CHAPTER 1

Design Overview

The Virtualized Multiservice Data Center (VMDC) architecture is based on the foundational design principles of modularity, high availability (HA), differentiated service support, secure multi-tenancy, and automated service orchestration (Figure 1-1).

Design Principles

These design principles provide streamlined turn-up of new services, maximized service availability, resource optimization, facilitated business compliance, and support for self-service IT models. These benefits maximize operational efficiency and enable private and public cloud providers to focus on their core business objectives.

Figure 1-1 VMDC Design Principles

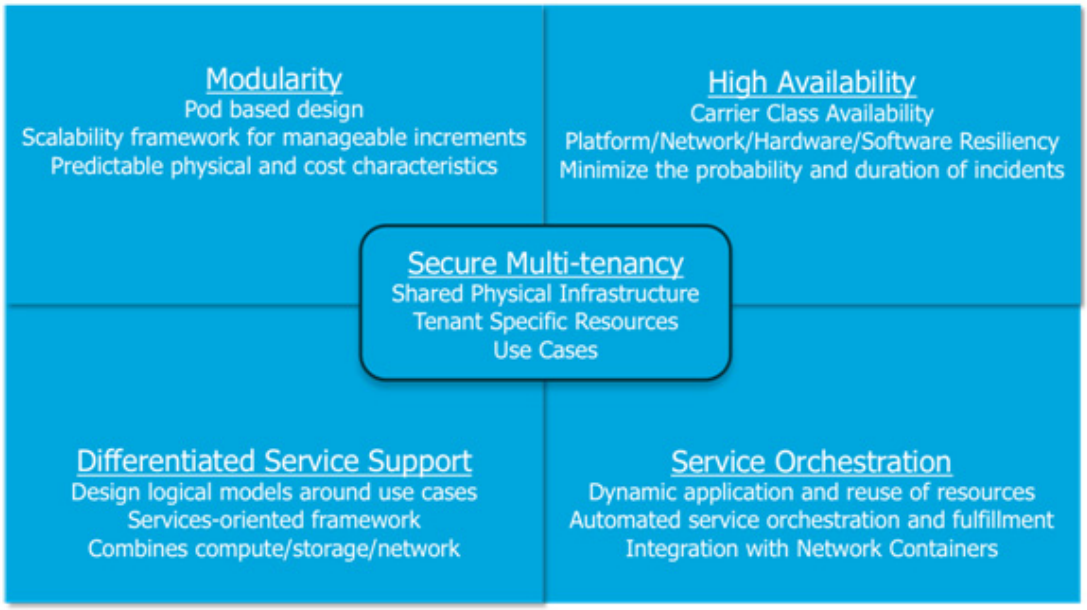

Modularity—Unstructured growth is at the root of many operational and CAPEX challenges for data center administrators. Defining standardized physical and logical deployment models is the key to streamlining operational tasks such as moves, adds and changes, and troubleshooting performance issues or service outages. VMDC reference architectures provide blueprints for defining atomic units of growth within the data center, called PoDs.

High Availability—The concept of public and private "Cloud" is based on the premise that the data center infrastructure transitions from a cost center to an agile, dynamic platform for revenue-generating services. In this context, maintaining service availability is critical. VMDC reference architectures are designed for optimal service resilience, with no single point of failure for the shared ("multi-tenant") portions of the infrastructure. As a result, great emphasis is placed upon availability and recovery analysis during VMDC system validation.

Differentiated Service—Generally, bandwidth is plentiful in the data center infrastructure. However, clients may need to remotely access their applications via the Internet or some other type of public or private WAN. Typically, WANs are bandwidth bottlenecks. VMDC provides an end-to-end QoS framework for service tuning based upon application requirements. This release adds consideration of a set of tools for application visiibility, control and optimization, enhancing the ability to provide application-centric differentiated services.

