

VMDC Architecture Overview

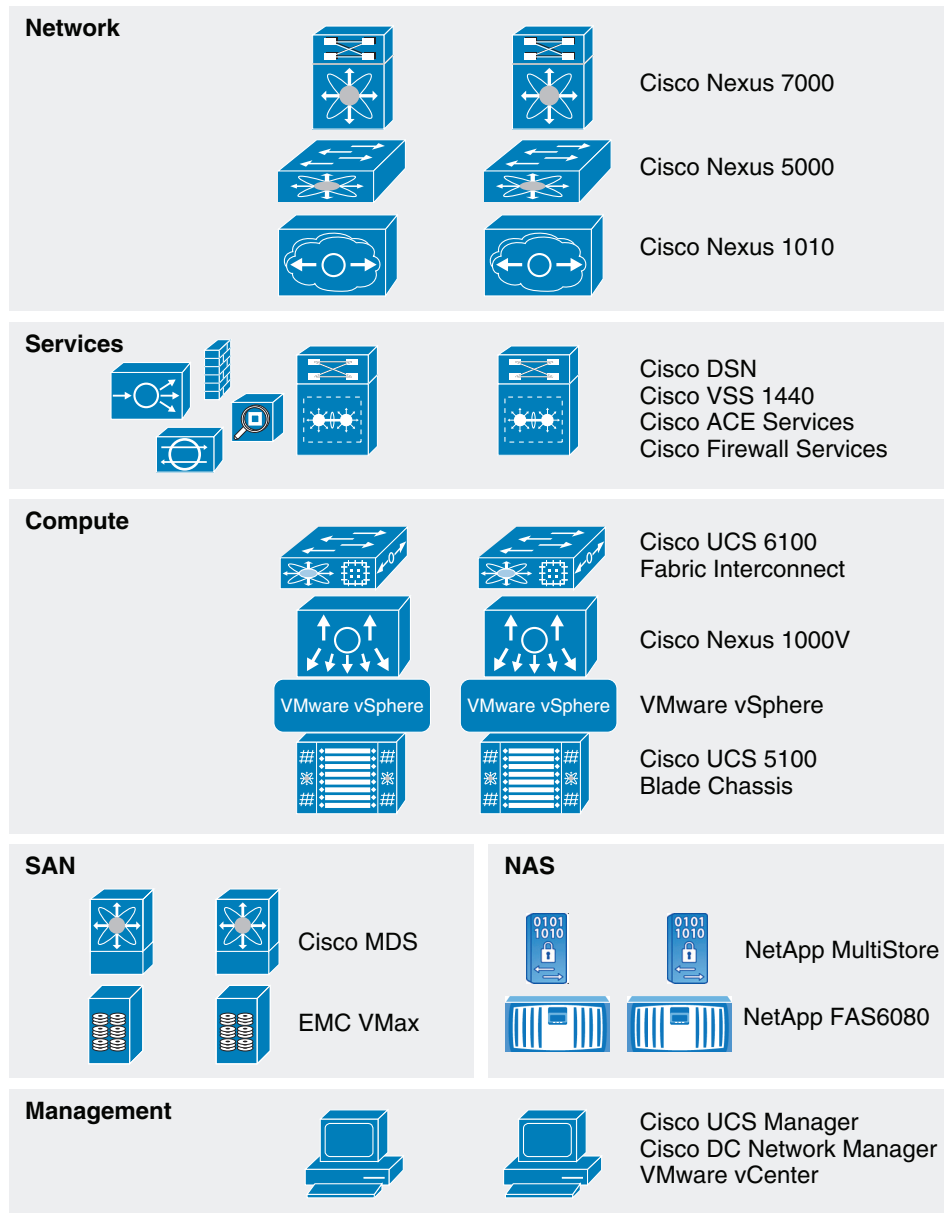
Cisco VMDC 2.1 is a reference architecture that helps customers design, deploy, and implement data center network, compute, storage and management while demonstrating the integration points between each. This architecture demonstrates how enterprise and mid-market customers can build highly resilient, 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 multsocket, 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

