Cisco Virtualized Multi-Tenant Data Center Framework

White Paper

Last Updated: May 18, 2012

Cisco Virtualized Multi-Tenant Data Center Framework

Introduction

The Cisco® Virtualized Multi-Tenant Data Center (VMDC) architecture is a set of specifications and guidelines for creating and deploying a scalable, secure, and resilient infrastructure that addresses the needs of cloud computing. To develop a trusted approach to cloud computing, Cisco VMDC combines the latest routing and switching technologies, advancements in cloud security and automation, and leading edge offerings from cloud ecosystem partners. Cisco VMDC enables service providers (SPs) to build secure public clouds and enterprises to build private clouds with the following benefits:

- Reduced time to deployment—Provides a fully tested and validated architecture that enables technology adoption and rapid deployment.
- Reduced risk—Enables enterprises and service providers to deploy new architectures and technologies with confidence.
- Increased flexibility—Enables rapid, on-demand workload deployment in a multi-tenant environment using a comprehensive automation framework with portal-based resource provisioning and management capabilities.
- Improved operational efficiency—Integrates automation with multi-tenant resource pools (compute, network, and storage) to improve asset use, reduce operational overhead, and mitigate operational configuration errors.

The Cisco VMDC architecture is part of a comprehensive set of cloud architectures and offerings from Cisco.

Corporate Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

Figure 1 Cisco Cloud Focus Areas

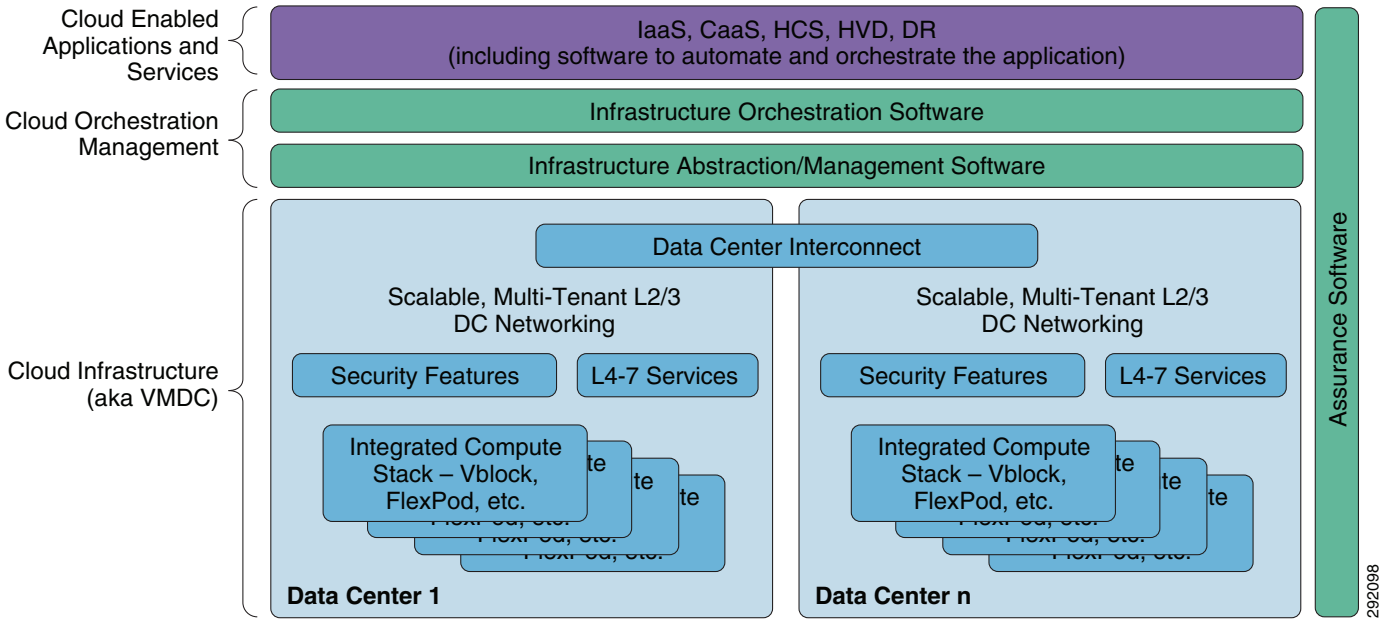

As depicted in Figure 1, Cisco cloud focus areas include:

- Cloud-enabled applications and services, such as Infrastructure as a Service (IaaS), Compute as a Service (CaaS), cloud-enabled Hosted Collaboration System (HCS), cloud-enabled Hosted Virtual Desktop (HVD), and Disaster Recovery and Business Continuity (DR/BC).
- Cloud Orchestration and Management subsystems configure, provision, monitor, and automate operational activities in the cloud.
- Cloud Infrastructure subsystems (based on Cisco VMDC) provide the network, compute, and storage capabilities, as well as security features, Layer 4-Layer 7 network-based services, and data center interconnections.

This document describes the Cisco VMDC cloud infrastructure architecture. Details of the Cisco VMDC architecture are provided in release-specific data sheets, design guides, implementation guides, and technical white papers. Details regarding other Cisco Cloud offerings, such as cloud management and automation and cloud enabled applications, are covered in their respective solution document sets.

Cisco VMDC Overview

The Cisco VMDC architecture is based on the following key cloud design concepts:

- Flexibility
- Multi-Tenant Consumption Models
- Service Differentiation
- Layered Security
- High Availability
- Comprehensive Service Management
- Application Enablement

Flexibility

The Cisco VMDC architecture is designed to enable different deployment models at different scale to facilitate gradual growth and expansion as needed by either small enterprises or large service providers. The basic modular building block of the architecture is a "point of delivery" (pod), which contains standardized compute, storage, and network components with no service specific dependency. Pods can be auto-discovered by the operation software and configured to the specific service profile as needed. A deployment can start with very small set of pods and grow as the business demands with no need for major overhaul or re-design. In this way Cisco VMDC helps administrators to scale their build-outs in predictable, logical units, easing their planning and capacity management and ultimately lowering their operational costs.

Multi-Tenant Consumption Models

Multi-tenant consumptions models refers to the ability of the data center to host multiple separate zones, each of which can serve a separate group of users with a specific service profile. Tenants can be organizations, departments, customers, enterprises, regions, etc. Using a rich set of security features that separate and secure tenant traffic and interactions, Cisco VMDC enables both the simple loose separation required in private clouds as well as the highly strict and secure separation required in public clouds.

In an enterprise context, tenants may be departments of an organization that operates a private cloud. Each department may have different computational needs, applications, or storage scale, while at the same time all of the departments need to be managed and maintained under a single unified operational domain. The secure separation requirements in this type of a private cloud may be simple traffic separation in addition to user access control. In a service provider context, tenants can be different enterprises that lease/use the SP cloud infrastructure in a shared public cloud environment. As in the private cloud, each of these enterprises can have very specific computation, application, and storage needs. However each tenancy must be enforced more strictly for security, usage, service level agreement (SLA) conformance, as well as independent operational models, so as to both satisfy the leasing enterprise requirements as well as the hosting SP business requirements. The Cisco VMDC architecture enables either model of tenancy and secure separation using a mix of physical and logical segmentation, monitoring, detection, and policy application.

Service Differentiation

The Cisco VMDC architecture enables a cloud provider to either define custom service grades or use pre-defined service grades to differentiate service offerings within the cloud. Various components of a service may include compute allocations in the form of CPU and virtual machine limits, storage and data protection allocations, network-based services such as VLAN segment allocations, quality of service (QoS) capabilities, security, disaster recovery and business continuity, and other application-level features. The Cisco VMDC architecture includes four pre-defined service tiers: Bronze, Silver, Gold, and Palladium. These tiers are not meant to be a strict definition of the levels of service that may be offered, but are simply representative service grades that were used in validation tests.

Layered Security

By embedding security at each layer of the data center, Cisco VMDC enables a rich and powerful set of tools for operators to secure their deployment, enabling highly secure multi-tenant deployments, one of the main hallmarks of cloud computing. Several features on the network devices and on the compute infrastructure combine to provide the robust security that gives organizations the confidence to use cloud computing infrastructure to address their business needs for application deployment. Figure 5 summarizes the security features that are used in the Cisco VMDC architecture.