Multi-Tenancy—As data centers transition to Cloud models, and from cost centers to profit center, services will naturally broaden in scope, stretching beyond physical boundaries in new ways. Security models must also expand to address vulnerabilities associated with increased virtualization. In VMDC, "multi-tenancy" is implemented using logical containers, also called "Cloud Consumer" models that are defined in these new, highly virtualized and shared infrastructures. These containers provide security zoning in accordance with Payment Card Industry (PCI), Federal Information Security Management Act (FISMA), and other business and industry standards and regulations. VMDC is certified for PCI and FISMA compliance.

Service Orchestration—Industry pundits note that the difference between a virtualized data center and a "cloud" data center is the operational model. The benefits of the cloud – agility, flexibility, rapid service deployment, and streamlined operations – are achievable only with advanced automation and service monitoring capabilities. The VMDC reference architectures include service orchestration and monitoring systems in the overall system solution. This includes best-of-breed solutions from Cisco (for example, Cisco Intelligent Automation for Cloud) and partners, such as BMC and Zenoss.

Deltas From VSA 1.0 and 1.0.1 Systems

VMDC VSA 1.0.1 leveraged FabricPath as the Unified Data Center fabric. FabricPath combines the stability and scalability of routing in Layer 2 (L2), supporting the creation of simple, scalable, and efficient L2 domains that apply to many network scenarios. Because traffic forwarding leverages the Intermediate System to Intermediate System (IS-IS) protocol, rather than Spanning Tree (STP), the bi-sectional bandwidth of the network is expanded, facilitating data center-wide workload mobility.

For a brief primer on FabricPath technology, refer to:

http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9402/white_paper_c11-687554.pdf

FabricPath benefits include:

In VSA 1.0.2 we return to an STP-based Layer 2 topology, as an interim step for deploying VSA on the new Cisco Nexus 9000 switching systems. As of this writing, the Nexus 9000 systems support "standalone" mode, running NX-OS code for layer 2 and layer 3 Data Center functionality; however, FabricPath is not a supported L2 technology. Of the various possible traditional STP protocol options supported on the Nexus 9000 systems—Spanning Tree Protocol (STP), IEEE 802.1w Rapid Spanning Tree (Rapid PVST+) or IEEE 802.1s Multiple Spanning Tree (MSTP) - we selected MSTP as the most

scalable, given that we envision requiring a large number of transit VLANs through the layer 2 domain for connections to per-tenant virtual routers (i.e., depending on the multi-tenant scale requirement), and this protocol decouples spanning tree instances from VLAN instances.

Note Certain application environments, especially those that generate high levels of broadcast, may not tolerate extremely large Layer 2 environments.

In this case, the VSA architecture mitigates high levels of broadcasts within the data center by logically bounding the MSTP L2 fabric with L3 devices: the PE/WAN edge router and multiple, per-tenant virtual routers.

Another minor change from previous VSA releases is the replacement of the clustered Cisco ASR9000 WAN Edge/PE routers with redundant ASR 1000 routers as an alternative option. This change targets Private Cloud deployment cases, where fewer remote sites are aggregating into the Data Center from the wide area backbone.

Finally, VSA 1.0.1 reintroduced the physical ASA security appliance, as an alternative perimeter firewall to the CSR or ASA1000v virtual firewalls as part of a hybrid physical/virtual security deployment option; however this was not a focus for VSA 1.0.2 and so although a valid architectural option was not included in scope for this incremental release.

VSA Differentiator: Virtual Network Services

Previous releases of VMDC addressed several methods of resiliently attaching redundant appliance or module-based service nodes to optimize service availability and efficient link path utilization, including Ether-channel, vPCs with Multi-Chassis EtherChannel on paired Virtual Switching Systems (VSSs), and vPCs on clustered (Cisco ASA) firewall appliances. However, service node implementation in VMDC VSA differs significantly from these releases in the following ways:

- **Placement**—In the Compute tier of the infrastructure, instead of the traditional aggregation layer
- **Form-Factor**—vApp, rather than physical
- **Application**—Dedicated per-tenant or organizational entity, rather than shared

These characteristics provide for a "pay as you grow" model with significant CAPEX savings in upfront deployment costs. In terms of availability implications, dedication of resources to a specific tenant means that strict 1:1 redundancy may no longer be the default mode of operation for these forms of service nodes. Rather, administrators now have greater flexibility to fine-tune redundant services and methods for those tenants or organizations who have mission-critical applications with high availability requirements.

