

Dell EMC VxBlock™ Systems for VMware NSX 6.3 and later Architecture Overview

Version 1.3

April 2019

Copyright © 2017- 2019 Dell Inc. All rights reserved.

Dell believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS-IS." DELL MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. USE, COPYING, AND DISTRIBUTION OF ANY DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA.

Dell EMC
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.DellEMC.com

CONTENTS

Introduction	5
Revision history	6
VMware NSX 6.3 and later on a VxBlock System	7
Introduction to VMware NSX 6.3 and later network virtualization.....	7
VMware vSphere cluster summary for VMware NSX.....	7
VMware NSX management cluster.....	8
VMware NSX Management cluster components.....	8
VMware NSX logical topology with local objects.....	9
Management cluster specifications.....	10
Hardware requirements.....	10
VMware vSphere cluster requirements.....	11
Custom resource pool requirements.....	11
Storage requirements.....	11
Networking requirements.....	11
VMware NSX edge cluster.....	11
Edge cluster components.....	11
Edge cluster specifications.....	12
Hardware requirements: Cisco UCS C-Series Rack Mount Servers...	13
Hardware requirements: Cisco UCS B-Series Blade Servers.....	15
VMware vSphere cluster.....	16
Custom resource pool.....	17
Storage requirements.....	17
Networking requirements.....	17
Logical topology of VMware NSX 6.3 and later local objects.....	18
Network requirements for the Cisco UCS Servers.....	18
VXLAN tunnel end point.....	20
VMware virtual network.....	20
Logical network for the Cisco UCS C-Series Rack Mount Servers using four ESGs.....	21
Logical network for the Cisco UCS C-Series Rack Mount Servers (Edge Cluster) using more than four ESGs.....	22
Logical network for the Cisco UCS B-Series Blade Servers (edge cluster).....	23
VMware NSX compute cluster.....	24
Hardware requirements.....	24
VXLAN tunnel end point.....	25
VMware virtual network for the compute cluster.....	25
Logical network for the Cisco UCS B-Series Blade Servers (compute cluster).....	25

Introduction

This document describes the high-level design of VMware NSX network virtualization technologies for standalone VxBlock Systems.

This document can be used as a reference and read in any order.

This document covers VMware NSX with VMware vSphere running on Cisco UCS C-Series Rack Mount servers and B-Series Blade Servers for the edge VMware vSphere cluster. Refer to the Release Certification Matrix for more information about supported hardware and software with VMware NSX.

The target audience for this document includes sales engineers, field consultants, and advanced services specialists who want to deploy a virtualized infrastructure using VMware NSX.

The [Glossary](#) provides related terms, definitions, and acronyms.

Revision history

Date	Document revision	Description of changes
April 2019	1.3	Updated to support Cisco Nexus C220 M5 servers at the NSX edge.
May 2018	1.2	Changed references to NSX 6.3 to NSX 6.3 and later.
March 2018	1.1	Updated the graphic in <i>Logical network for the Cisco UCS B-Series Blade Servers (edge cluster)</i> .
December 2017	1.0	Initial release

VMware NSX 6.3 and later on a VxBlock System

Introduction to VMware NSX 6.3 and later network virtualization

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

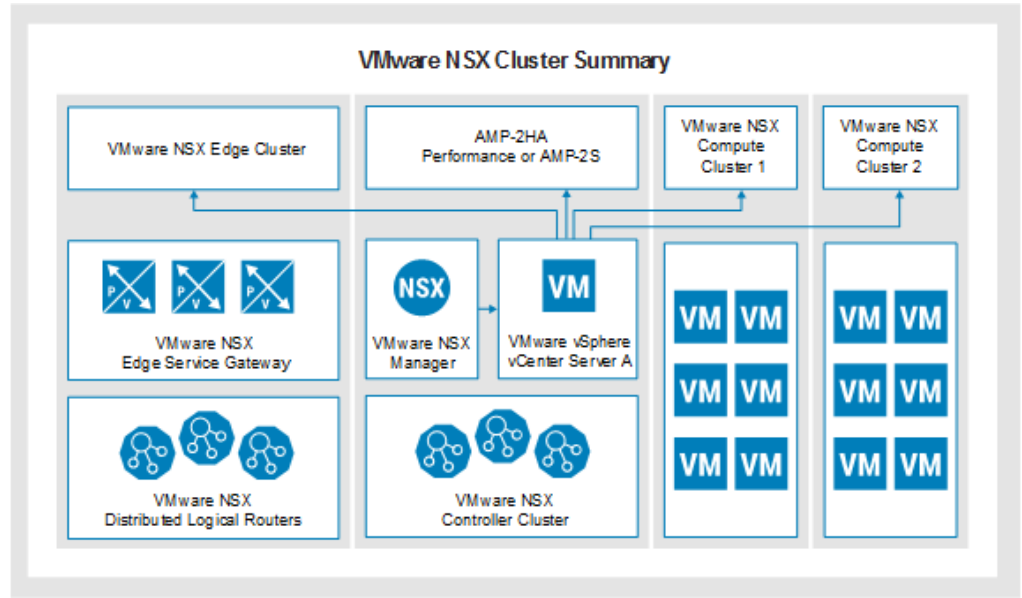
With VMware NSX, virtual networks are programmatically provisioned and managed independently of the underlying hardware. VMware NSX reproduces the entire network model in software, enabling a network topology to be created and provisioned in seconds. Network virtualization abstracts Layer 2 switching 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 following table describes the VMware vSphere clusters for VMware NSX:

Cluster	Description
Management	<p>Includes the VMware NSX manager appliance and three VMware NSX controllers that reside in the second generation of the Advanced Management Platform, specifically, AMP-2HA Performance or AMP-2S. 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:</p> <ul style="list-style-type: none"> • Allows the VMware NSX manager appliance and the VMware NSX controllers to be deployed into the management cluster. • 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 either the Cisco UCS C-Series Rack Mount servers or the B-Series Blade servers.
Edge	Includes the ESGs that provide external connectivity to the physical network and the DLRs that provide routing and bridging.
Compute	Includes the production VMs. There can be more than one compute cluster.

The following illustration provides the components and functions in each VMware NSX cluster:



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

VMware NSX management cluster

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.

VMware NSX Management cluster components

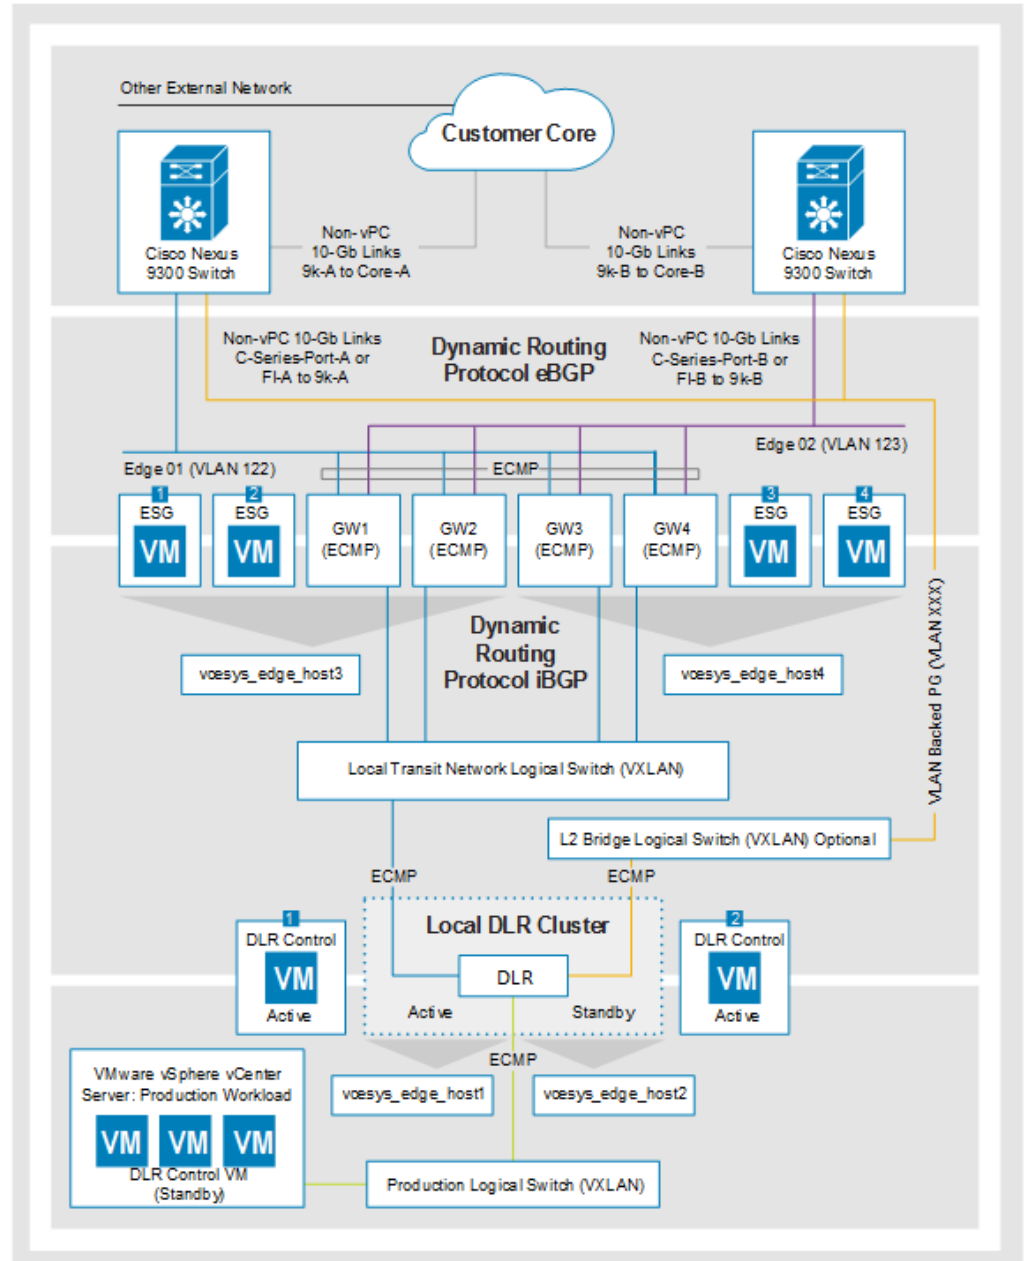
The VMware NSX management cluster on AMP-2HA Performance or AMP-2S consists of the VMware NSX manager appliance, VMware NSX controllers, and universal components.

Management component	Description
VMware NSX manager appliance	This central management component of VMware NSX 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 6.5. 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 hybrid mode to support both small and large virtual extensible LAN (VXLAN) environments.

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:



Management cluster specifications

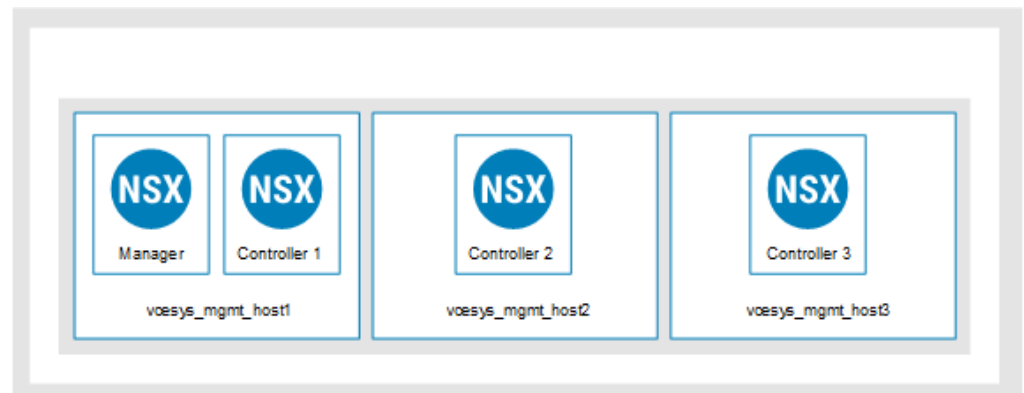
Use these specifications for the VMware NSX manager and VMware NSX controller.

Specification	VMware NSX manager	VMware NSX controllers
Quantity	1 to 8 VMs (With VMware NSX 6.3 and later, more than one VMware NSX manager can be deployed in a single domain.)	3 VMs
Location	Management cluster	Management cluster
Hardware	AMP-2HA Performance or AMP-2S with 3 servers min	AMP-2HA Performance or AMP-2S with 3 servers min
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 20 Gb disk
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

Hardware requirements

VxBlock Systems support VMware NSX virtual networking with AMP-2HA Performance or AMP-2S with three servers minimum. No other AMP type is supported with VMware NSX. AMP-2HA Performance or AMP-2S enables the VMware NSX manager and the three VMware NSX controllers to be 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:



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:

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.

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.

Storage requirements

The management cluster does not require a specific disk layout other than the standard disk layout of the AMP.

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

Networking requirements

No special network requirements besides AMP-2HA Performance or AMP-2S are required. 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.

VMware NSX edge cluster

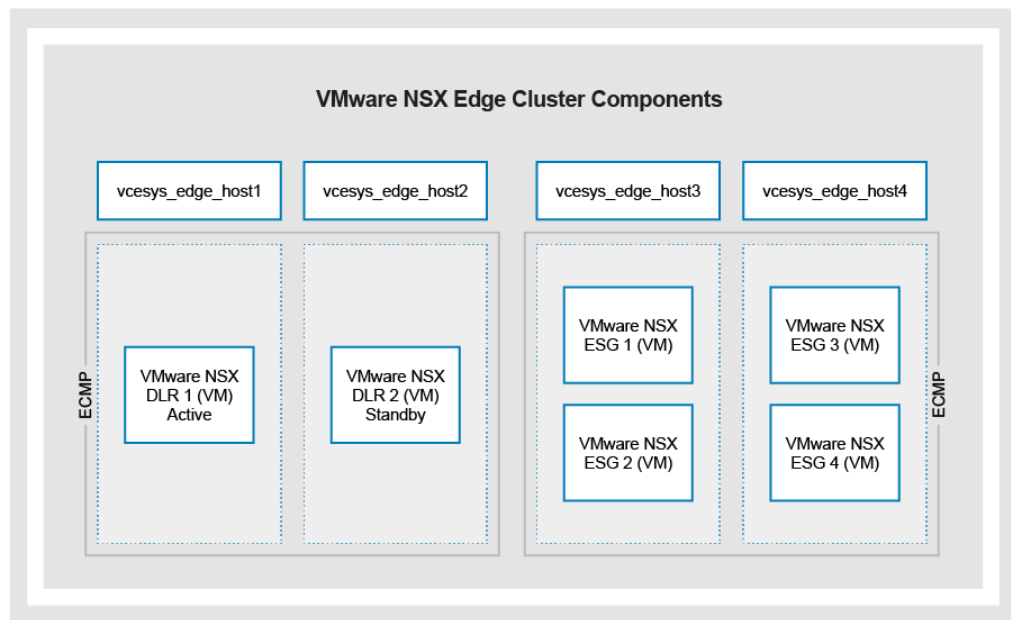
The VMware NSX edge cluster connects to the physical network and provides routing and bridging. The edge cluster supports either the Cisco UCS C-Series Rack Mount servers (recommended) or B-Series Blade Servers.

Edge cluster components

The following table explains VMware NSX edge cluster components:

Component	Description
ESG	Provides connectivity to the physical network.
DLRs	Provides routing and bridging.

The following illustration shows the VMware NSX components that belong to the edge cluster and which components are assigned to each VMware vSphere ESXi host. The illustration shows the minimum number of four ESGs.



Edge cluster specifications

The following table lists the specifications for the VMware NSX ESGs and VMware NSX DLR:

Specification	VMware NSX ESG	VMware NSX DLR
Quantity	4 VMs (4 or 6 are optional) Active/active with ECMP	2 VMs Active/active with ECMP
Location	Edge cluster	Edge cluster
Hardware	Cisco UCS C-Series Rack Mount Servers (Recommended) or Cisco UCS B-Series Blade Servers	Cisco UCS C-Series Rack Mount Servers (Recommended) or Cisco UCS B-Series Blade Servers
Size	4 vCPU 1 Gb RAM 2.15 Gb Disk	4 vCPU 512 Mb RAM 2.15 Gb Disk

Specification	VMware NSX ESG	VMware NSX DLR
Network	Two external interfaces for connections to north-south physical switches One internal interface for the connection to the local DLR	One uplink interface for connection to the ESG (for local DLRs)
Availability	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 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.

Hardware requirements: Cisco UCS C-Series Rack Mount Servers

The following table describes the hardware requirements for VMware NSX on the Cisco UCS C-Series Rack Mount Servers.

Component	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. However, the number of CPU sockets and Intel X520 or XXV710-DA2 cards and the amount of memory depends on the number of ESGs.
Performance enhancements	Each edge server uses a dual-port Intel NIC card to support full line rate speeds for the ESGs (VXLAN offloading). Two Intel NIC cards are installed in each edge server if more than four ESGs are deployed. Having a second Intel NIC card adds the additional bandwidth needed for the additional ESGs.
CPU and memory	<p>Configuration A (Typical for most installations)</p> <ul style="list-style-type: none"> • 4 ESGs • Single socket CPU per edge server • 64 Gb of memory per edge server <p>Configuration B</p> <ul style="list-style-type: none"> • 6 or 8 ESGs • Dual socket CPU per edge server • 128 Gb of memory per edge server
UCS-FEX	Each edge server includes a Cisco VIC card to take advantage of the Cisco SingleConnect technology. Each edge server connects directly to the Cisco Fabric Interconnect. Cisco UCS Manager can manage the Cisco UCS C-Series Rack Mount servers using service profiles. Only the VMware vSphere core management infrastructure uses the service profiles.
Physical cabling	Server to Cisco Nexus 9300 Series Switch connectivity:

Component	Requirements
	<p>Each edge server uses the Intel NIC card(s) to connect to the Cisco Nexus 9300 Series Switches. The number of ESGs determines if one or two Intel NIC cards are installed in each server. The physical cable configuration types are shown below:</p> <p>Configuration A (Typical for most installations)</p> <ul style="list-style-type: none"> • 4 ESGs • 1 Intel X520 card per edge server • Intel X520 port A connected to Cisco Nexus 9300 Series Switch A • Intel X520 port B connected to Cisco Nexus 9300 Series Switch B • 2 connections per edge server • 8 connections from all four edge servers to the Cisco Nexus 9300 Series Switches (four per switch) <p>Configuration B</p> <ul style="list-style-type: none"> • 6 or 8 ESGs • 2 Intel X520 or XXV710-DA2 cards per edge server • Intel NIC-A port A connected to Cisco Nexus 9300 Series Switch A • Intel NIC-A port B connected to Cisco Nexus 9300 Series Switch B • Intel NIC-B port A connected to Cisco Nexus 9300 Series Switch A • Intel NIC-B port B connected to Cisco Nexus 9300 Series Switch B • 4 connections per edge server. • 16 connections from all four edge servers to the Cisco Nexus 9300 Series Switches (eight per switch) <p>Server to Cisco UCS Fabric Interconnect Connectivity</p> <p>Each edge server uses the Cisco VIC 1227, 1387, or 1457 card to connect to the Cisco Fabric Interconnects. The physical cable configuration types are:</p> <ul style="list-style-type: none"> • VIC 1227, 1387, or 1457 port 1 connects to Cisco Fabric Interconnect A. • VIC 1227 or 1387 port 2 or 1457 port 3 connects to Cisco Fabric Interconnect B. • 2 connections per edge server • 8 connections from all four edge servers to the Cisco UCS Fabric Interconnects (four per Cisco Fabric Interconnect)

Hardware requirements: Cisco UCS B-Series Blade Servers

The following table describes the hardware requirements for VMware NSX on the Cisco UCS B-Series Blade Servers.

Component	Requirements
Cisco UCS B-Series Blade Servers	<p>The edge cluster uses four to six Cisco UCS B-Series Blade Servers, depending on the number of ESGs deployed, as specified below:</p> <p>Configuration A (Typical for most installations)</p> <ul style="list-style-type: none"> • 4 ESGs • 4 Cisco UCS B-Series Blade Servers <p>Configuration B</p> <ul style="list-style-type: none"> • 6 ESGs • 5 Cisco UCS B-Series Blade Servers <p>Configuration C</p> <ul style="list-style-type: none"> • 8 ESGs • 6 Cisco UCS B-Series Blade Servers
Performance enhancements	<p>Cisco does not support Cisco VIC cards that provide full line rate speeds. Each edge Cisco UCS B-Series Blade Server uses a Cisco VIC 1340 and 1380 card to support the VXLAN offloading capabilities, which provide VXLAN TCP checksum and segmentation to improve CPU performance. In addition, using two VIC cards allows the load balancing of traffic types to provide extra network performance. Using the VXLAN offloading capabilities provides up to half the full line rate at speeds of about 5 Gb of bandwidth per ESG. To provide full line speeds, use the Cisco UCS C-Series Rack Mount Server for the edge cluster.</p> <p>Note: The VXLAN offload feature is not enabled due to a driver limitation. In the meantime RSS is enabled on the card to provide the best performance.</p>
Physical cabling	<p>Each edge server uses the Cisco VIC 1340 and 1380 cards to virtually connect to the Cisco Fabric Interconnects. Physical 10 Gb cables connect the Cisco Fabric Interconnects to the Top of Rack (TOR) Cisco 9300 Series Switches in a non-virtual port channel used for external traffic. (A Cisco limitation exists where dynamic IP routing with peers on the virtual port channel (vPC) VLAN is not supported.) The number of ESGs determines the number of 10-Gb cable links, as follows:</p> <p>Configuration A (Typical for most installations)</p> <ul style="list-style-type: none"> • 4 ESGs • 4 10 Gb cables connect the Cisco Fabric Interconnect A to Cisco Nexus 9300 Series Switch A • 4 10 Gb cables connect the Cisco Fabric Interconnect B to Cisco Nexus 9300 Series Switch B • A total of eight connections exist from the Cisco Fabric Interconnects to the Cisco Nexus 9300 Series Switches (four per switch).

Component	Requirements
	<p>Configuration B</p> <ul style="list-style-type: none"> • 6 ESGs • 6 10 Gb cables connect the Cisco Fabric Interconnect A to Cisco Nexus 9300 Series Switch A • 6 10 Gb cables connect the Cisco Fabric Interconnect B to Cisco Nexus 9300 Series Switch B • A total of 12 connections exist from the Cisco Fabric Interconnects to the Cisco Nexus 9300 Series Switches (six per switch). <p>Configuration C</p> <ul style="list-style-type: none"> • 8 ESGs • 8 10 Gb cables connect the Cisco Fabric Interconnect A to Cisco Nexus 9300 Series Switch A • 8 10 Gb cables connect the Cisco Fabric Interconnect B to Cisco Nexus 9300 Series Switch B • A total of 16 connections exist from the Cisco Fabric Interconnects to the Cisco Nexus 9300 Series Switches (eight per switch).

VMware vSphere cluster

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:

Rule	ESGs	DLR	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 DLRs, DRS affinity rules are applied to the edge cluster to allow one local and one universal DLR as a pair to be assigned to its own VMware vSphere ESXi 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 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.	Anti-affinity rules are applied to the edge cluster for each controller on separate VMware vSphere ESXi hosts.

The edge cluster does not require custom resource pools.

Custom resource pool

The edge cluster does not require custom resource pools.

Storage requirements

The edge cluster has the following storage requirements:

- **Data stores:** The VMware vSphere ESXi hosts that are connected to the edge cluster use the VNX, XtremIO, Unity, or VMAX 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 VNX, XtremIO, Unity, or VMAX storage arrays.

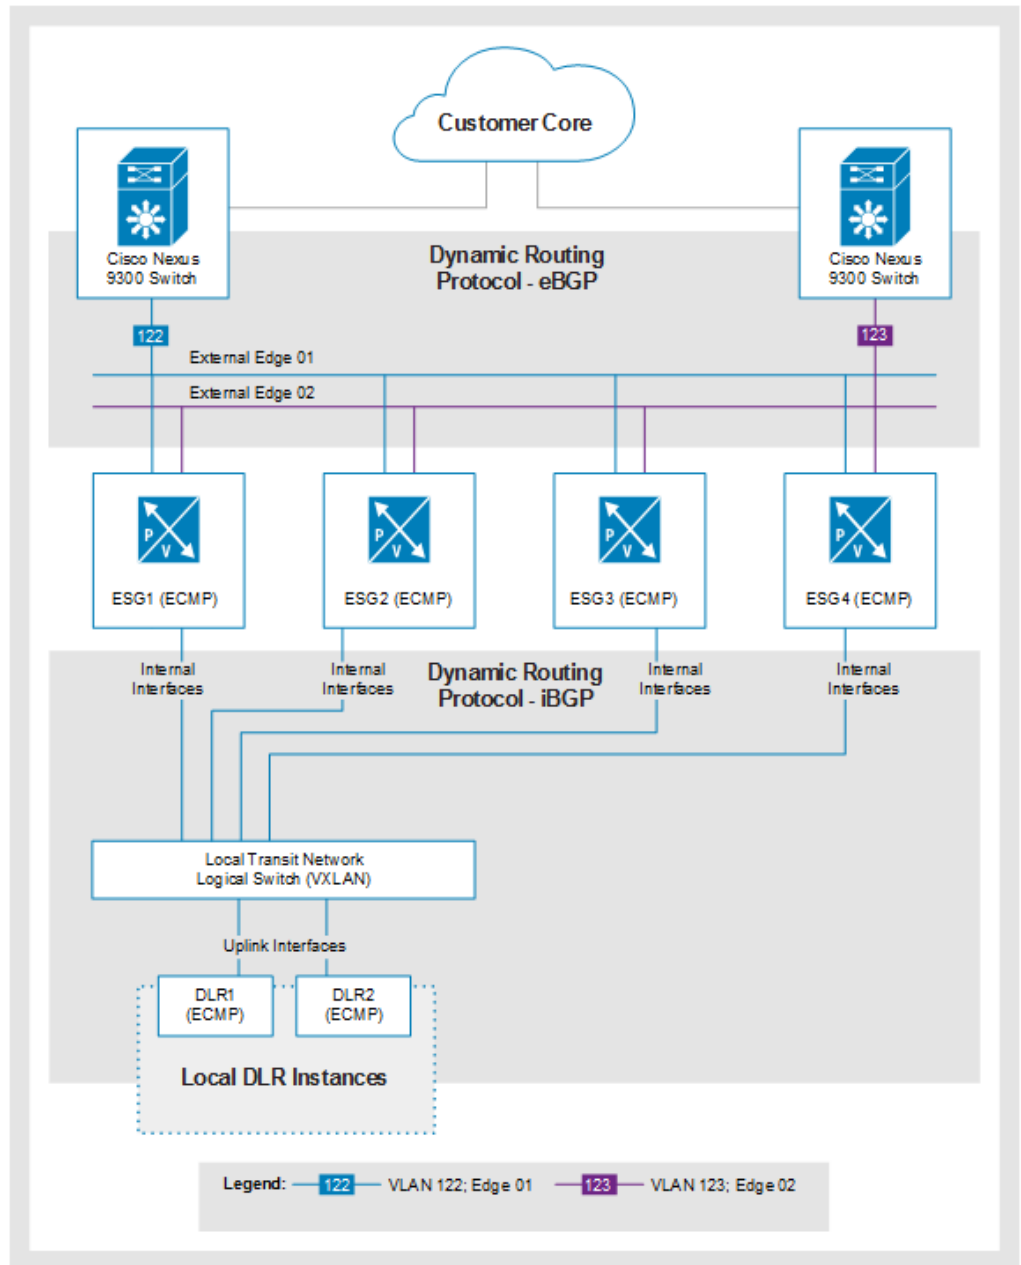
Networking requirements

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

Connectivity type	Description
External	<p>North-south connectivity exists between the ESGs and the Cisco Nexus 9300 Series Switches.</p> <p>Each ESG uses two separate uplink interfaces that connect to an uplink edge distributed port group independently on the edge VMware VDS.</p>
Internal and uplink	<p>The internal connectivity between the ESG and the DLR is as follows:</p> <ul style="list-style-type: none"> • Each ESG uses one internal interface to connect to the internal VXLAN local transit logical switch to reach the local DLR. • The local DLR uses an uplink interface to connect to the VXLAN local transit logical switch to reach the ESG.
Dynamic routing and Layer 3 termination	<p>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 switch virtual interfaces (SVI):</p> <ul style="list-style-type: none"> • Edge01 SVI on switch A • Edge02 SVI on switch B <p>The DLRs use iBGP to peer with the ESGs.</p> <p>All other VLANs internal to the VxBlock System terminate at the top of rack (ToR) Cisco Nexus 9300 Series switches.</p>
Layer 2 bridging	<p>Layer 2 bridging is an optional configuration to support VXLAN-to-VLAN connectivity. Layer 2 bridging works with a local DLR.</p>

Logical topology of VMware NSX 6.3 and later local objects

The NSX edge cluster uses the local DLR as shown in the following illustration:



Network requirements for the Cisco UCS Servers

The following table describes the network requirements for Cisco UCS C-Series Rack Mount Servers and Cisco UCS B-Series Blade Servers:

Component	Cisco UCS C-Series Rack Mount Servers	Cisco UCS B-Series Blade Servers
VLAN IDs	<p>VMware NSX requires three VLAN/SVIs on the Cisco Nexus 9300 Series Switches:</p> <ul style="list-style-type: none"> Two external edge VLAN/SVIs are used for external traffic (on/off ramp) One transport VLAN is used to pass VXLAN traffic. <p>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.</p>	
	<p>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.</p>	<p>With Cisco UCS B-Series Blade Servers, the external edge traffic and VXLAN transport traffic VLAN IDs must be created in the Cisco UCS Manager.</p>
VXLAN Tunnel End Points (VTEPs)	<p>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 created.</p> <p>Because there is more than one VTEP on each host, the Load Balance SRCID mode is enabled to load balance VXLAN traffic.</p>	
	<p>Although the link aggregation control protocol (LACP) is supported on the Cisco UCS C-Series Rack Mount Servers, Load Balance SRCID mode ensures a consistent compute blade configuration.</p> <p>The number of ESGs determines the number of uplinks (1 or 2 Intel NIC cards) created on the edge VMware VDS.</p> <p>Configuration A (Typical for most installations)</p> <ul style="list-style-type: none"> 4 ESGs 1 Intel NIC card per host, which provides two dvUplinks 2 VTEP/VMkernel distributed port groups per edge host <p>Configuration B</p> <ul style="list-style-type: none"> 6 or 8 ESGs 2 Intel NIC cards per host, which provides four dvUplinks 	<p>LACP is not supported for the Cisco B-Series Blade Servers.</p> <p>Regardless of the number of ESGs, two uplinks are created on the edge VMware VDS.</p> <p>Configuration A (Typical for most installations)</p> <ul style="list-style-type: none"> 4 ESGs 4 Cisco UCS B-Series Blade Servers <p>Configuration B</p> <ul style="list-style-type: none"> 6 ESGs 5 Cisco UCS B-Series Blade Servers <p>Configuration C</p> <ul style="list-style-type: none"> 8 ESGs 6 Cisco UCS B-Series Blade Servers

Component	Cisco UCS C-Series Rack Mount Servers	Cisco UCS B-Series Blade Servers
	<ul style="list-style-type: none"> 4 VTEP/VMkernel distributed port groups per edge 	

VXLAN tunnel end point

The number of VXLAN Tunnel End Points (VTEPs) deployed to each VMware vSphere ESXi 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.

For the Cisco UCS C-Series Rack Mount Servers, LACP is supported. However, Load Balance SRCID mode ensures a consistent compute blade configuration.

The number of ESGs determines the number of uplinks (1 or 2 Intel cards) created on the edge VMware VDS.

Configuration A Typical for most installations	Configuration B
<ul style="list-style-type: none"> Four ESGs One Intel X520 card per host, which provides two dvUplinks Two VTEP/VMkernel distributed port groups per edge host 	<ul style="list-style-type: none"> Six or eight ESGs Two Intel X520 or XXV710-DA2 cards per host, which provides four dvUplinks Four VTEP/VMkernel distributed port groups per edge host

VMware virtual network

For Cisco UCS C-Series Rack Mount Servers, two VDS are created for the edge cluster. For Cisco UCS B-Series Blade Servers, three VDS are created for the edge cluster.

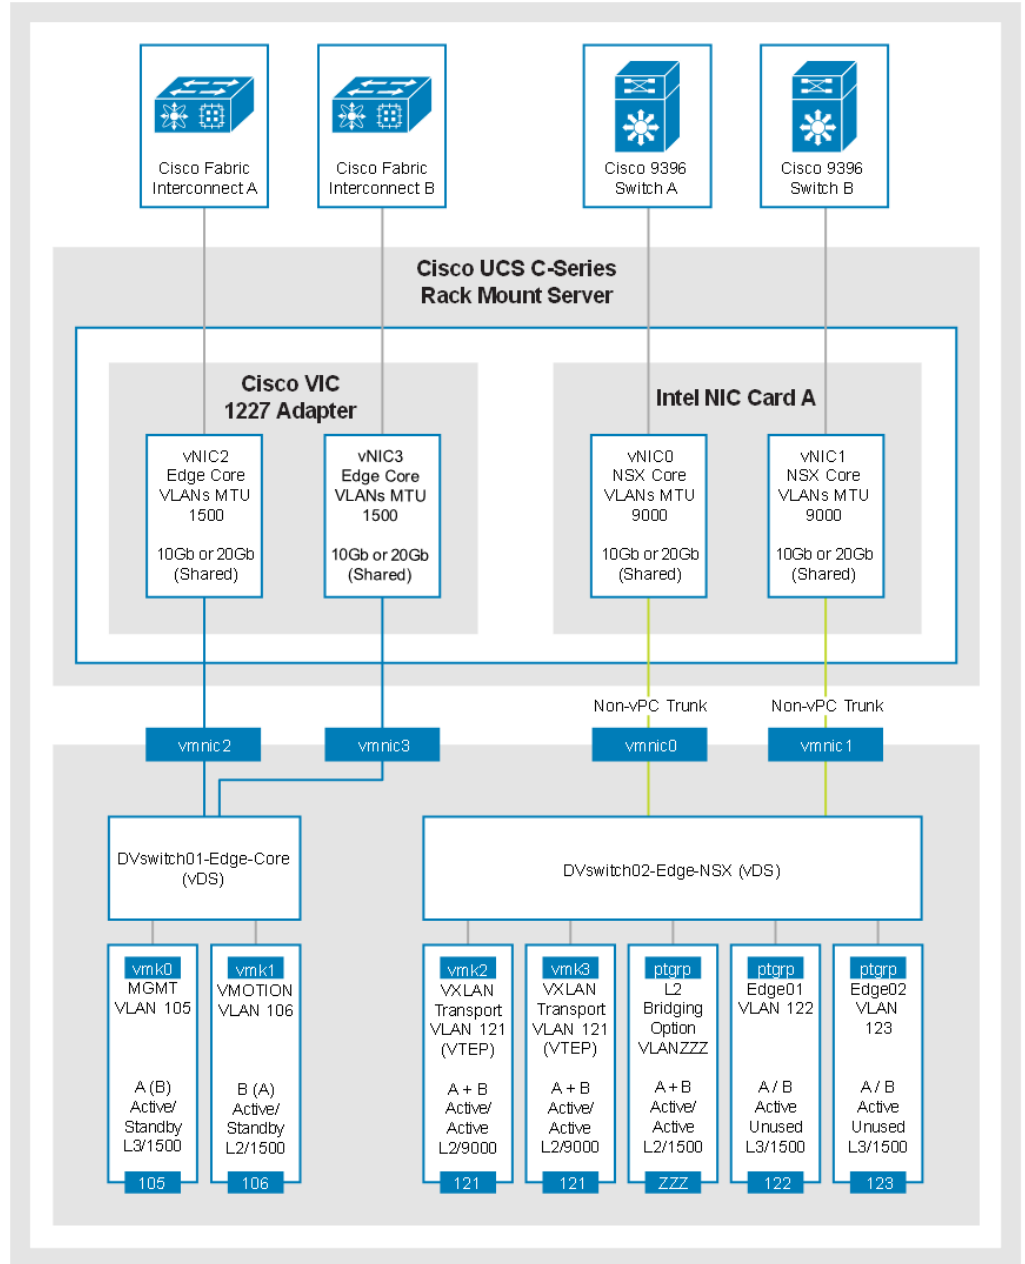
The following table describes each VDS:

Cisco UCS C-Series Rack Mount Servers	Cisco UCS B-Series Blade Servers
<p>DVswitch01-Edge-Core manages the VMware vSphere ESXi management and VMware vSphere vMotion traffic types.</p> <p>VMware vSphere ESXi Management is on a VMware VDS instead of a VMware VSS to improve VMware NSX network performance.</p>	<p>DVswitch01-Edge-Core manages the VMware vSphere ESXi management and VMware vSphere vMotion traffic types.</p> <p>VMware vSphere ESXi Management is on a VMware VDS instead of a VMware VSS to improve VMware NSX network performance.</p>
<p>DVswitch02-Edge-NSX manages the external edge (Edge01 and Edge02), transport, and the optional Layer 2 bridging traffic types.</p> <ul style="list-style-type: none"> 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. 	<p>DVswitch02-Edge-NSX-Core manages the VXLAN transport (east-west) and optional Layer 2 bridging traffic types.</p> <p>Jumbo frames (9000) are enabled on the DVswitch-NSX-Edge switch and on the VXLAN transport distributed port group for VXLAN transport traffic.</p>

Cisco UCS C-Series Rack Mount Servers	Cisco UCS B-Series Blade Servers
<ul style="list-style-type: none"> 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. <p>Jumbo frames (9000) are enabled on the DVswitch-Edge-NSX switch and on the VXLAN transport distributed port group for VXLAN transport traffic.</p>	
	<p>DVswitch03-NSX-External manages edge traffic for north-south connectivity.</p>

Logical network for the Cisco UCS C-Series Rack Mount Servers using four ESGs

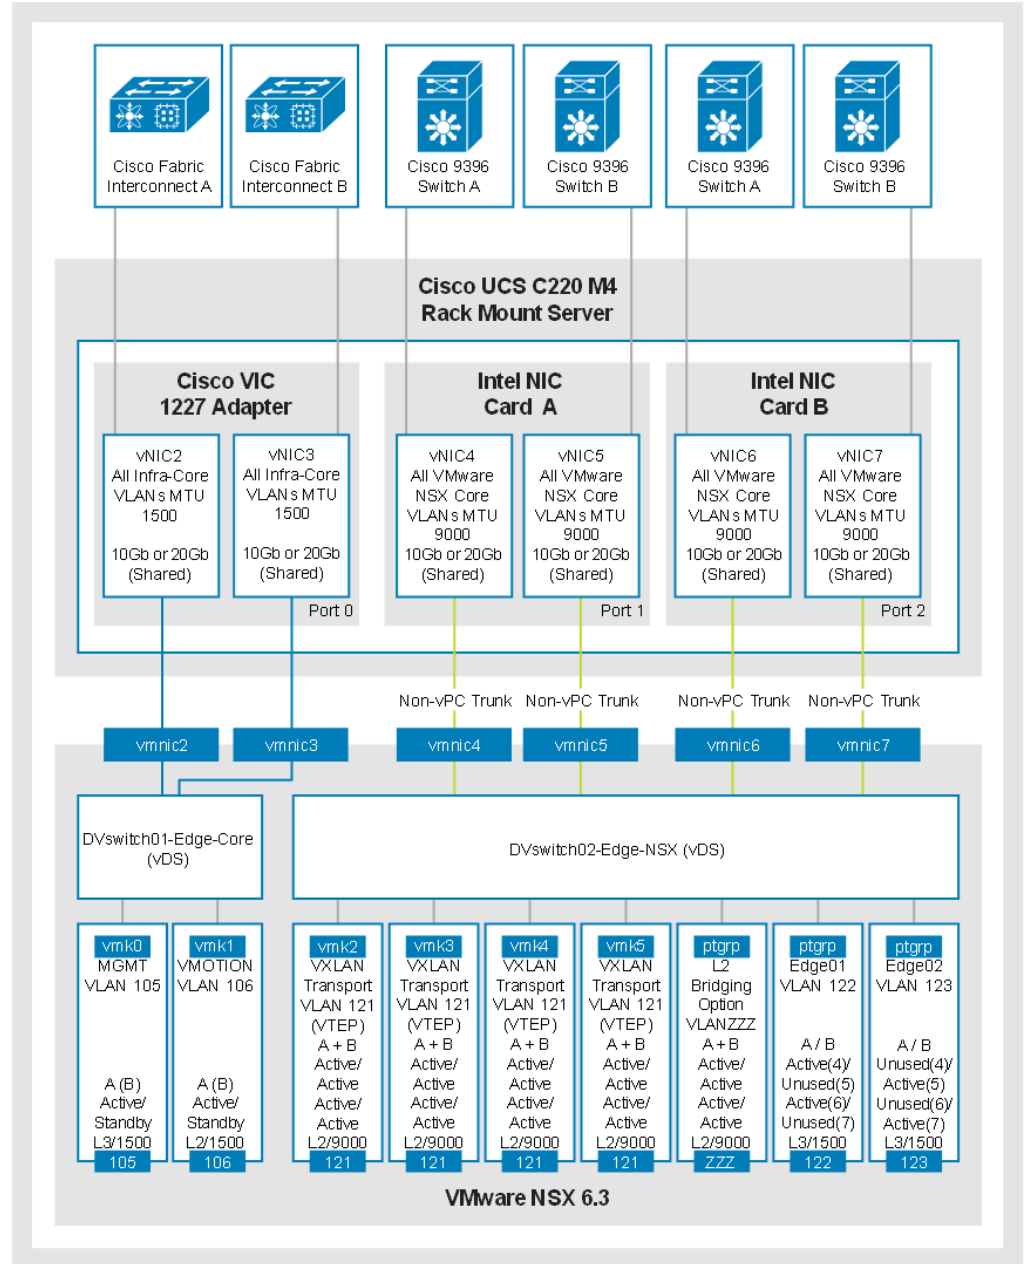
The following illustration shows the network layout for the hosts running on Cisco C-Series Rack Mount Servers in the edge cluster with four ESGs:



Logical network for the Cisco UCS C-Series Rack Mount Servers (Edge Cluster) using more than four ESGs

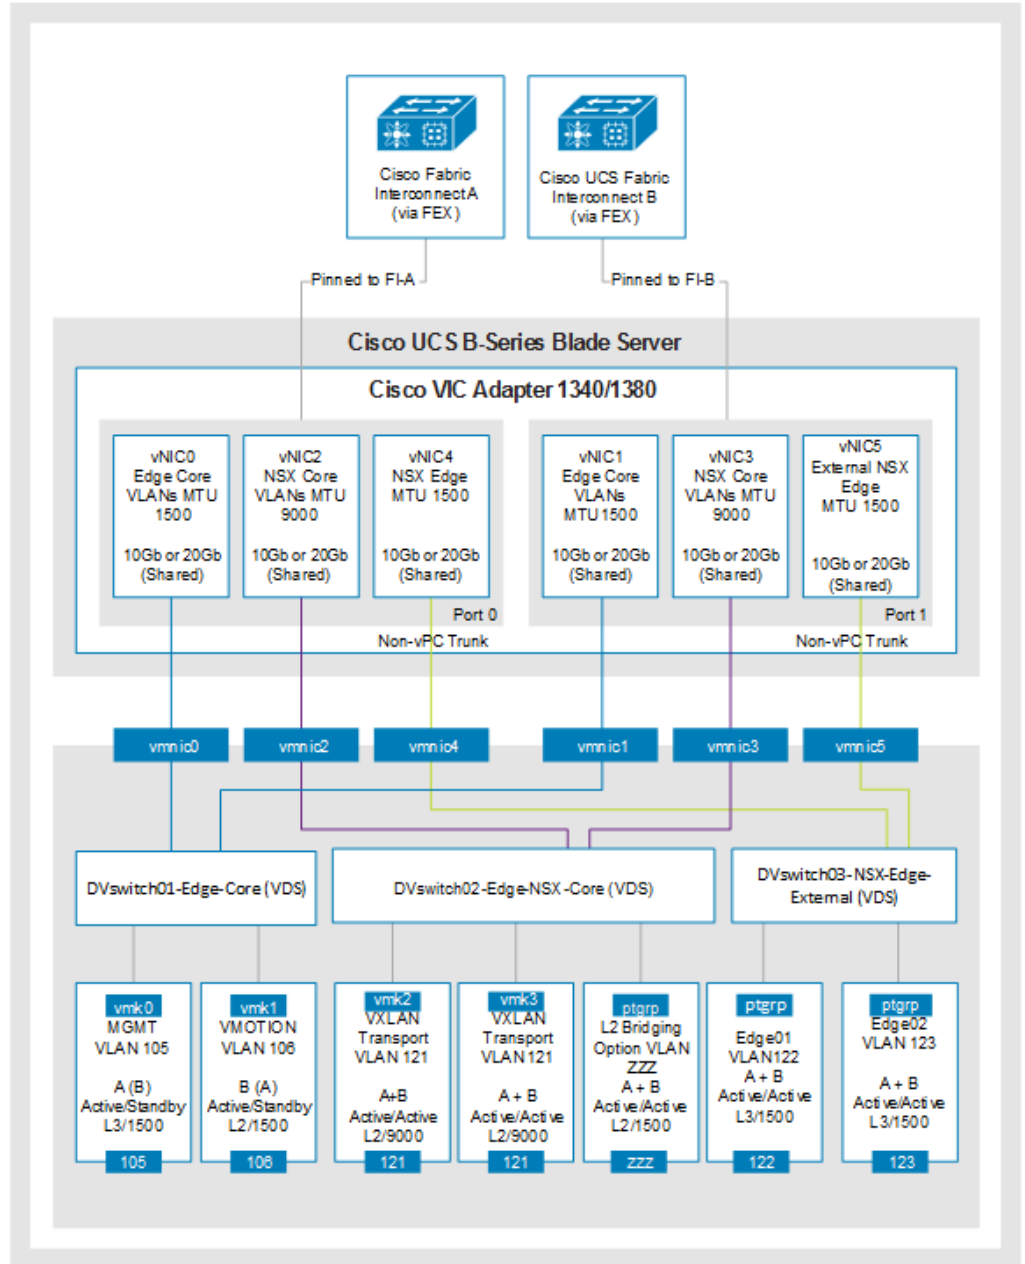
Adding more than four ESGs to the design requires a second Intel X520 card within each of the edge Cisco rack mount servers. This means four uplinks instead of two are added to the VMware vSphere Distributed Switch.

The number of VTEPs is based on the number of uplinks on the VMware vSphere Distributed Switch, which is four VTEPs in this design. 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 or eight ESGs:



Logical network for the Cisco UCS B-Series Blade Servers (edge cluster)

The following illustration shows the VMware virtual network layout for the hosts running on Cisco UCS B-Series Blade Servers in the edge cluster:



VMware NSX compute cluster

The VMware NSX compute cluster contains all of the production VMs.

Hardware requirements

For best performance, make sure each NSX compute cluster has network adapters that support VXLAN offload capabilities, such as the Cisco VIC 1340 and 1380.

VXLAN tunnel end point

The number of VXLAN Tunnel End Points (VTEPs) deployed to each VMware vSphere ESXi depends on the number of dvUplinks configured on the VMware VDS that has the transport distributed port group created. 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 blade servers design has two VXLAN Tunnel End Points (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.

VMware virtual network for the compute cluster

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).

Note: By default, the VMware vSphere ESXi management traffic resides on the VMware standard switch. However, you can put the VMware vSphere ESXi management traffic on the VMware VDS.

Logical network for the Cisco UCS B-Series Blade Servers (compute cluster)

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