High Availability

The Cisco VMDC architecture is designed to optimize service up time by enabling availability and fault tolerance at all layers of the data center through physical redundancy best practices and virtualized failover features at the network, compute, and storage layers:

- At the network layers, physical redundancy starts with hardware redundancy within a node, redundant nodes, and redundant links, with optimized failover convergence end-to-end. Other network-based features that enable high availability include, but are not limited to, the use of virtual port channels for Layer 2 multi-pathing, Multiple Spanning Tree protocol, Hot-Standby Router Protocol, BGP with Non Stop Forwarding and Bi-directional Forwarding Detection.
- At the compute layer, in addition to physical redundancy of Cisco Unified Computing System™ (UCS™) servers, the following features are leveraged: UCS end-host mode, Cisco Nexus® 1000V and Mac-pinning, redundant VSMs in active-standby mode, high availability within the cluster, and automated disaster recovery plans.
- At the storage layers, the Cisco VMDC architecture leverages best practice methodologies for SAN HA, prescribing full hardware redundancy at each device in the I/O path from host to SAN, beginning at the server, with dual port adapters per host. Redundant paths from the hosts feed into dual, redundant MDS SAN switches (i.e., with dual supervisors) and then into redundant SAN arrays with tiered RAID data protection.

In addition to a comprehensive HA design, end-to-end availability of the Cisco VMDC architecture is validated in real world scale labs to ensure deployment quality.

Comprehensive Service Management

The Cisco VMDC architecture is closely integrated with the service orchestration and service assurance subsystems that provide automation of configuration and provisioning for both the provider offering services through the cloud as well as the consumers using these services. Service orchestration is a multi-domain configuration abstraction layer on top of the data center infrastructure. It enables a portal-based configuration model in which the subscriber can select from a defined number of service options and host applications as virtual machines. Based on these selections, configuration actions are performed on the devices to achieve the service represented in the portal. This self-service, portal-based model offers customization per customer and reduces the number of manual tasks required of the IT department. Service orchestration also automates the configuration across many devices based on the services advertised through the portal. The orchestrator used in this architecture is a BMC Atrium Orchestrator. Cisco VMDC offers an open framework so many orchestration vendors may readily manage the architecture.

Application Enablement

An objective of the Cisco VMDC architecture is to enable faster and smoother application deployment. To achieve this objective, the architecture is validated with specific application-level requirements such as QoS. In turn, cloud-enabled application solutions, such as hosted collaboration and hosted virtual desktop, are validated on the Cisco VMDC infrastructure.

Cisco VMDC Architecture

This section describes the primary components of the Cisco VMDC architecture shown in Figure 1:

- Modular Building Blocks—Pods and Integrated Compute Stacks
- Hierarchical Network Design for Scalable Layer 2-Layer 3 Networking
- Security Features
- Services Layer—Layer 4-Layer 7 Services
- Cisco Data Center Interconnect

Modular Building Blocks—Pods and Integrated Compute Stacks

The first and smallest building block in the Cisco VMDC architecture is a generic integrated compute stack (ICS [storage and compute]) based on existing models, such as:

- VCE Vblock™ Infrastructure Packages, which combine Cisco UCS, Nexus, and MDS with EMC® storage components to provide multiple, fixed-sized configuration blocks
- Cisco FlexPod™, which combines Cisco UCS and Nexus with NetApp® storage to provide variable configuration of compute and storage

The Cisco VMDC architecture is not limited to a specific ICS definition but can be extended to include other compute and storage stacks. Both enterprises and service providers can build and deploy their ideal cloud platform using the ICS design, implementation, and operational best practices described in Cisco VMDC.

Multiple ICSes can be connected to each other with a network infrastructure to create a pod structure which constitutes the second modular unit of the Cisco VMDC architecture. Pods can be easily used to add incremental capacity in a data center. Each pod is discovered by the system, integrated into the resource pools, and assigned workloads as needed. The Cisco VMDC architecture specifies two pod designs:

- Compact pods are designed using a centralized service node architecture on a collapsed aggregation/core and top-of-rack access design.
- Large pods are designed using a distributed service node design on a collapsed access/aggregation layer and an end-of-row access design.

Either pod model can be used in large or small deployments based on operator preference.

