White Paper

Cisco UCS and EMC® VNX™ 5300 with Microsoft Private Cloud Fast Track 2.0

July 2012

Contents

Introduction

The Microsoft Private Cloud Fast Track Program is a joint effort between Microsoft and its partners to help organizations quickly develop and implement private clouds, while reducing both the cost and the risk. It provides a reference architecture that combines Microsoft software, consolidated guidance, and validated configurations with partner technology—such as computing power, network and storage architectures, and value-added software components.

The private cloud model provides much of the efficiency and agility of cloud computing along with the increased control and customization achieved through dedicated private resources. With the Microsoft Private Cloud Fast Track Program (Fast Track), Microsoft and its partners provide organizations both the control and the flexibility required to reap the full benefits of the private cloud.

Cisco and EMC Private Cloud Fast Track Description

The Cisco and EMC jointMicrosoft Private Cloud Fast Track solution is a reference archtitecure which fulfils and delivers on the value of the Microsoft Provate Cloud Fast Track program requirements. This document describes the solution architecture in detail and includes content from Cisco, EMC, and Microsoft.

Cisco and EMC with Microsoft Private Cloud Fast Track utilizes the core capabilities of Windows Server 2008 R2 SP1, Hyper-V, and Microsoft System Center 2012 to deliver a private cloud infrastructure as a service (IaaS) offering. The solution also includes software from Cisco and EMC to form a complete solution that is ready for your enterprise.

Business Value

The Cisco and EMC with Microsoft Private Cloud Fast Track solution provides a reference architecture for building private clouds that addresses an organization's unique requirements. Each Fast Track solution helps organizations implement private clouds with increased ease and confidence. Among the benefits of the Microsoft Private Cloud Fast Track Program are faster deployment, reduced risk, and a lower cost of ownership.

Reduced risk:

- Tested, end-to-end interoperability of compute, storage, and network
- Predefined, out-of-box solutions based on a common cloud architecture that has already been tested and validated
- High degree of service availability through automated load balancing

Lower cost of ownership:

- A cost-optimized, platform and software-independent solution for rack system integration
- High performance and scalability with Windows Server 2008 R2 SP1 operating system with Hyper-V technology
- Minimized backup times and fulfilled recovery time objectives for each business critical environment

Technical Benefits

The Microsoft Private Cloud Fast Track Program integrates best-in-class hardware implementations with Microsoft's software to create a Reference Implementation. This solution has been co-developed by Microsoft, Cisco and EMC and and has gone through a validation process. As a Reference Implementation, Microsoft, Cisco and EMC have taken the work of building a private cloud that is ready to meet a customer's needs.

Faster deployment:

- End-to-end architectural and deployment guidance
- Streamlined infrastructure planning due to predefined capacity
- Enhanced functionality and automation through deep knowledge of infrastructure
- Integrated management for virtual machine (VM) and infrastructure deployment
- Self-service portal for rapid and simplified provisioning of resources

Technical Overview

NIST Definition of Cloud Computing

Note: The following text is a verbatim copy of the NIST Definition of Cloud Computing v15.

Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model is composed of five essential characteristics, three service models, and four deployment models.

Essential Characteristics:

On-demand self-service. A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service provider.

Broad network access. Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, tablets, laptops, and workstations).

Resource pooling. The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or data center). Examples of resources include storage, processing, memory, and network bandwidth.

Rapid elasticity. Capabilities can be elastically provisioned and released, in some cases automatically, to scale rapidly outward and inward commensurate with demand. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be appropriated in any quantity at any time.

Measured service. Cloud systems automatically control and optimize resource use by leveraging a metering capability[1] at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.

Service Models:

[1] Typically this is done on a pay-per-use or charge-per-use basis.

Software as a Service (SaaS). The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure[2]. The applications are accessible from various client devices through either a thin client interface, such as a web browser (e.g., web-based email), or a program interface. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

Platform as a Service (PaaS). The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services, and tools supported by the provider[3]. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.

Infrastructure as a Service (IaaS). The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, and deployed applications; and possibly limited control of select networking components (e.g., host firewalls).

Deployment Models:

Private cloud. The cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (e.g., business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises.

Community cloud. The cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises.

Public cloud. The cloud infrastructure is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider.

Hybrid cloud. The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds).

Overview of Microsoft Private Cloud

Private cloud is a computing model that uses resources that are dedicated to your organization. As Figure 1 illustrates, a private cloud shares many of the characteristics of public cloud computing, including resource pooling, self-service, elasticity, and metered-by-use service delivered in a standardized manner, with the additional control and customization available from dedicated resources.

[2] A cloud infrastructure is the collection of hardware and software that enables the five essential characteristics of cloud computing. The cloud infrastructure can be viewed as containing both a physical layer and an abstraction layer. The physical layer consists of the hardware resources that are necessary to support the cloud services being provided, and typically includes server, storage, and network components. The abstraction layer consists of the software deployed across the physical layer which manifests the essential cloud characteristics. Conceptually, the abstraction layer sits above the physical layer.

[3] This capability does not necessarily preclude the use of compatible programming languages, libraries, services, and tools from of the sources.

Figure 1. The Attributes of Private Cloud Computing

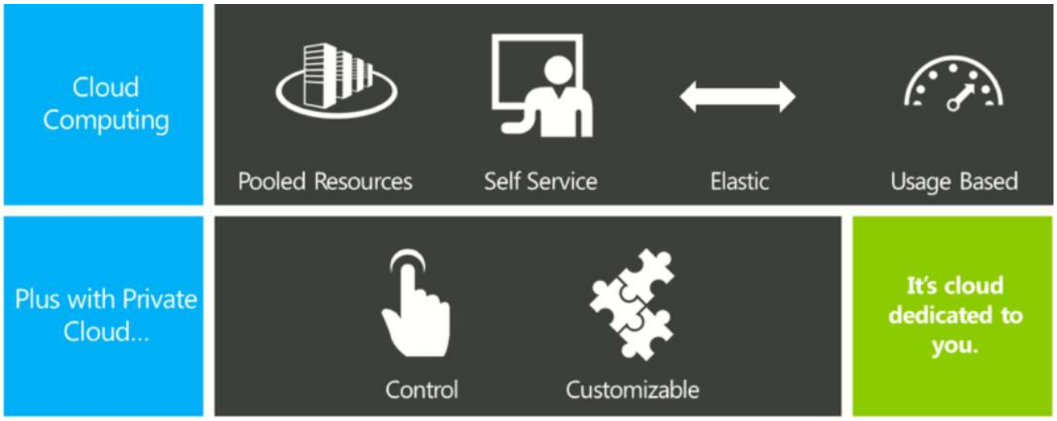

While virtualization is an important technological component of private cloud computing, the key differentiator is the continued abstraction of computing resources from infrastructure and the machines (virtual or otherwise) used to deliver those resources. Only with this abstraction can customers achieve the benefits of private cloud, including improved agility and responsiveness, reduced TCO, and increased business alignment and focus. Most importantly, a private cloud promises to exceed the cost effectiveness of a virtualized infrastructure through higher workload density and greater resource utilization.

The Microsoft Private Cloud is a unique and comprehensive offering, built on four foundations:

- **All About the App:** An application-centric cloud platform that helps you focus on business value.
- **Cross-Platform from the Metal Up:** Microsoft Private Cloud provides cross-platform support for multihypervisor environments, operating systems, and application frameworks.
- **Foundation for the Future:** Microsoft Private Cloud lets you go beyond virtualization to a true cloud computing platform.
- **Cloud on Your Terms:** Microsoft private cloud lets you consume cloud services on your terms, providing you with the choice and flexibility of a hybrid cloud model through common management, virtualization, identity services, and developer tools.

For further reading, please refer to the Microsoft Private Cloud Overview.

Private Cloud Architecture Principles

Resource Pooling

Resource optimization is a principle that drives efficiency and cost reduction and is primarily achieved through resource pooling. Abstracting the platform from the physical infrastructure enables optimization of resources through shared use. Allowing multiple consumers to share resources results in higher resource utilization and a more efficient and effective use of the infrastructure. Optimization through abstraction provides critical support for many of the Microsoft private cloud principles and ultimately helps drive down costs and improve agility.

Elasticity and Perception of Infinite Capacity

From a consumer's perspective, cloud services appear to have infinite capacity. The consumer can use as much or as little of the service as needed. Just as a consumer uses as much electricity as they need from the electric utility provider, so end users can consume cloud-based services on demand.

This utility mindset requires that capacity planning be proactive so that requests can be satisfied on demand. Applying this principle reactively and in isolation often leads to inefficient use of resources and unnecessary costs. Combined with other principles, such as using incentives to encourage desired consumer behavior, this principle allows for a balance between the cost of unused capacity and the need for agility.

Perception of Continuous Availability

From the consumer's perspective, cloud services appear to always be available when needed. The consumer should never experience an interruption of service, even if failures occur within the cloud environment. To achieve this perception, a provider must have a mature service management approach combined with inherent application resiliency and infrastructure redundancies in a highly automated environment. Much like the perception of infinite capacity, this principle can only be achieved in conjunction with the other Microsoft private cloud principles.

Predictability

Predictability is critical for any cloud-computing system, whether you are a consumer or provider. From the vantage point of the consumer, cloud services should be consistent; they should have the same quality and functionality any time they are used.

A provider must deliver an underlying infrastructure that ensures a consistent experience to the hosted workloads in order to achieve this predictability. This consistency is achieved through the homogenization of underlying physical servers, network devices, and storage systems.

From the service management perspective, providers drive predictability by standardizing their service offerings as well as by standardizing their processes. Predictability is essential for delivering service quality.

Metering/Chargeback: Taking a Service Provider's Approach to Delivering IT

Historically, when IT has been asked to deliver a service to the business, they purchase the necessary components and then build an infrastructure specific to the service requirements. This approach results in longer time to market and increased costs, so that often business expectations of agility and cost reduction are not met. Further compounding the problem, this model is often used when an existing service needs to be expanded or upgraded.

In contrast, when you take a "service provider's approach" by providing infrastructure as a service, everything is transformed. IT can now take advantage of a shared resource model that enables economies of scale. Combined with the other principles, the "as a service" model makes it possible for IT to achieve much greater agility in delivering services than it could previously.

Multitenancy

Multitenancy occurs when the infrastructure is logically subdivided and provisioned to different organizations or organizational units. The traditional example is a hosting company that provides servers to multiple customer organizations. Increasingly, enterprises and other entities that have a centralized IT organization use the multitenacy model to provide services to multiple business or internal organizational units, treating each as a customer or tenant.

Security and Identity

Security for the Microsoft private cloud is founded on:

- Protected infrastructure
- Application access
- Network access

Protected infrastructure uses security technologies as well as identity technology to help ensure that hosts, information, and applications are secured across all scenarios in the data center. Protected infrastructure includes physical (on premises) resources and virtual (on premises and cloud) resources.

Application access helps to ensure that IT can extend vital application access not only to internal users, but also to vital business partners and cloud users.

Network access uses an identity-centric approach to help ensure that users, whether they're at the office or in remote locations, have secure access on numerous devices, so that business gets done the way it should to maintain productivity.

Most important is that the secure data center uses a common, integrated technology to ensure that users have simple access with a common identity across all environments. It's also crucial that management is integrated across physical, virtual, and cloud environments so that business can take advantage of all capabilities without requiring additional significant financial investments.

Dynamic Data Center IaaS Reference Model

Infrastructure as a service (IaaS) is the application of the private cloud architecture principles we just described to deliver infrastructure. As the cloud ecosystem matures, products broaden and deepen in features and capabilities. System architects can use the reference model shown in Figure 2 as a guide to developing a holistic solution spanning all the layers required for mature IaaS. Note that this model is used as a reference tool only: based on experience operating private clouds in real-world environments, technical reference architecture emphasizes some pieces are emphasized more than others in the model shown in Figure 2.

Figure 2. The Dynamic Data Center (DDC) Reference Model: IaaS View

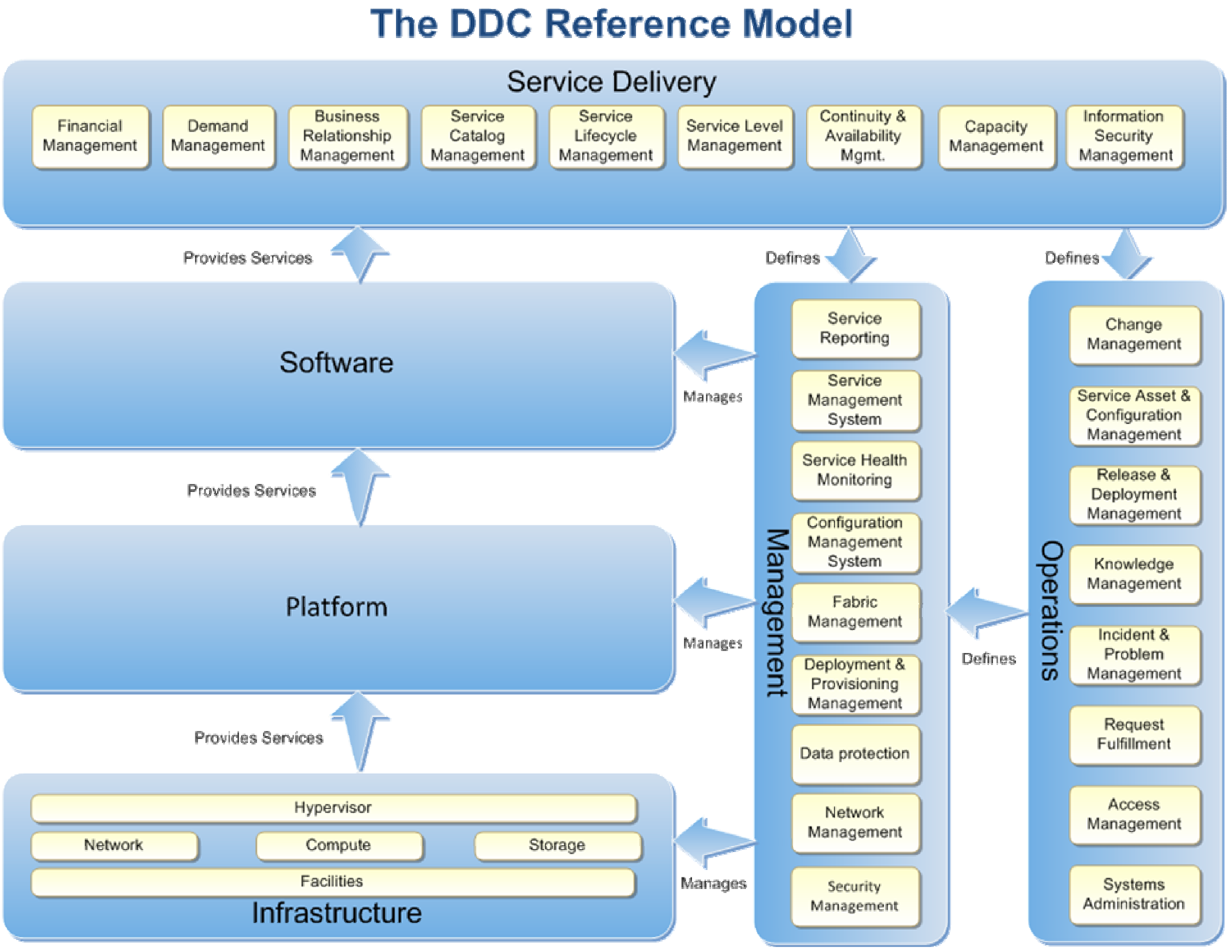

The DDC Reference Model is split into the following layers:

The software, platform, and infrastructure layers represent the technology stack, where each provides services to the layer above.

The service operations and management layers represent the process perspective and include the management tooling required to implement aspects of the process.

The service delivery layer represents the alignment between business and information technology (IT).

The DDC Reference Model is a deliberate attempt to blend technology and process perspectives (for example, Information Technology Infrastructure Library [ITIL] and Microsoft Operations Framework [MOF]) because cloud computing is as much about service management as it is about the technologies involved in it.

For further reading, please see: Private Cloud Reference Model

Conceptual Architecture

One of the goals of the layered infrastructure is to support the development of complex workflow and automation over time by creating a collection of simple automation tasks, assembling them into procedures that are managed by the management layer, and then creating workflows and process automation that are controlled by the orchestration layer.

Fabric

Scale Units

In a modular architecture, the concept of a scale unit refers to the point to which a module can increase in capability—in other words, the point to which it can scale—before another module is required. For example, an individual server is a scale unit. It can be expanded to a certain point in terms CPU and RAM, but beyond its maximums, an additional server is required to continue scaling. Each scale unit also has associated certain amount of physical installation labor, configuration labor, and so on, associated withit. In general, the bigger the scale unit—for example, a full rack of preconfigured servers—the more labor overhead can be minimized.

It is critical to know the scale limits of all components, both hardware and software, so that you can determine the best scale units for the architecture overall. Scale units make it possible to document all the requirements (space, power, HVAC, connectivity, etc.) required for implementation.

Servers

The hardware-architecture choices that are available to data center architects are constantly evolving. Choices range from rack-mounted servers to tightly integrated, highly redundant blade systems or container models. The same spectrum of choices exists for storage and networking equipment.

Server scale limits are well published and include the number and speed of CPU cores, maximum amount and speed of RAM, the number and type of expansion slots, and so on. Particularly important are the number and type of onboard I/O ports as well as the number and type of supported I/O cards. Both Ethernet and Fibre Channel expansion cards often provide multiport options where a single card can have four ports. Additionally, in blade server architectures, there are often limitations in the I/O cards supported and number of I/O card combinations. It is important to be aware of these limitations as well as the oversubscription ratio between blade I/O ports and any blade chassis switch modules.

A single server is not typically a good scale unit for a private cloud solution because of the amount of overhead required to install and configure an individual server.

Storage

Storage architecture is a critical design consideration for private cloud solutions. The topic is challenging as storage solutions are rapidly evolving in terms of new standards, protocols, and implementations. Storage and supporting storage networking are critical to the overall performance of the environment and overall TCO of the environment. Storage architectures today have several layers, including the storage array(s), the storage network, the storage protocol, and for virtualization, the file system utilizing the physical storage.

One of the primary objectives of the private cloud solution is to enable rapid provisioning and deprovisioning of virtual machines. Doing so at large scale requires tight integration with the storage architecture and robust automation. Provisioning a new virtual machine on an already existing logical unit number (LUN) is a simple operation; however, tasks like provisioning a new LUN and adding it to a host cluster are relatively complicated and also greatly benefit from automation.

