

# NSX-T 3.0 Edge Design

## Step-by-Step UI Workflow

Version: 1.0

## Table of Contents

<b>NSX-T EDGE DESIGN STEP-BY-STEP UI WORKFLOW</b>	<b>3</b>
<b>1 Introduction</b>	<b>3</b>
<b>2 Edge Node Deployment Prerequisites</b>	<b>5</b>
2.1 VLAN's and IP's for Edge	5
2.2 NSX-T Transport zones	5
2.3 Edge Multi-TEP Uplink Profile	5
2.3.1 For Edge VM and BareMetal Edge with 2 Uplinks:	7
2.3.2 For BareMetal Edge with 4 PNIC	8
2.4 Compute Manager - vCenter - For Edge VM's only	9
2.5 VDS/VSS Port Groups (PG) - For Edge VM's only	10
2.5.1 VLAN Port Group:	10
2.5.2 First VLAN Trunk Port Group:	11
2.5.3 Second VLAN Trunk Port Group:	12
<b>3 Deploy and configure Edge with Single-NVDS Multi-TEP</b>	<b>14</b>
3.1 Edge VM Deployment and Configuration Workflow	14
3.2 BareMetal Edge Deployment and Configuration Workflow	18
3.2.1 For BareMetal Edge with 2 PNIC	19
3.2.2 For BareMetal Edge with 4 PNIC	20
<b>4 Preparing NSX-T for Layer 2 External (North-South) connectivity</b>	<b>22</b>
4.1 Edge Uplink Profile with 2 additional Named Teaming Policy	22
4.2 Associate Named Teaming Policy with the VLAN Transport Zone	23
4.3 Create External VLAN Segments with Named Teaming Policy	24
<b>5 Preparing NSX-T for Layer 3 External (North-South) connectivity</b>	<b>26</b>
5.1 Create NSX-T Tier-0 Gateway	26
5.2 Create External Interfaces on Tier-0 Gateway to connect to Physical Switches	27
5.3 Configure dynamic routing (BGP) between Tier-0 Gateway and physical switches	28
<b>6 Other Relevant Guides</b>	<b>32</b>

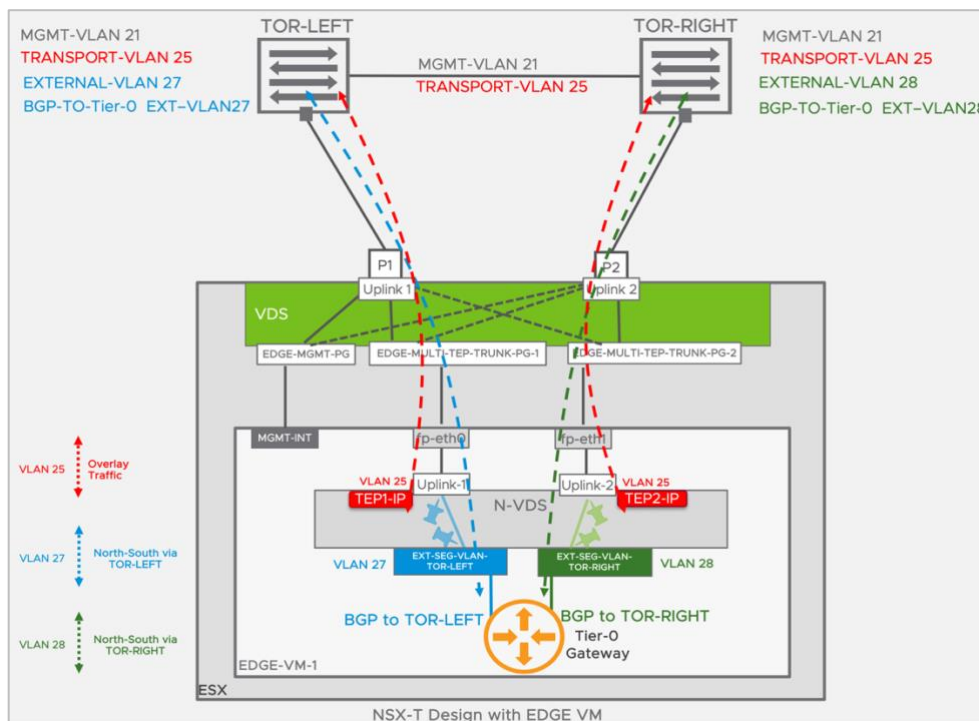
# NSX-T Edge Design Step-by-Step UI Workflow

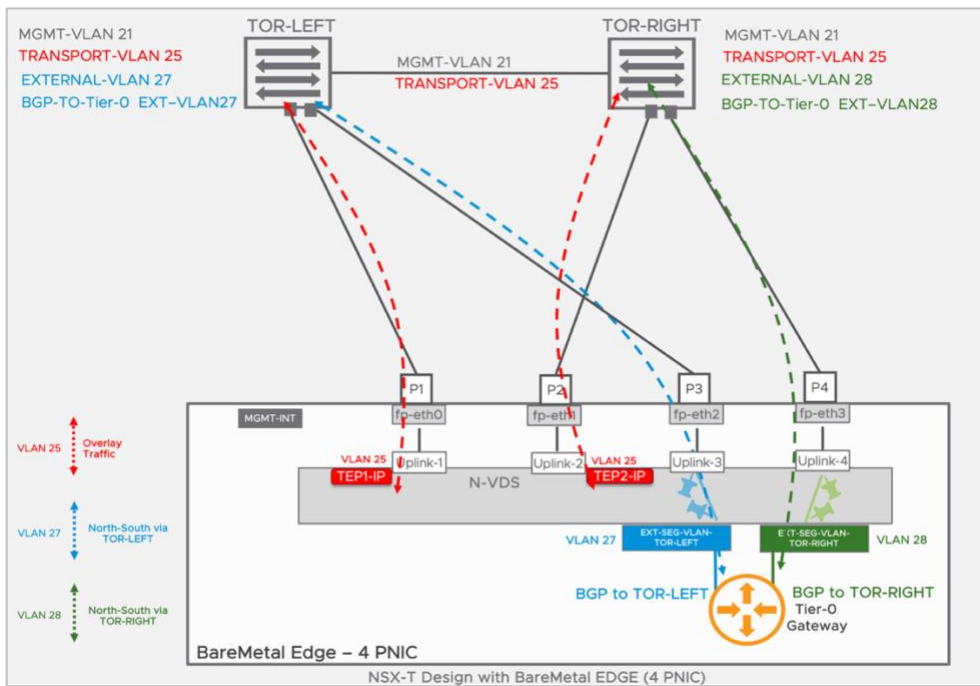
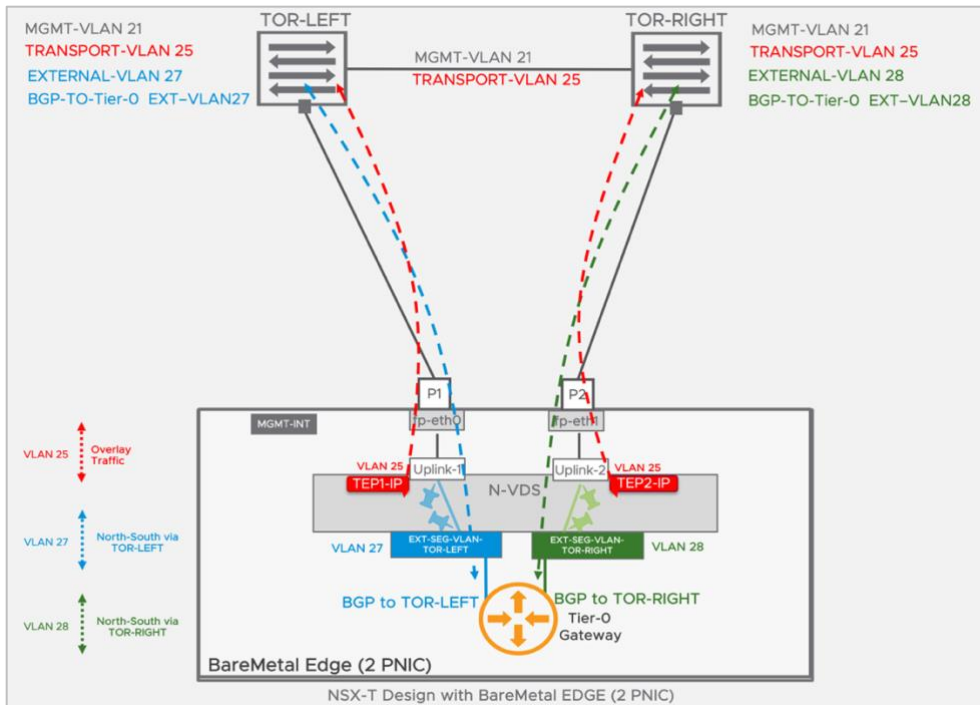
## 1 Introduction

This document is an informal document that walks through the step-by-step deployment and configuration workflow for NSX-T Edge Single N-VDS Multi-TEP design. This document assumes that the user is aware of the NSX-T basic concept. More information on the NSX-T concept, Edge Design can be found on the NSX-T design guide linked [here](#). NSX-T Edge Design options are covered in section 4.7, this document mainly focuses on design recommended in section 4.7.1.2 and 4.7.2.2.