Figure 2 depicts the pod and ICS constructs and Figure 3 depicts a collection of pods that together make up a data center. Using the ICS and pod constructs, the Cisco VMDC architecture enables the creation of flexible data centers which can easily expand.

Figure 2 Pods and Integrated Compute Stacks

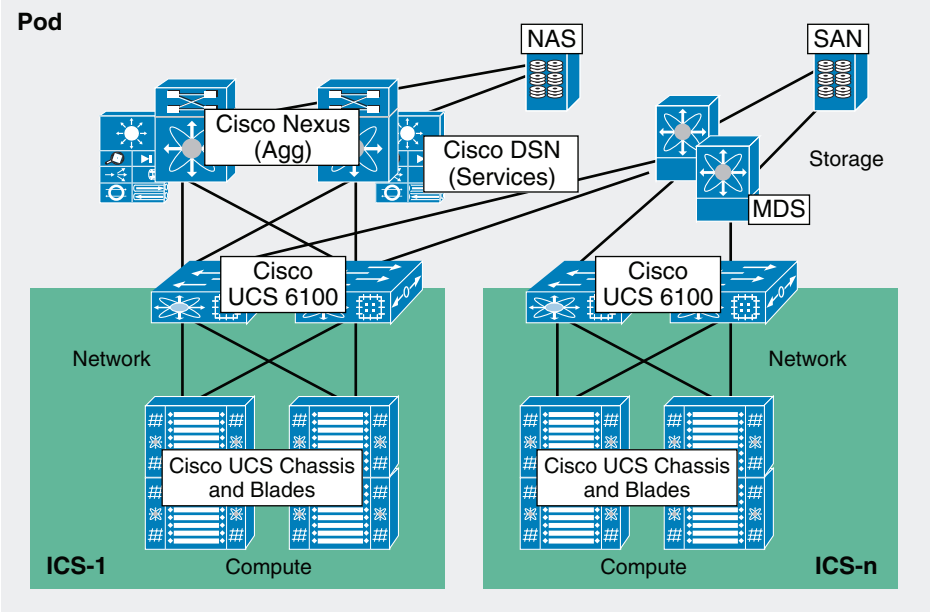

Figure 3 Scalable Data Center Design Using Modular Building Blocks

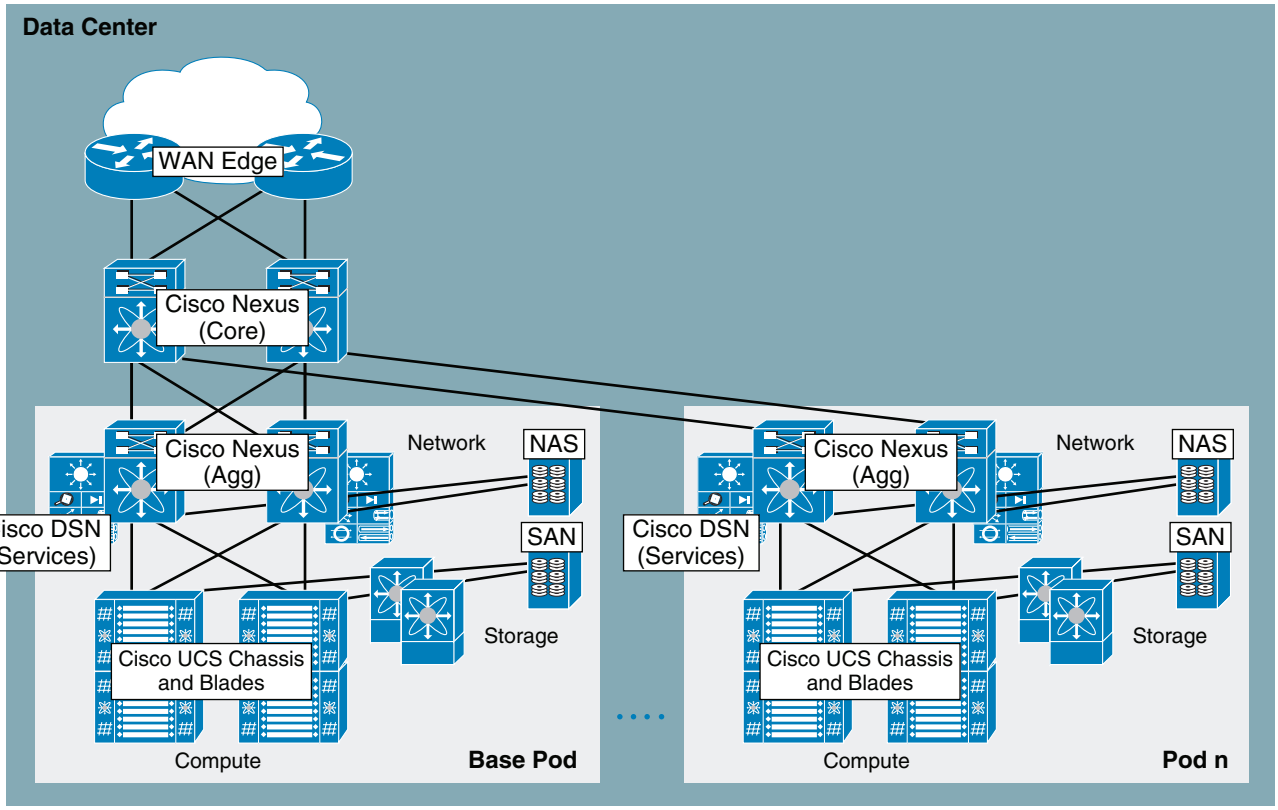

The Cisco VMDC architecture facilitates the movement of workloads from ICS to ICS within a pod, from ICS to ICS in different pods, and to an ICS in a different data center.

Hierarchical Network Design for Scalable Layer 2-Layer 3 Networking

Hierarchical designs have been commonly used in networking to achieve scale and high availability and are similarly used in Cisco VMDC to create a robust network framework for the data center. The Cisco VMDC network is organized into core, aggregation, and access layers, in a way similar to a campus network design with the exception that the term aggregation layer replaces the term distribution layer.

Core Layer

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 using 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.

Aggregation Layer

The aggregation layer of the data center network 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.

Access Layer

The access layer of the network provides connectivity for server farm 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 server farm 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 data center network is less concerned with distributing network access across multiple geographically disparate wiring closets and is more focused on aggregating server resources and providing an insertion point for shared data center services. This model uses redundant switches at each layer of the network topology for device-level failover creating 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 are 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. Figure 4 illustrates the network topology of a Cisco VMDC large pod design.

