ARISTA DESIGN GUIDE

NSX™ for vSphere with Arista CloudVision®

Version 1.0

August 2015

Table of Contents

1 Executive Summary

Enterprise data centers are already realizing the tremendous benefits of server and storage virtualization solutions to consolidate infrastructure, reduce operational complexity, and dynamically scale application infrastructure. However, the data center network has not kept pace and remains rigid, complex, and closed to innovation--a barrier to realizing the full potential of virtualization and the software defined data center (SDDC).

VMware NSX network virtualization delivers for networking what VMware has already delivered for compute and storage. It enables virtual networks to be created, saved, deleted, and restored on demand without requiring any reconfiguration of the physical network. The result fundamentally transforms the data center network operational model, reduces network provisioning time from days or weeks to minutes and dramatically simplifies network operations.

Many data centers have workloads that have not been virtualized, or cannot be virtualized. In order to integrate them into the SDDC architecture, NSX is partnering with Arista and other vendors to provide the capability of extending virtual networking to the physical by leveraging software or hardware Layer 2 or Layer 3 gateways. As a platform, NSX provides partners the capability of integrating their solution and build on the top of the existing functionalities. NSX enables an agile overlay infrastructure for Public and Private cloud environments. By leveraging Arista's robust and resilient underlay infrastructure and CloudVision platform, customers will be able to drastically speed business services, mitigate operational complexity, and reduce costs.

This document is second in the series (Refer to NSX Design Guide for the first document) and will focus on Arista Layer 2 gateway integration with NSX. In addition, this document provides guidance for deploying VMware NSX with Arista infrastructure. It discusses the fundamental building blocks of NSX with VMware ESXi and Arista CloudVision along with step-by-step configurations of both platforms.

2 Extending virtual networks to physical space with Layer-2 gateways

NSX operates efficiently using a "network hypervisor" layer, distributed across all the hosts. However, in some cases, certain hosts in the network are not virtualized and cannot implement natively the NSX components. NSX provides thus the capability to bridge or route toward external, non-virtualized networks. This document is more specifically focusing on the bridging solution, where a Layer 2 gateway extends a logical Layer 2 network to a physical Layer 2 network.

The main functionality that a Layer 2 gateway achieves is:

- Map an NSX logical switch to a VLAN. The configuration and management of the Layer 2 gateway is embedded in NSX.
- Traffic received on the NSX logical switch via a tunnel is decapsulated and forwarded to the appropriate port/VLAN on the physical network. Similarly, VLAN traffic in the other direction is encapsulated and forwarded appropriately on the NSX logical switch.

2.1.1 Software Gateway

NSX includes natively a software version of the Layer 2 gateway functionality, with a data plane entirely implemented in the kernel of the Hypervisor, for maximum performance.

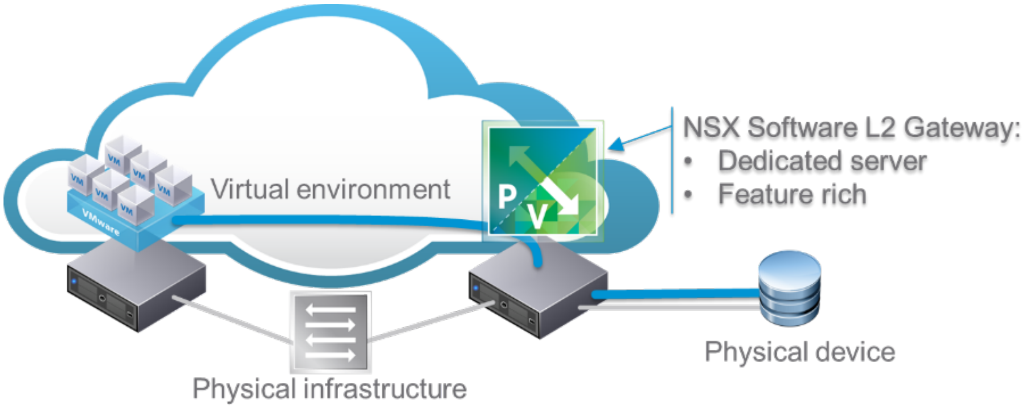

Figure 1: Software Layer 2 gateway

2.1.2 Hardware Gateway

NSX as a platform allows the integration of third party components to achieve layer-2 gateway functionality in hardware.

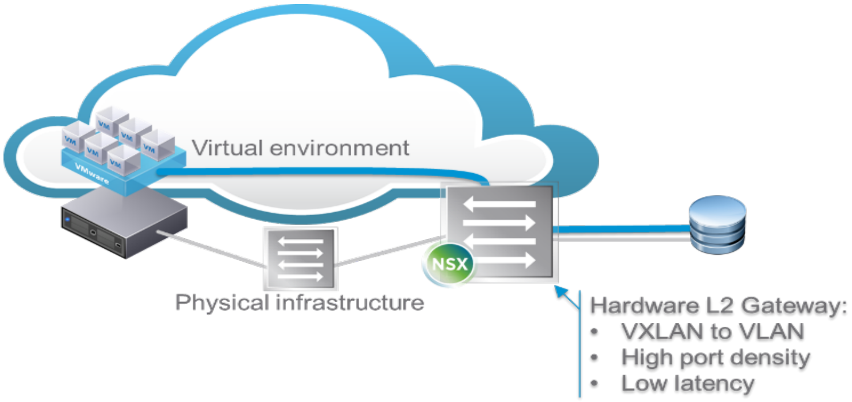

Figure 2: Hardware Layer 2 gateway

This form factor is beneficial to deployments where bandwidth and port density cannot be conveniently achieved by the software gateway.

The hardware L2 gateway is a physical device provided by a VMware partner and that can be controlled directly by NSX. This gateway will encapsulate/decapsulate the VXLAN traffic from NSX in order to allow virtual workload to interact with non-virtualized devices. The gateway does bridging between a logical switch on the NSX side and a VLAN on the physical side.

Several components are involved in the connection of the hardware gateway to NSX. They are represented in Figure 3.

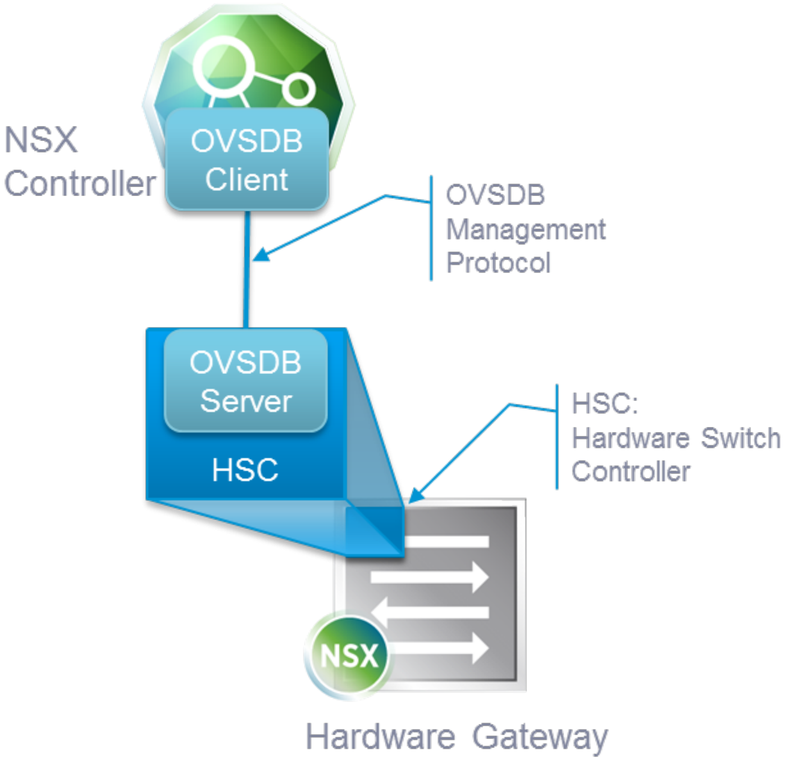

Figure3: Connection Hardware Gateway/NSX Controller

The NSX controller is responsible for handling the interaction with the hardware gateway. For this purpose, a connection is established between the NSX controller and a dedicated piece of software called the Hardware Switch Controller (HSC). The HSC can be embedded in the hardware gateway or can run as a separate standalone appliance. The HSC can control one or several hardware gateways. Arista leverages CloudVision platform as HSC, which acts as a single point of Integration to NSX for all Arista hardware gateways.

The HSC runs an OVSDB server, and the NSX controller connects as an OVSDB client. OVSDB is the Open vSwitch Database Management Protocol detailed in RFC 7047. It is an open source project that provides the capability of managing a database remotely.

This communication channel between the NSX Controller and the Hardware gateway will be used mainly to propagate two kinds of information:

2.1.2.1 Configuration information

The NSX controller will push the administrator-configured association between Logical Switch and Physical Switch/Port/VLAN to the Hardware Gateway via the HSC. The NSX Controller will also advertise a list of Replication Service Nodes (RSNs) that the Hardware Gateway will leverage to forward Broadcast, Unknown unicast or Multicast (BUM) traffic.

2.1.2.2 Run-time information

The NSX Controller will advertise to the HSC the list of Hypervisor VTEPs relevant to the Logical Switches configured on the Hardware Gateway. The NSX Controller will also advertise to the HSC the association between the MAC address of the VMs in the virtual network and the VTEP through which they can be reached.

Note that there can be several NSX controllers in an NSX deployment, providing redundancy and scale-out. The tasks mentioned as being performed by the NSX Controller are in fact shared across all the NSX Controllers in the network. The HSC will connect to all controllers.

2.1.3 NSX Integration with Arista CloudVision and Hardware gateway

The Arista CloudVision platform provides network-wide visibility and a single point of integration to NSX.

CloudVision’s foundation is an infrastructure service, sharing and aggregating working state of physical switches running Arista EOS software to provide network visibility and central coordination. State from each participating physical EOS node is registered to CloudVision using the same publish/subscribe architecture of EOS’s system database (SysDB). By communicating to each participating switch instance using a high performance binary API, CloudVision will actively synchronize state relevant to network-wide operational tasks. As an example, CloudVision’s VXLAN Control Service (VCS) aggregates network-wide VXLAN state for integration and orchestration with VMware NSX. CloudVision can be run as a standalone VM or a cluster of VMs for high availability. CloudVision uses same EOS software image as any other Arista switches.

Using CloudVision as the integration point allows for changes in the network topology to be abstracted away from NSX. In addition, CloudVision provides software and hardware version independency from joint certification. Since CloudVision runs the same EOS as any other Arista switches, customers need to only certify the CloudVision version with NSX. CloudVision, in turn, provides the aggregate VXLAN state of the physical network to NSX for the most effective physical to virtual synchronization in today's data center. This abstraction improves controller scaling by using only one touch point to control all Arista switches, abstracting physical network configuration details and letting network software handle network specific events while letting NSX focus on the higher level orchestration. CloudVision also provides redundant hardware L2 Gateways for NSX with the MLAG with VXLAN functionality. MLAG with VXLAN on Arista switches provides non-blocking active-active forwarding and redundancy with hitless failover in an event of switch failure.

In operation, Arista CloudVision will register with the NSX controller and will use the OVSDB protocol to synchronize topology information, MAC to VXLAN Endpoints, and VXLAN ID bindings with NSX. CloudVision will appropriately program the Arista switch or MLAG (Multi-Chassis Link Aggregation) switch pairs as the NSX hardware gateway. This hardware gateway integration allows for nearly instantaneous synchronization of state between physical and virtual VXLAN Tunnel Endpoints during any network change or workload modification event.

VMware's NSX Manager front-ends the entire NSX network virtualization platform . Users can also manage and operate workloads spanning virtual and non-virtualized systems from NSX as a single pane of glass.

Arista's CloudVision platform provides a set of services that simplifies monitoring, management and NSX integration with Arista switches in the virtualized data center. Users can provision Arista VTEPs as gateway nodes via the NSX Manager UI. This speeds up service delivery and helps businesses better address their needs in a programmatic and automated way across data centers.

Deployment of CloudVision Platform requires two steps:

1. Enable VXLAN Control Service (VCS) on Arista ToR switches and on CloudVision
2. Enable Hardware Switch Controller (HSC) Service on CloudVision

2.1.3.1 VXLAN Control Service

Arista CloudVision and Top of Rack switches internally runs VXLAN Control Service (VCS) through which each hardware VTEP share states between each other in order to establish VXLAN tunnels without the need for a multicast control plane.