The figure below shows the Edge design with Single N-VDS Multi-TEP for reference, Edge VM, BareMetal Edge with 2 PNIC and BareMetal Edge with 4 PNIC. This document walks through the step-by-step configuration workflow for this design using NSX-T 3.0 UI. This workflow is broken down into following 3 sub-workflows:

1. Deploy and configure the Edge node with Single-NVDS Multi-TEP.
2. Preparing NSX-T for Layer 2 External (North-South) connectivity
3. Preparing NSX-T for Layer 3 External (North-South) connectivity





## 2 Edge Node Deployment Prerequisites

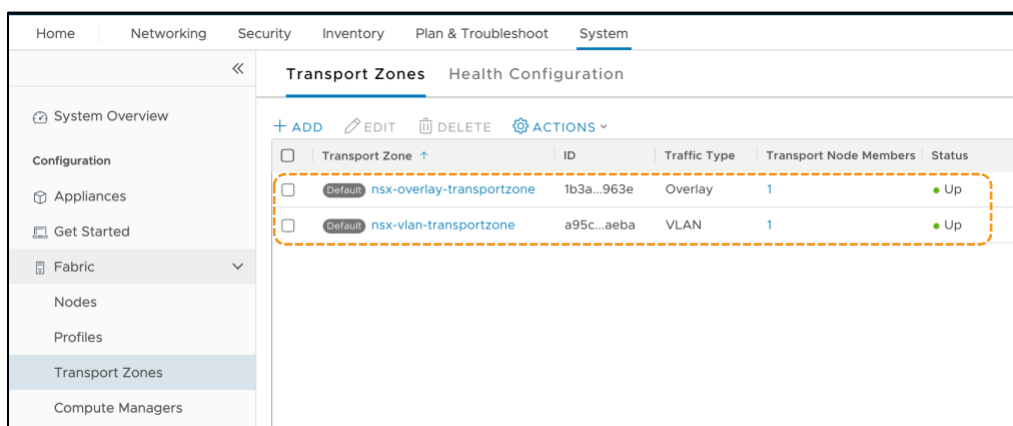
Before user starts deploying and configuring Edge Nodes/Cluster make sure following prerequisites are taken care of.

### 2.1 VLAN's and IP's for Edge

- **Management Interface** – One VLAN and an IP.
- **Overlay Tunnel Endpoints (TEP)** – 1 VLAN for Overlay Transport and N number of TEP IP's for Edge Multi-TEP deployment. N equal to number of uplinks used for Overlay Traffic. The design shown above uses 2 uplinks so two IP's needed for TEP.
- **External-Interfaces(s)**- This is needed with Tier-0 Gateway to connect to physical switches for North-South connectivity – BGP and data path. 2 VLAN's and 4 IP's for a 2 Tier-0 Gateway SR with 2 uplinks each.

### 2.2 NSX-T Transport zones

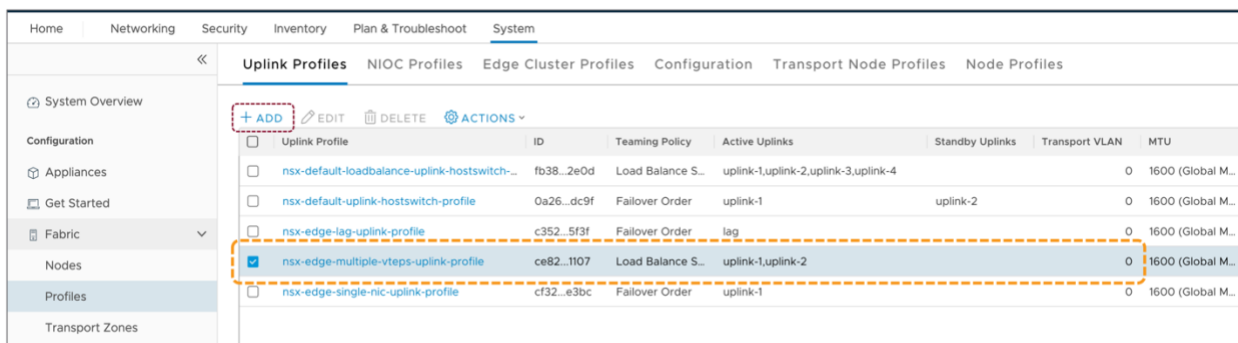
Need to have overlay and VLAN transport zone. The user can use the system created default transport zones or create its own. Recommendation is to use default transport zones.



Transport Zone	ID	Traffic Type	Transport Node Members	Status
Default nsx-overlay-transportzone	1b3a...963e	Overlay	1	Up
Default nsx-vlan-transportzone	a95c...aeba	VLAN	1	Up

### 2.3 Edge Multi-TEP Uplink Profile

Add an uplink profile which is similar to default edge profile highlighted below “[nsx-edge-multiple-vteps-uplink-profile](#)” with update to overlay transport VLAN id used (e.g 25). This VLAN used to send Overlay Traffic to and from Edge.



Uplink Profile	ID	Teaming Policy	Active Uplinks	Standby Uplinks	Transport VLAN	MTU
nsx-default-loadbalance-uplink-hostswitch-profile	fb38...2e0d	Load Balance S...	uplink-1,uplink-2,uplink-3,uplink-4		0	1600 (Global M...
nsx-default-uplink-hostswitch-profile	0a26...dc9f	Failover Order	uplink-1	uplink-2	0	1600 (Global M...
nsx-edge-lag-uplink-profile	c352...5f3f	Failover Order	lag		0	1600 (Global M...
nsx-edge-multiple-vteps-uplink-profile	ce82...1107	Load Balance S...	uplink-1,uplink-2		0	1600 (Global M...
nsx-edge-single-nic-uplink-profile	cf32...e3bc	Failover Order	uplink-1		0	1600 (Global M...

### New Uplink Profile

Name \*

Description

LAGs

+ ADD DELETE

<input type="checkbox"/> Name *	LACP Mode	LACP Load Balancing *	Uplinks *	LACP Time Out
No LAGs found				

Teamings

+ ADD CLONE DELETE

<input checked="" type="checkbox"/> Name *	Teaming Policy *	Active Uplinks *	Standby Uplinks
<input checked="" type="checkbox"/> [Default Teaming]	Load Balance Source	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.

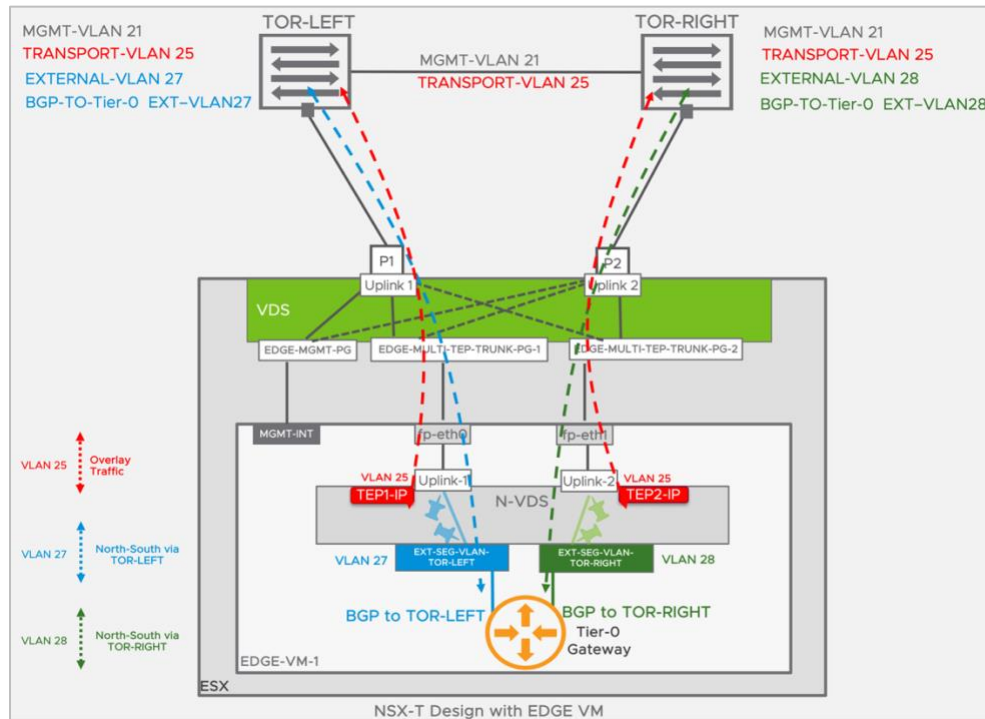
Transport VLAN

MTU

CANCEL ADD

Home	Networking	Security	Inventory	Plan & Troubleshoot	System	POLICY	MANAGER
Uplink Profiles NIOC Profiles Edge Cluster Profiles Configuration Transport Node Profiles Node Profiles							
+ ADD EDIT DELETE ACTIONS							
<input type="checkbox"/>	Uplink Profile	ID	Teaming Policy	Active Uplinks	Standby Uplinks	Transport VLAN	MTU
<input type="checkbox"/>	demo-nsx-edge-multiple-vteps-uplink-profile	a815...10...	Load Balance Source	uplink-1,uplink-2		25	1600 (Glo...
<input type="checkbox"/>	demo-nsx-host-uplink-hostswitch-profile	7237...2...	Load Balance Source	uplink-1,uplink-2		25	1600 (Glo...
<input type="checkbox"/>	nsx-default-loadbalance-uplink-hostswitch...	fb38...2...	Load Balance Source	uplink-1,uplink-2,uplink...		0	1600 (Glo...
<input type="checkbox"/>	nsx-default-uplink-hostswitch-profile	0a26...d...	Failover Order	uplink-1	uplink-2	0	1600 (Glo...
<input type="checkbox"/>	nsx-edge-lag-uplink-profile	c352...5...	Failover Order	lag		0	1600 (Glo...
<input type="checkbox"/>	nsx-edge-multiple-vteps-uplink-profile	ce82...11...	Load Balance Source	uplink-1,uplink-2		0	1600 (Glo...
<input type="checkbox"/>	nsx-edge-single-nic-uplink-profile	cf32...e...	Failover Order	uplink-1		0	1600 (Glo...