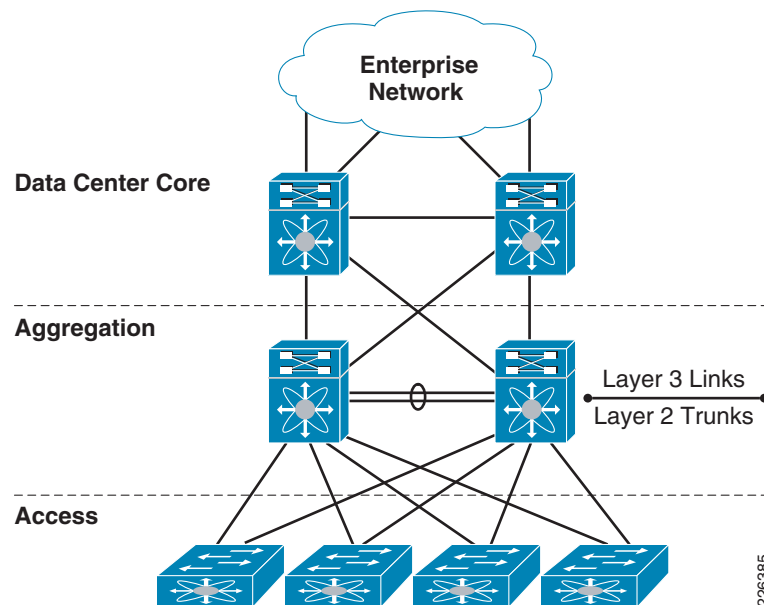
Hierarchical network design has been commonly used in networking for many years. This model uses redundant switches at each layer of the network topology for device-level failover that creates a highly available transport between end nodes. Data center networks often require additional services beyond 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



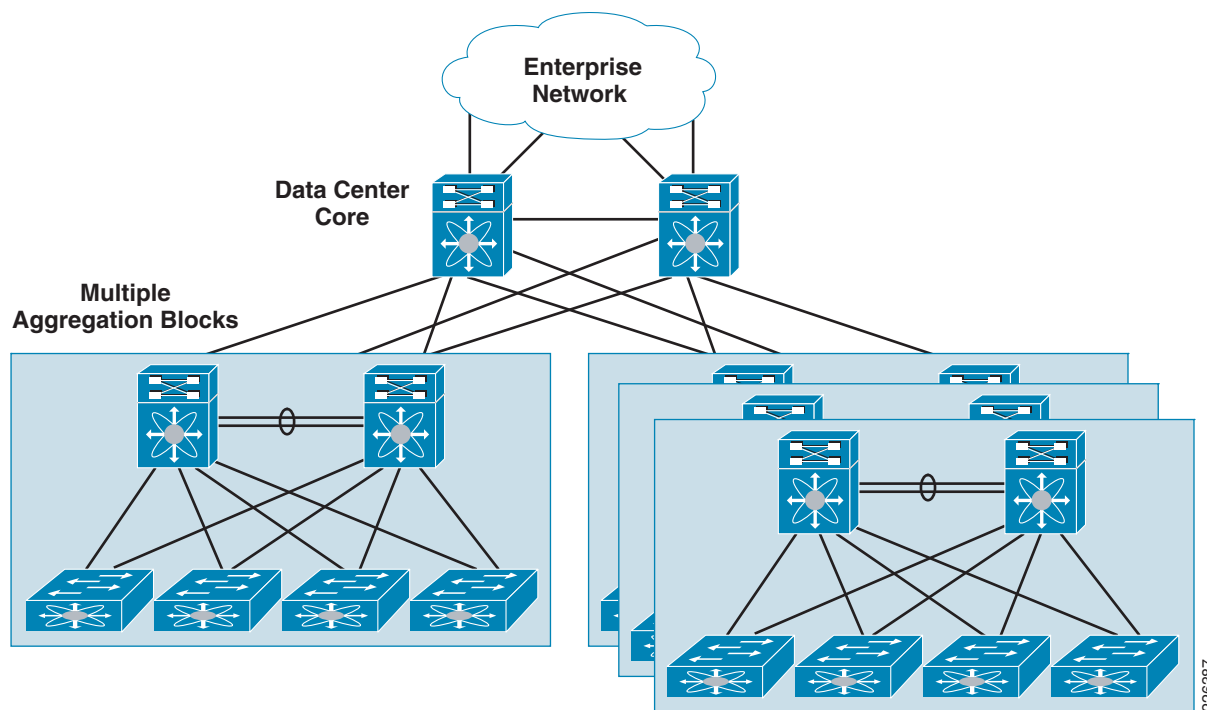
Core Layer

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, 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.