The diagram below is a graphical representation of a of typical Spine/Leaf network. The Leaf / Top-Of-Rack (ToR) switches are acting as the hardware VTEP gateways. In this deployment the TORs act as the VTEPs with a CloudVision instance running within a VM. All Arista switches use VCS to share their MAC to VTEP binding with CloudVision.

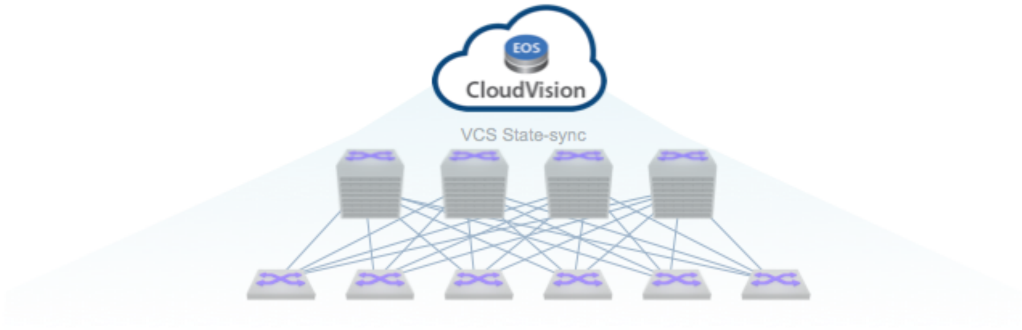

Figure 4: VXLAN State Synchronization with VCS

2.1.3.2 Hardware Switch Controller Service

The Hardware Switch controller Service (HSC) provides an integration point between the OVSDB controllers and the VCS -- it provides a means for software and hardware switches to exchange state.

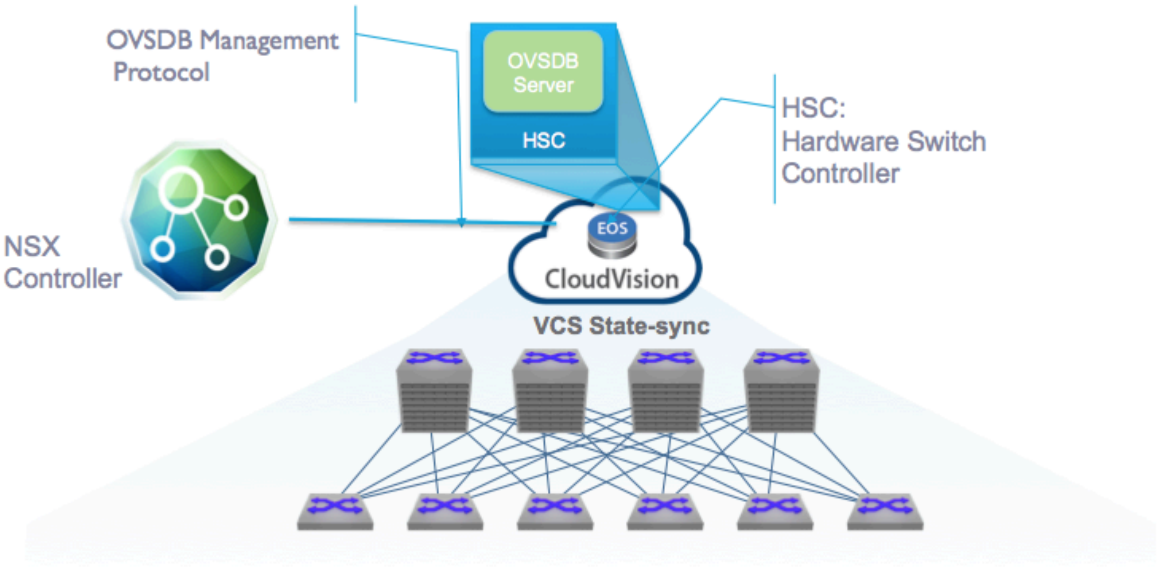

Figure 5: CloudVision as HSC

In this deployment, the CloudVision is modeled as a *single* L2 Gateway in the transport network with *multiple* VXLAN Transport Connectors providing the overlay. Each Arista VTEP managed by CloudVision is mapped to a VXLAN Transport Connector in the single L2 Gateway in the Transport Node.

Details about configuring aforementioned steps are covered in the next section of this paper.

3 Configuration

This section provides an overview of the configuration required on the Arista switches and NSX platform including integration of hardware gateway.

This document assumes readers have a functional knowledge of NSX and deploying Arista routing and switching infrastructure. Readers are strongly advised to read the VMware® NSX for vSphere Network Virtualization Design Guide for additional context; it provides a detailed characterization of NSX operations, components, design, and best practices for deploying NSX.

Specifically, the goal of this document is to provide guidance for running NSX with Arista switches deployed as leaf-spine architecture with layer-3 topology. In addition, this document also covers deployment of Arista switches as NSX hardware gateways for virtual to physical workload communication.

The section covers three critical aspects of the design:

- Connectivity requirements for NSX including VMkernel networking, VLAN allocation and configuration, VXLAN Tunnel End-Point (VTEP) configuration, and layer-3 peering and routing configurations
- Arista switch connectivity with NSX including registering CloudVision with NSX
- NSX Components and Cluster connectivity

3.1 NSX VMkernel Networking Requirements

In a traditional ESXi environment, several infrastructure VLANs are provisioned. The three VLANs defined for the VMkernel interfaces in this example are shown in the table below:

Traffic Types	Functions	VLAN ID
Management	ESXi and NSX management plane	100
vMotion	VM mobility	101
IP Storage VLAN	Applications & infrastructure data store connectivity	102

Table 1: Infrastructure Traffic Types and VLAN

NSX uses an additional infrastructure VMkernel interface for the VXLAN encapsulated guest traffic that is tunneled over the physical network.

VXLAN VLAN: During the NSX configuration phase an additional VMkernel interface is created for VXLAN traffic. Note that the configuration of the four VMkernel interfaces used in this example is a one time-task. By leveraging VXLAN tunnels, further definition of logical networks used for carrying VM traffic is achieved entirely independently from the physical network, eliminating the need to define additional VLANs for accommodating VM growth. The VLAN Switch Virtual Interface (SVI) termination is either at the aggregation layer or at the access layer, depending on the topology deployed with Arista physical network.

Traffic Type	Function	VLAN ID
VXLAN VLAN	Overlay VXLAN VTEP Connectivity	103

Table 2: VXLAN VLAN for VM Guest Traffic

Additional VLANs are needed for:

- **L3 ECMP Connectivity**: Two VLANs are typically required for allowing north-south traffic from the NSX domain to the physical world.
- **Bridging**: When using the Layer 2 Gateway functionality, some Logical Switches are extended to dedicated VLANs in the physical infrastructure.

3.1.1 Demystifying the VXLAN and IP Connectivity

VXLAN decouples the connectivity for the logical space from the physical network infrastructure. Devices connected to logical networks can leverage the entire set of network services (load balancer, firewall, NAT, etc.) independently from how the underlying physical infrastructure is configured. This helps solve many of the challenges of traditional data center deployments, such as agile & programmatic application deployment, vMotion across layer-3 boundaries, as well as multi-tenancy support to overcome the VLAN limitation of 4094 logical segments.

3.1.1.1 NSX VXLAN Capabilities

NSX VXLAN implementation offers critical enhancements to VXLAN:

- NSX enables multicast free VXLAN with the help of the controller. Removing multicast from the underlay network greatly simplifies physical network configuration. The distributed discovery and efficient management of control plane functions (MAC, VTEP and ARP table management) are relegated to highly available clustered controllers.

- NSX enables VXLAN encapsulation at the kernel level in the ESXi host. This VXLAN encapsulated frame is nothing but a generic IP packet, which is then routed by Arista switch forwarding traffic, based on the destination VTEP IP address. In addition, the underlay does not see the explosion of MAC addresses or the intensive configuration requirement for ad-hoc connectivity in the virtual domain.

The net effect of these enhancement is shown in below figure where ESXi encapsulate/decapsulate the VXLAN header with multicast-free replication of BUM (Broadcast Unknown Multicast) traffic.

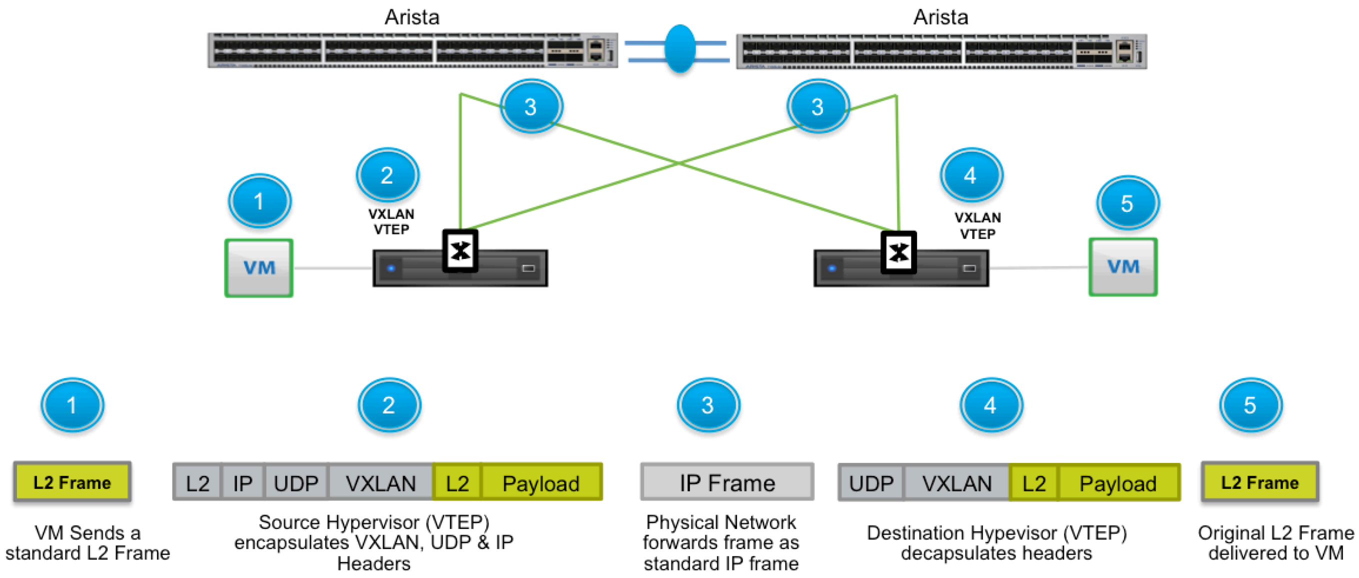
Arista
Arista
3
3
2
VXLAN VTEP
4
VXLAN VTEP
1
5
VM
VM
1
2
3
4
5
L2 Frame
L2
IP
UDP
VXLAN
L2
Payload
IP Frame
UDP
VXLAN
L2
Payload
L2 Frame
VM Sends a standard L2 Frame
Source Hypervisor (VTEP) encapsulates VXLAN, UDP & IP Headers
Physical Network forwards frame as standard IP frame
Destination Hypevisor (VTEP) decapsulates headers
Original L2 Frame delivered to VM

Figure6: VXLAN Encapsulation & Decapsulation at ESXi Kernel

These enhancements to VXLAN simplify the underlay physical network. For additional details about VXLAN, packet flow for various layer-2 control plane discovery, and connectivity, please refer to the VMware® NSX for vSphere Network Virtualization Design Guide.

3.1.2 VXLAN and VDS Connectivity with Arista Switches

VXLAN connectivity consists of two components: transport zone and VTEP. The transport zone is a collection of ESXi clusters participating in a single VXLAN domain. VTEP (VXLAN Tunnel End Point) is a logical interface (VMkernel) that connects to the transport zone for encapsulation/decapsulation of VM guest traffic as shown in Figure 1.

In addition, Arista's Layer-2 VTEP services enable VXLAN-to-VLAN (or physical port) capabilities at Top of Rack (ToR) layer. The Arista's Layer-2 gateway integrates with NSX through data plane (VXLAN) and control plane (OVSDB) protocols while unifying the management plane. This integration delivers Layer-2 gateway services that discover non-virtualized workloads in the data center, enabling seamless communication with virtualized workload by linking VXLAN tunnels to VLANs in the physical network. Users can also manage and operate workloads spanning virtual and non-virtualized systems from NSX as a single pane of glass.

