Dell EMC VxBlock System 1000

for VMware NSX 6.4 Architecture Overview

June 2021
Rev. Document revision 1.3

DELL Technologies

Notes, cautions, and warnings

NOTE: A NOTE indicates important information that helps you make better use of your product.

CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

Contents

Revision history

Date	Document revision	Description of changes
June 2021	1.3	Updated for VMware vSphere 7.0.
September 2019	1.2	Added support for VMware vSphere 6.7 with NSX 6.4. Added support for Intel Xeon 5218 CPU on Cisco UCS C-Series M5 servers on the edge cluster.
March 2019	1.1	Added support for Cisco UCS C220 M5 Servers on the edge cluster.
October 2018	1.0	Initial version

1

Introduction to VMware NSX 6.4 network virtualization

VMware NSX network virtualization is part of the software-defined data center that offers cloud computing on VMware virtualization technologies.

With VMware NSX, virtual networks are programmatically provisioned and managed independently of the underlying hardware. VMware NSX reproduces the entire network model in software, enabling a network topology to be created and provisioned in seconds. Network virtualization abstracts Layer 2 switching and Layer 3 routing operations from the underlying hardware, similar to what server virtualization does for processing power and operating systems.

VMware vSphere cluster summary for VMware NSX

The management cluster differs between AMP-3S and AMP-VX. The following sections describe the VMware vSphere clusters for VMware NSX with AMP-3S and AMP-VX.

Management cluster with AMP-3S

In AMP-3S, the VMware NSX management appliance and three VMware NSX controllers reside in the Advanced Management Platform.

The VMware NSX controllers provide failover and scalability. The management cluster also includes the VMware vCenter Server, which attaches to the edge and compute clusters and provides the following functionality:

- Allows the VMware NSX manager appliance to be deployed into the management cluster.
- Allows the VMware NSX controllers to be deployed in the management cluster (AMP-3S) or on the edge cluster (AMP-VX).
- Allows the VMware NSX edge services gateway (ESG) and VMware NSX distributed logical routers (DLRs) to be deployed in the edge cluster, which resides on the Cisco UCS C-Series Rack Mount Servers.

Management cluster with AMP-VX

In AMP-VX, the NSX managers reside in the management cluster and the three VMware NSX controllers reside in edge cluster of the production VMware vCenter Server.

- The VMware NSX management appliance is deployed into the Management cluster through the management VMware vCenter Server.
- The VMware NSX controller is deployed into edge cluster through the production VMware vCenter Server.
- The VMware NSX ESG and VMware NSX DLRs are deployed in the edge cluster, which resides on the Cisco UCS C-Series Rack Mount Servers through the production VMware vCenter Server.

Edge cluster

The edge cluster includes the ESG VMs that provide external connectivity to the physical network and the DLRs that provide routing and bridging.

Compute cluster

The compute cluster includes the production VMs. There can be more than one compute cluster.

The following illustration shows the components and functions in each VMware NSX cluster on a VxBlock System 1000 with AMP-3S:

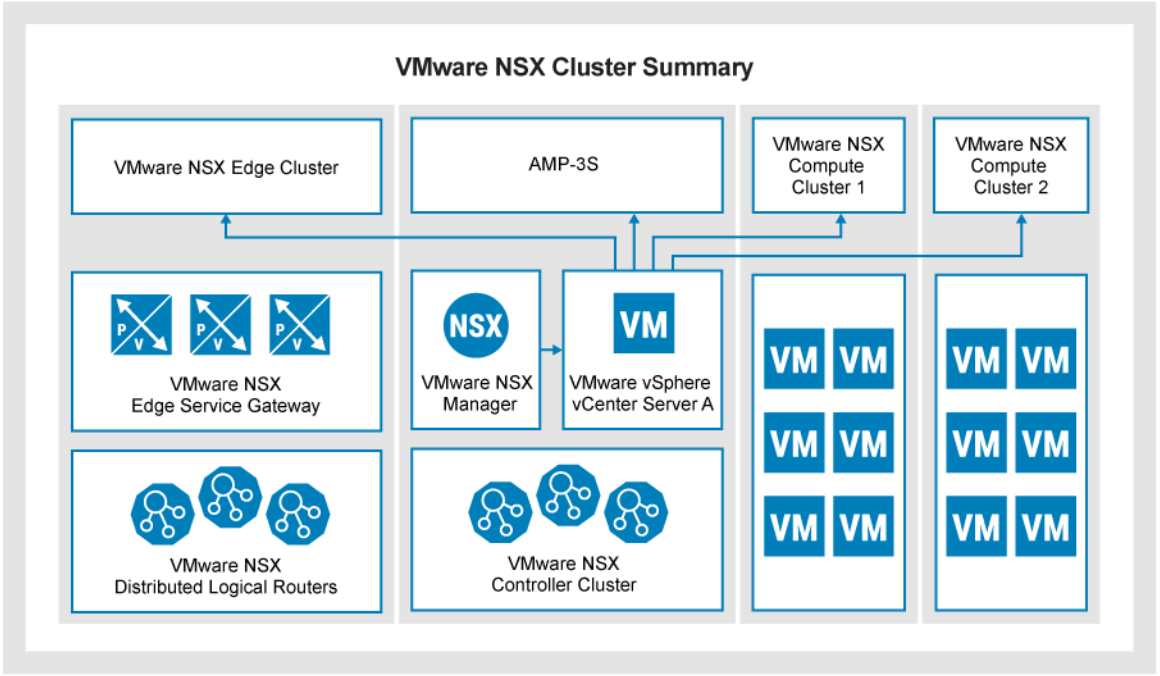

NOTE: More than one compute cluster can exist in the VMware vCenter Server.

The following illustration shows the components and functions in each VMware NSX cluster on a VxBlock System 1000 with AMP-VX:

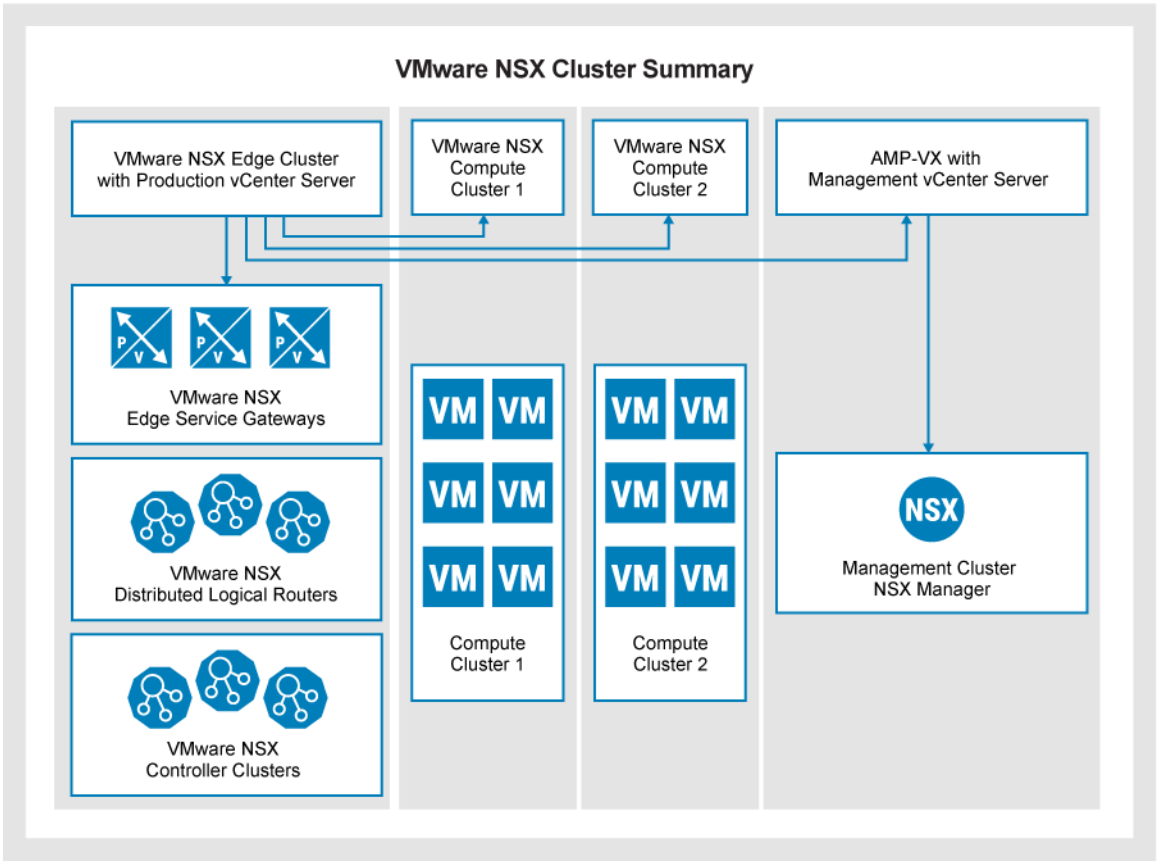

Cross VMware vCenter Server option

VMware NSX 6.4 running on VMware vSphere 6.x allows the data plane to span multiple VMware vCenter Server instances in the same domain. Universal objects (transport zone, logical switches, DLRs, and transit network logical switch) are created on the primary VMware vCenter Server instance and replicated across the secondary VMware vCenter Servers. Up to seven secondary VMware vCenter Servers are supported in a cross VMware vCenter deployment.