Networking

Many network architectures include a tiered design with three or more tiers, such as core, distribution, and access. Designs are driven by the port bandwidth and the quantity of ports required at the edge, as well as the ability of the distribution and core tiers to provide higher-speed uplinks to aggregate traffic. Additional considerations include Ethernet broadcast boundaries and limitations, and spanning treeor other loop avoidance technologies.

A dedicated management network is a frequent feature of advanced data center virtualization solutions. Most virtualization vendors recommend that hosts be managed via a dedicated network so that there is not competition with guest traffic needs and to provide a degree of separation for security and ease of management purposes. This typically implies dedicating one network interface card (NIC) per host and port per network device to the management network.

With advanced data center virtualization, a frequent use case is to provide isolated networks dedicated to different owners such as particular departments or applications. Multitenant networking refers to using technologies such as VLANs or IPsec isolation techniques to provide dedicated networks that utilize a single network infrastructure or wire.

Managing the network environment in an advanced data center virtualization solution can present challenges. Ideally, network settings and policies are defined centrally and applied universally by the management solution. In the case of IPsec-based isolation, this can be accomplished using Microsoft Active Directory and group policy to control firewall settings across the hosts and guest settings, as well as the IPsec policies controlling network communications.

For VLAN-based network segmentation, several components, including the host servers, host clusters, Virtual Machine Manager (VMM), and the network switches, must be configured correctly to enable both rapid provisioning and network segmentation. With Hyper-V and host clusters, identical virtual networks must be defined on all nodes in order for a virtual machine to be able to failover to any node and maintain its connection to the network. At large scale, this can be accomplished via PowerShell scripting.

Virtualization

The virtualization layer is one of the primary enablers in environments with greater IT maturity. The decoupling of hardware, operating systems, data, applications, and user state opens a wide range of options for better management and distribution of workloads across the physical infrastructure. Hypervisor-based virtualization technologies make it possible for the virtualization layer to migrate running VMs from one server to another with zero downtime. Hypervisor virtualization also provides many other features, giving the system a rich set of capabilities. These capabilities can be utilized by the automation, management, and orchestration layers to maintain desired states (such as load distribution) or to proactively address decaying hardware or other issues that would otherwise cause faults or service disruptions.

As with the hardware layer, it is critical that the virtualization layer can be managed by the automation, management, and orchestration layers. The abstraction of software from hardware that virtualization provides moves the majority of management and automation tasks into the software space, instead of requiring people to perform manual operations on physical hardware.

Management

Fabric Management

In fabric management, discrete capacity pools of compute (servers), storage, and network resources are treated as a single fabric. The fabric is subdivided into capacity clouds, or resource pools, that carry certain characteristics such as delegation of access and administration, service-level agreements (SLAs), cost metering, and so on. This enables the centralization and automation of complex management functions that can be carried out in a highly standardized and repeatable fashion, which increases availability and lowers operational costs.

Process Automation and Orchestration

Orchestration— is the management of all the automation and management component which must be implemented as the interface between the IT organization and the infrastructure. Orchestration provides the bridge between IT business logic, such as "deploy a new web-server VM when capacity reaches 85 percent," and the dozens of steps in an automated workflow that are required to actually implement such a change.

Ideally, the orchestration layer provides a graphical interface in which complex workflows that consist of events and activities across multiple management-system components can be combined, so as to form an end-to-end IT business process, such as automated patch management or automatic power management. The orchestration layer must provide the ability to design, test, implement, and monitor these IT workflows

Service Management System

A service management system is a set of tools designed to facilitate service management processes. Ideally, these tools should be able to integrate data and information from the entire set of tools found in the management layer. The system should also be capable of processing and presenting the data as needed. At a minimum, the service management system should be linked to the configuration management system (CMS), commonly known as the configuration management database (CMDB), and should log and track incidents, problems, and changes. It is also preferable that the service management system be integrated with the service health modeling system so that incident tickets can be auto-generated.

User Self-Service

Self-service capability is a characteristic of private cloud computing and must be present in any implementation. The intent is to get users as close as possible to a self-service capability by presenting them with the options available for provisioning in an organization. The capability may be basic — for instance, allowing users to provision a virtual machine with only with a predefined configuration—it may be more advanced, allowing user to set the base configuration or even to configure a platform’s capability or service.

Self-service capability is a critical business driver that enables members of an organization to become more agile in responding to business needs using IT capabilities aligned with internal business IT requirements and governance.

This means that the interface between IT and the business are abstracted to a simple, well-defined, and approved set of service options and that these options are presented as a menu in a portal or available from the command line. The business selects these services from the catalog, begins the provisioning process, and is notified upon completion; the business is then only charged for what they actually use. This is analogous to capability available on public cloud platforms.

Service Delivery

Unlike traditional virtualized environments, private clouds are delivered as a service. This implies that you have a managed service delivery capability. What services do you offer, and what features or elements do they contain? Management of service delivery is discussed in this section.

Service Catalog

Service catalog management involves defining and maintaining a catalog of services offered to consumers. This catalog lists the following:

- The classes of service available
- The requirements to be eligible for each service class
- The service level attributes and targets that are included with each service class (such as availability targets)
- The cost model for each service class

The service catalog may also include specific virtual machine (VM) templates, such as a high-compute template, designed for different workload patterns. Each template will define the VM configuration specifics, such as the amount of allocated CPU, memory, and storage.

Capacity Management

Capacity management defines the processes necessary to achieve the perception of infinite capacity. Capacity must be managed to meet existing and future peak demand while controlling under-utilization. Business relationship and demand management are key inputs into effective capacity management and require a service provider's approach. Predictability and optimization of resource usage are primary principles in achieving capacity management objectives.

Availability Management

Availability management defines processes necessary to achieve the perception of continuous availability. Continuity management defines how risk will be managed in a disaster scenario to make sure minimum service levels are maintained. The principles of resiliency and automation are fundamental here.

Service-Level Management

Service-level management is the process of negotiating service-level agreements (SLAs) and making sure the agreements are met. SLAs define target levels for cost, quality, and agility by service class as well as the metrics for measuring actual performance. Managing SLAs is necessary for achieving the perception of infinite capacity and continuous availability. This, too, requires a service provider's approach by IT.

Service Lifecycle Management

Service Lifecycle Management takes an end-to-end management view of a service. A typical journey starts with identification of a business need, and continues through business relationship management, to the time when that service becomes available. Service strategy drives service design. After launch, the service is transitioned to operations and refined through continual service improvement. Taking a service provider's approach is key to successful service lifecycle management.

Operations

After the system is deployed, it must be operated correctly. The processes and tools described in this section help support the proper post-deployment operation of the overall system.

Change Management

Change Management process is responsible for controlling the lifecycle of all changes. The primary objective of Change Management is to eliminate or at least minimize disruption while desired changes are made to services. Change Management focuses on understanding and balancing the cost and risk of making the change versus the benefit of the change to either the business or the service. Driving predictability and minimizing human involvement are the core principles for achieving a mature Service Management process and ensuring changes can be made without impacting the perception of continuous availability.

Incident and Problem Management

The goal of incident management is to resolve events that are impacting, or threaten to impact, services as quickly as possible and with minimal disruption. The goal of problem management is to identify and resolve root causes of incidents that have occurred, as well as to identify and prevent or minimize the impact of incidents that may occur.

Configuration Management

Configuration management is the process responsible for ensuring that the assets required to deliver services are properly controlled, and that accurate and reliable information about those assets is available when and where it is needed. This information includes details of how the assets have been configured and the relationships between assets.

This typically requires a configuration management database (CMDB to store configuration records throughout their lifecycle. The configuration management system maintains one or more configuration management databases, and each database stores the attributes of configuration items and relationships with other configuration items.

Reference Architecture

The Cisco and EMC architecture is highly modular. Although each customer’s components might vary in its exact configuration, after a Cisco and EMC configuration system is built, it can easily be scaled as requirements and demands change. This includes both scaling up (adding additional resources within a Cisco UCS® chassis and/or EMC VNX array) and scaling out (adding additional Cisco UCS chassis and/or EMC VNX array).

The joint Cisco and EMC solution validated with Microsoft Private Cloud includes EMC® VNX™ 5300 storage, Cisco Nexus® 5000 Series network switches, the Cisco Unified Computing System™ (Cisco UCS) platforms, and Microsoft virtualization software in a single package (see Figure 3). The computing and storage can fit within one data center rack, with networking residing in a separate rack or deployed according to a customer’s data center design. Because of port density, the networking components can accommodate multiple configurations of this kind.

Figure 3. Cisco and EMC Reference Configuration

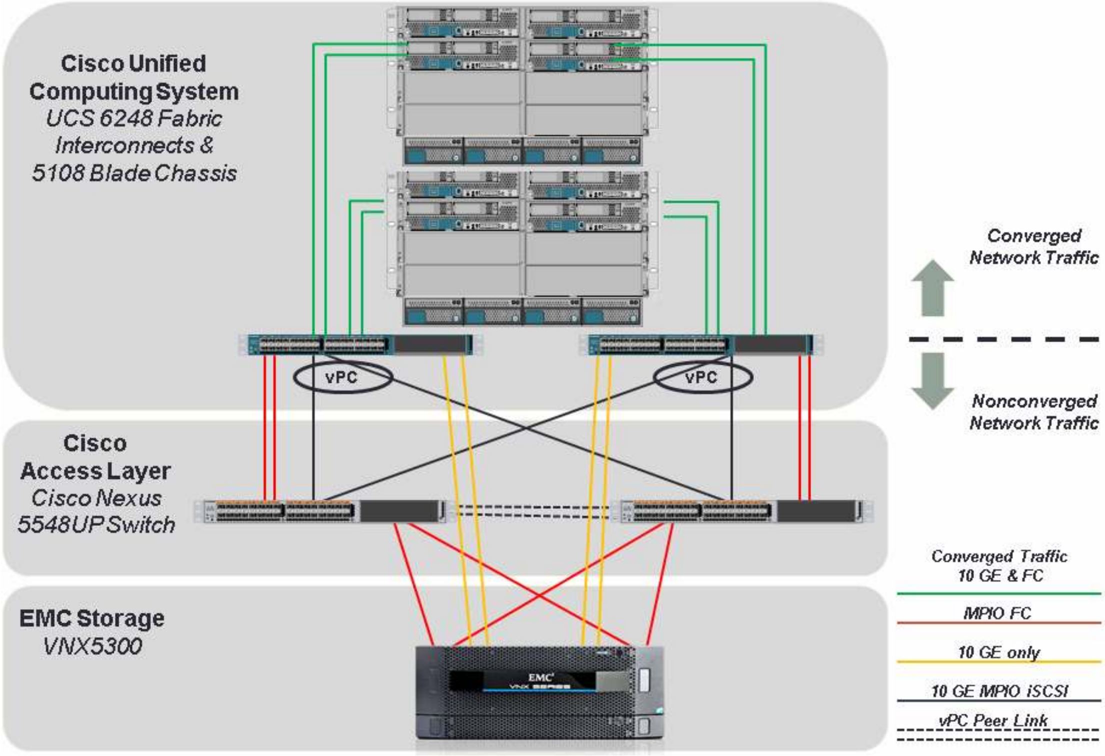
Cisco Unified Computing System
UCS 6248 Fabric Interconnects & 5108 Blade Chassis
vPC
vPC
Converged Network Traffic
Nonconverged Network Traffic
Cisco Access Layer
Cisco Nexus 5548UP Switch
EMC Storage
VNX5300
EMC
VNX SERIES
Converged Traffic 10 GE & FC
MPIO FC
10 GE only
10 GE MPIO iSCSI
vPC Peer Link

The reference configuration in Figure 3 contains the following components:

Two Cisco Nexus 5548UP Switches

Two Cisco 6248UP Fabric Interconnects

Two Cisco UCS 5108 Blade Chassis with two fabric extenders per chassis

Eight Cisco UCS B200 M2 Blades Servers, four in each chassis

One EMC VNX5300 Unified Storage Platform

Storage is provided by an EMC VNX5300 (high availability [HA] within a single chassis) with accompanying disk shelves. All systems and fabric links are configured to provide redundancy for end-to-end high availability. Server virtualization is provided by Microsoft’s Hyper-V hypervisor. This is a base design that can be easily reconfigured to adapt to changing or different business needs. For example, room is left in the chassis for additional blades.

The remainder of this document details the overall architecture and components of the reference architecture in Figure 3.

Use Cases

Service Models

Figure 4 depicts the taxonomy of cloud services, and helps clearly define the separation of responsibilities in each service model. Please see the section NIST Definition and Cloud Computing in this guide for more information on the service models.

Figure 4. Taxonomy of Cloud Services

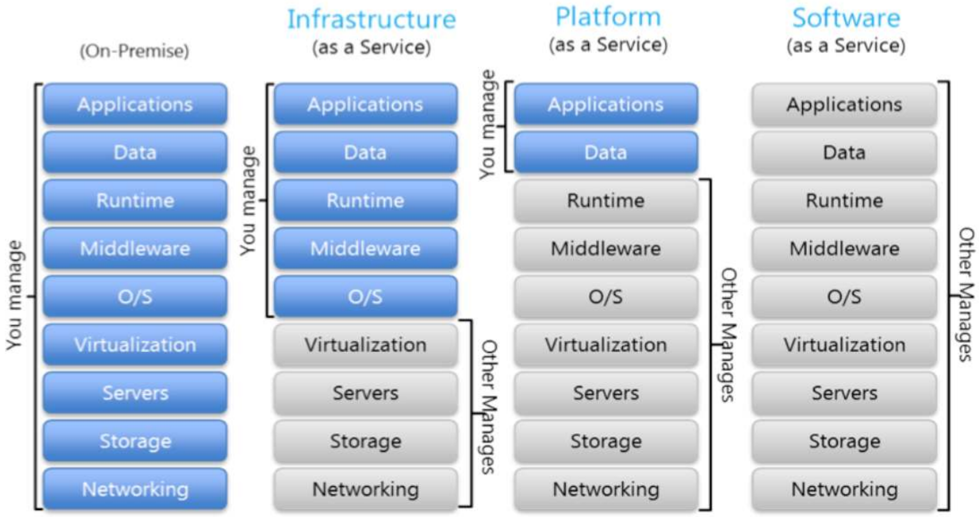

Infrastructure-as-a-Service

Infrastructure as a service (IaaS) abstracts hardware (server, storage, and network infrastructure) into a pool of computing, storage, and connectivity capabilities that are delivered as services for a usage-based (metered) cost. The goal of IaaS is to provide a flexible, standard, and virtualized operating environment that can become a foundation for platform as a service (PaaS) and software as a service (SaaS).

IaaS is usually seen to provide a standardized virtual server. The consumer takes responsibility for configuration and operations of the guest operating system (OS), software, and database (DB). Compute capabilities (such as performance, bandwidth, and storage access) are also standardized.

Service levels cover the performance and availability of the virtualized infrastructure. The consumer takes on the operational risk that exists above the infrastructure, i.e. those applications and services the consumer places on top of the infrastructure.

The Cisco and EMC with Microsoft Private Cloud Fast Track solution aims primarily to deliver infrastructure as a service while enabling platform as a service and softwareasaservice.

Data Center Consolidation and Virtualization

Data center consolidation and virtualization allows enterprise customers to migrate physical and virtual machines to Microsoft Hyper-V technology virtualization and Hyper-V-based cloud environments. This assists in reducing capital and operational expenses and improving manageability of both virtual and physical environments through utilization of the Microsoft System Center family of products.

Goals:

- Reduce cost of facilities, hardware, and licensing of alternative solutions through deploying a highly standardized Hyper-V, network, and storage infrastructure
- Reduce server sprawl and implement a more holistic and robust management solution
- Transition from organically grown virtualized environments to a private cloud solution to implement new capabilities and grow the business

Virtual Desktop Infrastructure

Virtual desktop infrastructure (VDI) enables IT staff to deploy desktops in virtual machines on secure and centralized hardware. A centralized and optimized virtual desktop enables users to access and run their desktops and applications wherever they may be, while IT is able to build a more agile and efficient IT infrastructure. Flexible Windows desktop scenarios give organizations the ability to choose the client computing solutions that best meet the unique needs of their businesses.

Fabric Logical Architecture

The logical architecture has two parts. The first is the "fabric," which is the physical infrastructure (servers, storage, and network) that will host and run all customer/consumer virtual machines. The second is "fabric management," which is a set of virtual machines comprising the Microsoft SQL Server and Microsoft System Center management infrastructure. The recommended practice is to have two or more Hyper-V host servers in a dedicated host cluster for the fabric management VMs and then have separate clusters for the fabric. For smaller scale deployments, the fabric management VMs can be hosted on the fabric itself.

Fabric

Figure 5 depicts the minimum requirements for high-level fabric. The requirements are categorized in compute, storage, and network layers. The minimums and recommendations are designed to balance cost vs. density and performance.