Also, update the Uplink profile with 2 named teamed policy. NSX-T named teaming policy with dedicated uplink is needed to make sure NSX-T Tier-0 Gateway can have deterministic and optimal point-to-point Layer 3 peering with TOR, in multi-physical top-of-rack (TOR) switch topology. Below Figure shows the named teaming policy usage with Blue and Green Layer 3 connectivity.

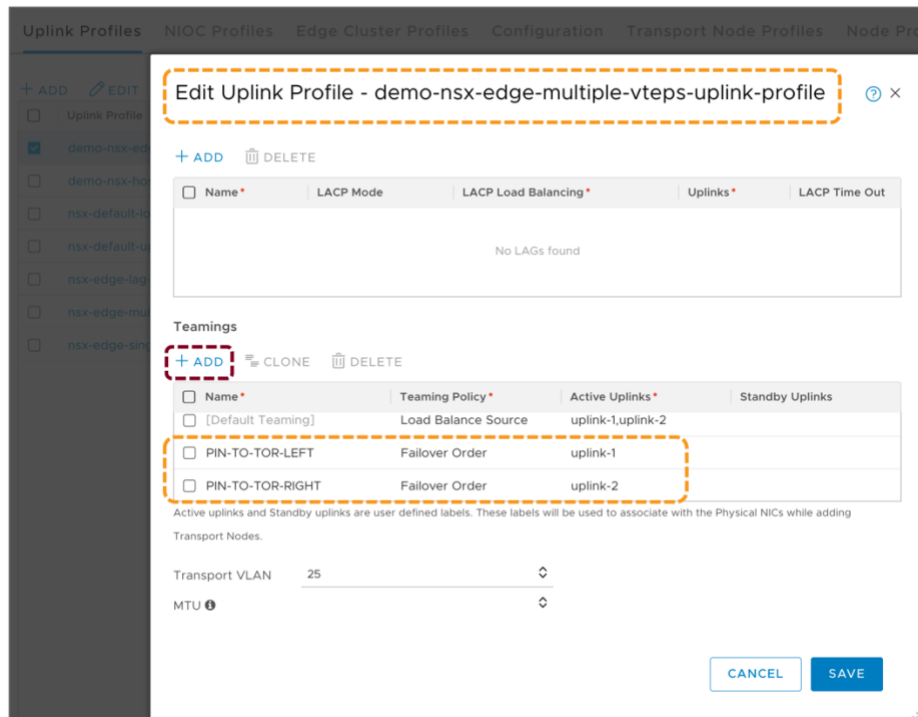


Based on number of Edge Active uplinks for named Teaming policy will change as described below.

### 2.3.1 For Edge VM and BareMetal Edge with 2 Uplinks:

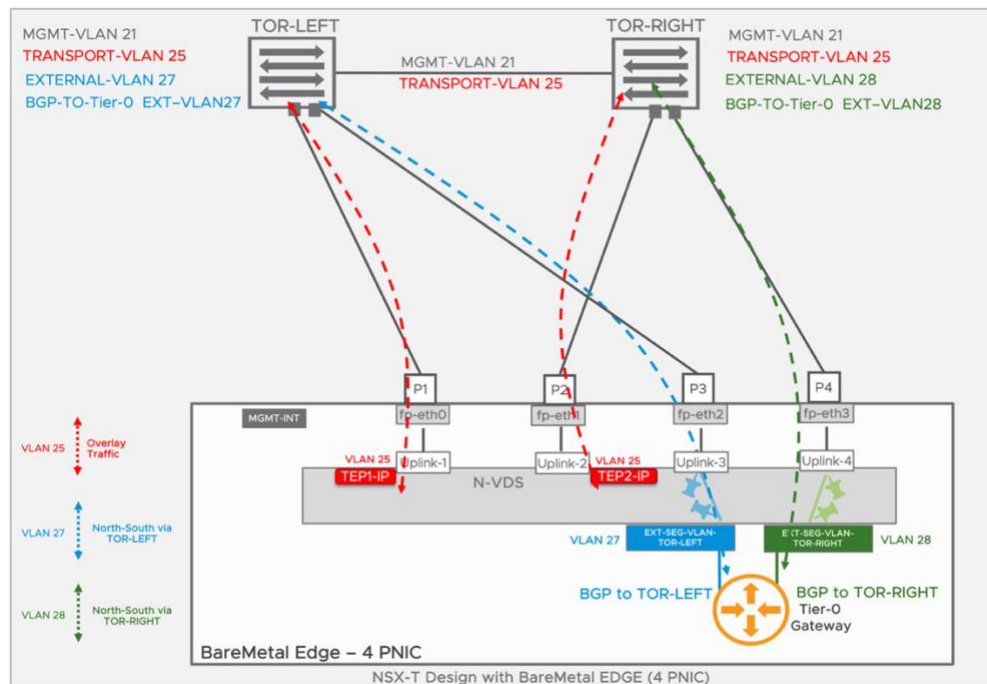
With 2 PNIC BareMetal Edge both Overlay traffic and North-South traffic share same 2 uplinks/PNICS. Overlay traffic uses default teaming policy to load balance traffic on both the uplinks. For North-South traffic need to add following named teaming policy:

- Under Teamings - Add 2 Teaming Policies: one with Active Uplink as **“uplink-1”** and other with **“uplink-2”**.
- Make a note of the policy name used, as we would be using this in the next section. In this example they are **“PIN-TO-TOR-LEFT”** and **“PIN-TO-TOR-RIGHT”**.



### 2.3.2 For BareMetal Edge with 4 PNIC

With 4 PNIC BareMetal Edge first 2 PNIC's are dedicated for Overlay traffic and other 2 are used for North-South traffic as shown in below figure.



For BareMetal Edge with 4 PNIC create a separate uplink profile as below. Profile creation follows same workflow except that user would use **uplink-3** and **uplink-4** as



Active uplinks respectively instead of uplink-1 and uplink-2, only for named teaming policy. Keep the default teaming as it is with uplink-1 and uplink-2. Make a note of the policy name used as we would be using this in the next section.

?

×

New Uplink Profile

Name \*

demo-nsx-BM-edge-4-pric-uplink-profile

Description

LAGs

+ ADD

DELETE

<input type="checkbox"/>	Name *	LACP Mode	LACP Load Balancing *	Uplinks *	LACP Time Out
No LAGs found					

Teamings

+ ADD

CLONE

DELETE

<input type="checkbox"/>	Name *	Teaming Policy *	Active Uplinks *	Standby Uplinks
<input type="checkbox"/>	[Default Teaming]	Load Balance Source	uplink-1,uplink-2	
<input type="checkbox"/>	PIN-TO-TOR-LEFT	Failover Order	uplink-3	
<input type="checkbox"/>	PIN-TO-TOR-RIGHT	Failover Order	uplink-4	

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

25

↕

MTU ⓘ

↕

CANCEL

ADD

## 2.4 Compute Manager - vCenter - For Edge VM's only

Register vCenter as a compute manager on the NSX-T manager. This helps to deploy and configure Edge VM using a built-in NSX-T UI wizard.

Alternatively, Edge VM's can be deployed using OVA directly without vCenter. However, this needs additional steps for Edge to join the management plane. Once joined to the management plane, the NSX-T UI wizard can be used to configure NSX. More details in the NSX-T installation guide.

Home	Networking	Security	Inventory	Plan & Troubleshoot	System	POLIC
------	------------	----------	-----------	---------------------	--------	-------

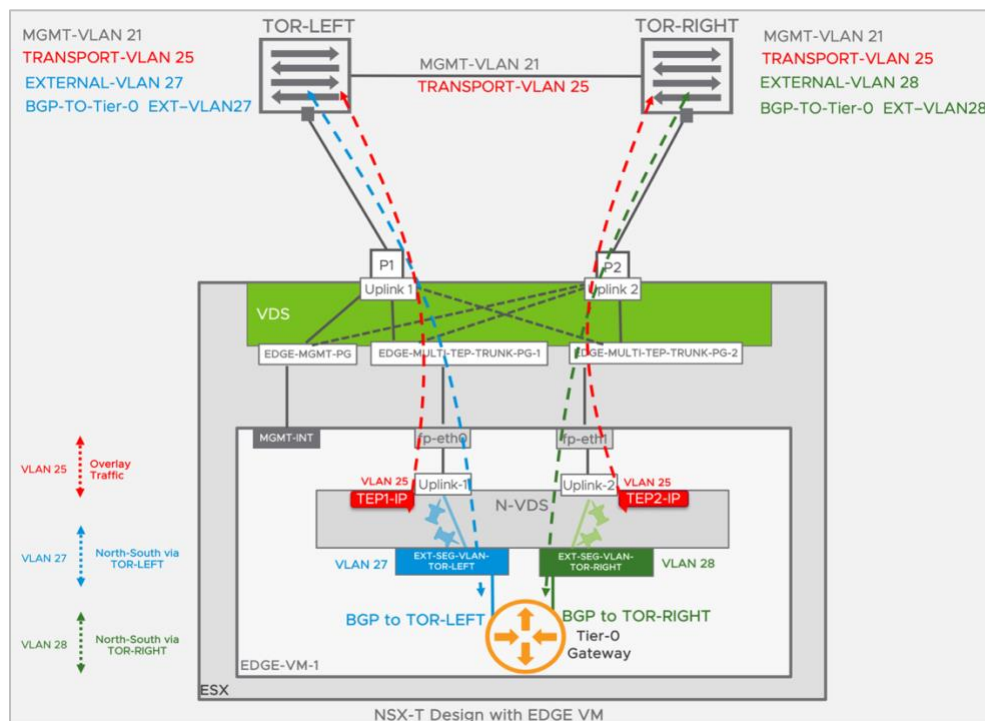
  

Compute Managers							
+ ADD	EDIT	DELETE	ACTIONS				
Search							
<input type="checkbox"/>	Compute Manager	ID	FQDN or IP Addr	Type	Registration Status	Version	Connection Status
<input type="checkbox"/>	DG-vCenter-01	2831...	10.114.213.4	vCenter	Registered	6.7.0	Up

## 2.5 VDS/VSS Port Groups (PG) - For Edge VM's only

This section assumes that Edge VM is hosted on ESX not configured as NSX-T transport Node. Edge VM optionally can be hosted on ESX part of NSX-T Transport-Node. In that case user need to create equivalent NSX-T VLAN segments, which is not covered in this document. If Edge VM is connected to NSX-Segments, in that case it is a requirement that Overlay Transport VLAN used for ESX host is different than the Transport VLAN used for Edge.

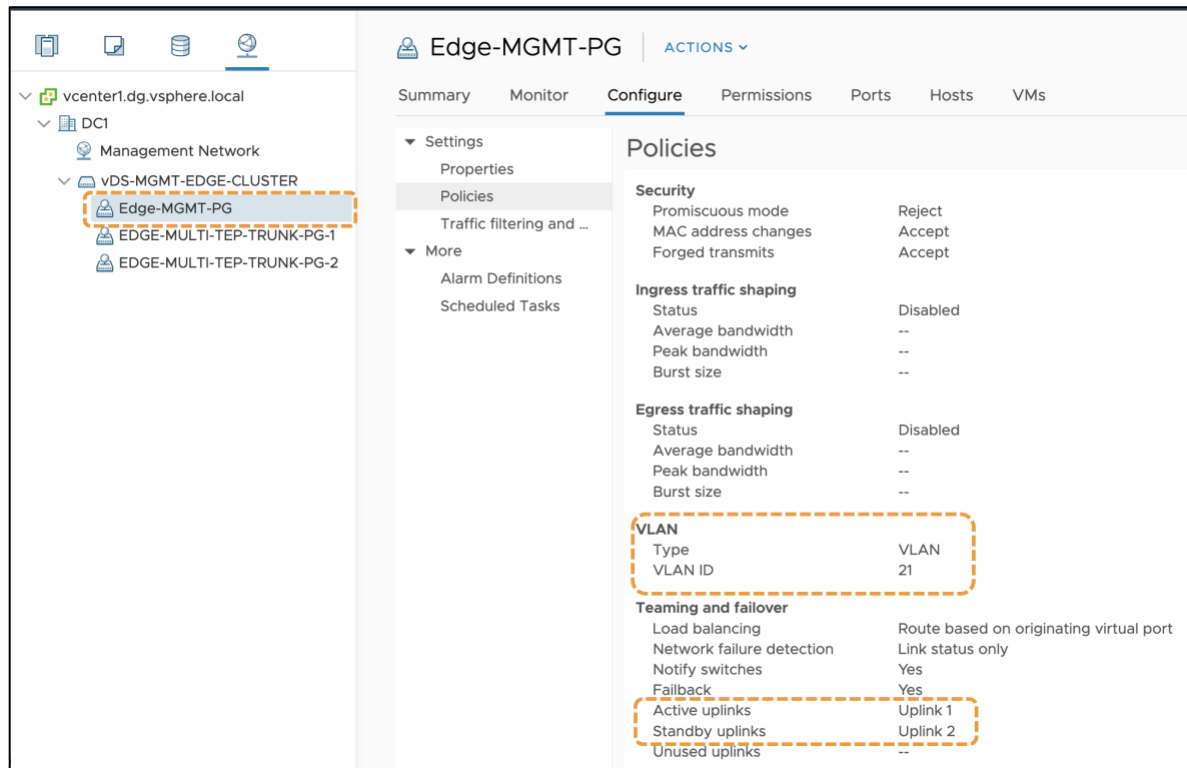
Need a total of 3 VSS/VDS port groups to connect Edge VM management and data path vNics as shown in the Edge VM design below:



### 2.5.1 VLAN Port Group:

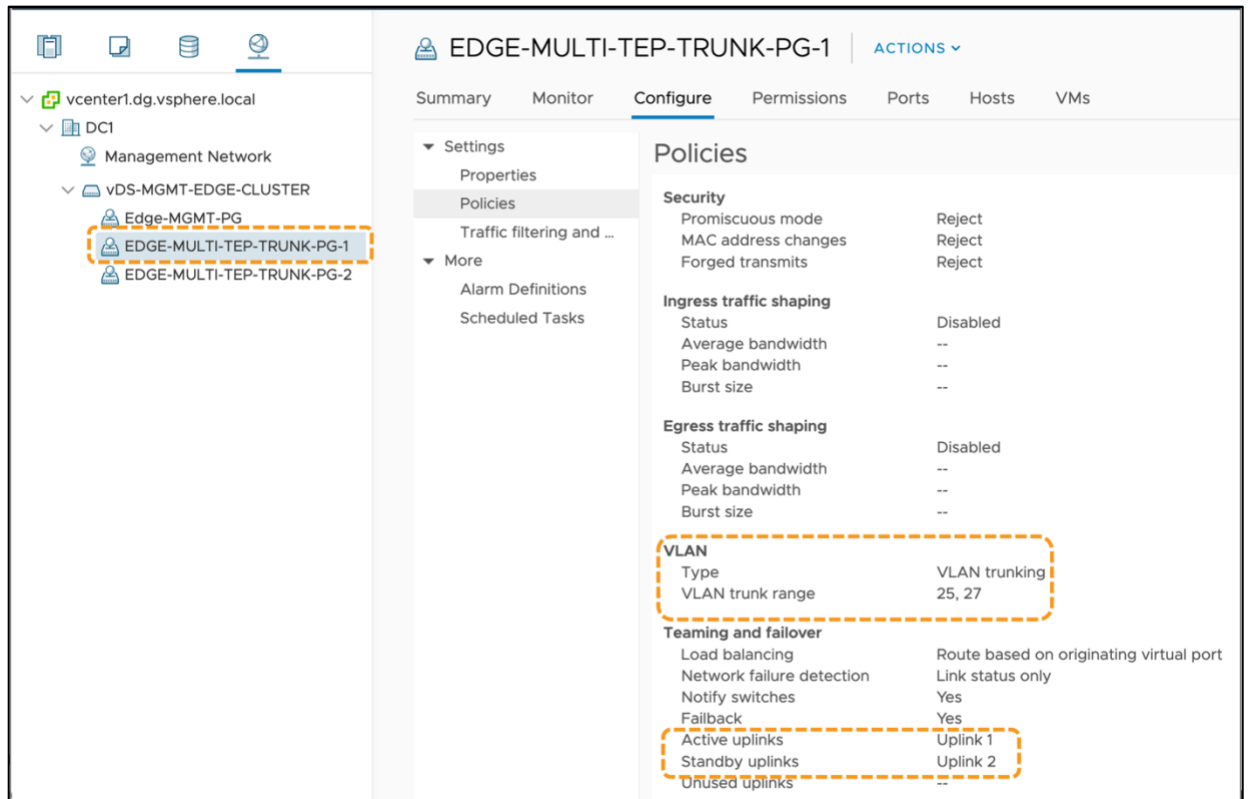
- With management VLAN ID (e.g. 21) for Edge Management.

- Active/Standby Teaming with Uplink-1/vmnic1 as Active and Uplink2/vmnic2 as standby.
- The first Edge VM interface (eth0), which is for management, is connected to this PG.



### 2.5.2 First VLAN Trunk Port Group:

- With Overlay Transport/TEP VLAN ID ( e.g. 25) and a VLAN ID (e.g. 27), **used to connect Tier-0 Gateway to Left Top Of Rack (TOR) Switch** for North-South connectivity.
- Active/Standby Teaming with **Uplink-1/vmnic1 as Active** and Uplink2/vmnic2 as standby.
- The second Edge VM interface, which is the first data path interface (fp-eth0), is connected to this PG.



### 2.5.3 Second VLAN Trunk Port Group:

- With Overlay Transport/TEP VLAN ID ( e.g. 25) and a VLAN ID (e.g. 28), **used to connect Tier-0 Gateway to Right Top Of Rack (TOR) Switch** for North-South connectivity.
- Active/Standby Teaming with **Uplink-2/vmnic2 as Active** and Uplink1/vmnic1 as standby.
- The third Edge VM interface, which is the second data path interface (fp-eth1), is connected to this PG.

▼ vcenter1.dg.vsphere.local

▼ DC1

Management Network

▼ vDS-MGMT-EDGE-CLUSTER

Edge-MGMT-PG

EDGE-MULTI-TEP-TRUNK-PG-1

EDGE-MULTI-TEP-TRUNK-PG-2

EDGE-MULTI-TEP-TRUNK-PG-2

ACTIONS ▼

SummaryMonitorConfigurePermissionsPortsHostsVMs

▼ Settings

Properties

Policies

Traffic filtering and ...

▼ More

Alarm Definitions

Scheduled Tasks

Policies

Security

Promiscuous modeAccept

MAC address changesAccept

Forged transmitsAccept

Ingress traffic shaping

StatusDisabled

Average bandwidth--

Peak bandwidth--

Burst size--

Egress traffic shaping

StatusDisabled

Average bandwidth--

Peak bandwidth--

Burst size--

VLAN

TypeVLAN trunking

VLAN trunk range25, 28

Teaming and failover

Load balancingRoute based on originating virtual port

Network failure detectionLink status only

Notify switchesYes

FailbackYes

Active uplinksUplink 2

Standby uplinksUplink 1

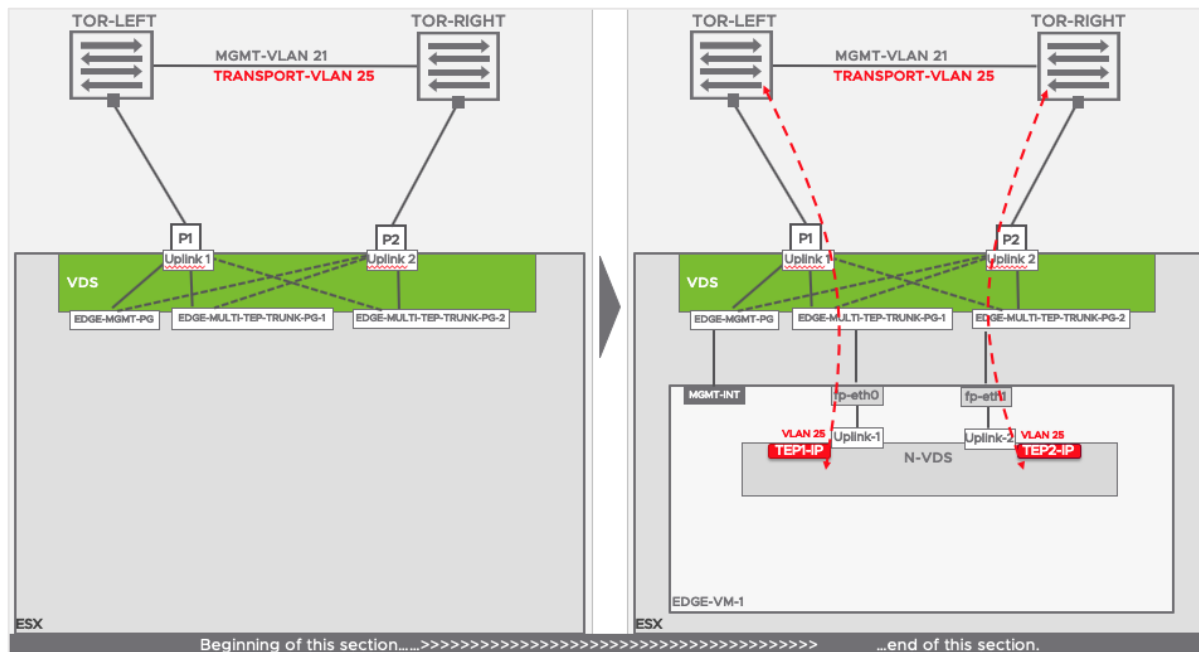
Unused uplinks--

### 3 Deploy and configure Edge with Single-NVDS Multi-TEP


NSX-T Edges comes in two form factors- VM and Bare metal. This section, walks through the Edge VM deployment and configuration first, followed by the BareMetal Edge configuration.

#### 3.1 Edge VM Deployment and Configuration Workflow

In this section, we will deploy Edge VM and configure NSX using NSX-T UI wizard. The figure below shows the start and end configuration state for this section.



Once the prerequisites are taken care of, on NSX-T UI navigate to System-> Fabric-> Nodes -> Edge Transport Nodes. And use the “ADD EDGE VM” wizard to deploy and configure Edge VM.



1. Provide Name and Description for the Edge VM and select form-factor based on the deployment need.

## Add Edge VM

- Name and Description**
- Credentials
- Configure Deployment
- Configure Node Settings
- Configure NSX

### Name and Description

Name \*

EDGE-VM-A01

Host name/FQDN \*

EDGE-VM-A01

Enter Fully Qualified Domain Name (FQDN)  
e.g. subdomain.example.com

Description

Form Factor \*

☐ Small

☒ Medium

☐ Large

☐ Extra Large

2 vCPU

4 vCPU

8 vCPU

16 vCPU

4 GB RAM

8 GB RAM

32 GB RAM

64 GB RAM

200 GB Storage

200 GB Storage

200 GB Storage

200 GB Storage

> Advanced Resource Reservations

CANCEL

NEXT

- Provide a password for built-in local user admin, root and audit. Optionally, enable the SSH login for the local users.

## Add Edge VM

- Name and Description
- Credentials**
- Configure Deployment
- Configure Node Settings
- Configure NSX

### Credentials

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

CLI Credentials

CLI User Name \*

admin

CLI Password \*

\*\*\*\*\*

CLI Confirm Password \*

\*\*\*\*\*

Allow SSH Login

☒ Yes

Root Credentials

System Root Password \*

\*\*\*\*\*

System Root Confirm Password \*

\*\*\*\*\*

Allow Root SSH Login

☐ No

> Audit Credentials

CANCEL

PREVIOUS

NEXT

- Select vCenter, cluster, and resources to deploy the Edge VM.

### Add Edge VM

- Name and Description
- Credentials
- Configure Deployment**
- Configure Node Settings
- Configure NSX

### Configure Deployment

Compute Manager \*

DG-vCenter-01

Cluster \*

MgtEdge\_Cluster

Resource Pool

Host

10.114.217.4

Datastore \*

NFS-DG-LAB-A01

Did not find expected? Try refresh to fetch latest datastores from [System](#)

CANCEL

PREVIOUS

NEXT

- Provide following Edge Management related info:
  - Edge Management IP, Gateway, NTP and DNS server.
  - Select Edge Management Port group created on vCenter for management vNic. Click to list and chose the port-groups available on chosen vCenter.

### Add Edge VM

- Name and Description
- Credentials
- Configure Deployment
- Configure Node Settings**
- Configure NSX

### Configure Node Settings

IP Assignment \*

☐ DHCP
 ☒ Static

Management IP \*

10.114.213.27/27

Default Gateway

10.114.213.1

Management Interface \*

Edge-MGMT-PG (Distributed...

Edge-MGMT-PG (Distributed Virtual Portgroup)

Search Domain Names

DNS Servers

10.114.213.16

NTP Servers

ntp.eng.vmware.com

CANCEL

PREVIOUS

NEXT

- This is the main NSX configuration page where the user configures Single N-VDS Multi-TEP configuration.
  - Select N-VDS and transport zone** - select default switch “nvds1” and default transport-zone – both overlay and VLAN.
  - Select Uplink Profile** – select the multi-TEP uplink profile created, as part of prerequisite, from the drop-down list.



C) **Provide TEP IP addresses**- Need 2 TEP IP's per Edge node, which associates with the 2 uplink interfaces – uplink-1, uplink-2. The user has the option to use Static IP's or NSX-T IP pool for TEP IP assignment.

D) **Associate Edge uplinks to the VDS Trunk port-groups** created as part of the prerequisite. Uplink-1 is associated with first Trunk PG and uplink-2 is associated with the second trunk Port-group.

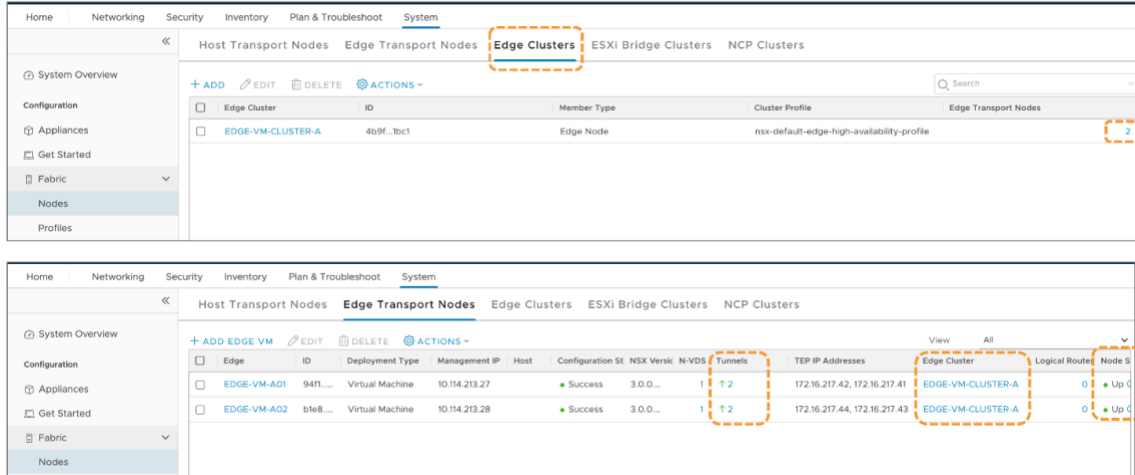
The screenshot shows the 'Add Edge VM' configuration wizard, specifically the 'Configure NSX' step. The wizard is divided into two main sections: 'Add Edge VM' on the left and 'Configure NSX' on the right. The 'Add Edge VM' section has a sidebar with steps 1 through 5, with step 5 'Configure NSX' selected. The 'Configure NSX' section has a '+ ADD SWITCH' button and a 'New Node Switch' section. The 'New Node Switch' section contains the following fields and options:

- Edge Switch Name:** nvds1
- Transport Zone:** nsx-vlan-transportzone, nsx-overlay-transportzone (with a dropdown arrow)
- Uplink Profile:** demo-nsx-edge-multiple-vteps-uplink-profile (with a dropdown arrow and 'OR Create New Transport Zone' link)
- IP Assignment:** Use Static IP List (with a dropdown arrow and 'OR Create New Uplink Profile' link)
- Static IP List:** 172.16.217.41, 172.16.217.42
- Gateway:** 172.16.217.33
- Subnet Mask:** 255.255.255.240
- Teaming Policy Switch Mapping:** A table with columns 'Uplinks' and 'DPDK Fastpath Interfaces'. It shows two rows: 'uplink-1 (active)' and 'uplink-2 (active)', both mapped to 'EDGE-MULTI-TEP-TRUNK-PG-1 (Distributed Virtual Portgroup)' and 'EDGE-MULTI-TEP-TRUNK-PG-2 (Distributed Virtual Portgroup)'. The 'uplink-1 (active)' row also shows 'fp-eth0' and 'fp-eth1' interfaces.

Annotations A, B, C, and D are placed on the left side of the 'Configure NSX' section, pointing to the 'Edge Switch Name', 'Uplink Profile', 'Static IP List', and 'Teaming Policy Switch Mapping' sections respectively. At the bottom right, there are 'CANCEL', 'PREVIOUS', and 'FINISH' buttons.

Similarly, deploy second Edge VM with its own name, management IP, TEP IP's, and keeping the rest of the options the same.

Once EDGE VM's are deployed, create an Edge Cluster, and add above created Edge VM's into the cluster. Once Edge VM's are in Cluster, it creates Overlay Tunnels between their TEP interfaces. Edge Transport Nodes page shows the health of the Edge under Node Status.



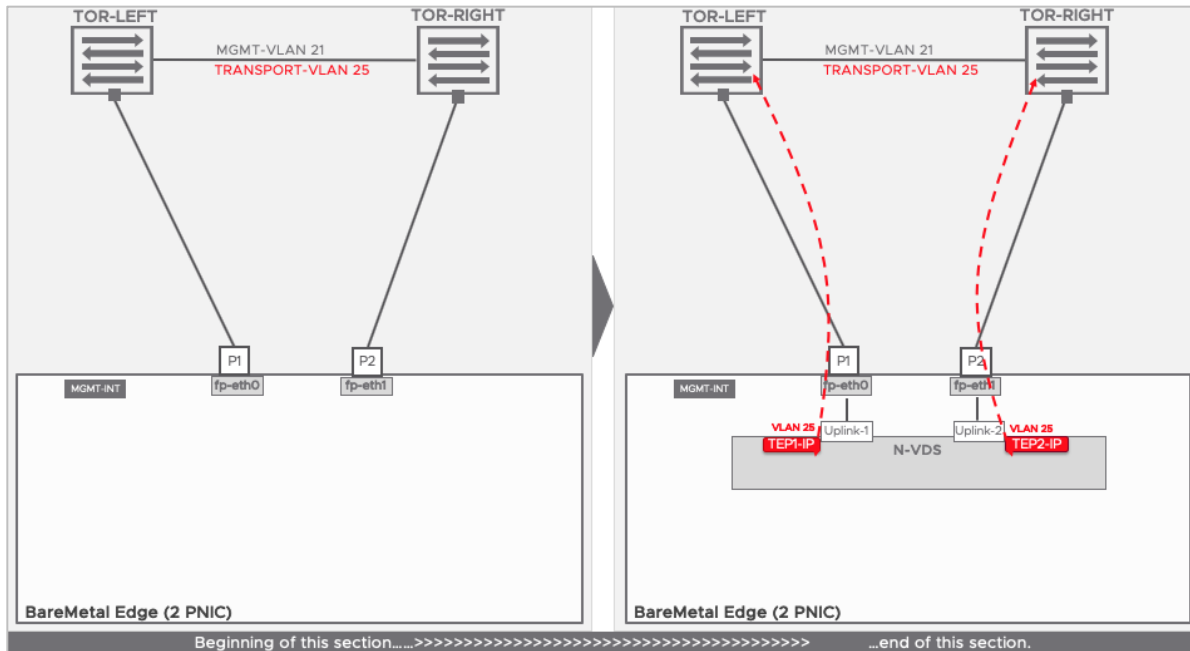
### 3.2 BareMetal Edge Deployment and Configuration Workflow

Unlike Edge VM installation, the BareMetal (BM) Edge deployment needs to be done using ISO image. Once, installation is done BareMetal edge can be joined to management plane by using CLI from the console or SSH session, as shown below. More info in the NSX-T installation guide.


```
nsx-edge> set hostname EDGE-BM-001
EDGE-BM-001>

EDGE-BM-001> join management-plane <NSX-T Mgr IP> username admin password <password> thumbprint 7b54c8a293de673c1d99fa7932899903b19f9c200060245e27d69fc28a38805
Node successfully registered as Fabric Node: 9078d30e-87a6-11ea-9e03-005056ad2994
EDGE-BM-001>
```

Once BareMetal Edge joins the management-plane NSX-T UI would have BM Edge as part of Edge Transport Nodes with configuration state as “Configure NSX”. The user can now use the “Configure NSX” wizard from UI to complete the BM Edge NSX configuration. The figure below shows BareMetal Edge before and after completing NSX configuration.



To configure NSX on BareMetal Edge, click on the “Configure NSX” as shown below.



In this configure NSX wizard select:

- N-VDS and transport zone** - select default switch “nvds1” and default transport-zone – both overlay and VLAN.
- Uplink Profile** – select the multi-TEP uplink profile created, as part of prerequisite, from the drop-down list.
- Provide TEP IP addresses** - Need 2 TEP IP’s per Edge, which associates with the 2 uplink interfaces – uplink-1, uplink-2. The user has the option to use Static IP’s or NSX-T IP pool for TEP IP assignment.
- Associate Edge uplinks to the fast-path interfaces (fp-ethx)** of BareMetal edge, as below based on number of PNICS.

### 3.2.1 For BareMetal Edge with 2 PNIC

Map uplinks to fp-eth interfaces as below. When user click on the NICs cell it provides drop down list of all available fp-ethx interfaces.

- uplink-1 to fp-eth0
- uplink-2 to fp-eth1

×

Edit Edge Transport Node - EDGE-BM-B01

Name \*

EDGE-BM-B01

Description

+ ADD SWITCH

▼ New Node Switch

Edge Switch Name \*

nvsd1

Transport Zone \*

nsx-vlan-transportzone

nsx-overlay-transportzone

Uplink Profile \*

demo-nsx-edge-multiple-vsteps-uplink-profile

IP Assignment \*

Use Static IP List

Static IP List \*

172.16.217.45

172.16.217.46

Gateway \*

172.16.217.33

Subnet Mask \*

255.255.255.240

OR Create New Transport Zone

OR Create New Uplink Profile

Teaming Policy Switch Mapping

Uplinks	Virtual NICs
🔗 uplink-1 (active)	fp-eth0
🔗 uplink-2 (active)	fp-eth1

CANCEL

SAVE

### 3.2.2

For BareMetal Edge with 4 PNIC

Bare Metal Edge with 4 PNIC topology is shown as in the Figure below.

The diagram illustrates the BareMetal Edge with 4 PNIC topology. It shows two TOR switches (TOR-LEFT and TOR-RIGHT) connected to a central N-VDS switch. The N-VDS switch has four physical network interface cards (P1, P2, P3, P4) connected to the TOR switches. The N-VDS switch also has four virtual network interface cards (fp-eth0, fp-eth1, fp-eth2, fp-eth3) connected to the TOR switches. The N-VDS switch is connected to a BGP to TOR-LEFT Gateway and a BGP to TOR-RIGHT Tier-0 Gateway. The diagram shows various VLANs and their connections: MGMT-VLAN 21, TRANSPORT-VLAN 25, EXTERNAL-VLAN 27, BGP-TO-Tier-0, and EXT-VLAN27. The diagram also shows the connections between the N-VDS switch and the TOR switches, including the use of uplink profiles and static IP lists.

Map uplink to fp-eth interfaces as below. When user click on the NICs cell it provides drop down list of all available fp-ethx interfaces.

- uplink-1 to fp-eth0
- uplink-2 to fp-eth1
- uplink-3 to fp-eth2
- uplink-4 to fp-eth3

Edit Edge Transport Node - EDGE-BM-B02

Name \* EDGE-BM-B02

Description

+ ADD SWITCH

▼ New Node Switch

Edge Switch Name \* nvdsl

Transport Zone \* nsx-overlay-transportzone nsx-vlan-transportzone

Uplink Profile \* demo-nsx-bm-edge-4-pric-uplink-profile

IP Assignment \* Use Static IP List

Static IP List \* 172.16.217.39 172.16.217.40

Gateway \* 172.16.217.33

Subnet Mask \* 255.255.255.240

Teaming Policy Switch Mapping

Uplinks	Virtual NICs
uplink-1 (active)	fp-eth0
uplink-2 (active)	fp-eth1
uplink-3 (active)	fp-eth2
uplink-4 (active)	fp-eth3

CANCEL SAVE

Similarly, configure second BareMetal Edge with its own TEP IP's, and keeping the rest of the options the same.


Once BareMetal Edges are deployed, create an Edge Cluster, and add above created BareMetal Edge into the cluster. Once BM Edge VM's are in Cluster, it creates Overlay Tunnels between their TEP interfaces. Edge Transport Nodes page shows the health of the Edge under Node Status.

#### 4 Preparing NSX-T for Layer 2 External (North-South) connectivity

This section walks through the NSX-T Layer 2 networking configuration to have North-South connectivity between NSX-T virtualized networks to Physical network. This is achieved using the following 3 configuration steps:

1. Update Edge Uplink Profile with 2 additional Named Teaming Policies.
2. Associate newly added Named Teaming Policy with the VLAN Transport Zone.
3. Create External VLAN Segments and associate it with Named Teaming Policy

The figure below shows before and after state of Layer 2 External Connectivity.



#### 4.1 Edge Uplink Profile with 2 additional Named Teaming Policy

NSX-T named teaming policy with dedicated uplink is needed to make sure NSX-T Tier-0 Gateway can have deterministic and optimal point-to-point Layer 3 peering with TOR, in multi-physical top-of-rack (TOR) switch topology.

The prerequisite section above has the workflow to create these. Make sure the Edge uplink profile used has 2 named teaming policies and make a note of the teaming policy name used. In this example it is “PIN-TO-TOR-LEFT” and “PIN-TO-TOR-RIGHT”.

On NSX-T UI, navigate to Uplink Profiles page under System -> Fabric -> Profiles and click on the Edge uplink profile to get the teaming policy name.

Uplink Profiles   NIOC Profiles   Edge Cluster Profiles   Configuration   Transport Node Profiles   Node Profiles

**Edit Uplink Profile - demo-nsx-edge-multiple-vsteps-uplink-profile** ⓘ ×

+ ADD   DELETE

<input type="checkbox"/> Name *	LACP Mode	LACP Load Balancing *	Uplinks *	LACP Time Out
No LAGs found				

**Teamings**

+ ADD   CLONE   DELETE

<input type="checkbox"/> Name *	Teaming Policy *	Active Uplinks *	Standby Uplinks
<input type="checkbox"/> [Default Teaming]	Load Balance Source	uplink-1,uplink-2	
<input type="checkbox"/> PIN-TO-TOR-LEFT	Failover Order	uplink-1	
<input type="checkbox"/> PIN-TO-TOR-RIGHT	Failover Order	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   25   ⇅

MTU ⓘ   ⇅

CANCEL   SAVE

#### 4.2 Associate Named Teaming Policy with the VLAN Transport Zone

On NSX-T UI, navigate to Transport Zone page under System -> Fabric. And edit the VLAN Transport Zone used with Edge configuration.

- Under Uplink Teaming Policy Names- add names of the Teaming Policy created above - "PIN-TO-TOR-LEFT" and "PIN-TO-TOR-RIGHT".

The screenshot shows the 'Edit Transport Zone' dialog for 'nsx-vlan-transportzone'. The dialog has the following fields and options:

- Name:** nsx-vlan-transportzone (highlighted with a dashed orange box)
- Description:** (empty text box)
- Switch Name:** nsxHostSwitch
- Traffic Type:**
  - ☐ Overlay
  - ☒ VLAN
- Uplink Teaming Policy Names:**
  - PIN-TO-TOR-LEFT (with a close button 'x')
  - PIN-TO-TOR-RIGHT (with a close button 'x')

At the bottom right of the dialog are 'CANCEL' and 'SAVE' buttons.

#### 4.3 Create External VLAN Segments with Named Teaming Policy

On NSX-T UI, navigate to Segments page under Networking and add a VLAN segment.

- Provide segment name, VLAN transport zone name, used earlier, from the list.
- Provide VLAN id (e.g. 27) used for Top of Rack Left connectivity.
- **Important config:** Chose Uplink Teaming Policy “PIN-TO-TOR-LEFT” to pin all traffic from this segment to uplink-1 and to TOR-Left.

Similarly, create the second segment with VLAN id (e.g. 28) associated with TOR Right connectivity. Make sure to associate Uplink Teaming Policy “PIN-TO-TOR-RIGHT”, to pin all traffic from this segment to uplink-2 and to TOR-Right.

NSX-T named uplink teaming policy with dedicated uplinks needs to be used to make sure NSX-T Tier-0 Gateway can have deterministic and optimal point-to-point Layer 3 peering with TOR, in multi-physical top-of-rack (TOR) switch topology.



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
EXT-SEG-Vlan27	None	nsx-vlan-transportzone	Gateway CIDR IPv4 CIDR e.g. 10.22.12.2/23 Gateway CIDR IPv6 CIDR e.g. fc7e:f206:db42::1/48 SET DHCP CONFIG	Set			

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 also need IP Sec

VLAN \*

27 X

Enter Fully Qualified Domain Name

Domain Name

Enter Fully Qualified Domain Name

Metadata Proxy

Set

Address Bindings

Set

Description

Description

VPN Tunnel ID

Uplink Teaming Policy

PIN-TO-TOR-LEFT

IP Address Pool

Select IP Pool

Replication Mode

Hierarchical Two-Tier replication

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

SEGMENT PROFILES

DHCP STATIC BINDINGS

SAVE

CANCEL


Home	Networking	Security	Inventory	Plan & Troubleshoot	System
SEGMENTS					
ADD SEGMENT					
Segment Name		Connectivity	Transport Zone	Subnets	
EXT-SEG-Vlan27		None	nsx-vlan-transportzone	VLAN	
L2 VPN		Not Set	VPN Tunnel ID		
VLAN		27	Uplink Teaming Policy		
Domain Name		Not Set	IP Address Pool		
Edge Bridges		0	Metadata Proxy		
Address Bindings		Not Set	Replication Mode		
Description		Not Set	Tags		
SEGMENT PROFILES					
DHCP STATIC BINDINGS					
Segment Name		Connectivity	Transport Zone	Subnets	
EXT-SEG-Vlan28		None	nsx-vlan-transportzone	VLAN	
L2 VPN		Not Set	VPN Tunnel ID		
VLAN		28	Uplink Teaming Policy		
Domain Name		Not Set	IP Address Pool		
Edge Bridges		0	Metadata Proxy		
Address Bindings		Not Set	Replication Mode		
Description		Not Set	Tags		
SEGMENT PROFILES					
DHCP STATIC BINDINGS					

## 5 Preparing NSX-T for Layer 3 External (North-South) connectivity

This section walks through the NSX-T networking configuration to have Layer 3 North-South connectivity between NSX-T networks to Physical network. This is achieved using the following 3 configuration steps:

1. Create NSX-T Tier-0 Gateway
2. Create External Interfaces on Tier-0 Gateway to connect to Physical Switches.
3. Configure dynamic routing (BGP) between Tier-0 Gateway and physical switches.

The figure below shows before and after finishing the Layer 3 External connectivity configuration.



## 5.1 Create NSX-T Tier-0 Gateway

On NSX-T UI, navigate to Tier-0 Gateways page under Networking and add a Tier-0 Gateway.

- Provide gateway name, HA Mode based on the design.
- Select the Edge Cluster created earlier from the drop-down list to deploy Tier-0 gateway Services.

Home | Networking | Security | Inventory | Plan & Troubleshoot | System

Tier-0 Gateways

ADD GATEWAY

Tier-0 Gateway Name	HA Mode	Linked Tier-1 Gateways
DC-Tier-0-GW	Active Active	

Edge Cluster: EDGE-VM-CLUSTER-A

Additional Settings

Route Distinguisher for VRF Gateways

Tags: Tag (Required) Scope (Optional) Max 30 allowed. Click (+) to save.

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

INTERFACES

ROUTING

BGP

ROUTE RE-DISTRIBUTION

SAVE CANCEL

5.2 Create External Interfaces on Tier-0 Gateway to connect to Physical Switches

Edit Tier-0 Gateways to add External Interfaces to connected to physical switches. In this scenario we will be creating:

- 2 External Interfaces on every Edge Node.
- One each belonging to each of the External Segments created, to connect to TOR-Left (EXT-SEG-Vlan27) and to TOR-Right (EXT-SEG-Vlan28).

In the below example, an external Interface is created on EDGE-VM-A01 and picked the External VLAN Segment, which is pinned to uplink-1 and to TOR-Left switch.

Tier-0 Gateways

ADD GATEWAY

DC-Tier-0-GW Active Active

IP Address Management No Dynamic IP Allocation

Edge Cluster EDGE-VM-CLUSTER-A

Additional Settings

Route Distinguisher for VRF Gateways

EVPN Settings

Tags Tag (Required) Scope (Optional) Max 30 allowed. Click (+) to save.

INTERFACES

External and Service Interfaces

MULTICAST

BGP

ROUTE RE-DISTRIBUTION

Set Interfaces

Tier-0 Gateways DC-Tier-0-GW Interfaces

ADD INTERFACE

Name	Type	IP Address / Mask	Connected To (Segment)
EXT-INT-A01-1	External	10.114.213.53/29	EXT-SEG-Vlan27

Enter IP Address Mask

CIDR e.g. IPv4 172.16.10.1/24 or IPv6 fc7e:f206:db42::1/48

Edge Node EDGE-VM-A01

PIM Disabled

Tags Tag (Required) Scope (Optional) Max 30 allowed. Click (+) to save.

SAVE CANCEL

Similarly create other 3 external interfaces on Tier-0 gateway. Below is the screenshot with all 4 external interfaces with its IP, External segment and Edge Node.

Once External interfaces are up and running, user can start configuring dynamic routing BGP on these interfaces, which is captured in next section.

Set Interfaces

Tier-0 Gateways

DC-Tier-0-GW

#interfaces

ADD INTERFACE

COLLAPSE ALL

Search


	Name	Type	IP Address / Mask	Connected To(Segment)	Status
⋮	EXT-INT-A01-1	External	10.114.213.53/29	EXT-SEG-Vlan27	Success
	Edge Node	EDGE-VM-A01	MTU	Not Set	VIEW STATISTICS
	Tags	0	PIM	Disabled	
	ND Profile	default	URPF Mode	Strict	
⋮	EXT-INT-A01-2	External	10.114.213.61/29	EXT-SEG-Vlan28	Success
	Edge Node	EDGE-VM-A01	MTU	Not Set	VIEW STATISTICS
	Tags	0	PIM	Disabled	
	ND Profile	default	URPF Mode	Strict	
⋮	EXT-INT-A02-1	External	10.114.213.54/29	EXT-SEG-Vlan27	Success
	Edge Node	EDGE-VM-A02	MTU	Not Set	VIEW STATISTICS
	Tags	0	PIM	Disabled	
	ND Profile	default	URPF Mode	Strict	
⋮	EXT-INT-A02-2	External	10.114.213.62/29	EXT-SEG-Vlan28	Success
	Edge Node	EDGE-VM-A02	MTU	Not Set	VIEW STATISTICS
	Tags	0	PIM	Disabled	
	ND Profile	default	URPF Mode	Strict	

CLOSE

5.3 Configure dynamic routing (BGP) between Tier-0 Gateway and physical switches  
NSX-T BGP Configuration involves following:

- Base BGP configuration – Local AS, ECMP, Timers, etc.
- First BGP neighbor to TOR-Left switch, use corresponding External Interfaces IP as a BGP source
- Second BGP neighbor to TOR-Right switch, use corresponding External Interfaces IP as a BGP source
- Enable BGP route Re-Distribution on Tier-0 Gateway to advertise the Tier-0 and Tier-1 routes.

On TOR-Left and TOR-Right switches- configure 2 BGP neighbors each with External Interfaces IP, in the given External Segment pinned to that switch, as a neighbor.



BGP configuration done under Tier-0 gateway. Under Tier-0 gateway > BGP >

- Set base BGP configuration - Local-as, timers, ECMP, etc.
- Click “Set” next to BGP neighbors to configure BGP neighbors.

The screenshot shows the 'Tier-0 Gateways' configuration page. The 'DC-Tier-0-GW' is configured with 'Active Active' HA Mode. The 'Edge Cluster' is set to 'EDGE-VM-CLUSTER-A'. The 'INTERFACES' section shows 4 external and service interfaces. The 'ROUTING' section is expanded, showing 'BGP' settings. The 'Local AS' is set to '65001'. The 'BGP' checkbox is checked. The 'Graceful Restart' is set to 'Helper Only' with a timer of '180' seconds. The 'Route Aggregation' is set to 'Set'. The 'BGP Neighbors' section is highlighted with a red dashed box, and the 'Set' button is circled in red. Other settings include 'Inter SR BGP' (On), 'ECMP' (On), 'Multipath Relax' (On), and 'Graceful Restart Stale Timer' (600 seconds).

- The below example shows the BGP neighbor configuration to TOR-Left switch with its IP, Remote-AS, and source addresses as External Interfaces belonging to Segment connected to TOR-Left switch. Enable BFD, which is recommended for faster convergence.

Set BGP Neighbors

Tier-0 Gateways DC-Tier-0-GW #Neighbors 1

ADD BGP NEIGHBOR EXPAND ALL Search

IP Address	BFD	Remote AS number	Route Filter	Allows-in	Status
10.114.213.49	Enabled	6 E.g. 65000	1	Disabled	
Source Addresses: 10.114.213.53 X, 10.114.213.54 X Select IP Address					
Max Hop Limit: 1					
Graceful Restart: Helper Only					
Description: Description					

> TIMERS & PASSWORD

SAVE CANCEL

- Similarly configure other BGP neighbors to TOR-Right with the corresponding neighbor and source IP's. Below is the configuration snapshot of BGP configuration with status showing "Success".

Set BGP Neighbors

Tier-0 Gateways DC-Tier-0-GW #Neighbors 2

ADD BGP NEIGHBOR COLLAPSE ALL Search

IP Address	BFD	Remote AS number	Route Filter	Allows-in	Status
10.114.213.49	Enabled	6	1	Disabled	Success
Source Addresses: 10.114.213.53, 10.114.213.54 View Less					
Max Hop Limit: 1					
Graceful Restart: Helper Only					
Description: Not Set					
> TIMERS & PASSWORD					
10.114.213.57	Enabled	6	1	Disabled	Success
Source Addresses: 10.114.213.61, 10.114.213.62 View Less					
Max Hop Limit: 1					
Graceful Restart: Helper Only					
Description: Not Set					
> TIMERS & PASSWORD					

- BGP neighbor/status details can be found by clicking on the info icon in the status column. On BGP status page, the down arrow after Edge Node name can be used to select other edge node and BGP neighbor status on that Edge node.

## BGP Connectivity Status | 10.114.213.49

Edge Node

5f2e7f68-87b7-11ea-a62c-005056ad7564

Connection State	ESTABLISHED
Source Address	10.114.213.54
Local Port	179
Remote Port	59473
Messages Received	2009
Messages Sent	1740
Time since established	36028 milliseconds
Total In Prefix Count	285
Total Out Prefix Count	285
Connection Drop Count	1
Established Connection Count	2
Announced Capabilities	addPath IPv4 Unicast Rx, IPv4 Unicast, routeRefresh, gracefulRestart, 4byteAs
Negotiated Capability	IPv4 Unicast, routeRefresh
Neighbor Router ID	10.114.211.212

CLOSE

- Enable BGP route Re-Distribution on Tier-0 Gateway to advertise the Tier-0 and Tier-1 routes to BGP neighbors.

## 6 Other Relevant Guides

[NSX-T Design Guide](#)

[ESXi Transport Node Preparation](#)