For more information on the Cross VMware vCenter Server option, refer to the following sections of the *Cross-vCenter NSX Installation Guide*:

- Benefits of cross-vCenter NSX
- Support matrix for NSX service in cross-vCenter NSX

Local objects (as opposed to universal objects) reside on a single VMware vCenter Server. Local objects do not replicate across multiple VMware vCenter Servers, but can use such features as Layers 4 through 7 services and Layer 2 bridging.

VMware NSX logical topology with local objects

This local VMware vCenter Server environment supports Layers 4 through 7 services and Layer 2 bridging.

The following illustration shows the local objects (DLRs and local transit network logical switch) integrated in the VMware NSX single VMware vCenter Server topology:

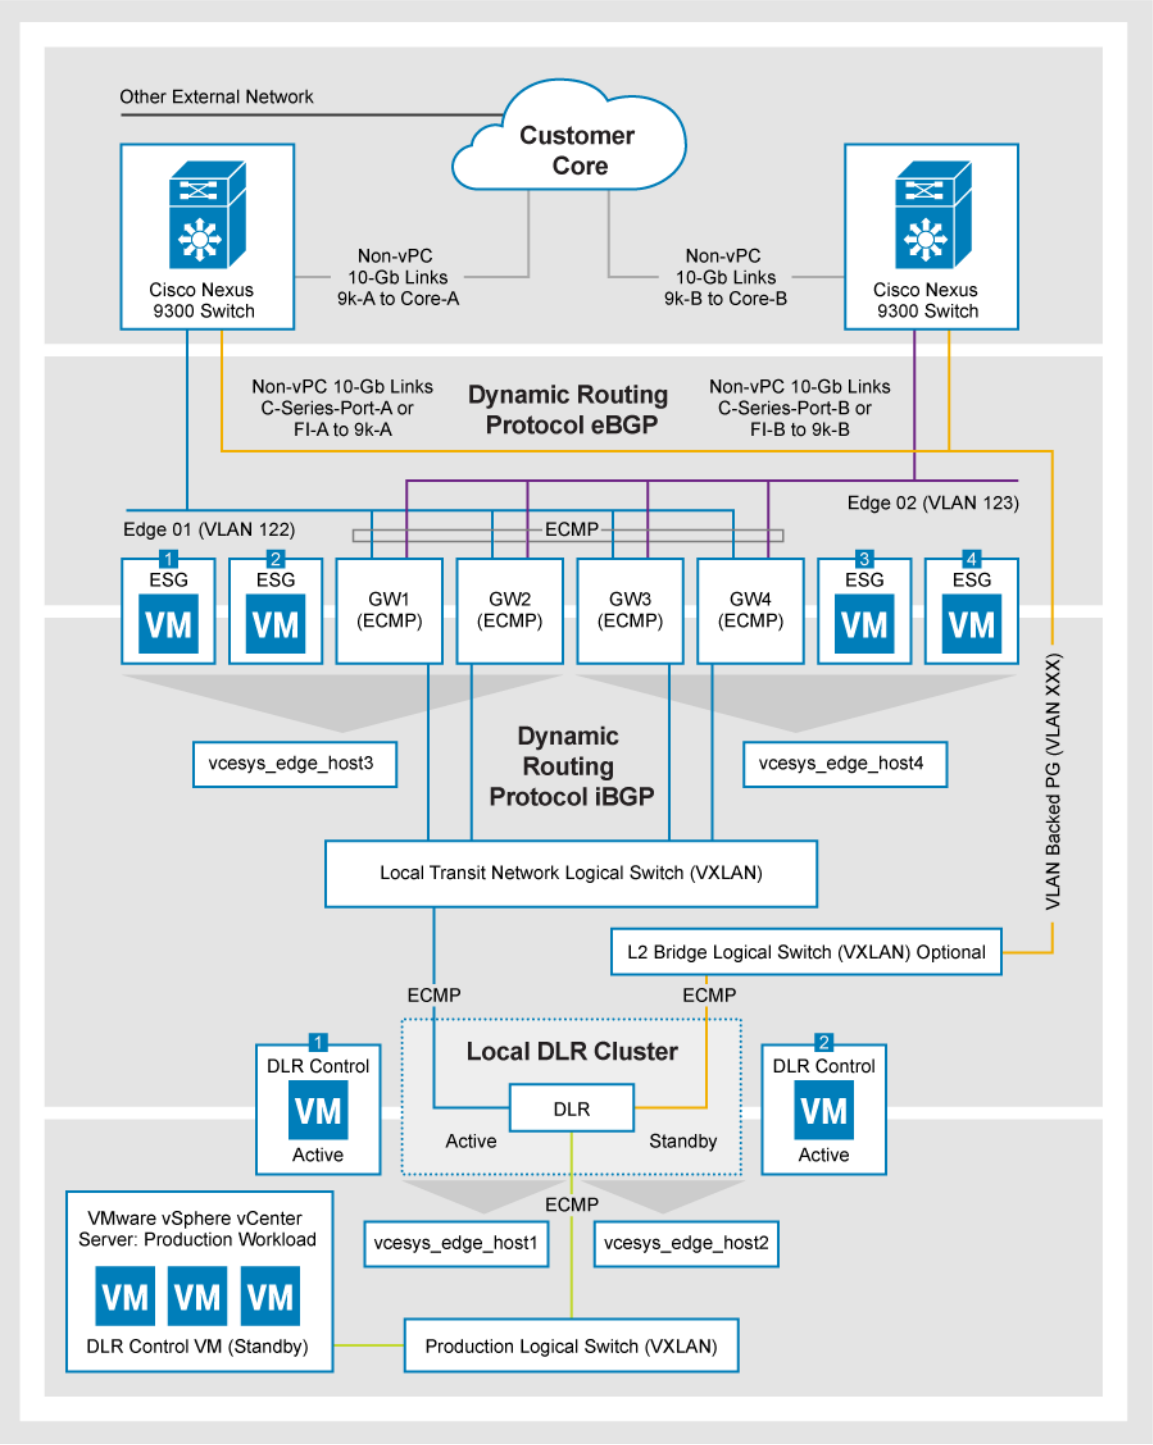

VMware NSX local topology with universal objects (cross VMware vCenter Server option)

The following illustration shows a cross VMware vCenter Server with universal objects, DLRs and universal transit network logical switch, integrated in the VMware NSX topology. This topology cannot support Layers 4 through 7 services and Layer 2 bridging. However, the universal production logical switch spans VMware vCenter Servers A and B.