Figure 4 Hierarchical Network Design in Cisco VMDC Large Pod

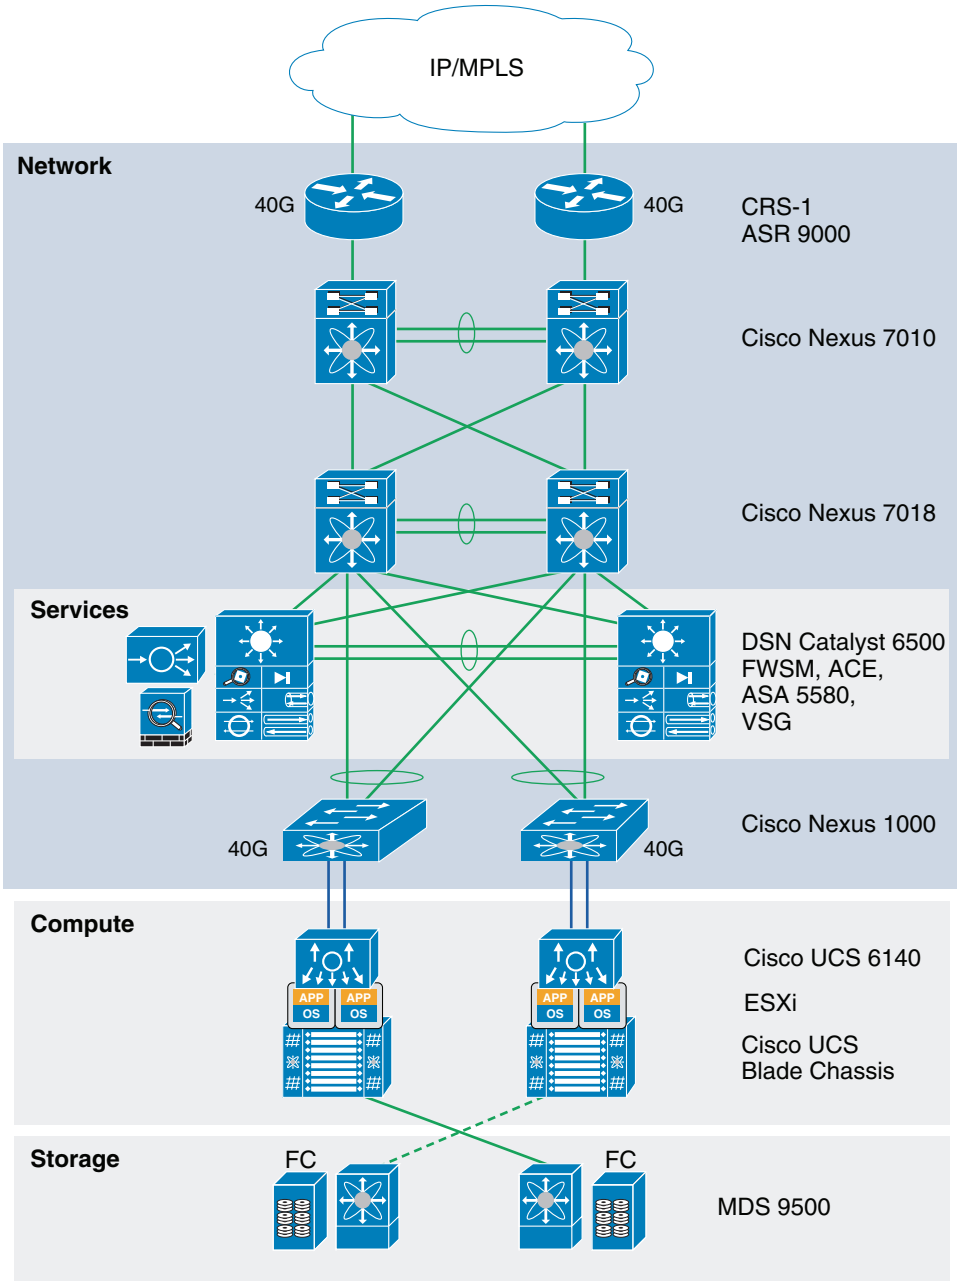

Security Features

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 Cisco VMDC environment. Virtualization at the various levels in the Cisco VMDC architecture provides logical separation in network, compute, and storage resources.