VMDC Virtualized Containers

The VMDC architecture can support multiple virtual containers, referred to as cloud consumer models. These models are described in greater detail later in this document, and in previous release material:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/VMDC/2.2/collateral/vmdcConsumerModels.pdf

Because this release is based on unique, dedicated per-tenant security, load balancing and optimization services, for validation purposes VMDC VSA 1.0.2 focuses only on containers that do not feature shared (multi-tenant) security/services zones. High-level representations of these are highlighted in green in

Figure 1-2.

Note The "Gold" container does not feature a shared zone, but is considered to be a subset of the "Expanded Gold" container.

Figure 1-2 VMDC Containers

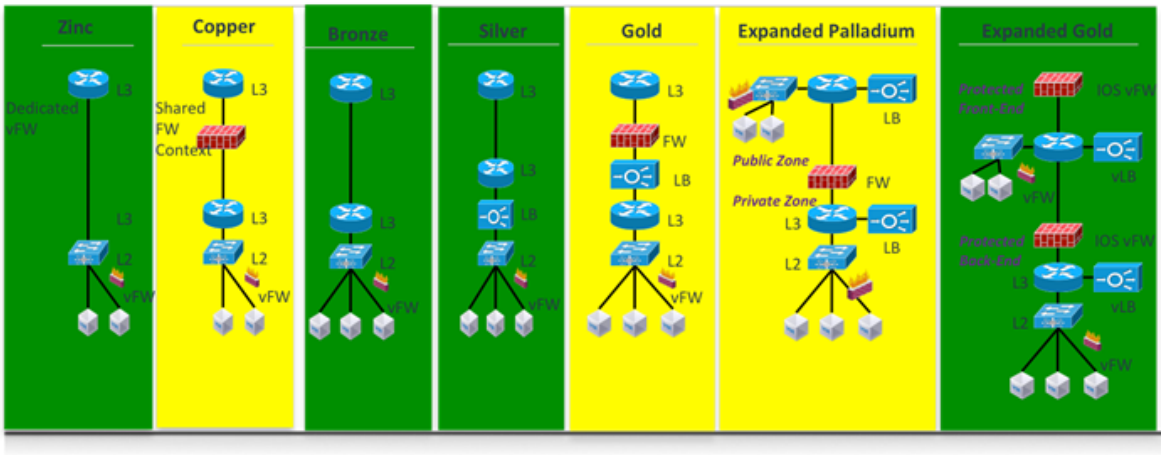

As you move from left to right Figure 1-2, the validated VMDC VSA containers, which are based upon real-world, commonly deployed N-tiered application and security models, become increasingly complex, growing from single to multiple security zones and policy enforcement points and from application of single to multiple types of services. VMDC VSA features additional dedicated network service options, such as network analysis and optimization. Although not shown in Figure 1-2, these were validated as part of the Expanded Gold container in VSA 1.0. In VSA 1.0.1, we evolved the system architecture by extending these tenancy models down into the storage layer of the infrastructure, leveraging new storage abstraction and isolation technology in the form of NetApp's Storage Virtual Machines.

Solution Components

The following sections describe the network components used in the VMDC VSA solution (summarized in Table 1-1) and provide a snapshot of the intra-DC and overall system end-to-end network topology model validated in VMDC VSA 1.0.2 (Figure 1-3 and Figure 1-4).