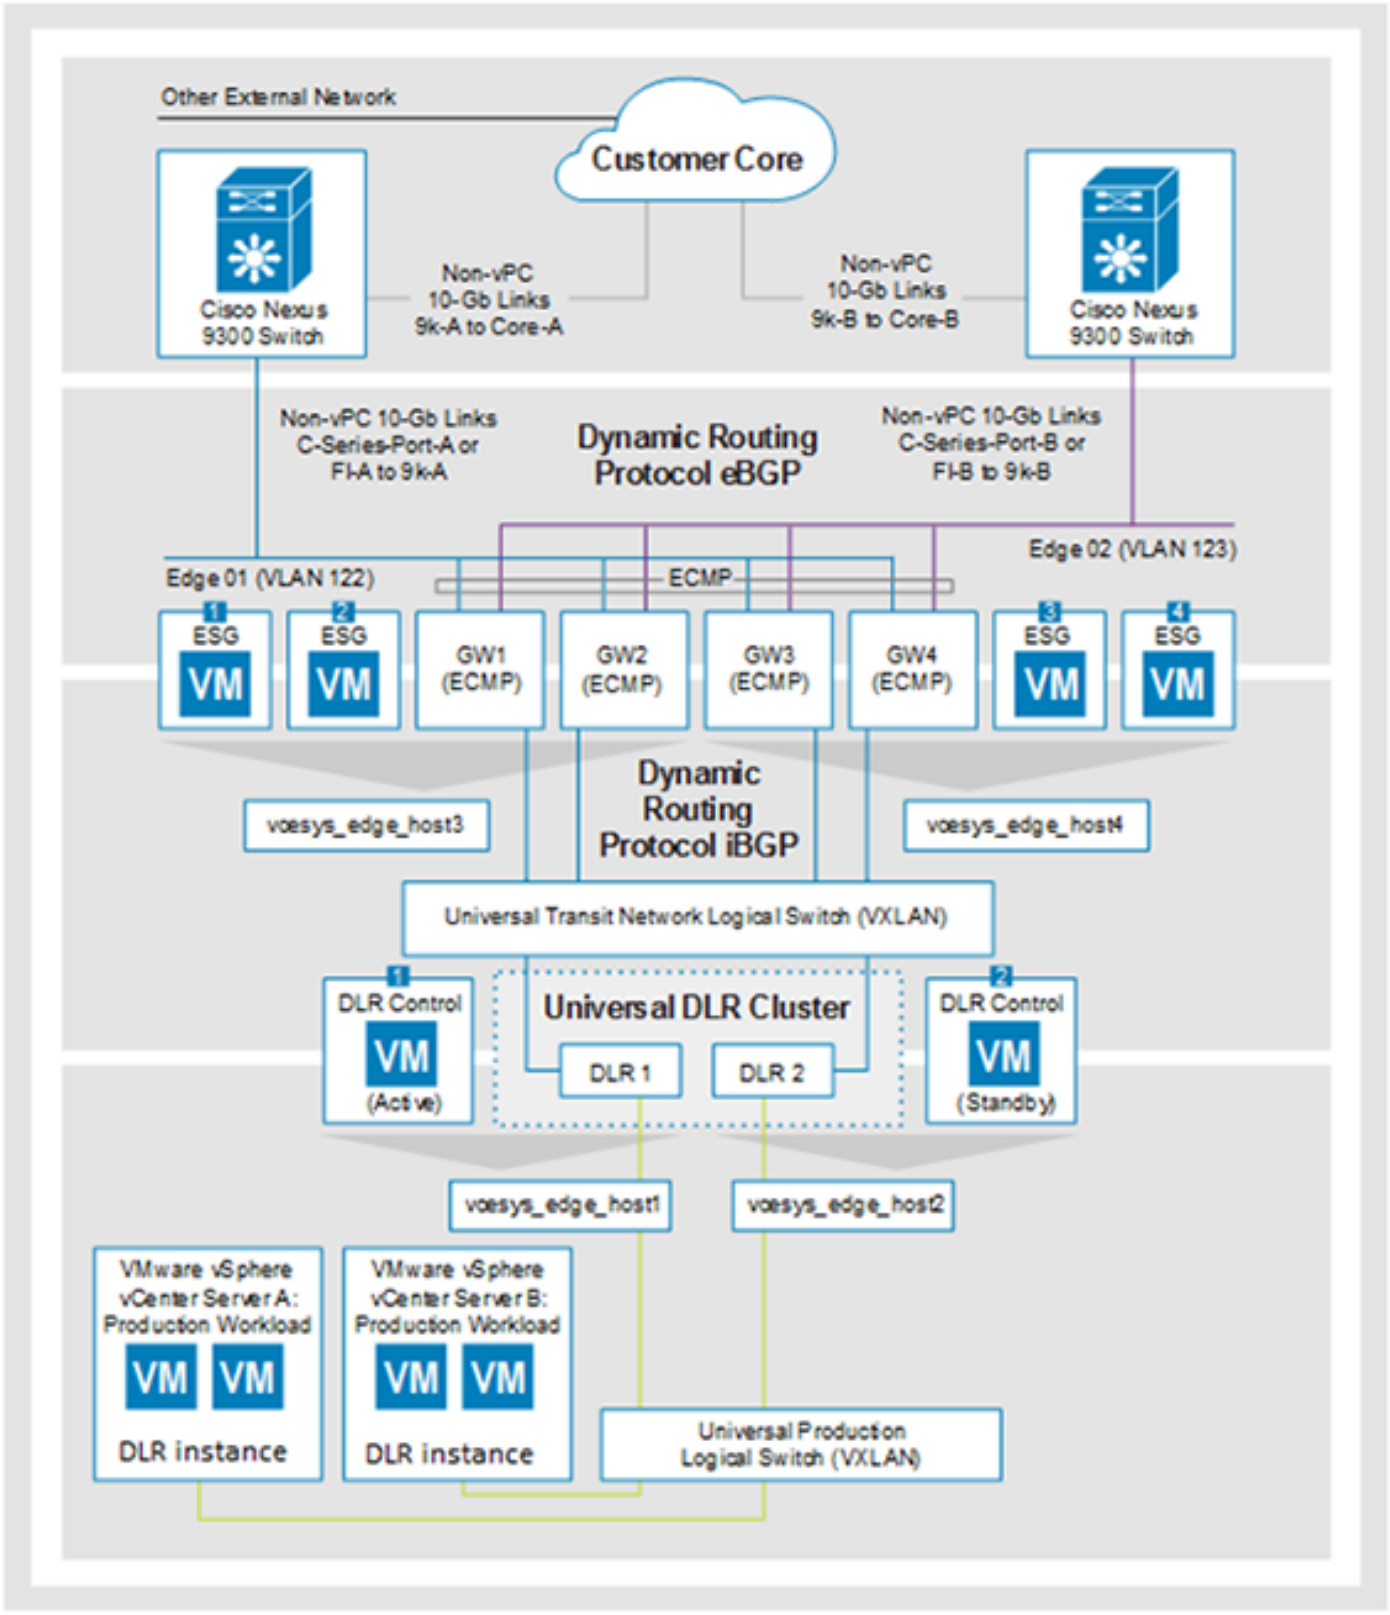
Other External Network
Customer Core
Cisco Nexus 9300 Switch
Non-vPC 10-Gb Links 9k-A to Core-A
Non-vPC 10-Gb Links 9k-B to Core-B
Cisco Nexus 9300 Switch
Non-vPC 10-Gb Links C-Series-Port-A or FI-A to 9k-A
Dynamic Routing Protocol eBGP
Non-vPC 10-Gb Links C-Series-Port-B or FI-B to 9k-B
Edge 01 (VLAN 122)
ECMP
Edge 02 (VLAN 123)
ESG VM
ESG VM
GW1 (ECMP)
GW2 (ECMP)
GW3 (ECMP)
GW4 (ECMP)
ESG VM
ESG VM
Dynamic Routing Protocol iBGP
voesys_edge_host3
voesys_edge_host4
Universal Transit Network Logical Switch (VXLAN)
DLR Control VM (Active)
Universal DLR Cluster
DLR 1
DLR 2
DLR Control VM (Standby)
voesys_edge_host1
voesys_edge_host2
VMware vSphere vCenter Server A: Production Workload
VM VM
DLR instance
VMware vSphere vCenter Server B: Production Workload
VM VM
DLR instance
Universal Production Logical Switch (VXLAN)

2

VMware NSX management cluster

The VMware NSX management cluster differs between AMP-VX and AMP-3S.

In AMP-3S, the management cluster consists of the management and control planes for VMware NSX. The VMware NSX manager handles the management plane and the VMware NSX controllers handle the control plane.

In AMP-VX, the management cluster consists of the management plane and holds the NSX manager components and resides in the management VMware vCenter Server. Controllers reside in the edge cluster and the edge and compute clusters reside in the production VMware vCenter Server.

Management cluster components

Management cluster components differ between the AMP-VX and AMP-3S.

The following table shows the VMware NSX management components for a VxBlock System 1000 with AMP-3S:

Management component	Description
VMware NSX management appliance	The VMware NSX management appliance is paired with a dedicated VMware vCenter Server. Use the web client to configure the NSX manager. A single domain can span multiple VMware vCenter Servers running VMware vSphere 6.x. One primary VMware NSX manager and up to seven secondary VMware NSX managers are supported in a single domain.
VMware NSX controllers	Three VMware NSX controllers are joined in a cluster to provide failover and scalability. VMware NSX controllers are deployed in hybrid mode to support both small and large virtual extensible LAN (VXLAN) environments. In a cross-vCenter Server configuration, the controllers are deployed only on the primary VMware NSX manager. The controller metadata is replicated across the secondary VMware NSX managers to allow synchronization of the control plane across the VMware vCenter Servers.
Universal components	The universal components consist of the universal transport zone, universal logical switches, and universal DLR. These components are created on the primary VMware NSX manager and not on the secondary VMware NSX manager.