Figure 5 Security Framework of the Cisco VMDC Architecture

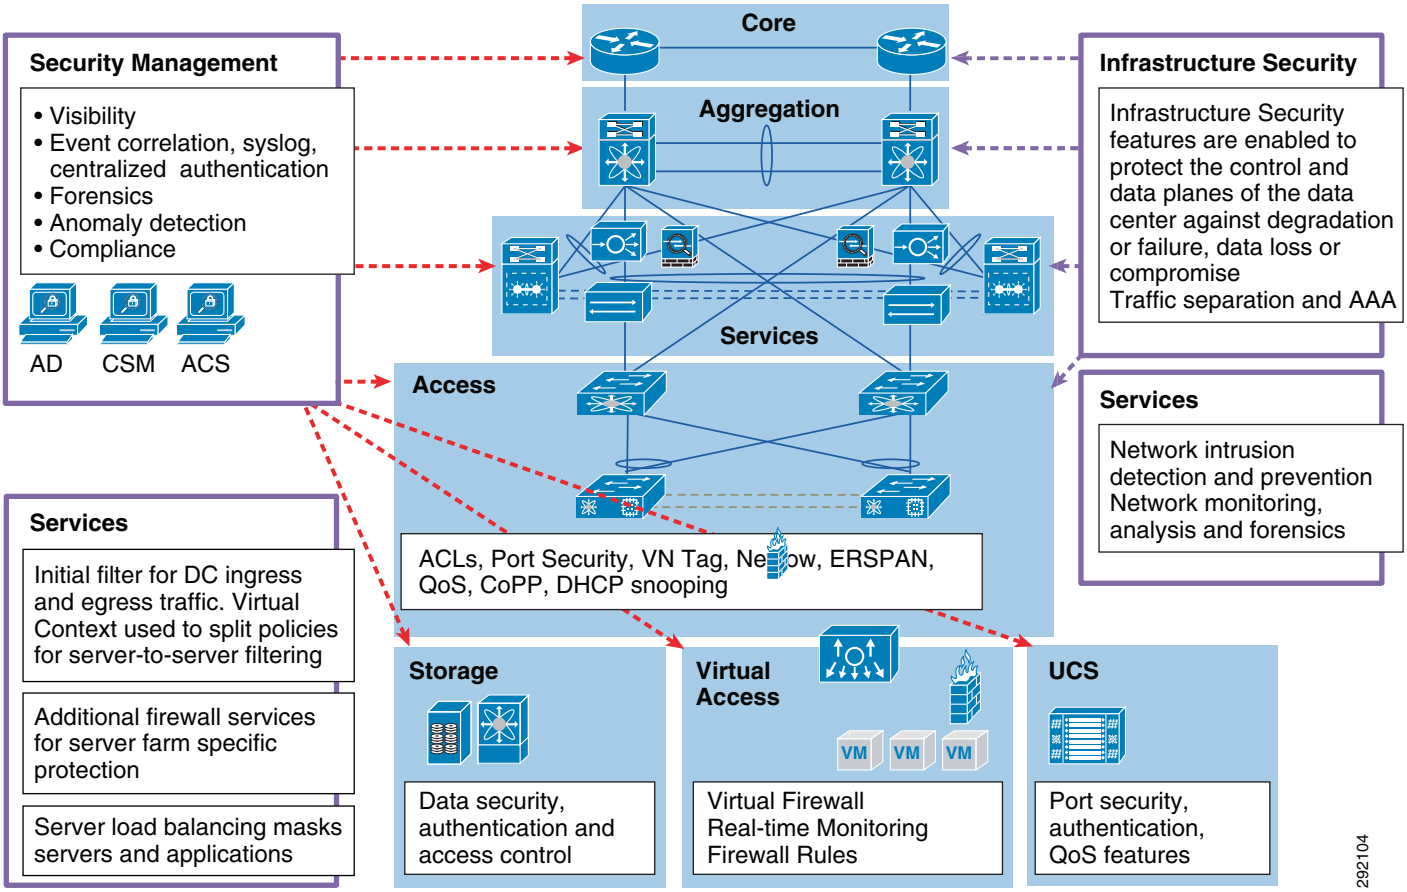

Services Layer—Layer 4-Layer 7 Services

The Cisco 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. Each tenant using the Cisco VMDC infrastructure is entitled to some compute, network, and storage resource SLA. 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.

Network and security services such as firewalls, server load balancers, intrusion prevention systems, application-based firewalls, and network analysis modules are typically deployed at the data center services layer. In the Cisco VMDC 2.0 large pod design, the Data Center Services Node (DSN),

composed of dual Cisco Catalyst® 6500 switches in a Virtual Switching System (VSS) mode, is utilized as the services layer. The Cisco Firewall Services Module (FWSM) and Application Control Engine (ACE) service modules in the DSN provide firewall and server load-balancing (SLB) services. The Cisco ASA5580 connected to the DSN provides secure remote access (IPsec-VPN and SSL-VPN) services so that remote clients can securely connect to the cloud resources.

Cisco Data Center Interconnect

Cisco Data Center Interconnect (DCI) solutions in the Cisco VMDC architecture enable cloud deployments to meet business continuity and corporate compliance objectives through disaster recovery, high availability, and data security mechanisms over geographically-distributed, multi-site data centers. These solutions transparently extend LAN and SAN connectivity and provide accelerated, highly-secure data replication, server clustering, and workload mobility between geographically dispersed data centers. DCI solution options include:

- Point-to-point or point-to-multipoint interconnection using virtual switching system (VSS), virtual PortChannel (vPC), and optical technologies
- Point-to-point interconnection using Ethernet over Multiprotocol Label Switching (EoMPLS) natively (over an MPLS core) and over a Layer 3 IP core
- Point-to-multipoint interconnections using virtual private LAN services (VPLS) or advanced VPLS (A-VPLS) natively (over an MPLS core) or over a Layer 3 IP core

The DCI capabilities are described and validated in detail in the Cisco DCI set of solutions (http://www.cisco.com/en/US/netsol/ns975/index.html).