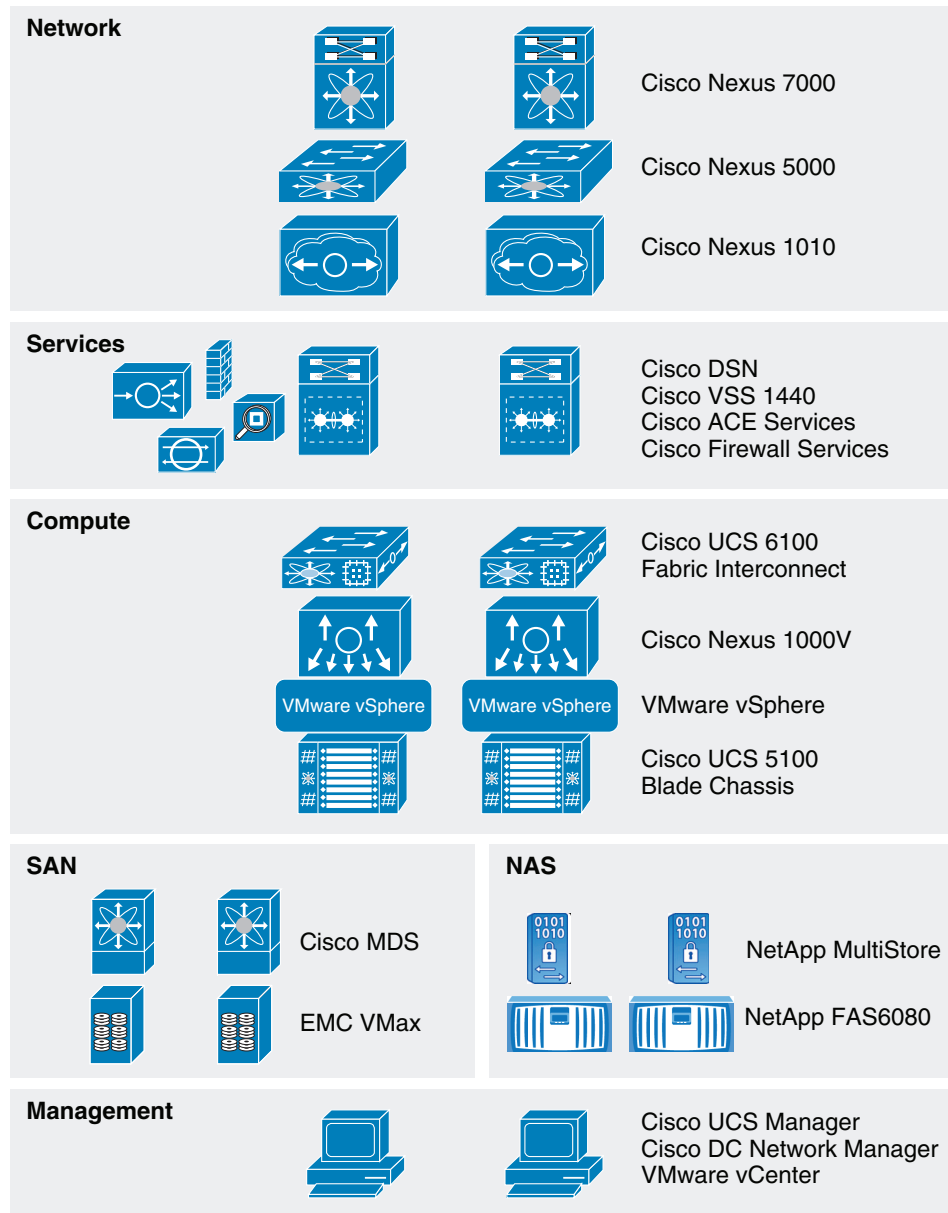
Modular Building Blocks

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:

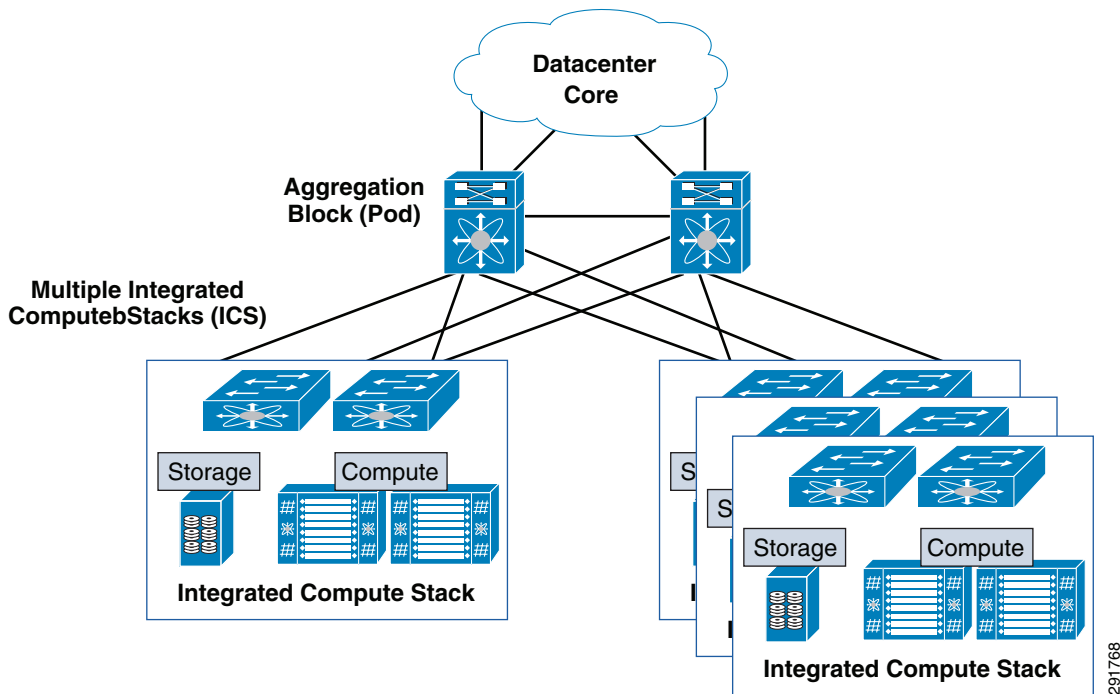
- Minimize operational impact; reduce total cost of ownership (TCO)
- Flexible, multi-vendor architecture
- Pretested and validated IT infrastructure
- Private cloud foundation

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

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

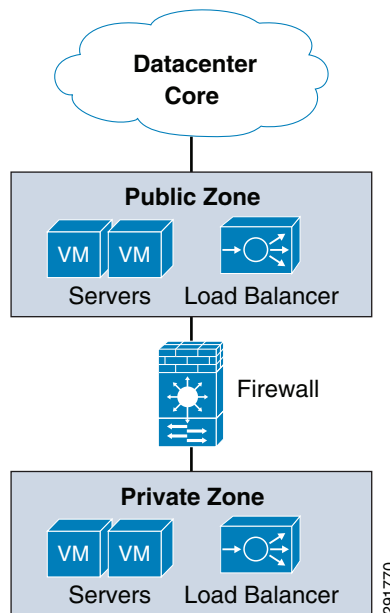
Virtualized Multi-Tenancy

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:

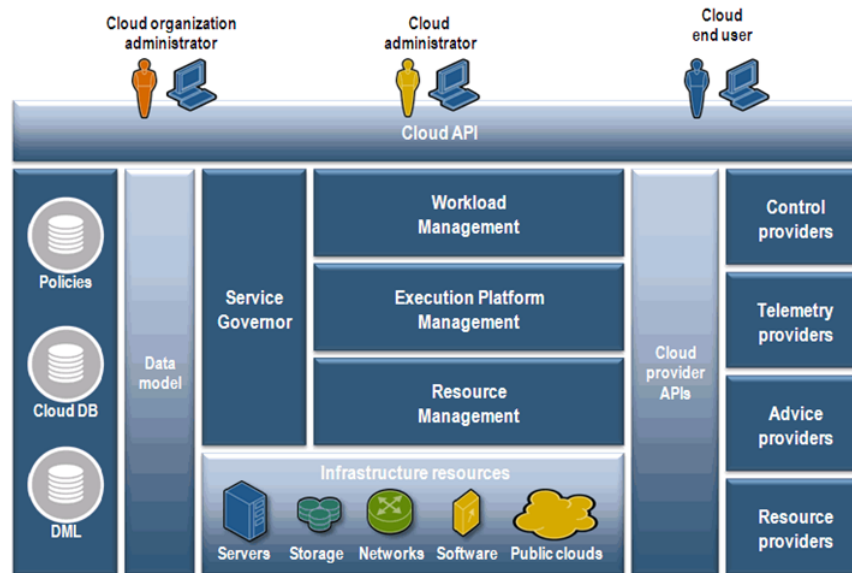
- **Virtual Server**—Memory, CPU, Capacity management, Storage allocation

- **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.