



Preface

The Cisco Virtualized Multiservice Data Center (VMDC) system provides design and implementation guidance for enterprises deploying private cloud services, and for service providers (SPs) building public and virtual private services. With the goal of providing an end-to-end system architecture, VMDC integrates Cisco and third-party products in the cloud computing ecosystem.

Audience

This guide is intended for, but not limited to, system architects, network design engineers, system engineers, field consultants, advanced services specialists, and customers who want to understand how to deploy a public or private cloud data center infrastructure. This guide assumes that the reader is familiar with the basic concepts of IP protocols, quality of service (QoS), and high availability (HA), and that readers are aware of general system requirements and have a basic understanding of enterprise or SP network and data center architectures.

Overview

VMDC, Cisco's reference architecture for cloud deployment, has been widely adopted by numerous SPs and enterprises worldwide. In this and previous releases, VMDC has provided design guidance for scalable, secure, resilient, public and private cloud infrastructures serving multiple consumers or tenants:

- In the data center portion of the architecture, VMDC 2.X designs were centered on traditional hierarchical infrastructure models incorporating leading Cisco platforms and Layer 2 (L2) resilience technologies such as Virtual Port Channel (vPC), providing network containers or tenancy models of different sizes and service profiles, with necessary network based services and orchestration and automation capabilities to accommodate the various needs of cloud providers and consumers.
- VMDC 3.X systems releases introduced Cisco FabricPath for intra-DC networks, as an optional L2 alternative to a hierarchical vPC-based design. FabricPath removes the complexities of Spanning Tree Protocol (STP) to enable more extensive, flexible, and scalable L2 designs. Customers leveraging VMDC reference architecture models can choose between vPC-based and FabricPath-based designs to meet their particular requirements.

In this release there is a high degree of focus on cloud provider use cases, building a new logical topology model around the creation of virtual private cloud tenant containers in the shared data center infrastructure. Future or incremental, inter-related releases will incorporate additional cloud consumer models specific to enterprise and private cloud use cases. For example, the VSA 1.0.1 release addresses a hybrid consumer model, comprising physical and virtual security service appliances used together as part of a per-consumer or per-tenant network service set and examines extension of the tenancy container into NetApp storage systems using Storage Virtual Machines (SVMs) for isolation.

The Virtual Services Architecture models can be implemented on either a 2.X (classical Ethernet) or 3.X (FabricPath) VMDC infrastructure. However, in this VSA 1.0.2 release we focus on fundamental implications of an all-virtual approach, and have opted to do so over a simple classical Ethernet layer 2 topology provided by Nexus 9000 Data Center switching systems running in "standalone" mode with NX-OS 6.1(2) software.

Problem Statement

The architecture described in this guide addresses the following customer challenges:

1. **Tenancy Scale and Tenancy Model Enhancements:** Previous VMDC systems releases leveraged various network abstraction technologies, for example, virtual LANs (VLANs) and virtual routing and forwarding (VRF), for tenant isolation, including separated routing and forwarding. Each abstraction technology impacts logical scale and control plane overhead. In a traditional hierarchical DC network model, the pressure point from a network scalability and control plane perspective is at the aggregation layer of the infrastructure, with the number of route peers, VRFs, VLANs, and MAC capacity supported by aggregation nodes presenting key multi-dimensional scalability factors.

Previous VMDC systems releases used hypervisor-enforced abstraction and separation for data-at-rest tenancy. In the VSA system, an additional layer of storage abstraction, isolation and delegation for secure multi-tenancy was introduced in release 1.0.1, using the Storage Virtual Machine (SVM) and Storage QoS capabilities of NetApp's clustered Data ONTAP. Tenants can be provisioned to use either storage resources from a single, shared SVM, or some or all tenants can be provisioned with their own unique SVM for an additional layer of isolation. A per-tenant SVM model also provides the option of delegating management of those storage resources to the tenants themselves. Storage QoS can proactively prevent or reactively mitigate the impact of one tenant's high workload on other tenants' applications.

The virtual services architectural (VSA) models discussed in VSA 1.0 and 1.0.2 present several alternatives, addressing tenancy scale using either a centralized provider edge (PE) and distributed, per-tenant virtual customer edge (vCE) routing model or a distributed, single-tenant virtual Provider Edge (vPE) logical model which relies on extension of MPLS to the Compute Layer of the Data Center infrastructure and Inter-AS Option B for scalable extension of MPLS VPNs across Autonomous Systems. In the vCE scenario tenancy scale is increased to the number of eBGP peers (or alternatively, static routes) supported by the PE nodes. As of this writing, this is 5000 per pair of redundant ASR 9000 Series PE routers. In the distributed single-tenant vPE model, tenancy scale will be a function of the VPNv4 route scale of the Data Center InterAS border router.



Note This VSA 1.0.2 release leveraged the vCE logical model as the basis for the system.

2. **Complexity:** Current VMDC architecture models feature a relatively high degree of management complexity because service appliances are shared across multiple tenants, and are allocated in logical "slices" (contexts) by automation systems. The VSA model reduces service orchestration

complexity, removing cross-tenant dependencies for L4-L7 service allocation. The VSA models represent simpler logical topologies compared to the back-to-back VRF-Lite method employed in VMDC 2.X releases to create rigorous (VRF-based) tenant isolation.