Table 1-1 VMDC VSA 1.0 Solution Component Matrix

Function	Components
Network	Cisco ASR 9000, ASR 1000, ISRG2 3945, CSR Cisco Nexus 7009, 7004 (Nexus 7018 and Nexus 7010 not in SUT but valid architectural option) Sup2E, F2E and Sup2, F2 series 1 and 10 Gbps Ethernet cards Cisco Nexus 5548 Cisco Nexus Fabric Extender 2248TPE
Services (Physical Form Factor)*	Cisco ASA 5585-X with SSP-60 * Only applicable to VSA 1.0.1
Services (vApp Form Factor)	Citrix NetScaler VPX Server Load Balancer Cisco Netscaler 1000v Server Load Balancer Cisco vWAAS Cisco vNAM Cisco Nexus 1100 (Services Chassis)
Security Services (vApp Form Factor)	Cisco IOS XE 3.10 ZBF (for example, on CSR) Cisco ASA 1000V Virtual Security Gateway
Compute	Cisco Unified Computing System (UCS) Cisco UCS 6296UP Fabric Interconnect Cisco Fabric Extender 2208XP IO Module UCS 5108 Blade Server Chassis UCS B200/230/440-M2 and B200-M3 Blade Servers C200/240-M2/M3L Servers UCS M81KR Virtual Interface card UCS P81E Virtual Interface card UCS Virtual Interface card 1280, 1240
Virtualization	VMware vSphere VMware ESXi 5.1 Hypervisor Cisco Nexus 1000V (virtual access switch)
Storage Fabric*	Cisco MDS 9513 (1/2/4/8 Gbps 24-Port FC Module; 18/4-Port Multiservice Module; Sup-2A; 24-port 8 Gbps FC Module; 18-port 4 Gbps FC Module) * *Not Applicable to this release.*

Table 1-1 VMDC VSA 1.0 Solution Component Matrix (continued)

Function	Components
Storage Array	NetApp FAS 6250[1]
Orchestration/ Management*	Domain Management: • UCS Manager • CIMC • Nexus 1000V Virtual Supervisor Module • Cisco Virtual Network Management Center • vWAAS Central Manager (vCM) • VMware vCenter 5.1 • Fabric Manager • Ontap 8.1.2 Service Assurance: * *CLSA VMDC VSA 1.0 not in scope* Orchestration: * *BMC CLM and CIAC not in scope*

1. Refer to NetApp storage array product family information: http://www.netapp.com/us/products/storage-systems/index.aspx