For a given cluster, only one VDS is responsible for VXLAN connectivity. The cluster design in section below goes in details of VDS design recommendation. However, there are three critical design requirements for VXLAN connectivity: VLAN ID for VXLAN, and VDS uplink Configuration.

- **VLAN ID for VXLAN:** At the NSX configuration phase, the VTEP(s) are defined with transport zone VLAN ID; this VLAN port-group is dynamically created during the cluster VXLAN preparation. For a VLAN that supports VXLAN transport zone, a specific configuration for a VLAN ID is required based on the physical topology. NSX requires the VDS dvUplink configuration to be consistent per VDS and thus for VLAN ID for the VXLAN transport zone has to be the same regardless of layer-2 or layer-3 topology. The detailed configuration VLAN ID mapping to a specific topology is described in section 3.2.3.

- **VDS Uplink Configuration**: The NSX creates a dvUplink port-group for VXLAN that must be consistent for any given VDS and NIC teaming policy for VXLAN port-group must be consistent across all hosts belonging to the VDS. Typically one VTEP is sufficient; however multiple VTEPs are also supported. The number of VTEP(s) supported is determined by a combination of the number of host interfaces exposed to VDS and uplink teaming mode configurations as shown in the table below.

Teaming and Failover Mode	NSX Support	Multiple VTEPs Created	Uplink Behavior 2 x 10G	Arista Port Configurations
Route Based on Originating Port	✓	Yes	Both Active	Standard
Route Based on Source MAC Hash	✓	Yes	Both Active	Standard
LACP	✓	No	Flow based	MLAG Port-Channel - LACP
Route Based on IP Hash (Static EtherChannel)	✓	No	Flow based	MLAG Port-Channel – LACP mode OFF
Explicit Failover Order	✓	No	Only one link is active	Standard
Route Based on Physical NIC Load (LBT)	×	N/A	N/A	

Table 3: VDS Teaming Mode and NSX Support

As one can notice from Table 3, selecting LACP or Static EtherChannel teaming mode limits the choice for selecting team-modes per traffic types (port-groups for management, vMotions, VXLAN). With LACP or Static EtherChannel, only one VTEP per host is created. Any other teaming modes typically require a VTEP for each physical uplink. The only exception is that LBT (Load Based Teaming) mode is not supported for VTEP VMkernel.

The table above also shows the port-configuration mode for Arista switches relative to the uplink teaming mode. Notice that LACP and Static EtherChannel modes require VLAN based MLAG (Mutli-chasis LAG) on Arista Switches.

3.2 Arista Connectivity with NSX

Arista switches are capable of supporting multiple types of leaf-spine topologies (Layer-2 or Layer-3). Arista, however, strongly recommends using layer-3 leaf Spine Network. The VMware® NSX for vSphere Network Virtualization Design Guide goes into detail on leaf-spine topology attributes and hence is not discussed here since most of the recommendations apply to Arista switches. In leaf-spine design, the layer-3 is terminated at the ToR and thus all the VLANs originating from ESXi hosts terminate on Arista ToR. Typically in layer-3 design this means the VLAN ID is irrelevant and can be kept unique or the same for a type of traffic per rack, as long as the VLAN ID maps to the unique subnet.

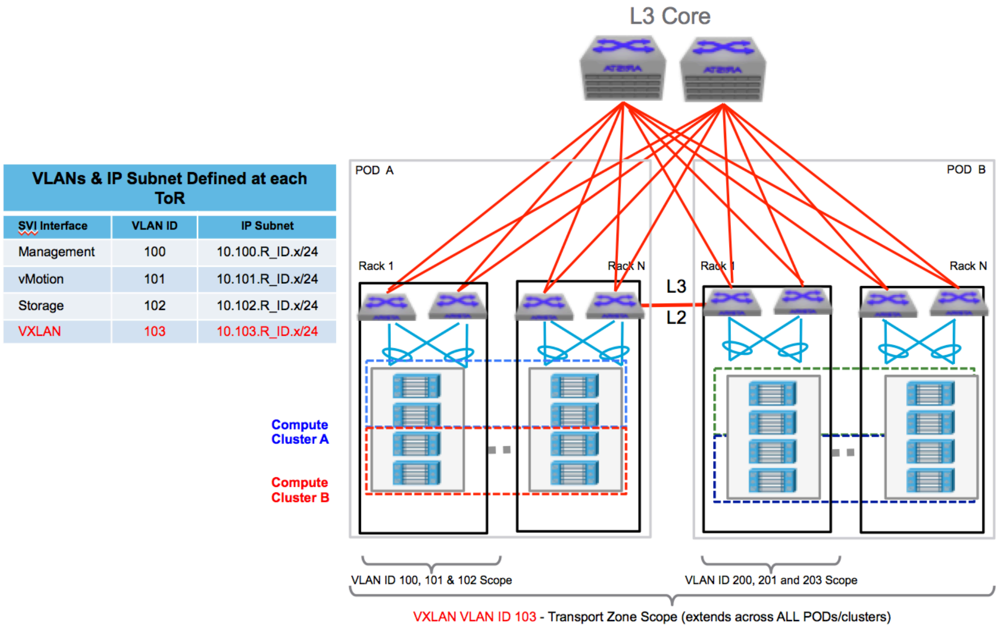

VLANs & IP Subnet Defined at each ToR

SVI Interface	VLAN ID	IP Subnet
Management	100	10.100.R_ID.x/24
vMotion	101	10.101.R_ID.x/24
Storage	102	10.102.R_ID.x/24
VXLAN	103	10.103.R_ID.x/24

L3 Core
POD A
POD B
Rack 1
Rack N
Rack 1
Rack N
L3
L2
Compute Cluster A
Compute Cluster B
VLAN ID 100, 101 & 102 Scope
VLAN ID 200, 201 and 203 Scope
VXLAN VLAN ID 103 - Transport Zone Scope (extends across ALL PODs/clusters)

Figure 7: Layer-3 POD Design with VXLAN

VLANs & IP Subnet Defined at each ToR		
SVI Interface	**VLAN ID**	**IP Subnet**
Management	100	10.100.R_ID.x/24
vMotion	101	10.101.R_ID.x/24
Storage	102	10.102.R_ID.x/24
VXLAN	103	10.103.R_ID.x/24

Table 5: SVI to VLAN Mapping for Layer-3 POD Design

The VTEPs configured by NSX are automatically created and leverage the same VLAN ID on the uplinks of all the hosts attached to the VDS.
In other words, for the VXLAN VTEP, the VLAN ID remains the same for every ToR; however, the subnet that maps to this VLAN can be different on a per-ToR basis. One can keep the VLAN ID for the rest of the traffic types to be the same for every rack. This simplifies the configuration for every rack and only requires configuration once. As an example, this is depicted in the table above, which can be repeated for each ToR configuration with unique subnet identified by rack ID.

This section covers the details of connectivity requirements for various NSX components and clusters required for integration with Arista switches. Supporting NSX for Arista switches requires few simple steps:

1. Enabling support for jumbo frames
2. Configurations supporting IP forwarding and routing support
3. VLAN configuration requirements based on physical network topology
4. Configuring Arista CloudVision for enabling Hardware VTEP

3.2.1 Arista EOS Jumbo Frame Configurations

Arista switches support jumbo frame; however it is not enabled by default. The jumbo frame configuration steps are different for layer-2 and layer-3 interfaces.

3.2.1.1 Configuration steps for layer-2 interface

Change MTU to 9214 on each layer-2 interface via the “l2 mtu 9214” interface command. Sample CLI is show for each type of interface in below table:

```
Arista(config-if-<interface>)#{no|default} l2 mtu <configured value in bytes>
```

3.2.1.2 Configuration steps for layer-3 interface

Change MTU to 9214 on each layer-3 interface via the “mtu 9214” interface command. Sample CLI is show for each type of interface in below table:

```
Arista(config-if-<interface>)#{no|default} mtu <configured value in bytes>
```

3.2.2 Configuration Support for IP Forwarding

The IP forwarding feature requires defining the SVI interface with an IP address and enabling the routing protocol of choice (OSPF - Open Shortest Path First or BGP - Border Gateway Protocol). The SVI configuration is also enabled with the First Hop Redundancy Protocol (FHRP) to provide redundancy for ToR failure.

Arista recommends using Layer-3 leaf Spine architecture with BGP as the routing protocol. This is primarily based on the rapid convergence properties of BGP and its ability to scale out without requiring a significant increase in routing protocol ‘chatter’. It is also a deterministically peered routing protocol meaning that a simple cabling error cannot result in an adverse traffic pattern or malformed topology.

The routing protocol configuration and its interaction with NSX routing domains is further described in Edge cluster connectivity in section 4.3.3.

3.2.3 Arista CloudVision Configuration

As described in section 2.1.5, deployment of CloudVision Platform requires two steps:

3. Enable VXLAN Control Service (VCS) on Arista ToR switches and on CloudVision
4. Enable Hardware Switch Controller (HSC) Service on CloudVision

3.2.3.1 VXLAN Control Service

Arista EOS, both for CloudVision and for Top of Rack switches internally runs VXLAN Control Service (VCS) through which each hardware VTEP shares VXLAN state with each other. This collective state can then be relayed to NSX Controller and vice versa.

3.2.3.1.1 Configuration on Arista Top of Rack Switches

The following configuration is necessary on each TOR switch that needs to connect to the VXLAN Control Service running on CloudVision.

```
ToR1(config)#management cvx
ToR1(config-mgmt-cvx)#server host 172.27.6.248
ToR1(config-mgmt-cvx)#no shutdown
```

The IP address above is the *management IP address of the CloudVision or the IP address that CloudVision is listening on for client connections.*

Next you should configure the VXLAN interface as described in the EOS Command Guide.

```
ToR1(config)#interface vxlan 1
ToR1(config-if-Vx1)#vxlan controller-client
```

3.2.3.1.2 Configuration on Arista CloudVision

The following instructions assume that the CloudVision is running in a VM.
For instructions on running EOS on the VM please look at this article on Arista EOS Central

Enable the controller agent on the CloudVision:

```
cvx(config)#cvx
cvx(config-cvx)#no shutdown
```

Once the CloudVision is enabled and it is connected to all the switches, the VXLAN Control service should be enabled.

Enable the Vxlan Service agent:

```
cvx(config)#cvx
cvx(config-cvx)#service vxlan
cvx(config-cvx)#no shutdown
cvx(config-cvx)#exit
cvx2(config)#
```

In addition, on each switch the following commands could be issued to see what reachability has been advertised (to the NSX Controller) and received (from the NSX controller).

```
ToR1#show vxlan controller address-table advertised

ToR1#show vxlan controller address-table received

ToR1#show vxlan flood vtep
```

3.2.3.2 Hardware Switch Controller Service

The Hardware Switch controller Service (HSC) provides an integration point between the NSX controllers and the VCS -- it provides a means for software and hardware switches to exchange state

Following configuration is required on Arista CloudVision:

3.2.3.2.1 Enable the HSC service agent on the CloudVision

```
cvx (config)#cvx
cvx(config-cvx)#no shutdown
cvx(config-cvx)#service hsc
cvx(config-cvx-hsc)#no shutdown
```

OVSDB allows MAC addresses to be distributed among VTEPs, obviating the need for data plane MAC learning. VMware NSX takes advantage of this facility. The default behavior is to use control plane MAC learning.

Next we need to point the HSC service to an NSX Controller. If there is more than one controller, the first may register the others automatically as is the case with NSX.

3.2.3.2.2 Pointing CloudVision to a NSX Controller:

```
cvx(config)#cvx
cvx(config-cvx)#service hsc
cvx(config-cvx-nsx)#manager 10.10.100.2
```

CloudVision connects to the controller via OVSDB with or without SSL authentication. If authentication is not required, CloudVision's management IP in the Management network can be used. Otherwise, the SSL certificate generated by the CVX's OVSDB instance must be provided. The SSL certificate generated by the HSC service is as follows and will be used to register CVX with an NSX controller.

3.2.3.2.3 Displaying HSC SSL certificate:

```
cvx(config)#show hsc certificate
…
-----BEGIN CERTIFICATE-----
…
----END CERTIFICATE-----
cvx(config)#
```

The connection to the OVSDB Controller can be verified as follows.