The following table shows the VMware NSX management components for the VxBlock System 1000 with AMP-VX:

Management component	Description
VMware NSX management appliance	The VMware NSX management appliance is paired with a dedicated VMware management vCenter Server. Use the web client to configure the NSX manager. A single domain can span multiple VMware vCenter Servers running VMware vSphere 6.x. One primary VMware NSX manager and up to seven secondary VMware NSX managers are supported in a single domain.

Management cluster specifications

Management cluster specifications are for the VMware NSX manager and VMware NSX controller.

The following table shows the NSX manager and controller specifications for the **VxBlock System 1000 with AMP-3S**:

Specification	VMware NSX manager	VMware NSX controllers
Quantity	1 to 8 VMs	3 VMs

Specification	VMware NSX manager	VMware NSX controllers
	(With VMware NSX 6.4, more than one VMware NSX manager can be deployed in a single domain.)	
Location	Management cluster	Management cluster (Anti-affinity rules are enabled on AMP cluster.)
Hardware	AMP-3S	AMP-3S
Size	4 vCPU (8 vCPU if there are more than 256 hypervisors) 16 Gb RAM (24 Gb RAM if there are more than 256 hypervisors) 60 Gb disk	4 vCPU (2048 MHz reservation) 4 Gb RAM 28 Gb disk
Edge cluster	The DLR and ESG reside in the edge cluster.	NA
Management cluster	NSX manager and controllers reside in the management cluster.	NA
Network	vcesys_esx_mgmt (105) or amx_esx_mgmt (205), depending on which is available on the AMP platform	vcesys_esx_mgmt (105) or amx_esx_mgmt (205), depending on which is available on the AMP platform
Availability	VMware High Availability	VMware High Availability
Distribution	OVA	OVA

The following table shows the NSX manager and controller specifications for the **VxBlock System 1000 with AMP-VX**:

Specification	VMware NSX manager
Quantity	1 to 8 VMs (With VMware NSX 6.4, more than one VMware NSX manager can be deployed in a single domain.)
Location	Management cluster
Hardware	AMP-VX
Size (per NSX manager instance)	4 vCPU (8 vCPU if there are more than 256 hypervisors) 16 Gb RAM (24 Gb RAM if there are more than 256 hypervisors) 60Gb disk
Edge cluster	DLR, ESGs, and controllers reside in the edge cluster.
Management cluster	NSX managers reside in the management cluster.
Network	vcesys_esx_mgmt (105) or amx_esx_mgmt (205), depending on which is available on the AMP platform
Availability	VMware High Availability
Distribution	OVA

Management cluster hardware requirements

VxBlock System 1000 support VMware NSX virtual networking with AMP-VX or AMP-3S with a minimum of three servers. No other AMP type is supported with VMware NSX. The VMware NSX manager and the three VMware NSX controllers are dedicated to a VMware vSphere ESXi host for redundancy and scalability. No special cabling is required.

The following illustration shows the VMware NSX manager and the three VMware NSX controllers with a three-node AMP-3S:

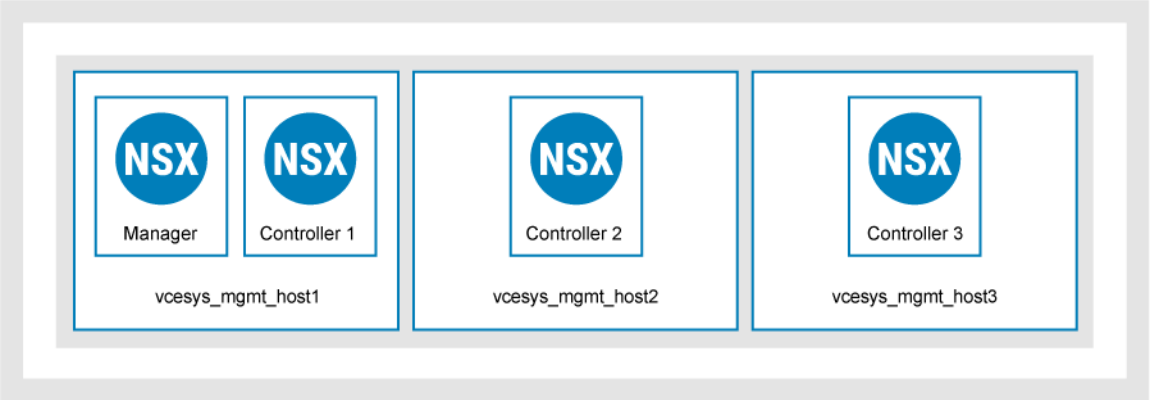

The following illustration shows the three VMware NSX managers for a cross-vCenter deployment with AMP-VX:

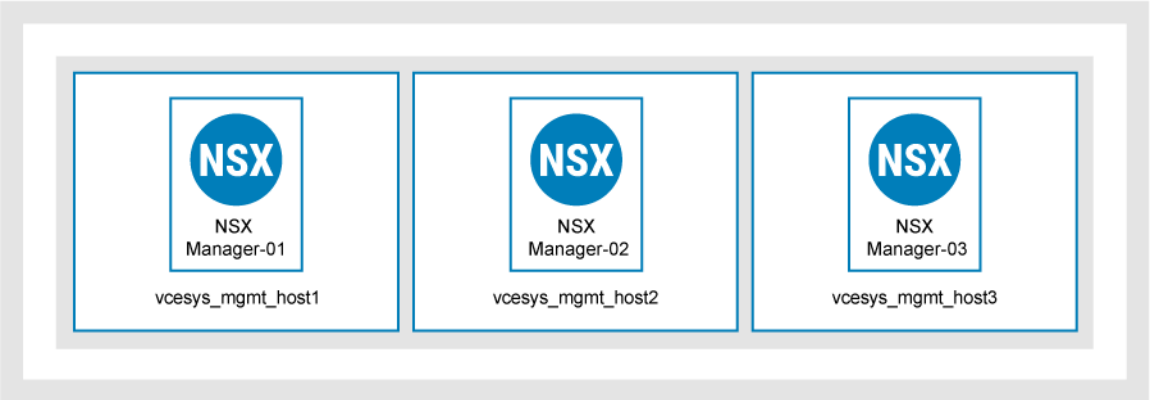

VMware vSphere supported versions

VMware NSX-V supports VMware vCenter and VMware ESXi versions 6.5, 6.7, and 7.0. The addition of VMware vSphere 7.0 requires VMware NSX-V version 6.4.8 or above. Customers that are upgrading to VMware vSphere 7.0 are required to upgrade their current version of VMware NSX-V to at least NSX-V 6.4.8 or the latest RCM level before upgrading to VMware vSphere 7.

VMware vSphere cluster requirements

The management cluster requires VMware High Availability (HA) and VMware vSphere Distributed Resource Scheduler (DRS) to provide VM protection against a VMware vSphere ESXi host failure and to balance VM workloads in the cluster.

The following table shows the rules applied to the DRS for AMP-3S:

Rule	VMware NSX manager	VMware NSX controllers
DRS affinity rules (same host)	Affinity rules are not applied to the management cluster, regardless of the numbers of VMware NSX managers.	Affinity rules are not applied to the management cluster, because each controller should not be on the same VMware vSphere ESXi host.
Anti-affinity rules (separate host)	Anti-affinity rules are applied to the management cluster if more than one VMware NSX manager exists.	Anti-affinity rules are applied to the management cluster for each controller on separate VMware vSphere ESXi hosts.

The following table shows the rules applied to the DRS for AMP-VX:

Rule	VMware NSX manager
DRS affinity rules (same host)	Affinity rules are not applied to the management cluster, regardless of the numbers of VMware NSX managers.
Anti-affinity rules (separate host)	Anti-affinity rules are applied to the management cluster if more than one VMware NSX manager exists.

Management cluster custom resource pool requirements

The VMware NSX management cluster does not require custom resource pools. However, for heavy workloads, create memory reservations for the VMware NSX manager.

Management cluster storage requirements

The management cluster does not require a specific disk layout other than the standard disk layout of the AMP-VX or AMP-3S.

For AMP-3S, the VMware vSphere ESXi hosts that are connected to the management cluster use the management storage array. The VMware NSX components, including NSX manager and controllers, are deployed across three separate data stores to protect against LUN corruption and to improve performance and resilience.

For AMP-VX, the VMware vSphere ESXi hosts that are connected to the management cluster use the vSAN cluster. VMware NSX manager is deployed in the management cluster and it uses the vSAN cluster as a data store to protect against LUN corruption and to improve performance and resilience. VMware NSX manager is handled by management VMware vCenter Server.