Figure 1-3 VMDC VSA 1.0.2 Intra-DC Topology

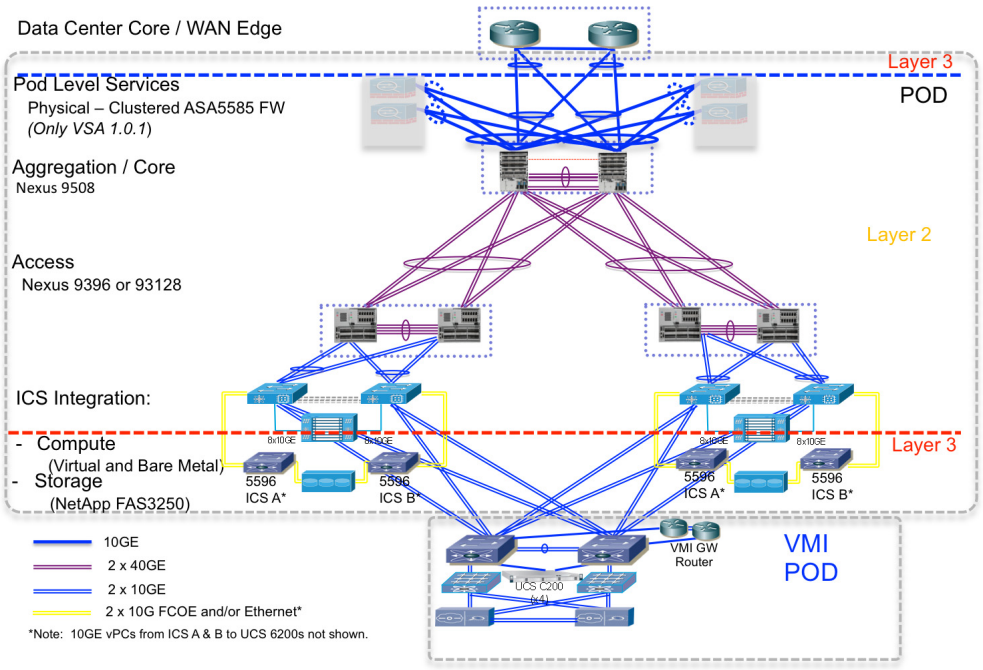

Figure 1-4 VMDC VSA 1.0 End-to-End Topology

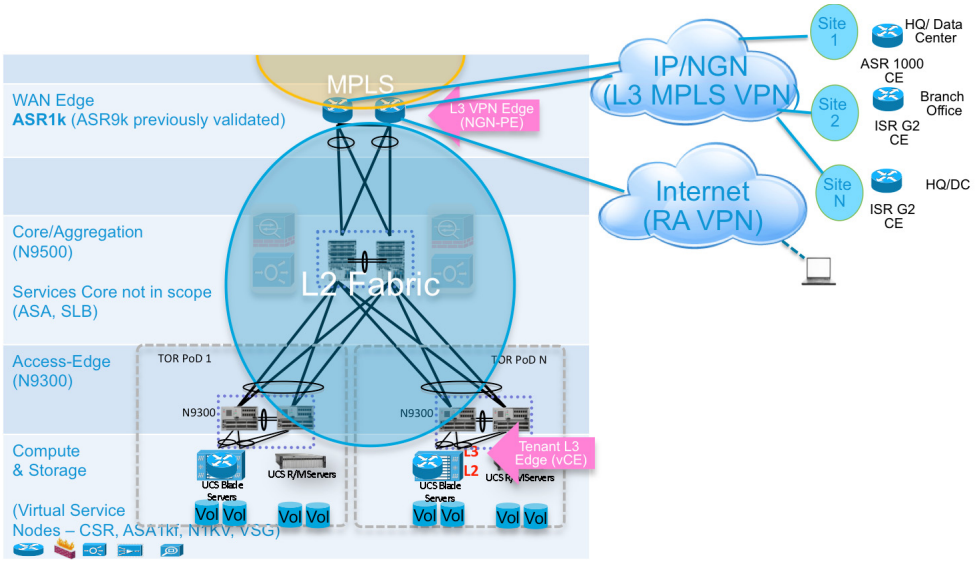

The virtual service model may easily be utilized in scaled down form at Enterprise remote sites to provide private cloud services as part of a Public Provider managed service offering. In this case, the remote sites in the preceding diagram in this context would be centrally controlled via out of band management paths. The private clouds can be tailored to fit application and services requirements, ranging in size from a Flexpod or Vblock to a small C-Series chassis "pod-in-a-box" entry point (Figure 1-5).

Figure 1-5 Remote Site Private Cloud

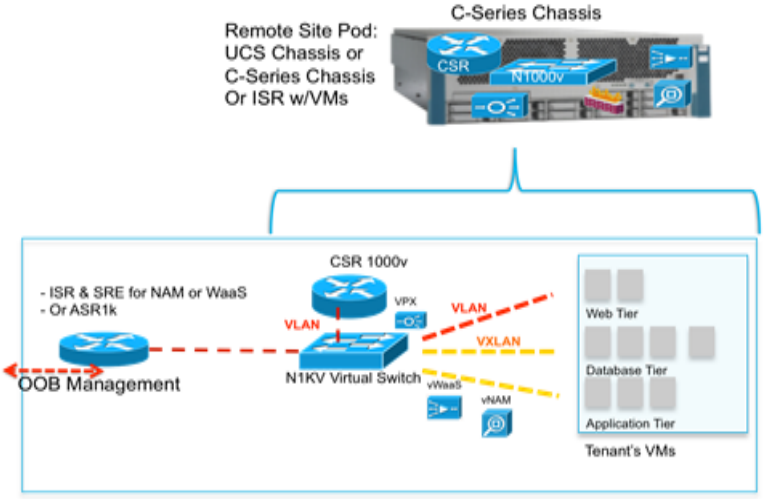

VMDC Change Summary

The following release change summary is provided for clarity.