3.2.3.2.4 Displaying connection to HSC:

```
cvx#show hsc status
HSC is enabled.

Manager                   Connected       State
-------------------------- --------------- -------------
ssl:10.10.100.2:6632      True            ACTIVE (142s)

cvx(config)#show hsc status detail
HSC is enabled.

Manager                   Connected       State
-------------------------- --------------- --------------
ssl:10.10.100.2:6632      True            ACTIVE (1824s)

Hsc agent is Running, OVSDB connection state is Connected.
```

Having established a connection to the OVSDB controller, the HSC Service will publish the inventory of switches managed by CloudVision to OVSDB. For inventory to succeed, LLDP must be enabled on each TOR VTEP:

3.2.3.2.5 Enable LIDP

```
Tor1(config)#lldp run
```

Note: LLDP is enabled by default on all Arista switches.

3.2.4 Registering Arista CloudVision with NSX

Having configured the hardware VTEP, the NSX controller needs to be configured to communicate with CloudVision and obtain models for its inventory of hardware switches.

The user must configure the HSC of their Hardware Gateway with the NSX controller IP address. Note that there are typically several redundant NSX Controllers; only one of them needs to be specified (the others will be automatically discovered.). Refer to section 3.2.4 for CloudVision configuration.

Once the administrator has pointed the HSC to an NSX Controller, they also need to enter the certificate collected in part 3.2.3.2.3 above that will be used by the OVSDB Client on the Controller to connect to the server on the HSC. .

From there, the registration of Arista CloudVision in the NSX UI is relatively straightforward in vCenter: Navigate to the Networking & Security/Service Definition tab. Then select the Hardware Devices menu:

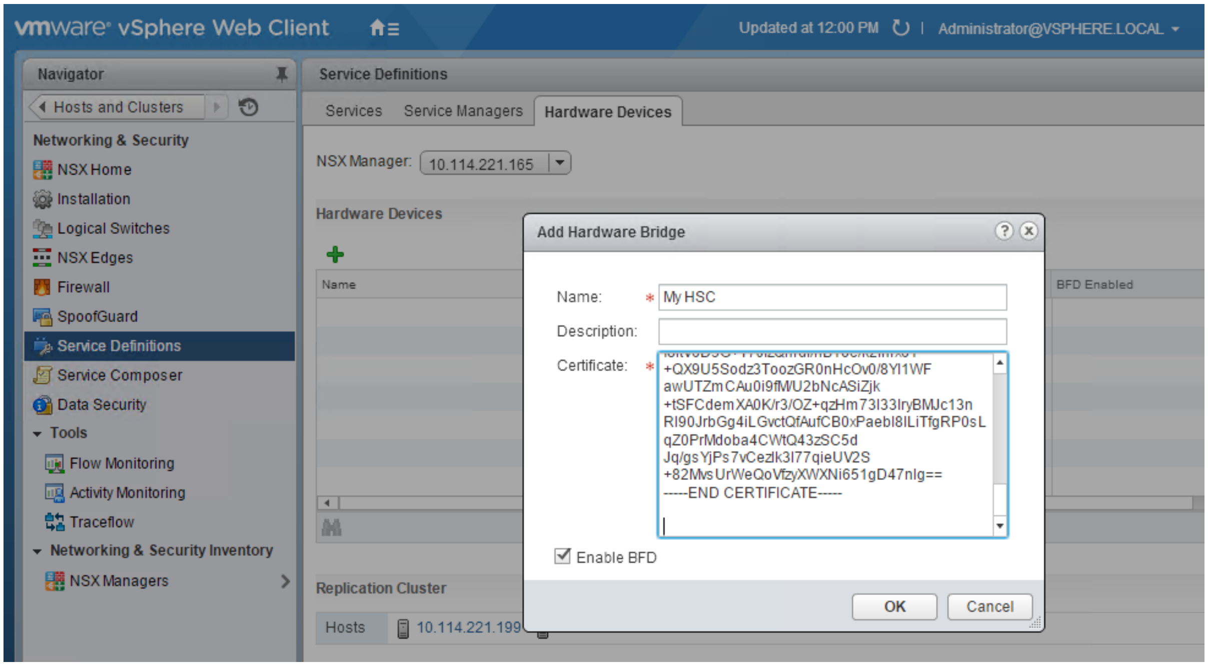

NSX Manager: 10.114.221.165

Name: * My HSC

Description:

Certificate: *

```
+QX9U5Sodz3ToozGR0nHcOv0/8YI1WF
awUTZmCAu0i9fM/U2bNcASiZjk
+tSFCdemXA0K/r3/OZ+qzHm73I33IryBMJc13n
RI90JrbGg4iLGvctQfAufCB0xPaebI8ILiTfgRP0sL
qZ0PrMdoba4CWtQ43zSC5d
Jq/gsYjPs7vCezIk3I77qieUV2S
+82MvsUrWeQoVfzyXWXNi651gD47nIg==
-----END CERTIFICATE-----
```

☑ Enable BFD

Figure8: Registration of a Hardware Gateway

Here, you will add a new profile by clicking on the “plus” sign in the Hardware Devices section. Figure above is showing the resulting pop-up window that has been populated with a profile name and the certificate retrieved from the HSC. Note that BFD is enabled by default, meaning the Arista switches will establish BFD sessions to the Replication Service Nodes. This is critical for protecting against the silent failure of an RSN and VMware will only support configurations running BFD.

Once this configuration is done, the Hardware Gateway should show up as available, as pictured in Figure 9 below:

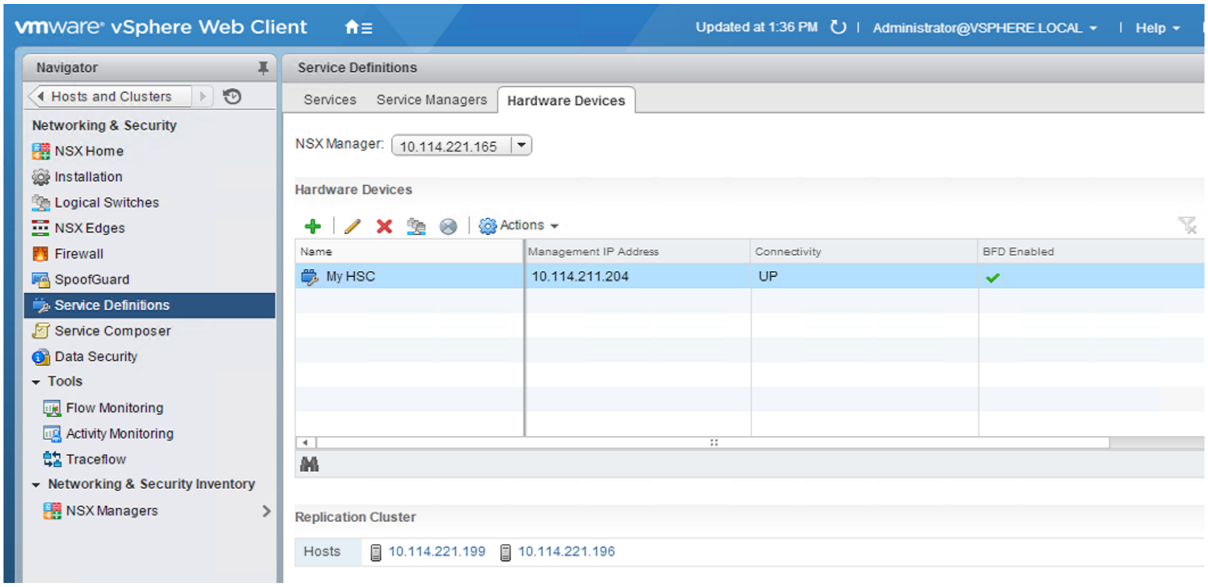

NSX Manager: 10.114.221.165

Hardware Devices

Name	Management IP Address	Connectivity	BFD Enabled
My HSC	10.114.211.204	UP	✓

Replication Cluster

Hosts: 10.114.221.199 10.114.221.196

Figure9: Successful registration

Note also that the Replication Cluster must also be configured on the same screen. The Replication Cluster is the set of Hypervisors that will act as RSNs. The administrator will select a subset of the Hypervisors in the network.

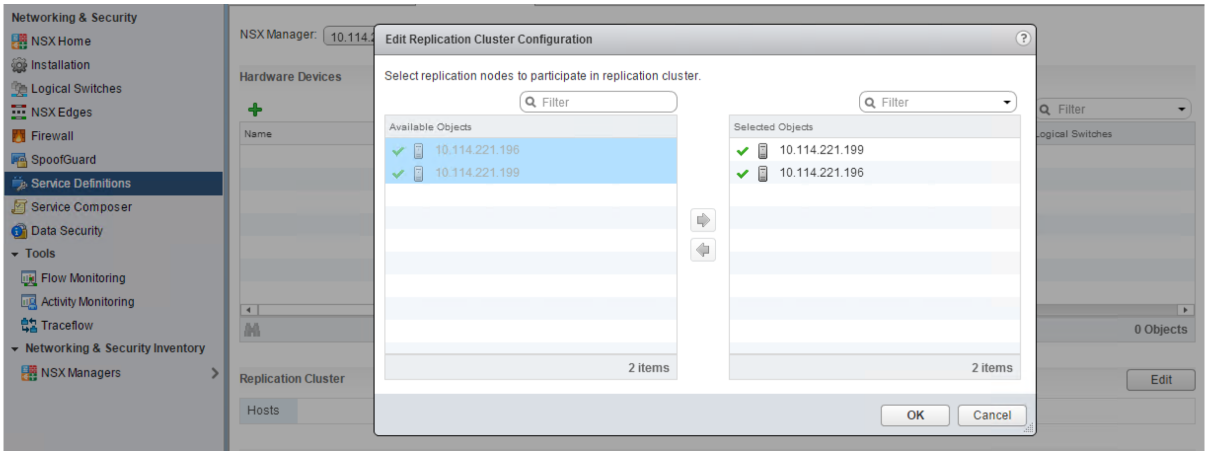

Edit Replication Cluster Configuration

Select replication nodes to participate in replication cluster.

Available Objects
✓ 10.114.221.196
✓ 10.114.221.199
2 items

Selected Objects
✓ 10.114.221.199
✓ 10.114.221.196
2 items

Figure10: Replication Service Node Selection

3.2.4.1 Building a logical switch to a physical switch/physical port/VLAN

Once Arista CloudVision is added to NSX, a Logical Switch can programmatically be mapped to any physical port/VLAN advertised by this gateway. This section will illustrate the mapping of a logical switch to a particular port, leveraging vCenter UI.

First, the administrator selects a Logical Switch in the Network & Security/Logical Switches tab.

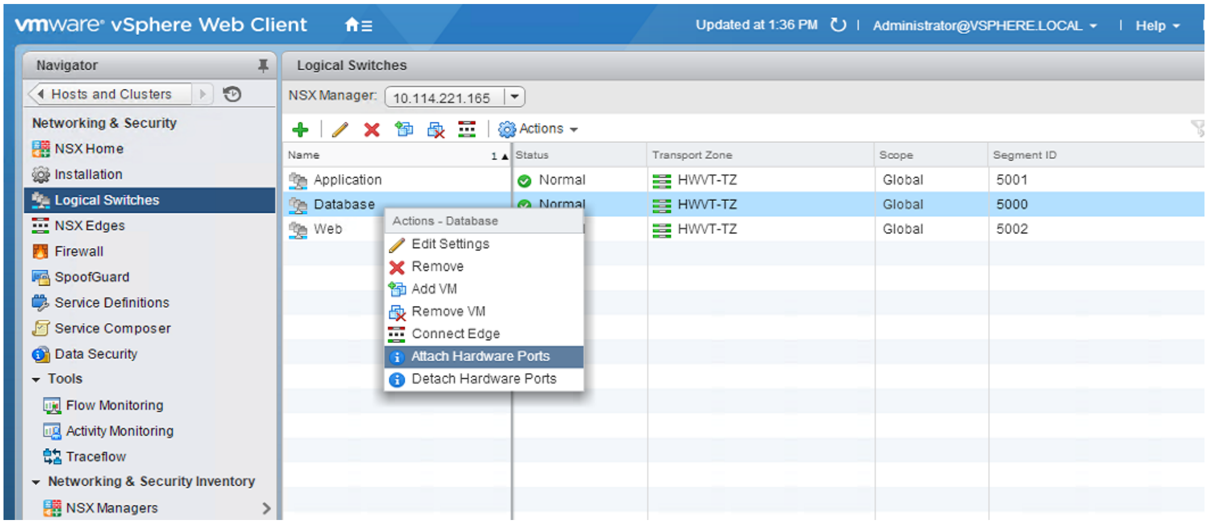

NSX Manager: 10.114.221.165

Name	Status	Transport Zone	Scope	Segment ID
Application	Normal	HWVT-TZ	Global	5001
Database	Normal	HWVT-TZ	Global	5000
Web		HWVT-TZ	Global	5002

Figure11: Logical Switch Selection

Figure above shows a menu that is dropped down by right clicking on the “Database” Logical Switch entry in the table. From there, selecting “Attach Hardware Ports” will open a pop-up (represented in Figure below) specifying a port binding.