Edge cluster resources (DLR, perimeter gateways, and controllers) use the external storage array and are handled by the production VMware vCenter Server.

Management cluster networking requirements

The management cluster has no special network requirements for AMP-3S and AMP-VX.

For AMP-3S, the VMware NSX management traffic (VMware NSX managers) and control plane traffic (VMware NSX controllers) are on the same network segment as the VMware vSphere ESXi management traffic to improve performance.

For AMP-VX, the management VMware vCenter Server handles the VMware NSX management traffic. The edge cluster and compute cluster resources are on the same network segment and handled by the production VMware vCenter Server. AMP-VX provides the connectivity between the management VMware vCenter Server and production VMware vCenter Server as part of the default AMP-VX configuration.

3

VMware NSX edge cluster

The VMware NSX edge cluster connects to the physical network and provides routing and bridging. The edge cluster supports the Cisco UCS C-Series Rack Mount Servers.

Edge cluster components

VMware NSX edge cluster components differ between the AMP-3S and the AMP-VX.

The following table describes VMware NSX edge cluster components for AMP-VX. For AMP-3S, the same components apply, except the VMware NSX controller, which is deployed in the management cluster.

Component	Description
ESG	Provides connectivity to the physical network.
DLRs	Provide routing and bridging.
Universal components	The only universal object in the edge cluster is the DLR. The ESGs support multiple internal connections that connect to a local and universal DLR at the same time.
NSX controllers	Control virtual networks and overlay transport tunnels. NSX controllers are the central control point for all logical switches in a network. NSX controllers collect the information of all VMs, hosts, logical switches, and VXLANs.

The following illustration shows the edge cluster VMware NSX components for AMP-VX. The illustration shows the components that are assigned to each VMware vSphere ESXi host and the minimum of four ESGs.

For AMP-3S, the same edge cluster components apply, except the VMware NSX controller, which is deployed in the management cluster.

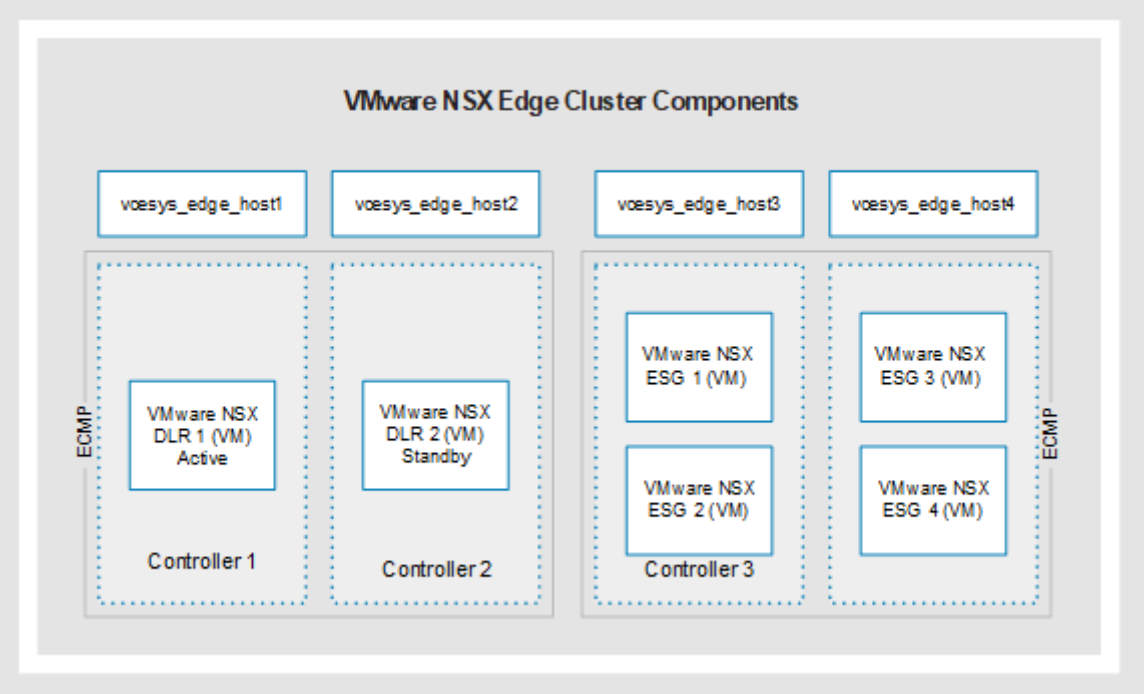

Edge cluster specifications

The specifications for the VMware NSX ESGs and VMware NSX DLR differ for AMP-3S and AMP-VX.

The following table lists the specifications for the VMware NSX ESGs and VMware NSX DLR for AMP-VX. For AMP-3S, the same edge cluster components apply, except the VMware NSX controller, which is deployed in the management cluster.

Specification	VMware NSX ESG	VMware NSX controllers	VMware NSX DLR
Quantity	4 to 6 VMs Active/active with ECMP	3 VMs	2 VMs Active/active with ECMP
Location	Edge cluster	Edge cluster	Edge cluster
Hardware	Cisco UCS C-Series Rack Mount Servers	Cisco UCS C-Series Rack Mount Servers	Cisco UCS C-Series Rack Mount Servers
Size	2 vCPU 1 Gb RAM 1.1 Gb disk	4 vCPU (2048 MHz reservation) 4 Gb RAM 20 Gb disk	1 vCPU 512 Mb RAM 1.1 Gb disk
Network	Two external interfaces connect to north-south physical switches. One internal interface connects to the local DLR and one internal interface connects to the universal DLR (if the universal DLR is deployed).	vcesys_esx_mgmt (105) or amx_esx_mgmt (205), depending on which is available on the AMP platform	One uplink interface connects to the ESG (for local DLRs and universal DLRs)
Availability	VMware HA	VMware HA	VMware HA
Distribution	Anti-affinity rules are enabled to separate the ESG VM pairs in the edge cluster. Affinity rules maintain two ESGs per host, except for the first two hosts.	Anti-affinity rules enabled on edge cluster.	Anti-affinity rules are enabled to keep both DLR VMs separate in the edge cluster. Anti-affinity rules are enabled across the first two hosts in the edge cluster.

Edge cluster hardware requirements: Cisco UCS C-Series Rack Mount Servers

The edge cluster uses four Cisco UCS C-Series Rack Mount Servers, regardless of the number of ESGs.

Each edge server uses two dual-port Intel NIC cards to support full 10 Gb line rate speeds for the ESGs (VXLAN offloading). Two Intel cards are installed in each edge server.

The following table describes the hardware specifications for the Cisco UCS C-Series M4 and M5 Rack Mount Servers that are used for the NSX edge cluster:

Component	Cisco M4 based edge node	Cisco M5 based edge node	Notes
CPU	Intel Xeon E5-2660 v3 (Haswell) 2.6Ghz, 10C, 25 MB Cache	Intel Xeon 5115 (Skylake) 2.4Ghz, 10C, 13.75 MB Cache Intel Xeon 5218 (Cascade Lake)	Intel Xeon 5115 is similar to the Intel Xeon E5-2600 v3, because it has the same number of cores and has similar clock speed and cost.

Component	Cisco M4 based edge node	Cisco M5 based edge node	Notes
		2.3Ghz, 16C, 22 MB Cache	Intel Xeon 5218 allows capacity to support future releases with new capabilities.
Memory	128 GB (8 x 16 GB DDR4-2400)	96 GB (6 x 16 GB DDR4-2666)	The workload on a node does not exceed 48 GB. The memory allows for expansion and future versions of code.
Network Interface Card (NIC)	Intel X520 Dual-port 10 GB adapter	Intel XXV710-DA2 Dual-port 25 GB PCIe adapter	VxBlock 1000 does not support 25 GB connections to the data plane. However, 25 GB capabilities in the NIC allow for future support. The XXV710 cards also support GENEVE offload and DPDK, which enables flexibility in future releases.
VIC (Virtual Interface Card)	Cisco VIC 1227 (2x 10 GB SFP+ mLOM)	Cisco VIC 1387 (2x 40 GB QSFP mLOM) OR Cisco VIC 1457 (4x 10/25 GB SFP28 mLOM)	The VIC card provides management and storage capabilities to the node. The second-generation VIC cards, such as the VIC 1227, are not supported on the M5 compute platform. The VIC 1457 provides a longer service life for the node. It also costs less than the VIC1387. FRUs and expansions can use a node with the VIC 1387 mLOM. The VIC 1387 ensures RCM compatibility for existing deployments. If you are not running UCS 4.x, you can add NSX edge nodes without doing a full RCM upgrade. QSA adapters convert to the SFP+ form factor as necessary.