Service Management and Automation

The Cisco VMDC architecture is complemented by a set of cloud automation and orchestration capabilities through partnership with BMC. These management and automation capabilities are described and validated in detail in other Cisco solutions. The BMC Cloud Lifecycle Management is a cloud service delivery environment which includes a service catalog that defines service offerings, a self-service portal for procuring resources, and management capabilities to control the cloud. BMC software increases service agility and reduces complexity through autoprovisioning and configuration of the end-to-end infrastructure supporting each cloud offering. The specific capabilities of the BMC software include:

- Virtual machine hosting and network orchestration with best-in-class scalability
- Multi-tenant partitioning through which customer and virtual machine traffic is fully segmented
- Various levels of security options with Layer 4 to Layer 7 integration rules
- Application delivery services through content load-balancing options
- Secure Shell (SSH) and SSL termination

These management and automation capabilities are described and validated in detail in the Cisco Cloud-O set of solutions (http://www.cisco.com/en/US/netsol/ns1103/networking_solutions_solution_category.html).

Solution Validation Scope

Cisco VMDC and related solutions are tested and validated for end-to-end functionality, real world scale, and optimized performance, including:

- Data center end-to-end functionality verification for SAN and NAS designs—End-to-end feature/integration validation including QoS for all data center network layers from access to WAN edge on all platforms, ESX/VM provisioning, bootup, and maintenance as well as SAN/NAS storage design verification.
- Disaster recovery scenario validation—Transparent movement of data center workloads for business continuance (active-backup scenario).
- Automation validation—Validation of service orchestration, portal, and service catalog validation with element manager integration for compute and network.
- Data Center Services functionality validation—Validation of service tier offerings with data center services node (firewall, load balancing)
- Failover scenario validation—Validation of redundancy designs (with baseline steady state traffic)—routing, vPC/MEC, ECMP, VSS, HSRP, Active-Active service modules, clustering
- Security validation—End-to-end security validation on various components
- Scalability verification—Multi-dimensional scalability (VLAN, MAC, HSRP, routes, contexts, VM) within scope of architecture

Cisco VMDC Releases

- Release 1.0, 1.1—Introduces foundational architecture for deploying virtualized and multi-tenant data centers for cloud-based services. It supports high availability, elasticity, and resiliency of virtualized compute, network, and storage services.
- Release 2.0—Enhances and expands release 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 two pod scale points, compact and large.
- Release 2.1—Generalizes and simplifies the release 2.0 architecture for a multi-tenant virtualized data centers used for private cloud. Improvements include multicast support, simplified network design, jumbo frame support, improved convergence, performance, and scalability for private cloud, QoS best practices, and increased design flexibility with multi-tenant design options.
- Release 2.2—Builds on top of releases 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.

Cisco VMDC Documentation

The Cisco VMDC architecture conforms to the Cisco Validated Design™ (CVD) guidelines for end-to-end system validation and documentation. CVD status indicates that the solution has undergone rigorous design analysis and validation testing to ensure end-to-end completeness of functionality as well as overall quality. The results of these design and validation steps are documented in the following sets of documents which get iterated for every release as needed:

- Solution Data Sheets—Very brief documents highlighting the most prominent details of a release. These documents can be used as a quick at-a-glance reference of a given release. They include a list of specific component hardware and software releases, basic solution topology, validated scale profiles, service profiles, etc.
- Solution Design Guides—Describe significant architectural considerations, choices, and decisions that are made to define a specific release. Architectural considerations include platform choices, end-to-end functionality, scale, availability, security, and manageability of a given solution.
- Solution Implementation Guides—Describe specific profiles that get tested and validated in the lab with the respective test results. The validation tests include those listed in the solution validation scope above.
- Technical White Papers—Cover specific technical or business details in more depth.

For More Information

For more information about Cisco VMDC and access to the relevant documents, visit http://www.cisco.com/go/vmdc or contact your Cisco account representative.