configuration window at step 3: 'Configure deployment'. The sidebar on the left shows steps 1 and 2 completed with green checkmarks, and step 3 is highlighted. The main area contains the following configuration options:

- Datacenter:** A dropdown menu showing 'DC_Lab_NSX-v'.
- Appliance Size:** Four radio buttons: 'Compact', 'Large' (selected), 'X-Large', and 'Quad Large'.
- Enable auto rule generation:** A checked checkbox. Below it is the text: 'Enable auto rule generation, to automatically generate service rules to allow flow of control traffic.'
- NSX Edge Appliances:** A section with a table. Above the table are icons for adding (+), editing (pencil), and deleting (x).

Resource Pool	Host	Datastore	Folder
Cluster-Mgt_Ed...		OpenFiler	

Figure 48 – Logical Centralized Router creation, third pane

New NSX Edge

- ✓ 1 Name and description
- ✓ 2 CLI credentials
- ✓ 3 Configure deployment
- ✓ 4 **Configure interfaces**
- 5 Default gateway settings
- 6 Firewall and HA
- 7 Ready to complete

Configure interfaces

Configure interfaces of this NSX Edge

+ ✎ ✕

vNIC#	Name	IP Address	Subnet Prefix Length	Connected To
0	External	20.20.20.2*	24	vDS-External-...
1	Transit-01	172.16.1.1*	29	Transit-Netwo...

Figure 49 – Logical Centralized Router creation, fourth pane

New NSX Edge

- ✓ 1 Name and description
- ✓ 2 CLI credentials
- ✓ 3 Configure deployment
- ✓ 4 Configure interfaces
- 5 **Default gateway settings**
- 6 Firewall and HA

Default gateway settings

☒ **Configure Default Gateway**

vNIC: * External

Gateway IP: * 20.20.20.1

MTU: 1500

Figure 50 – Logical Centralized Router creation, fifth pane

New NSX Edge

- ✓ 1 Name and description
- ✓ 2 CLI credentials
- ✓ 3 Configure deployment
- ✓ 4 Configure interfaces
- ✓ 5 Default gateway settings
- 6 **Firewall and HA**
- 7 Ready to complete

Firewall and HA

☒ **Configure Firewall default policy**

Default Traffic Policy: ☒ Accept ☐ Deny

Logging: ☐ Enable ☒ Disable

Configure HA parameters

Configuring HA parameters is mandatory for HA to work.

Figure 51 – Logical Centralized Router creation, sixth pane

Configure Dynamic Routing on Logical Distributed and Centralized Routers

Dynamic routing configuration on Logical Distributed Router

- 1) Enable Dynamic Routing:
 - a) From **NSX Home** -> **NSX Edges**, select the Logical Distributed Router and navigate to **Manage** -> **Routing** -> **Global Configuration**, and click Edit Dynamic Routing Configuration.
 - b) Accept the default Router ID and Publish change (don't click "Enable OSPF" here because a Protocol Address needs to be defined first)



Figure 52 – Logical Distributed Router Dynamic Routing configuration

- 2) Enable OSPF
 - a) Navigate to **Manage** -> **Routing** -> **OSPF**, click Edit:
 - o Enable OSPF checkbox
 - o Add a Protocol Address of 192.168.10.3
 - o Forwarding Address of 192.168.10.2
 - b) and **Publish change**

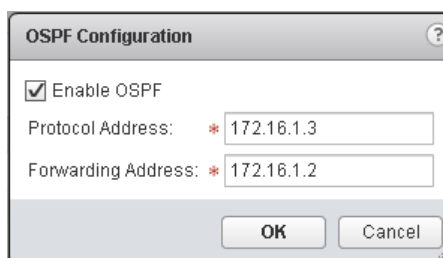
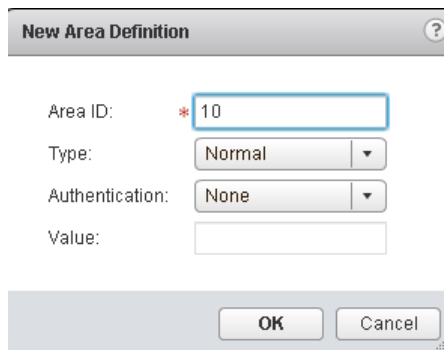


Figure 53 – Logical Distributed Router OSPF configuration

- 3) Configure OSPF
 - a) Navigate to **Manage** -> **Routing** -> **OSPF**, click Edit:
 - b) Add a new Area Definition with the default values:



New Area Definition

Area ID: * 10

Type: Normal

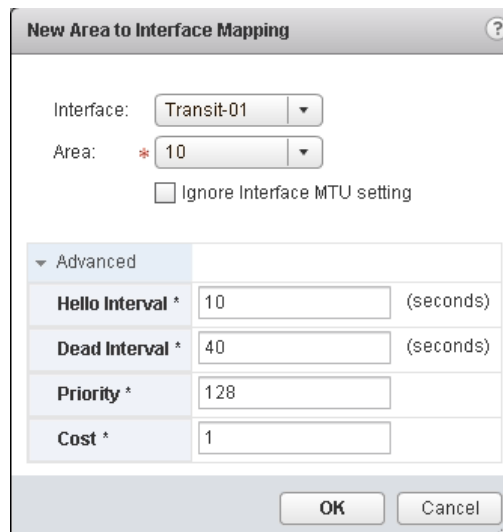
Authentication: None

Value:

OK Cancel

Figure 54 – Logical Distributed Router OSPF area configuration

4) Add the Area to Interface Transit-Uplink and Publish change:



New Area to Interface Mapping

Interface: Transit-01

Area: * 10

☐ Ignore Interface MTU setting

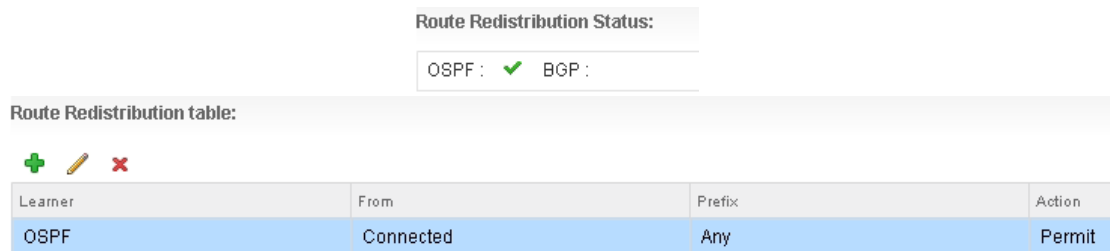
Advanced

Hello Interval *	10	(seconds)
Dead Interval *	40	(seconds)
Priority *	128	
Cost *	1	

OK Cancel

Figure 55 – Logical Distributed Router OSPF area interface configuration

5) Validate Route Redistribution for connected networks is permitted:



Route Redistribution Status:

OSPF : ✓ BGP :

Route Redistribution table:


+ ✎ ✕

Learner	From	Prefix	Action
OSPF	Connected	Any	Permit

Figure 56 – Logical Distributed Router dynamic routing route redistribution

Dynamic routing configuration on Logical Centralized Router

- 1) Enable Dynamic Routing
 - a) From [NSX Home](#) -> [NSX Edges](#), select the Logical Centralized Router and navigate to [Manage](#) -> [Routing](#) -> [Global Configuration](#), Click Edit Dynamic Routing Configuration
 - b) Accept the default Router ID and Publish change (don't click "Enable OSPF" here because a Protocol Address needs to be defined first)



Edit Dynamic Routing Configuration

Router ID: * External - 20.20.20.2

☒ Enable OSPF

☐ Enable BGP

☐ Enable IS-IS

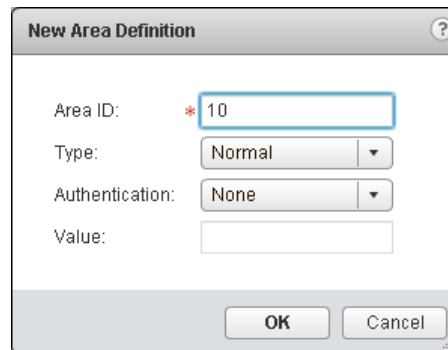
☐ Enable Logging

Log Level: Info

Save Cancel

Figure 57 – Logical Centralized Router Dynamic Routing configuration

- 2) Configure OSPF
 - a) Navigate to **Manage -> Routing -> OSPF**, click Edit:
 - b) Add a new Area Definition with the default values:



New Area Definition

Area ID: * 10

Type: Normal

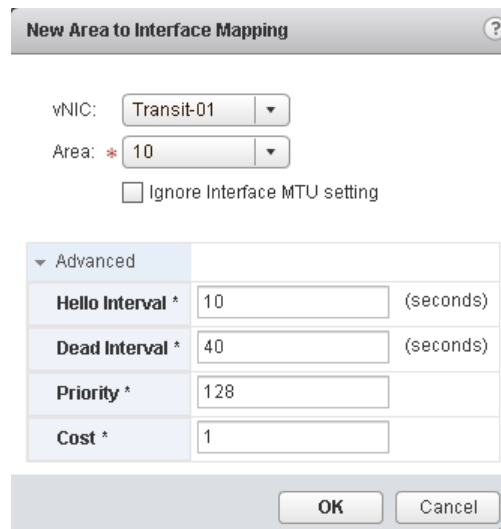
Authentication: None

Value:

OK Cancel

Figure 58 – Logical Centralized Router OSPF area configuration

- 3) Add the Area to Interface Transit-Uplink and Publish change:



New Area to Interface Mapping

vNIC: Transit-01

Area: * 10

☐ Ignore Interface MTU setting

Advanced	
Hello Interval *	10 (seconds)
Dead Interval *	40 (seconds)
Priority *	128
Cost *	1

OK Cancel

Figure 59 – Logical Centralized Router OSPF area interface configuration

- 4) Add Route Redistribution for connected networks and static routes and Publish change:

Route Redistribution Status:			
OSPF : ✓ ISIS : BGP :			
Route Redistribution table:			
<div> <div>+</div> <div></div> <div>x</div> </div>			
Learner	From	Prefix	Action
OSPF	Static routes,Connected	Any	Permit

Figure 60 – Logical Centralized Router dynamic routing route redistribution

Validate that dynamic routes are being learned

```

vShield-edge-1-0> show ip ospf neighbor
Neighbor ID      Priority     Address        Dead Time   State
20.20.20.2       128         172.16.1.1     37          Full/BDR
vShield-edge-1-0> show ip route

Codes: 0 - OSPF derived, i - IS-IS derived, B - BGP derived,
C - connected, S - static, L1 - IS-IS level-1, L2 - IS-IS level-2,
IA - OSPF inter area, E1 - OSPF external type 1, E2 - OSPF external type 2

Total number of routes: 10

O    E2  0.0.0.0/0          [110/1]       via 172.16.1.1
C    10.0.1.0/24       [0/0]         via 10.0.1.1
C    10.0.1.1/32       [0/0]         via 0.0.0.0
C    10.0.2.0/24       [0/0]         via 10.0.2.1
C    10.0.2.1/32       [0/0]         via 0.0.0.0
C    10.0.3.0/24       [0/0]         via 10.0.3.1
C    10.0.3.1/32       [0/0]         via 0.0.0.0
O    E2  20.20.20.0/24       [110/0]       via 172.16.1.1
C    172.16.1.0/29     [0/0]         via 172.16.1.3
C    172.16.1.3/32     [0/0]         via 0.0.0.0

```

Figure 61 – OSPF status on Logical Distributed Router

```

vShield-edge-2-0> show ip ospf neighbor
Neighbor ID      Priority     Address        Dead Time   State
172.16.1.2       128         172.16.1.3     31          Full/DR
vShield-edge-2-0> show ip route

Codes: 0 - OSPF derived, i - IS-IS derived, B - BGP derived,
C - connected, S - static, L1 - IS-IS level-1, L2 - IS-IS level-2,
IA - OSPF inter area, E1 - OSPF external type 1, E2 - OSPF external type 2

Total number of routes: 8

S    0.0.0.0/0          [1/1]         via 20.20.20.1
O    E2  10.0.1.0/24         [110/1]       via 172.16.1.2
O    E2  10.0.2.0/24         [110/1]       via 172.16.1.2
O    E2  10.0.3.0/24         [110/1]       via 172.16.1.2
C    20.20.20.0/24     [0/0]         via 20.20.20.2
C    20.20.20.2/32     [0/0]         via 0.0.0.0
C    172.16.1.0/29     [0/0]         via 172.16.1.1
C    172.16.1.1/32     [0/0]         via 0.0.0.0

```

Figure 62 – OSPF status on Logical Centralized Router

Validate communication from internal to Centralized Router external interface

```
root@web-01:~# ping 20.20.20.2
PING 20.20.20.2 (20.20.20.2) 56(84) bytes of data:
64 bytes from 20.20.20.2: icmp_req=1 ttl=63 time=1.08 ms
64 bytes from 20.20.20.2: icmp_req=2 ttl=63 time=1.12 ms
64 bytes from 20.20.20.2: icmp_req=3 ttl=63 time=0.958 ms
```

Figure 63 – Communication from web-01 to Centralized Router external interface

Create many-to-one NAT (for traffic initiated from Web-Tier01 to external)

- 1) Add a NAT IP address to a Centralized Router external interface. From [NSX Home](#) -> [NSX Edges](#), select the Centralized Distributed Router and navigate to [Manage](#) -> [Settings](#) -> [Interfaces](#), Click Edit External interface and add IP address 20.20.20.3

vNIC#	1 ▲ Name	IP Address	Subnet Prefix Length	Connected To	Type
0	External	20.20.20.2* 20.20.20.3	24	vDS-External-PG	Uplink
1	Transit-01	172.16.1.1*	29	Transit-Network-01	Internal

Figure 64 – NAT IP address on External interface

- 2) Configure many-to-one NAT. Navigate to [Manage](#) -> [NAT](#), Add DNAT and Publish change.

Add SNAT Rule

Applied On: External

Original Source IP/Range: * 10.0.1.0/24

Translated Source IP/Range: * 20.20.20.3

Description: DNAT Web-Tier-01

☒ Enabled

☐ Enable logging

OK Cancel

Figure 65 – DNAT configuration for Web-Tier-01 subnet

Validate communication from Web-Tier-01 to Internet

```
root@web-01:~# ping 30.30.30.211
PING 30.30.30.211 (30.30.30.211) 56(84) bytes of data.
64 bytes from 30.30.30.211: icmp_req=1 ttl=61 time=1.85 ms
64 bytes from 30.30.30.211: icmp_req=2 ttl=61 time=2.31 ms
64 bytes from 30.30.30.211: icmp_req=3 ttl=61 time=2.25 ms
```