Figure 5. Private Cloud Fabric Infrastructure

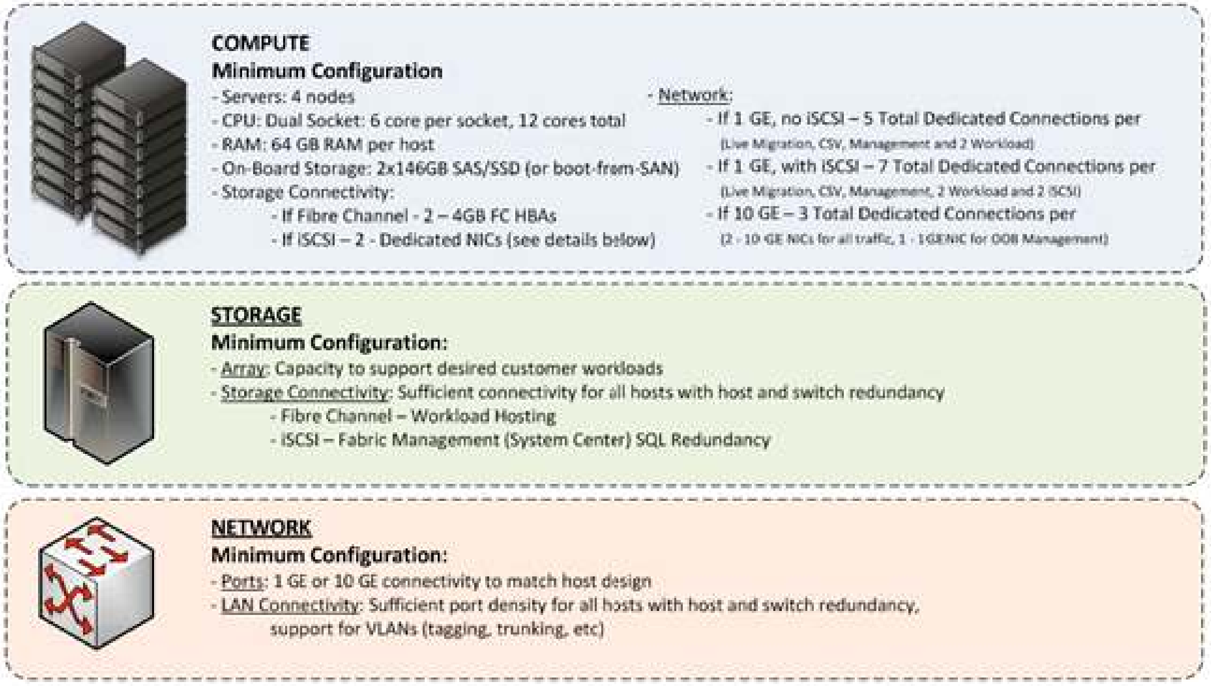

COMPUTE
Minimum Configuration
- Servers: 4 nodes
- CPU: Dual Socket: 6 core per socket, 12 cores total
- RAM: 64 GB RAM per host
- On-Board Storage: 2x146GB SAS/SSD (or boot-from-SAN)
- Storage Connectivity:
  - If Fibre Channel - 2 – 4GB FC HBAs
  - If iSCSI – 2 - Dedicated NICs (see details below)
- Network:
  - If 1 GE, no iSCSI – 5 Total Dedicated Connections per
    (Live Migration, CSV, Management and 2 Workload)
  - If 1 GE, with iSCSI – 7 Total Dedicated Connections per
    (Live Migration, CSV, Management, 2 Workload and 2 iSCSI)
  - If 10 GE – 3 Total Dedicated Connections per
    (2 - 10 GE NICs for all traffic, 1 - 1GENIC for OOB Management)

STORAGE
Minimum Configuration:
- Array: Capacity to support desired customer workloads
- Storage Connectivity: Sufficient connectivity for all hosts with host and switch redundancy
  - Fibre Channel – Workload Hosting
  - iSCSI – Fabric Management (System Center) SQL Redundancy

NETWORK
Minimum Configuration:
- Ports: 1 GE or 10 GE connectivity to match host design
- LAN Connectivity: Sufficient port density for all hosts with host and switch redundancy, support for VLANs (tagging, trunking, etc)

Server Architecture

The host server architecture is a critical component of the virtualized infrastructure, as well as a key variable in the consolidation ratio and cost analysis. The ability of the host server to handle the workload of a large number of consolidation candidates increases the consolidation ratio and helps provide the desired cost benefit. See note below.

The system architecture of the host server refers to the general category of the server hardware itself. Examples include rack-mounted servers, blade servers, and large symmetric multiprocessor servers (SMP). The primary tenet to consider when selecting system architectures is that each Hyper-V host will contain multiple guests with multiple workloads. Processor, RAM, storage, and network capacity are critical, as well as high I/O capacity and low latency.

Note: Capacity planning and management are required to verify that the services running within the cloud do not exceed the platform's capacity.

Server and Blade Design

The Cisco and EMC solution is based on Cisco's Unified Computing System incorporating blade servers. The following sections describe the components of UCS.

Cisco Unified Computing System Overview

The Cisco Unified Computing System is a next-generation data center platform that unites computing, networking, storage access, and virtualization resources into a cohesive system designed to reduce TCO and increase business agility. The system integrates a low-latency, lossless 10 Gigabit Ethernet (10 GE) unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multichassis platform in which all resources participate in a unified management domain.

Figure 6 shows the various hardware components, switches, fabric interconnects, chassis, blades, and interface cards, that comprise UCS. Figure 7 shows a logical view of the architecture of UCS. Interface cards, Ethernet or converged network, reside in the blade. Blades reside in a chassis. The chassis contains fabric extenders to connect to the Fabric Interconnect. The UCS Manager software runs with the operating system of the Fabric Interconnect, which can then manage a domain of all components within multiple chassis.

Figure 6. Cisco Unified Computing Hardware Components

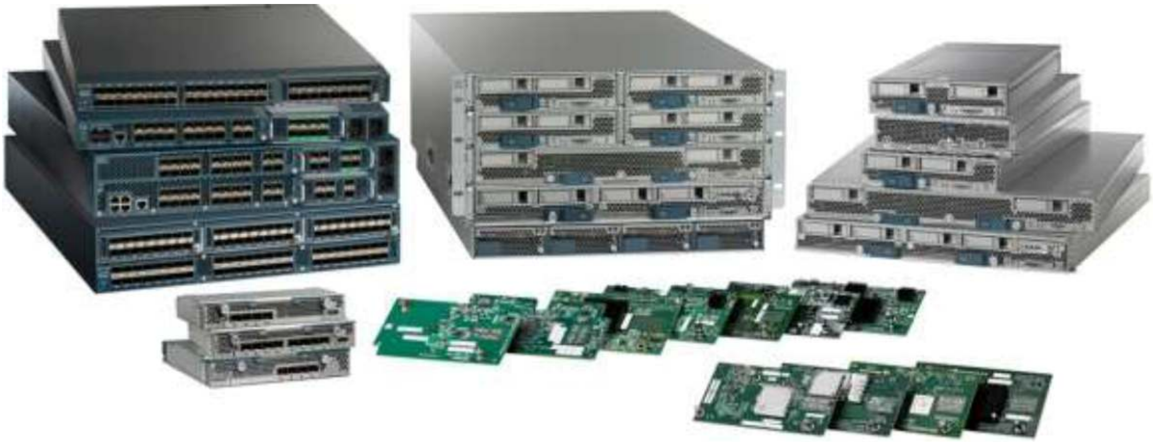

Figure 7. Unified Computing System Logical Architecture

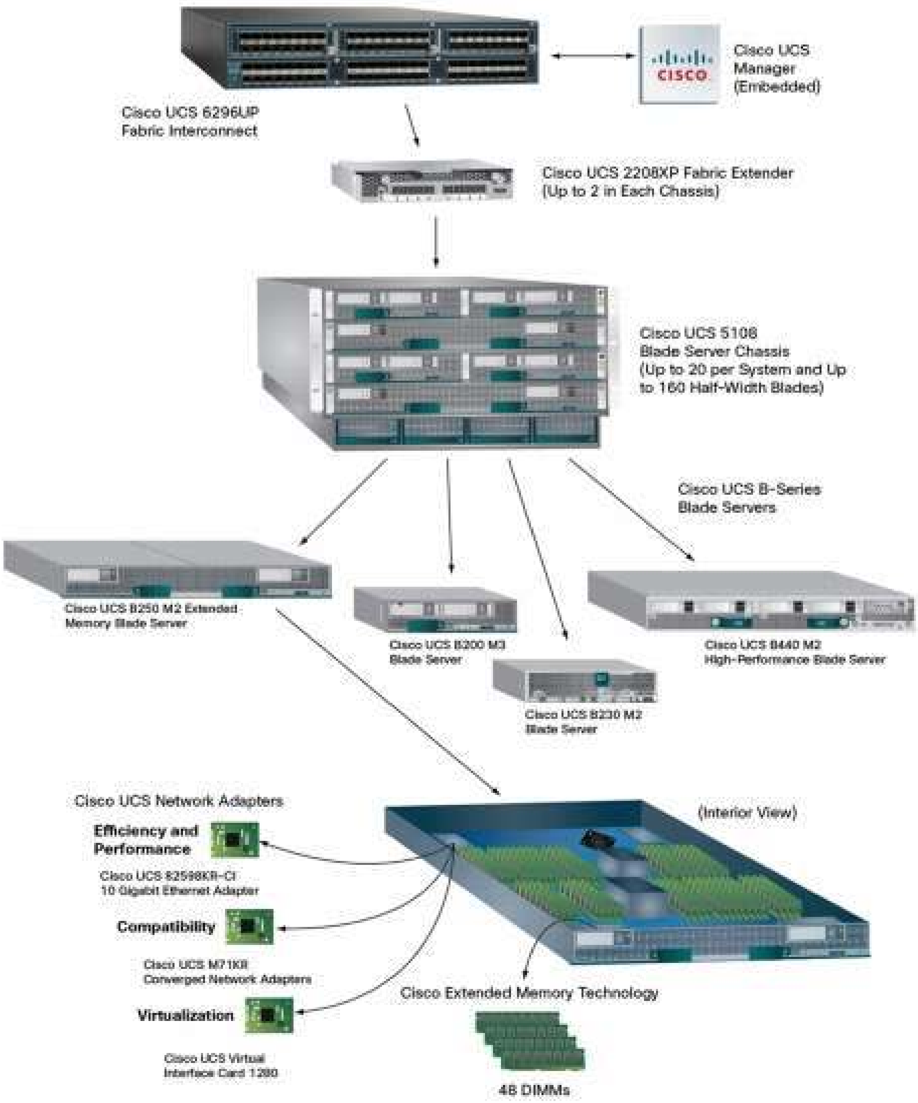
Cisco UCS Manager (Embedded)
Cisco UCS 6296UP Fabric Interconnect
Cisco UCS 2208XP Fabric Extender (Up to 2 in Each Chassis)
Cisco UCS 5108 Blade Server Chassis (Up to 20 per System and Up to 160 Half-Width Blades)
Cisco UCS B-Series Blade Servers
Cisco UCS B250 M2 Extended Memory Blade Server
Cisco UCS B200 M3 Blade Server
Cisco UCS B440 M2 High-Performance Blade Server
Cisco UCS B230 M2 Blade Server
Cisco UCS Network Adapters
Efficiency and Performance
Cisco UCS 82598KR-CI 10 Gigabit Ethernet Adapter
Compatibility
Cisco UCS M71KR Converged Network Adapters
Virtualization
Cisco UCS Virtual Interface Card 1280
(Interior View)
Cisco Extended Memory Technology
48 DIMMs

Cisco UCS 6200 Series Fabric Interconnect

Cisco UCS 6200 Series Fabric Interconnects (see Figure 7) are a core part of the Cisco Unified Computing System, providing both network connectivity and management capabilities for the system. The Cisco UCS 6200 Series offers wired-rate, low-latency, lossless 10 Gigabit Ethernet (10 GE); Fibre Channel over Ethernet (FCoE); and FC functions.

The Cisco UCS 6200 Series provides the management and communications backbone for the Cisco UCS B-Series Blade Servers and the 5100 Series Blade Server Chassis. All chassis, and therefore all blades, attached to the Cisco UCS 6200 Series Fabric Interconnects become part of a single, highly available management domain. In addition, by supporting unified fabric, the Cisco UCS 6200 Series provides both LAN and storage area network (SAN) connectivity for all blades within its domain.

From a networking perspective, the Cisco UCS 6200 Series uses a cut-through architecture that supports deterministic, low-latency, wired-rate 10 GE on all ports; switching capacity of 2 terabits (Tb); and 320-Gbps bandwidth per chassis, independent of packet size and enabled services.

Cisco low-latency, lossless 10 Gigabit Ethernet unified network fabric capabilities increase the reliability, efficiency, and scalability of Ethernet networks. The fabric interconnect supports multiple traffic classes over a lossless Ethernet fabric from the blade through the fabric interconnect. Significant TCO savings come from an FCoE-optimized server design in which NICs, host bus adapters (HBAs), cables, and switches can be consolidated.

Unified Fabric with FCoE: I/O Consolidation

The Cisco UCS 6200 Series is built to consolidate LAN and SAN traffic onto a single unified fabric, which saves the capital expenditures and operating expenses associated with multiple parallel networks, different types of adapter cards, switching infrastructure, and cabling within racks. Unified ports support allows either base or expansion module ports in the fabric interconnect to support direct connections from Cisco UCS to existing native FC SANs. The capability to connect FCoE to native FC protects existing storage system investments while dramatically simplifying in-rack cabling.

Cisco UCS Manager

The Cisco UCS 6200 Series hosts and runs Cisco UCS Manager in a highly available configuration that enables the fabric interconnects to fully manage all Cisco UCS elements. Connectivity to the Cisco UCS 5100 Series blade chassis is maintained through the Cisco UCS 2100 or 2200 Series fabric extenders in each blade chassis.

The Cisco UCS 6200 Series interconnects support out-of-band management through a dedicated 10/100/1000Mbps Ethernet-management port, as well as in-band management. Cisco UCS Manager typically is deployed in a clustered active-passive configuration on redundant fabric interconnects that are connected through dual 10/100/1000 Ethernet clustering ports.

Optimization for Virtualization

For virtualized environments, the Cisco UCS 6200 Series supports Cisco virtualization-aware networking and Cisco Data Center Virtual Machine Fabric Extender (VM-FEX) architecture. Cisco Data Center VMFEX allows the fabric interconnects to provide policy-based VM connectivity with network properties that move with the virtual machine and a consistent operational model for both physical and virtual environments.

Cisco UCS 6248UP 48-Port Fabric Interconnect

The Cisco UCS 6248UP 48-port fabric interconnect (Figure 8) is a one-rack unit (1RU) 10 Gigabit Ethernet, FCoE, and FC switch that offers up to 960-Gbps throughput and up to 48 ports. The switch has 32 1/10-Gbps fixed Ethernet, FCoE, and FC ports and one expansion slot.

Figure 8. Cisco UCS 6248UP 48-Port Fabric Interconnect

Expansion Module Option for Cisco UCS 6200 Series Fabric Interconnect

The Cisco UCS 6200 Series supports an expansion module that can be used to increase the number of 10 Gigabit Ethernet, FCoE, and FC ports (Figure 9). This unified port module provides up to 16 ports that can be configured for 10 Gigabit Ethernet, FCoE, and 1,2,4, or 8-Gbps native FC using the Small Form-Factor Pluggable (SFP) or SFP+3 interface for transparent connectivity with existing FC networks.

Figure 9. Expansion Module for the Cisco UCS 6200 Series Fabric Interconnect

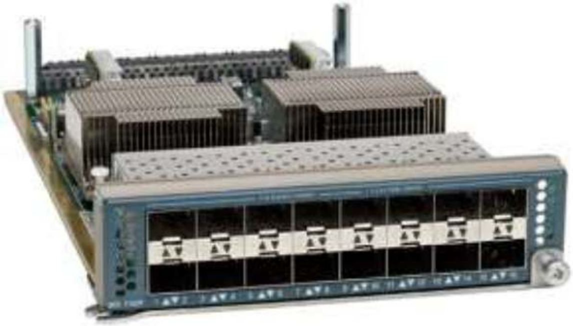

Server Design Principles

The server design must provide a high level of availability throughout. This includes, for example, features such as redundant power distribution units (PDUs), storage path, networking, and disks. To provide this level of availability, it is necessary to use two or more storage controllers to support multipathing on the I/O side.

Use multiple network adapters, multiport network adapters, or both on each host server. For converged designs, network technologies that provide teaming or virtual NICs can be used if redundancy is provided through NIC teaming or a hardware-level failover solution and if multiple virtual NICs or VLANs can be presented to the hosts for traffic segmentation and bandwidth control. The following network connections are required:

- One network dedicated to management purposes on the host machine
- One network dedicated to the cluster shared volumes (CSVs) and cluster communications network
- One network dedicated to the live-migration network
- One or more networks dedicated to the guest virtual machines (use 10-Gbps network adapters for highest consolidation)
- One network dedicated to iSCSI with multipath I/O (MPIO)

For the recommended configuration by quantity and type of NIC, refer to the Hyper-V: Live Migration Network Configuration Guide in the Microsoft TechNet Library.

The Cisco Unified Computing System uses a 10 Gigabit Ethernet unified fabric as the underlying I/O transport mechanism. Paired with the unified fabric in this architecture, Cisco Virtual Interface Card (VIC) technology has been specified for this solution (Figure 10). The Cisco VIC is a standards-compliant converged network adapter (CNA) that allows both traditional Ethernet as well as FC traffic to share a common physical transport. An added advantage of using this card is the ability to dynamically carve out and assign large numbers of NICs and HBAs from the same physical card. By using two fully redundant 10 Gigabit Ethernet upstream ports, up to 128 I/O devices can be presented to the PCIe bus and to the operating system running on top of it. These I/O devices can be mixed between NICs and HBAs at the administrator’s discretion.

Figure 10. Figure 10 Cisco UCS M81KR Virtual Interface Card

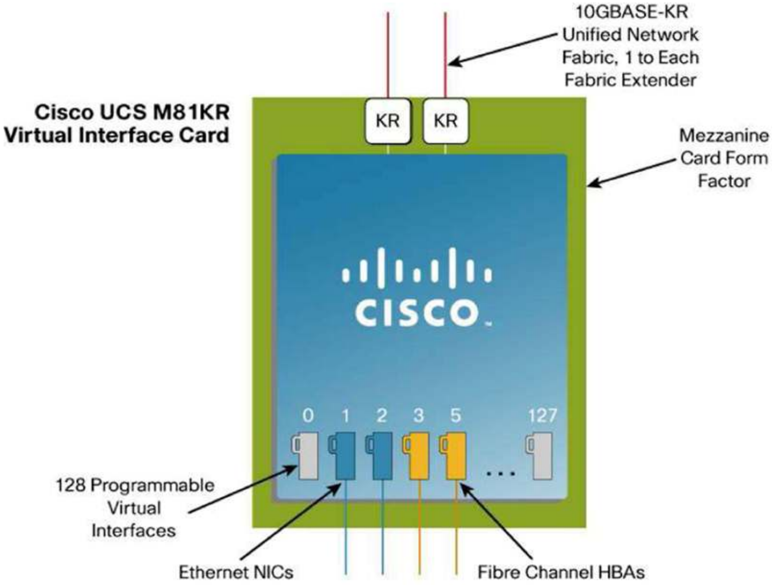

Fabric failover, which allows NIC redundancy to be managed below the operating system, is an additional innovation in the Cisco UCS solution. In the event of an upstream fabric failure, the NICs assigned to the failed fabric automatically failover to use the remaining fabric. This occurs without the need for special teaming drivers, which frequently introduce instability to a system.

In addition to redundancy, there is still a requirement to provide a platform that has optimal virtualization support through hardware virtualization capabilities, such as Intel® virtualization (Intel-VT) as the baseline, along with the Second Level Address Translation (SLAT) capabilities of Windows Server 2008 R2 Hyper-V to maximize performance. All these features are available with the Cisco UCS compute platform.

The implementation of a stateless computing model is a unique feature of the Cisco UCS architecture. By building a single point of management for all system components across as many as 160 server blades, Cisco has created a single-pane-of-glass view for the entire compute infrastructure. This unified view of the Cisco UCS is shown in Figure 11.

Figure 11. Cisco UCS: A Single Point of Management for All System Components

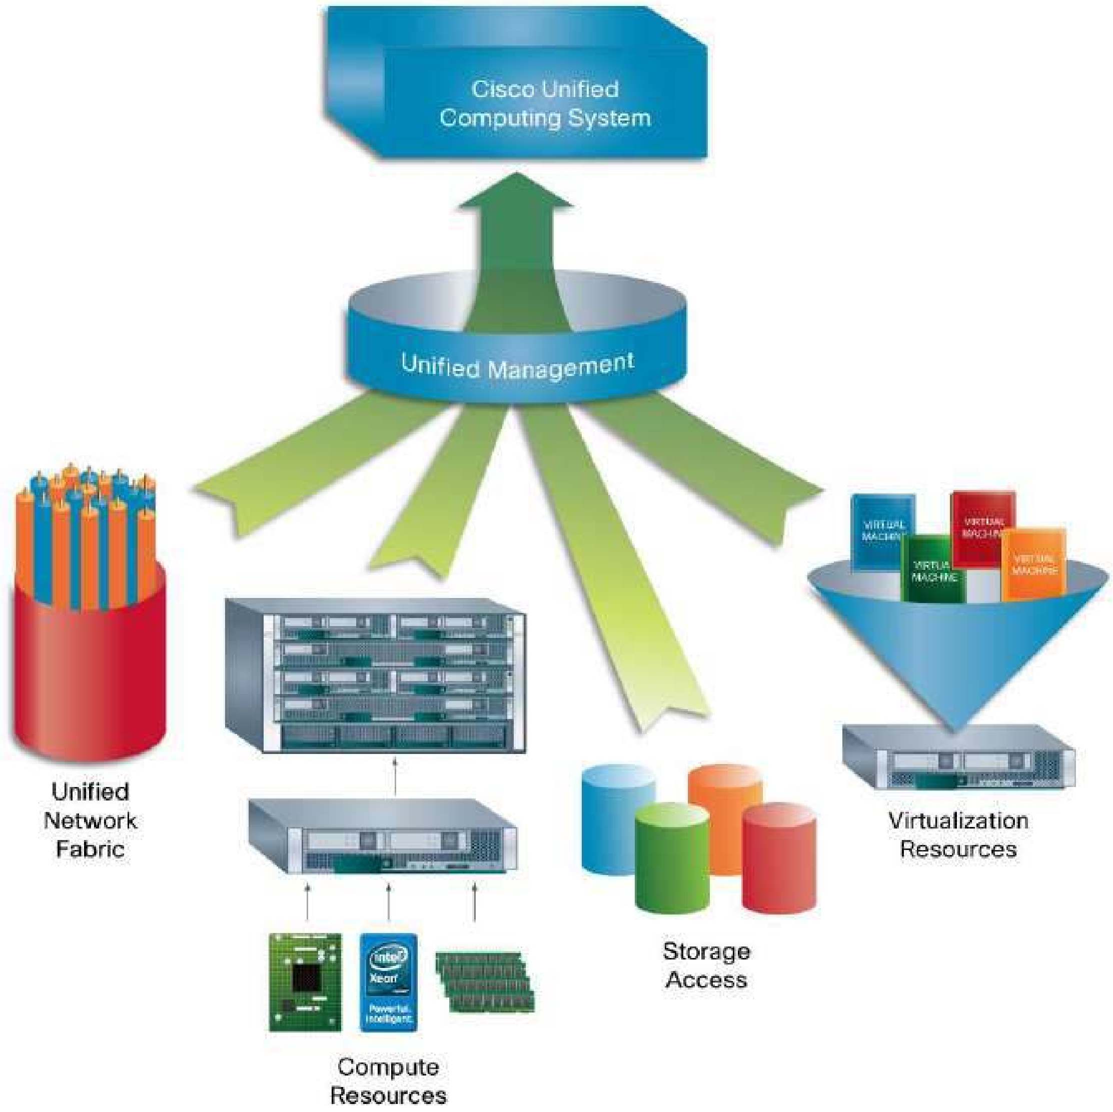

This unified management is then enhanced by abstracting away the configuration of individual services into atomic containers called service profiles. Service profiles contain all the components that identify a physical server to an operating system or hypervisor instance. Figure 12 shows the types of settings contained within a service profile.

Figure 12. Settings in a Service Profile

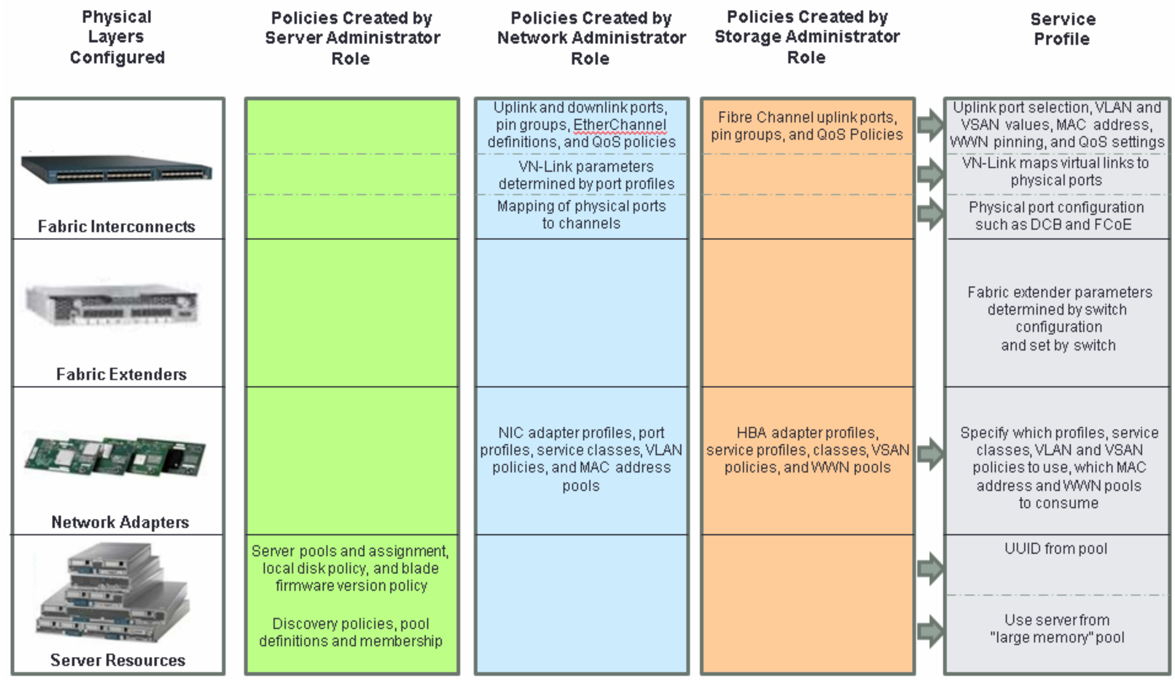

Physical Layers Configured	Policies Created by Server Administrator Role	Policies Created by Network Administrator Role	Policies Created by Storage Administrator Role	Service Profile
Fabric Interconnects		Uplink and downlink ports, pin groups, EtherChannel definitions, and QoS policies VN-Link parameters determined by port profiles Mapping of physical ports to channels	Fibre Channel uplink ports, pin groups, and QoS Policies	Uplink port selection, VLAN and VSAN values, MAC address, WWN pinning, and QoS settings VN-Link maps virtual links to physical ports Physical port configuration such as DCB and FCoE
Fabric Extenders				Fabric extender parameters determined by switch configuration and set by switch
Network Adapters		NIC adapter profiles, port profiles, service classes, VLAN policies, and MAC address pools	HBA adapter profiles, service profiles, classes, VSAN policies, and WWN pools	Specify which profiles, service classes, VLAN and VSAN policies to use, which MAC address and WWN pools to consume
Server Resources	Server pools and assignment, local disk policy, and blade firmware version policy Discovery policies, pool definitions and membership			UUID from pool Use server from "large memory" pool

Because the service profile is an atomic container, the entire personality of a server can be dynamically moved around the compute farm as needed for maximum flexibility.

This logical state abstraction can be managed entirely by using the Cisco UCS Manager API. The API supports every function within the Cisco UCS Manager management schema, and it is fully documented on Cisco Developer Network-UCS Manager. Additionally, the API is used by a broad variety of industry partners, which enables rich integration. This architecture features integration through the Microsoft System Center Operations Manager (SCOM) management pack for Cisco UCS, as well as through the Cisco UCS Manager Windows PowerShell toolkit.

Server Storage Connectivity

To achieve maximum flexibility and agility during both host and virtual machine provisioning events, a diskless architecture has been specified. This architecture relies on boot-from-SAN technology for the host operating systems and the use of iSCSI for the guest virtual machine data volumes.

This reference architecture uses a unified fabric architecture that allows both FCoE and 10 Gigabit Ethernet iSCSI over a common network fabric. To provide the proper quality of service (QoS), advanced QoS features are core to the fabric implementation of this architecture.

Storage Architecture

The storage design for any virtualization-based solution is a critical element. Storage is typically responsible for a large percentage of the solution’s overall cost, performance, and agility.

The EMC® VNX™ series utilizes the newest technology available to optimize customer storage environments. The use of Intel’s new Xeon 5600 processor provides faster clock speeds and more cores to handle the daily demands of storage I/O. Increased memory capacity, the optional use of Flash drives in FAST VP (Fully Automated Storage

Tiering for Virtual Pools) or FAST Cache, a new 4-lane 6 Gbps / lane SAS back end, and denser disk enclosures all help provide optimal scalability and efficiency.

The convergence of block and file storage (formerly SAN and network-attached storage [NAS]) in the VNX series achieves the simplicity that is necessary to effectively perform management operations.

The VNX series is designed for a wide range of environments that include mid-tier through enterprise. VNX provides offerings that include file only, block only, and unified (block and file) implementations. The VNX5300 Unified platform is implemented within the Fast Track solution.

EMC Unified VNX Storage

The VNX series is EMC's next generation of midrange products. It supporting various block and file protocols within a single system. The VNX operating environment offers significant advancements in efficiency, simplicity, and performance.

Benefits of the VNX series include:

- Automated tiering with Fully Automated Storage Tiering for Virtual Pools (FAST VP)
- FAST Cache for real-time performance acceleration
- Unified management with Unisphere™, which delivers a cohesive, unified user experience
- Industry-leading efficiency features including thin provisioning, compression, and de-duplication
- Snapshots and clones for local protection and alternate uses of data copies
- Application-consistent recovery through the integration of Replication Manager
- Disaster recovery using EMC's industry-leading RecoverPoint technology
- Scalability and Performance through its ability to scale to 1,000 drive in a single system, latest Intel processing technology, and high-bandwidth SAS backend
- Investment protection by providing UltraFlex Modules for file and block based host connectivity

VNX Hardware Components

VNX storage systems include the following components:

- Storage processors (SPs) support block data with UltraFlex I/O technology that supports Fibre Channel, iSCSI, and FCoE protocols. File connectivity is provided through the Data Movers of which up to 8 can be clustered. The Data Movers are connected to the Storage Processors via 8Gbps FC interconnect.

The VNX5300 uses a Disk Processing Enclosure (DPE) for core block processing. The 3U DPE provides two storage processor units and 15x3.5" or 25x2.5" drives. UltraFlex I/O units provide the ability to add required connectivity for host systems, including native Fibre Channel, iSCSI or FCOE.

Write cache is mirrored across the CMI bus for added protection and high available. Cache is destaged to persistent disk upon power failure or other fault scenarios, and is saved to a predefined Vault area on disks within the DPE. In total. the VNX5300 supports a maximum of 125 drives configured through the use of a combination of the DPE and multiple disk-array enclosers (DAEs).

- Up to 8 x 2U X-Blades (also referred to as Data Movers) are housed within a Data Mover Enclosure (DME) accessing data from the back end and providing host access using the same UltraFlex I/O technology that supports the NFS, CIFS, MPFS, and pNFS protocols. The X-Blades in each array are scalable and provide redundancy to ensure that no single point of failure exists.

- Standby power supplies are 1U in size and provide enough power to each storage processor to ensure that any data in flight is de-staged to the vault area in the event of a power failure. This ensures that no writes are lost. Upon restart of the array, the pending writes are reconciled and persisted to disk.
- Up to two Control stations, each 1U in size, provide management functions primarily on the file side including the failover of X-Blades.
- Disk-array enclosures (DAE) house the drives used in the array. They allow the expansion of each array's drive count to provide the storage required for expanding needs over the time of each implementation. There are two flavors: the 3U DAE holds up to 15 x 3.5" drives, the 2U DAE holds up to 25 x 2.5" drives. Supported drive types can differ between the two.
- Both the DPE and DME utilize dual-power supplies that are integrated into the chassis, and provide redundant power in the event of a single power supply failure.
- 4-port 8-Gb optical Fibre Channel UltraFlex I/O modules are used for front-end connectivity to hosts on storage processors. They are also used on X-Blades for connectivity to storage processors.
- 4-port 1-Gb copper iSCSI modules are available on X-Blades for NAS services.
- 2-port 1-Gb copper plus 2-port 1-Gb optical are available on X-Blades for NAS services.
- 4-port 1 Gigabit Ethernet iSCSI or TCP Offload Engine (TOE) are available on storage processors for front-end connectivity to hosts.
- 2-port 10 Gigabit Ethernet or FCoE are available on storage processors for front-end connectivity to hosts.
- 2-port 10 Gigabit Ethernet iSCSI is available on storage processors for front-end connectivity to hosts. It is used on X-Blades for NAS services.

Storage Options

Not all workloads have the same availability requirements nor achieve their requirements in the same way. In the case of data center architecture, it is possible to classify workloads into either stateful or stateless. A stateful workload has data that is specific to that VM; if the data were lost, it would become unavailable. A stateless workload uses data stored elsewhere in the data center and can achieve high availability through resiliency in the application. An example of a stateless workload is a front-end web server farm.

Most data centers run more stateful workloads, and therefore SAN storage will be used throughout. This provides the necessary persistence of data to the configured storage, and provides protection in the event of individual server failures, etc.

Once the workload type is determined, the performance and availability characteristics of the specific workload should be analyzed to determine the storage characteristics required (performance, redundancy, and so on).

The design includes an EMC VNX5300 storage array, which provides a fully compliant shared storage infrastructure. Support for both native FC and iSCSI Ethernet connectivity is provided through the implementation of dedicated, redundant I/O modules within the array.

The VNX5300 configured for Fast Track incorporates a DPE with 25 x 2.5-inch drive form factor. It includes four onboard 8 Gbps Fibre Channel ports and two 6 Gbps SAS ports for back-end connectivity on each storage processor. There is a micro-DB9 port and service LAN port available, which are used as a backup method of connectivity for system maintenance. Figure 13 shows the view of the VNX5300 storage processing enclosure from the rear. Note that the I/O modules shown do not accurately represent the Fast Track configuration.

Figure 13. VNX5300 Storage Processor Enclosure Rear View

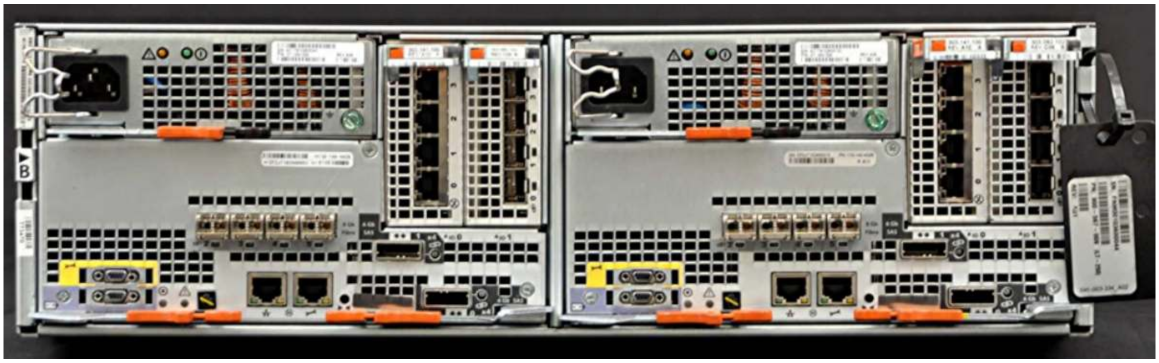

Each SP in the enclosure has a power supply module and two UltraFlex I/O module slots, as shown in 14. Both I/O module slots can be populated on this model. Any slots without I/O modules will be populated with blanks (to ensure proper airflow). The fault/status light of the DPE form factor is located on the rear, which is different from previous models and VNX models using a SPE form factor.

Figure 14. Expanded view of a single Storage Processing (SP) Component

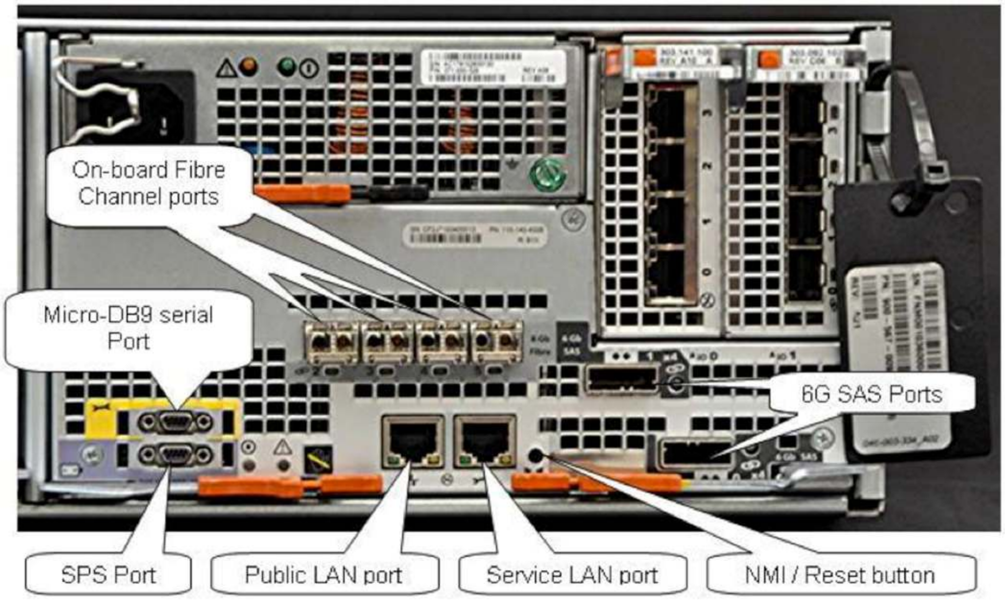

The front of the DPE houses the first tray of drives, as shown in 15. At the top of the enclosure, facing front, from left to right, is the enclosure fault LED (amber), air channel leading to the temperature sensor, and the enclosure power LED (blue).

Figure 15. VNX5300 Storage Processor Enclosure Front View

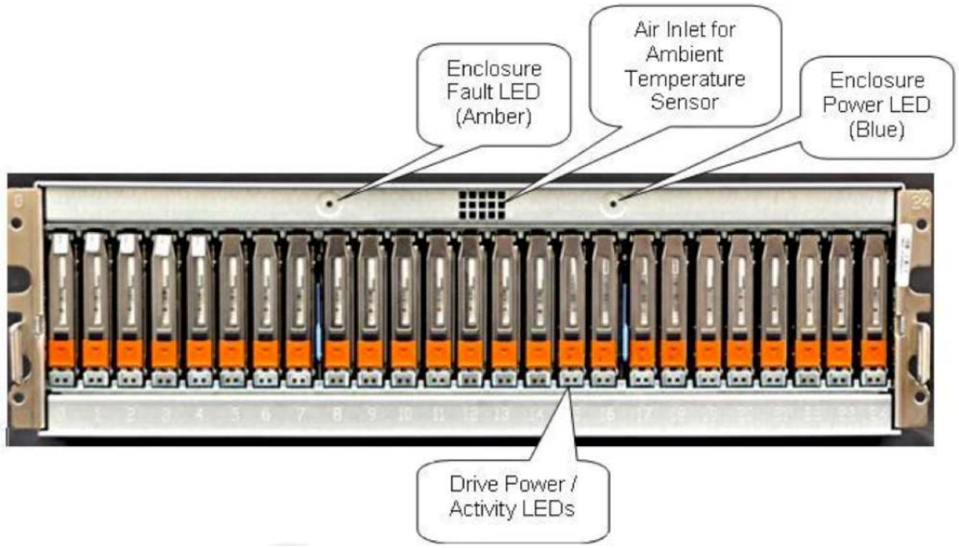
Enclosure Fault LED (Amber)
Air Inlet for Ambient Temperature Sensor
Enclosure Power LED (Blue)
Drive Power / Activity LEDs

The VNX5300 configured for the Fast Track reference architecture is configured with 75 300GB SAS drives. The storage connectivity provided to the UCS environment is comprised of Fibre Channel connectivity from the on-board 8-Gb FC connections on each Storage Processor, as well as iSCSI connectivity provided by an additional 10-Gb iSCSI I/O module within each storage processor enclosure.

SAN Storage Protocols

Block- Level Versus File-Based Storage

In Windows Server 2008 R2 SP1, file-based storage is not supported for Hyper-V host clusters. Hyper-V host clusters require block-based shared storage, accessible to each host in the cluster.

Storage connectivity is provided through block-based protocols: FCoE and iSCSI at the management and compute server levels, and as FC and iSCSI at the array level. Connectivity is provided through the dual Cisco Nexus 5548UP Switches, which also provide the necessary translations and infrastructure connectivity.

Comparison of iSCSI, FC, and FCoE

Fibre Channel has historically been the storage protocol of choice for enterprise data centers for a variety of reasons, including performance and low latency. These considerations have offset Fibre Channel's typically higher costs. In the last several years, Ethernet's continually advancing performance from 1 GBps to 10 GBps and beyond have led to great interest in storage protocols that is the Ethernet transport, such as iSCSI, and recently, Fibre Channel over Ethernet (FCoE).

A key advantage of protocols that use the Ethernet transport is the ability to take advantage of a "converged" network architecture where a single Ethernet infrastructure serves as the transport for both local-area network (LAN) and storage traffic. This can reduce costs in several ways, such as the elimination of dedicated Fibre Channel switches and a reduction in cabling, which can also be a significant cost in large data center environments.

As a technology, FCoE brings the benefits of an Ethernet transport while retaining the advantages of the Fibre Channel protocol and the ability to use Fibre Channel storage arrays.

Several enhancements to standard Ethernet are required for FCoE. This is commonly referred to as enhanced Ethernet or data center Ethernet. These enhancements require Ethernet switches capable of supporting enhanced Ethernet, and the configured Cisco Nexus swicthes support all requirements.

For Hyper-V, iSCSI-capable storage provides an advantage in that it is the protocol that can also be utilized by Hyper-V guest virtual machines for guest clustering. This provides the necessary storage connectivity to configure Windows Server Failover Clustering within the guest Virtual Machines, and provides the necessary levels of High Availability for implemetations such as SQL Server,

A common practice in large-scale virtualization deployments is to utilize both Fibre Channel and iSCSI. Fibre Channel provides the host storage connectivity, and iSCSI is used only by guests that require in-OS iSCSI connectivity, such as a guest cluster. In this case, both Ethernet and block storage I/O will be sharing the same pipe, segregation and traffic prioritization is achieved by VLANs and QoS within the Nexus and UCS environment.

Storage Network

Both iSCSI and FCoE utilize an Ethernet transport for storage networking. This provides another architecture choice in terms of whether to utilize a dedicated Ethernet network with separate switches, cables, paths and so on or whether to use a “converged” network where multiple traffic types are run over the same cabling and infrastructure.

Figure 16 illustrates the differences between a traditional architecture on the left with separate Ethernet and Fibre Channel switches, each with redundant paths, and a converged architecture where both Ethernet and Fibre Channel (via FCoE) utilize the same set of cables while still providing redundant paths. The converged architecture requires fewer switches and cables. In the converged architecture, the switches must be capable of supporting enhanced Ethernet.

Figure 16. Converged Architecture Versus LAN and Fibre Channel SAN

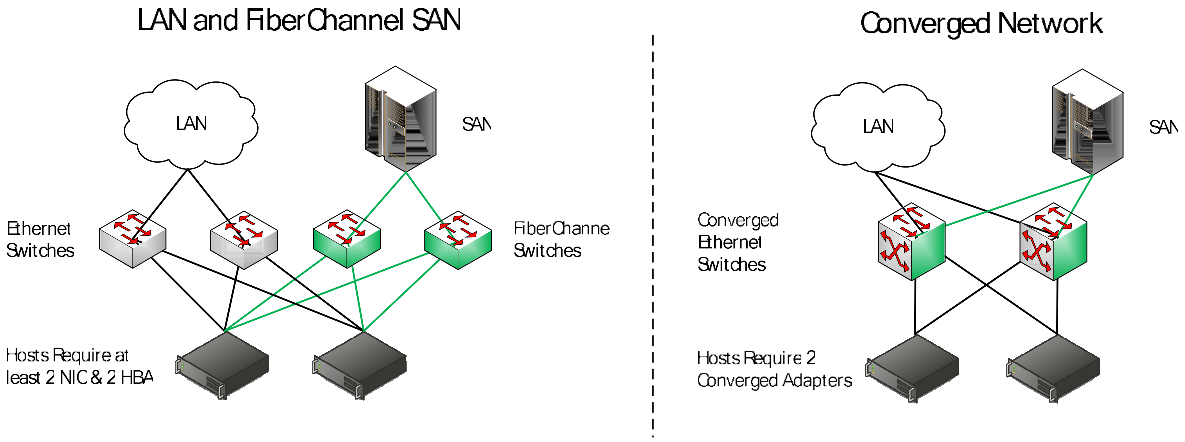
LAN and FiberChannel SAN
LAN
SAN
Ethernet Switches
FiberChanne Switches
Hosts Require at least 2 NIC & 2 HBA
Converged Network
LAN
SAN
Converged Ethernet Switches
Hosts Require 2 Converged Adapters

The Cisco and EMC design implements a highly available solution for storage connectivity. Virtual network and HBA interfaces are presented to each blade server to provide the necessary network and storage connectivity.

Storage devices are accessible as block devices either as FC target devices (FCoE) or as iSCSI target devices through the network infrastructure. In all cases, connectivity is provided in a redundant manner. EMC’s PowerPath® host software, implemented in conjunction with Windows MPIO provides redundant scalable connectivity to block storage devices independent of the transport protocol.

Cluster Shared Volumes

Windows Server 2008 R2 includes the first version of Windows Failover Clustering to offer a distributed file access solution. Cluster Shared Volumes (CSV) functionality in Windows Server 2008 R2 is exclusively for use with the Hyper-V role and enables all nodes in the cluster to access the same cluster storage volumes at the same time. CSV uses standard Windows NT file system (NTFS) and has no special hardware requirements beyond supported block-based shared storage.

CSV provides not only shared access to the disk, but also storage path I/O fault-tolerance (dynamic I/O redirection). In the event the storage path on one node becomes unavailable, the I/O for that node will be rerouted via Server Message Block (SMB) protocol through another node. There is a performance impact while running this state; it is designed for use as a temporary failover path while the primary dedicated storage path is brought back online. This feature can use any cluster communications network and further increases the need for high-speed networks.

CSV maintains metadata information about the volume access and requires that some I/O operations take place over the cluster communications network. One node in the cluster is designated as the coordinator node and is responsible for these disk operations. Virtual machines, however, have direct I/O access to the volumes and y use only the dedicated storage paths for disk I/O, unless a failure scenario occurs as described earlier.

CSV Limits

Table 1 lists the limitations that are imposed by the NTFS file system and that are inherited by CSV.

Table 1. CSV Limits

CSV Parameter	Limitation
Maximum Volume Size	256 TB
Maximum # Partitions	128
Directory Structure	Unrestricted
Maximum Files per CSV	4+ Billion
Maximum VMs per CSV	Unlimited

CSV Requirements

- All cluster nodes must use Windows Server 2008 R2 (with or without SP 1) Enterprise or Datacenter.
- All cluster nodes must use the same drive letter for the system disk.
- All cluster nodes must be on the same logical network subnet. Virtual LANs (VLANs) are required for multisite clusters running CSV.
- NT LAN Manager (NTLM) authentication in the local security policy must be enabled on cluster nodes.
- SMB must be enabled for each network on each node that will carry CSV cluster communications.
- “Client for Microsoft Networks” and “File and Printer Sharing for Microsoft Networks” must be enabled in the network adapter’s properties to enable all nodes in the cluster to communicate with the CSV.
- The Hyper-V role must be installed on any cluster node that may host a VM.

CSV Volume Sizing

Because all cluster nodes can access all CSV volumes simultaneously, we can now use standard LUN allocation methodologies based on performance and capacity requirements of the workloads running within the VMs themselves. Generally speaking, isolating the VM operating system I/O from the application data I/O is a good start, in addition to application-specific I/O considerations, such as segregating databases and transaction logs and creating SAN volumes and/or storage pools that factor in the I/O profile itself (that is, random read and write operations versus sequential write operations).

CSV's architecture differs from other traditional clustered file systems, which frees it from common scalability limitations. As a result, there is no special guidance for scaling the number of Hyper-V Nodes or VMs on a CSV volume other than ensuring that the overall I/O requirements of the expected VMs running on the CSV are met by the underlying storage system and storage network. While rare, disks and volumes can enter a state where a chkdisk is required, which with large disks may take a long time to complete, causing downtime of the volume during this process somewhat proportional to the volume's size and number of files residing on the disk.

Each enterprise application you plan to run within a VM may have unique storage recommendations and even perhaps virtualization-specific storage guidance. That guidance applies to use with CSV volumes as well. The important thing to keep in mind is that all VM's virtual disks running on a particular CSV will contend for storage I/O.

Also worth noting is that individual SAN LUNs do not necessarily equate to dedicated disk spindles. A SAN storage pool or RAID array may contain many LUNs. A LUN is simply a logic representation of a disk provisioned from a pool of disks. Therefore, if an enterprise application requires specific storage IOPS or disk response times, you must consider all the LUNs in use on that storage pool. An application that would require dedicated physical disks if it were not virtualized may require dedicated storage pools and CSV volumes running within a VM.

The Cisco and EMC design provides full support for Microsoft Cluster Shared Volumes, and includes the necessary isolated VLAN infrastructure for Live Migration requirements.

CSV Design Patterns

There are several CSV design patterns that are common in Hyper-V deployments. They are discussed in this section.

Single CSV per Cluster

In the single CSV per cluster design pattern (Figure 17), the SAN is configured to present a single large LUN to all the nodes in the host cluster. The LUN is configured as a CSV in Failover Clustering. All VM-related files (virtual hard disks [VHDs]), configuration files, and so on) belonging to the VMs hosted on the cluster are stored on the CSV. Optionally, data deduplication functionality provided by the SAN can be utilized (if supported by the SAN vendor).

Figure 17. Single CSV per Cluster Design Pattern

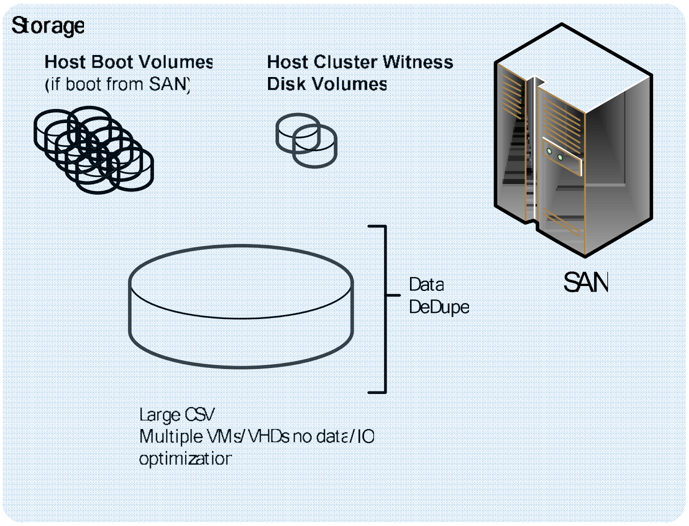

Multiple CSVs per Cluster

In the multiple CSVs per cluster design pattern (Figure 18), the SAN is configured to present two or more large LUNs to all the nodes in the host cluster. The LUNs are configured as a CSV in Failover Clustering. All VM-related files (VHDs, configuration files, and so on) belonging to the VMs hosted on the cluster are stored on the CSVs. Optionally, data deduplication functionality provided by the SAN can be utilized (if supported by the SAN vendor).

The Cisco and EMC Fast Track solution is configured with multiple CSVs per cluster.

Figure 18. Multiple CSVs per Cluster Design Pattern

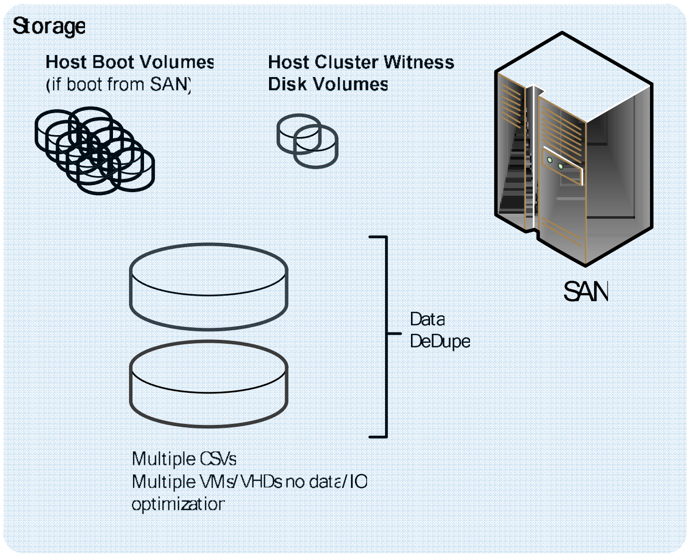

For both the single and multiple CSV patterns, each CSV has the same I/O characteristics, so each individual VM has all its associated VHDs stored on one of the CSVs (Figure 19).

Figure 19. VM Virtual Hard Drives on Common CSV

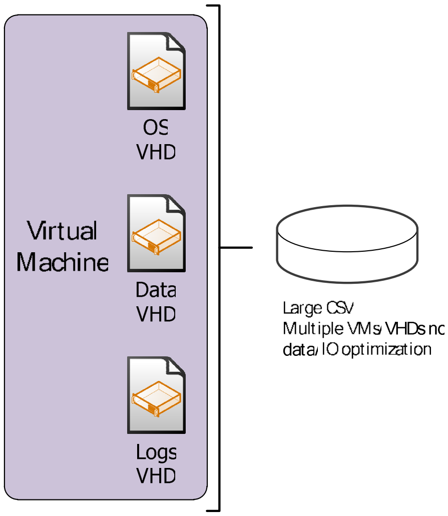

Multiple I/O Optimized CSVs per Cluster

In the multiple I/O optimized CSVs per cluster design pattern (Figure 20), the SAN is configured to present multiple LUNs to all the nodes in the host cluster, but the LUNs are optimized for particular I/O patterns such as fast sequential read performance, or fast random write performance. The LUNs are configured as CSV in Failover Clustering. All VHDs belonging to the VMs hosted on the cluster are stored on the CSVs but targeted to the most appropriate CSV for the given I/O needs.

Figure 20. Multiple I/O Optimized CSVs per Cluster Design Pattern

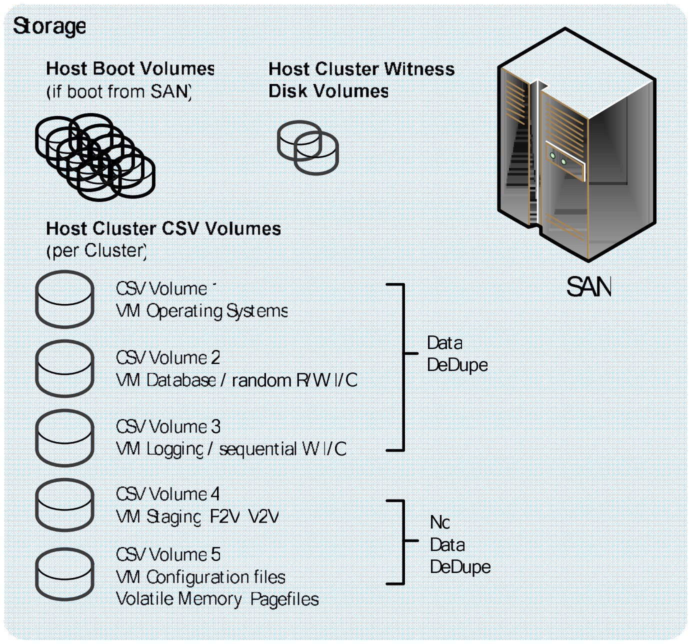

In the multiple I/O optimized CSVs per cluster design pattern, each individual VM has all its associated VHDs stored on the appropriate CSV per required I/O requirements (Figure 21).

Figure 21. VM Virtual Hard Drives on I/O Optimized CSV

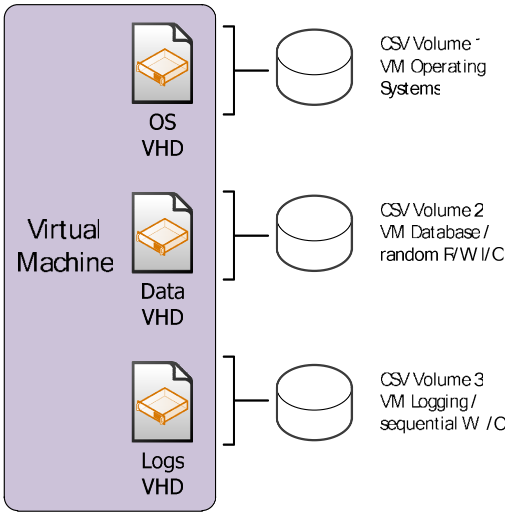

Note that a single VM can have multiple VHDs, and each VHD can be stored on a different CSV (provided that all CSVs are available to the host cluster that the VM is created on).

SAN Design

High Availability

All components within the Cisco and EMC design provide full redundancy (Figure 22).

The EMC VNX5300 is a dual-storage processor design, providing full power and connectivity redundancy within and across the storage processors. All I/O connectivity is provided in a redundant and scalable mode to the core infrastructure. All disk enclosures utilized in the configuration also provide full redundant power and SAS connectivity.

For FC connectivity, a port from each storage processor is connected to each Cisco Nexus 5548UP Switch. A failure of a single Cisco Nexus switch would result in loss of only half of the connectivity to the servers. MPIO with EMC PowerPath at the server, dynamically deconfigures the failed routes. Loss of a VNX5300 storage processor would result in failover of the LUNs owned by the faulted storage processor to the surviving storage processor. In tandem, MPIO with EMC PowerPath at the server will dynamically deconfigure the failed routes, providing transparent failover of the disk resources at the host level.

For converged network connectivity, each VNX5300 storage processor is connected to each 6248UP controller by dual 10 Gigabit Ethernet interfaces. A failure of a single Cisco UCS 6248UP Controller would result in loss of only half the connectivity to the servers. MPIO with EMC PowerPath at the server dynamically deconfigures the failed iSCSI routes. Loss of a VNX5300 storage processor would result in a failover of all iSCSI LUNs to the surviving storage processor. MPIO with EMC PowerPath at the server dynamically deconfigures the failed routes.

When the failure of any component has been resolved, MPIO with EMC PowerPath will automatically reconfigure the paths to full operational state.

Figure 22. High-Availability, Full-Redundancy SAN Design

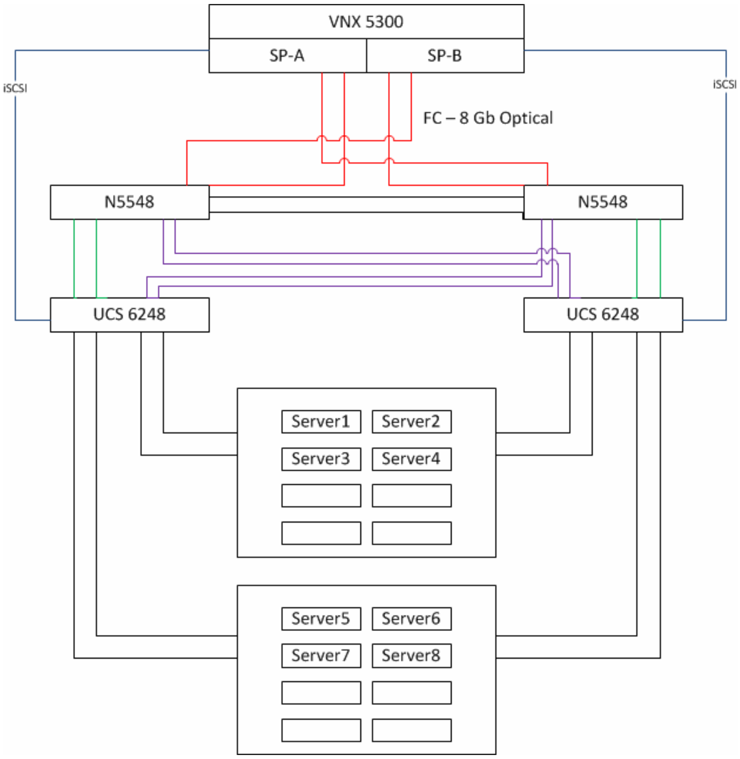

Performance

The array utilized in the design is based on a VNX5300 system with 75 x 300-GB SAS drives. The VNX system is able to scale 125 drives and can support a mix of SAS and Flash storage devices. In the solution, the array is able to scale to tens of thousands of I/Os.

Host connectivity is provided in a redundant, scalable manner. FC connectivity is provided through four FC connections to dual Cisco Nexus 5548UP Switches (FC connectivity at 8 Gb). The UCS environment is connected to the Cisco Nexus 5548UP Switches via FCoE connections for both storage and network connectivity. iSCSI connectivity is provided through direct connections from the VNX5300 GigE controllers to the Cisco UCS 6248UP Fabric Interconnects.

Sizing of the VNX array can be adjusted to support the anticipated workload from the combined compute and management nodes. This also allows for the implementation of Fast Cache technology, that enables significantly higher perfromance levels.

Storage performance is a complicated mix of drive, interface, controller, cache, protocol, SAN, HBA, driver, and operating system considerations. The overall performance of the storage architecture is typically measured in terms of maximum throughput, maximum I/O operations per second (IOPS), and latency or response time. While each of the three factors is important, IOPS and latency are the most relevant to server virtualization.

Most modern SANs utilize a combination of high-speed disks, slower-speed disks, and large memory caches. Storage controller cache can improve performance during burst transfers or when the same data is accessed frequently by storing it in the cache memory, which is typically several orders of magnitude faster than the physical disk I/O. Storage system cache is most effectively utilized to service write workloads from host systems. Write operations that represent updates to blcok contents are stored directly within SAN cache. For the VNX series arrays, write cache is a mirrored resource across storage processors. As a result, write operations can be immediately satisfied from the host perspective, and are destaged to disk as a subsequent process. Write data is protected by backup battery supplies within the VNX array, and will be destaged immediately to disk in the event of complete power failure.

Most modern SANs use a combination of high-speed disks, slower-speed disks, and large memory caches. EMC VNX has implemented solid state disk technology to fill a widening gap between ever faster memory and spinning disk that reached a performance plateau. Making effective use of drives that offer the best $/GB and drives that offer the best $/Performance requires careful planning and tools to place the right data onto the right tier at the right time. EMC VNX FAST Suite does just that, by monitoring workload characteristics within the array, and dynamically placing data on the requisite storage tier. While Fast Cache nor the VNX Fast Suite are included in the base Fast Track configuration, they can be added as options within the configuration to provide higher levels of performance.

Drive Types

The Cisco and EMC Fast Track solution defined the base VNX5300 array to be configured with 75 x 300-GB SAS drives (10K rpm). The base configuration can be ammended to implement alternate disk storage types, for example, replacing the 300 GB drives with 600 GB drives. Additionally, inclusion of optional features such as Fast Cache or EMC Fast Suite can provide significant solution advantages.

RAID Array Design

The VNX5300 array can be configured with multiple RAID protection schemes to provide a robust set of perfromance characteristics for a variety of user workloads. In the base configuration, multiple pools of RAID-5 and RAID-1/0 storage pools are configured. These pools are utilized to support the core infrastructire services as well as provide the basis for general user workload requirements. It is fully expected that storage pools for user workloads will be customized in each implementation.

The VNX5300 array provides additional levels of redundancy for events such as drive failures. The Cisco and EMC design includes hot spare disk drives. In the event of drive failures, two scenarios are possible. Often the VNX sniffer software detects a failing drives and will initiate a proactive copy to a spare, avoiding a longer rebuild process. In situations where a drives fails before this occurs, a regular rebuild from the parity information will occur.

Cisco and EMC utilized existing best practices for core components, including RAID 5 for general- purpose storage and RAID 1/0 for locations utilized for SQL Server transaction log space.

Multipathing

In all cases, multipathing should be used. Generally, storage vendors will build a device-specific module (DSM) on top of Microsoft's Windows Server 2008 R2 SP1 MPIO software.

Cisco and EMC implement EMC PowerPath multipathing solutions for all server deployments. EMC PowerPath is a supported DSM for Windows MPIO. PowerPath implements support for both iSCSI and FC/FCoE connectivity.

Fibre Channel

Fibre Channel connectivity is implemented within the design as the one of the block-based connections from the VNX5300 array to the two Cisco Nexus 5548UP Switches. Redundancy within the environment is provided by multiple redundant connections from each storage processor within the VNX5300 to each of the Cisco Nexus 5548UP Switches.

Each Cisco Nexus 5548UP Switch is implemented as a separate FC/FCoE fabric for zoning purposes. Zoning is maintained on each switch for all local connections (that is, all FCoE and FC ports within that switch).

Masking is fully implemented within the VNX5300, allowing complete segregation of LUNs where necessary, or allowing for shared access in Failover Clustering configurations. Masking operations are managed either through the usage of the UniSphere management interface for initial configuration, or via ESI/PowerShell functionality.

iSCSI

In the solution design, iSCSI connectivity from the VNX5300 to the Cisco UCS 6248UP Fabric Interconnects is established through dedicated connections. A dedicated VLAN environment is then configured within the 6248UP Fabric Interconnects for all iSCSI connectivity.

The Cisco UCS management infrastructure provides quality of service mechanisms to guarantee VLAN access in the event of competition of VLAN resources.

Encryption and Authentication

The EMC VNX5300 storage array ensures that only specified initiators are able to log in to front-end ports. Additionally, masking operations within the VNX arrays ensure that only designated initiators are able to access LUN resources.

Challenge Handshake Authentication Protocol (CHAP) authentication is implemented by the VNX5300 environment on iSCSI connectivity, and is utilized in conjunction with the Microsoft host iSCSI initiator.

Jumbo Frames

Jumbo frames are implemented within the proposed solution. The VNX5300, Cisco Nexus 5548UP Switches, and Cisco UCS infrastructure support jumbo frames. Jumbo frames are utilized within the VMs for any iSCSI connectivity implemented.

Data Deduplication

The EMC VNX5300 array implements snapshot technology, which is used to provide local replica management. In such implementations only changed blocks are stored in a separate, designated allocation pool. Unchanged blocks are referenced by indirection pointers back to the original source storage allocations. Implementations that utilize “gold master” images or templates can derive significant storage utilization benefits from such implementations.

Thin Provisioning

EMC VNX storage arrays fully support “thin” devices through the implementation of EMC virtual provisioning technology. Virtual provisioning can be implemented as appropriate in the configuration. By default, storage pools defined within the Fast Track environment will be configured as thin capable, and LUNs created from the pools may be defined as thin LUNs.

Volume Cloning

Volume cloning is another common practice in virtualization environments. This can be used for both host and VM volumes, dramatically increasing host installation times and VM provisioning times.

Cisco and EMC utilize array-based replication (cloning) for duplication of templates or gold master images. The implementation utilizes System Center Virtual Machine Manager Rapid Provisioning.

Volume Snapshots

SAN volume snapshots are a common method of providing a point-in-time, instantaneous backup of a SAN volume or LUN. These snapshots are typically block-level and only use storage capacity as blocks change on the originating volume. Some SANs provide tight integration with Hyper-V integrating both the Hyper-V VSS Writer on hosts and volume snapshots on the SAN. This integration provides a comprehensive and high-performing backup and recovery solution.

Cisco and EMC have included the EMC Avamar Backup/Recovery solution as an optional offering in the Fast Track solution. The EMC Avamar product fully integrates with Hyper-V volume shadow copy services (VSS) operations.

Storage Tiering

The EMC VNX5300 storage array supports the optional implementation of EMC Fully Automated Storage Tiering functionality. EMC supports the inclusion of FAST technology as appropriate within the Fast Track configuration, and customers may include a Flash Drive Tier to enhance overall system performance, if required by deployment needs.

Tiering storage is the practice of physically partitioning data into multiple distinct classes based on price, performance, or other attributes. Data may be dynamically moved among classes in a tiered storage implementation based on access activity or other considerations.

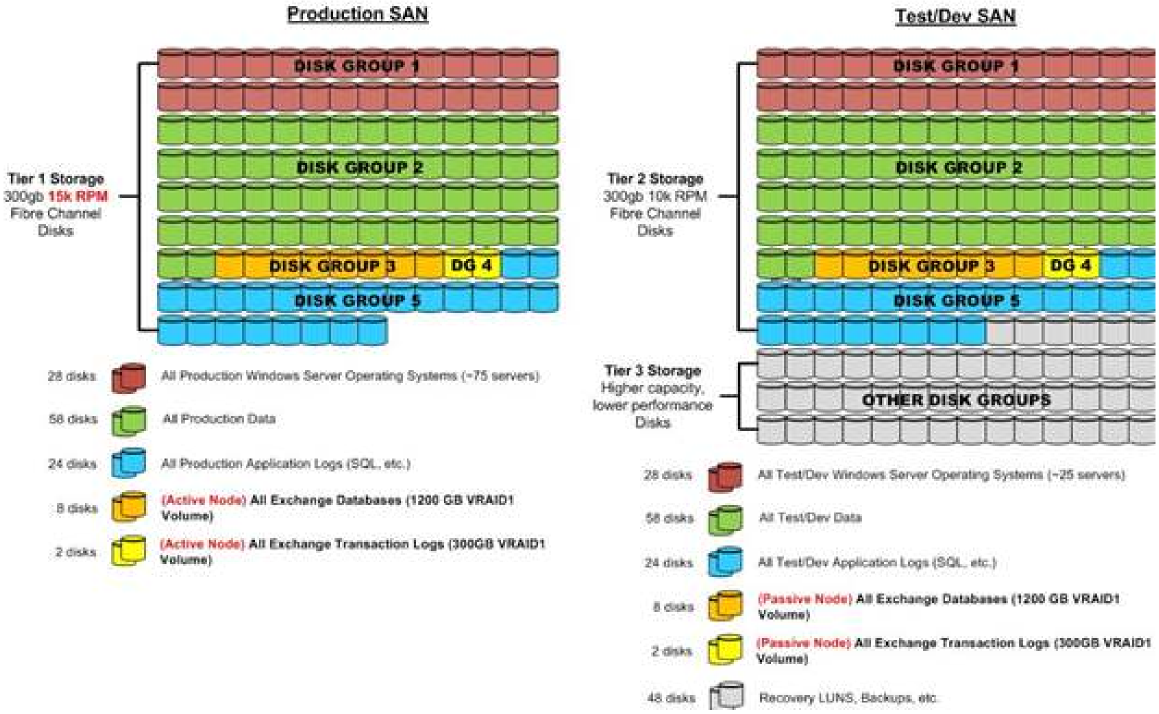
Production SAN
Test/Dev SAN
DISK GROUP 1
DISK GROUP 2
DISK GROUP 3
DG 4
DISK GROUP 5
Tier 1 Storage
300gb 15k RPM
Fibre Channel
Disks
DISK GROUP 1
DISK GROUP 2
DISK GROUP 3
DG 4
DISK GROUP 5
Tier 2 Storage
300gb 10k RPM
Fibre Channel
Disks
Tier 3 Storage
Higher capacity,
lower performance
Disks
OTHER DISK GROUPS
28 disks All Production Windows Server Operating Systems (~75 servers)
58 disks All Production Data
24 disks All Production Application Logs (SQL, etc.)
8 disks (Active Node) All Exchange Databases (1200 GB VRAID1 Volume)
2 disks (Active Node) All Exchange Transaction Logs (300GB VRAID1 Volume)
28 disks All Test/Dev Windows Server Operating Systems (~25 servers)
58 disks All Test/Dev Data
24 disks All Test/Dev Application Logs (SQL, etc.)
8 disks (Passive Node) All Exchange Databases (1200 GB VRAID1 Volume)
2 disks (Passive Node) All Exchange Transaction Logs (300GB VRAID1 Volume)
48 disks Recovery LUNS, Backups, etc.

Storage Automation

One of the objectives of the Cisco and EMC with Microsoft Private Cloud solution is to enable rapid provisioning and deprovisioning of virtual machines. Doing so at large scale requires tight integration with the storage architecture and robust automation. Provisioning a new virtual machine on an already existing LUN is a simple operation; however, provisioning a new CSV LUN, adding it to a host cluster, and so on, are relatively complicated tasks that must be automated. System Center Virtual Machine Manager 2012 (VMM 2012) enables end-to-end automation of this process through SAN integration using the Storage Management Initiative Specification (SMI-S) protocol.

Historically, many storage vendors have designed and implemented their own storage management systems, APIs, and command line utilities. This has made it a challenge to use a common set of tools, scripts, and so on across heterogeneous storage solutions.

For the robust automation that is required in an advanced data center virtualization, a SAN solution supporting SMI-S is required. Preference is also given to SANs supporting standard and common automation interfaces such as PowerShell.

EMC provides a SMI-S Provider infrastructure that supports the EMC VNX5300 and is complaint to System Center Virtual Machine Manager 2012 requirements. The solution implements the SMI-S provider to facilitate SCVMM 2012 automation.

In addition, the ESI/PowerShell framework can also be utilized in the environment as appropriate. For example, initial server provisioning can be automated through the usage of the ESI/PowerShell framework.

Network Architecture

Many network architectures include a tiered design, with three or more tiers such as core, distribution, and access. Designs are driven by the port bandwidth and quantity required at the edge, as well as the ability of the distribution and core tiers to provide higher-speed uplinks to aggregate traffic. Additional considerations include Ethernet broadcast boundaries and limitations, and spanning tree and other loop-avoidance technologies.

Cisco UCS offers a unique perspective on server-focused network architecture. An integrated network strategy at the core of Cisco UCS provides 10 Gigabit Ethernet connectivity to all components. Coupling this fabric with the stateless, policy-driven server architecture described previously in this document allows vast simplification of the physical infrastructure typically deployed in a new server build-out.

