

Cisco Virtualized Multi-Tenant Data Center, Version 2.1 Design Guide

Last Updated: October 14, 2011



Cisco
Validated
Design



CCDE, CCENT, CCSI, Cisco Eos, Cisco Explorer, Cisco HealthPresence, Cisco IronPort, the Cisco logo, Cisco Nurse Connect, Cisco Pulse, Cisco SensorBase, Cisco StackPower, Cisco StadiumVision, Cisco TelePresence, Cisco TrustSec, Cisco Unified Computing System, Cisco WebEx, DCE, Flip Channels, Flip for Good, Flip Mino, Flipshare (Design), Flip Ultra, Flip Video, Flip Video (Design), Instant Broadband, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn, Cisco Capital, Cisco Capital (Design), Cisco:Financed (Stylized), Cisco Store, Flip Gift Card, and One Million Acts of Green are service marks; and Access Registrar, Aironet, AITouch, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Lumin, Cisco Nexus, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, Continuum, EtherFast, EtherSwitch, Event Center, Explorer, Follow Me Browsing, GainMaker, iLYNX, IOS, iPhone, IronPort, the IronPort logo, Laser Link, LightStream, Linksys, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, PCNow, PIX, PowerKEY, PowerPanels, PowerTV, PowerTV (Design), PowerVu, Prisma, ProConnect, ROSA, SenderBase, SMARTnet, Spectrum Expert, StackWise, WebEx, and the WebEx logo are registered trademarks of Cisco and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1002R)

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Cisco Virtualized Multi-Tenant Data Center, Version 2.1, Design Guide
© 2011 Cisco Systems, Inc. All rights reserved.



CONTENTS

Preface vii

Purpose of This Document	vii
Audience	vii
Problem Identification	ii-vii
Solution Objectives	viii
Change Summary	viii
Related Documentation	ix
About Cisco Validated Designs	x

VMDC Architecture Overview 1-1

Key Components	1-1
Network	1-2
Cisco Nexus 7000	1-3
Cisco Nexus 5000	1-3
Cisco Nexus 1000V	1-3
Cisco Network Analysis Module (NAM) Virtual Service Blade	1-4
Services	1-4
Cisco Data Center Services Node (DSN)	1-4
Compute	1-5
Cisco UCS and UCSM	1-6
VMware vSphere and vCenter Server	1-6
Storage	1-7
Cisco MDS 9513	1-7
Cisco Management Interface	1-7
EMC Symmetrix VMAX	1-8
NetApp FAS6080 Filer	1-8
Service Orchestration	1-8
BMC CLM 2.1	1-8
Hierarchical Network Design Reference Model	1-8
Core Layer	1-9
Aggregation Layer	1-10
Services Layer	1-11
Access Layer	1-11
Virtual Access Layer Edge	1-12

Modular Building Blocks	1-12
Pod	1-12
Integrated Compute Stack (ICS)	1-14
Vblock	1-14
FlexPod	1-14
Virtualized Multi-Tenancy	1-15
Virtual Private Data Center Concept	1-15
Differentiated Services	1-15
Service Orchestration	1-16
BMC CLM 2.1	1-16
BMC CLM 2.1 Solution Architecture	1-17
BMC CLM 2.1 Enhancements	1-17
VMDC Design Considerations	2-1
High Availability	2-1
Network Availability	2-1
Aggregation and Access Layer Availability	2-2
Nexus 1010 Deployment Options	2-3
Services Availability	2-4
Active-Active Mode with Multiple Virtual Contexts	2-5
Virtual Access Availability	2-6
Nexus 1010 Manager High Availability	2-7
VSM High Availability	2-7
Compute Availability	2-8
Storage Availability	2-10
Virtualized Multi-Tenancy	2-10
Flexible Tenant Model	2-10
Network and Services Separation	2-13
Compute Separation	2-14
Storage Separation	2-14
Storage Area Network (SAN)	2-14
Network Attached Storage (NAS)	2-14
Performance and Scalability	2-15
Validated Scale	2-15
Understanding Tenant Scalability	2-17
Per Tenant Multicast Support	2-18
Anycast RP for PIM-SM	2-20
PIM Sparse Mode (PIM-SM)	2-20
IGMP Snooping	2-20
IGMP Snooping Querier	2-20

Jumbo Frame Support	2-21
Platform Specific Limits	2-22
Service Assurance	2-22
Traffic Engineering	2-23
MAC Pinning	2-26
Quality of Service Framework	2-27
Classification and Marking	2-28
Queuing	2-30
Network Analysis	2-31
NetFlow	2-31
Encapsulated Remote Switched Port Analyzer (ERSPAN)	2-32
Bill of Materials As Validated	A-1



Preface

The Cisco Virtualized Multi-tenant Data Center (VMDC) is a reference architecture for cloud ready infrastructure and is a design that is validated in a lab environment. This guide describes the design of the Cisco VMDC architecture and identifies environment-specific considerations to be addressed prior to deployment. It also discusses the problems solved by this architecture and describes the four pillars of a cloud-ready, multi-tenancy environment. This design guide focuses on infrastructure elements but does not address automation and orchestration considerations.

This preface contains the following topics:

- [Purpose of This Document, page vii](#)
- [Audience, page vii](#)
- [Solution Objectives, page viii](#)
- [Related Documentation, page ix](#)
- [About Cisco Validated Designs, page x](#)

Purpose of This Document

This document identifies the design considerations and validation efforts required to design and deploy a cloud-ready infrastructure that serves as a foundation for either Infrastructure as a Service (IaaS) offerings or application environments deployed on a shared infrastructure.

Audience

The target audience for this guide includes, but is not limited to, sales engineers, field consultants, professional services, IT managers, partner engineering, and customers who want to deploy a Cisco VMDC-based cloud ready infrastructure.

Problem Identification

Today's traditional IT model suffers because resources are located in different, unrelated silos, which leads to low utilization, gross inefficiency, and an inability to respond quickly to changing business needs. Enterprise servers reside in one area of the data center and network switches and storage arrays

in another. In many cases, different business units own much of the same type of equipment, use it in much the same way, in the same data center row, and yet require separate physical systems to separate their processes and data from each other.

This separation is often ineffective, complicates the delivery of IT services, and sacrifices business activity alignment. As the IT landscape changes, cost reduction pressures, focus on time to market, and employee empowerment are compelling enterprises and IT providers to develop innovative strategies to address these challenges.

By deploying a Cisco VMDC infrastructure, each business unit can be a tenant and benefit from the transparency of the virtual environment that still "looks and feels" like the traditional physically separate topology.

From the tenant viewpoint, each system is separate with its own network and storage; however, the separation is not provided by a server rack, but by a Cisco VMDC environment. The servers, networks, and storage are securely separated and in some cases, more so than in a traditional environment.

Solution Objectives

The Cisco VMDC architecture is a blueprint for organizations that either want to start moving toward or move all the way toward a cloud infrastructure. This design addresses the following key requirements:

- It creates a shared infrastructure that avoids parallel underutilized assets.
- It provides a transition from a single tenant model per dedicated infrastructure to a multi-tenant model using a shared infrastructure.
- Using a shared environment, it matches the isolation and security of a dedicated environment.
- It scales in overall infrastructure and in individual tenant segments.

The secure cloud architecture extends end-to-end control of the tenant environment, from compute platform through network connectivity, storage resources, and data management. This architecture enables Service Providers and Enterprises to securely offer their users unprecedented control over their entire application environment. Unique isolation technologies combined with extensive management flexibility deliver the cloud computing benefits that IT providers require to confidently provide high levels of security and service for multi-tenant customers and consolidated application environments.

Change Summary

Cisco VMDC 2.1 is based on Cisco's general multi-tenancy architecture and improves the Cisco VMDC 2.0 Compact Pod design. The Cisco VMDC 2.0 Compact Pod validated design documents are located at the following URLs:

Cisco VMDC 2.0 Design Guide

http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/VMDC/2.0/design_guide/vmdcDesignGuideCompactPoD20.html

Cisco VMDC 2.0 Deployment Guide

http://www.cisco.com/en/US/partner/docs/solutions/Enterprise/Data_Center/VMDC/2.0/implementation_guide/vmdcImplementationGuideCompactPod20.html

Table ii-1 summarizes the high-level differences.

Table ii-1 Summary of Changes between Cisco VMDC 2.1 and Cisco VMDC 2.0

Requirement	Cisco VMDC 2.1	Cisco VMDC 2.0
End-to-End Network Architecture	Services on the stick design modification (Core/Aggregation handoff)	
Enterprise centric services integration	Services sandwich design (Aggregation/Sub-aggregation)	
Service Orchestration	Orchestration requirements addressed separately	Service Orchestration and network-compute-workload automation with BMC AO, BBSA, BBNA, UCSM, and VCenter
SLA Assurance	Enterprise multi-tenancy SLA with QoS and alignment with WAN/Campus QoS requirements	Preliminary QoS guidelines based on VM role
Applications	Functional multicast validation for end-to-end DC components covering clustering and VRF enabled multicast requirements	Multicast applications not validated
Products/Monitoring	Cisco Nexus 1010 integration and Cisco Network Analysis Module (NAM) capability validation	No major monitoring capability
Other Critical Features	Jumbo MTU support and jumbo frame validation	Jumbo frame support not validated

Related Documentation

The Cisco VMDC design recommends that general Cisco data center design best practices be followed as the foundation for IaaS deployments. The following Cisco Validated Design (CVD) companion documents provide guidance on such a foundation:

Data Center Design—IP Network Infrastructure

http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/DC_3_0/DC-3_0_IPInfra.html

Data Center Service Patterns

http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/DC_3_0/dc_serv_pat.html

Security and Virtualization in the Data Center

http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/DC_3_0/dc_sec_design.html

Designing Secure Multi-Tenancy into Virtualized Data Centers

http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns743/ns1050/landing_dcVDDC.html

Enhanced Secure Multi-Tenancy Design Guide

http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/Virtualization/secureldg_V2.html

The following VMDC solution document provide additional details on the solution:

Cisco VMDC 1.1 Design and Deployment Guide

http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/VMDC/vmdcDdg11.pdf

Cisco VMDC Solution Overview

http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/solution_overview_c22-602978.html

Cisco VMDC Solution White Paper

http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns836/white_paper_c11-604559.html

About Cisco Validated Designs

The Cisco Validated Design Program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information visit www.cisco.com/go/validateddesigns.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.



CHAPTER 1

VMDC Architecture Overview

Cisco VMDC 2.1 is a reference architecture that helps customers design, deploy, and implement data center network, compute, storage and management while demonstrating the integration points between each. This architecture demonstrates how enterprise and mid-market customers can build highly resilient, scalable, secure and manageable systems that can support any type of application workload, whether client-server, virtualized or “big data-oriented”. By providing prescriptive reference designs for each component of the system, enterprises can more rapidly deploy existing and emerging technologies with minimal risk while maintaining investment protection. This solution helps customers evolve towards a simple, agile and efficient data center as they evaluate emerging trends such as cloud.

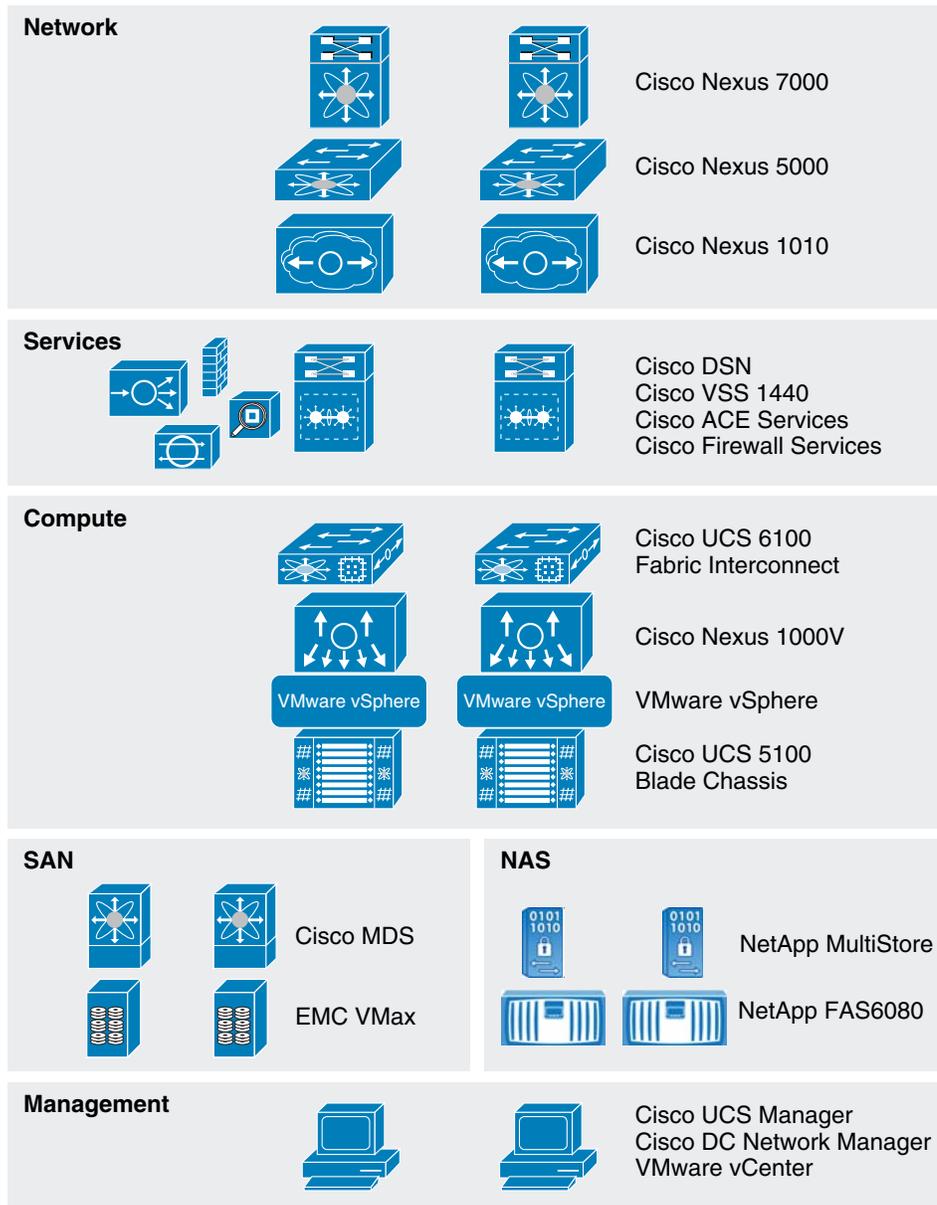
Cisco VMDC 2.1 leverages the following key cloud building concepts:

- **Key Components** At the core of the VMDC 2.1 solution are the Cisco Nexus family of products, the Cisco Datacenter Services Node, and the Cisco Unified Computing System. These key Cisco products provide a wide range of business benefits including convergence, scalability, and intelligence to the datacenter.
- **Hierarchical Network Design Model** This model uses redundant switches at each layer of the network topology for device-level failover that creates a highly available transport between end nodes using the network.
- **Modular Building Blocks** The resource pools consist of three main components: network, compute, and storage. Each of these components is virtualized so that each cloud tenant appears to have its own set of physical resources.
- **Virtualized Multi-Tenancy** Virtualization allows the logical division of a shared pool of network, compute, and storage resources among multiple tenants.
- **Service orchestration** Service orchestration is an add on deployment option that uses a set of tools and APIs to automate the provisioning process using a predefined workflow. Service orchestration is presented as a web portal from which an end user can request specific resources from the datacenter.

Key Components

The VMDC 2.1 solution includes key components at each layer in the datacenter design. The complete list of components at each layer is illustrated in [Figure 1-1](#).

Figure 1-1 Key Components of the Cisco VMDC 2.1 Solution



291769

Network

The following components were used in the network layer of VMDC 2.1:

[Cisco Nexus 7000](#)

[Cisco Nexus 5000](#)

[Cisco Nexus 1000V](#)

[Cisco Network Analysis Module \(NAM\) Virtual Service Blade](#)

Cisco Nexus 7000

As Cisco's flagship switching platform, the Cisco Nexus 7000 Series is a modular switching system designed to deliver 10 Gigabit Ethernet and unified fabric in the data center. This new platform delivers exceptional scalability, continuous operation, and transport flexibility. It is primarily designed for the core and aggregation layers of the data center.

The Cisco Nexus 7000 Platform is powered by Cisco NX-OS (<http://www.cisco.com/en/US/products/ps9372/index.html>), a state-of-the-art operating system, and was specifically designed with the unique features and capabilities needed in the most mission-critical place in the network, the data center.

For more information, see: <http://www.cisco.com/en/US/products/ps9402/index.html>.

Cisco Nexus 5000

The Cisco Nexus 5000 Series (<http://www.cisco.com/en/US/products/ps9670/index.html>), part of the Cisco Nexus Family of data center class switches, delivers an innovative architecture that simplifies data center transformation. These switches deliver high performance, standards-based Ethernet and FCoE that enables the consolidation of LAN, SAN, and cluster network environments onto a single Unified Fabric. Backed by a broad group of industry-leading complementary technology vendors, the Cisco Nexus 5000 Series is designed to meet the challenges of next-generation data centers, including dense multisolet, multicore, virtual machine-optimized deployments, where infrastructure sprawl and increasingly demanding workloads are commonplace.

The Cisco Nexus 5000 Series is built around two custom components: a unified crossbar fabric and a unified port controller application-specific integrated circuit (ASIC). Each Cisco Nexus 5000 Series Switch contains a single unified crossbar fabric ASIC and multiple unified port controllers to support fixed ports and expansion modules within the switch.

The unified port controller provides an interface between the unified crossbar fabric ASIC and the network media adapter and makes forwarding decisions for Ethernet, Fibre Channel, and FCoE frames. The ASIC supports the overall cut-through design of the switch by transmitting packets to the unified crossbar fabric before the entire payload has been received. The unified crossbar fabric ASIC is a single-stage, nonblocking crossbar fabric capable of meshing all ports at wire speed. The unified crossbar fabric offers superior performance by implementing QoS-aware scheduling for unicast and multicast traffic. Moreover, the tight integration of the unified crossbar fabric with the unified port controllers helps ensure low latency lossless fabric for ingress interfaces requesting access to egress interfaces.

For more information, see: <http://www.cisco.com/en/US/products/ps9670/index.html>.

Cisco Nexus 1000V

The Nexus 1000V switch is a software switch on a server that delivers Cisco VN-Link services to virtual machines hosted on that server. It takes advantage of the VMware vSphere framework to offer tight integration between server and network environments and help ensure consistent, policy-based network capabilities to all servers in the data center. It allows policy to move with a virtual machine during live migration, ensuring persistent network, security, and storage compliance, resulting in improved business continuance, performance management, and security compliance. Last but not least, it aligns management of the operational environment for virtual machines and physical server connectivity in the data center, reducing the total cost of ownership (TCO) by providing operational consistency and visibility throughout the network. It offers flexible collaboration between the server, network, security, and storage teams while supporting various organizational boundaries and individual team autonomy.

The Nexus 1010 Virtual Services Appliance hosts the Cisco Nexus 1000V Virtual Supervisor Module (VSM) and supports the Cisco Nexus 1000V Network Analysis Module (NAM) Virtual Service Blade to provide a comprehensive solution for virtual access switching. The Cisco Nexus 1010 provides dedicated hardware for the VSM, making the virtual access switch deployment much easier for the network administrator.

For more information on Nexus 1000V,
see: <http://www.cisco.com/en/US/partner/products/ps10785/index.html>.

For more information on Cisco VN-Link technologies
see: <http://www.cisco.com/en/US/netsol/ns894/index.html>.

Cisco Network Analysis Module (NAM) Virtual Service Blade

The NAM offers flow-based traffic analysis of applications, hosts, and conversations, performance-based measurements on application, server, and network latency, quality of experience metrics for network-based services and problem analysis using deep, insightful packet captures. The Cisco NAM includes an embedded, Web-based Traffic Analyzer GUI that provides quick access to the configuration menus and presents easy-to-read performance reports on Web for different types of services and traffic. The Cisco NAM line of products improves visibility into and monitors the performance of the many physical and virtual layers within the data center.

For more information,
see: http://www.cisco.com/en/US/products/ps5740/Products_Sub_Category_Home.html.

Services

The following components were used in the services layer of VMDC 2.1:

- [Cisco Data Center Services Node \(DSN\)](#)

Cisco Data Center Services Node (DSN)

The Cisco DSN is a dedicated Cisco Catalyst® 6500 Series services chassis housing three Cisco FWSMs and one Cisco ACE Module, providing up to 15 Gbps of secure load-balancing system throughput. The Cisco DSN enables cloud services by integrating firewall security and application delivery along with third-party solutions and monitoring.

Cisco Catalyst 6500 Virtual Switching System 1440

The Cisco Catalyst 6500 Series Virtual Switching System (VSS) 1440 allows for the merging of two physical Cisco Catalyst 6500 Series Switches together into a single, logically-managed entity. The key enabler of a VSS 1440 is the Virtual Switching Supervisor 720-10G. Once a VSS 1440 is created it acts as a single virtual Catalyst switch delivering the following benefits:

- **Operational Manageability**

Two Catalyst 6500s share a single point of management, single gateway IP address, and single routing instance eliminating the dependence on First Hop Redundancy Protocols (FHRP) and Spanning Tree Protocols.

- **Availability**

Delivers deterministic, sub-200 millisecond Layer 2 link recovery through inter-chassis stateful failovers and the predictable resilience of Etherchannel.

- **Scalability**

Scales system bandwidth capacity to 1.4 Tbps by activating all available bandwidth across redundant Catalyst 6500 switches.

The VSS platform fully supports the use of Cisco integrated service modules such as the Cisco Application Control Engine (ACE), Firewall Services Module, and Network Analysis Module. In addition, the VSS platform is capable of supporting both gigabit and ten gigabit Ethernet devices allowing for network based services via a variety of appliance form factors.

Cisco Firewall Services Module

The Cisco Firewall Services Module (FWSM) is a stateful firewall residing within a Catalyst 6500 switching platform. The integrated module employs the power, cooling and space available in the chassis to provide data center security services. The FWSM module offers device level redundancy and scalability through multiple virtual security contexts. Each virtual security context may be transparently introduced at the Layer 2 network level or as a router "hop" at Layer 3. With either deployment model, the security policies associated with each virtual context are consistently applied to protect the related data center networks.

For more information,

see: <http://www.cisco.com/en/US/products/hw/modules/ps2706/ps4452/index.html>.

Cisco Application Control Engine (ACE)

The Cisco Application Control Engine (ACE) module and application platforms perform server load balancing, network traffic control, service redundancy, resource management, encryption and security, and application acceleration and optimization, all in a single network device. The Cisco ACE technologies provide device and network service level availability, scalability, and security features to the data center.

The Cisco ACE offers the following device level services:

- Physical redundancy with failover capabilities for high availability
- Scalability through virtualization allows ACE resources to be logically partitioned and assigned to meet specific tenant service requirements
- Security via access control lists and role-based access control

Network service levels support the following:

- Application availability through load balancing and health monitoring of the application environments
- Scalability of application load balancing, health monitoring, and session persistence policies as all are locally defined within each ACE virtual partition
- Security services including ACLs and transport encryption (SSL/TLS) between the ACE virtual context, client population, and associated server farm

For more information,

see: http://www.cisco.com/en/US/products/ps5719/Products_Sub_Category_Home.html.

Compute

The following components were used in the compute layer of VMDC 2.1:

- [Cisco UCS and UCSM](#)
- [VMware vSphere and vCenter Server](#)

Cisco UCS and UCSM

The Cisco Unified Computing System is a revolutionary new architecture for blade server computing. The Cisco UCS is a next-generation data center platform that unites compute, network, storage access, and virtualization into a cohesive system designed to reduce total cost of ownership (TCO) and increase business agility. The system integrates a low-latency, lossless 10 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multi-chassis platform in which all resources participate in a unified management domain. Managed as a single system whether it has one server or 320 servers with thousands of virtual machines, the Cisco UCS decouples scale from complexity. The Cisco UCS accelerates the delivery of new services simply, reliably, and securely through end-to-end provisioning and migration support for both virtualized and non-virtualized systems.

UCS Components

The Cisco Unified Computing System is built from the following components:

- **Cisco UCS 6100 Series Fabric Interconnects**
(<http://www.cisco.com/en/US/partner/products/ps10276/index.html>) is a family of line-rate, low-latency, lossless, 10-Gbps Ethernet and Fibre Channel over Ethernet interconnect switches.
- **Cisco UCS 5100 Series Blade Server Chassis**
(<http://www.cisco.com/en/US/partner/products/ps10279/index.html>) supports up to eight blade servers and up to two fabric extenders in a six rack unit (RU) enclosure.
- **Cisco UCS 2100 Series Fabric Extenders**
(<http://www.cisco.com/en/US/partner/products/ps10278/index.html>) bring unified fabric into the blade-server chassis, providing up to four 10-Gbps connections each between blade servers and the fabric interconnect.
- **Cisco UCS B-Series Blade Servers**
(<http://www.cisco.com/en/US/partner/products/ps10280/index.html>) adapt to application demands, intelligently scale energy use, and offer best-in-class virtualization.
- **Cisco UCS B-Series Network Adapters**
(<http://www.cisco.com/en/US/partner/products/ps10280/index.html>) offer a range of options, including adapters optimized for virtualization, compatibility with existing driver stacks, or efficient, high-performance Ethernet.
- **Cisco UCS Manager** (<http://www.cisco.com/en/US/partner/products/ps10281/index.html>) provides centralized management capabilities for the Cisco Unified Computing System.

For more information, see: <http://www.cisco.com/en/US/partner/netsol/ns944/index.html>.

VMware vSphere and vCenter Server

VMware vSphere and vCenter Server offer the highest levels of availability and responsiveness for all applications and services with VMware vSphere, the industry's most reliable platform for data center virtualization. Optimize IT service delivery and deliver the highest levels of application service agreements with the lowest total cost per application workload by decoupling your business critical applications from the underlying hardware for unprecedented flexibility and reliability.

VMware vCenter Server provides a scalable and extensible platform that forms the foundation for virtualization management (<http://www.vmware.com/solutions/virtualization-management/>). VMware vCenter Server, formerly VMware VirtualCenter, centrally manages VMware vSphere (<http://www.vmware.com/products/vsphere/>) environments, allowing IT administrators dramatically improved control over the virtual environment compared to other management platforms. VMware vCenter Server:

- Provides centralized control and visibility at every level of virtual infrastructure.
- Unlocks the power of vSphere through proactive management.
- Is a scalable and extensible management platform with a broad partner ecosystem.

For more information, see <http://www.vmware.com/products/>.

Storage

The following components were used in the storage layer of VMDC 2.1:

- [Cisco MDS 9513](#)
- [Cisco Management Interface](#)
- [EMC Symmetrix VMAX](#)
- [NetApp FAS6080 Filer](#)

Cisco MDS 9513

The Cisco MDS 9513 Multilayer Director allows you to deploy high-performance SANs using a high-performance, protocol-independent switch fabric. It provides uncompromising high availability, security, scalability, ease of management, and transparent integration of new technologies for extremely flexible data center SAN solutions. The Cisco MDS 9513 is compatible with first-, second-, and third-generation Cisco MDS 9000 Family switching modules.

For more information, see: <http://www.cisco.com/en/US/products/hw/ps4159/index.html>.

Cisco Management Interface

The following Cisco management interfaces were used in the storage layer of VMDC 2.1:

- [Cisco Device Manager](#)
- [Cisco Fabric Manager](#)

Cisco Device Manager

Device Manager is a management solution for Cisco MDS 9000 Family switch chassis. It graphically depicts installed switching modules, the supervisor modules, and the status of each port within each module, the power supplies, and the fan assemblies. Device Manager provides two views, Device View and Summary View. Use Summary View to monitor interfaces on the switch. Use Device View to perform the following switch-level configurations:

- Configure zones for multiple VSANs
- Manage ports, port channels, and trunking
- Manage SNMPv3 security access to switches
- Manage CLI security access to the switch
- Manage alarms, events, and notifications
- Save and copy configuration files and software image
- View hardware configuration
- View chassis, module, port status, and statistics

Cisco Fabric Manager

Fabric Manager is a management solution for the MDS family of switches, the Nexus 5000 SAN features, and the UCS Fabric Interconnect with limited support. It provides a robust centralized management station for SAN and unified fabric-enabled devices such as the MDS family of switches and the Nexus 5000. Using Fabric Manager, you can perform the tasks needed during a device's deployment cycle, such as discovery, inventory, configuration, performance monitoring, and troubleshooting.

The tables in the Fabric Manager Information pane correspond to dialog boxes in Device Manager. While Device Manager shows values for a single switch, Fabric Manager shows values for multiple switches. However, for verifying or troubleshooting device-specific configuration, Device Manager provides more detailed information than Fabric Manager.

For more information, see:

http://www.cisco.com/en/US/partner/docs/switches/datacenter/mds9000/sw/5_0/configuration/guides/fund/fm/fmfund_5_0_1.html.

EMC Symmetrix VMAX

EMC Symmetrix VMAX provides high-end SAN storage for the virtual data center.

For more information, see: <http://www.emc.com/products/detail/hardware/symmetrix-vmax.htm>.

NetApp FAS6080 Filer

The NetApp FAS6080 provided Enterprise Class Network Attached Storage (NAS) Solution over fully redundant 10 Gigabit Ethernet LANs.

For more information, see <http://www.netapp.com/us/products/storage-systems/fas6000/fas6000.html>.

Service Orchestration

The following components were used for Service Orchestration of VMDC 2.1:

- [BMC CLM 2.1](#)

BMC CLM 2.1

BMC Cloud Lifecycle Management provides the foundation for a strong, flexible, and valuable Cloud infrastructure that supports IT operations and delivers exceptional service quality to the business.

For more information, see:

<http://www.bmc.com/products/product-listing/cloud-lifecycle-planning-management-software.html>

Hierarchical Network Design Reference Model

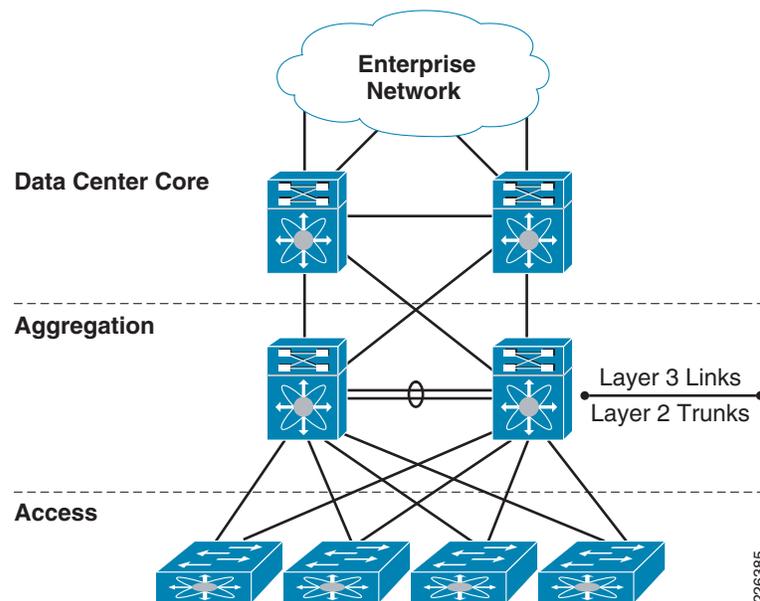
Hierarchical network design has been commonly used in networking for many years. This model uses redundant switches at each layer of the network topology for device-level failover that creates a highly available transport between end nodes. Data center networks often require additional services beyond basic packet forwarding, such as server load balancing, firewall, or intrusion prevention. These services

might be introduced as modules populating a slot of one of the switching nodes in the network or as standalone appliance devices. Each service approach also supports the deployment of redundant hardware to preserve the high availability standards set by the network topology.

A structured data center environment uses a physical layout that correlates tightly to the hierarchy of the network topology. Decisions on cabling types and the placement of patch panels and physical aggregation points must match the interface types and densities of the physical switches being deployed. In a new data center build-out, the two can be designed simultaneously, also taking into consideration the constraints of power and cooling resources. When seeking to avoid significant new investment within an existing data center facility, an architect must consider the pre-existing physical environment of cabling, power, and cooling when selecting switching platforms. Careful planning in conjunction with networking requirements and an eye toward flexibility for the future is critical when designing the physical data center environment. Taking a modular approach to data center design provides flexibility and scalability in both network topology design and utilization of physical resources.

Figure 1-2 illustrates the primary network switching layers of the hierarchical network design reference model for the data center environment. The overall hierarchical model is similar to the reference topology for enterprise campus design, but the term aggregation layer replaces the term distribution layer. The data center network is less concerned with distributing network access across multiple geographically disparate wiring closets and is focused aggregating server resources and providing an insertion point for shared data center services.

Figure 1-2 Hierarchical Network Design Reference Model



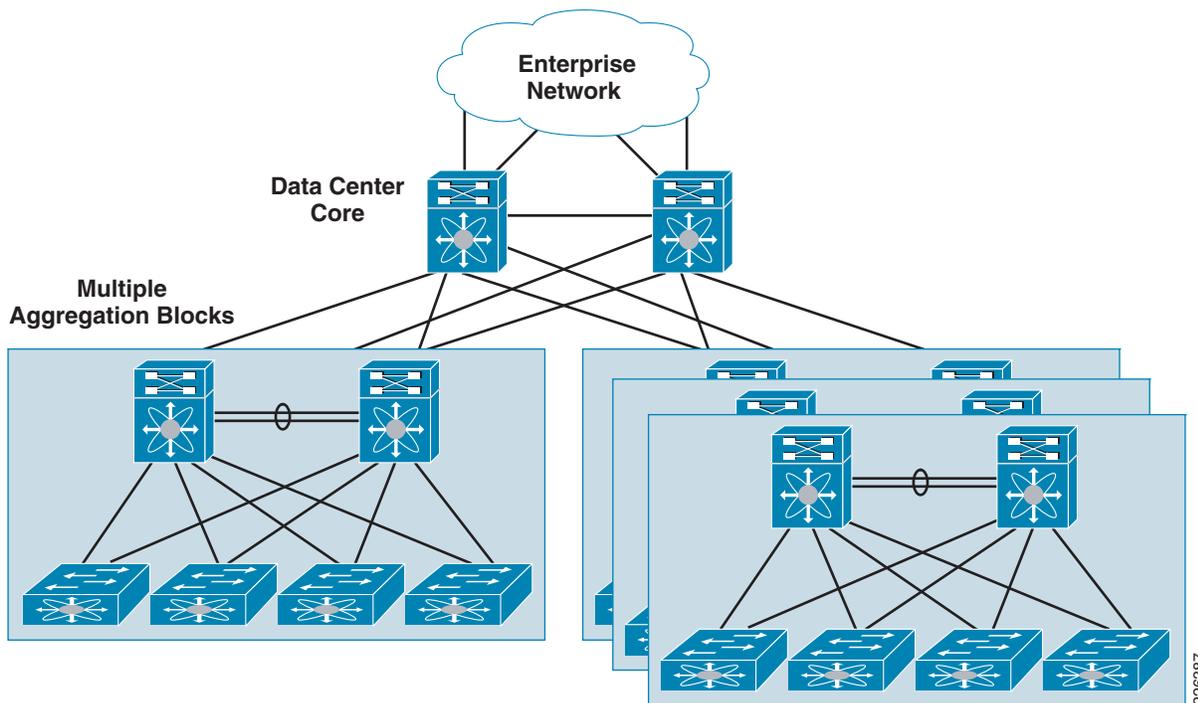
Core Layer

The hierarchical network design model gains much of its stability and high availability characteristics by splitting out switching nodes based on their function, and providing redundant switching units for each functional layer required. The core of a data center network is typically broken out into a pair of high performance, highly available chassis-based switches. In larger or geographically dispersed network environments, the core is sometimes extended to contain additional switches. The recommended approach is to scale the network core continuing to use switches in redundant pairs. The

primary function of the data center network core is to provide highly available, high performance Layer-3 switching for IP traffic among the other functional blocks of the network, such as campus, Internet edge and WAN. By configuring all links connecting to the network core as point-to-point Layer-3 connections, rapid convergence around any link failure is provided, and the control plane of the core switches is not exposed to broadcast traffic from end node devices or required to participate in STP for Layer-2 network loop prevention.

In small-to-medium enterprise environments, it is reasonable to connect a single data center aggregation block, or pod, directly to the enterprise switching core for Layer-3 transport to the rest of the enterprise network. Provisioning a separate, dedicated pair of data center core switches provides additional insulation from the rest of the enterprise network for routing stability and also provides a point of scalability for future expansion of the data center topology. As the business requirements expand and dictate two or more aggregation blocks serving separate pods or zones of the data center, a dedicated data center core network provides for scale expansion without requiring additional Layer-3 interfaces to be available on the enterprise core. An illustration of scaling the data center topology with a dedicated core and multiple aggregation blocks is provided in [Figure 1-3](#).

Figure 1-3 *Scaling the Data Center with a Dedicated Core*



Aggregation Layer

The Aggregation layer of the data center provides connectivity for the Access layer switches in the server farm, and aggregates them into a smaller number of interfaces to be connected into the Core layer. In most data center environments, the Aggregation layer is the transition point between the purely Layer 3 routed Core layer, and the Layer 2-switched Access layer. 802.1Q trunks extend the server farm VLANs between Access and Aggregation layers. The Aggregation layer also provides a common connection point to insert services into the data flows between clients and servers, or between tiers of servers in a multi-tier application.

The preferred devices in this distribution layer of the solution are Nexus 7000 Series switches. From a physical perspective, the Nexus 7000 provides more than enough slot and port density to support the surrounding core, services, and access layer devices within the topology. In addition, the Nexus devices offer a rich set of Layer 2, Layer 3, and virtualization features permitting a new level of segmentation and control within the Aggregation layer of the data center. In fact, the Nexus 7000 VDC construct allows enterprises to consolidate multiple distribution blocks into a pair of Nexus 7000 switches without sacrificing any of the functionality highlighted earlier.

Services Layer

The VMDC reference architecture provides an open flexible model for integrating network services like server load balancing (SLB) and firewall security. These services can be integrated using either appliances or service modules. The VMDC architecture supports both models, however VMDC 2.1 focuses on integration of the Cisco Data Center Services Node.

The Cisco® Data Center Service Node (DSN) complements the Cisco Nexus® 7000 Series Switches in the data center and offers the choice to host specific integrated network services relevant in a given data center. Examples of network services include the Cisco Firewall Services Module (FWSM) and the Cisco ACE Application Control Engine Module, for server load balancing. This services node-based solution offers proven enterprise products enabling customers to use a common architecture and easily integrate the solution with existing network infrastructure

Cisco DSN uses a dual-homed approach for data path connectivity to redundant aggregation-layer switches. This approach decouples the service modules from dependence on a specific aggregation switch. Because the Cisco DSN is self-contained, it provides operational flexibility for the system maintenance that may be required for the aggregation-layer switches or the Cisco DSN. From a high-availability perspective, if one of the aggregation switches or Cisco DSNs fails, traffic can continue to flow through the other aggregation switch to the active Cisco DSN without the need of any failover event in the service modules themselves.

A major advantage of the Cisco DSN is the capability to introduce new services in a controlled manner using predictable traffic patterns. The Cisco DSN consists of a Cisco Catalyst 6500 Series Switch using service modules that are dedicated to security and server load-balancing functions. The Cisco DSN can be directly attached to an aggregation-layer switch, such as a Cisco Nexus 7000 Series Switch, or it can use the Cisco DSN as the aggregation layer if ports are available. The primary goal of the Cisco DSN is to provide higher performance, reliability, and manageability by transparently applying network services in the data center to create a more flexible, functional, and secure server farm.

For more information, please refer to these links:

http://www.cisco.com/en/US/products/ps9336/products_tech_note09186a0080a7c72b.shtml

<http://www.cisco.com/go/vss/>

Access Layer

The access layer of the network provides connectivity for serverfarm end nodes residing in the data center. Design of the access layer is tightly coupled to decisions on server density, form factor, and server virtualization that can result in higher interface count requirements. Traditional data center access layer designs are strongly influenced by the need to locate switches in a way that most conveniently provides cabling connectivity for racks full of server resources. The most commonly used traditional approaches for data center serverfarm connectivity are end-of-row, top-of-rack, and integrated switching. Each design approach has pros and cons, and many enterprises use multiple access models in the same data center facility as dictated by server hardware and application requirements.

The Cisco Nexus 5000 Series switches provide high-density 10-Gigabit Ethernet connectivity and innovative storage integration capabilities for the support of FCoE. With a Layer-2 capable implementation of NX-OS, the Nexus 5000 is optimized for the evolving data center access layer. For customers requiring a density of 1-Gigabit Ethernet server connectivity, the Nexus 2000 Fabric Extenders may be deployed in conjunction with a Nexus 5000 Series switch and treated as a single virtual chassis in the access layer of the data center topology. This approach may be used to provide ToR switching to multiple racks of servers, with all management functions for the Nexus 2000 Fabric Extenders centralized into their associated Nexus 5000 Series switch. The Nexus 5000 Series can also be placed middle-of-row (MoR) to provide 10-Gigabit Ethernet interfaces to nearby servers.

Virtual Access Layer Edge

The evolution of networking technology in the data center is most evident at the access layer of the network and within the server farm. Several options for building the data center access layer introduce switch virtualization that allows the function of the logical Layer-2 access layer to span multiple physical devices. The virtual access-layer is a logical layer inside the server fabric providing connectivity with virtualized server hardware, hypervisor and VMs with additional functionality of policy management (separation, ACL, etc.), mobility, and service assurance capability.

The Nexus 1000V virtual distributed switch allows the network architect to provide a consistent networking feature set across both physical servers and virtualized servers. The Nexus 1000V operates as a virtualized chassis switch, with Virtual Ethernet Modules (VEMs) resident on the individual virtualized servers managed by a central Virtual Supervisor Module (VSM) that controls the multiple VEMs as one logical modular switch. The VSM provides a centralized point of configuration and policy management for the entire virtual distributed switch. Both the Cisco Nexus 2000 Fabric Extenders and the Cisco Nexus 1000V represent variations on the evolving capabilities of the data center virtual-access sub-layer.

Modular Building Blocks

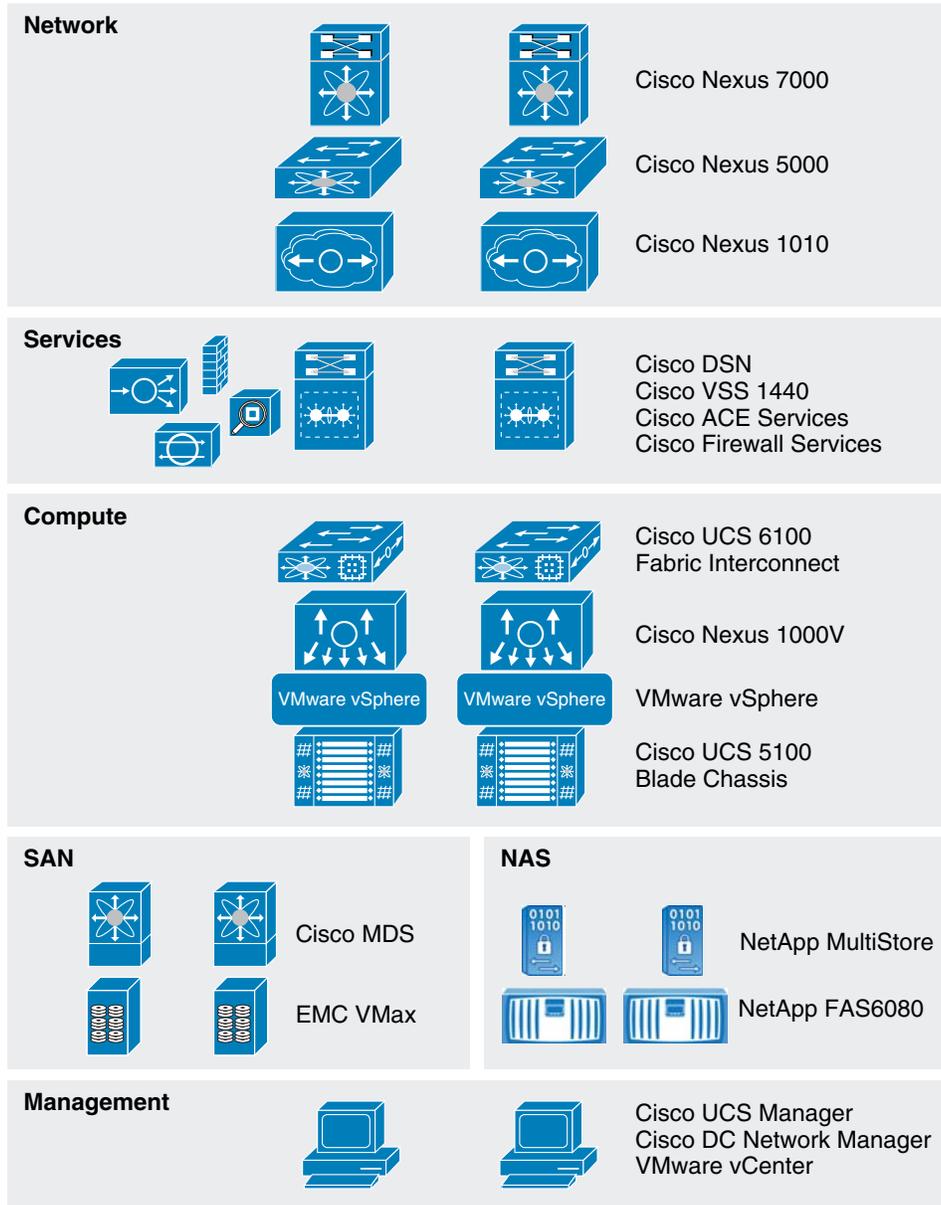
Cisco VMDC 2.1 provides a modular solution that addresses the needs of larger enterprise and smaller service provider data centers. This architectural consistency enables providers to select the design that best suits their immediate needs, while providing the ability to scale to meet future needs without retooling or retraining staff. Within a hierarchical design, this ability to scale is based on two modular building blocks: the pod and the integrated compute stack (ICS). The fundamental business drivers for adopting pod and ICS modularity are as follows:

- Minimize operational impact; reduce total cost of ownership (TCO)
- Flexible, multi-vendor architecture
- Pretested and validated IT infrastructure
- Private cloud foundation

Pod

A pod identifies modular unit of data center components. This modular architecture provides a predictable set of resource characteristics (network, compute, and storage resource pools, power, and space consumption) per unit that is added repeatedly as needed. In this discussion, the aggregation layer switch pair and services layer nodes are the foundation of the pod. To complete the pod architecture one or more integrated compute stacks are added ([Figure 1-4](#)).

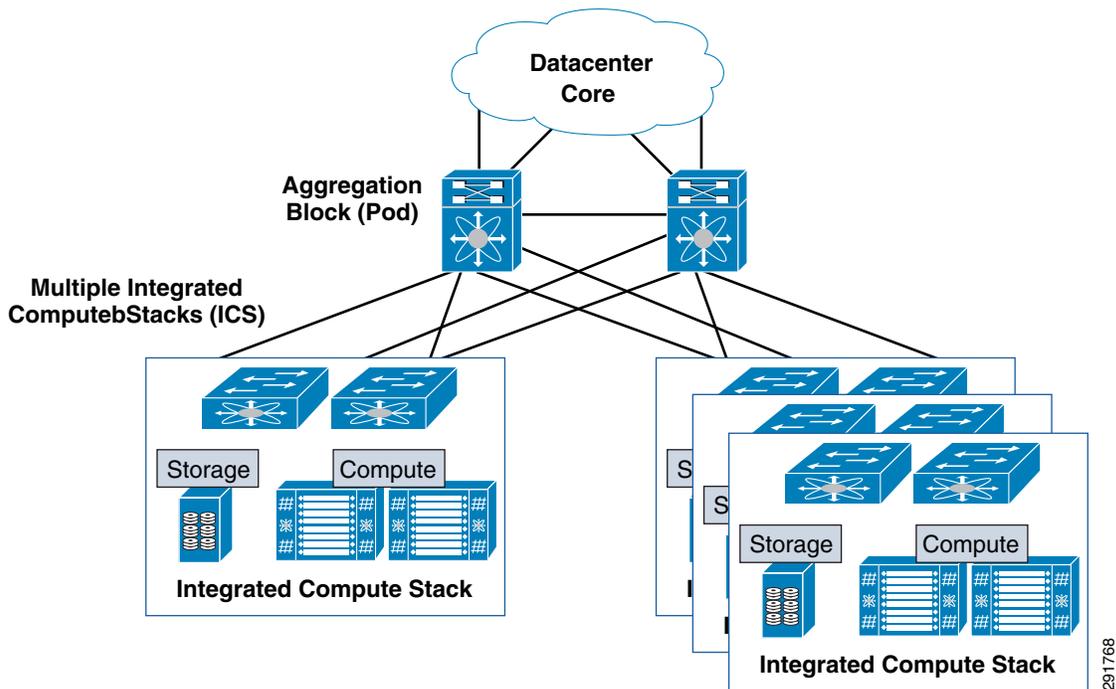
Figure 1-4 VMDC 2.1 Basic Pod Components



291769

To scale a pod, customers can add additional integrated compute stacks (see [Figure 1-5](#)). You can continue to scale in this manner until the pod resources are exceeded.

Figure 1-5 Expanding a Pod with Multiple Integrated Compute Stacks



Pods can be interconnected with each other in the same physical data center or between data centers technologies such as xPLS or Overlay Transport Virtualization (OTV).

Integrated Compute Stack (ICS)

An integrated compute stack can include network, compute, and storage resources in a second smaller repeatable unit (see Figure 1-6). In this discussion, the access layer switch pair, storage, and compute resources are contained within an integrated compute stack. The architectural blueprint and logical overlay of Cisco VMDC 2.1 is independent of ICS infrastructure components.

Vblock

The Vblock combines Cisco UCS with EMC storage components to provide multiple, fixed-sized configuration blocks. The technical overview and detailed information is described at the following URL:

http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns743/ns1050/landing_vblock.html

FlexPod

The FlexPod ICS combines Cisco UCS with NetApp storage to provide variable configuration of compute and storage components based on work load. The current configurations are detailed at the following URL:

http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns743/ns1050/landing_flexpod.html

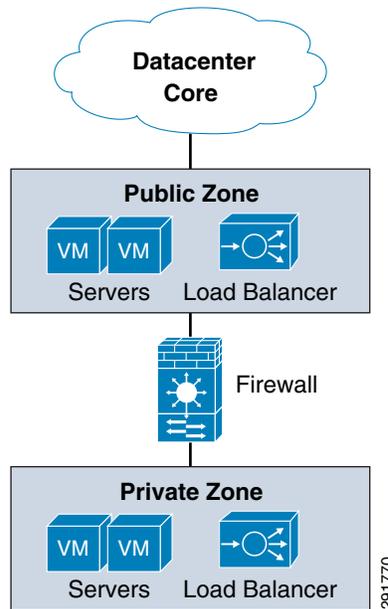
Virtualized Multi-Tenancy

Traditionally, a dedicated infrastructure would be deployed for each tenant that it hosted. This approach, while viable for a multi-tenant deployment model, does not scale well because of cost, complexity to manage, and inefficient use of resources. Deploying multiple tenants in a common infrastructure yields more efficient resource use and lower costs. However, each tenant may require path isolation for security and privacy from others sharing the common infrastructure. Therefore, logical separation or virtualization is a fundamental concept for multi-tenancy in the VMDC environment. Virtualization at the various levels in the VMDC 2.1 architecture provides logical separation in the network, compute, and storage resources.

Virtual Private Data Center Concept

In VMDC 2.1 a tenant can be defined as an external partner or subsidiary or an internal department or business unit, such as engineering or human resources. Each tenant is given a virtual private datacenter construct which contains a public server farm (public zone), a firewall protected private server farm (private zone), and load balancing services available in each zone.

Figure 1-6 VMDC 2.1 Virtual Data Center Tenant Construct



Differentiated Services

Each tenant using the VMDC infrastructure is entitled to some compute, network, and storage resource SLA (Service Level Agreement). One tenant may have higher SLA requirements than another based on a business model or organizational hierarchy. For example, tenant A may have higher compute and network bandwidth requirements than tenant B, while tenant B may have a higher storage capacity requirement. The objective is to ensure that tenants within this environment receive their subscribed

SLAs while their data, communication, and application environments are securely separated, protected, and isolated from other tenants. Cisco VMDC relies on the following key concepts to deliver a solution that meets the requirements of these groups.

- **Availability** allows the infrastructure to meet the expectation of compute, network, and storage to always be available even in the event of single hardware failure. Like the Secure Separation requirement, each layer has its own manner of providing a high availability configuration that works seamlessly with adjacent layers. Security and availability are best deployed in a layered approach.
- **Secure Separation** ensures one tenant cannot disrupt other tenants' resources, such as virtual machine (VM), network bandwidth, tenant data, or storage. It also ensures protection against data loss, denial of service attacks, and unauthorized access. Each tenant must be securely separated using techniques such as access control, virtual storage controllers, VLAN segmentation, and firewall rules. Secure separation also implies a defense-in-depth security approach with policy enforcement and protection at each layer.
- **Service Assurance** provides isolated compute, network, and storage performance during both steady state and non-steady state operation. For example, the network and the UCS blade architecture can provide each tenant with a certain bandwidth guarantee using Quality of Service (QoS); resource pools within VMware help balance and guarantee CPU and memory resources.
- **Management** is required to rapidly provision and manage resources and view resource availability. Domain and element management provides comprehensive administration of the shared resources that comprise the Virtual Multi-tenant Data Center architecture. The demarcation point for managing this design is defined by the interactive and programmable interfaces delivered by Cisco, and partners. The administrative interfaces and APIs in this portfolio address infrastructure components such as UCS Manager, and Data Center Network Manager. These element managers and their associated open APIs provide the foundation for delivering cohesive service lifecycle orchestration with solution partners.

Service Orchestration

Service orchestration is an add on deployment option that uses a set of tools and APIs to automate the provisioning process by using a predefined workflow. Service orchestration is presented as a web portal from which an end user can request specific resources from the datacenter.

BMC CLM 2.1

BMC Cloud Lifecycle Management (CLM) 2.1 solution provides a comprehensive set of capabilities for orchestrating and managing cloud environments. It enables on boarding and pooling of resources for compute, storage, and networking, and creating policies to manage those pools. It provides functionality to provision pods, network containers, physical servers, and virtual server instances. It also provides the ability for end-users, through a portal, to place service requests to create and manage their network containers and server instances. BMC CLM 2.1 is fully multitenant aware. It can support simultaneous use of the cloud environment by multiple tenants that can request, deploy, and operate services independently.

BMC CLM 2.1 solution deploys Cisco's Virtualized Multi-Tenant Data Center (VMDC) 2.1 design that includes the new container model that has been validated as an out of the box blueprint and workflow model. The infrastructure service resources validated in this solution in Cisco labs are:

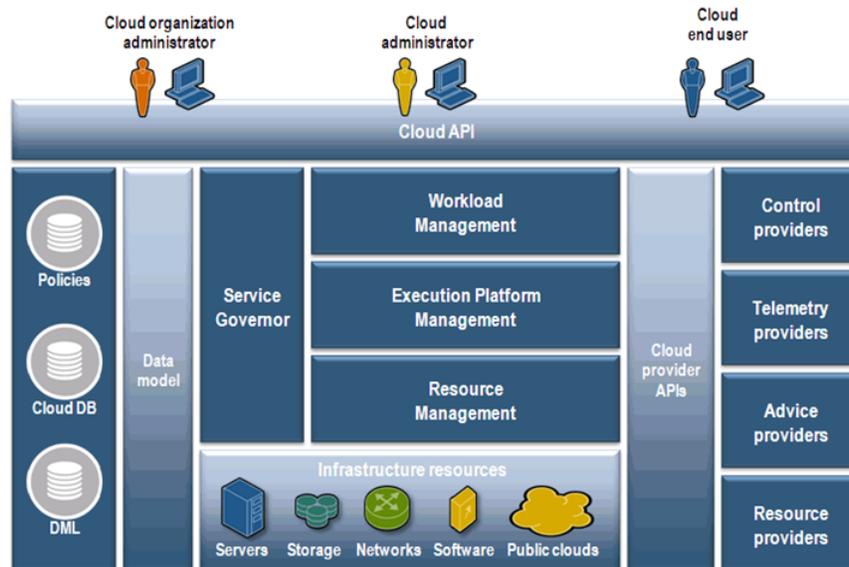
- **Virtual Server**—Memory, CPU, Capacity management, Storage allocation

- **Network Virtualization**—VLAN, VRF, SVI, Virtual Context, Virtual Address, Firewalling, load balancing, ACL filtering
- **Storage**—Multi-pathing, Storage classification, Tiering

BMC CLM 2.1 Solution Architecture

BMC CLM 2.1 has an entirely restructured architecture that provides the foundation for scaling your cloud and for configuring multiple data centers (Figure 1-7).

Figure 1-7 BMC Cloud Lifecycle Management Architecture



BMC CLM 2.1 Enhancements

CLM 2.1 offers several new features backed by a new architecture.

1. **Comprehensive Cloud Administration**—CLM 2.1 has a centralized cloud administration portal, the BMC Cloud Lifecycle Management Administration Console that allows administrators to manage all aspects of their cloud environments.
2. **Enhanced End-User Portal**—CLM 2.1 provides cloud end users with a more powerful portal, with more available management options and detailed information about service instances.
3. **Service Catalog and Service Offerings**—BMC CLM 2.1 features a new Service Catalog for defining service offerings, options, and pricing.
4. **Service Blueprints**—BMC CLM 2.1 supports single and multi-tier applications through service blueprints. Service blueprints represent all of the components of a service offering (such as type, logical mapping to resources, size, and so on) that reside behind the offering that users select in the BMC Cloud Lifecycle Management My Cloud Services Console. Service blueprints can also define the applications that must be installed as part of a service.
5. **End-to-End Support for Multi-Tier Applications**—CLM 2.1 can provision applications that span multiple server instances. It automates the deployment of multi-tier applications.

6. **Service Governor**—The Service Governor enables intelligent, policy-based placement of cloud services.
7. **New Installation Planner**—A new installation planner improves the installation experience by simplifying the solution installation and providing the initial steps for cloud configuration. It also reduces the time and effort to deploy BMC CLM 2.1.
8. **New Solution Architecture**—An entirely restructured architecture provides the foundation for scaling your cloud and for configuring multiple data centers. It allows integration with third-party applications and customization of existing BMC Cloud Lifecycle Management functionality through a new REST API based solution architecture.



CHAPTER 2

VMDC Design Considerations

The Cisco VMDC 2.1 release highlights key design areas of network, compute, and storage with focus on the following:

- [“High Availability” section on page 2-1](#)
- [“Virtualized Multi-Tenancy” section on page 2-10](#)
- [“Performance and Scalability” section on page 2-15](#)
- [“Service Assurance” section on page 2-22](#)

High Availability

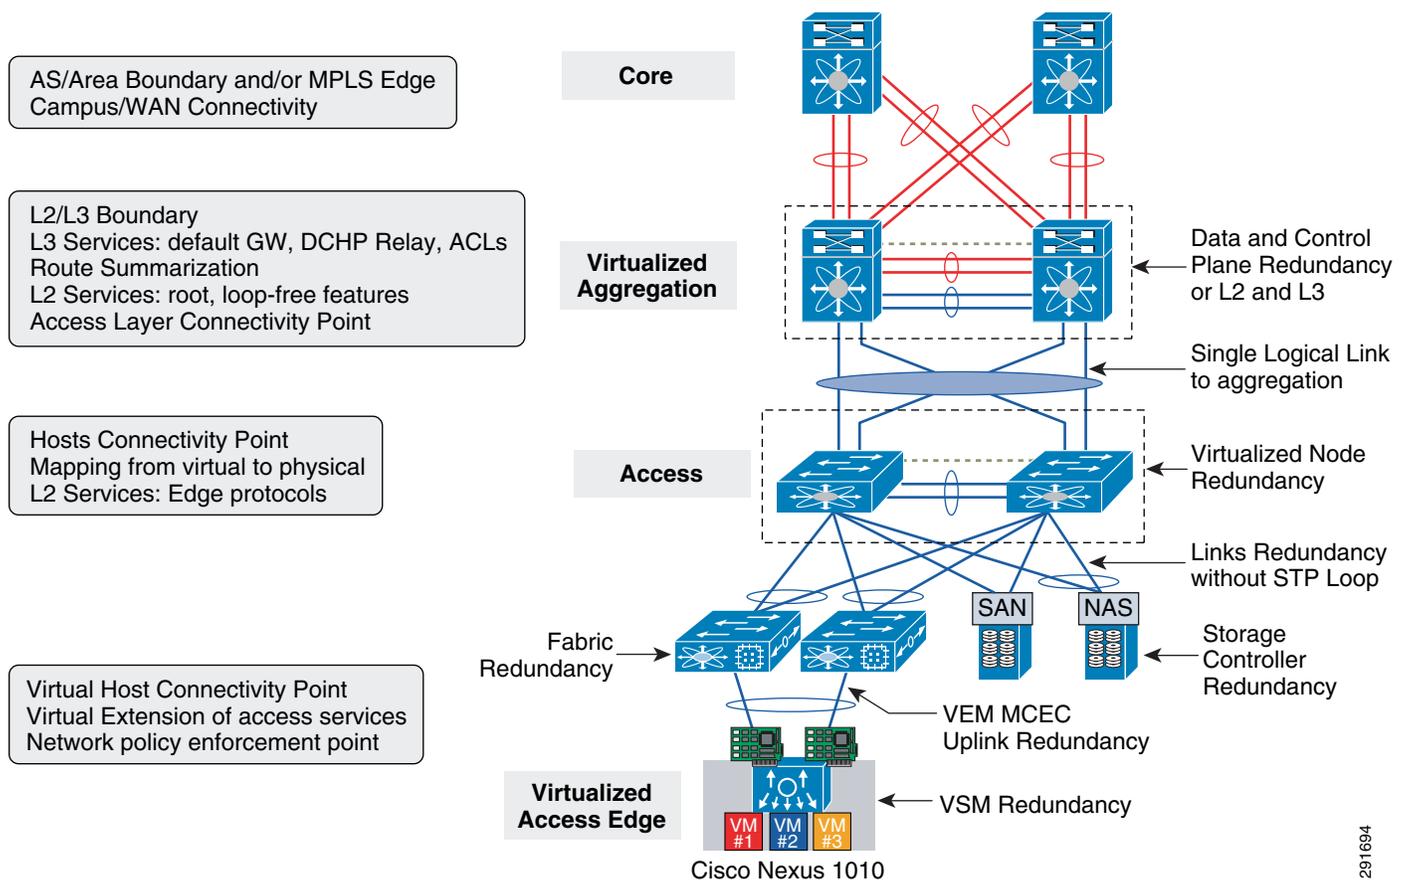
High Availability is key design point for building a virtualized cloud environment. Eliminating planned downtime and preventing unplanned downtime are key aspects in the design of the multi-tenant shared infrastructure. This section covers high availability design considerations and best practices related to the network, compute, and storage components in the VMDC 2.1 architecture.

Network Availability

Network availability is paramount to any organization running a virtualized data center service. It is strategic for disaster planning, as well as everyday operations, and ensures that tenants can reliably access application servers.

The VMDC 2.1 infrastructure design is based on a three-tier model (core, aggregation, and access) as depicted in [Figure 2-1](#).

Figure 2-1 Physical Topology



Aggregation and Access Layer Availability

To achieve high availability in the datacenter several design points as well as key features available with the Nexus 7000 at the aggregation layer and Nexus 5000 at the access layer should be used:

- **Device Redundancy**—The core, aggregation, and access layers are typically composed of at least two physical devices, each with redundant power supplies, fans, fabric cards, supervisors, and line cards.
- **Supervisor Redundancy**—Redundant supervisors can be installed in a single chassis to allow continuous system operation. Cisco NX-OS provides continuous system operation, permitting maintenance, upgrades, and software certification without service interruption. The combination of process modularity, hitless In-Service Software Upgrade (ISSU) capability, and stateful graceful restart mitigates the effects of software upgrades and other operations.
- **Link Redundancy**—The physical links between all layers of the network should be distributed across the redundant linecards to ensure failure of a single module does not adversely impact the overall topology.
- **Virtual Port Channels (vPC)**—The virtual PortChannel (vPC) feature allows one end of a PortChannel to be split across a pair of Cisco Nexus 7000 Series Switches. vPC provides Layer 2 multipathing through the elimination of Spanning Tree Protocol blocked ports in dual-homed connections. vPC enables fully used bisectional bandwidth and simplified Layer 2 logical topologies without the need to change the existing management and deployment models.

- **Multi-Chassis Ether Channels (MEC)**—Multi-Chassis Ether Channels were used to connect the aggregation layer to the services layer. MEC allows for redundant routed paths between the aggregation switches and the services switches.
- **Virtual Route and Forwarding (VRF)**—Redundant VRF instances provide Layer 3 services for their associated tenant segments.
- **Fast Convergence**—Network convergence is optimized by providing tools and functions to make both failover and fallback transparent and fast. For example, Cisco NX-OS provides Spanning Tree Protocol enhancements such as Bridge Protocol Data Unit (BPDU) guard, loop guard, root guard, BPDU filters, and bridge assurance to help ensure the health of the Spanning Tree Protocol control plane; Unidirectional Link Detection (UDLD) Protocol; NSF graceful restart of routing protocols; millisecond timers for First-Hop Resiliency Protocol (FHRP); Shortest-Path First (SPF) optimizations such as link-state advertisement (LSA) pacing and incremental SPF; IEEE 802.3ad link aggregation with adjustable timers; and Bidirectional Forwarding Detection (BFD).

Nexus 1010 Deployment Options

Deployment of the Cisco Nexus 1010 offers many benefits. First, because the Cisco Nexus 1010 appliance is owned and operated by the network team, deployment no longer depends on collaboration with the network, storage, and virtualization operations teams. Instead, the Cisco Nexus 1010 can be installed and deployed in the same way as any networking device.

Another benefit is the flexibility of placement: the Cisco Nexus 1010 can be inserted into the network at various locations. The previous section discussed the four options for connecting the Cisco Nexus 1010 to the network. These methods can be used in various areas of the network. Typically, Cisco Nexus 1010 appliances are deployed in a central management domain. Often, this is where other network appliances, such as the Cisco Application Control Engine (ACE), Cisco Wide Area Application Services (WAAS), the NAM, etc. are deployed.

- For more information on the Cisco Nexus 1010 deployment follow the link below:
- http://www.cisco.com/en/US/partner/prod/collateral/switches/ps9441/ps9902/white_paper_c07-603623.html

The Cisco Nexus 1010 has six Gigabit Ethernet interfaces available for network connectivity: two Gigabit Ethernet LAN interfaces on the motherboard, and four Gigabit Ethernet interfaces available through a PCI card. Four types of traffic flows through these interfaces: management, control, packet, and data traffic. The Cisco Nexus 1010 does not reside in the data path of normal virtual machine data traffic. However, when the Cisco Nexus 1000V NAM Virtual Service Blade is deployed, data traffic from the selected virtual machines will flow to the Cisco Nexus 1010 to be analyzed. The decision to use or not use the NAM is one factor that influences which network connectivity option should be used to connect the Cisco Nexus 1010 to the network.

Figure 2-2 Color Code for Various Cisco Nexus 1010 Traffic



The six interfaces on the Cisco Nexus 1010 can be connected to the network in four ways. The best connectivity option for the Cisco Nexus 1010 in a particular situation depends on the customer's needs and requirements. In VMDC 2.1 Option 3 from the deployment guide was chosen as the preferred deployment option. This option uses the two LOM interfaces for management traffic, and the four interfaces on the PCI card are used carry control, packet, and data traffic. In this configuration, the two

management interfaces should be connected to two separate upstream switches for redundancy. In addition, the four ports used for control, packet, and data traffic should be divided between two upstream switches for redundancy (Figure 2-3).

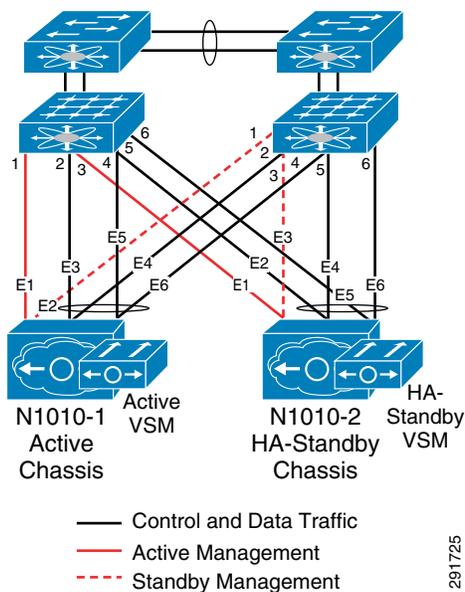
Figure 2-3 Nexus 1010 LOM Interfaces for Management and Other 4 NICs for Control, Packet and Data Traffic



This option is ideal for deployments including a NAM within the Cisco Nexus 1010 but requiring a separate management network. Because control traffic is minimal most of the bandwidth available on the four Gigabit Ethernet interfaces will be used for NAM traffic.

In VMDC 2.1 the Cisco Nexus 1010 appliance is connected to Cisco Nexus 2000 Series Fabric Extenders which connect to the Cisco Nexus 5000 Series. Because the Cisco Nexus 1010 uses Gigabit Ethernet interfaces to connect to the network, the fabric extender provides an optimal connectivity solution.

Figure 2-4 Cisco VMDC 2.1 Nexus 1010 High Availability Deployment



Services Availability

The recommended platform for service modules is the Data Center Services Node (DSN). It is comprised of a Catalyst 6500 in Virtual Switch System (VSS) mode forming a resilient architecture. Integrated service modules, such as the Cisco ACE or FWSM, ASA-SM, or standalone devices, such as the ASA 5500 series or IPS 4200 platforms, may attach directly to the DSN via multi-chassis Etherchannel to create a highly available design.

Active-Active Mode with Multiple Virtual Contexts

With VSS, the service modules will be in active-active mode, with each virtual context in active-standby mode on the designated service modules of each Cisco DSN.

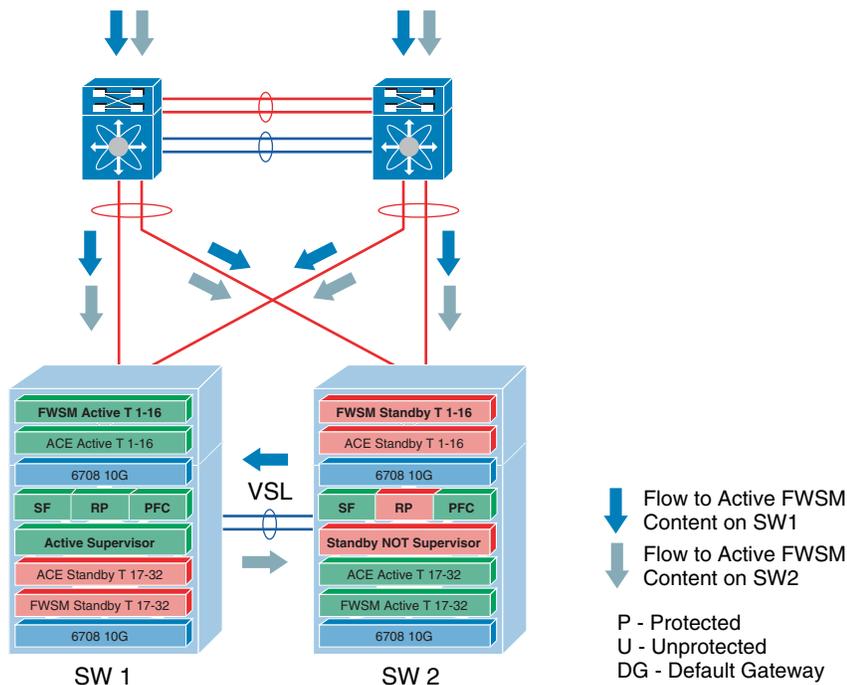
This model uses the virtualization capabilities of the Cisco FWSM and Cisco ACE Module to distribute a portion of the traffic across both services chassis. The traffic is not automatically balanced equally across the devices; however, the network administrator can assign different server farm subnets to specific contexts, based on expected load or on other factors. Routing virtualization is also used in the active-active model through the implementation of VRF instances in the aggregation switches.

The active-active design model allows the Cisco FWSM and Cisco ACE in the Cisco DSN to support an active context, optimizing resources in each Cisco DSN through load distribution across the Cisco DSN pair (VSS). To achieve an active-active design, failover groups are defined for each service module. Failover groups contain virtual contexts and determine which physical Cisco FWSM and Cisco ACE will be active for the particular group. Each module is assigned a primary and secondary priority status for the failover group. The fault-tolerant interface between the Cisco FWSM and Cisco ACE on each chassis uses a separate physical connection between chassis. Since the Cisco DSN is a VSS configuration, all configured VLANs are carried across the virtual switch links (VSLs). As a result, no separate links are needed for fault-tolerant links or stateful connectivity.

With the virtualization capabilities of the Cisco Catalyst 6500 Series services modules, separate contexts can be created that behave like separate virtual devices. The first Cisco FWSM and Cisco ACE are primary for the first context and standby for the second context. The second Cisco FWSM and Cisco ACE are primary for the second context and secondary for the first context. This setup allows modules on both sides of the designs to be primary for a portion of the traffic, and it allows the network administrator to optimize network resources by distributing the load across the topology instead of having one set of modules nearly idle in a pure-standby role.

In an active-active design, network administrators must properly plan for failure events in which one service module supports all the active contexts. If the total traffic exceeds the capacity of the remaining service module, the potential to lose connections exists; thus, it is important to size the VSL accordingly. It is a best practice for the bandwidth of the VSL to be equal to the total amount of uplink traffic coming into a single chassis.

Figure 2-5 Active-Active Services Chassis with Virtual Contexts



To achieve high availability in the services layer, many of the features used for basic network availability are utilized in addition to some key features available with the Nexus 6500:

- **Device Redundancy**—The services layer is typically composed of two physical devices, each with redundant power supplies, fans, line cards, and possibly redundant supervisor modules.
- **Virtual Switching System (VSS)**—The Cisco® Catalyst® 6500 Series Switches Virtual Switching System (VSS) 1440 is a network system virtualization technology that pools two Cisco Catalyst 6500 series switches with Virtual Switching Supervisor 720-10G VSS into a single virtual switch. In a VSS, the data plane and switch fabric of both supervisor engines are active at the same time in both chassis, thereby providing a combined system switching capacity of 1440Gbps.
- **Multi-Chassis EtherChannels (MEC)**—Multi-Chassis EtherChannels were used to connect the aggregation layer to the services layer. MEC allows for redundant routed paths between the aggregation switches and the services switches.
- **Virtual Route and Forwarding (VRF)**—Redundant VRF instances provide Layer 3 services for their associated tenant segments.
- **Service Modules in Active/Active Mode with Multiple Virtual Contexts**—With VSS, the service modules should be deployed in active-active mode, with each virtual context in active-standby mode on the designated service modules of each Catalyst 6500.

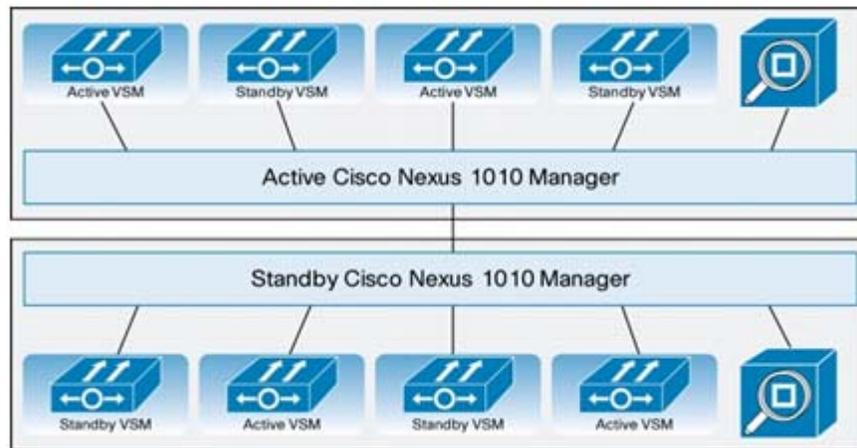
Virtual Access Availability

The Nexus virtual distributed switch (Nexus 1000V and Nexus 1010) form the virtual access layer.

Nexus 1010 Manager High Availability

Two redundant Cisco Nexus 1010 appliances should be deployed to achieve high availability, with one Cisco Nexus 1010 used as the primary appliance, and the second Cisco Nexus 1010 used as the secondary appliance. The two appliances will run in an active-standby setup to offer high availability from both the management and deployment sides. Figure 2-6 shows how high availability is built into the Cisco Nexus 1010 Manager.

Figure 2-6 Nexus 1010 Manager High Availability



If one Cisco Nexus 1010 were to fail, management would automatically failover to the other Cisco Nexus 1010 without disruption of traffic or operations. For two Cisco Nexus 1010 appliances to form a high-availability pairing, the control VLAN and domain ID of both Cisco Nexus 1010 appliances must match.

Another high-availability feature built into the Cisco Nexus 1010 is the capability of the Cisco Nexus 1010 Manager to automatically distribute the placement of the active VSMs across the two appliances. This feature helps balance the distribution of traffic and reduces the potential fault domain.

VSM High Availability

High availability is also configured for the redundant virtual services blades that are created on the Cisco Nexus 1010.

Not all virtual services blades are active on the active Cisco Nexus 1010. As long as the active and standby Cisco Nexus 1010 appliances are connected, access through a serial connection is maintained to any virtual service. When one Cisco Nexus 1010 fails, the remaining Cisco Nexus 1010 becomes active and all virtual services in the standby state on that Cisco Nexus 1010 become active on their own.

A virtual service can be removed completely from both redundant Cisco Nexus 1010 appliances, or from only one. If one of a redundant pair of virtual services becomes unusable, it can be removed from just the Cisco Nexus 1010 on which it resides. This feature aids recovery by preserving the remaining virtual service in the pair. Removal of just the failed service may be necessary if a new instance of the service must be provisioned.

You should create redundant VSMs on the Cisco Nexus 1010 with the Cisco Nexus 1000V Series software image. The current version is bundled as an ISO image and included in the Cisco Nexus 1010 bootflash repository folder. The image is copied to a new VSM service when the VSM is created. After the first VSM is created, that software image can be used to create additional VSMs. Upgrading VSMs to a new release of the Cisco Nexus 1000V Series is available as needed.

For more information about VSM high availability, see the Cisco Nexus 1000V High Availability and Redundancy Configuration Guide, Release 4.0(4)SV1(3).

http://www.cisco.com/en/US/partner/docs/switches/datacenter/nexus1000/sw/4_0_4_s_v_1_3/high_availability/configuration/guide/n1000v_ha_preface.html.

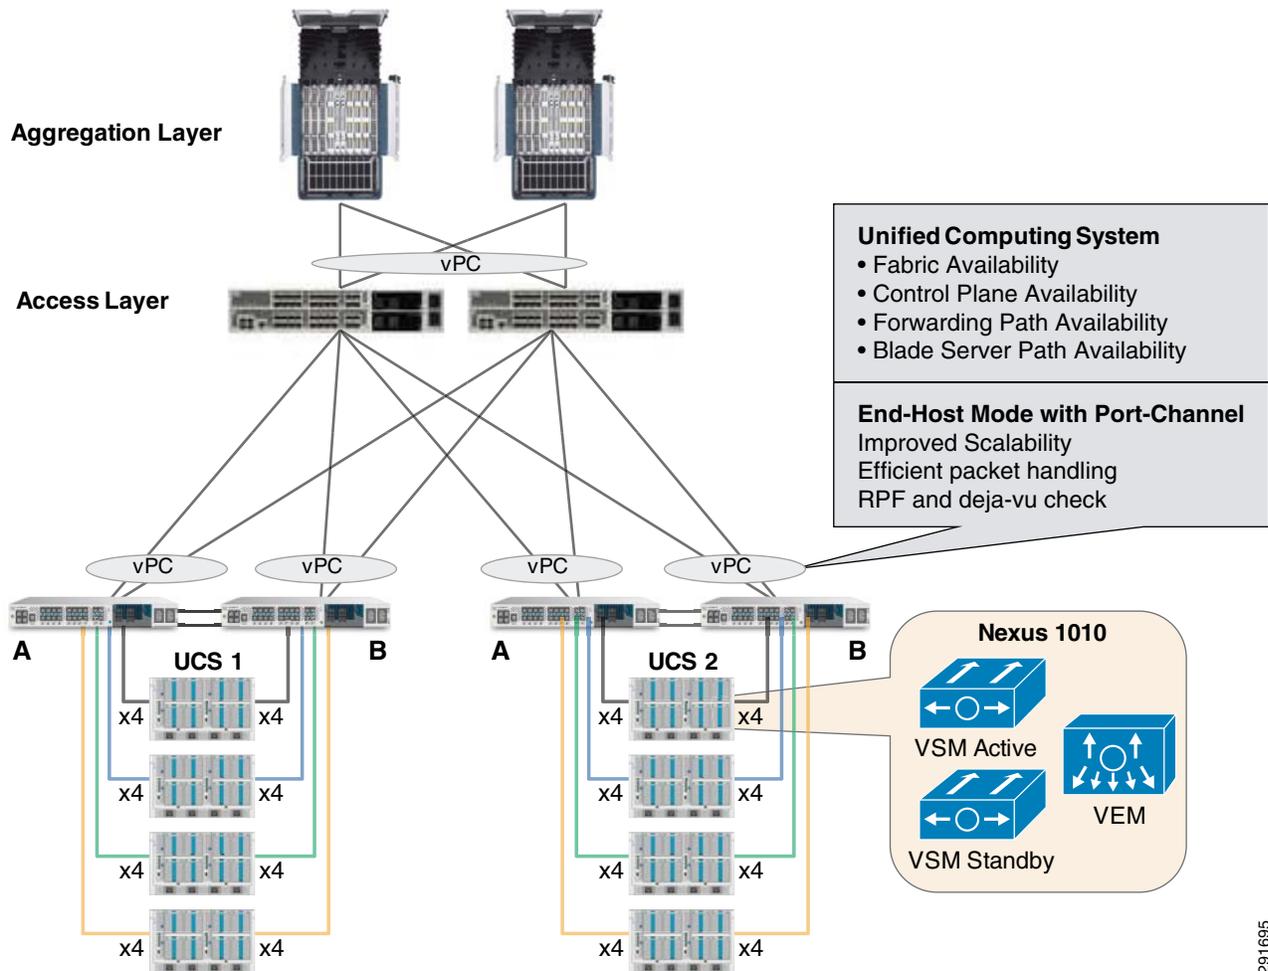
To provide high availability at the virtual access layer, the Cisco VMDC solution relies on the following features:

- Cisco Nexus 1010 offers availability features for large-scale networking. Within a single appliance Cisco Nexus 1010 offers process-level availability conferred by the modular nature of NX-OS, as well as Virtual Service Blade availability features such as restart-on-failure. Cisco Nexus 1000V switch VSM active/standby high availability is fully supported on Cisco Nexus 1010.
- Deploying dual Cisco Nexus 1010 appliances in a high availability cluster provides active/standby failover of Cisco Nexus 1010 Manager and Virtual Service Blades.
- Always deploy the Cisco Nexus 1000V Series VSM (virtual supervisor module) in pairs, where one VSM is defined as the primary module and the other as the secondary. The two VSMs run as an active-standby pair, similar to supervisors in a physical chassis, and provide high availability switch management. The Cisco Nexus 1000V Series VSM is not in the data path so even if both VSMs are powered down, the Virtual Ethernet Module (VEM) is not affected and continues to forward traffic.
- Virtual Port Channels (vPC)—The virtual PortChannel (vPC) feature allows one end of a PortChannel to be split across a pair of Cisco Nexus 7000 Series Switches. vPC provides Layer 2 multipathing through the elimination of Spanning Tree Protocol blocked ports in dual-homed connections. vPC enables fully used bisectional bandwidth and simplified Layer 2 logical topologies without the need to change the existing management and deployment models.

Compute Availability

The Cisco VMDC 2.1 solution relies on the Cisco UCS at the compute layer. The availability of the Cisco UCS fabric is depicted in [Figure 2-7](#).

Figure 2-7 Compute Redundancy in a VMDC 2.1 Pod



The UCS system provides redundancy at every level:

- **Fabric Availability** The UCS provides two independent fabric paths, A and B. In this design, the fabric failover is handled by the Nexus 1000V so this Cisco UCS feature is not used.
- **Control Plane Availability.** The UCS 6100 is enabled in active/standby mode for the control plane (UCS Manager) that manages the entire UCS system.
- **Forwarding Path Availability** It is recommended that each fabric interconnects (UCS 6100) be configured in end-host mode. Uplinks from each UCS 6100 are connected to a Nexus 5000 as port channels with LACP "active-active" mode. This port channel configuration is a best practice recommendation that provides scalability as well as reduces the CPU load when performing RPF and Déjà vu check on packets as there are fewer logical interfaces to process.
- **Blade Server Path Availability** Each blade server is enabled with a Cisco VIC adapter (M81KR - Converged Network Adaptor (CNA) that provides 10 Gbps connectivity to each fabric in the UCS 5108 chassis.

Storage Availability

In the storage layer, the design is consistent with the high availability model implemented at other layers in the infrastructure, which include physical and path redundancy.

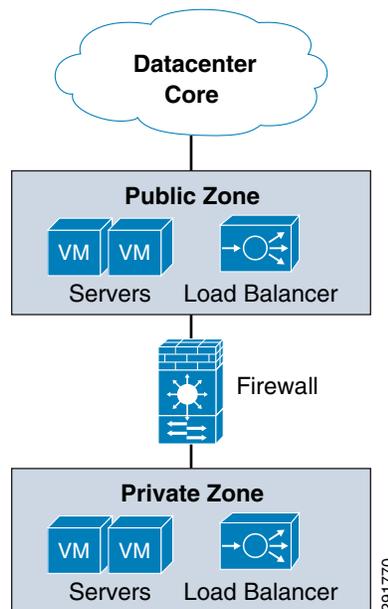
Virtualized Multi-Tenancy

Traditionally, a dedicated infrastructure would be deployed for each tenant that it hosted. This approach, while viable for a multi-tenant deployment model, does not scale well because of cost, complexity to manage, and inefficient use of resources. Deploying multiple tenants in a common infrastructure yields more efficient resource use and lower costs. However, each tenant may require path isolation for security and privacy from others sharing the common infrastructure. Therefore, logical separation or virtualization is a fundamental building block for multi-tenant environments. Virtualization at the various levels in the VMDC 2.1 architecture provides logical separation in the network, compute, and storage resources.

Flexible Tenant Model

In VMDC 2.1 a tenant can be defined as referenced as an external partner or subsidiary or an internal department or business unit, such as engineering or human resources. The basic tenant container is a two tier virtual private datacenter model which contains a public server farm (public zone), a firewall protected private server farm (private zone), and load balancing services available in each zone. The tenant container is designed a flexible model that can be adapted to fit any number of tenant specific requirements. [Figure 2-8](#) shows the basic tenant construct.

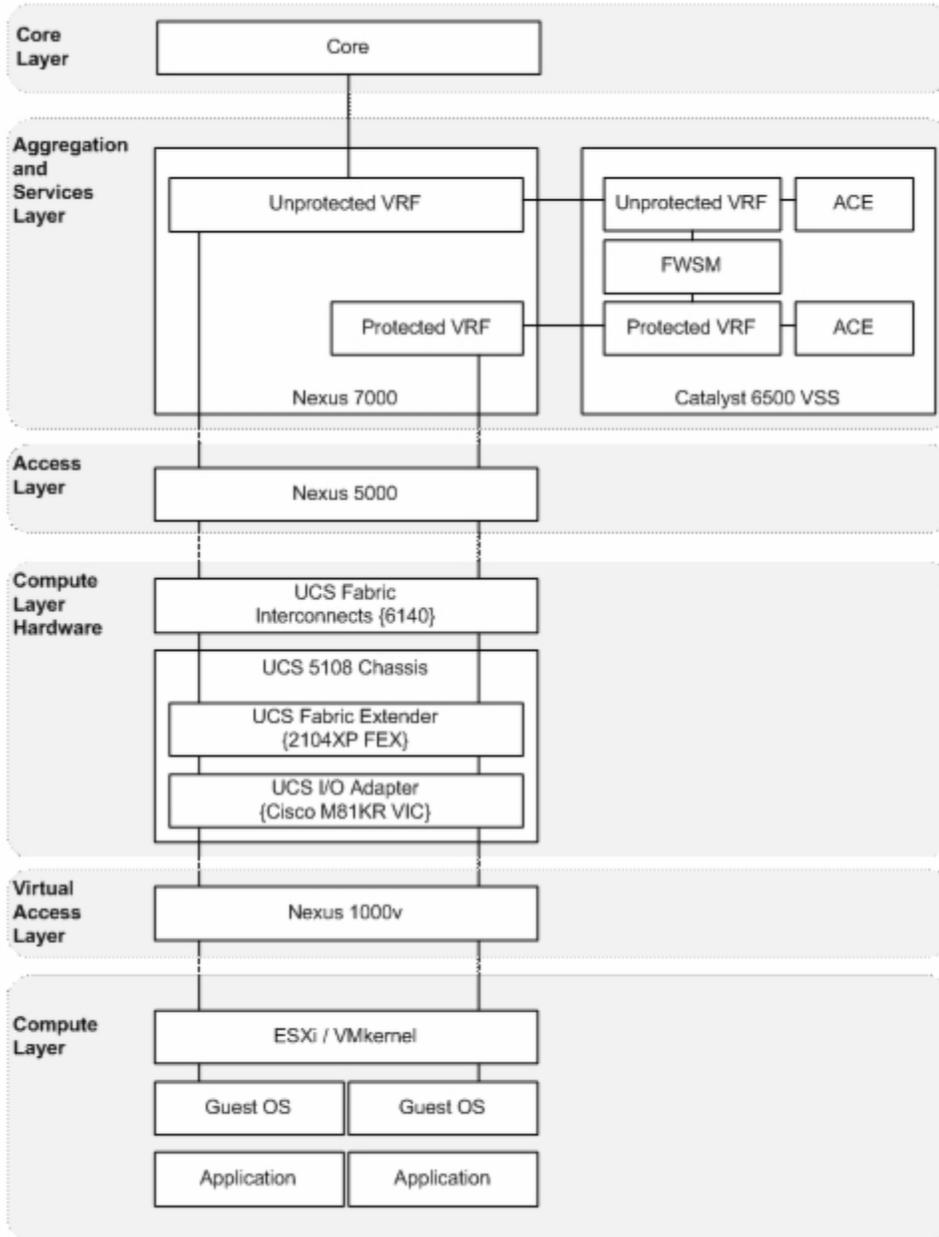
Figure 2-8 VMDC 2.1 Two Tier Virtual Private Data Center Tenant Model



The tenant virtual private datacenter is built using a combination of Layer 3 VRFs and Layer 2 VLANs to provide logical path isolation in the network. Each tenant's virtual datacenter is built with a unique pair of VRFs on the aggregation and services switches representing the public and private zones. The

VRFs contain per tenant routing information which is exchanged via OSPF. Multiple tenant VLANs in the Layer 2 domain of each zone are mapped to the corresponding VRFs. Figure 2-9 shows a block diagram with a single tenant construct as it is overlaid on the physical pod topology.

Figure 2-9 VMDC 2.1 Block Diagram

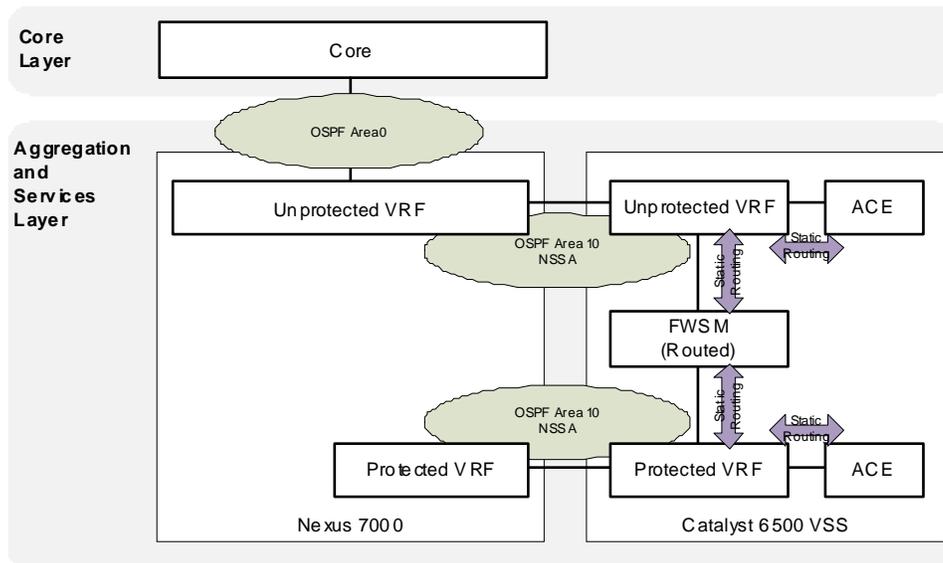


The VMDC 2.1 two tier virtual datacenter uses OSPF for each tenant as the interior gateway protocol. The remainder of the routing information is provided via static routes which are redistributed into OSPF at the Autonomous System Border Router (ASBR).

Not-so-stubby areas (NSSAs) are an extension of OSPF stub areas. Stub areas prevent the flooding of external link-state advertisements (LSAs) into NSSAs, relying instead on default routes to external destinations. NSSAs are more flexible than stub areas in that a NSSA can import external routes into the OSPF routing domain.

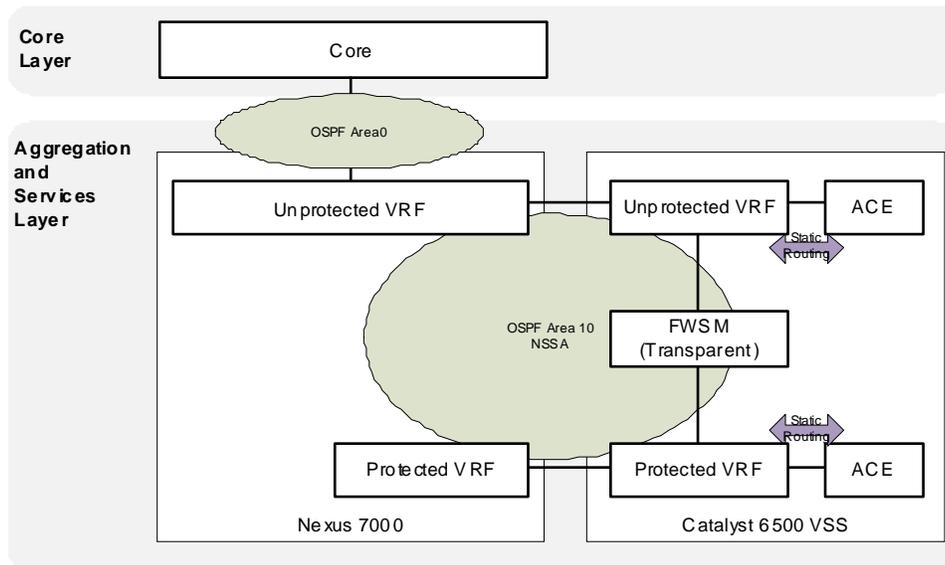
If the FWSM context is deployed in routed mode (recommended as the most flexible option) the public becomes a true NSSA with the connection to Area 0 and the private OSPF area is almost effectively a totally NSSA area given there is no connection to area 0 and a default static route is used to exit to the public zone. In this scenario there are 2 separate routing domains connected via static routes on the FWSM (Figure 2-10).

Figure 2-10 Tenant Routing with FWSM in Routed Mode



If the FWSM context is deployed in transparent mode the public and private interfaces form an OSPF adjacency and the OSPF NSSA is extended through the FWSM forming a single routing domain. In this case all routing information will be populated in both tenant zones (Figure 2-11).

Figure 2-11 Tenant Routing with FWSM in Transparent Mode



Several additional tenant models can be constructed using the basic building blocks of the two tier architecture.

- Multiple ACE contexts
- Single Tier Virtual Private Datacenter
- Two Tier Virtual Private Datacenter with single firewall and multiple Private Zones
- Two Tier Virtual Private Datacenter with multiple firewalls and multiple Private Zones
- Three Tier Virtual Private Datacenter

Network and Services Separation

The Cisco VMDC 2.1 solution assumes there is no need to interconnect between tenants so server-to-server communication between tenants is not required. Based on this assumption the tenant container requires path isolation and/or logical resource separation at each of the network layers in the architecture, Layer 3 (routed), Services (firewall and server load balancing), and Layer 2 (switched) domains. Therefore, the three types of domains must be virtualized and the virtual domains must be mapped to each other to keep traffic segmented. This mapping combines device virtualization with data path virtualization at the different layers in the network.

- **Aggregation Layer**—Layer 3 separation (VRF-Lite) and Layer 2 separation (VLAN)
- **Services Layer**—Layer 3 separation (VRF-Lite), Layer 2 separation (VLAN), and Virtual Device Contexts
- **Access Layer**—Layer 2 separation (VLAN)
- **Virtual Access Layer**—Layer 2 separation (VLAN)

Compute Separation

Virtualization in the compute layer introduces new challenges and concerns. A physical host now contains multiple logical servers (Virtual Machines) requiring some number of policies to be applied at the VM level. Also, new technologies, such as vMotion, introduced VM mobility within a cluster, where policies follow VMs as they are moved across switch ports and between hypervisor hosts.

To provide traffic isolation for virtual machines, the VMDC solution emphasizes the following techniques:

- **Port Profiles** Port profiles enable VLAN-based separation. Using features found in the Nexus 1000V, you create port profiles and apply them to virtual machine NICs via the VMware vCenter. Each port profile is a policy that can be applied to the VM. The policy settings include VLAN, uplink pinning, security, and policy information.
- **Virtual Adapters** Cisco UCS M81KR Virtual Interface Card (VIC) is a network interface consolidation solution. Traditionally, each VMware ESX server has multiple LAN and SAN interfaces to separate vMotion, service console, NFS, backup, and VM data. In this model, the ESXi host requires 10 adapters. Using the Cisco VIC, distinct virtual adapters are created for each traffic flow type using a single, two-port adapter.
- **VLAN Separation** Using the Cisco VIC features, you can create virtual adapters and map them to unique virtual machines and VMkernel interfaces through the hypervisor. In a multi-tenant scenario where distinct tenants reside on the same physical server and transmit their data over a shared physical interface, the infrastructure cannot isolate the tenant production data. However, Cisco VIC combined with VN-Link technology can isolate this data via VLAN-based separation. VLAN separation is accomplished when virtual adapters (up to 128) are mapped to specific virtual machines and VMkernel interfaces.

Storage Separation

To extend secure separation to the storage layer, VMDC 2.1 uses isolation mechanisms available in either SAN or NAS environments. Tenant separation can extend through the switches and into the storage arrays.

Storage Area Network (SAN)

The VMDC 2.1 architecture was validated using the Cisco MDS 9513 and EMC VMAX for Block Storage. This allows for Fiber Channel (FC) access separation at the switch port level (VSAN), logical path access separation via IVR, path level separation using WWN/Device Hard Zoning, and at the virtual media level inside the Storage Array (LUN Masking and Mapping).

Network Attached Storage (NAS)

The VMDC 2.1 architecture was validated using NetApp for NFS storage, which enables virtualized storage space such that each tenant (application or user) can be separated with use of IP spaces and VLANs mapped to network layer separation.

Performance and Scalability

Performance is a measure of the speed at which a system works. Scalability is the ability to grow in size or complexity without showing negative effects. Problems in either area may expose the enterprise to operating inefficiencies and potential failures of critical business components. Testing, monitoring, and tuning the environment ensures optimal performance and user satisfaction.

There are multiple dimensions that form overall datacenter scalability, including physical capacity, logical capacity, topology, and functionality. All of these aspects combined drive the overall data center footprint. The following list illustrates some key variables that contribute to the scaling and performance design considerations in a given datacenter:

- Pods
- Port Density
- Access switches
- Access ports
- Port channels
- vPCs
- VLANs
- MAC addresses
- STP logical ports
- Interfaces
- Routing adjacencies
- HSRP groups
- Routes
- Multicast routes
- CPUs
- Memory
- Oversubscription Ratios

Validated Scale

The VMDC 2.1 architecture scalability scope is derived from applying the above principles and approaches to establish a validated design. Although VMDC 2.1 provides a baseline reference for scalability, each implementation will vary in ways that change operational scalability at each layer.

Table 2-1 Architecture Scalability Scope

Device	Feature	Detail	32 Tenant
Nexus 7010	VRF	Each tenant requires 2 VRFs	64
	VLAN	Total Tenant VLANs	192
	MAC	Total MAC addresses	13000
	RIB	Routes in public zone	1312
		Routes in private zone	640

Table 2-1 Architecture Scalability Scope (continued)

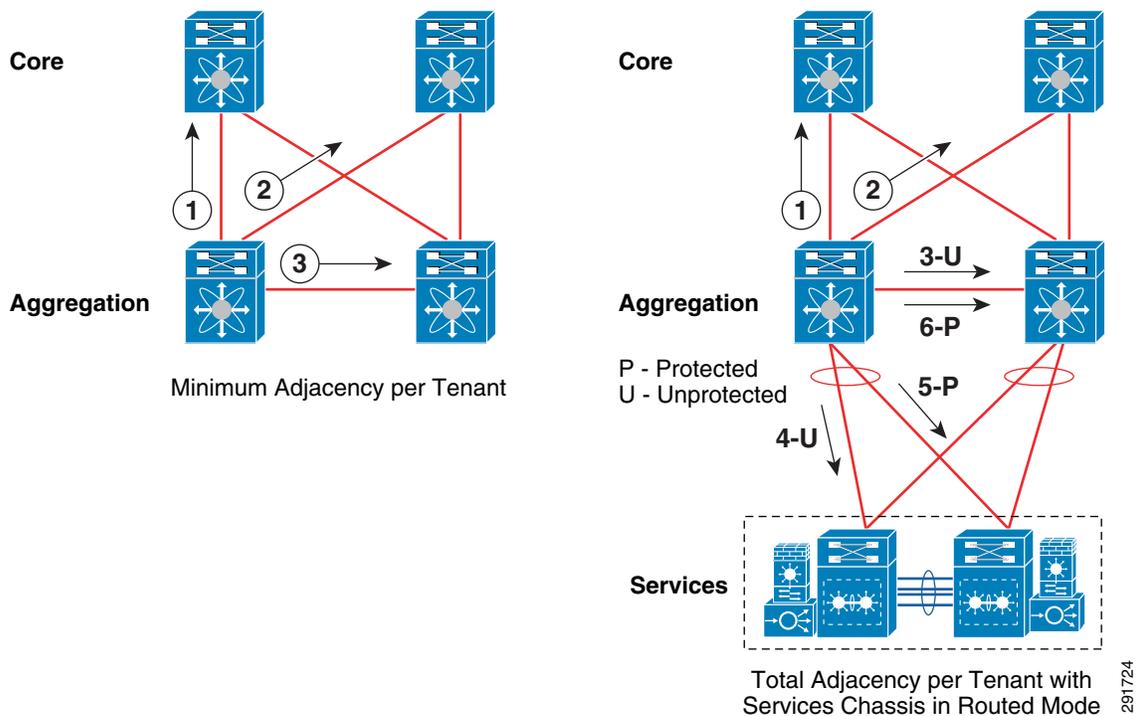
Device	Feature	Detail	32 Tenant
	OSPF	AFI	64
		Neighbor adjacencies in public zone	128
		Neighbor adjacencies in private zone	64
	Multicast	PIM adjacencies public zone only	24
		Total mroutes public zone only	128
		Total number of (*,G) routes public zone only	64
		Total number of (S,G) routes public zone only	64
Catalyst 6509	VRFs	Each tenant requires 2 VRFs	64
	VLAN	2 ACE VLANs / 2 FWSM VLANs	128
	RIB	Routes in public zone	832
		Routes in private zone	416
	OSPF	Processes	64
		Neighbor adjacencies	128
ACE	Context	2 ACE contexts per Tenant	64
	VIPs	4 VIPs per ACE context 32T	128
FWSM	Context	1 FW context per tenant	32
Nexus 5020	VLANs	3 Server VLANs per VRF	192
		Management VLANs	8
		NFS vFiler VLANs	32
	MAC	Total MAC addresses	13000
Nexus 61xx	VLANs	3 Server VLANs per VRF	192
		Management	8
		NFS vFiler VLANs	32
	MAC	Total MAC addresses	13000
Nexus 1000v	VLANs	3 Server VLANs per VRF	192
		Management	8
		NFS vFiler VLANs	32
	MAC	Total MAC addresses	13000
UCS	VM	Test VMs	128
		VMs per blade server ratio	4:1

Understanding Tenant Scalability

Scalability at the aggregation pod level is largely referenced by the number of tenants that can be accommodated. Each tenant is defined by a one or more virtual routing and forwarding (VRF) instances at the aggregation layer and hence requires a per-tenant control plane adjacency formed by the underlying routing protocol. The number of VRFs is a consideration but another key metric, as mentioned in the previous section, is the number of routing adjacencies supported on the aggregation platform.

The VMDC 2.1 design uses OSPF in both the public and private zones of the virtual datacenter tenant construct. This design requires a minimum of 3 adjacencies per tenant with the datacenter core, assuming the core is using OSPF. The VVMDC 2.1 design was implemented with 6 routing adjacencies per tenant on each aggregation switch in the pod as shown in Figure 2-12.

Figure 2-12 Per Tenant OSPF Adjacency Requirements



The current number of adjacencies supported per aggregation layer is 300 (specified at: http://cco.cisco.com/en/US/docs/switches/datacenter/sw/5_x/nx-os/unicast/configuration/guide/13_limits.html).

Hence, the maximum number of tenants that can be deployed with six adjacencies is 50. The VMDC 2.1 architecture was validated with 32 tenants, the same as the Cisco VMDC 2.0 Compact Pod Cisco Validated Design (CVD).

The number of tenants per aggregation layer can be improved using multiple methods. If the DSN connectivity is managed by static routing, then you can reduce the adjacency requirement by three per tenant. This reduction enables the same topology to support up to 100 tenants using OSPF as a routing protocol. In addition, BGP can be used as an alternative protocol to support up to 150 tenants using dynamic routing to the service chassis, or up to 250 tenants using static routes to the services chassis (Table 2-2).

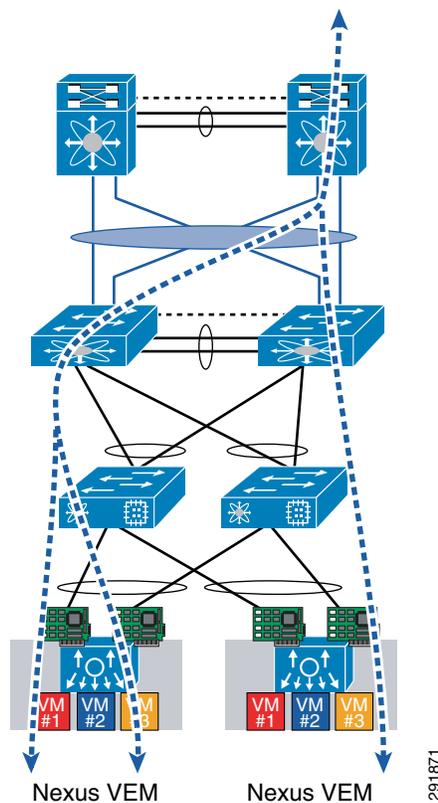
Table 2-2 Tenant Quantity Configurations

Number of Tenants	Core Facing Routing Configuration	Services Type	Services Facing Routing Configuration
50	OSPF	Service Chassis	OSPF
100	OSPF	Service Chassis or Appliance	Static
150	BGP	Services Chassis	BGP
250	BGP	Service Chassis or Appliance	Static

Per Tenant Multicast Support

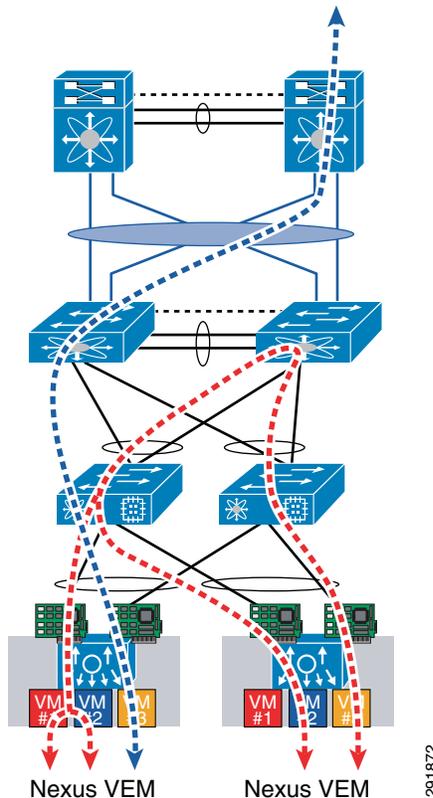
Multicast is supported in the VMDC architecture in two different implementations depending on the application requirements. The implementation does not include multicast configurations to support multicast flows through the services layer (DSN).

Multicast in the Public Zone Front-End VLANs supports end user applications that take advantage of multicast technologies including corporate communications, distance learning, and distribution of software, stock quotes, and news. This is shown in [Figure 2-13](#) and is accomplished using PIM and IGMP Snooping.

Figure 2-13 Multicast Flows in Public Zone Front-End VLANs

Multicast in the Public and Private Zone Back-End VLANs supports clustering technologies like Oracle Rack, Microsoft SQL Cluster, REHL cluster, and VERITAS cluster. This is shown in [Figure 2-14](#) and accomplished using IGMP Snooping and IGMP Snooping Querier.

Figure 2-14 Multicast Flows in Public and Private Zone Back-End VLANs



The multicast implementation in Cisco VMDC 2.1 is structured around the following features and configurations at specific locations in the topology:

- Core (Per Tenant)
 - PIM-SM (sparse mode)
 - Anycast RP using MSDP
- Per Tenant Public Zone - Intra and Inter VLAN
 - Static RP
 - PIM-SM (sparse mode) configured on the Aggregation Layer Nexus 7000 for Front End VLANs
 - IGMP Querier deployed at Access Layer Nexus 5000 for Back End VLANs
 - IGMP Snooping
- Per Tenant Private Zone - Intra VLAN only
 - IGMP Querier deployed at Access Layer Nexus 5000 for Back End VLANs
 - IGMP Snooping

Anycast RP for PIM-SM

The RP is a critical function for PIM-SM deployments. RP redundancy is always recommended. The best form of redundancy for PIM-SM is Anycast RP which is described in the document:

Anycast RP:

http://www.cisco.com/en/US/docs/ios/solutions_docs/ip_multicast/White_papers/anycast.html

VMDC 2.1 does not specify an RP redundancy design or specific RP location, the assumption is that the RP would be either at the core layer, somewhere else within the datacenter, or further out on the enterprise network. The reference of Anycast RP is for example purposes only as this is a typical method for ensuring RP reachability.

PIM Sparse Mode (PIM-SM)

PIM-SM uses a pull model to deliver multicast traffic. Only network segments with active receivers that have explicitly requested the data will receive the traffic.

PIM-SM distributes information about active sources by forwarding data packets on the shared tree. Because PIM-SM uses shared trees (at least, initially), it requires the use of a rendezvous point (RP). The RP must be administratively configured for each tenant requiring multicast support. For each tenant public zone (VRF) in the system, independent multicast system resources are maintained, including the MRIB.

IGMP Snooping

Every Cisco switch in the VMDC 2.1 solution supports IGMP snooping. IGMP snooping should always be enabled if you are running IP Multicast. Some platform and switch software combinations may not have IGMP snooping enabled by default. Make sure IGMP snooping is enabled before running any multicast streams.

IGMP snooping is an IP Multicast constraining mechanism that runs on a Layer 2 LAN switch. Without IGMP snooping enabled, all multicast traffic will be forwarded to all hosts connected to the switch. IGMP snooping will insure that only hosts that are interested in the data stream will receive it.

The Back End VLANs in both the public and private zones run multicast in a contained environment and not have it forwarded to the rest of the network. On these VLANs, PIM is not enabled on the routers so there is no IGMP querier elected.

IGMP Snooping Querier

For Back-End VLANs where PIM is not enabled on an interface because the multicast traffic does not need to be routed, you must configure an IGMP snooping querier to send membership queries. You define the querier in a VLAN that contains multicast sources and receivers but no other active querier.

When an IGMP snooping querier is enabled, it sends out periodic IGMP queries that trigger IGMP report messages from hosts that want to receive IP multicast traffic. IGMP snooping listens to these IGMP reports to establish appropriate forwarding.

For additional multicast information see the following links:

IP Multicast Best Practices

http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6552/ps6592/whitepaper_c11-474791.html

IP Multicast White Papers

http://www.cisco.com/en/US/products/ps6552/prod_white_papers_list.html

Jumbo Frame Support

It has been shown that a server can boost its performance and reduce its CPU utilization significantly by using jumbo frames during data transfers. For example, during the tests a server configured with two GbE NICs was shown to have increased its network throughput and decreased its CPU utilization by 44% when using 9 KB frames instead of standard 1518-byte frames.

This type of performance improvement is only possible when long data transfers are performed; for example, in applications such as:

- Server to Server communication (e.g., NFS transactions, vMotion, etc.)
- Server clustering
- High-speed data backups

In these scenarios jumbo frames are becoming a standard for high-speed transfers over the Ethernet medium.

A jumbo frame is basically anything bigger than 1522 bytes, with a common size of 9000 bytes, which is exactly six times the size of a standard Ethernet frame. With Ethernet headers, a 9k byte jumbo frame would be 9014-9022 bytes. This makes it large enough to encapsulate a standard NFS (network file system) data block of 8192 bytes, yet not large enough to exceed the 12,000 byte limit of Ethernet's error checking CRC (cyclic redundancy check) algorithm.

When designing for jumbo MTU services, consider the following factors:

- Understand jumbo frames are not a standard, so testing prior to implementation is critical
- Jumbos must be supported end to end for any application benefit
- Do not use anything above the common size of 9000 bytes
- Transit links can be configured with higher values to allow for any additional header overhead
- If available, use TCP Offload Engines with Jumbo Frames on server based NICs

In the VMDC 2.1 design, all paths below the Nexus 7000 were all enabled for jumbo frames. The full front end paths, to the services chassis and exiting the aggregation towards the core, campus, or WAN, were not enabled but could be accommodated with additional configuration. The links are illustrated in [Figure 2-15](#).

Figure 2-15 Jumbo Frame Enabled Links

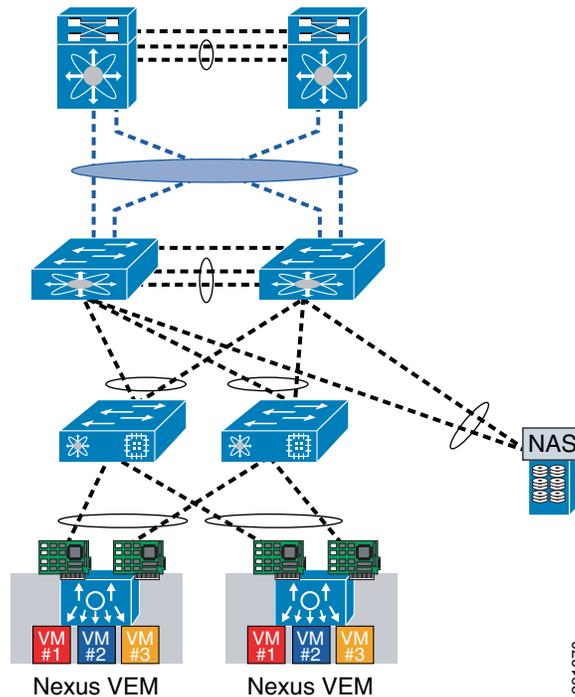


Table 2-3 lists maximum configurable jumbo MTU values for devices used in the VMDC 2.1 solution. All devices were configured with their maximum supported values.

Table 2-3 Maximum Jumbo MTU Sizes

Platform	Maximum Value
Nexus 7000	9216
Nexus 5020/5548	9216
UCS B Series	9216
Nexus 1000v	9000
ESXi 4.1 U1	9000
NetApp FAS6080	9000

Platform Specific Limits

The scalability numbers supported per device typically improve with each new release of hardware or software. Refer to the latest release notes and configuration maximums for each component in the VMDC 2.1 architecture.

Service Assurance

Service assurance is generally defined as the application of policies and processes ensuring that services offered over networks meet a pre-defined service quality level for an optimal subscriber experience. The practice of service assurance enables providers to control traffic flows and identify faults and resolve

those issues in a timely manner so as to minimize service downtime. The practice also includes policies and processes to proactively diagnose and resolve service quality degradations or device malfunctions before subscribers are impacted.

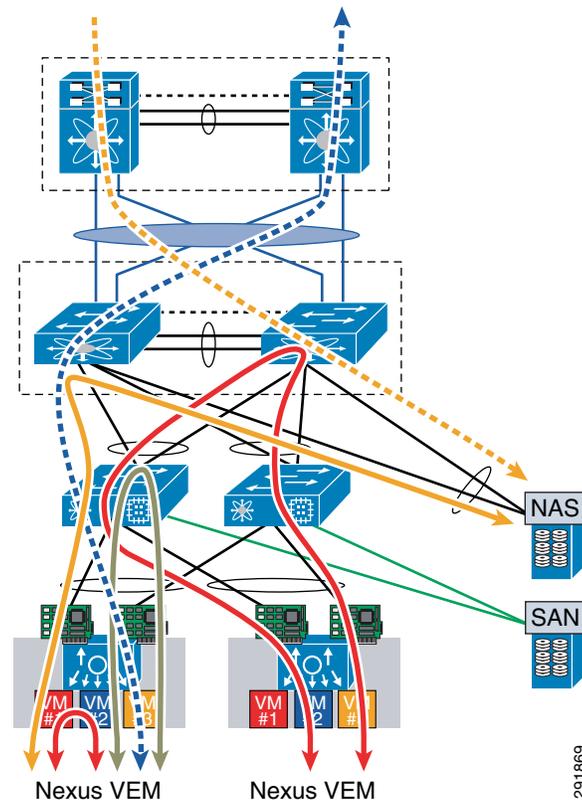
In VMDC 2.1 network service assurance encompasses the following concepts:

- Traffic Engineering
- Quality of Service (QoS) framework
- Network Analysis

Traffic Engineering

Traffic engineering is a method of optimizing the performance of a network by dynamically analyzing, predicting and regulating the behavior of data transmitted over that network. [Figure 2-16](#) shows some typical traffic patterns seen within the VMDC 2.1 pod.

Figure 2-16 Data Center Traffic Flows



PortChannels are typically deployed for redundancy and load sharing capabilities. Since the Cisco Nexus 1000V Series is an end-host switch, the network administrator can use a different approach than can be used on a physical switch, implementing a PortChannel mechanism in either of two modes:

- Standard PortChannel: The PortChannel is configured on both the Cisco Nexus 1000V Series and the upstream switches

- Special PortChannel: The PortChannel is configured only on the Cisco Nexus 1000V Series, with no need to configure anything upstream. There are two options available here, MAC Pinning and vPC Host Mode.

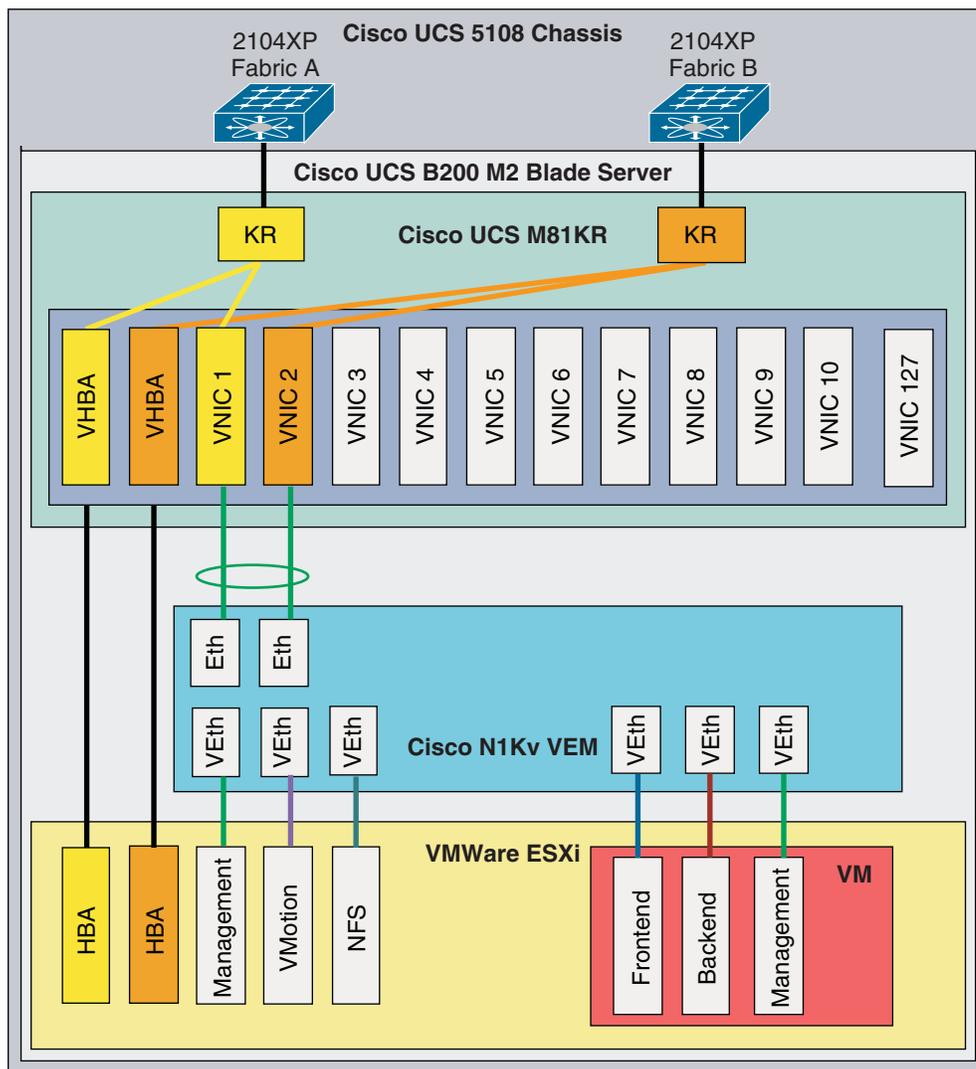
Regardless of the mode, PortChannels are managed using the standard PortChannel CLI construct, but each mode behaves differently.

For more information on the Nexus 1000v Port-Channel configurations follow this link:

http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9902/guide_c07-556626.html#wp9000299

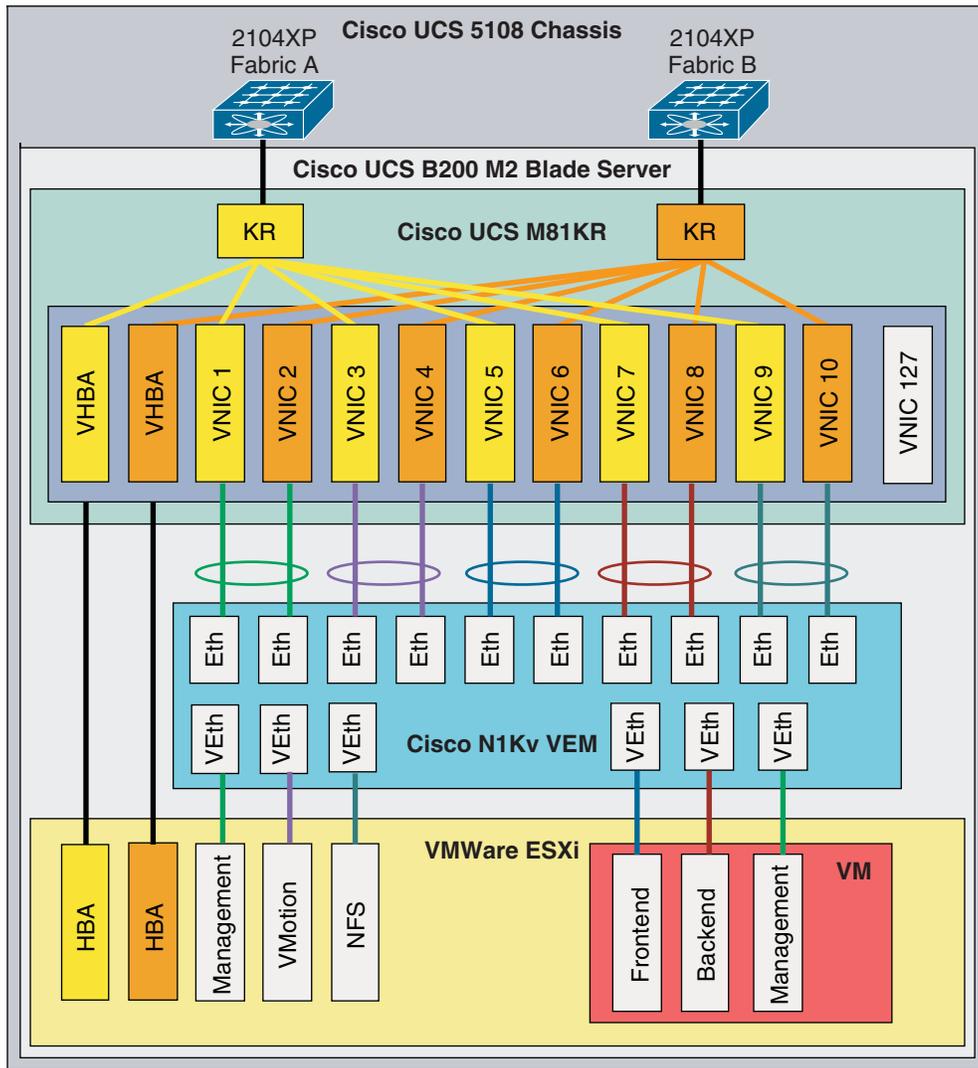
The VMDC 2.1 virtual access layer design focuses on the Special PortChannel option with vPC-Host Mode and then uses MAC Pinning to select specific links from the port channel. The design also expands on a typical single uplink portchannel setup to use a 5 PortChannel uplink configuration on the Nexus 1000v. This configuration allows a more granular approach to uplink management on the Nexus 1000v. The different options are shown in Figure 2-17 and Figure 2-18.

Figure 2-17 Nexus 1000v single Uplink PortChannel Model



291867

Figure 2-18 Nexus 1000v 5 uplink PortChannel Model



291734

Traffic engineering can be performed selectively by configuring the Nexus 1000v to select the target uplink with a manual configuration (static pinning) instead of the default. For example, front-end traffic that contains many diversified flows can use both members (fabrics) of the port-channel. On the other hand, back-end traffic, which has more diversity in terms of bandwidth/response time usage (VM-to-VM - inter fabric traffic flows, vMotion, backup, and so forth) may benefit by selecting a path such that it allows VM-to-VM traffic to remain within a single fabric where the Fabric Interconnect switches the traffic locally.

Table 2-4 Traffic Classification Example for MAC Pinning

Traffic Type	Classification Category	UCS Fabric	Mac-Pining Option	Rational
Front End Traffic	Tenant Data	Fabric A & B	Automatic	Load Share on all available uplinks, most traffic should be exiting the pod through the Aggregation Nexus 7000
Back End Traffic	Tenant Data	Fabric-A	Manual	Keep most back end traffic local switched on one Fabric Interconnect
vMotion	VMkernel/Control	Fabric-B	Manual	Keep vMotion traffic local switched on one Fabric Interconnect

MAC Pinning

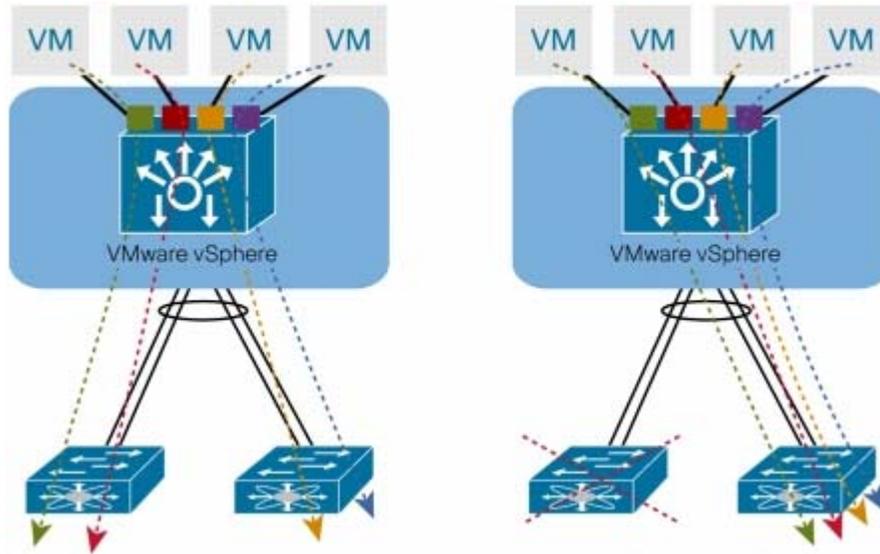
MAC pinning defines all the uplinks coming out of the server as standalone links and pins different MAC addresses to those links in a round-robin fashion. This approach helps ensure that the MAC address of a virtual machine will never be seen on multiple interfaces on the upstream switches. Therefore, no upstream configuration is required to connect the Cisco Nexus 1000V Series VEM to the upstream switches (Figure 2-19).

Furthermore, MAC pinning does not rely on any protocol to distinguish the different upstream switches, making the deployment independent of any hardware or design.

However, this approach does not prevent the Cisco Nexus 1000V Series from constructing a PortChannel on its side, providing the required redundancy in the data center in case of a failure. If a failure occurs, the Cisco Nexus 1000V Series will send a gratuitous Address Resolution Protocol (ARP) packet to alert the upstream switch that the MAC address of the VEM learned on the previous link will now be learned on a different link, enabling failover in less than a second.

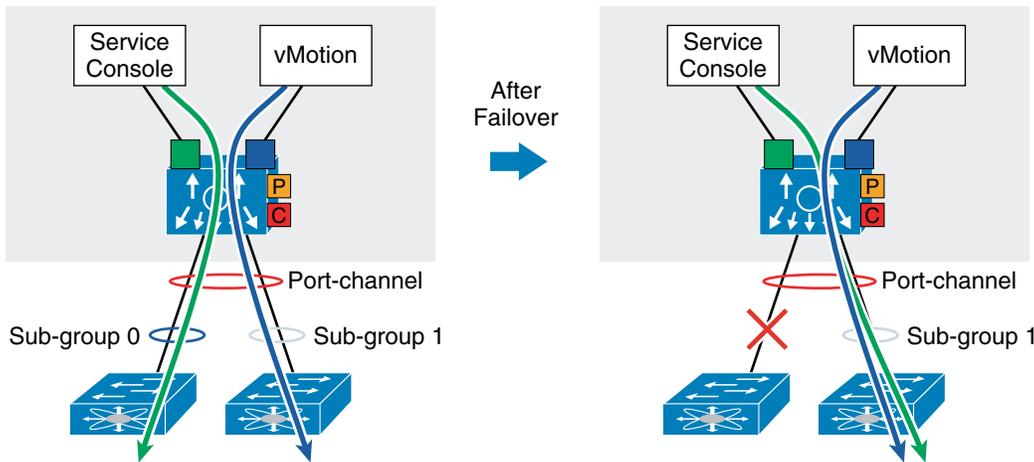
MAC pinning enables consistent and easy deployment of the Cisco Nexus 1000V Series since it does not depend on any physical hardware or any upstream configuration, and it is the preferred method for deploying the Cisco Nexus 1000V Series if the upstream switches cannot be clustered.

Figure 2-19 MAC-Pinning Details



In the case of a fabric failure the Nexus 1000 selects the available remaining fabric to recover the traffic. Figure 2-20 illustrates the fabric failover with sub-group mac-pinning.

Figure 2-20 Mac-Pinning Failover



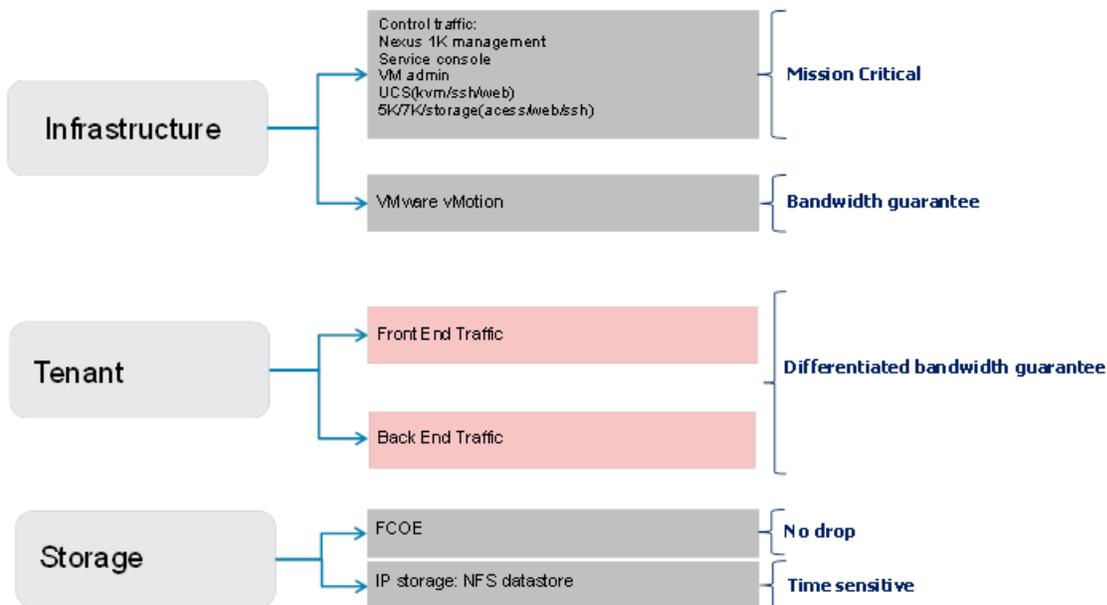
Quality of Service Framework

Quality of Service is a key to service assurance because it enables differentiated treatment of specific traffic flows. This differentiated treatment ensures that in the event of congestion or failure conditions, critical traffic is provided sufficient amount bandwidth to meet throughput requirements.

Traffic Flow Types illustrates the different traffic flow types defined in the VMDC 2.1 architecture. These traffic types are organized in infrastructure, tenant, and storage traffic categories.

- Infrastructure traffic comprises management and control traffic, including VMware service console and vMotion communication. This is typically set to the highest priority in order to maintain administrative communications during periods of instability or high CPU utilization.
- Tenant traffic is differentiated into Front End and Back End Traffic with service levels to accommodate various types of traffic requirements in each category.
- The VMDC 2.1 design incorporates both FC and IP-attached storage. As indicated in [Figure 2-21](#), storage requires two sub-categories, since these traffic types are treated differently throughout the network. FC traffic by definition requires a “no drop” policy, while NFS datastore traffic is sensitive to delay and loss.

Figure 2-21 Traffic Flow Types



To enable differentiated services the following QoS features leveraged in this design are as follows:

- Classification and Marking
- Queuing

Classification and Marking

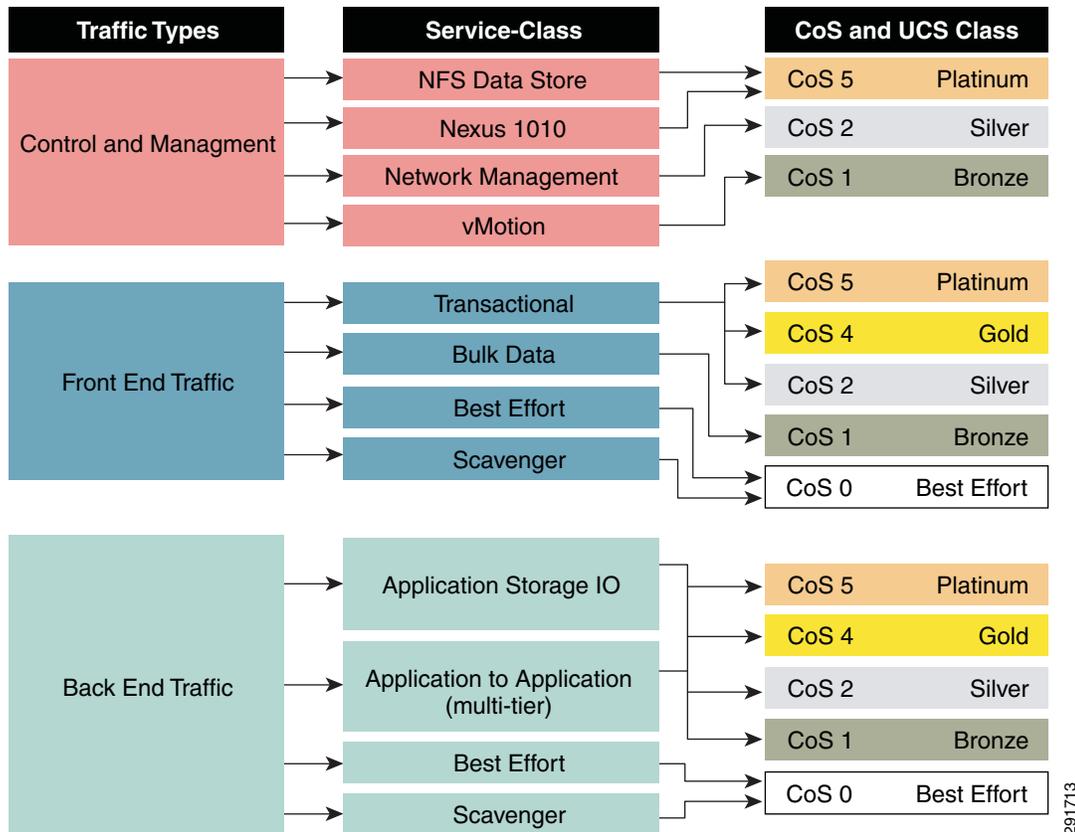
The process of classification is one of inspecting different fields in the Ethernet Layer 2 header, along with fields in the IP header (Layer 3) and the TCP/UDP header (Layer 4), to determine the level of service that should be applied to the frame as it transits the network devices. The process of marking rewrites the COS in the Ethernet header or the Type of Service bits in the IPv4 header if desired.

A key driver around understanding classification and marking in the datacenter is the integration of additional protocols into the data path in the future. In newer datacenter QoS models CoS 3 is reserved for loss-less data (FCoE), which is not specifically addressed in this design guide. However, in the WAN/Campus QoS services model, CoS 3 is used for VOIP signaling and may require remarking at the ingress boundary using the Nexus 7000.

In VMDC 2.1 the assumption is that the DSCP values will not be altered and that only the CoS values will be re-written so that traffic can be queued appropriately within the datacenter devices. The typical trust boundaries are not changed, however there is now the potential need to change CoS markings from what would normally be done with the default DSCP to CoS mapping values on the switches at each layer of the datacenter.

Figure 2-22 shows some example traffic classifications and corresponding CoS traffic markings used for the different traffic types found in a typical VMDC 2.1 architecture deployment.

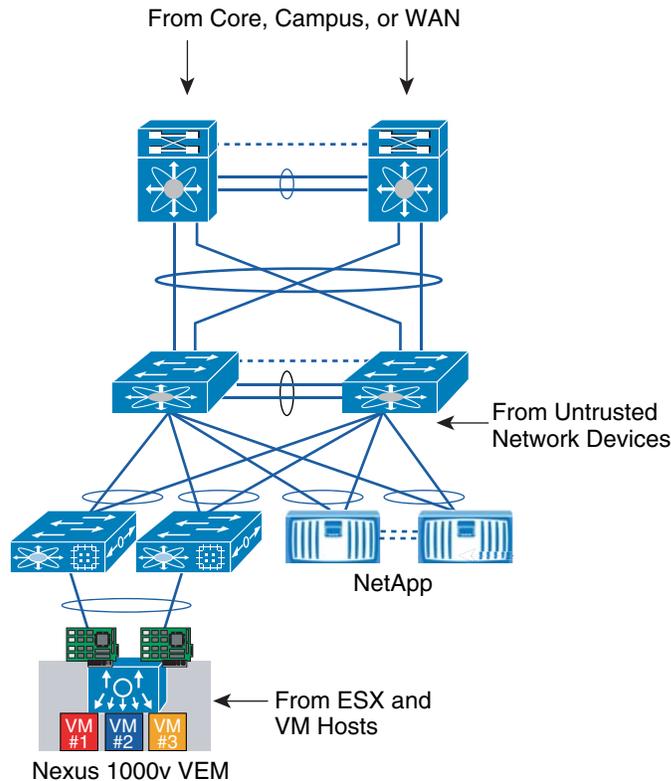
Figure 2-22 Example Traffic Types and CoS Traffic Markings



There are four main places of focus for setting appropriate CoS values:

1. At the aggregation layer, the Nexus 7000 trusts DSCP by default but may need to re-classify the ingress traffic from the Core, Campus, WAN or Services Chassis. The DSCP will be unchanged but a new egress CoS value can be marked as it leaves the Nexus 7000. Traffic egressing the Nexus 7000 to the Core, Campus, WAN may not have the expected correct CoS value but should be correctly classified at the next hop based on the underlying DSCP value.
2. At the services layer, the Catalyst 6500 is configured to trust DSCP, there is no need to re-classify traffic as the ingress CoS values will be ignored and an internal CoS value will be set based on the DSCP to CoS mapping table configuration. The egress CoS value will be set based on the mapping but traffic will be re-classified at ingress to the aggregation layer as mentioned above.
3. At the access layer, the Nexus 5000 only marks traffic for any attached devices that are untrusted or incapable of marking CoS values.
4. At the virtual access layer, the Nexus 1000v marks the DSCP value based on the given enterprise QoS model and marks the CoS value based on data center model.

Figure 2-23 Pod CoS Marking Locations



291712

Queuing

To provide differentiated treatment per defined traffic class in the event of congestion, minimum bandwidth guarantee must be defined in a policy and applied to an interface, sub-interface, or virtual circuit.

Table 2-5 represents the VMDC 2.1 example SLA framework for bandwidth guarantees as configured QoS policies at various layers. The implementation specifics vary due to expected traffic loads as well as differences in connectivity, interface types, and QoS scheduling and queuing capabilities across specific platforms in the infrastructure.

Table 2-5 Example SLA Bandwidth Guarantees

CoS Marking	UCS Traffic Type	UCS 6 Class BW%	N5K 5 Class BW%	N7K 5 Class BW%	C6K 4 Class BW%
5	Platinum	16%	Up to 100%	Up to 100%	Up to 100%
4	Gold	16%	20%	35%	40%
3	FCoE	9%	—	—	
2	Silver	16%	20%	15%	
1	Bronze	16%	20%	15%	5%
0	Default	27%	40%	20%	25%

Network Analysis

The use of network analysis devices is another service readily available in the VMDC 2.1 design. The Cisco Nexus 1000v NAM VSB is integrated with the Nexus 1010 Virtual Services Appliance to provide network and performance visibility into the Nexus 1000V switching deployment. The NAM VSB uses the embedded instrumentation, such as Netflow and Encapsulated Remote SPAN (ERSPAN) on the Nexus 1000V switch as the data source for traffic analysis, application response time, interface statistics, and reporting.

For more information on the Cisco Prime NAM for Nexus 1010 deployment follow the link below:

http://www.cisco.com/en/US/docs/net_mgmt/network_analysis_module_virtual_blade/4.2/install/guide/nexus/nx42_install.html

The following methods of data collection were used in the design:

- NetFlow
- Encapsulated Remote Switched Port Analyzer (ERSPAN)

NetFlow

NetFlow was developed by Cisco to provide better insight into the IP traffic on the network. NetFlow defines flows as records and exports these records to collection devices. NetFlow provides information about the applications in and utilization of the data center network. The NetFlow collector aggregates and assists network administrators and application owners to interpret the performance of the data center environment.

The use of NetFlow is well documented in a traditional network environment, but the Nexus 1000v provides this capability within the virtual network environment. Nexus 1000v supports NetFlow v9 and by default will use the management 0 interface as an export source.



Caution

The use of advanced features such as NetFlow will consume additional resources (i.e., memory and CPU, of your ESX host). It is important to understand these resource dynamics before enabling any advanced features.

Figure 2-24 is an output example that shows the Cisco NetFlow Collector reporting application statistics on the virtual Ethernet interfaces that reside on the Nexus 1000v. The Nexus 1000v may also monitor flows from the physical interfaces associated with the platform and VMkernel interfaces including VMotion traffic as seen in Figure 2-25.

Figure 2-24 Cisco NetFlow Collector Application Statistics

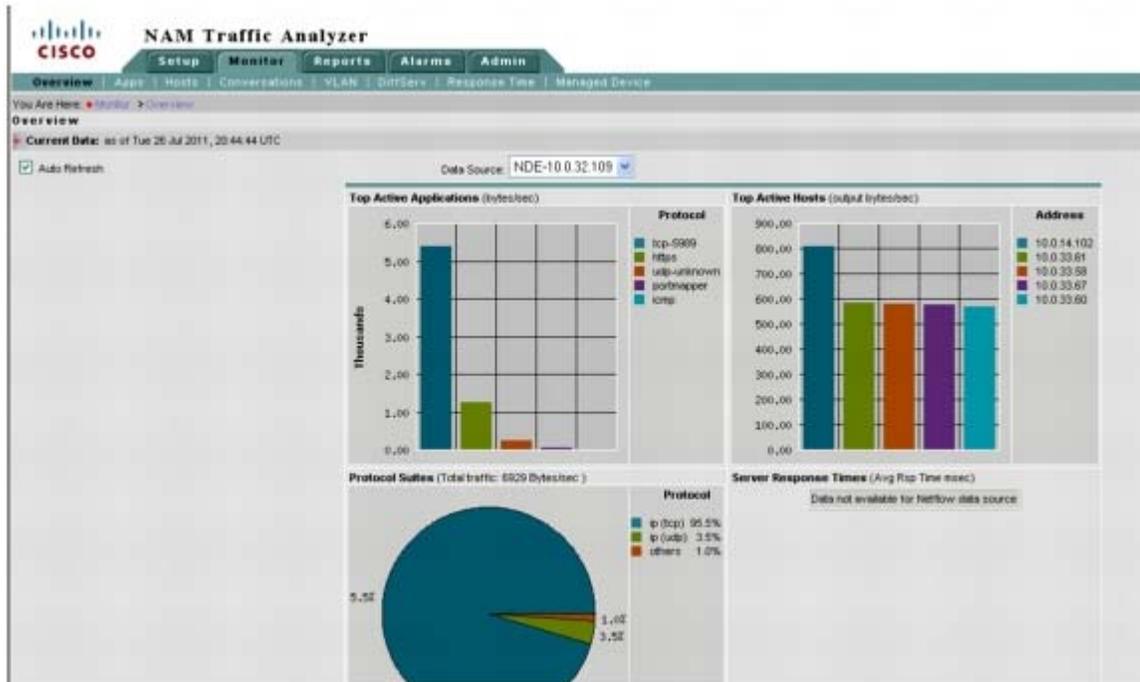
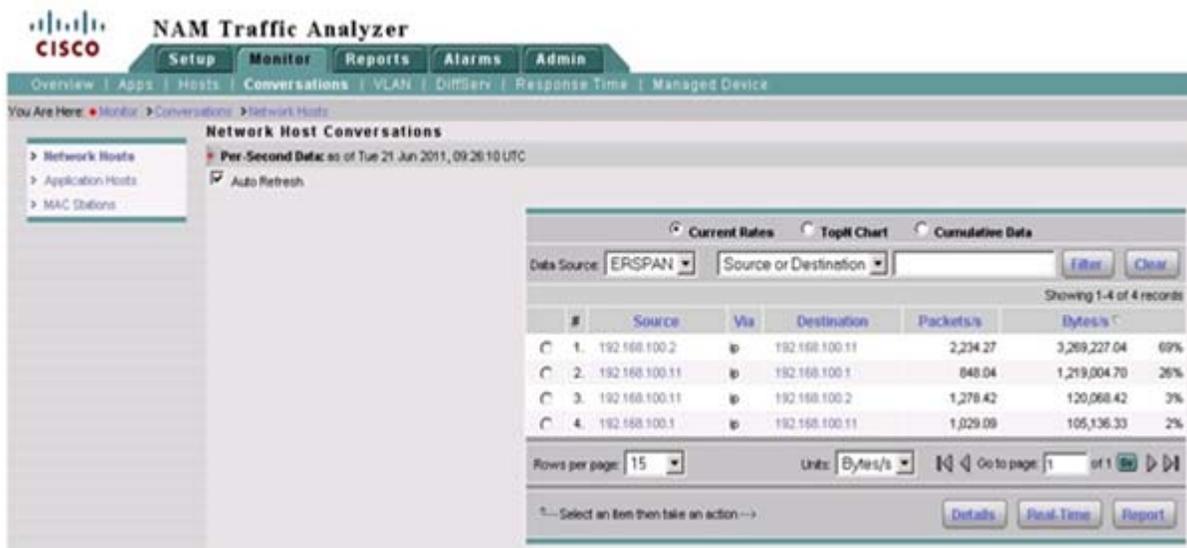


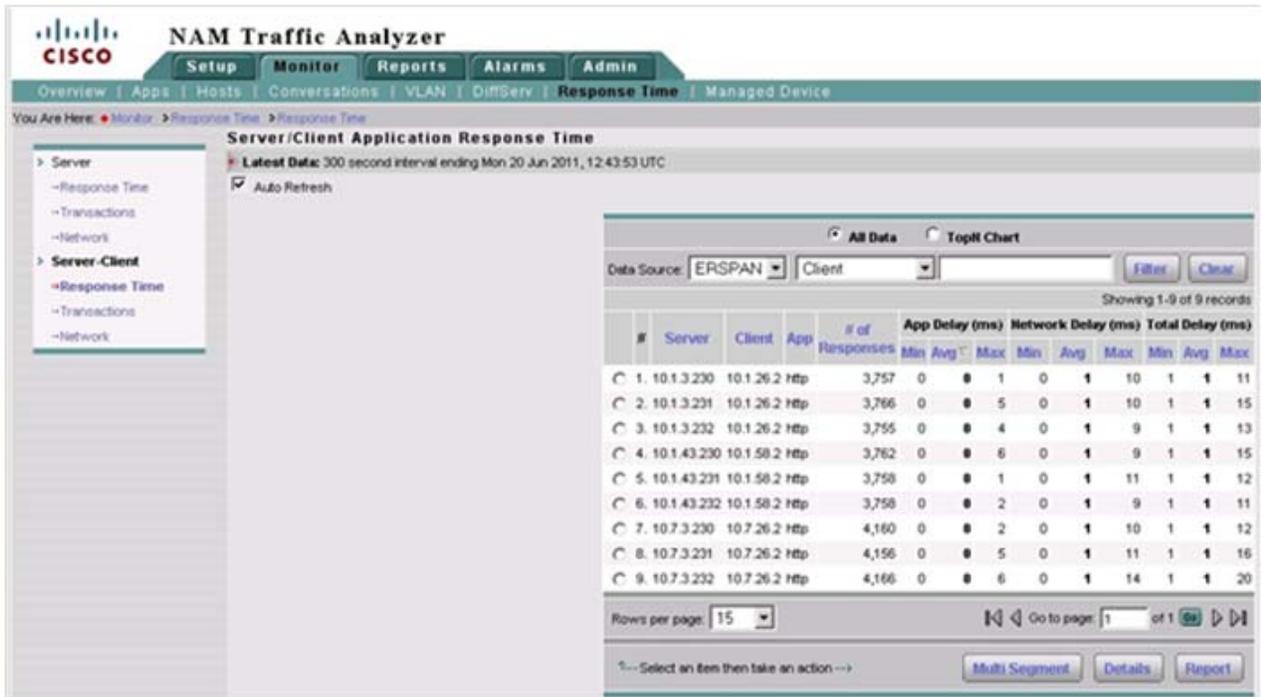
Figure 2-25 Cisco NetFlow Collector Nexus 1000v vMotion Results



Encapsulated Remote Switched Port Analyzer (ERSPAN)

ERSPAN allows for remote monitoring of network resources. ERSPAN uses GRE tunnels to route traffic to the appropriate destination. The Nexus 1000v supports ERSPAN, allowing network administrators to observe the traffic associated with the following:

Figure 2-27 View of Application Response Time Data Collected on N1Kv VEM Uplink





APPENDIX **A**

Bill of Materials As Validated

Table A-1 presents the product part numbers for the components, as well as the network infrastructure required to build out the Cisco VMDC solution, version 2.1.

Table A-1 Cisco VMDC Solution, Version 2.1. 10 Gbps Configuration

Part Number	Description	Quantity
Cisco Nexus 7010		
N7K-C7010-BUN	Nexus 7010 Bundle (Chassis,SUP1,(3)FAB1,(2)AC-6KW PSU)	2
N7K-SUP1	Nexus 7000 - Supervisor, Includes External 8 GB Log Flash	2
N7K-M108X2-12L	Nexus 7000 - 8 Port 10GbE with XL option (req. X2)	4
SFP-10G-SR	10GBASE-SR SFP Module	32
DC3 COMPACT FLASH	Compact Flash Option	2
N7K-CPF-2GB	Nexus Compact Flash Memory 2 GB (Expansion Flash - Slot 0)	2
N7KS1K9-50	Cisco NX-OS Release 5.0	2
Cisco Nexus 5020		
N5k-C5020P-BF	N5000 2RU Chassis no PS 5 Fan Modules 40 ports (req SFP+)	2
N5K-M1008	N5000 1000 Series Module 8xFC 4/2/1 G (req SFP)	2
SFP-10G-SR	10GBASE-SR SFP Module	
N5K PS OPT	Nexus 5 K Power Supply Options	2
N5K-PAC-1200W	Nexus 5020 PSU module, 100-240VAC 1200W	4
N5K CAB OPT	Power Cables	4
CAB-C13-C14-JMPR	Recessed receptacle AC power cord 27	4
N5K EXPAND	N5K Expansion Class	2
N5K-M1-BLNK	N5000 1000 Series Expansion Module Blank	2
N5KUK9-413N1.1	Nexus 5000 Base OS Software Rel 4.1(3)N1(1)	2
N5020-ACC-KIT	Generic accessory kit for Nexus 5020	2
SFP-H10GB-CU3M	<ul style="list-style-type: none"> 10GBASE-CU SFP+ Cable 3 Meter 	64
or	or	
SFP-H10GB-CU5M	<ul style="list-style-type: none"> 10GBASE-CU SFP+ Cable 5 Meter 	
Cisco Nexus 1010, Cisco Nexus 1000V, and VMware (Supporting 1 Pod)		

Table A-1 Cisco VMDC Solution, Version 2.1. 10 Gbps Configuration (continued)

Part Number	Description	Quantity
N1K-C1010	Nexus 1010 Appliance	2
N1K-C1010-NAM-4.2	Cisco NAM Software 4.2 for Cisco Nexus 1000V NAM Virtual Service Blade	2
L-N1KC1010-NAM4.2=	Cisco NAM Virtual Service Blade Software 4.2 for C1010 (eDelivery)	2
N1K-VSMK9-404S12	Nexus 1000V VSM on Physical Media	2
L-N1K-VLCPU-32	Nexus 1000V eDelivery CPU License Qty 32 (1YR Min Service)	2
L-N1K-VLCPU-16	Nexus 1000V Paper CPU License Qty 16-Pack	2
VMW-VCS-3A	VMware vCenter Server Standard, 3yr 24x7 support	1
VMW-VS-ENTP-3A	VMware vSphere Enterprise Plus (1 CPU), 3yr 24x7 support	64
Cisco Nexus 2148		
N2K-C2148T-1GE	N2K 1GE FEX, 1PS, 1 Fan Module, 48x1G-BaseT+4x10GE (req SFP+)	2
SFP-10G-SR	10GBASE-SR SFP Module	4
CAB-AC-250V/13A	North America, NEMA L6-20 250V/20A plug-IEC320/C13 receptacle	4
N2K-PAC-200W	N2K-C2148 Series FEX 200W AC Power Supply	2
SFP-H10GB-CU3M or SFP-H10GB-CU5M	<ul style="list-style-type: none"> • 10GBASE-CU SFP+ Cable 3 Meter or • 10GBASE-CU SFP+ Cable 3 5 Meter 	4
Cisco UCS B-Series - Support 1 Full Pod- 4 Clusters - 64 Servers (10-Gbps Compute Layer)		
N10-S6100	UCS 6120XP 20-port Fabric Interconnect/0 PSU/2 fans/no SFP+	4
B SW IMG OPT	Software Image Options	4
N10-MGT001	UCS Manager v1.0.1	4
B SLOT 0 OPT	Slot Options	4
N10-E0080	8-port 4Gb FC/Expansion module/UCS 6100 Series	4
B PWR SUP OPT	Power Supply Options	4
N10-PAC1-550W	550W power supply unit for UCS 6120XP/100-240VAC	8
B PWR CAB OPT	Power Cables	4
CAB-C13-C14-JMPR	Recessed receptacle AC power cord 27	8
B ACC KIT OPT	Accessory Kit Options	4
N10-SACCA	Accessory kit for UCS 6120XP Fabric Interconnect	4
N20-C6508	UCS 5108 Blade Server Chassis/0 PSU/8 fans/0 fabric extender	8
SC IO OPT	I/O Module Addons	8
N20-I6584	UCS 2104XP Fabric Extender/4 external 10Gb ports	16
SC PWR SUP OPT	Power Supply	8
N20-PAC5-2500W	2500W power supply unit for UCS 5108	32
SC PWR CAB OPT	Power Cables	8
CAB-C19-CBN	Cabinet Jumper Power Cord, 250 VAC 16A, C20-C19 Connectors	32

Table A-1 Cisco VMDC Solution, Version 2.1. 10 Gbps Configuration (continued)

Part Number	Description	Quantity
SC EXPANSION OPT	Expansion Options (Hidden)	8
N20-BBLKD	HDD slot blanking panel for UCS B-Series Blade Servers	128
N20-FAN5	Fan module for UCS 5108	64
N01-UAC1	Single phase AC power module for UCS 5108	8
N20-FW001	UCS 5108 Blade Server Chassis FW package/DO NOT PUBLISH	8
BLADE 0	Blade Options 0	64
N20-B6620-1	UCS B200 M1 Blade Server w/o CPU, memory, HDD, mezzanine	64
G PROC OPT	Processor Options	64
N20-X00001	2.93GHz Xeon X5570 95W CPU/8MB cache/DDR3 1333MHz	128
G MEM OPT	Memory Options	64
N01-M304GB1 or N01-M308GB2	<ul style="list-style-type: none"> 4GB DDR3-1333MHz RDIMM/PC3-10600/dual rank 1Gb DRAMs 8GB DDR3-1333MHz RDIMM/PC3-10600/dual rank 2Gb DRAMs (12 slots per blade) 	64 * 12 = 768
G MEZZ OPT	Mezzanine Options	64
N20-AC0002	UCS M81KR Virtual Interface Card/PCIe/2-port 10Gb	64
G EXPAND OPT	Expansion (Hidden)	64
N20-BHTS1	CPU heat sink for UCS B200 M1 Blade Server	128
Cisco Catalyst 6509-VSS Data Center Services Node		
WS-C6509-E	Catalyst 6500 Enhanced 9-slot chassis, 15RU, no PS, no Fan Tray	2
SV33AEK9-12233SXI	Cisco CAT6000-VSS720 IOS ADVANCED ENTERPRISE SERVICES SSH	2
VS-S720-10G-3CXL	Cat 6500 Supervisor 720 with 2 ports 10GbE MSFC3 PFC3C XL	2
VS-F6K-MSFC3	Catalyst 6500 Multilayer Switch Feature Card (MSFC) III	2
VS-F6K-PFC3CXL	Catalyst 6500 Sup720-10G Policy Feature Card 3CXL	2
VS-S720-10G	Catalyst 6500 Supervisor 720 with 2 10GbE ports	2
MEM-C6K-CPTFL1GB	Catalyst 6500 Compact Flash Memory 1GB	2
BF-S720-64MB-RP	Bootflash for SUP720-64MB-RP	2
CF-ADAPTER-SP	SP adapter for SUP720 and SUP720-10G	2
MEM-C6K-CPTFL1GB	Catalyst 6500 Compact Flash Memory 1GB	2
WS-C6509-E-FAN	Catalyst 6509-E Chassis Fan Tray	2
WS-CAC-6000W	Cat6500 6000W AC Power Supply	4
CAB-AC-2500W-US1	Power Cord, 250Vac 16A, straight blade NEMA 6-20 plug, US	8
ACE20-MOD-K9=	Application Control Engine 20 Hardware	2
ACE20 SW OPT	ACE Module Software Options	2
SC6K-3.0.0A16-ACE	ACE 3.0.0A1(6) Software Release	2
ACE20 PERF LIC OPT	Performance License Options	2

Table A-1 Cisco VMDC Solution, Version 2.1. 10 Gbps Configuration (continued)

Part Number	Description	Quantity
ACE-04G-LIC	Application Control Engine (ACE) 4Gbps License	2
ACE-VIRT-050	Application Control Engine Virtualization 50 Contexts	2
WS-X6708-10G-3CXL	C6K 8 port 10 Gigabit Ethernet module with DFC3CXL (req. X2)	4
WS-F6700-DFC3CXL	Catalyst 6500 Dist Fwd Card- 3CXL, for WS-X67xx	4
WS-X6708-10GE	Cat6500 8 port 10 Gigabit Ethernet module (req. DFC and X2)	4
X2-10GB-SR	10GBASE-SR X2 Module	20
GLC-T	1000BASE-T SFP	4
WS-SVC-NAM-2-250S	Cisco Catalyst 6500 and Cisco 7600 Network Analysis Module	2
SC-SVC-NAM-4.2	Cisco NAM 4.2 for Cat6500/C7600 NAM	2
WS-SVC-FWM-1-K9	Firewall blade for 6500 and 7600, VFW License Separate	6
SC-SVC-FWM-4.0-K9	Firewall Module Software 4.0 for 6500 and 7600, 2 free VFW	6
FR-SVC-FWM-VC-T2	Catalyst 6500 and 7600 virtual FW licensing for 50 VF	6
SF-FWM-ASDM-6.1F	Device Manager for FWSM 4.0 for Catalyst 6500 and 7600	2
MDS 9513		
MDS 9513	Cisco MDS 9513 Multilayer Director Switch	2
DS-X9224-96K9	MDS 9000 24-Port 8-Gbps Fibre Channel Switching Module with SFP and SFP+ LC connectors	4
DS-X9704	4-port 10Gbps Fibre Channel Switching Module	6
DS-X9530-SF2-K9	MDS 9500 Series Supervisor-2 module	4
DS-X9304-18K9	18-port Fibre Channel/4-port Gigabit Ethernet Multiservice (MSM-18/4) module	2
EMC V-Max		
SB-DE15-DIR	V-MAX 15SLT DR ENCL	16
SB-DB-SPS	V-MAX SB SPS	4
SB-FE80000	V-MAX 8M FC-NO PREM	4
SB-32-BASE	V-MAX BASE-32 GB	1
SB-ADD32NDE	V-MAX ADD ENGINE-32 GB	1
NF4103001B	V-MAX 4G 10 K300GB DRIVE	88
SB-PCBL3DHR	50A 3PH DELTA HBL-RSTOL	2
SB-ACON3P-50	ADPTR AC 3PH 50A W/3/4IN CONDUIT ADPTR	4
SB-CONFIG05	V-MAX CONFIG 05	1
PP-SE-SYM	PPATH SE SYM	1
ESRS GW 100	SECURE REMOTE SUPPRT GW	1
SYMVP-RN-OPN	SYMM VIRTUAL PROV RUNTIME	1
ENPTY-SB-BAS	V-MAX ENGINUITY BASE LICENSE	1
ENPTY-SB-C02	V-MAX ENGINUITY 1TB (15-25TB)	24
M-PRESW-001	PREMIUM SOFTWARE SUPPORT	1

Table A-1 Cisco VMDC Solution, Version 2.1. 10 Gbps Configuration (continued)

Part Number	Description	Quantity
NF4106001BU	V-MAX 4G 10K600GB DRV UPG	48
SYMVP-RN-OPN	SYMM VIRTUAL PROV RUNTIME	1
ENPTY-SB-UPG	V-MAX ENGINUITY BASE UPGRADE	1
SYM-MIGR-BAS	SYMMETRIX MIGRATION PKG BASE LICENSE	1
ENPTY-SB-C04	V-MAX ENGINUITY 1TB (41-60TB)	25
NetApp (Network Attached Storage)		
FAS6080A-IB-BS2-R5	FAS6080A,IB,ACT,ACT,HW/SW, 220V,R5	1
X1107A-R6-C	NIC 2-PORT BARE CAGE SFP+ 10GbE SFP+ PCIe, -c	2
DS4243-0724-12A-R5-C	DSK SHLF, 12x2.0TB,7.2K,SATA, IOM3,-C,45	2