- **VMDC 1.0, 1.1**—Introduces architecture foundation for deploying virtualized and multi-tenanted data centers for cloud-based services. It supports high availability, elasticity, and resiliency of virtualized compute, network, and storage services.
- **VMDC 2.0**—Expands VMDC 1.1 by adding infrastructure orchestration capability using BMC software's Cloud Lifecycle Management, enhances network segmentation and host security, uses integrated compute stacks (ICS) as building blocks for the PoD, and validates compact and large PoD scale points.
- **VMDC 2.1**—Generalizes and simplifies VMDC 2.0 architecture for a multi-tenant virtualized data center used for private cloud. Improvements include multicast support, simplified network design, jumbo frame support, improved convergence, performance, scalability for private cloud, QoS best practices, and increased design flexibility with multi-tenant design options.
- **VMDC 2.2**—Builds on top of VMDC 2.0 and 2.1 for a common release supporting public, private, and hybrid cloud deployments. Enhancements include “defense in depth” security, multi-media QoS support, and Layer 2 (VPLS) based DCI.
- **VMDC 2.3**—Further expands on topology models in previous 2.X releases, providing a more collapsed architectural model, offering smaller footprint and entry point option. Enhancements include introduction of a new “copper” tenancy container mode.
- **VMDC 3.0/3.0.1**—Introduces FabricPath as an L2 multi-pathing technology alternative for the intra and inter-pod Data Center Unified Fabric infrastructure, considering the implications of various methods of appliance or service module-based service insertion.
- **VSA 1.0**—Introduces the Virtual Services Architecture, comprising Network Function Virtualization for IaaS or ITaaS services, addressing tenancy scale and agility with per-tenant perimeter and second tier firewalls, and virtual server load balancing using Citrix Netscaler. Addresses enhanced application visibility and control with end to end QoS, and virtual network analysis (vNAM), and virtual Wide Area Application Services (vWaaS) appliances.
- **VSA 1.0.1**—Reintroduces one of the centralized network service appliances—the ASA5585 Firewall - to provide hybrid physical/virtual tiered defense in depth. Also extends cloud consumer models to include storage tenancy, and enhances storage resilience through clustering. Finally, adds an alternative L3 model for tenancy by extending MPLS to the compute tier of the infrastructure, where the CSR functions as a single-tenant virtual Provider Edge MPLS router.
- **VSA 1.0.2**—Expands the intra-DC topology options for VSA, leveraging an MSTP-based, classical Ethernet L2 fabric on standalone Nexus 9000 systems.

Related Documents

The following documents are available for reference and consideration.

- Cisco Virtualized Multi-tenant Data Center Design and Implementation Guides, Releases 1.0-2.2

  http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns743/ns1050/landing_vmdc.html#~releases
- Design Considerations for Classical Ethernet Integration of the Cisco Nexus 7000 M1 and F1 Modules

http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/VMDC/2.6/vmdcm1f1wp.html

- Virtualized Multi-tenant Data Center New Technologies - VSG, Cisco Nexus 7000 F1 Line Cards, and Appliance-Based Services VPLS and EoMPLS Based DCI Solution with nV Edge and vPC

  http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/VMDC/2.6/vmdctechwp.html

- Cisco VMDC 2.2 Design Guide

  http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/VMDC/2.2/design_guide/vmdcDesign22.html

- VMDC 3.0.1 Fabric Path-based Design Guide

  http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/VMDC/3.0.1/DG/VMDC_3.0.1_DG.html

- Data Center Interconnect over MPLS, Ethernet or IP Transport documents

  http://www.cisco.com/c/en/us/solutions/enterprise/data-center-designs-data-center-interconnect/index.html

  http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Data_Center/VMDC/DCI/1-0/DG/DCI.html

- Cloud Service Assurance for VMDC

  http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/dz_cloudservice.html