Rather than including localized Ethernet and FC switching in each chassis, all fabric aggregation is performed at a top-of-rack (ToR) type of device called the fabric interconnect. Cisco UCS 6200 Series Fabric Interconnects are a family of wired-rate, low-latency, lossless 10 Gigabit Ethernet, DCB, and FCoE interconnect switches that consolidate I/O at the system level. Based on the same switching technology as the Cisco Nexus 5500 Series Switches, the Cisco UCS 6200 Series Fabric Interconnects provide the additional features and management capabilities that make up the core of the Cisco UCS.

The fabric interconnects supply a unified fabric that connects every server in the system through wire-once 10 Gigabit Ethernet, FCoE downlinks, flexible 10 Gigabit Ethernet, and 1, 2, 4, or 8-Gbps FC uplinks (shown in the Figure 24). Out-of-band management, including switch redundancy, is supported through dedicated management and clustering ports. The interconnects feature front-to-back cooling, redundant front-plug fans and power supplies, and rear cabling that facilitates efficient cooling and serviceability. Typically deployed in active-active redundant pairs, the fabric interconnects provide uniform access to both networks and storage, which eliminates the barriers to deploying a fully virtualized environment based on a flexible, programmable pool of resources.

Figure 23. Unified Fabric

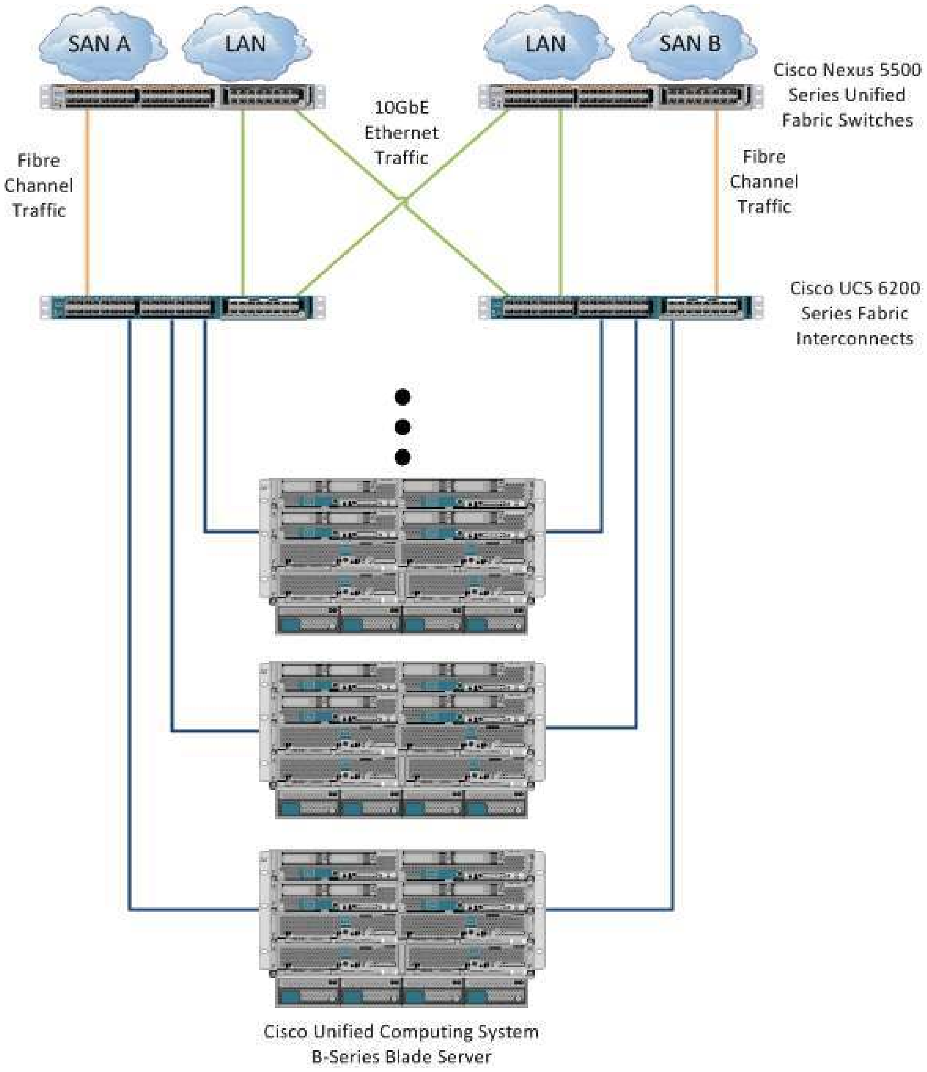
SAN A
LAN
LAN
SAN B
Cisco Nexus 5500 Series Unified Fabric Switches
10GbE Ethernet Traffic
Fibre Channel Traffic
Fibre Channel Traffic
Cisco UCS 6200 Series Fabric Interconnects
Cisco Unified Computing System B-Series Blade Server

As shown in Figure 24, the Cisco Unified Computing System can join any standards-compliant existing SAN and LAN infrastructure as it leaves the fabric interconnects.

High Availability and Resiliency

Providing redundant paths from the server through all the network tiers to the core is this solution’s joint best practice recommendation for high availability and resiliency. A variety of technologies (NIC teaming, Spanning Tree Protocol, and so on) can be used to create redundant path availability without looping.

Each network tier should include redundant switches. With redundant pairs of access-tier switches, individual switch resiliency is slightly less important, and therefore, the expense of redundant power supplies and other component redundancy may not be required. At the aggregation and core tiers, full hardware redundancy, in addition to device redundancy, is recommended due to the critical nature of those tiers.

However, sometimes devices fail, become damaged, or get misconfigured. In these situations, remote management and the ability to remotely power cycle all devices become important factors in restoring service rapidly.

Note: The network design must allow for the loss of any switch module or switch without dropping host server connectivity.

Cisco Unified Computing System

The Cisco Unified Computing System platform provides the following features that support high availability and resiliency:

- Redundant LAN and SAN fabrics
- Fabric failover hardware-based NIC teaming
- Port channel link aggregation, load balancing, and link fault tolerance
- Hot-swappable, field-replaceable power supplies, fan modules, and expansion modules
- 1+1 power redundancy
- N+1 fan module redundancy

Cisco Nexus 5548UP Switches

The Cisco Nexus platform provides the following high availability and resiliency features:

- Redundant LAN and SAN fabrics
- Virtual port channel
- Spanning Tree Protocol
- Link aggregation control protocol
- Cisco fabric path
- Hot-swappable field-replaceable power supplies, fan modules, and expansion modules
- 1+1 power redundancy
- N+1 fan module redundancy

Network Security and Isolation

The network architecture must enable both security and isolation of network traffic. A variety of technologies can be used individually or in concert to assist with security and isolation, including:

- **VLANs:** VLANs enable traffic on one physical LAN to be subdivided into multiple virtual LANs or broadcast domains. This is accomplished by configuring devices or switch ports to tag traffic with specific VLAN IDs. A VLAN trunk is a network connection that is able to carry multiple VLANs, with each VLAN tagged with specific VLAN IDs.
- **ACLs:** Access control lists (ACLs) enable traffic to be filtered or forwarded based on a variety of characteristics such as protocol, source and destination port, and many other characteristics. ACLs can be used to prohibit certain traffic types from reaching the network or to enable or prevent traffic from reaching specific endpoints.
- **IPsec:** IPsec enables both authentication and encryption of network traffic to protect from both man-in-the-middle attacks as well as network sniffing and other data collection activities.

- **QoS:** QoS enables rules to be set based on traffic type or attributes so that one form of traffic does not block all others (by throttling it) or to make sure critical traffic has a certain amount of bandwidth allocated.

Additional security and isolation technologies include:

- Standard and extended Layer 2 ACLs (MAC addresses, protocol type, and so on)
- Standard and extended Layer 3 to Layer 4 ACLs (IPv4 and IPv6, Internet Control Message Protocol [ICMP], and TCP)
- User Datagram Protocol (UDP) and so on
- VLAN-based ACLs (VACLs)
- Port-based ACLs (PACLs)
- Named ACLs
- Optimized ACL distribution
- ACLs on virtual terminals (VTYs)
- Dynamic Host Configuration Protocol (DHCP) snooping with option 82
- Dynamic Address Resolution Protocol (ARP) Inspection
- IP source guard
- DHCP relay
- Cisco TrustSec (CTS) (authentication and policy download from ACS)
- Ethernet port security

Network Automation

Remote interfaces and management of the network infrastructure through Secure Shell (SSH) or a similar protocol are important to both the automation and the resiliency of the data center network. Remote access and administration protocols can be used by management systems to automate complex or error-prone configuration activities. For example, adding a VLAN to a distributed set of access tier switches can be automated to avoid the potential for human error.

Cisco UCS Manager

Cisco UCS Manager offers the following features:

- A unified embedded management interface that integrates server, network, and storage access
- Policy and model-based management with service profiles that improve agility and reduce risk
- Autodiscovery to detect, inventory, manage, and provision system components that are added or changed
- A comprehensive open XML API that facilitates integration with third-party systems management tools
- Role-based administration that builds on existing skills and supports collaboration across disciplines

Cisco Nexus 5548UP Switch

The Cisco Nexus platform offers the following related features:

- Switch management using 10,100, or 1000-Mbps management or console ports
- CLI-based console to provide detailed out-of-band management
- In-band switch management
- Locator and beacon LEDs on Cisco Nexus 2000 Series

- Port-based locator and beacon LEDs
- Configuration synchronization
- Module preprovisioning
- Configuration rollback
- Secure Shell version 2 (SSHv2)
- Telnet
- Authentication, authorization, and accounting (AAA)
- AAA with role-based access control (RBAC)
- RADIUS
- Terminal Access Controller Access-Control System Plus (TACACS+)
- Syslog (8 servers)
- Embedded packet analyzer
- Simple Network Management Protocol (SNMPv1, v2, and v3) (IPv4 and IPv6)
- Enhanced SNMP MIB support
- XML (NETCONF) support
- Remote monitoring (RMON)
- Advanced Encryption Standard (AES) for management traffic
- Unified username and passwords across command-line interface (CLI) and SNMP
- Microsoft Challenge Handshake Authentication Protocol (MS-CHAP)
- Digital certificates for management between switch and RADIUS server
- Cisco Discovery Protocol Versions 1 and 2
- RBAC
- Switched Port Analyzer (SPAN) on physical, PortChannel, VLAN, and FC interfaces
- Encapsulated Remote SPAN (ERSPAN)
- Ingress and egress packet counters per interface
- Network Time Protocol (NTP)
- Cisco Generic Online Diagnostics (Cisco GOLD)
- Comprehensive bootup diagnostic tests
- Call Home feature of Cisco IOS Release 12.2SX
- Cisco Smart Call Home
- Cisco Fabric Manager
- Cisco Data Center Network Manager (DCNM)
- CiscoWorks LAN Management Solution (LMS)

Virtualization Architecture

Storage Virtualization

Storage virtualization, a concept in IT system administration, refers to the abstraction (separation) of logical storage from physical storage so that it can be accessed without regard to physical storage or heterogeneous structure. This separation allows systems administrators increased flexibility in how they manage storage for end users.

For more information about storage virtualization, visit: http://en.wikipedia.org/wiki/Storage_virtualization

Network Virtualization

In computing, network virtualization is the process of combining hardware and software network resources and network functionality into a single, software-based administrative entity, a virtual network. Network virtualization involves platform virtualization, often combined with resource virtualization.

Network virtualization is categorized as either external—combining many networks, or parts of networks, into a virtual unit—or internal, providing network-like functionality to the software containers on a single system. Whether virtualization is internal or external depends on the implementation provided by vendors who support the technology.

Various equipment and software vendors offer network virtualization by combining any of the following:

- Network hardware, such as switches and network adapters, also known as network interface cards (NICs)
- Networks, such as virtual LANs (VLANs) and containers such as virtual machines (VMs)
- Network storage devices
- Network media, such as Ethernet and Fibre Channel

For more information about network virtualization, visit:
http://www.snia.org/education/storage_networking_primer/stor_virt/

Server Virtualization

Hardware virtualization uses software to create a virtual machine (VM) that emulates a physical computer. This creates a separate OS environment that is logically isolated from the host server. By providing multiple VMs at once, this approach allows several operating systems to run simultaneously on a single physical machine.

Hyper-V technology is based on a 64-bit hypervisor-based microkernel architecture that enables standard services and resources to create, manage, and execute virtual machines. The Windows Hypervisor runs directly above the hardware and ensures strong isolation between the partitions by enforcing access policies for critical system resources such as memory and processors. The Windows Hypervisor does not contain any third-party device drivers or code, which minimizes its attack surface and provides a more secure architecture (see Figure 25).

Figure 24. Detailed Windows Hypervisor Architecture

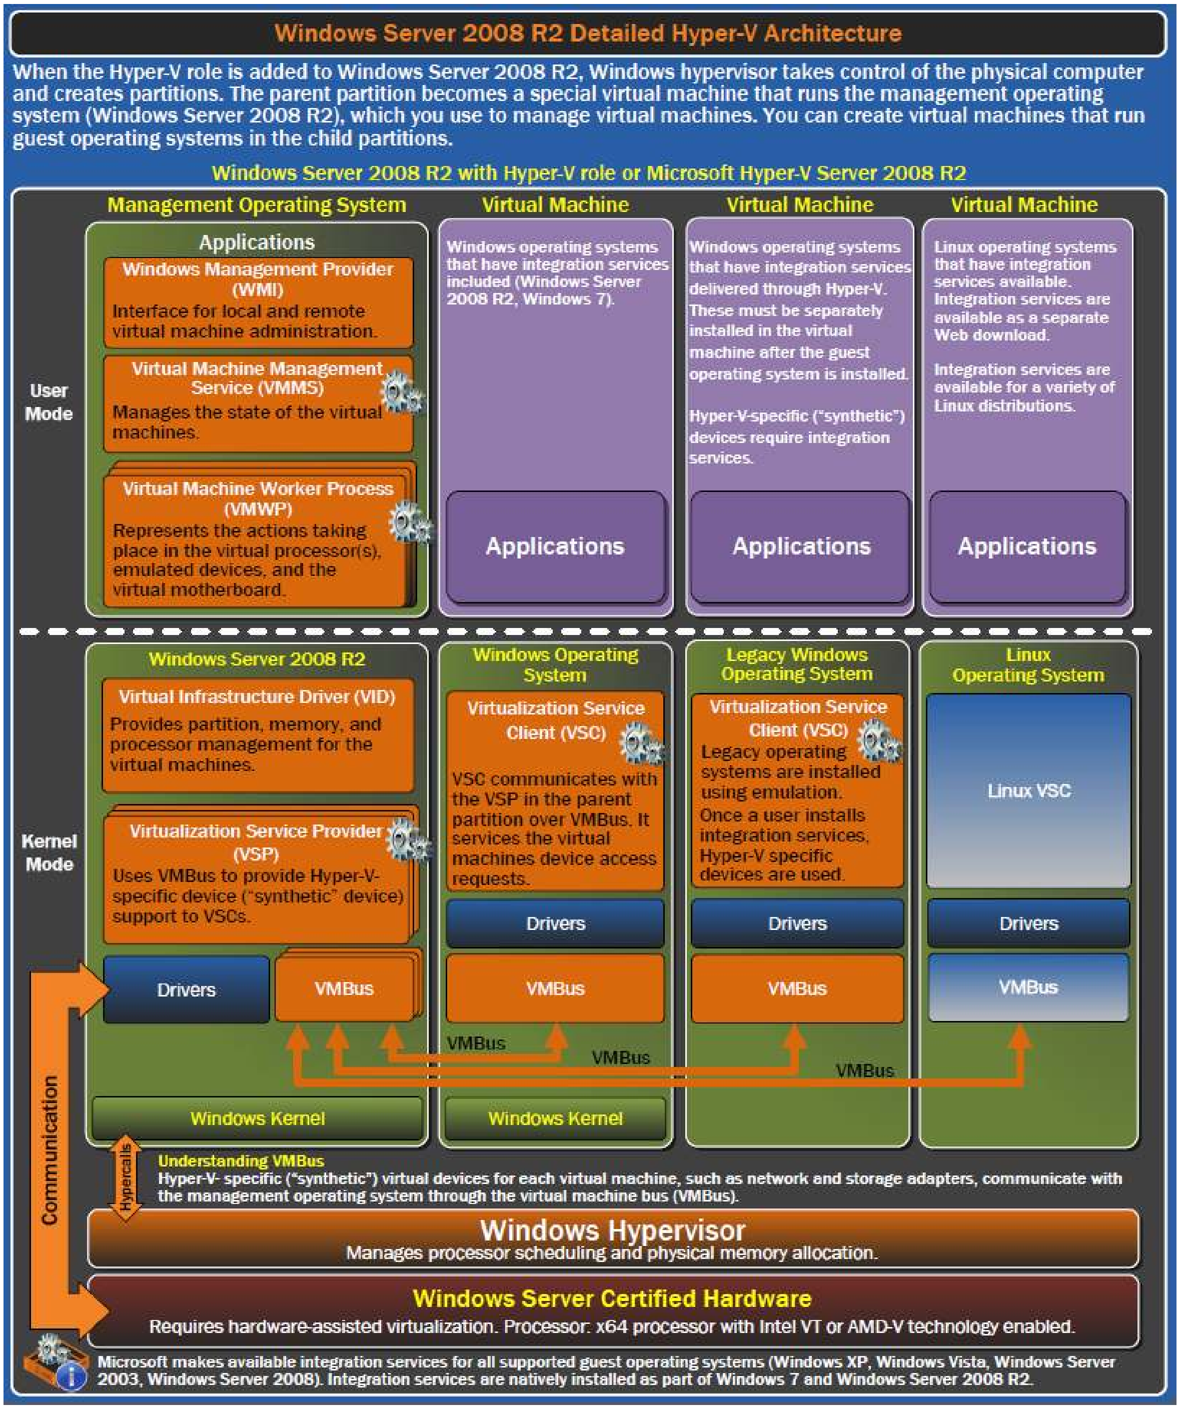

Windows Server 2008 R2 Detailed Hyper-V Architecture

When the Hyper-V role is added to Windows Server 2008 R2, Windows hypervisor takes control of the physical computer and creates partitions. The parent partition becomes a special virtual machine that runs the management operating system (Windows Server 2008 R2), which you use to manage virtual machines. You can create virtual machines that run guest operating systems in the child partitions.

Windows Server 2008 R2 with Hyper-V role or Microsoft Hyper-V Server 2008 R2

Management Operating System

Applications

Windows Management Provider (WMI)
Interface for local and remote virtual machine administration.

Virtual Machine Management Service (VMMS)
Manages the state of the virtual machines.