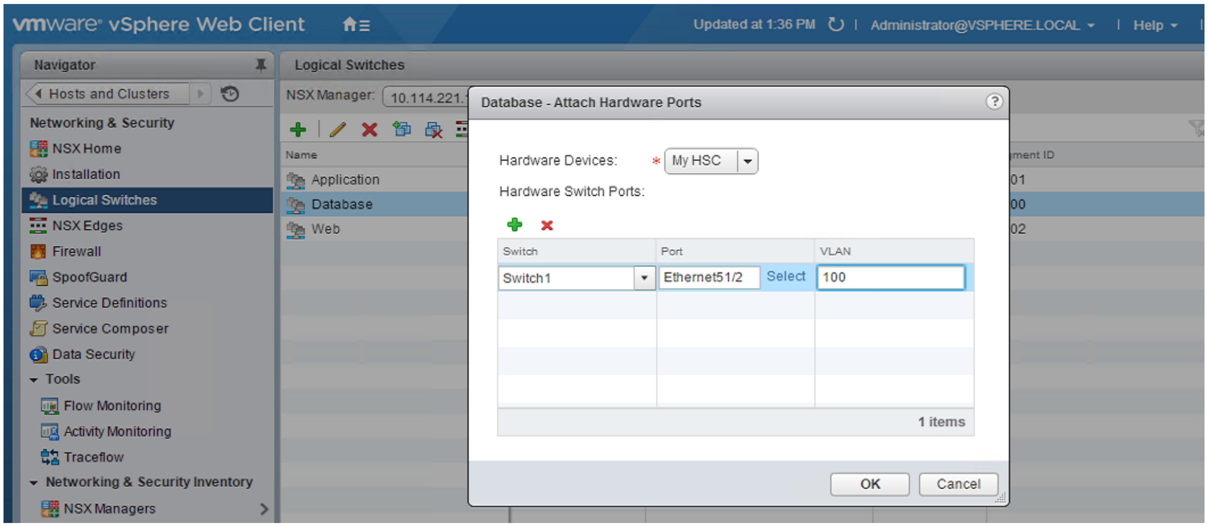

Hardware Devices: * My HSC

Hardware Switch Ports:

Switch	Port	VLAN
Switch1	Ethernet51/2	100

1 items

Figure12: Binding a Logical Switch to a physical switch, physical port and VLAN

Three columns are available in above Figure:

- Physical switch: remember that Arista CloudVision will control several hardware switches, so this selection is necessary to identify which one is concerned by this configuration
- Port: the HSC provides a list of physical ports available for binding on the physical switch.

- VLAN: specify which VLAN tag will be used on the particular port selected. A VLAN value of 0 represents an access port, where the extended Logical Switch traffic will be sent untagged on the port.

Once this selection is done, the Logical Switch is extended to the physical world at Layer 2 on the physical switch/physical port/VLAN specified. Note that several bindings can be achieved for a particular Logical Switch.

3.2.4.2 High Availability with MLAG

Arista supports MLAG towards the compute and VXLAN together on its wide variety of switches, which provides hardware L2 Gateways redundancy for NSX. MLAG with VXLAN on Arista switches provide non-blocking, active-active forwarding and redundancy with hitless failover in an event of switch failure. Arista CloudVision abstracts the details of Multi-Chassis LAG (MLAG) and present a pair of MLAG switches as a single VTEP.

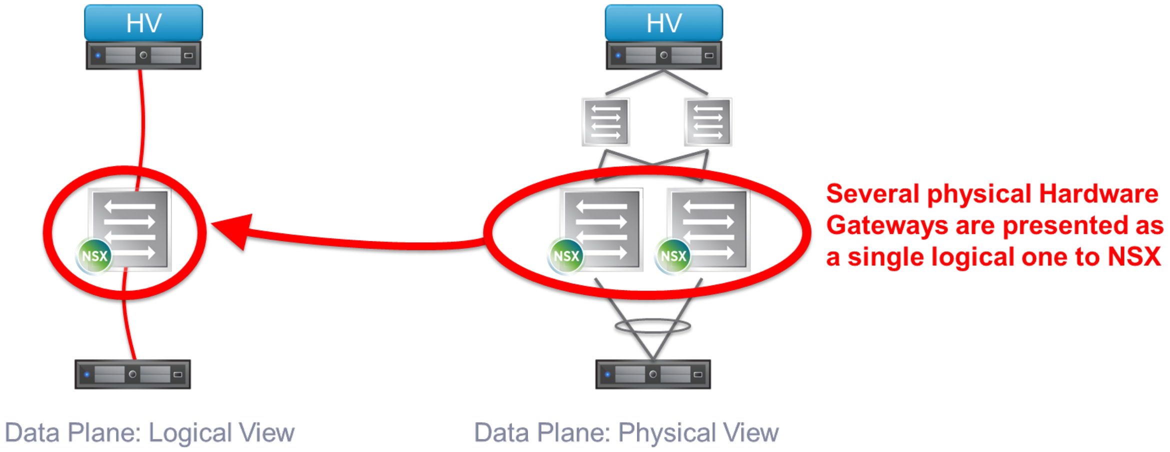
HV
NSX
HV
NSX
NSX
Several physical Hardware Gateways are presented as a single logical one to NSX
Data Plane: Logical View
Data Plane: Physical View

Figure13: MLAG-based redundancy

3.2.4.3 Impact on the scope of Layer 2 in the network

The fact that several Hardware Gateways can be active at the same time can also influence the network design. Typically, a Logical Switch is extended to a VLAN in order to provide connectivity to some service that cannot be virtualized. This service is usually redundant, meaning that its physical implementation spans several different racks in the data center. In the left part of Figure 14 below, some virtual machines attached a Logical Switch access physical servers through a Software Gateway. All the traffic from the Logical Switch to the VLAN 10, where the physical servers are located, have to go through a single bridging instance.

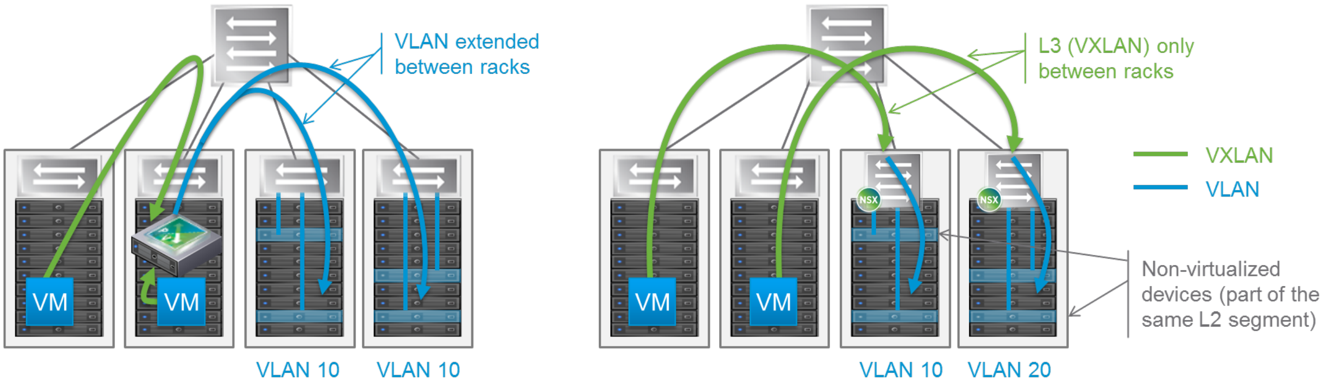

Figure 14: Network

This means that VLAN 10 has to be extended between racks in order to reach all the necessary physical servers. The trend in data center networking in the last few years has been to try to reduce as much as possible the span of Layer 2 connectivity in order to minimize its associated risks and limitations. The right side of Figure 17 shows how this can be achieved leveraging separate active Hardware Gateway. In this alternate design, each rack hosting physical servers is configured with a Hardware Gateway. Thanks to this model, there is no need to extend Layer 2 connectivity between racks, as Logical Switches can extend directly to the relevant Hardware Gateway when reaching physical servers. Note also that the VLANs defined in the racks have local significance (the example is showing that the Logical Switch is extended to VLAN 10 in one rack and VLAN 20 in the other.)

3.3 NSX Components and Cluster Connectivity

The NSX functions and component operation are defined in the VMware® NSX for vSphere Network Virtualization Design Guide. The reader is strongly advised to read the document in order to follow the rationale regarding connectivity to physical network. The NSX components are categorized in following table. The NSX components organization and functions are mapped to appropriate cluster. The VMware® NSX for vSphere Network Virtualization Design Guide calls for organizing NSX components, compute, and management of the virtualized environment. This organization principle is carried in the document and repeated to maintain ease of user readability.

Function	NSX Components	Recommended Clusters Designation
Management Plane	NSX Manager & vCenter Connectivity	Management Cluster
Control Plane	NSX Controller Cluster	Management Cluster* *Can be in Edge Cluster
	Logical Routers Control VM	Edge Cluster
Data Plane East-West	Compute and Edge VDS kernel components – VXLAN forwarding & DLR (Distributed Logical Router)	Compute & Edge Cluster
Data Plane North-South	Edge Service Gateway (ESG)	Edge Cluster
Bridging Traffic	DLR Control VM	Edge Cluster

Table 6: NSX Functions and Components Mapping to Cluster Type

The VMware® NSX for vSphere Network Virtualization Design Guide recommends building three distinct vSphere cluster types. The figure below shows an example of logical components of cluster design to the physical rack placement.

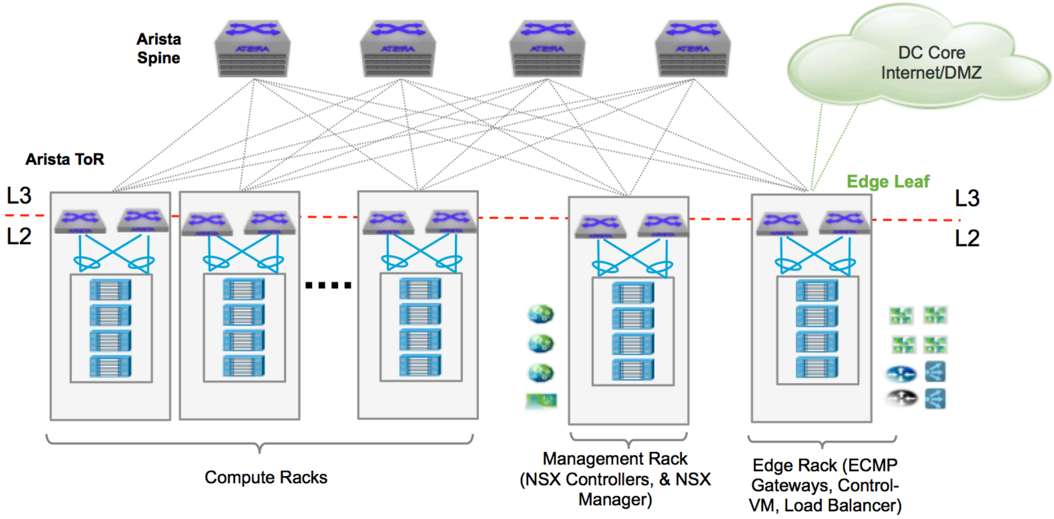
Arista Spine
DC Core Internet/DMZ
Arista ToR
L3
L2
Edge Leaf
L3
L2
Compute Racks
Management Rack (NSX Controllers, & NSX Manager)
Edge Rack (ECMP Gateways, Control-VM, Load Balancer)

Figure 15: Mapping Cluster Types to Functions

As shown in the diagram, edge and management clusters are distributed to separate physical racks and connect to separate ToR switches. For management and edge clusters, the resources are shared or split between two racks to avoid any single rack failure. This also enables scaling.

Note that for even in smaller configurations a single rack can be used to provide connectivity for the edge and management cluster. The key concept is that the edge cluster configuration is localized to a ToR pair to reduce the span of layer-2 requirements; this also helps localize the egress routing configuration to a pair of ToR switches. The localization of edge components also allows flexibility in selecting the appropriate hardware (CPU, memory and NIC) and features based on network-centric functionalities such as firewall, NetFlow, NAT and ECMP routing.

In order to provide a recommendation on connecting host belonging to different cluster types it is important to know the VDS uplink design option. These capabilities are described in section 3.3.3.2

The VMware® NSX for vSphere Network Virtualization Design Guide best practices document calls for a separate VDS for compute and edge cluster. This enables flexibility of choosing VDS uplink configuration mode per cluster type.

3.3.1 Management Cluster Connectivity

The management cluster consists of hosts supporting multiple critical virtual machines and virtual appliances. The NSX manager VM and controllers also typically deployed in management clusters requiring high availability (surviving the failure of the host or ToR/uplink). Typically, the management cluster is not prepared for VXLAN and thus connectivity from the ESXi host is VLAN based port-group on a separate VDS. In order to achieve maximum availability and load sharing, LACP teaming mode is typically recommended. Thus, the Arista switch ports connecting to management hosts require LACP. For Arista switches, this is achieved by enabling traditional layer-2 VLAN-based Multi-chassis LAG (MLAG). Typically, all the traffic types including management, vMotion, and IP storage are carried over LACP.

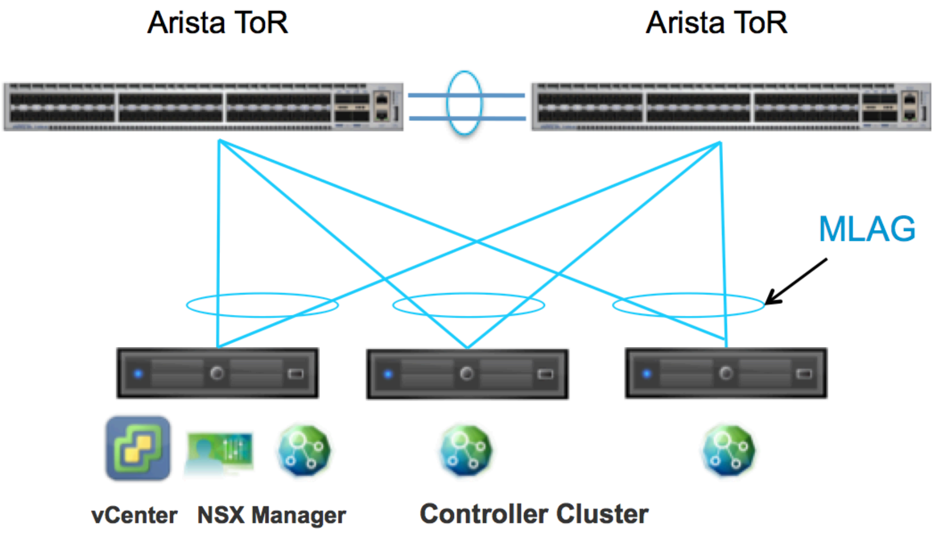

Figure 16: Arista MLAG Connectivity for Management Cluster

3.3.2 Compute Cluster Connectivity

NSX offers a clear departure from the traditional methods, in which the VLANs are only defined once for infrastructure traffic (VXLAN, vMotion, storage, management). The VM connectivity is defined programmatically without relying on the physical network as described in section 3 below. This decoupling enables a repeatable rack design where physical planning (power, space, cooling and cookie-cutter switch configuration) is streamlined. The physical network only requires robust forwarding and adequate bandwidth planning.

The compute cluster requires the most flexibility as it carries multiple types of traffic. Each type of traffic can have its own service level. For example, the storage traffic requires the lowest latency, as opposed to vMotion, which may require higher bandwidth.

Some workloads may have many sources and destinations and require granular load sharing by using multiple VTEPs. The flexibility of selecting teaming mode per traffic type and allowing multiple VTEPs for VXLAN (as described in the VDS uplink configuration section) are two primary reasons for *not* recommending LACP for the compute cluster host's connectivity to Arista switches.

3.3.3 Edge Cluster Connectivity

NSX ESG (Edge Services Gateway) is a multi-function VM, enabling services such as north-south routing, firewall, NAT, load balancing, and SSL-VPN. The capabilities and features are beyond the scope of this paper. Please refer to VMware® NSX for vSphere Network Virtualization Design Guide. This section covers necessary technical details that are pertinent to physical and logical connectivity required. The critical functions provided by the edge cluster hosting multiple edge VMs are:

- Providing on-ramp and off-ramp connectivity to physical networks (north-south L3 routing delivered by NSX edge virtual appliances)
- Supporting centralized logical or physical services (firewall, load-balancers, and logical router control VM, etc.)

The benefits of confining edge clusters to a pair of ToRs (or pair of racks) include:

- Reducing the need to stretch VLANs across compute clusters
- Localizing the routing configuration for N-S traffic, reducing the need to apply any additional configuration knobs for N-S routing on the compute ToRs
- Allowing network admins to manage the cluster workload that is network centric (operational management, BW monitoring and enabling network-centric features such as NetFlow, security, etc.)

3.3.3.1 Arista Switches and NSX Routing

This paper addresses connectivity for north south routing with the ECMP mode of the edge services gateway. The NSX edge gateway provides ECMP (Equal Cost Multi-path) based routing, which allows up to eight VMs presenting 8-way bidirectional traffic forwarding from NSX logical domain to the enterprise DC core or Internet. This represents up to 80 Gbps (8 x10GE interfaces) of traffic that can be offered from the NSX virtual domain to the external network in both directions. It's scalable per tenant, so the amount of bandwidth is elastic as on-demand workloads and/or multi-tenancy expand or contract. The configuration requirements to support the NSX ECMP edge gateway for N-S routing is as follows:

- VDS uplink teaming policy and its interaction with ToR configuration
- Requires two external VLAN(s) per pair of ToR
- Route Peering with Arista Switches

3.3.3.2 VDS uplink design with ESXi Host in Edge cluster

The edge rack has multiple traffic connectivity requirements. First, it provides connectivity for east-west traffic to the VXLAN domain via VTEP; secondly, it provides a centralized function for external user/traffic accessing workloads in the NSX domain via dedicated VLAN-backed port-group. This later connectivity is achieved by establishing routing adjacencies with the next-hop L3 devices. The figure below depicts two types of uplink connectivity from host containing edge ECMP VM.

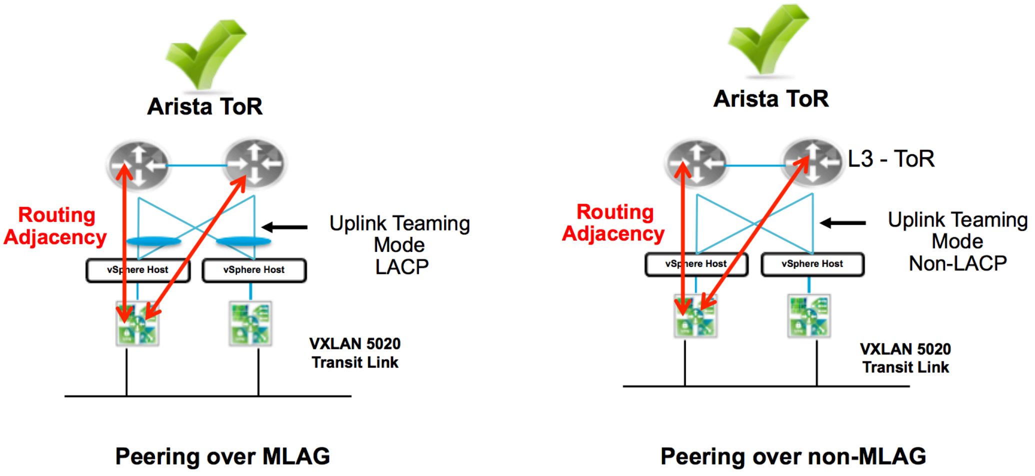

Figure 17: Arista Layer 3 Peering over MLAG Support

3.3.3.3 Edge ECMP Peering and VLAN Design

Once the uplink-teaming mode is determined, the next step is to provide design guidance around VLAN configuration and mapping to uplink as well peering to Arista switches.

The first decision is how many logical uplinks should be deployed on each NSX edge. The recommended design choice is to always map the number of logical uplinks to the number of VDS dvUplinks defined on NSX edge VM available on the ESXi servers hosting the NSX edge VMs. This means always map a VLAN (port-group) to a VDS dvUplink, which then maps to a physical link on the ESXi host that connects to the Arista switch, over which an edge VM forming a routing peer relationship with Arista switch.

In the example shown in below, NSX Edge ECMP VMs (E1-E8) are deployed on ESXi hosts with two physical uplinks connected to the Arista ToR switches. Thus, the recommendation is to deploy two logical uplinks on each NSX edge. Since an NSX edge logical uplink is connected to a VLAN-backed port-group, it is necessary to use two external VLAN segments to connect the physical routers and establish routing protocol adjacencies.

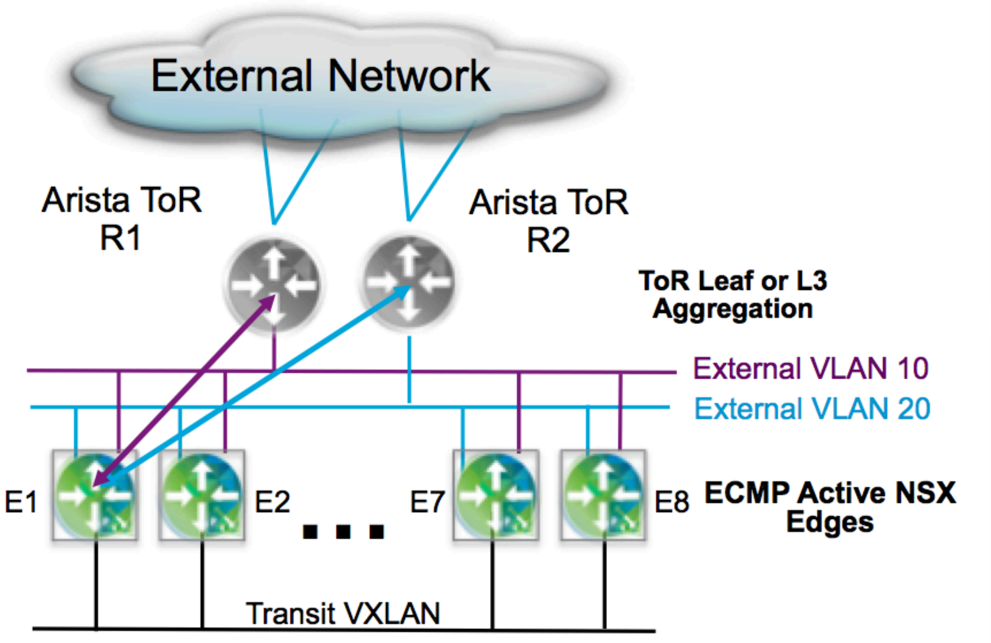

Figure 18: VLAN, dvUplink and Routing Peering Mapping

As shown in the figure above, each ECMP node peers over its respective external VLANs to exactly one Arista router. Each external VLAN is defined only on one ESXi uplink (in the figure above external VLAN10 is enabled on uplink toward R1 while external VLAN20 on the uplink toward R2). This is done so that under normal circumstances both ESXi uplinks can be concurrently utilized to send and receive north-south traffic, even without requiring the creation of a port-channel between the ESXi host and the ToR devices.

In addition, with this model a physical failure of an ESXi NIC would correspond to a logical uplink failure for the NSX edge running inside that host, and the edge would continue sending and receiving traffic leveraging the second logical uplink (the second physical ESXi NIC interface).

In order to build a resilient design capable of tolerating the complete loss of an edge rack, it is also recommended to deploy two sets of four edge gateways in two separate edge racks. The below table describes the necessary configuration with ECMP edge.

Port Group	VLAN	dvUplink 1	dvUplink 2	Load Balancing
VTEPs	XXX	Active	Active	SRC_ID
Edge-External-1	YYY	Active	NA	SRC_ID
Edge-External-2	ZZZ	NA	Active	SRC_ID

Table 7: Edge Cluster VDS Configuration

3.3.3.4 NSX Edge Routing Protocol Timer Recommendations

The NSX edge logical router allows dynamic as well as static routing. The recommendation is to use dynamic routing protocol to peer with Arista switches in order to reduce the overhead of defining static routes every time the logical network is defined. The NSX edge logical routers support OSPF, BGP and IS-IS routing protocol. The NSX edge ECMP mode configuration supports reduction of the routing protocol “hello and hold” timer to improve failure recovery of traffic in the case of node or link failure. The minimum recommended timer for both OSPF and BGP is shown in table below.

Routing Protocol	Keep Alive or Hello Timer	Hold Down Timer
OPSF	1	3
BGP	1	3

Table 8: Edge Cluster VDS Configuration

4 Benefits of NSX Architecture with Arista Infrastructure

NSX enables users to build logical services for networking and security without having to make configuration changes to the physical infrastructure. In this case, once the Arista switches are

configured to provide IP connectivity and the routing configuration is provisioned as described above, we can continue to deploy new services with NSX.
Let us look at some examples that show how applications can be deployed with NSX for network virtualization.

4.1 Logical Layer Connectivity

The figure below shows how logical layer-2 segments can be built. Here we can observe that servers in the physical infrastructure can be in different subnets, yet an overlay network enables VMs to be in the same subnet and layer-2 adjacent, essentially providing topology-independent connectivity and mobility beyond the structured topology constraint imposed by physical networking.

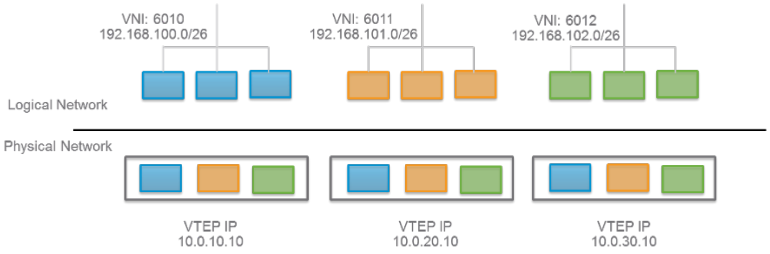

Figure 19: Logical Layer 2

NSX builds multicast-free VXLAN based overlay networks. One can extend layer-2 and IP subnets across servers connected to different ToR Arista switches in a layer-3 fabric. This layer-2 adjacency between the VMs can be established independently of the physical network configuration. New logical networks can be created on demand via NSX, decoupling the logical virtual network from the physical network topology.

4.2 Routing to Physical Infrastructure

In order to route from the logical network to the physical network, NSX can learn and exchange routes with the physical infrastructure in order to reach resources such as a database server or a non-virtualized application, which could be located on different subnet on a physical network.

NSX provides a scale-out routing architecture with the use of ECMP between the NSX distributed router and the NSX Edge routing instances as shown in the figure below. The NSX Edges can peer using dynamic routing protocols (OSPF or BGP) with the physical routers and provide scalable bandwidth. In the case of a Arista switch infrastructure, the routing peer could be a any Arista ToR.

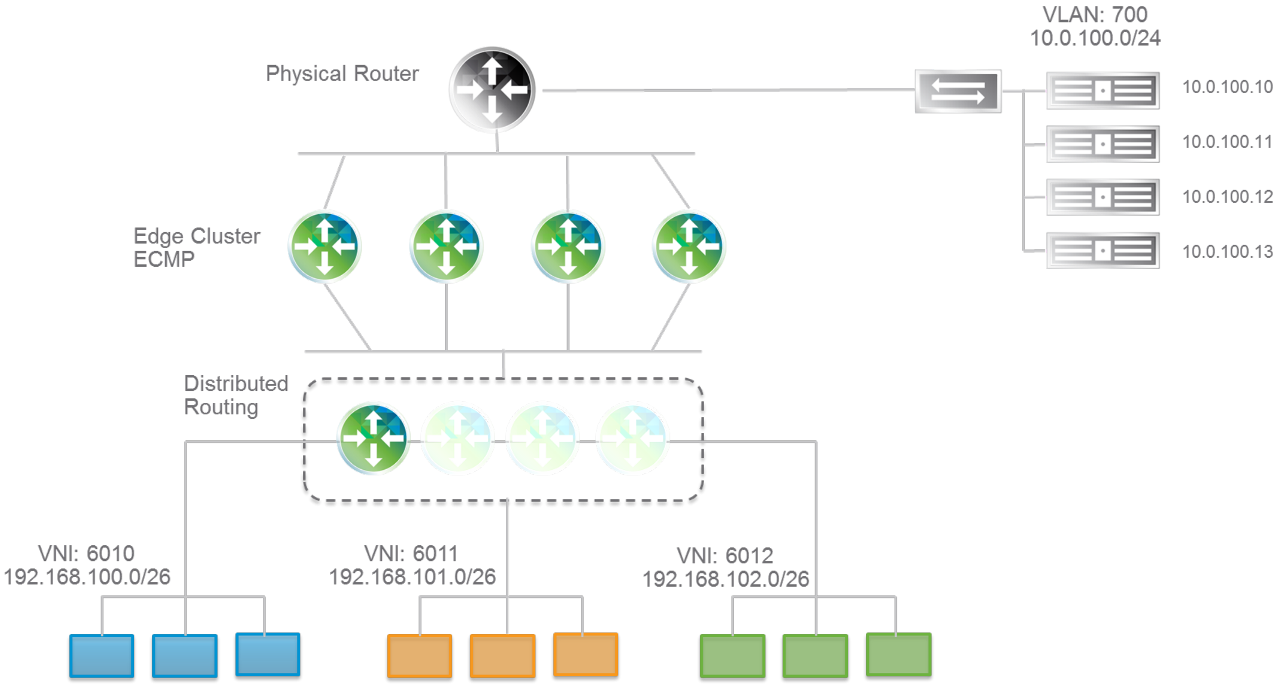

Figure 20: Routing from Logical to Physical Workloads

4.3 Simplified Operation and Scalability with Arista CloudVision

Arista's CloudVision® solution simplifies the integration of physical and virtual environments. CloudVision leverages a network-wide database, which collects the 'state' of the entire physical network and presents a single, open interface for VMWare NSX™, to integrate with the physical network. Using CloudVision as the integration point allows for the details of the physical network to be abstracted away from cloud orchestrators and overlay controllers. In addition, CloudVision simplifies the NSX integration effort because the NSX only needs to integrate with CloudVision itself. This allows customers to avoid the complication of certifying NSX with the many combinations of hardware and software versions. CloudVision in turn provides the aggregate state of the physical network for the most effective physical to virtual synchronization. This approach provides a simpler and more scalable approach for controller integration to the physical network.

CloudVision is build on open APIs, including OVSDB and JSON, which provide a standards-based approach for this integration

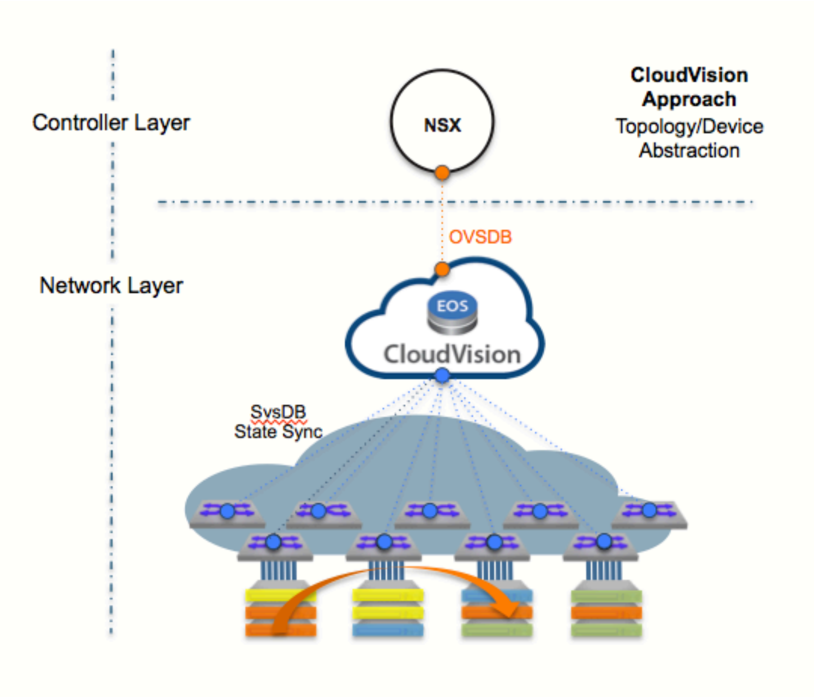

Figure 22: Single Point of Integration for NSX

4.4 Logical and Physical Network Visibility with Arista VMTracer

Arista's VM Tracer feature for NSX is natively integrated into Arista EOS. It automates discovery of directly connected virtual infrastructure, streamlining dynamic provisioning of related VLANs and port profiles on the network. Arista's switches utilize the VMware vCenter and NSX Manager APIs to collect provisioning information. VM Tracer then combines this information with data from the switch's database to provide a clear and concise mapping of the virtual to physical network. Customers can get real time tracking of logical switches and VMs from a single CLI on any Arista switch in the network.

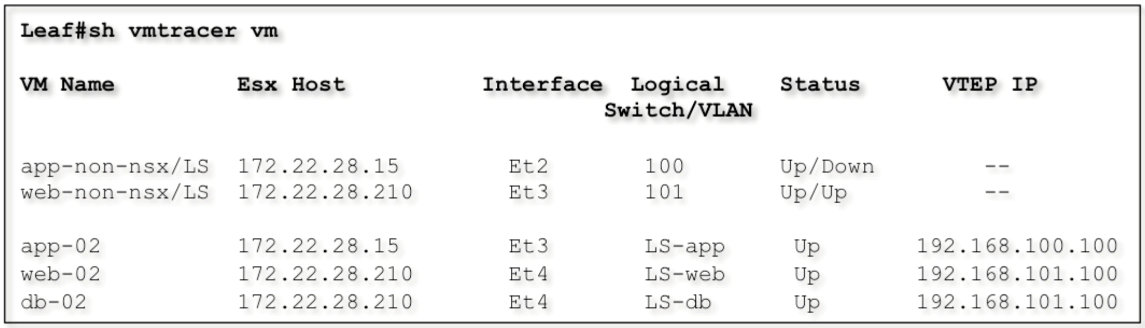

```
Leaf#sh vmtracer vm

VM Name          Esx Host           Interface  Logical     Status      VTEP IP
                                            Switch/VLAN

app-non-nsx/LS   172.22.28.15         Et2        100       Up/Down        --
web-non-nsx/LS   172.22.28.210        Et3        101       Up/Up          --

app-02           172.22.28.15         Et3        LS-app     Up       192.168.100.100
web-02           172.22.28.210        Et4        LS-web     Up       192.168.101.100
db-02            172.22.28.210        Et4        LS-db      Up       192.168.101.100
```

Figure 23: Logical and Physical visibility through Arista VMTracer

4.5 Security with Distributed Firewall

NSX by default enables the distributed firewall on each VM at the vNIC level. The firewall is always in the path of the traffic to and from VM. The key benefit is that it can reduce the security exposure at the root for east-west traffic and not at the centralized location. Additional benefits of distributed firewall include:

- Eliminating the number of hops (helps reduce bandwidth consumption to and from the ToR) for applications traversing to a centralized firewall
- Flexible rules sets (rules sets can be applied dynamically, using multiple objects available in vSphere such as logical SW, cluster and DC)
- Allowing the policy and connection states to move with VM vMotion
- Developing an automated workflow with programmatic security policy enforcement at the time of deployment of the VM via cloud management platform, based on exposure criteria such as tiers of security levels per client or application zone

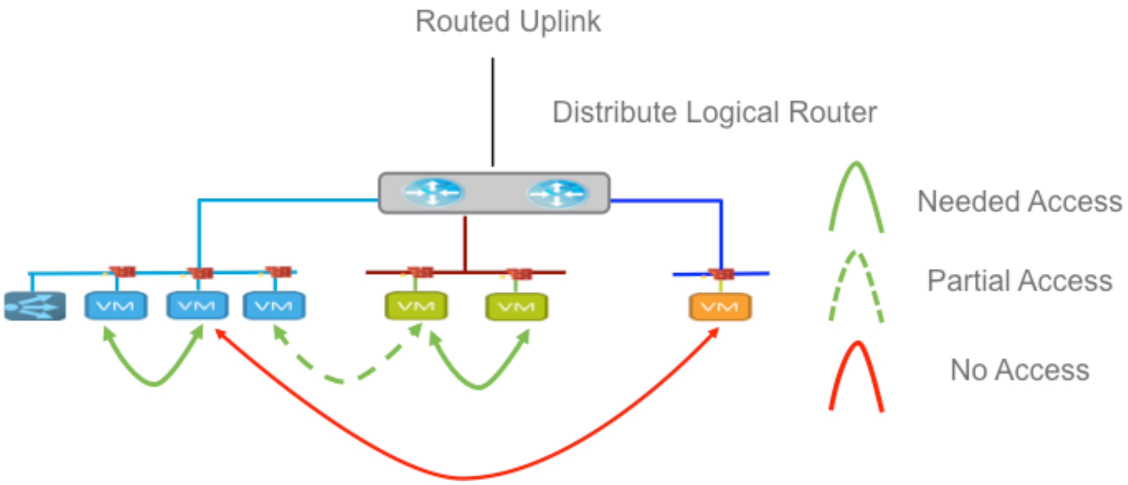

Figure 24: Micro-segmentation and Protection of Traffic

As shown in the figure above, the designer now has flexibility in building a sophisticated policy since policy is not tied to physical topology. The policy can be customized for inter- and intra-layer-2 segment(s), complete or partial access, as well as managing N-S rules sets that can be employed directly at the VM level with edge firewall being an option for the interdomain security boundary.

Micro-segmentation as shown in the figure above allows creating a PCI zone within a shared segment, allowing sophisticated security policies for desktops in a VDI environment as well as eliminating the scaling limitation of centralized access-control ACL management.

4.6 Flexible Application Scaling with Virtualized Load Balancer

Elastic application workload scaling is one of the critical requirements in today's data center. Application scaling with a physical load balancer may not be sufficient given the dynamic nature of self-service IT and DevOps style workloads. The load-balancing functionality natively supported in the edge appliance covers most of the practical requirements found in deployments. It can be deployed programmatically based on application requirements with appropriate scaling and features. The scale and application support level determines whether the load balancer can be configured with layer-4 or layer-7 services. The topology wise the load balancer can be deployed either in-line or in single-ARM mode. The mode is selected based on specific application requirements, however the single-ARM design offers extensive flexibility since it can be deployed near the application segment and can be automated with the application deployment.

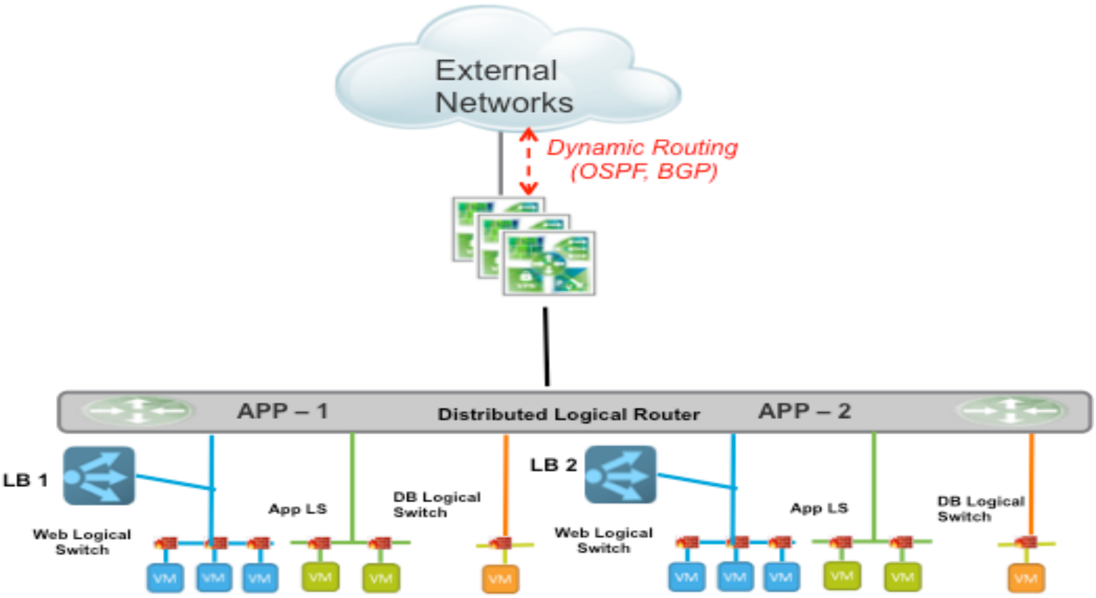

Figure 25: Logical Load Balancing per Application

The figure above shows the power of a software-based load-balancer in which multiple instances of the load-balancer serve multiple applications or segments. Each instance of the load-balancer is an edge appliance that can be dynamically defined via an API as needed and deployed in a high-availability mode. Alternatively, the load balancer can be deployed in an in-line mode, which can serve the entire logical domain. The in-line load-balancer can scale via enabling multi-tier edge per application such that each application is a dedicated domain for which first-tier edge is a gateway for an application, the second-tier edge can be an ECMP gateway to provide the scalable north-south bandwidth.

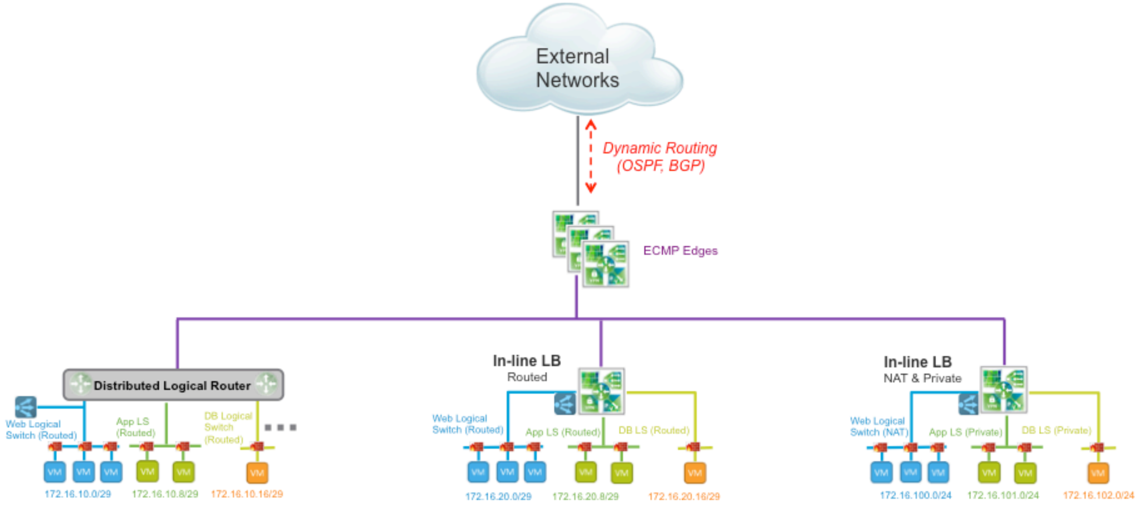

Figure 26: Scaling Application and Services with NSX

As one can observe from the figure above, the first application block on the left is allowing a single-ARM load-balancer with distributed logical routing. The center and the right block of the application allow an in-line load-balancer with either routed or NAT capability respectively. The second-tier edge is enabled with ECMP mode to allow the application to scale on demand from 10GB to 80GB and more.

5 The Arista Advantage

The rise of cloud networking has driven many recent networking innovations. The increased network efficiency brought by focus on automation, orchestration, and programmability techniques is in large part due to the benefits seen in cloud network environments and then re-applied at varying scale to other network environments. As a leader in cloud networking, Arista has applied the lessons learned from the cloud network environments to build a better network operating system, Arista EOS® (Extensible Operating System).

Arista's scale-out cloud network designs are underpinned on many other features of EOS® that we have not mentioned so far in this paper.

Smart System Upgrade (SSU)

Smart System Upgrade (SSU) is an EOS network solution designed to allow maintenance to be performed on any infrastructure element, without adversely impacting application traffic flow. SSU is built on the same cloud networking principles of simplicity, flexibility, and repeatability. SSU uses a simplified stateless approach to ensure that application traffic flow is not impacted.

Zero Touch Provisioning (ZTP) / Zero Touch Replacement (ZTR)

When initially deploying or replacing a switch, ZTP can be used to configure a switch without user intervention; it is as simple as rack, connect and power-on. Built to fully leverage the power of Arista EOS, ZTP provides a flexible solution, provisioning the network infrastructure without requiring a network engineer to be present at the time of initial install or replacement. The ZTP process runs by default at system boot, and based on administrator preferences can ensure the proper software image is installed and complete auto-provisioning of the switch is performed. Utilizing well-understood standards-based protocols, ZTP can leverage the same services that your servers already use, no retraining required. Used as a key component of a cloud data center, Arista ZTP/ZTR will help ensure that any switch is deployed running the proper version of software and with the proper configuration, and can be gracefully inserted into the network.

Automation Integration with eAPI

The foundation that enables EOS to effortlessly carry out workflow automation is eAPI. Providing a pragmatic interface to EOS, eAPI allows the network engineering and operations team to build robust automation scripts. Unlike many operational interfaces, eAPI utilizes the same commands and structure as the command line interface, simplifying the building of tools and scripts that interact with the system. These automation scripts return structured data from the target EOS device. This structured data, in a JSON format, is easily parse-able by any scripting language, providing the engineer or operator maximum flexibility in choice of tools.

Network Telemetry

Network Telemetry is a new model for faster troubleshooting from fault detection to fault isolation. Network Telemetry streams data about network state, including both underlay and overlay network statistics, to applications from Splunk, ExtraHop, Corvil and Riverbed. With critical infrastructure information exposed to the application layer, issues can be proactively avoided.

Openworkload

Open Workload is a network application enabling open workload portability, automation through integration with leading virtualization and orchestration systems, and simplified troubleshooting by offering complete physical and virtual visibility.

- Seamless Scaling - full support for network virtualization, connecting to NSX controllers
- Integrated Orchestration - interfaces to VMware NSX™ to simplify provisioning
- Workload Visibility to the VM-level, enabling portable policies, persistent monitoring, and

rapid troubleshooting of cloud networks

Designed to integrate with VMware, Arista's open architecture allows for integration with any virtualization and orchestration system

Advanced Event Management (AEM)

Advanced Event Management (AEM) provides a set of tools that can augment and enhance the capabilities of EOS, providing key indicators and actionable items for EOS resources in real time. Leveraging Event Manager, EOS can react and respond to various activities happening in the network. For instance, you can be notified of the link status or protocol state of an interface or set of interfaces. Upon receiving the notification, a set of predetermined tasks can be performed. Tasks might include things such as changing the links IGP metric or gracefully removing the node out of network. Event Monitor captures, catalogs and stores all moves, adds and deletes to ephemeral state tables in EOS. Using the information that is captured by Event

6 Conclusion

The VMware network virtualization solution addresses current challenges with physical network and computing infrastructure, bringing flexibility, agility and scale to VXLAN-based logical networks. Along with the ability to create on-demand logical networks using VXLAN, the NSX Edge gateway helps users deploy various logical network services such as firewall, DHCP or NAT. This is possible due to its ability to decouple the virtual network from the physical network and then reproduce the properties and services in the virtual environment.

In conclusion, Arista and VMware are delivering the industry's first scalable best-of-breed solution for network virtualization in the Software Defined Data Center. Cloud providers, enterprises and web customers will be able to drastically speed business services, mitigate operational complexity, and reduce costs. All of this is available now from a fully automated and programmatic SDDC solution that bridges the virtual and physical infrastructure.

ARISTA

Santa Clara—Corporate Headquarters
5453 Great America Parkway
Santa Clara, CA 95054
Tel: 408-547-5500
www.arista.com

Ireland—International Headquarters
4130 Atlantic Avenue
Westpark Business Campus
Shannon
Co. Clare, Ireland

Singapore—APAC Administrative Office
9 Temasek Boulevard
#29-01, Suntec Tower Two
Singapore 038989