3. Customer Evolution to NFV for IaaS: For years, customers have seen the transition from physical to virtual services as a foundation for an evolution toward “next-gen” data center service-oriented architectures, providing increased flexibility and agility through greater “software definability”.
4. Need for virtual appliance-based multi-tenancy design guidance. VMDC VSA provides a starting point, representing an opportunity to initially consider a couple of specific deployment models (vCE and vPE) out of several possible options for an "all-virtual" virtual private cloud instantiation, exploring end-to-end service differentiation, performance and impact on future automation requirements.
5. Need to address logical segmentation constraints of traditional 802.1q VLAN L2 domains through the application of virtual overlays. VMDC VSA presents a first look at the use of VXLANs for logical segmentation.
6. Specific to this release: Additionally, we address considerations of replacing the layer 2 FabricPath fabric utilized in VSA 1.0 and 1.0.1 with a traditional spanning-tree/Virtual Port-channel, classical Ethernet layer 2 fabric provided by Nexus 9000 switching systems.

VMDC VSA 1.0.2 addresses the following use cases:

- Data center and PoD design
- Split N-tiered applications
- Multi-tenancy (including Virtual Extensible LAN (VLXAN)-based logical segmentation), matched with SVMs for multi-tenant allocation of storage resources
- Application-centric instrumentation (statistics collection, network analysis, WAN optimization, Performance Agent)
- Storage Clustering for HA

Solution Proposal

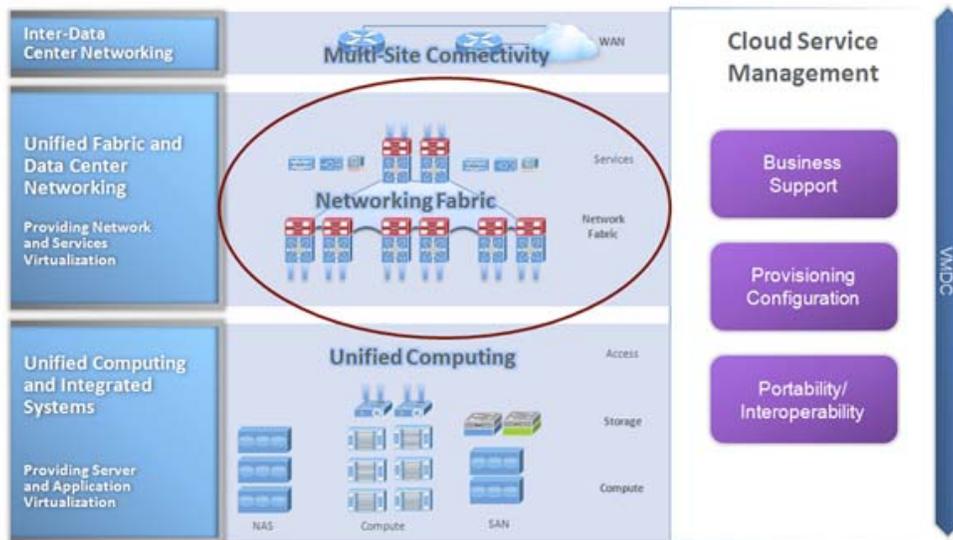
To address the identified requirements, we modified the Unified Computing component of the VMDC architecture, shifting virtualized service functions from the Unified Fabric/Data Center Networking portions of the infrastructure. [Figure II-1](#) shows a high level view of the overall VMDC system.

In general, the solution comprises three modular layers:

1. Unified Computing and Integrated Systems (UCIS), providing server and application virtualization, typically consisting of FlexPods and Vblocks.
2. Unified Fabric and Data Center Networking (UFDC), providing network and network based services virtualization.
3. Data Center Interconnect (DCI), providing seamless multi-site connectivity.

The solution is complemented by Cloud Service Management components that enable end to end provisioning and orchestration, along with monitoring and assurance.

Figure II-1 High Level VMDC Solution



VMDC VSA shifts network service functions to the UCIS in a way that leverages existing design guidance for the Unified Fabric and DCI layers, along with previous guidance for the UCIS layers in terms of compute and storage for **application** workloads. We maintained the following assumptions:

- Previous design guidance for UCIS (FlexPod, Vblock) components applies. VMDC VSA validation was performed on the latest FlexPod release, including storage clustering and tenancy. Applications validated on FlexPod or Vblock continue to function on the overall VMDC architecture.
- Previous design guidance for DCI components applies. VMDC VSA focuses mainly on single-site, intra-PoD Virtual Private Cloud (VPC) design alternatives.
- There are no top-level management and orchestration components in VMDC VSA 1.0 through 1.0.2; orchestration and service assurance are addressed in separate, parallel VMDC programs.
- Cisco (XaaS) applications, such as Unified Communications (UC), Hosted Collaboration Solution (HCS), Media Data Center, Video Surveillance, and TelePresence, use the VMDC architecture as the infrastructure for validation. At this writing, VMDC 2.3 is the latest release used for these validations. No specific Cisco application validations are in the scope of VMDC VSA 1.0.2. However, given the level of validation already performed, we are confident that these applications will work in VMDC VSA 1.X infrastructures without major issues.