Edge cluster VMware vSphere requirements

The edge cluster requires VMware HA and VMware vSphere DRS to provide VM protection against a VMware vSphere ESXi host failure and to balance VM workloads in the cluster.

The following table provides the DRS rules for AMP-3S:

Rule	ESGs	DLRs	VMware NSX controllers
Affinity rules (same host)	Affinity rules are applied to the edge cluster to allow each pair of VMware NSX ESGs to be assigned to its own VMware vSphere ESXi host.	For local DLRs, DRS affinity rules do not need to be configured, because only two DLRs exist in HA mode. For local and universal DLRs, DRS affinity rules are applied to the edge cluster to allow	Affinity rules are applied to the management cluster, because each controller should be on different VMware vSphere ESXi hosts on the AMP servers.

Rule	ESGs	DLRs	VMware NSX controllers
		one local and one universal DLR as a pair to be assigned to its own VMware vSphere ESXi host.	
Anti-affinity rules (separate host)	Anti-affinity rules are applied to the edge cluster to each pair of VMware NSX ESGs to not cross the VMware vSphere ESXi hosts.	For local DLRs, DRS anti-affinity rules are applied to the edge cluster to each VMware NSX DLR to not cross the VMware vSphere ESXi hosts. For local and universal DLRs, DRS anti-affinity rules are applied to the edge cluster to each pair of VMware NSX DLRs to not cross the VMware vSphere ESXi hosts.	Anti-affinity rules are applied to the edge cluster for each controller on separate VMware vSphere ESXi hosts.

The following table provides the DRS rules for AMP-VX:

Rule	VMware NSX controllers
Affinity rules (same host)	Affinity rules are not applied to the edge cluster, because each controller should not be on the same VMware vSphere ESXi host.
Anti-affinity rules (separate host)	Anti-affinity rules are applied to the edge cluster for each controller on separate VMware vSphere ESXi hosts.

Edge cluster custom resource pool

The edge cluster does not require custom resource pools.

Edge cluster storage requirements

Edge cluster storage requirements differ between AMP-3S and AMP-VX.

The edge cluster has the following storage requirements for AMP-3S and AMP-VX:

- Data stores: The VMware vSphere ESXi hosts that are connected to the edge cluster use the storage arrays. The VMware NSX components, including the ESGs and DLRs, are deployed across two separate data stores to protect against LUN corruption and improve performance and resilience.
- Disk layout: No specific disk layout is necessary. VMware NSX supports the standard disk layout of the storage arrays.

Edge cluster networking requirements

The edge cluster has internal and external networking requirements.

The following table describes the network connectivity for the VMware NSX edge components:

Connectivity type	Description
External	North-south connectivity exists between the ESGs and the Cisco Nexus 9300 Series Switches. Each ESG uses two separate uplink interfaces that connect to an uplink edge distributed port group independently on the edge VMware VDS.
Internal and uplink	The internal connectivity between the ESG and the DLR is as follows:

Connectivity type	Description
	• Each ESG uses one internal interface to connect to the internal VXLAN local transit logical switch to reach the local DLR. If NSX uses the cross-vCenter Server option, a second internal interface connects to the internal VXLAN universal transit logical switch to reach the universal DLR. (By default, the local and universal DLRs are deployed only in the primary VMware vCenter Server.) • The local DLR uses an uplink interface to connect to the VXLAN local transit logical switch to reach the ESG. If NSX uses the cross-vSphere vCenter Server option, an uplink interface connects the universal DLR to the uplink VXLAN universal transit logical switch to reach the ESG.
Dynamic routing and Layer 3 termination	The ESGs use Equal Cost Multi-Pathing (ECMP) to increase bandwidth by load balancing traffic across equal cost multiple paths and to provide fault tolerance for failed paths. The ESGs use eBGP to peer with the ToR Cisco Nexus 9300 Series switches. On the Cisco Nexus 9300 Series switches, the two edge VLANs have the following switch virtual interfaces (SVIs): • Edge01 SVI on switch A • Edge02 SVI on switch B The DLRs use iBGP to peer with the ESGs. All other VLANs internal to the VxBlock System 1000 terminate at the ToR Cisco Nexus 9300 Series switches.
Layer 2 bridging	Layer 2 bridging is an optional configuration to support VXLAN-to-VLAN connectivity. Layer 2 bridging works with a local DLR. Layer 2 bridging is not supported with a universal DLR.

Edge cluster with local and universal objects

The NSX edge cluster uses the local DLR and the cross VMware vCenter Server option.

The following illustration shows the NSX edge cluster with the local DLR and the cross VMware vCenter Server option:

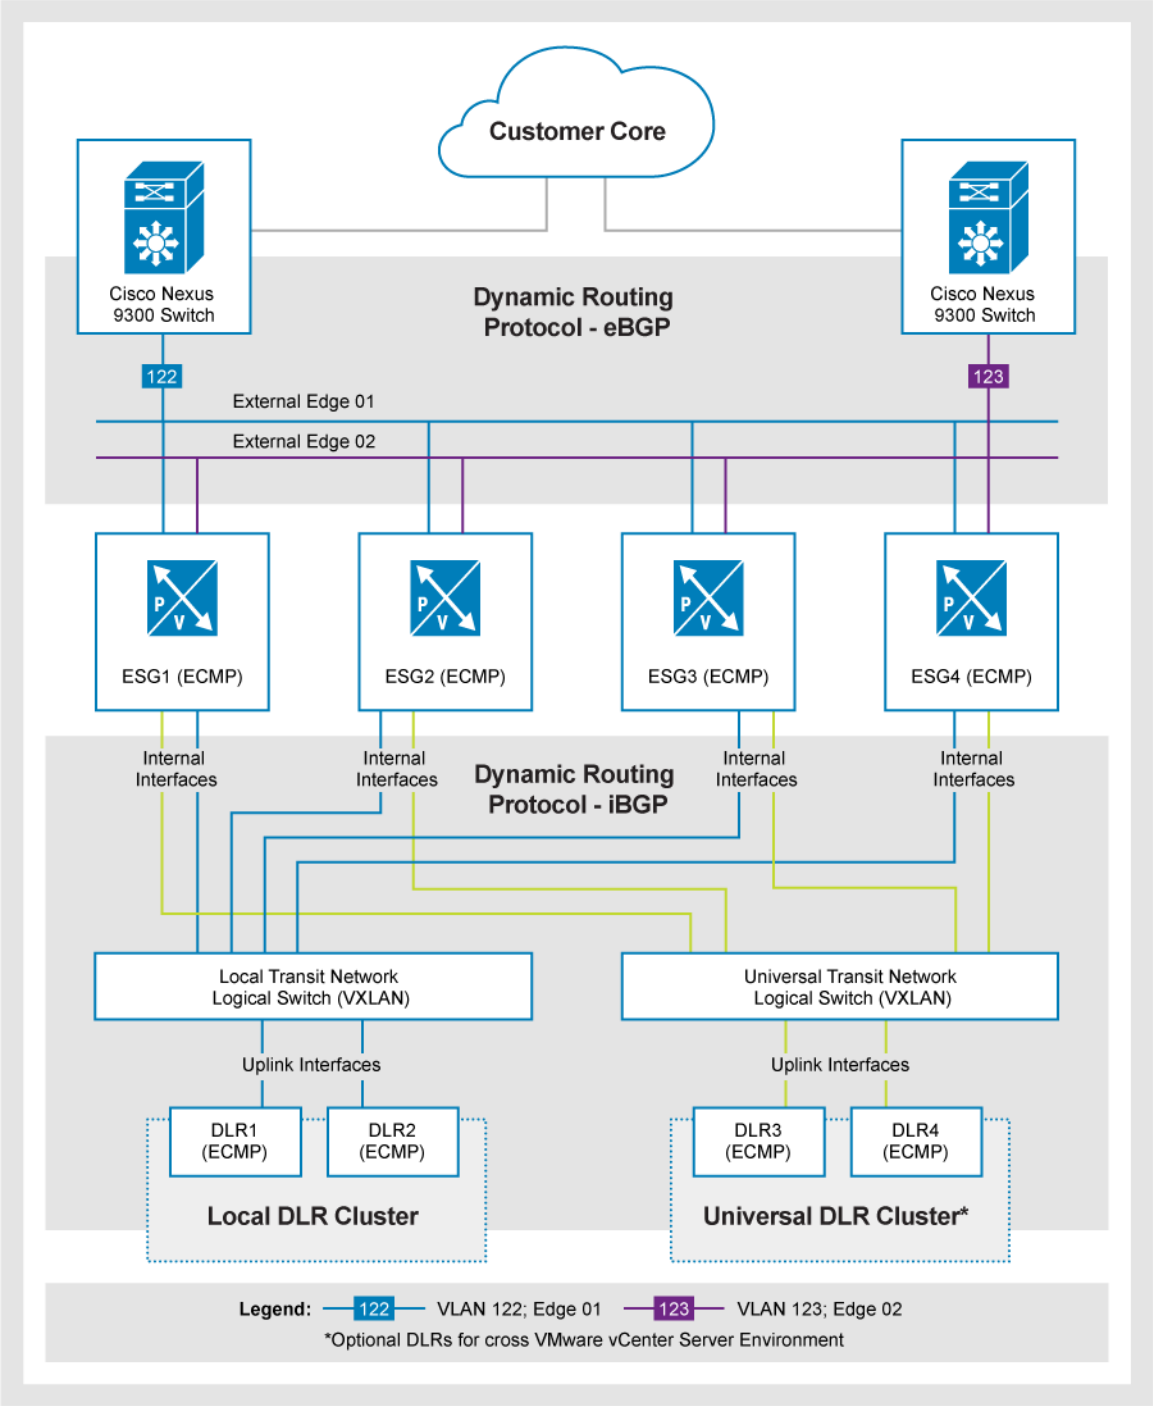

Edge cluster network requirements for the Cisco UCS Servers

The edge cluster has network requirements for the Cisco UCS Servers.

The following table describes the edge cluster network requirements for Cisco UCS C-Series Rack Mount Servers:

Component	Cisco UCS C-Series Rack Mount Servers
VLAN IDs	VMware NSX requires three VLAN/SVIs on the Cisco Nexus 9300 Series Switches: • Two external edge VLAN/SVIs for external traffic (on/off ramp)

Component	Cisco UCS C-Series Rack Mount Servers
	• One transport VLAN for VXLAN traffic The external traffic traverses north-south between the edge servers and the Cisco Nexus 9300 Series Switches. The transport VLAN is Layer 2 and does not have an SVI on the Cisco Nexus 9300 Series Switches. With Cisco UCS C-Series Rack Mount Servers, the external edge traffic VLAN IDs do not need to be created in the Cisco UCS Manager. However, because the compute blades pass VXLAN traffic, the VXLAN Transport VLAN ID must be added to the Cisco UCS Manager.
VXLAN Tunnel End Points (VTEPs)	The number of VTEPs deployed to each VMware vSphere ESXi host depends on the number of dvUplinks configured on the VMware VDS that has the transport distributed port group. Because there is more than one VTEP on each host, the Load Balance SRCID mode is enabled to load balance VXLAN traffic. Although link aggregation control protocol is supported on the Cisco UCS C-Series Rack Mount Servers, Load Balance SRCID mode ensures a consistent compute blade configuration. The number of ESGs determines the number of uplinks created on the edge VMware VDS. Configuration A (Typical for most installations) • 4 ESGs • 2 Intel NIC cards per host, which provides two dvUplinks • 2 VTEP/VMkernel distributed port groups per edge host Configuration B • 6 or 8 ESGs • 2 Intel NIC cards per host, which provides four dvUplinks • 4 VTEP/VMkernel distributed port groups per edge

Edge cluster VMware virtual network

The edge cluster Cisco C-Series Rack Mount Servers require two VMware Virtual Distributed Switches (VDS).

For Cisco UCS C-Series Rack Mount Servers, two VDS are created for the edge cluster:

- DVswitch01-Edge-Core manages the VMware vSphere ESXi management and VMware vSphere vMotion traffic types. VMware vSphere ESXi management is on a VMware VDS instead of a VMware vSphere Standard Switch to improve VMware NSX network performance.
- DVswitch02-Edge-NSX manages the external edge (Edge01 and Edge02), transport, and the optional Layer 2 bridging traffic types.
  - For four ESGs, VMware VDS requires two dvUplinks to connect to the Cisco Nexus 9300 Series Switches. This creates two VTEP VMkernel distributed port groups per edge host.
  - For more than four ESGs, VMware VDS uses four dvUplinks to connect to the Cisco Nexus 9300 Series Switches. This creates four VTEP VMkernel distributed port groups per edge host. Jumbo frames (9000) are enabled on the DVswitch-Edge-NSX switch and on the VXLAN transport distributed port group for VXLAN transport traffic.

Edge cluster network with Cisco UCS C-Series Rack Mount Servers and four ESGs

The following illustration shows the network layout for the hosts running on Cisco C-Series Rack Mount Servers in the edge cluster with four ESGs:

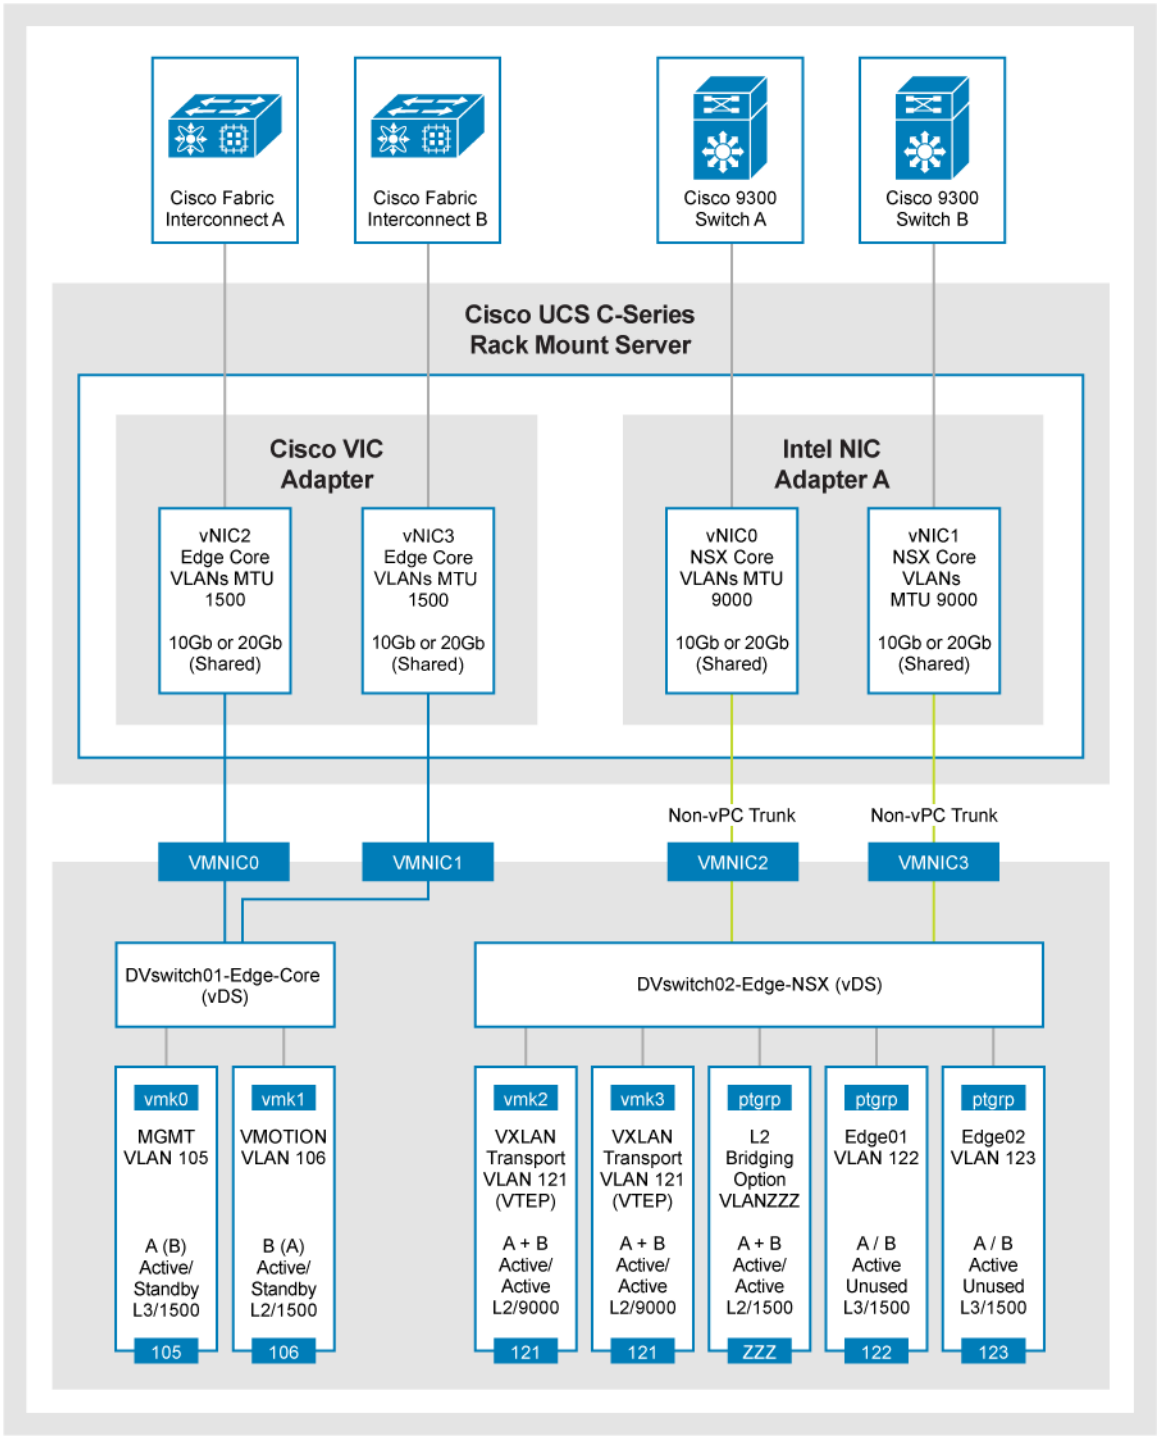

Edge cluster network with Cisco UCS C-Series Rack Mount Servers and more than four ESGs

The number of VTEPs is based on the number of uplinks on the VMware VDS, which is four VTEPs.

The following illustration shows the network layout for the hosts running on Cisco UCS C-Series Rack Mount Servers in the edge cluster with six ESGs:

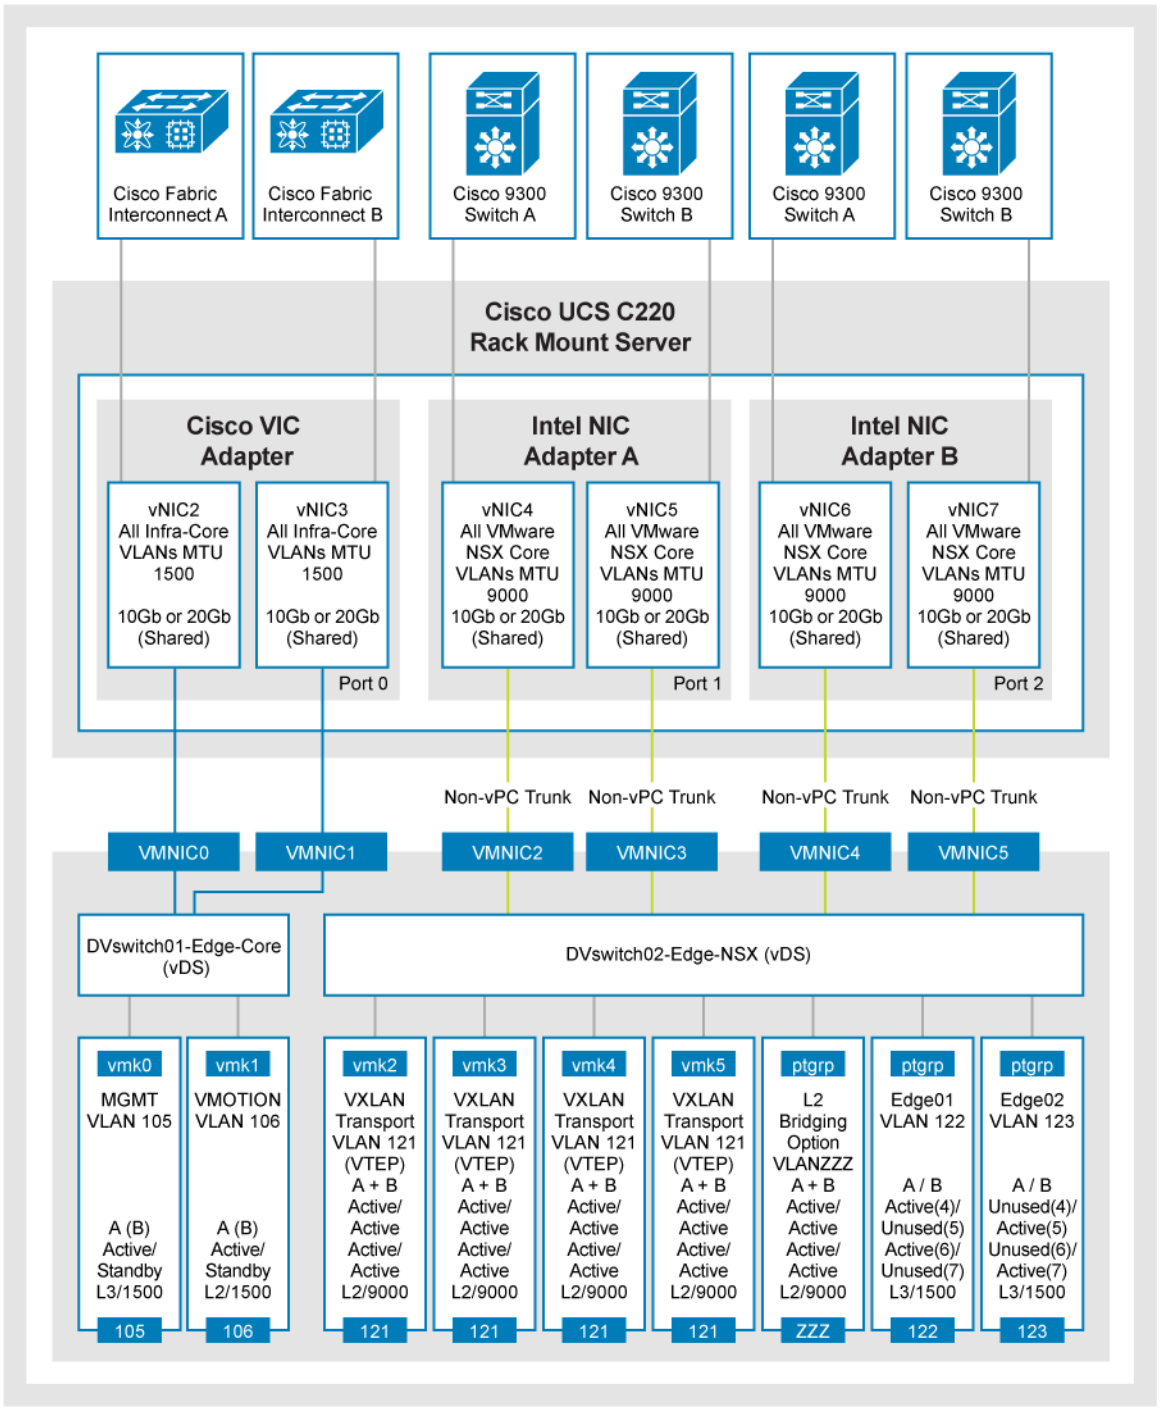

4

VMware NSX compute cluster

The VMware NSX compute cluster includes all the production VMs.

Compute cluster hardware requirements

For best performance, make sure each NSX compute cluster has network adapters that support VXLAN offload capabilities, such as the VIC 1300 or 1400 series of adapters.

Compute cluster VTEPs

The number of VTEPs deployed on each VMware vSphere ESXi server depends on the number of dvUplinks configured on the VMware VDS that has the transport distributed port group.

Because more than one VTEP is on each host, Load Balance SRCID mode is enabled to load balance VXLAN traffic.

The Cisco UCS B-Series Servers have two VTEPs deployed to each VMware vSphere ESXi host within the compute cluster.

LACP is not supported. Regardless of the number of ESG deployed, the number of dvUplinks is two.

Compute cluster VMware virtual network

More than one compute cluster can exist in the VMware vCenter Server.

A single VMware VDS spans across multiple compute clusters. However, additional VMware VDS can be deployed for a compute cluster or a set of compute clusters.

The single compute VDS manages the VMware vSphere vMotion, NFS, and VXLAN transport traffic types. The VXLAN transport NFS port groups and the VMware VDS are configured for jumbo frames (MTU 9000).

By default, the VMware vSphere ESXi management traffic resides on the VMware standard switch then moves to the VMware VDS.

Compute cluster network with Cisco UCS B-Series Blade Servers

The following illustration shows the VMware VM layout for the hosts in the compute cluster: