CloudStack | COLLABORATION CONFERENCE

NSX Integration on Apache CloudStack

Nicolas Vazquez, Shapeblue

nicolas.vazquez@shapeblue.com

November 20 - 22, 2024 | Madrid, Spain

Agenda

- SDN Overview
- VMware NSX
- NSX Integration on CloudStack 4.20
- CloudStack Zone Creation
- NSX-backed VPCs
- Demo
- Conclusions
- Questions

About me

- Nicolas Vazquez
- nicolas.vazquez@shapeblue.com
- nvazquez@apache.org

- Senior Software Engineer at Shapeblue
- Apache CloudStack Committer and PMC Member
- Dad (x1), tennis & football fan

SDN Overview

SDN Overview

Decoupled Architecture:

SDN separates the control plane (decision-making) from the data plane (forwarding), enabling centralized network control and programmability.

Virtualization:

SDN leverages virtualization to create software-based network functions, making the network more agile and adaptable.

Programmability:

SDN provides APIs and programming interfaces to automate network configuration, management, and optimization, reducing manual intervention.

Centralized Control:

The SDN controller acts as a central brain, managing and orchestrating the entire network infrastructure.

SDN Overview

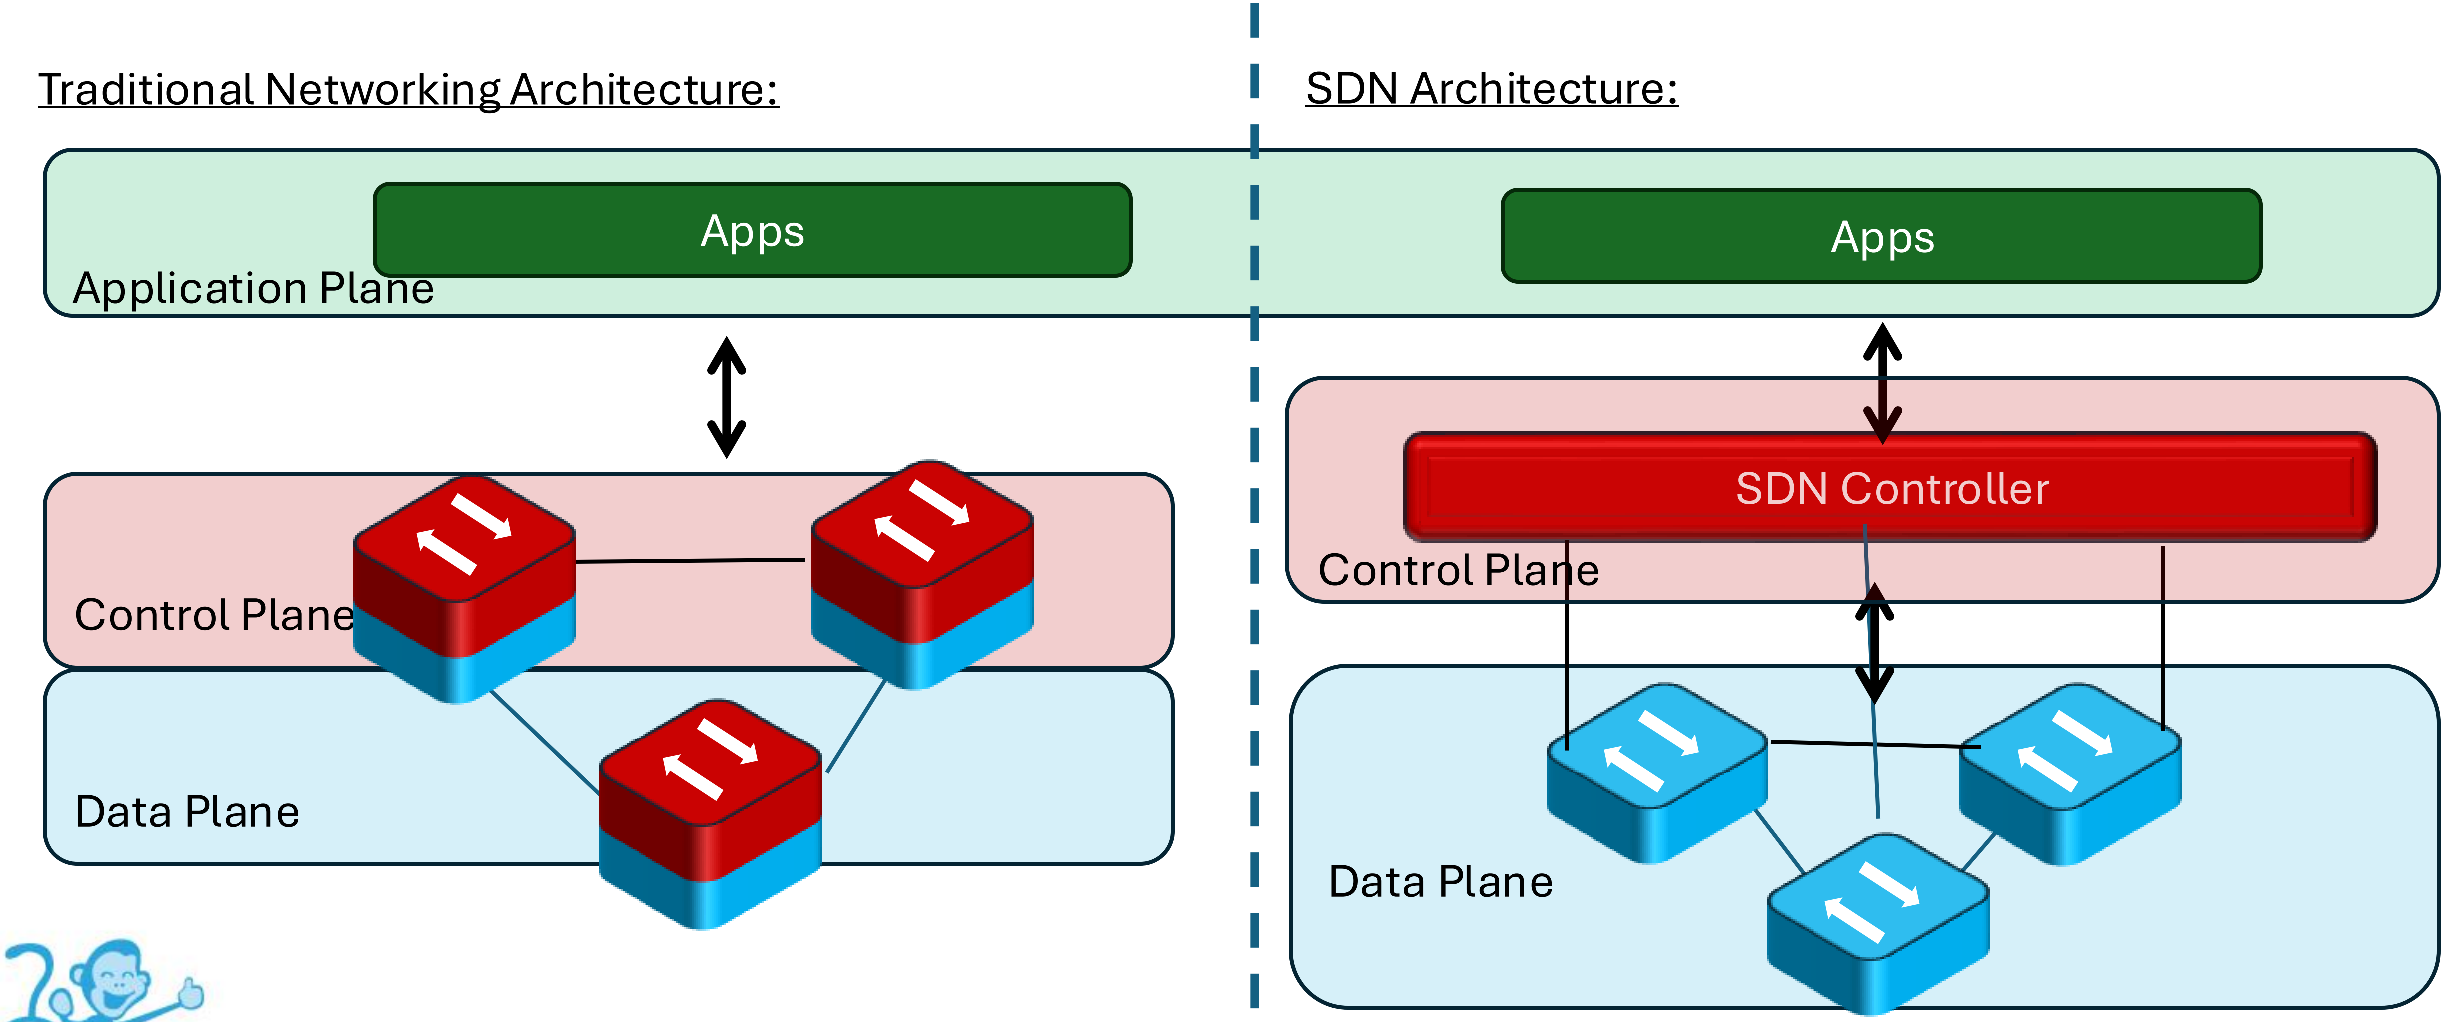

Recently Supported SDNs on CloudStack

Tungsten Fabric

Since: 4.18.0

VMware NSX

Since: 4.20.0

Netris

In progress

VMware NSX

VMware NSX Components

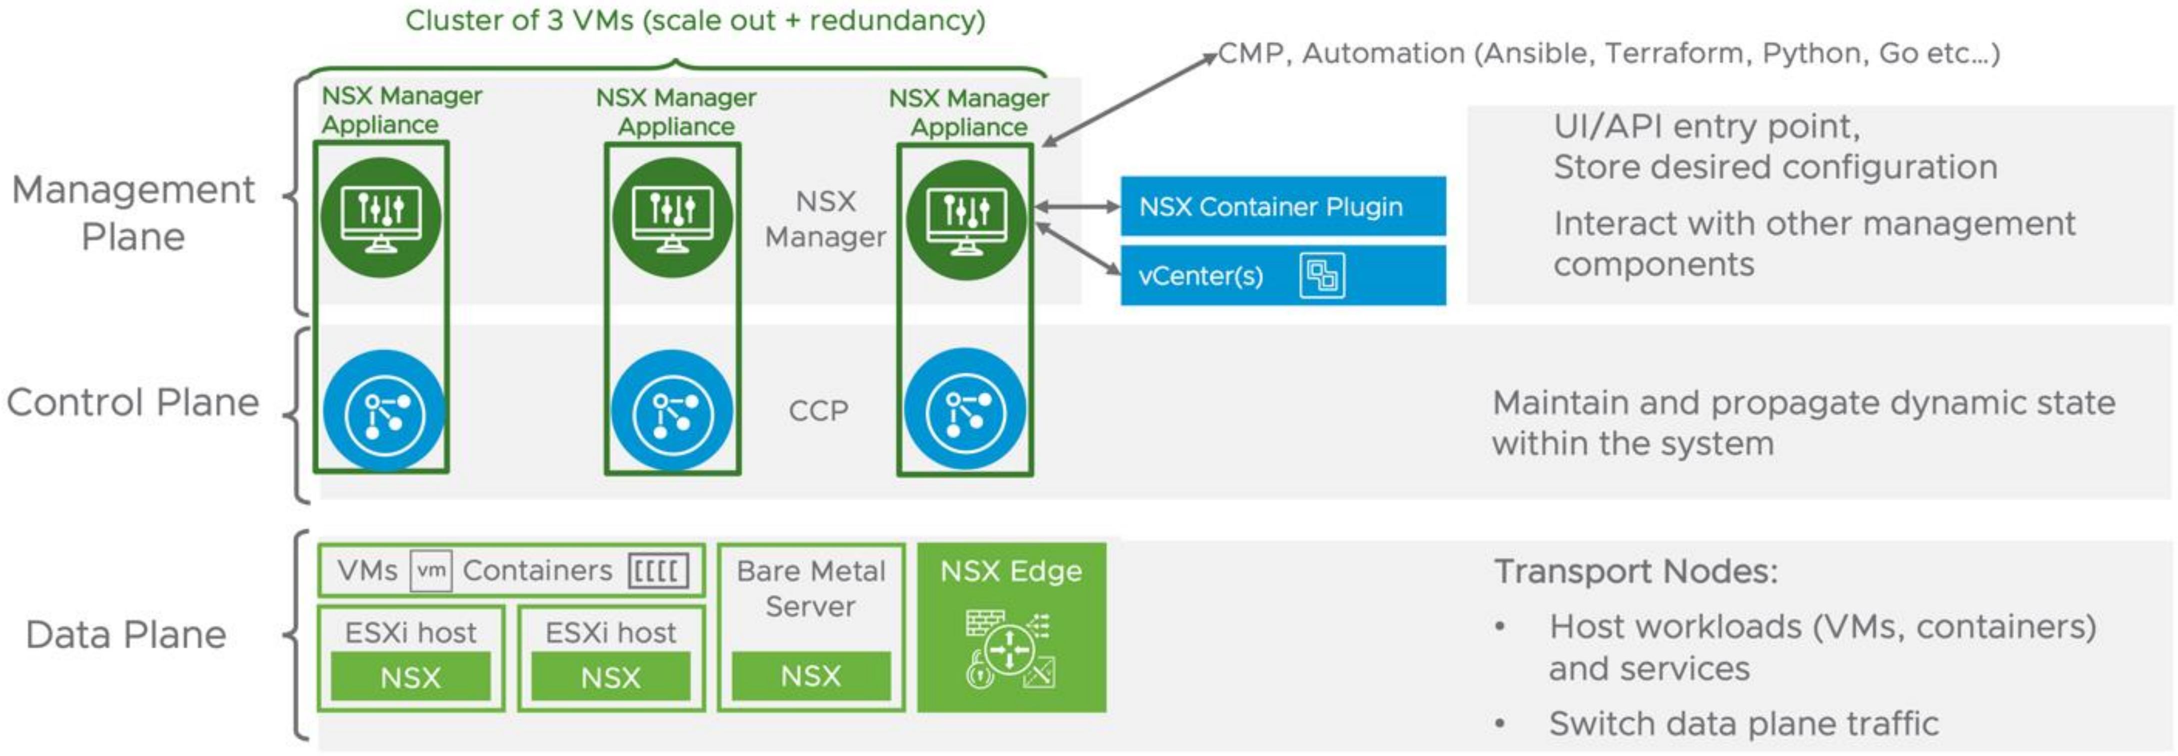
Cluster of 3 VMs (scale out + redundancy)
CMP, Automation (Ansible, Terraform, Python, Go etc…)
Management Plane
NSX Manager Appliance
NSX Manager Appliance
NSX Manager Appliance
NSX Manager
NSX Container Plugin
vCenter(s)
UI/API entry point,
Store desired configuration
Interact with other management components
Control Plane
CCP
Maintain and propagate dynamic state within the system
Data Plane
VMs
vm
Containers
ESXi host
NSX
ESXi host
NSX
Bare Metal Server
NSX
NSX Edge
Transport Nodes:
• Host workloads (VMs, containers) and services
• Switch data plane traffic

From: https://nsx.techzone.vmware.com/resource/nsx-reference-design-guide#nsx-architecture-components

VMware NSX Components

- Management Plane:
  - NSX Manager UI
  - NSX REST API
  - Manage Policies
- Control Plane:
  - CCP (Central): Receives information from the Management Plane and sends to LCP
  - LCP (Local): Monitors Data Plane and notifies changes to CCP

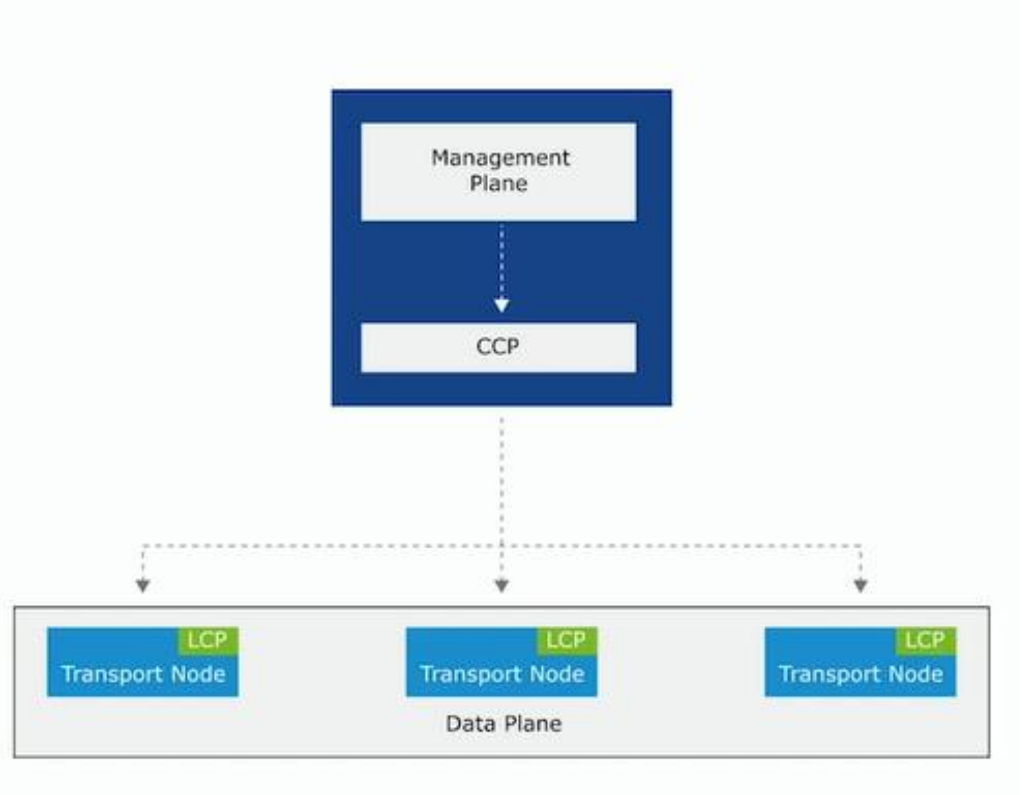

VMware NSX Components

- Data Plane:
  - Forwards the packets based on the configuration pushed by the control plane
  - Transport Nodes:
    - ESXi hosts
    - Baremetal hosts
    - NSX Edge Nodes

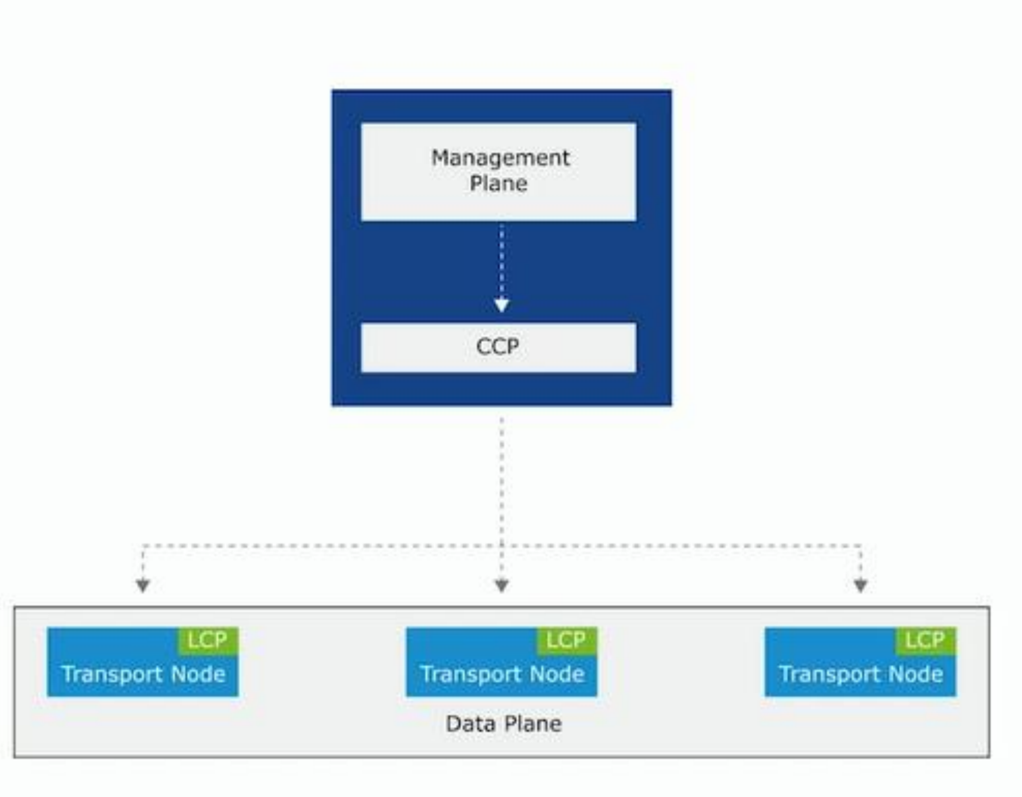

Vmware NSX Components – Multi-Tier

- Tier-0 Gateways:
  - South ↔ North routing between physical/external network and internal cloud
- Tier-1 Gateways:
  - East ↔ West routing
- Segments:
  - Virtual Layer-2 Domains (logical switches)
  - Can be VLAN or Overlay backed