Figure 66 – ping from Web VM to Internet

Note: The VM traffic flow in the fabric is shown below:

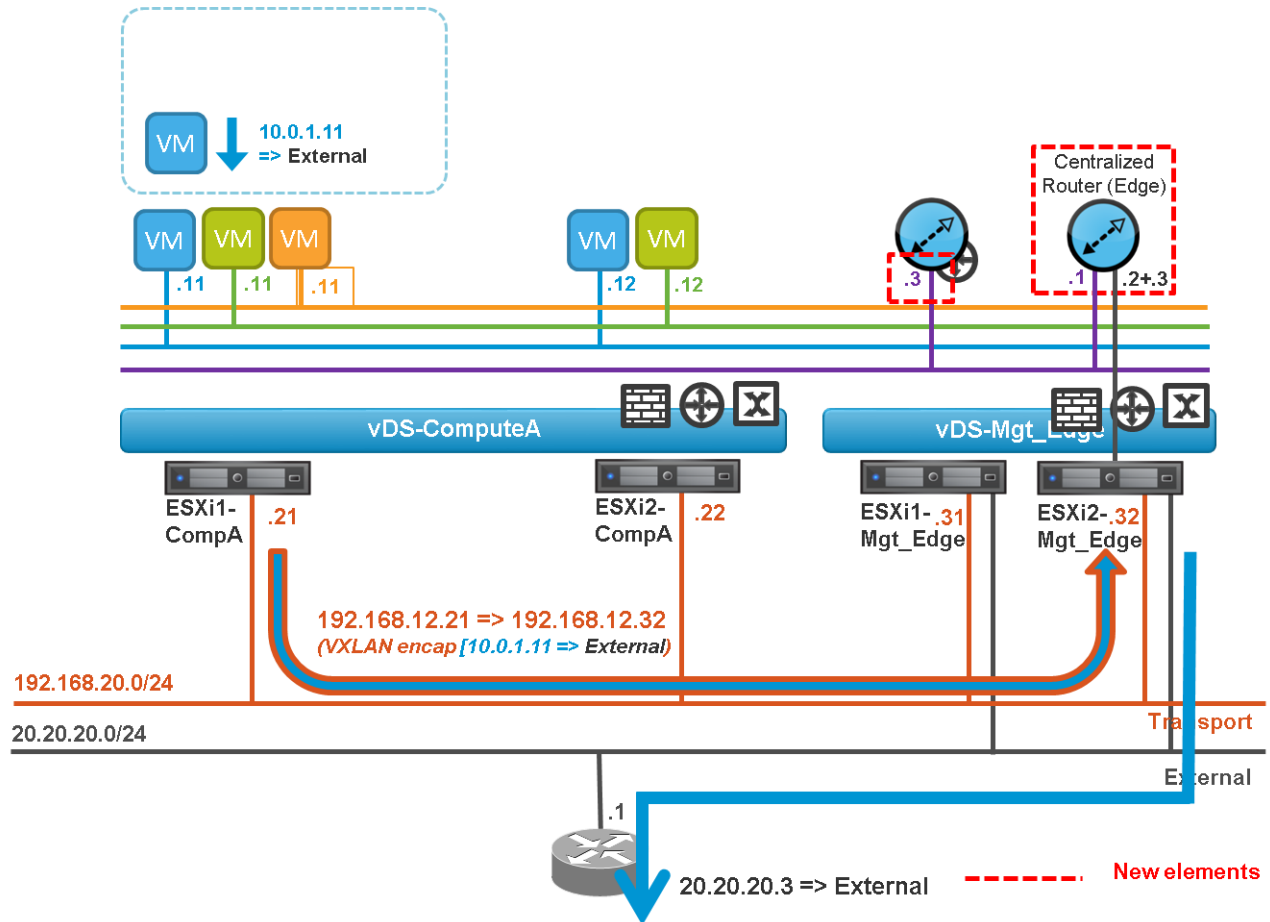


Figure 67 – Centralized Logical Router traffic flow

Logical Load Balancing

Goal of the Logical Load Balancing lab

In this step, you'll create a Logical Load Balancer (Edge) in one-arm mode.

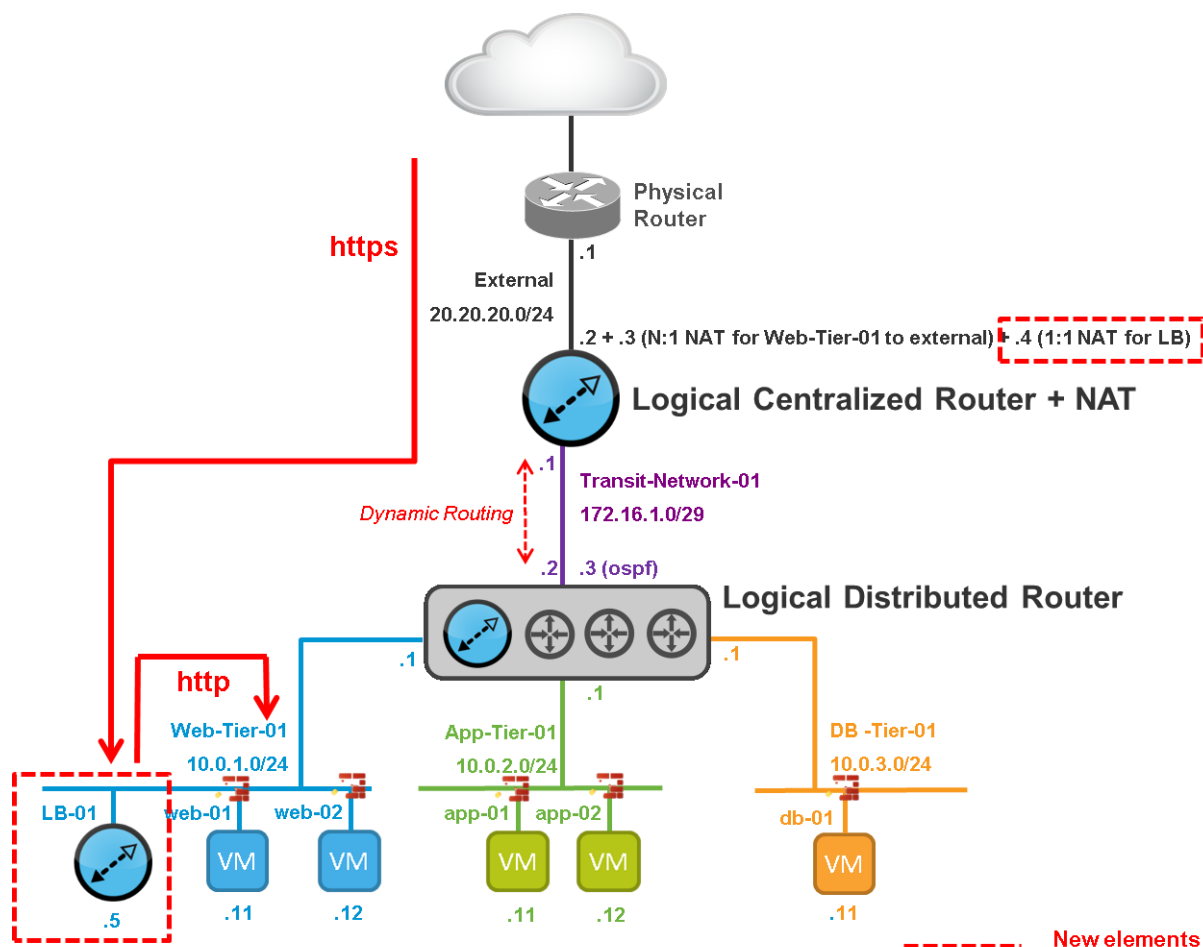


Figure 68 – Logical View Logical Load Balancer

The end-users access the VIP over https. The load balancer terminates https and talks to the servers over http.

Create one new Load Balancer

From **NSX Home -> NSX Edges**, create one Edge Service Gateway with one interface (LIF)

- Uplink to Web-Tier-01 with an IP of 10.0.1.5/24

The screenshot shows the 'New NSX Edge' configuration wizard. The left sidebar lists seven steps: 1 Name and description (selected), 2 CLI credentials, 3 Configure deployment, 4 Configure interfaces, 5 Default gateway settings, 6 Firewall and HA, and 7 Ready to complete. The main panel is titled 'Name and description'. It contains the following fields and options:

- Install Type:** Radio buttons for 'Edge Services Gateway' (selected), 'Logical (Distributed) Router', and 'Enable High Availability' (checkbox).
- Name:** Text field containing 'LoadBalancer-01'.
- Hostname:** Empty text field.
- Description:** Empty text area.
- Tenant:** Empty text field.

Below the 'Logical (Distributed) Router' option, there is a description: 'Provides Distributed Routing and Bridging capabilities.'

Figure 69 – Logical Distributer Router creation, first pane

The screenshot shows the 'New NSX Edge' configuration wizard, third pane: 'Configure deployment'. The left sidebar shows steps 1 through 7, with step 3 'Configure deployment' selected. The main panel contains the following configuration options:

- Datacenter:** Dropdown menu showing 'DC_Lab_NSX-v'.
- Appliance Size:** Radio buttons for 'Compact', 'Large' (selected), 'X-Large', and 'Quad Large'.
- Enable auto rule generation:** Checked checkbox.
- NSX Edge Appliances:** A table with columns: Resource Pool, Host, Datastore, and Folder.

The table contains one row of data:

Resource Pool	Host	Datastore	Folder
Cluster-Mgt_Ed...		OpenFiler	

Figure 70 – Logical Distributer Router creation, third pane

New NSX Edge

- 1 Name and description
- 2 CLI credentials
- 3 Configure deployment
- 4 Configure interfaces**
- 5 Default gateway settings
- 6 Firewall and HA

Configure interfaces

Configure interfaces of this NSX Edge

vNIC#	Name	IP Address	Subnet Prefix Length	Connected To
0	Web-01	10.0.1.5*	24	Web-Tier-01

Figure 71 – Logical Distributer Router creation, fourth pane

New NSX Edge

- 1 Name and description
- 2 CLI credentials
- 3 Configure deployment
- 4 Configure interfaces
- 5 Default gateway settings**
- 6 Firewall and HA
- 7 Ready to complete

Default gateway settings

☒ **Configure Default Gateway**

vNIC: * Web-01

Gateway IP: * 10.0.1.1

MTU: 1500

Figure 72 – Logical Distributer Router creation, fifth pane

New NSX Edge

- 1 Name and description
- 2 CLI credentials
- 3 Configure deployment
- 4 Configure interfaces
- 5 Default gateway settings
- 6 Firewall and HA**
- 7 Ready to complete

Firewall and HA

☒ **Configure Firewall default policy**

Default Traffic Policy: ☒ Accept ☐ Deny

Logging: ☐ Enable ☒ Disable

Configure HA parameters

Configuring HA parameters is mandatory for HA to work.

Figure 73 – Logical Distributer Router creation, sixth pane

Configure the Load Balancer

1. Enable Load Balancing
 - a. From **NSX Home** -> **NSX Edges**, select the Logical Load Balancer and navigate to **Manage** -> **Load Balancer** -> **Global Configuration**, click Edit and enable load balancer.

Edit Load balancer global configuration

☒ **Enable Loadbalancer**

☐ **Enable Service Insertion**

☐ **Acceleration Enabled**

☐ **Logging**

Log Level: Info

OK **Cancel**

Figure 74 – Enable Load Balancing

2. Create a self-signed certificate by navigating to **Manage -> Settings->Certificates**, add a new self-signed certificate clicking on:
 - a. Actions – Generate CSR

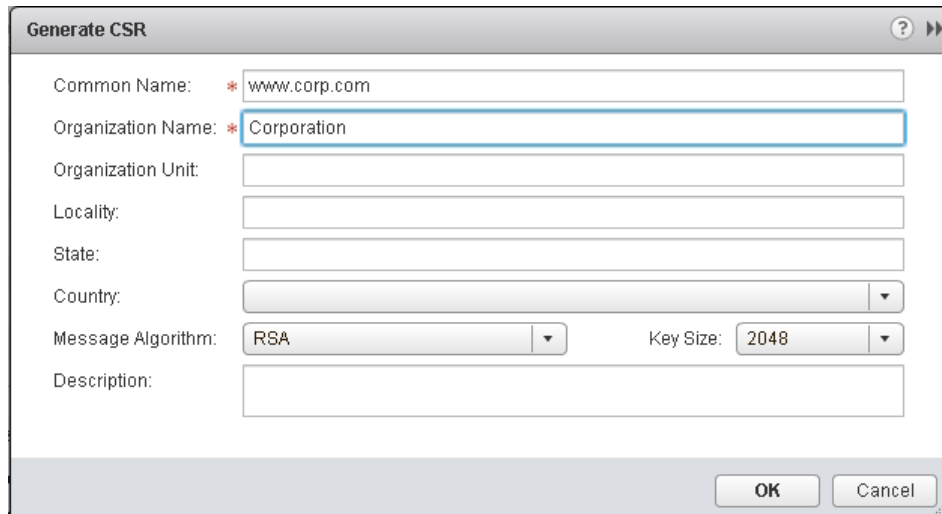
A screenshot of the 'Generate CSR' dialog box. It contains several input fields: 'Common Name' with the value 'www.corp.com', 'Organization Name' with the value 'Corporation', 'Organization Unit', 'Locality', 'State', and 'Country' (a dropdown menu). Below these are 'Message Algorithm' set to 'RSA' and 'Key Size' set to '2048'. There is also a 'Description' field. At the bottom right are 'OK' and 'Cancel' buttons.

Figure 75 – Certificate Signing Request (CSR)

- b. Actions – Self Sign Certificate

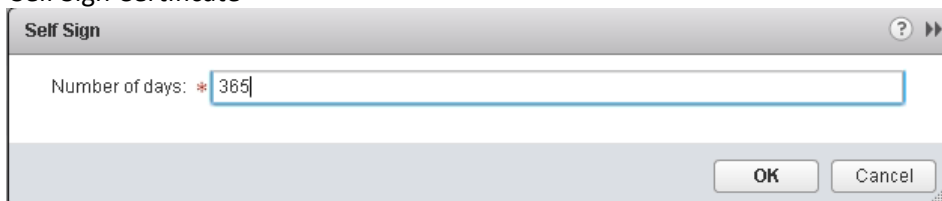
A screenshot of the 'Self Sign' dialog box. It contains a single input field 'Number of days' with the value '365'. At the bottom right are 'OK' and 'Cancel' buttons.

Figure 76 – Self Signing Certificate

3. Create an Application Profile
 - a. Navigate to **Manage -> Load Balancer -> Application Profiles**, add a new Application Profile with the following values:

New Profile

Name: Profile-HTTPS-Web-01

Type: ☐ TCP ☐ HTTP ☒ HTTPS

☐ Enable SSL Passthrough

HTTP Redirect URL:

Persistence: None

Cookie Name:

Mode:

☐ Insert X-Forwarded-For HTTP header

☐ Enable Pool Side SSL

Virtual Server Certificates **Pool Certificates**

Service Certificates CA Certificates CRL

	Common Name	Issuer	Validity
<input checked="" type="radio"/>	www.corp.com	www.corp.com	Tue May 20 2014

Figure 77 – Application Profile creation

4. Create a Server Pool by navigating to **Manage -> Load Balancer -> Pools**, adding a new Pool with the following values:

Edit Pool

Name: * Pool-Web-Tier-01

Description:

Algorithm: ROUND-ROBIN

Monitors: default_http_monitor

Members:

+ ✎ ✖


Enabled	Name	IP Address	Weight	Monitor Port	Port	Max Conn...	Min Conne...
✓	web-01	10.0.1.11	1	80	80	0	0
✓	web-02	10.0.1.12	1	80	80	0	0

☐ Transparent

OK Cancel

Figure 78 – Server Pool creation

5. Create VIP by navigating to **Manage -> Load Balancer -> Virtual Servers**, add a new VIP with the following values:



New Virtual Server

☒ Enabled

Name: * VIP-Web-Tier-01

Description: HTTPS SSL-Offload VIP

IP Address: * 10.0.1.5

Protocol: HTTPS

Port: * 443

Default Pool: Pool-Web-Tier-01

Application Profile: * Web-01-Profile

Figure 79 – VIP creation

Update the Distributed Firewall rules to allow Load Balancer-to-Web server communication

From **NSX Home -> Firewall**, update the rule “Inter Web-Tier-01” with the IP@ of the Load Balancer 10.0.1.5:

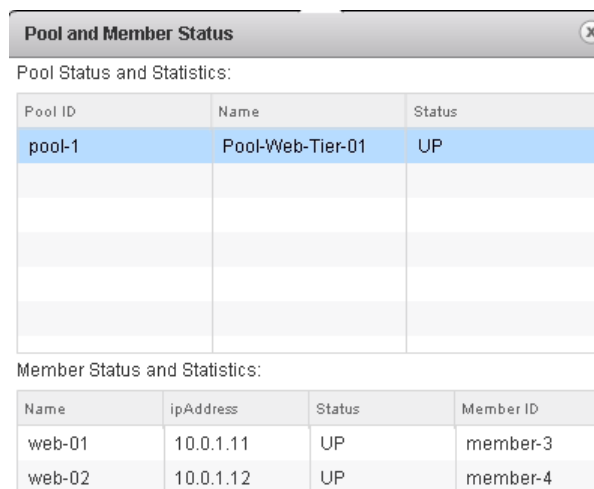
Note: You have to add the IP@ of the load balancer because the Edge doesn’t have the VM Tools.

No.	Name	Source	Destination	Service	Action	Applied To
▼ Application-01 (Rule 1 - 8)						
1	External access	any	Web-Tier-01	HTTPS	Allow	Web-Tier-01
2	Inter Web-Tier-01	10.0.1.5 Web-Tier-01	Web-Tier-01	SSH HTTP ICMP Echo	Allow	Web-Tier-01

Figure 80 – Updated Distributed Firewall rules

Validate that the Server Pool is UP

From **NSX Home -> NSX Edges**, select the Logical Load Balancer and navigate to **Manage -> Load Balancer -> Pools**, click **Show Pool Statistics** and validate the VIP is UP



Pool and Member Status

Pool Status and Statistics:

Pool ID	Name	Status
pool-1	Pool-Web-Tier-01	UP

Member Status and Statistics:

Name	ipAddress	Status	Member ID
web-01	10.0.1.11	UP	member-3
web-02	10.0.1.12	UP	member-4

Figure 81 – Server Pool status

Create a one-to-one NAT rule on the External Edge Router (for traffic initiated from external to load balancer)

1. Add NAT IP address to Centralized Router external interface

From [NSX Home](#) -> [NSX Edges](#), select the Logical Centralized Router and navigate to [Manage](#) -> [Settings](#) -> [Interfaces](#), click Edit External interface and add IP address 20.20.20.4

vNIC#	1 ▲	Name	IP Address	Subnet Prefix Length	Connected To	Type
0		External	20.20.20.2*			
			20.20.20.3			
			Show All			
1		Transit-01	172.16.1.1*			
2		vnic2				
3		vnic3				
4		vnic4				

Assigned IP Addresses: External	
List of IP Addresses assigned to External	
IP Address	Subnet Prefix Length
20.20.20.2*	24
20.20.20.3	
20.20.20.4	

Figure 82 – NAT IP address on External interface for VIP

2. Configure one-to-one NAT:

Navigate to [Manage](#) -> [NAT](#), Add DNAT and Publish change

Edit DNAT Rule	
Applied On:	External
Original IP/Range:	* 20.20.20.4
Protocol:	any
Original Port/Range:	any
Translated IP/Range:	* 10.0.1.5
Translated Port/Range:	any
Description:	DNAT for VIP Web-Tier-01
<input checked="" type="checkbox"/> Enabled	
<input type="checkbox"/> Enable logging	
<div>OK Cancel</div>	

Figure 83 – DNAT configuration for VIP

Check that external network hosts can communicate to VIP

```
root@VM-Internet:~# curl -k https://20.20.20.4/
<html><body><h1>It works!</h1>
<p>This is the default web page for this server.</p>
<p>The web server software is running but no content has been added, yet.</p>
</body></html>
```

Figure 84 – HTTPS access to VIP from external

Below, we depict the VM traffic flow in the fabric.

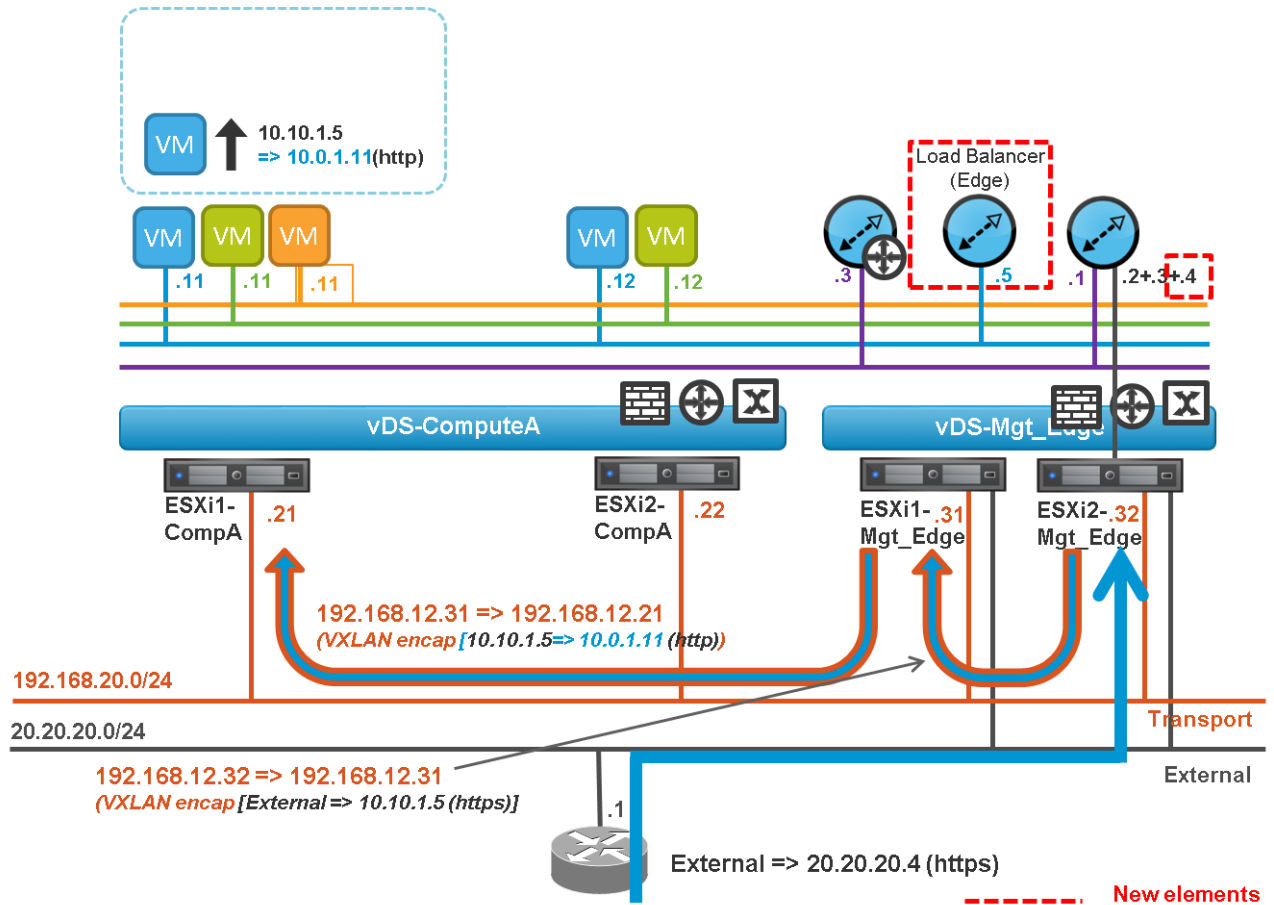


Figure 85 – Load Balancer traffic flow

Getting Help and More Information

NSX-v Documentation

In addition to this document, you can read the following documents for help setting up NSX-v. All are available from https://www-stage.vmware.com/support/pubs/nsx_pubs.html:

- NSX for vSphere Installation and Upgrade Guide
- NSX for vSphere Administration Guide
- NSX for vSphere API Reference Guide
- NSX for vSphere Command Line Interface Reference

Contacting the NSX Technical Services Team

You can reach the NSX technical services team at <http://www.vmware.com/support.html>.