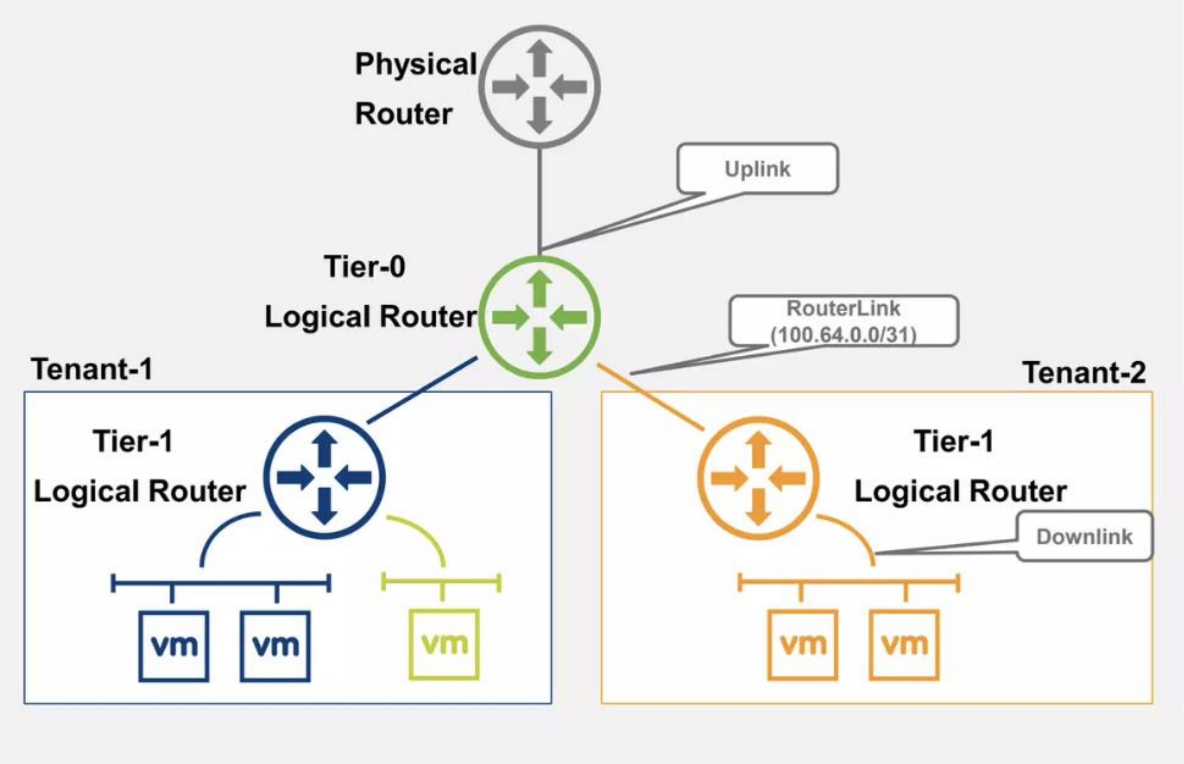

VMware NSX Components – Transport Zones

- Transport Zones:
  - Define the boundaries of NSX logical switches and segments across the network
  - Transport Nodes:
    - ESXi hosts: Support vSphere Distributed Switch (VDS)
    - Baremetal hosts: Support N-VDS host switch type.
    - NSX Edge Nodes: routing and connectivity services to external networks to the NSX deployment
  - Type: VLAN or Overlay
  - Segments are accessible across different nodes inside the same transport zone

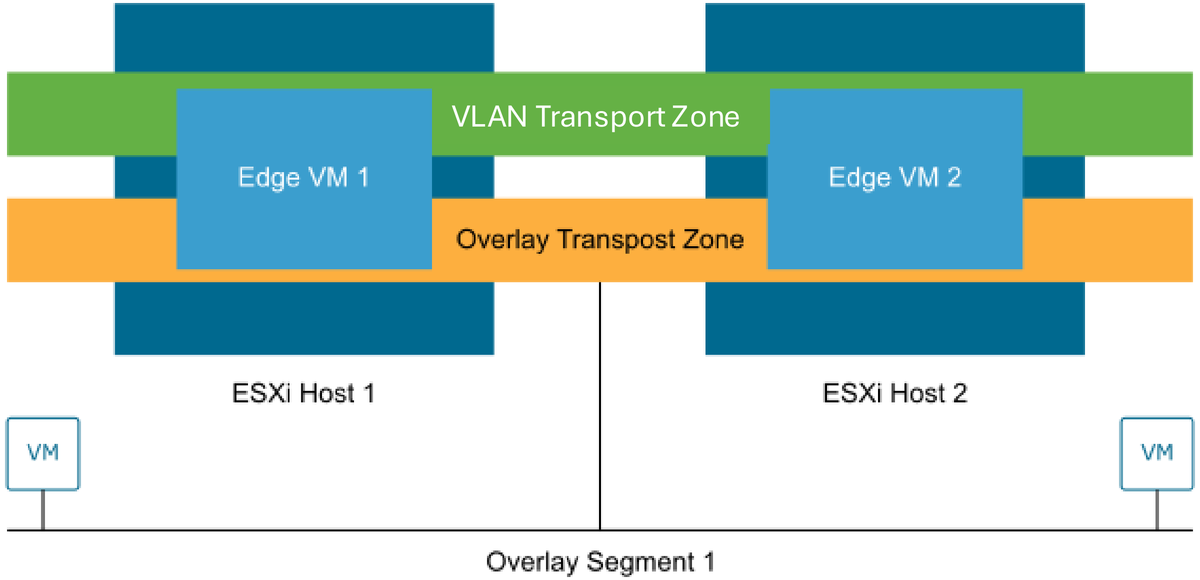

NSX Integration on CloudStack 4.20

Prerequisites - NSX

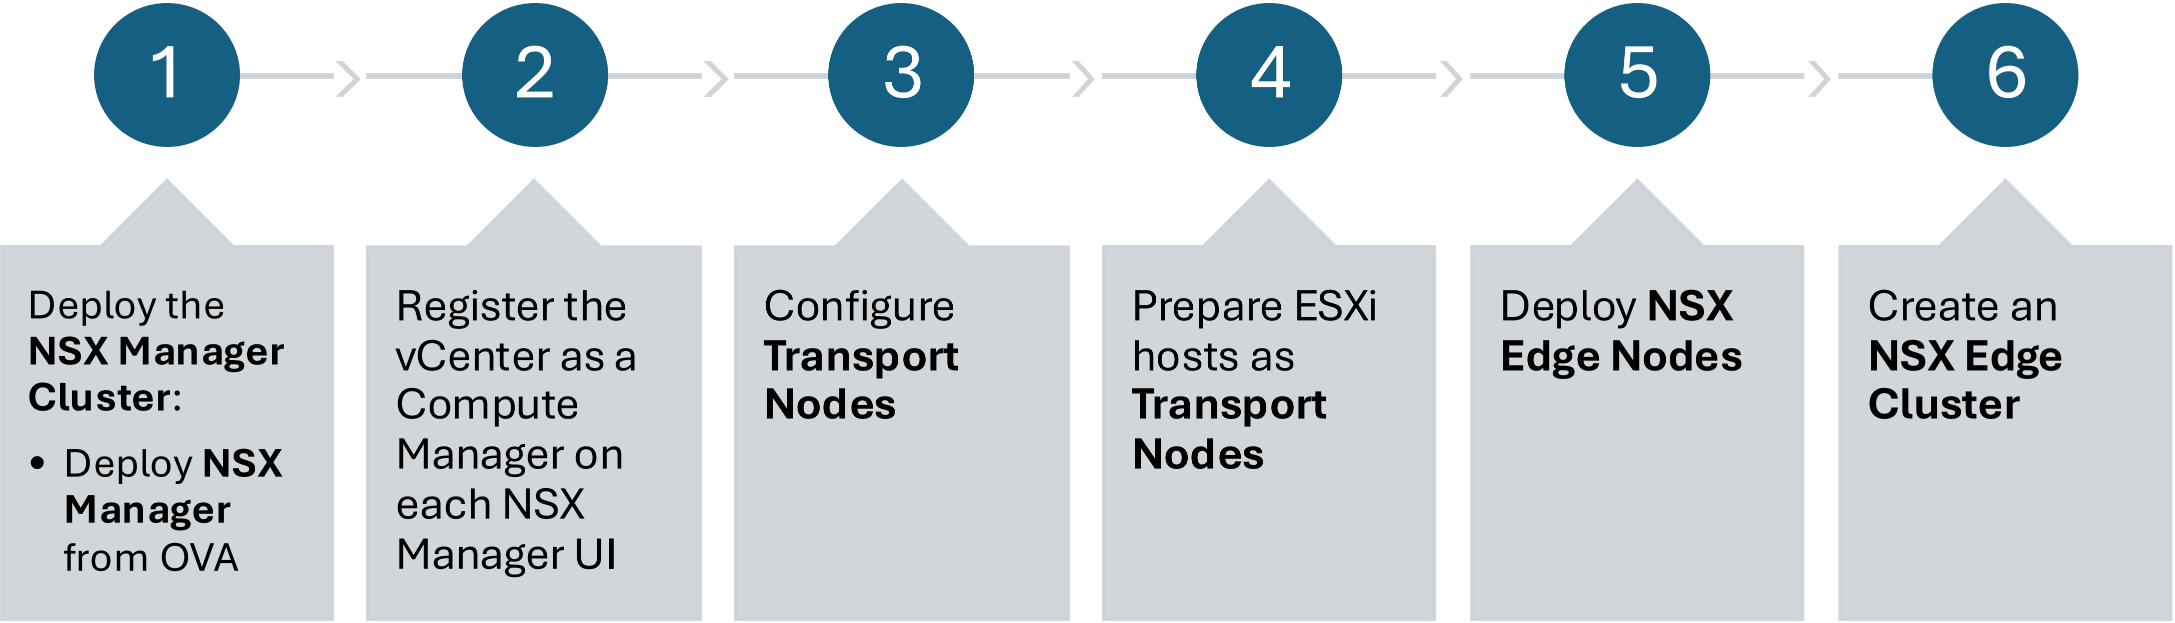

1. Deploy the **NSX Manager Cluster:**
   - Deploy **NSX Manager** from OVA
2. Register the vCenter as a Compute Manager on each NSX Manager UI
3. Configure **Transport Nodes**
4. Prepare ESXi hosts as **Transport Nodes**
5. Deploy **NSX Edge Nodes**
6. Create an **NSX Edge Cluster**

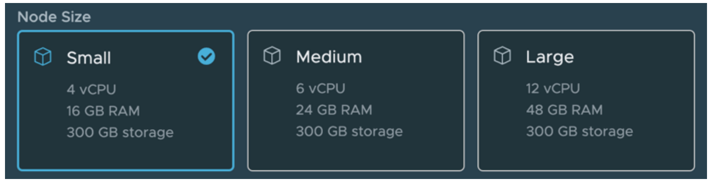

Node Size

☑ Small	Medium	Large
4 vCPU	6 vCPU	12 vCPU
16 GB RAM	24 GB RAM	48 GB RAM
300 GB storage	300 GB storage	300 GB storage

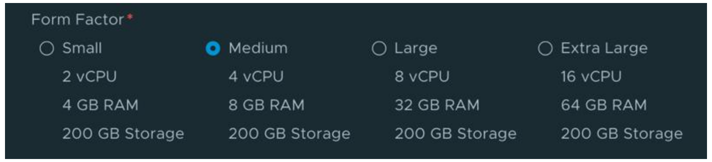

Form Factor*

○ 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

NSX Integration on CloudStack

- Introduced on CloudStack 4.20.0
- Supported Hypervisor: VMware
- Supported NSX version: 4.1.0

- Alex Mattioli – alex.mattioli@shapeblue.com
- Lucian Burlacu – lucian.burlacu@shapeblue.com
- Pearl D’Silva – pearl.dsilva@shapeblue.com
- Nicolas Vazquez – nicolas.vazquez@shapeblue.com

NSX Integration: VPC network functionalities

Routing between VPC network tiers (NSX segments)

Access Lists (ACLs) between VPC tiers and "public" network (TCP, UDP, ICMP) both as global egress rules and "public" IP specific ingress rules.

ACLs between VPC network tiers (TCP, UDP, ICMP)

Port Forwarding between "public" networks and VPC network tier

External load balancing – between VPCs network tiers and "public" networks (runs on Edge Cluster)

Internal load balancing – between VPC network tiers

Password injection, UserData and SSH Keys

External, Internal DNS

DHCP

Kubernetes host orchestration, supporting CKS on VPCs

NSX Integration

- Global settings:
  - ‘nsx.plugin.enable’: Enable the NSX plugin (false by default)
  - ‘vmware.management.portgroup’: Management Network for ESXi hosts
- Zone Creation:
  - Requires at least 2 physical networks:
    - Guest and Public traffic – Isolation type: NSX – distributed vSwitch
    - Management Traffic – Isolation type: VLAN – distributed vSwitch
  - Requires defining 2 Public IP ranges:
    - Public Traffic: used for System VMs and VRs (non NSX traffic)
    - NSX Public Traffic: for VPCs services (SNAT, DNAT, LB, etc)

CloudStack Zone Creation

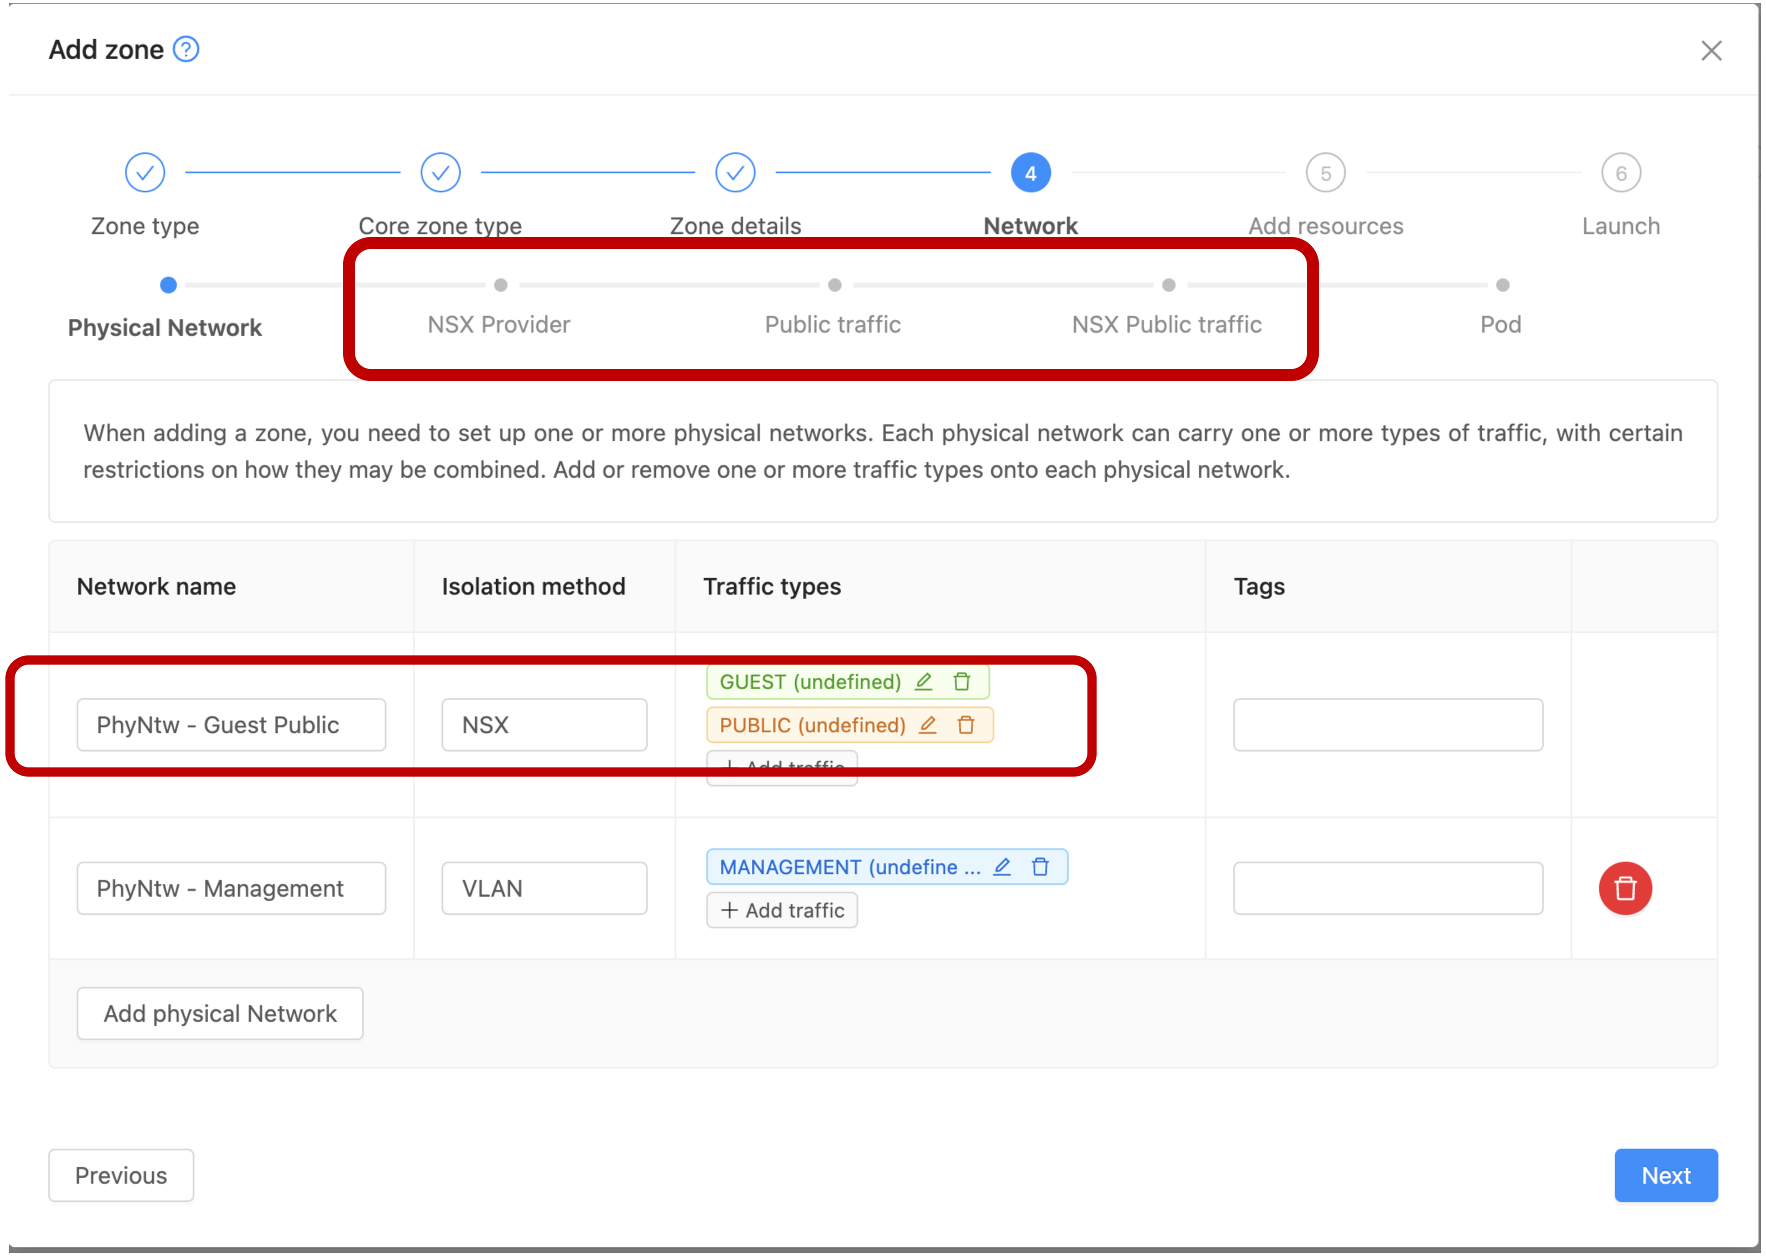

Add zone ?

✓ Zone type — ✓ Core zone type — ✓ Zone details — 4 Network — 5 Add resources — 6 Launch

● Physical Network — ● NSX Provider — ● Public traffic — ● NSX Public traffic — ● Pod

When adding a zone, you need to set up one or more physical networks. Each physical network can carry one or more types of traffic, with certain restrictions on how they may be combined. Add or remove one or more traffic types onto each physical network.

Network name	Isolation method	Traffic types	Tags	
PhyNtw - Guest Public	NSX	GUEST (undefined) PUBLIC (undefined) + Add traffic		
PhyNtw - Management	VLAN	MANAGEMENT (undefine ... + Add traffic		

Add physical Network

Previous

Next

Management Traffic – VLAN isolation

CloudStack Zone Creation:

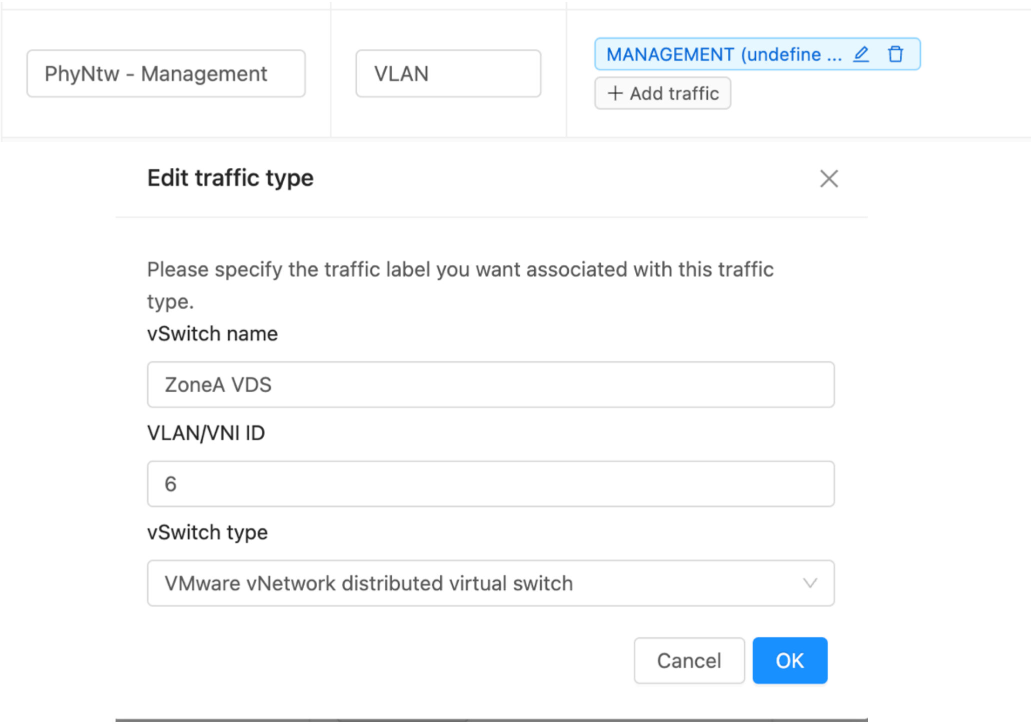

PhyNtw - Management | VLAN | MANAGEMENT (undefine ...

+ Add traffic

Edit traffic type

Please specify the traffic label you want associated with this traffic type.

vSwitch name: ZoneA VDS

VLAN/VNI ID: 6

vSwitch type: VMware vNetwork distributed virtual switch

vCenter Networking:

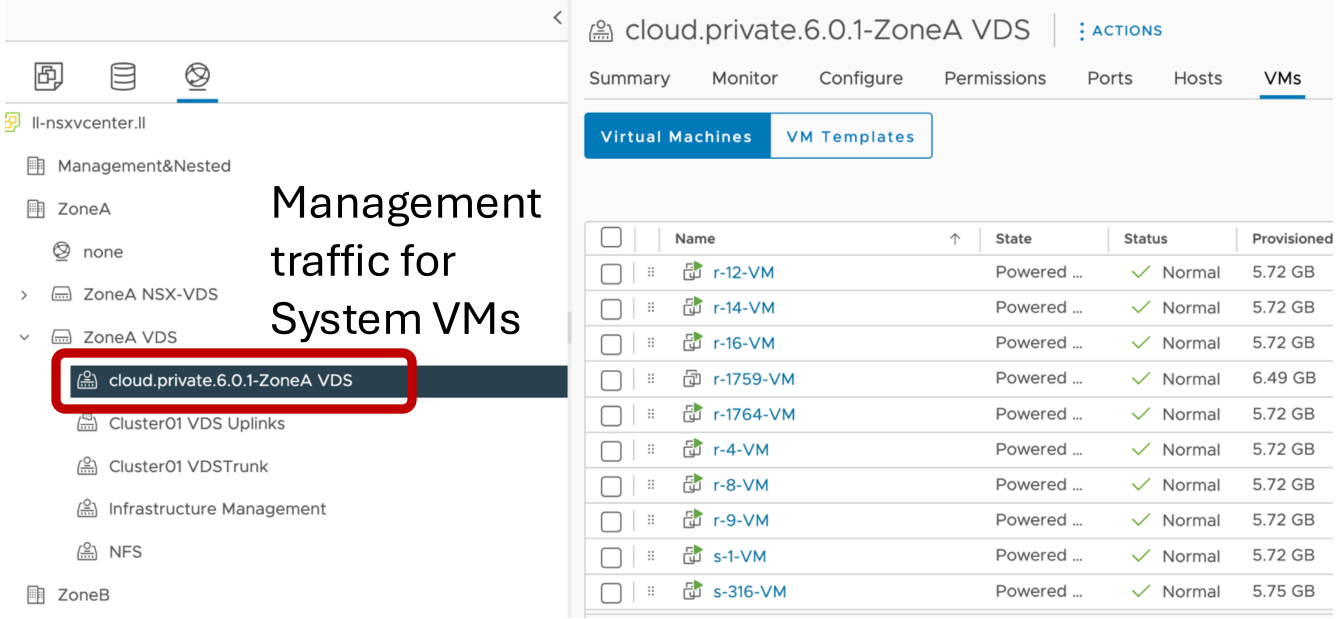

- ll-nsxvcenter.ll
  - Management&Nested
  - ZoneA
    - none
    - ZoneA NSX-VDS
    - ZoneA VDS
      - cloud.private.6.0.1-ZoneA VDS
      - Cluster01 VDS Uplinks
      - Cluster01 VDSTrunk
      - Infrastructure Management
      - NFS
  - ZoneB

Management traffic for System VMs

cloud.private.6.0.1-ZoneA VDS

Name	State	Status	Provisioned
r-12-VM	Powered ...	✓ Normal	5.72 GB
r-14-VM	Powered ...	✓ Normal	5.72 GB
r-16-VM	Powered ...	✓ Normal	5.72 GB
r-1759-VM	Powered ...	✓ Normal	6.49 GB
r-1764-VM	Powered ...	✓ Normal	5.72 GB
r-4-VM	Powered ...	✓ Normal	5.72 GB
r-8-VM	Powered ...	✓ Normal	5.72 GB
r-9-VM	Powered ...	✓ Normal	5.72 GB
s-1-VM	Powered ...	✓ Normal	5.72 GB
s-316-VM	Powered ...	✓ Normal	5.75 GB

Guest & Public Traffic – NSX isolation

CloudStack Zone Creation:

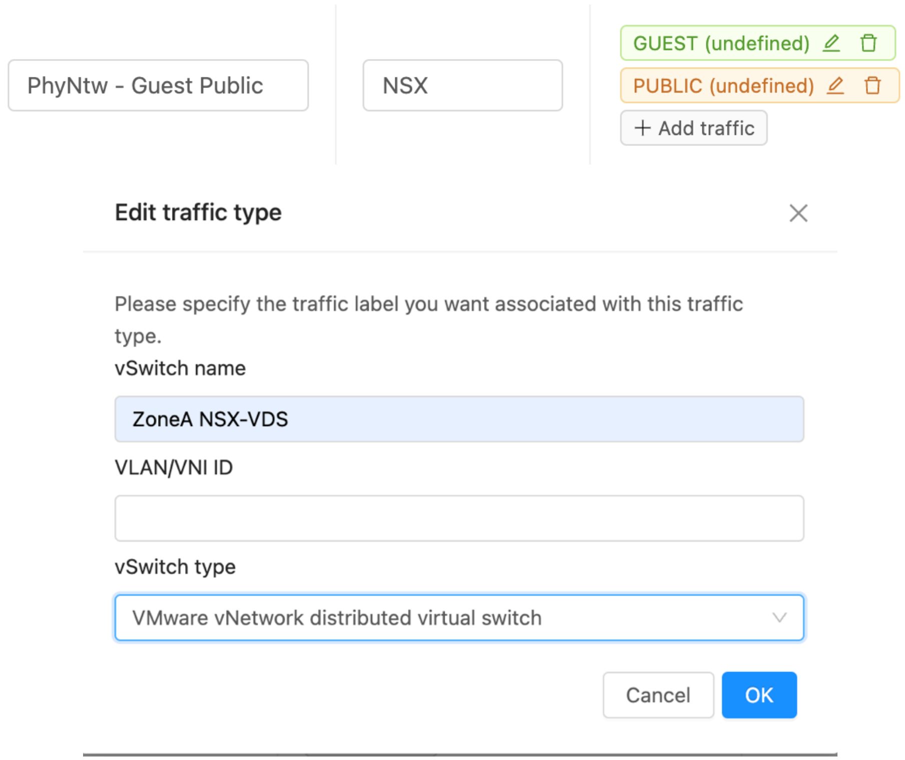

PhyNtw - Guest Public

NSX

GUEST (undefined)

PUBLIC (undefined)

+ Add traffic

Edit traffic type

Please specify the traffic label you want associated with this traffic type.

vSwitch name: ZoneA NSX-VDS

VLAN/VNI ID:

vSwitch type: VMware vNetwork distributed virtual switch

Cancel | OK

vCenter Networking:

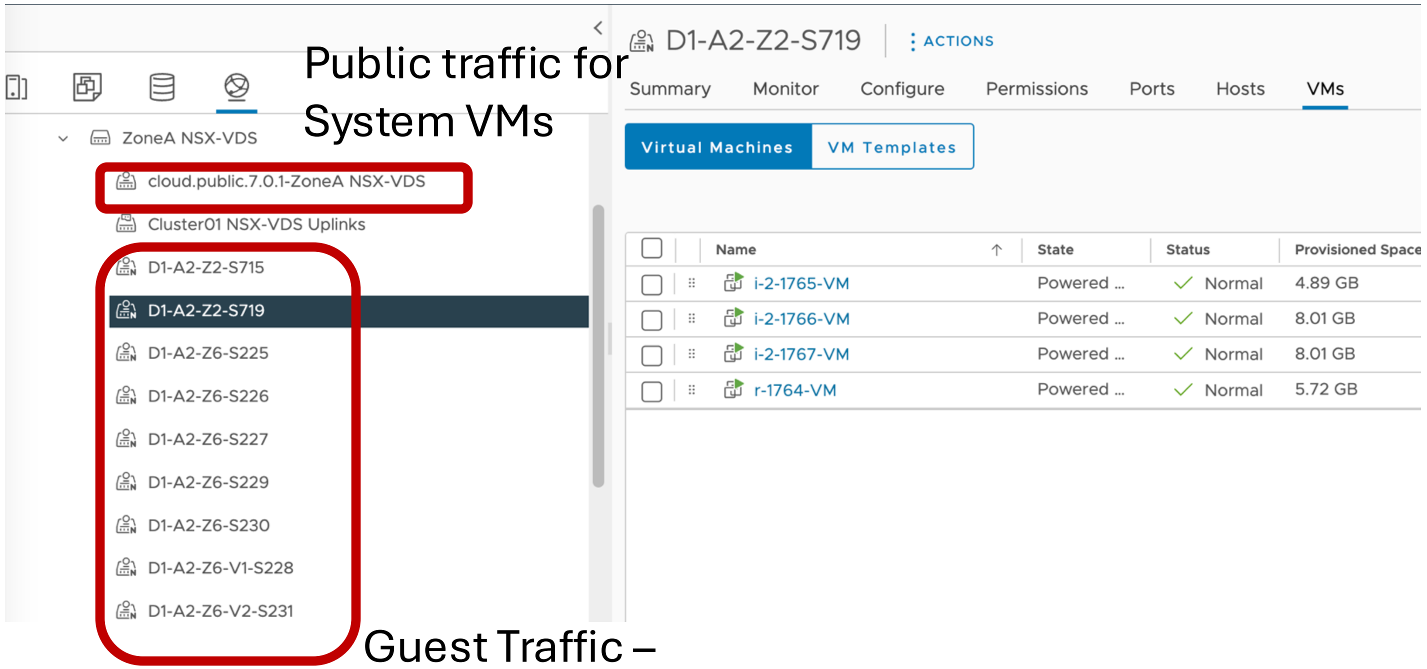

ZoneA NSX-VDS

- cloud.public.7.0.1-ZoneA NSX-VDS
- Cluster01 NSX-VDS Uplinks
- D1-A2-Z2-S715
- D1-A2-Z2-S719
- D1-A2-Z6-S225
- D1-A2-Z6-S226
- D1-A2-Z6-S227
- D1-A2-Z6-S229
- D1-A2-Z6-S230
- D1-A2-Z6-V1-S228
- D1-A2-Z6-V2-S231

D1-A2-Z2-S719 : ACTIONS

Name ↑	State	Status	Provisioned Space
i-2-1765-VM	Powered ...	✓ Normal	4.89 GB
i-2-1766-VM	Powered ...	✓ Normal	8.01 GB
i-2-1767-VM	Powered ...	✓ Normal	8.01 GB
r-1764-VM	Powered ...	✓ Normal	5.72 GB

Public traffic for System VMs

Guest Traffic – NSX segments

NSX Manager/Provider Information

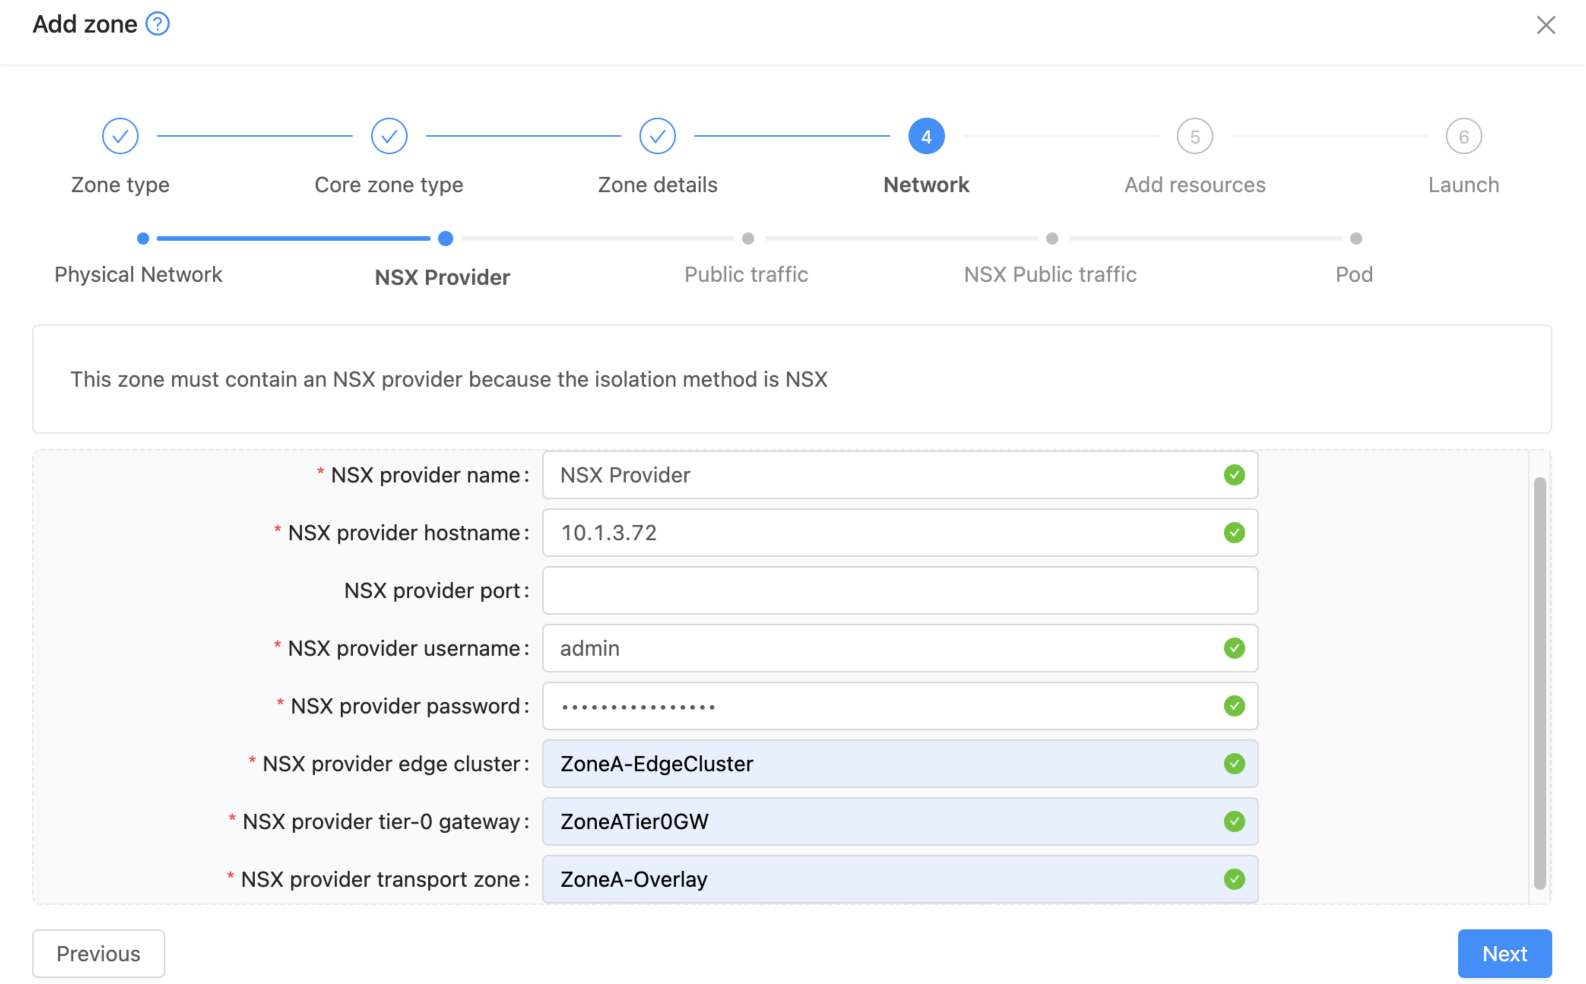

This zone must contain an NSX provider because the isolation method is NSX

* NSX provider name :	NSX Provider ✓
* NSX provider hostname :	10.1.3.72 ✓
NSX provider port :	
* NSX provider username :	admin ✓
* NSX provider password :	•••••••••••••••• ✓
* NSX provider edge cluster :	ZoneA-EdgeCluster ✓
* NSX provider tier-0 gateway :	ZoneATier0GW ✓
* NSX provider transport zone :	ZoneA-Overlay ✓

Public Traffic – System VMs and NSX Ranges

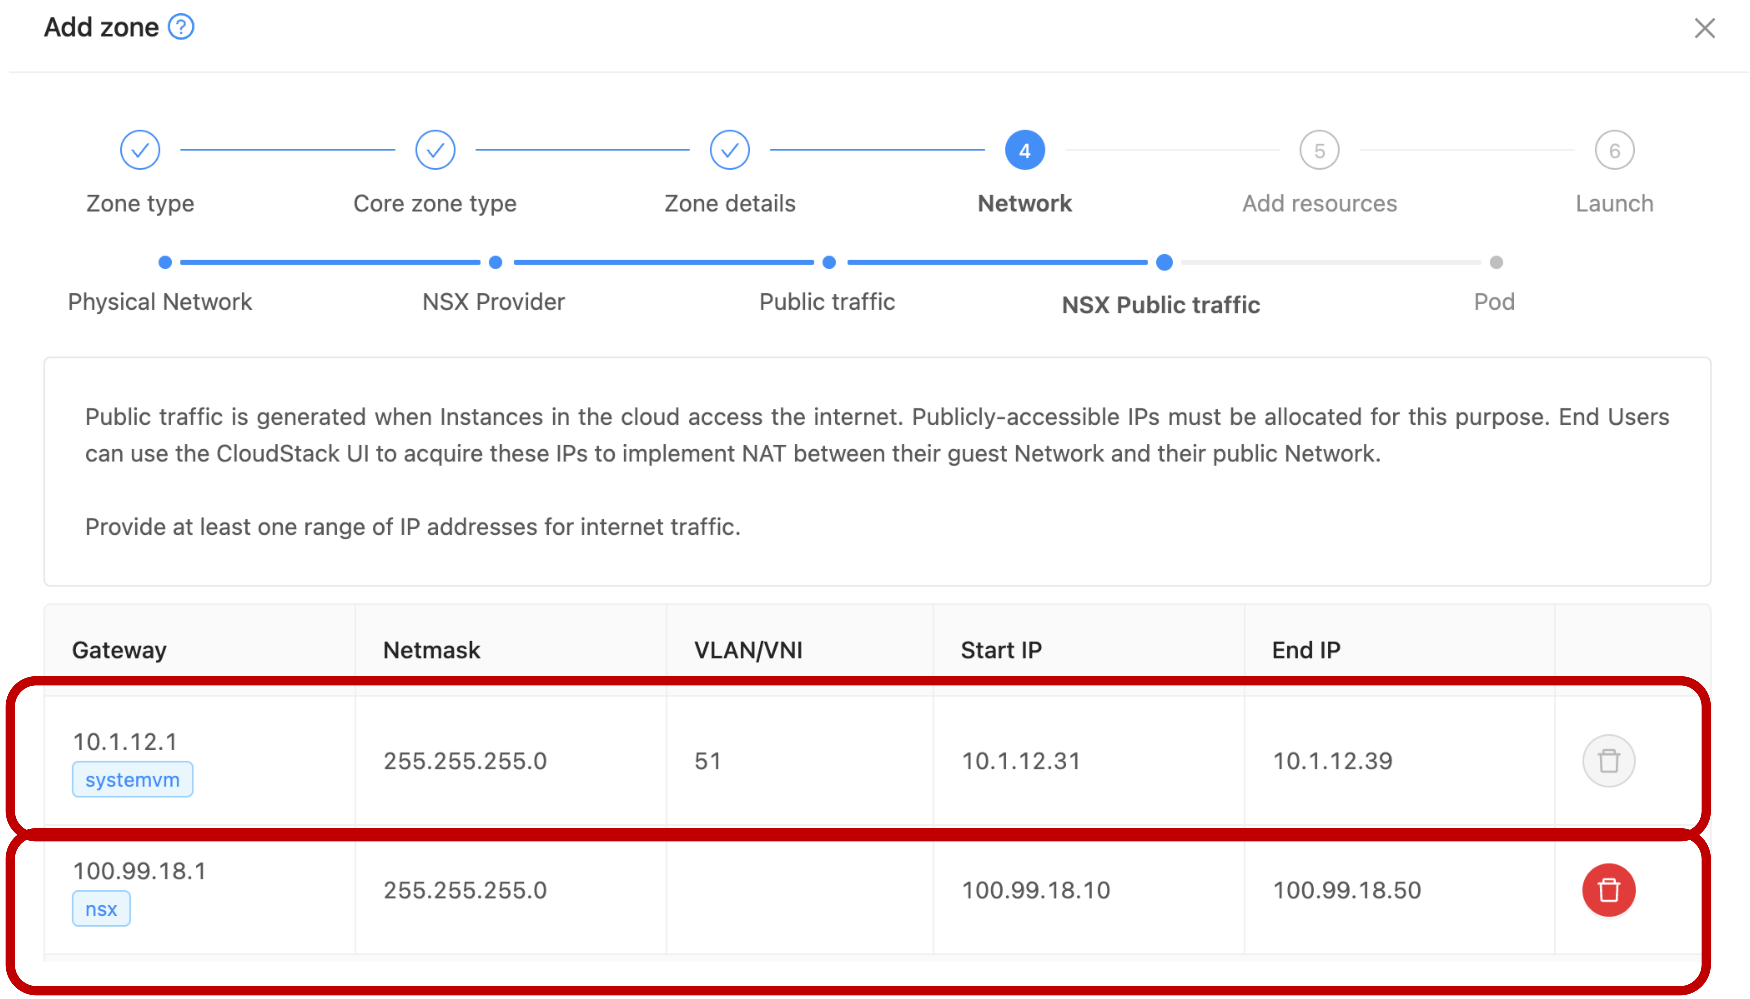

Add zone

Public traffic is generated when Instances in the cloud access the internet. Publicly-accessible IPs must be allocated for this purpose. End Users can use the CloudStack UI to acquire these IPs to implement NAT between their guest Network and their public Network.

Provide at least one range of IP addresses for internet traffic.

Gateway	Netmask	VLAN/VNI	Start IP	End IP	
10.1.12.1 systemvm	255.255.255.0	51	10.1.12.31	10.1.12.39	
100.99.18.1 nsx	255.255.255.0		100.99.18.10	100.99.18.50	

Zone Creation Summary

- At least 2 physical networks:
  - Guest and Public traffic – Isolation type: NSX – distributed vSwitch
  - Management Traffic – Isolation type: VLAN – distributed vSwitch
- At least 2 Public IP ranges:
  - System VMs and VRs Public Traffic
  - NSX Public Traffic: for VPCs services (SNAT, DNAT, LB, etc)
- NSX Manager information:
  - Hostname
  - Credentials
  - Edge Zone
  - Tier-0 GW
  - Transport Zone

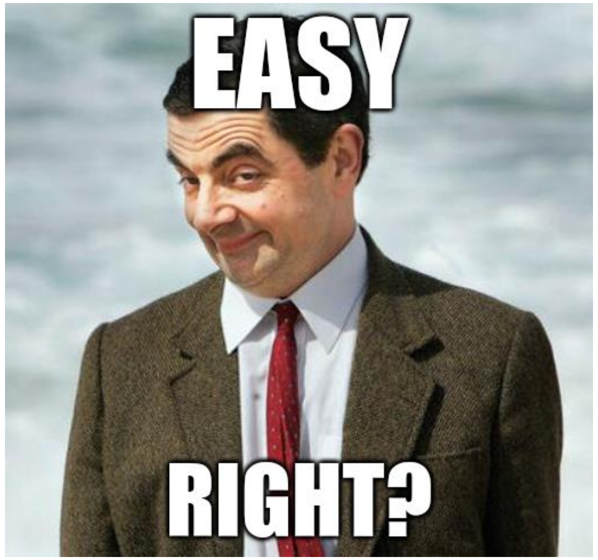

NSX-backed CloudStack VPCs

VPCs

- Tier-1 GW is the VPC Router
- Each VPC Network Tier is an NSX Segment
- Virtual Router is a helper VM
  - Provides UserData, Password Injection, SSH Keys Injection
  - VR **is not** a gateway for any VPC network tier
  - VR is assigned a random free guest IP on each VPC network tier

VPC:

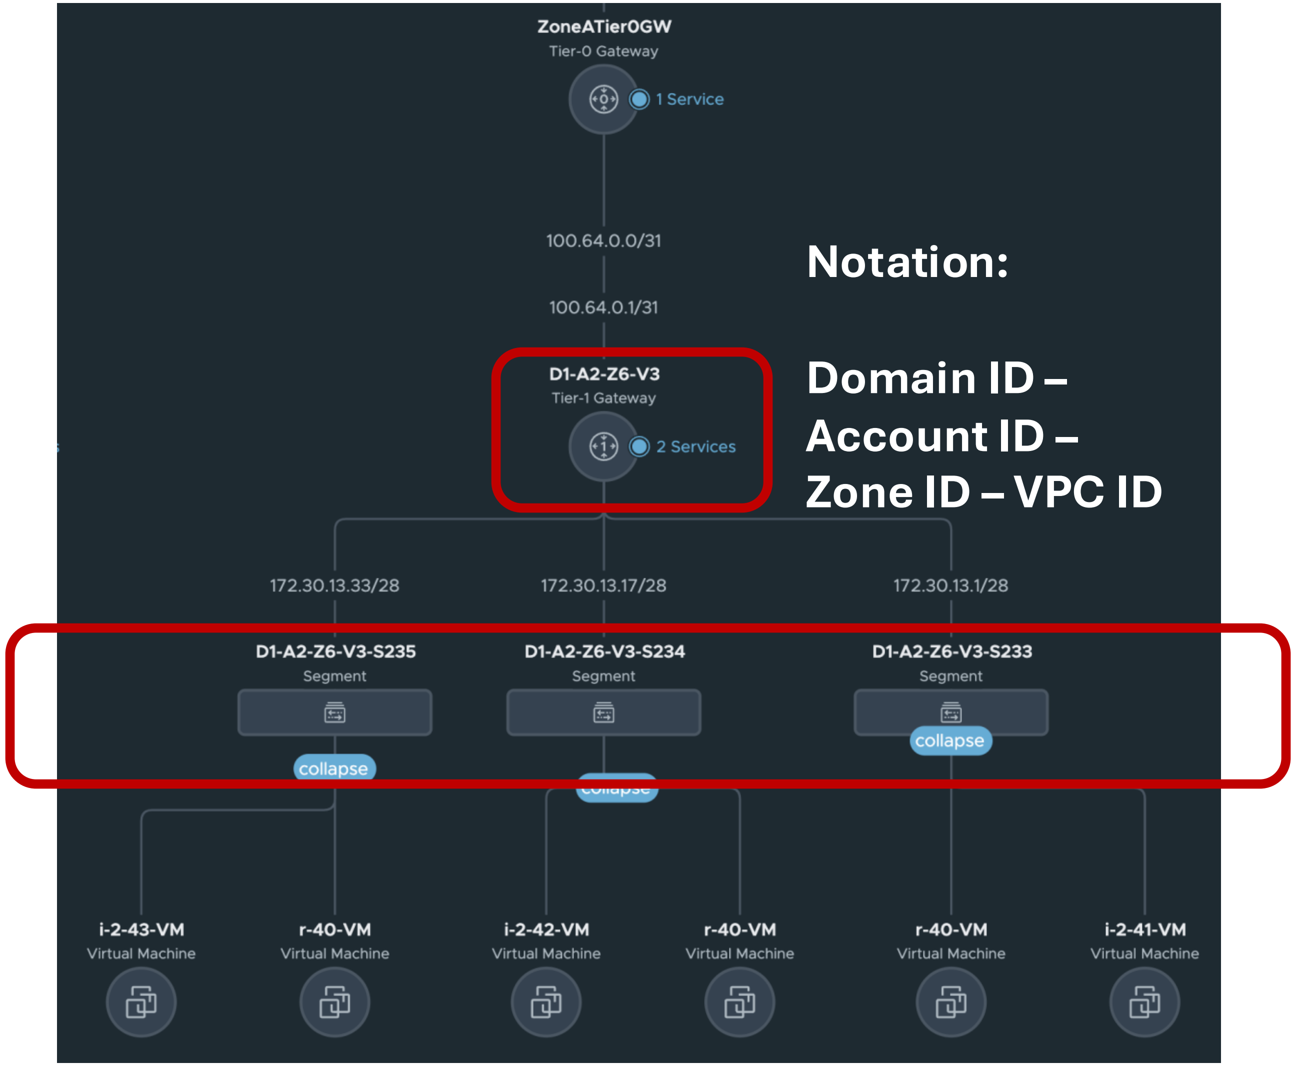

Notation:

Domain ID – Account ID – Zone ID – VPC ID

Notation:

Domain ID – Account ID – Zone ID – VPC ID – Network ID

Demo

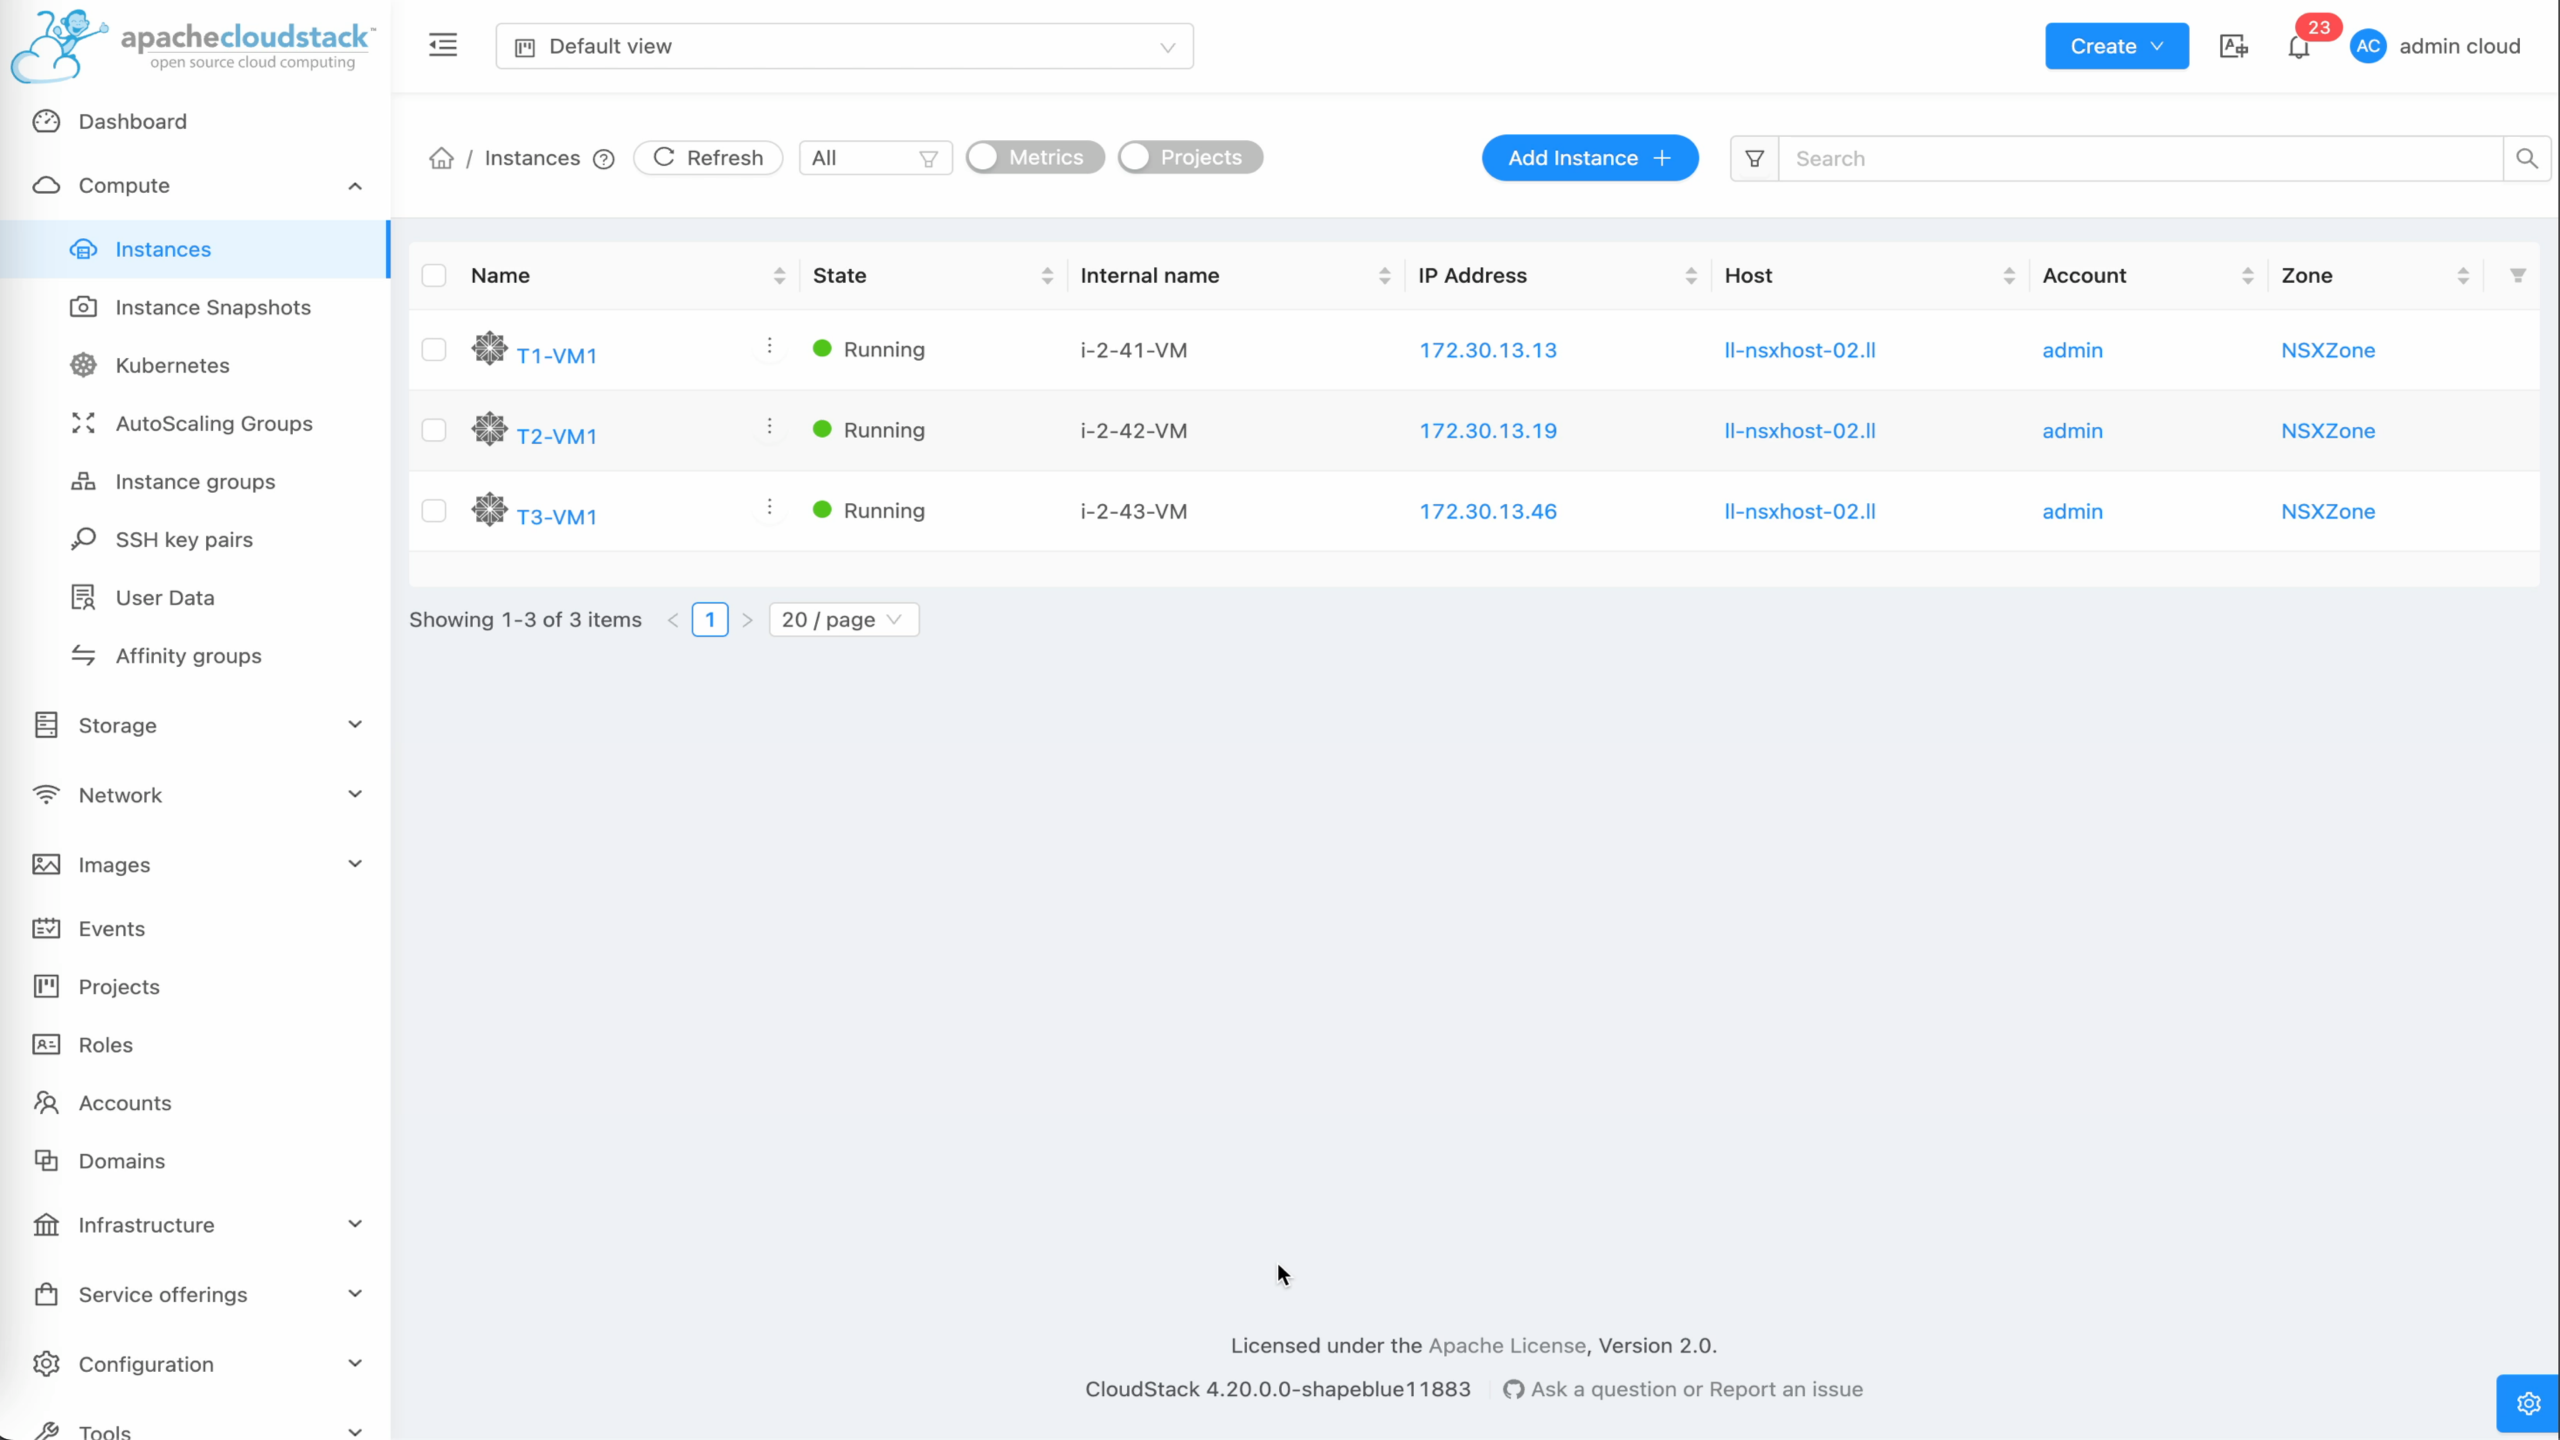

Name	State	Internal name	IP Address	Host	Account	Zone
T1-VM1	● Running	i-2-41-VM	172.30.13.13	ll-nsxhost-02.ll	admin	NSXZone
T2-VM1	● Running	i-2-42-VM	172.30.13.19	ll-nsxhost-02.ll	admin	NSXZone
T3-VM1	● Running	i-2-43-VM	172.30.13.46	ll-nsxhost-02.ll	admin	NSXZone

Showing 1-3 of 3 items

Licensed under the Apache License, Version 2.0.

CloudStack 4.20.0.0-shapeblue11883 Ask a question or Report an issue

Conclusions

- NSX 4.1.0 supported in CloudStack from version 4.20.0
- Documentation: https://docs.cloudstack.apache.org/en/latest/plugins/nsx-plugin.html
- Isolated Networks follow the same logic as VPCs with one tier
- Kubernetes Clusters are Supported!

Thank you!

#CSCollab24
@CloudStack