Virtual Machine Worker Process (VMWP)
Represents the actions taking place in the virtual processor(s), emulated devices, and the virtual motherboard.

Virtual Machine

Windows operating systems that have integration services included (Windows Server 2008 R2, Windows 7).

Virtual Machine

Windows operating systems that have integration services delivered through Hyper-V. These must be separately installed in the virtual machine after the guest operating system is installed.

Hyper-V-specific ("synthetic") devices require integration services.

Virtual Machine

Linux operating systems that have integration services available. Integration services are available as a separate Web download.

Integration services are available for a variety of Linux distributions.

Windows Server 2008 R2

Virtual Infrastructure Driver (VID)
Provides partition, memory, and processor management for the virtual machines.

Virtualization Service Provider (VSP)
Uses VMBus to provide Hyper-V-specific device ("synthetic" device) support to VSCs.

Windows Operating System

Virtualization Service Client (VSC)
VSC communicates with the VSP in the parent partition over VMBus. It services the virtual machines device access requests.

Legacy Windows Operating System

Virtualization Service Client (VSC)
Legacy operating systems are installed using emulation.
Once a user installs integration services, Hyper-V specific devices are used.

Linux Operating System

Understanding VMBus
Hyper-V- specific ("synthetic") virtual devices for each virtual machine, such as network and storage adapters, communicate with the management operating system through the virtual machine bus (VMBus).

Windows Hypervisor
Manages processor scheduling and physical memory allocation.

Windows Server Certified Hardware
Requires hardware-assisted virtualization. Processor: x64 processor with Intel VT or AMD-V technology enabled.

Microsoft makes available integration services for all supported guest operating systems (Windows XP, Windows Vista, Windows Server 2003, Windows Server 2008). Integration services are natively installed as part of Windows 7 and Windows Server 2008 R2.

In addition to the Windows Hypervisor, there are two other major elements to consider in Hyper-V: a parent partition and child partition. The parent partition is a special virtual machine that runs Windows Server 2008 R2 SP1, controls the creation and management of child partitions, and maintains direct access to hardware resources. In this model, device drivers for physical devices are installed in the parent partition. In contrast, the role of a child partition is to provide a virtual machine environment for the installation and execution of guest operating systems and applications.

Please see this detailed poster for more information: http://www.microsoft.com/download/en/details.aspx?id=2688

Windows Server 2008 R2 SP1 and Hyper-V Host Design

The recommendations in this section adhere to the support statements in the following article:

Requirements and Limits for Virtual Machines and Hyper-V in Windows Server 2008 R2
http://technet.microsoft.com/en-us/library/ee405267(WS.10).aspx

Licensing

Certain versions of Windows Server 2008 R2 (namely Standard, Enterprise, and Datacenter editions) include “virtualization use rights,” which is the right and license to run a specified number of Windows-based virtual machines. Windows Server 2008 R2 Standard edition includes use rights for one running virtual machine. Windows Server 2008 R2 Enterprise Edition includes use rights for up to four virtual machines. This does not limit the number of guests that the host can run; it means that licenses for four Windows Server guests are included. To run more than four, you simply need to ensure you have valid Windows Server licenses for the additional virtual machines.

In contrast to the two other Windows Server editions, Windows Server 2008 R2 Datacenter Edition includes unlimited virtualization use rights. From a licensing standpoint, this allows you to run as many Windows Server guests as you like on the licensed physical server.

The Cisco and EMC solution utilizes Microsoft Windows Server 2008 R2 SP1 Datacenter Edition for all physical server infrastructure.

Operating System Configuration

The following outlines the general considerations for the Hyper-V host operating system. Note that these are not meant to be installation instructions but rather the process requirements and process order.

Hyper-V requires specific hardware. To install and use the Hyper-V role, you will need the following:

- An x64-based processor. Hyper-V is available in 64-bit editions of Windows Server 2008 R2 – specifically, the 64-bit editions of Windows Server 2008 Standard, Windows Server 2008 Enterprise, and Windows Server 2008 Datacenter. Hyper-V is not available for 32-bit (x86) editions or Windows Server 2008 for Itanium-Based Systems. However, the Hyper-V management tools are available for 32-bit editions.
- Hardware-assisted virtualization. This is available in processors that include a virtualization option—specifically processors with Intel Virtualization Technology (Intel VT) or AMD Virtualization (AMD-V) technology. Hardware-enforced Data Execution Prevention (DEP) must be available and enabled. Specifically, you must enable Intel XD bit (execute disable bit) or AMD NX bit (no execute bit).
- Use Windows Server 2008 R2 SP1, either Full or Server Core installation option. Note: there is no upgrade path from Server Core to Full or vice-versa: make this selection carefully.
- Use the latest hardware device drivers.
- Hyper-V parent partition OS must be domain-joined.
- Hyper-V server role and Failover Clustering features are required
- Apply relevant Windows updates, including OOB updates not offered on Microsoft Update.

- Hyper-V Update List for Windows Server 2008 R2: http://social.technet.microsoft.com/wiki/contents/articles/hyper-v-update-list-for-windows-server-2008-r2/rss.aspx
- All nodes, networks, and storage must pass Cluster Validation Wizard.

Memory and Hyper-V Dynamic Memory

Dynamic Memory is a Hyper-V feature that helps you use physical memory more efficiently. With Dynamic Memory, Hyper-V treats memory as a shared resource that can be reallocated automatically among running virtual machines. Dynamic Memory adjusts the amount of memory available to a virtual machine, based on changes in memory demand and values that you specify. Dynamic Memory is available for Hyper-V in Windows Server 2008 R2 Service Pack 1 (SP1). You can make the Dynamic Memory feature available by applying the service pack to the Hyper-V role in Windows Server 2008 R2 or to Microsoft Hyper-V Server 2008 R2.

For a complete description of Dynamic Memory features, settings, and design considerations, refer to the Hyper-V Dynamic Memory Configuration Guide at: http://technet.microsoft.com/en-us/library/ff817651(WS.10).aspx. This guide provides the specific OS, service pack, and integration component levels for supported operating systems. The guide also contains the minimum recommended startup RAM setting for all supported operating systems.

In addition to the general guidance above, specific applications or workloads, particularly those with built in memory management capability such as SQL or Exchange, may provide workload specific guidance. The Fast Track Reference Architecture utilizes SQL 2008 R2. The SQL product group has published best practices guidance for Dynamic Memory in Running SQL Server with Hyper-V Dynamic Memory: http://msdn.microsoft.com/en-us/library/hh372970.aspx.

Storage Adapters

Unlike network adapters, storage adapters are certified by both the operating system and the storage provider. In this solution, the Cisco UCS platform has been carefully tuned to be compatible with both Windows Server and the EMC VNX5300 storage platform.

Fibre Channel and iSCSI Configuration

The solution uses redundant Fibre Channel connections to storage connected to the Hyper-V host systems. The Fibre Channel connections are used for boot from SAN for the hosts and Cluster Shared Volumes for storage of the virtual machines virtual hard drives. Redundant iSCSI connections are provided for shared storage for clustered virtual machines.

MPIO Configuration

Microsoft MPIO architecture supports iSCSI, Fibre Channel, and serial attached storage (SAS) SAN connectivity by establishing multiple sessions or connections to the storage array.

Multipathing solutions use redundant physical path components—adapters, cables, and switches —to create logical paths between the server and the storage device. In the event that one or more of these components fails, causing the path to fail, multipathing logic uses an alternative path for I/O so that applications can still access their data. Each network interface card (in the iSCSI case) or HBA should be connected by using redundant switch infrastructures to provide continued access to storage in the event of a failure in a storage fabric component.

EMC PowerPath provides support for both Fiber Channel/FCoE and iSCSI connectivity utilizing Windows MPIO infrastructure. Any additional documentation required will be provided.

Network Adapters

Network adapters and the way in which the network is configured have a direct correlation to the health and stability of your Windows failover cluster. This solution provides a Microsoft best practice networking environment designed for maximum performance and availability.

Protocol Bindings for Adapters

Table 2 lists protocol bindings for Adapters

Table 2. Protocol Bindings for Adapters

Setting	Mgmt Network Adapter	Heartbeat Network Adapter	Live Migration Network Adapter	iSCSI Network Adapter	Guest VM Network Adapter
Client for Microsoft Networks	Y	N	Y	N	N
File and Printer Sharing	Y	N	Y	N	N
Microsoft Virtual Network Switch Protocol	N	N	N	N	Y
Internet Protocol Version 6	O	O	O	O	N
Internet Protocol Version 4	Y	Y	Y	Y	N
Link-Layer Topology Discovery Mapper I/O Driver	Y	N	N	N	N
Link-Layer Topology Discovery Responder	Y	N	N	N	N

Performance Settings

The following Hyper-V R2 network performance improvements should be tested and considered for production use:

- TCP Checksum Offload is recommended. It benefits both CPU and overall network throughput performance, and is fully supported when performing a Live Migration.
- Support for Jumbo Frames was also introduced with Windows Server 2008. Hyper-V in Windows Server 2008 R2 simply extends this capability to VMs. So just as in physical network scenarios, Jumbo Frames add the same basic performance enhancements to virtual networking. That includes up to six times larger payloads per packet, which improves not only overall throughput but also reduces CPU utilization for large file transfers.
- Virtual machine queue (VMQ) technology essentially allows the host’s network interface card (NIC) to do direct memory access (DMA) on packets directly into individual VM memory stacks. Each VM device buffer is assigned a VMQ, which avoids needless packet copies and route lookups in the virtual switch. The result is less data in the host’s buffers and an overall performance improvement in I/O operations.

- The Cisco and EMC design implements a multi-VLAN environment to provide segregation of network interfaces. Support is provided for cluster communication and iSCSI networks that are isolated from other networks via VLAN implementation. Quality of service mechanisms within the switched environment are used to manage impact to critical VLANs.

NIC Teaming Configurations

NIC teaming can be utilized to enable multiple, redundant NICs and connections between servers and access tier network switches. Teaming can be enabled via hardware or software-based approaches. Teaming can enable multiple scenarios including path redundancy, failover, and load balancing.

Cisco UCS provides hardware-based NIC teaming called fabric failover. Each NIC can be enabled for fabric failover and pinned to either fabric A or fabric B. Fabric failover is transparent to the Windows operating system and the Hyper-V network switch. No extra software needs to be installed to benefit from fabric failover hardware-based NIC teaming.

Hyper-V Host Failover Cluster Design

A Hyper-V host failover cluster is a group of independent servers that work together to increase the availability of applications and services. The clustered servers (called nodes) are connected by physical cables and by software. If one of the cluster nodes fails, another node begins to provide service (a process known as failover). In case of a planned migration (called a Live Migration), users experience no perceptible service interruption.

The host servers are one of the critical components of a dynamic, virtual infrastructure. Consolidation of multiple workloads onto the host servers requires that those servers be highly available. Windows Server 2008 R2 provides advances in failover clustering that enable high availability and Live Migration of virtual machines between physical nodes.

Host Failover Cluster Topology

The Microsoft Private Cloud Fast Track program provides for a design pattern that relies on a single Hyper-V host cluster to provide both the management infrastructure and the compute nodes for running the private cloud virtual machines. This is the solution that Cisco and EMC have validated.

Each host cluster can contain up to 16 nodes. Host clusters require some form of shared storage such as a Fibre Channel or iSCSI SAN.

Host Cluster Networks

A variety of host cluster networks are required for a Hyper-V failover cluster. The network requirements enable high availability and high performance. The specific requirements and recommendations for network configuration are published on TechNet in the Hyper-V: Live Migration Network Configuration Guide: http://technet.microsoft.com/en-us/library/ff428137(WS.10).aspx

Management Network

A dedicated management network is required so hosts can be managed via a dedicated network such that there is no competition with guest traffic needs. A dedicated network provides a degree of separation for security and ease of management purposes. This typically implies dedicating a network adapter per host and port per network device to the management network. Additionally, Cisco offers a separate out-of-band management capability that enables remote management by Cisco UCS of the host operating system. For more information, visit: Cisco.com UCS Manager CLI Configuration Guide.

iSCSI Network

If using iSCSI, a dedicated iSCSI network is required so that storage traffic is not in contention with any other traffic. This typically implies dedicating two network adapters per host and two ports per network device to the storage network. For all iSCSI storage connections, EMC's PowerPath software provides an MPIO configuration using two separate VLANs.

CSV/Cluster Communications Network

Usually, when the cluster node that owns a virtual hard disk (VHD) file in CSV performs disk I/O, the node communicates directly with the storage—for example, through a storage area network (SAN). However, storage connectivity failures sometimes prevent a given node from communicating directly with the storage. To maintain function until the failure is corrected, the node redirects the disk I/O through a cluster network (the preferred network for CSV) to the node where the disk is currently mounted. This is called CSV redirected I/O mode.

Live Migration Network

During Live Migration, the contents of the memory of the virtual machine running on the source node need to be transferred to the destination node over a LAN connection. To provide a high-speed transfer, a dedicated redundant 10-Gbps Live Migration network is required. In this case, all Ethernet networking uses the 10-GB converged network that is provided by the Cisco Nexus 5548UP Switches that are specified as part of this solution. This significantly reduces the time required to evacuate the virtual machines from a host, with zero downtime during maintenance or Windows updates. QoS can be used so that sufficient bandwidth is reserved for this network.

Virtual Machine Networks

The virtual machine networks are dedicated to virtual machine LAN traffic. The VM network can be two or more Gigabit Ethernet networks, one or more network created via NIC teaming, or virtual networks created from shared 10 Gigabit Ethernet NICs.

Host Failover Cluster Storage

Cluster Shared Volumes (CSV) is a feature that simplifies the configuration and management of Hyper-V virtual machines in failover clusters. With CSV, on a failover cluster that runs Hyper-V, multiple virtual machines can use the same LUN (disk) yet failover (or move from node to node) independently of one another. CSV provides increased flexibility for volumes in clustered storage—for example, it allows you to keep system files separate from data to optimize disk performance, even if the system files and the data are contained within virtual hard disk (VHD) files. If you choose to use Live Migration for your clustered virtual machines, CSV can also provide performance improvements for the live migration process. CSV is available in versions of Windows Server 2008 R2 and of Microsoft Hyper-V Server 2008 R2 that include failover clustering. The CSV functionality is implemented at the parent partition level where appropriate.

Hyper-V Guest VM Design

Standardization is a key tenet of private cloud architectures. This also applies to virtual machines. A standardized collection of virtual machine templates can both drive predictable performance and greatly improve capacity planning capabilities. As an example, Table 3 illustrates what a basic VM template library would look like.

Table 3. Example of a Basic VM Template Library

Template	Specs	Network	OS	Unit Cost
Template 1 – Small	1 vCPU, 2GB Memory, 50gb Disk	VLAN 20	WS 2003 R2	1
Template 2 – Med	2 vCPU, 4GB Memory, 100gb Disk	VLAN 20	WS 2003 R2	2

Template	Specs	Network	OS	Unit Cost
Template 3 – X-Large	4 vCPU, 8GB Memory, 200gb Disk	VLAN 20	WS 2003 R2	4
Template 4 – Small	1 vCPU, 2GB Memory, 50gb Disk	VLAN 10	WS 2008	1
Template 5 – Med	2 vCPU, 4GB Memory, 100gb Disk	VLAN 10	WS 2008	2
Template 6 – X-Large	4 vCPU, 8GB Memory, 200gb Disk	VLAN 10	WS 2008	4

Note: Use standard documented virtual machine configurations for all virtual machines, management, and tenants used for fabric management or for workload deployment by tenants.

Virtual Machine Storage

This sections describes the different types of virtual hard drives that can be used by Hyper-V.

Microsoft recommends using only fixed VHDs for production. The following quote is from Windows TechNet:

> ***"Why are fixed VHDs recommended for production?"***

Fixed VHDs are recommended for production instead of dynamically expanding or differencing VHDs for the following reasons:

> ***The I/O performance is highest for fixed VHDs because the file is not dynamically expanded.***
>
> ***When a dynamically expanding disk is expanded, the host volume could run out of space and cause the write operations to fail. Using fixed VHDs prevents this from happening.***
>
> ***The file data will not become inconsistent due to lack of storage space or power loss. Dynamically expanding and differencing VHDs depend on multiple write operations to expand the file. The internal block allocation information can become inconsistent if all I/O operations to the VHD file and the host volume are not complete and persist on the physical disk. This can happen if the computer suddenly loses power.***

Dynamically Expanding Disks

Dynamically expanding virtual hard disks provide storage capacity as needed to store data. The size of the VHD file is small when the disk is created and grows as data is added to the disk. The size of the VHD file does not shrink automatically when data is deleted from the virtual hard disk. However, you can compact the disk to decrease the file size after data is deleted by using the Edit Virtual Hard Disk Wizard.

Fixed Size Disks

Fixed virtual hard disks provide storage capacity by using a VHD file that is in the size specified for the virtual hard disk when the disk is created. The size of the VHD file remains "fixed" regardless of the amount of data stored. However, you can use the Edit Virtual Hard Disk Wizard to increase the size of the virtual hard disk, which increases the size of the VHD file. If you allocate the full capacity at the time of creation, fragmentation at the host level is not an issue (fragmentation inside the VHD itself must be managed within the guest).

Differencing Disks

Differencing virtual hard disks provide storage to enable you to make changes to a parent virtual hard disk without altering that disk. The size of the VHD file for a differencing disk grows as changes are stored to the disk.

Pass-Through Disks

Hyper-V enables virtual machine guests to directly access local disks or SAN LUNs that are attached to the physical server without requiring the volume to be presented to the host server. The virtual machine guest accesses the disk directly (utilizing the disk's global unique identifier [GUID]) without having to use the host's file system. Given that the performance difference between fixed-disk and pass-through disks is now negligible, the decision is based on manageability.

For instance, if the data on the volume will be very large (hundreds of gigabytes), a VHD is hardly portable at that size given the extreme amounts of time it takes to copy. Also, bear in mind the backup scheme. With pass-through disks, the data can be backed up only from within the guest. When utilizing pass-through disks, there is no VHD file created; the LUN is used directly by the guest. Since there is no VHD file, there is no dynamic sizing capability or snapshot capability.

In-guest iSCSI Initiator

Hyper-V can also utilize iSCSI storage by directly connecting to iSCSI LUNs utilizing the guest's virtual network adapters. This is mainly used for access to large volumes, volumes on SANs which the Hyper-V host itself is not connected to, or for guest clustering. Guests cannot boot from iSCSI LUNs accessed through the virtual network adapters without utilizing a third-party iSCSI initiator.

Virtual Machine Networking

Hyper-V guests support two types of virtual network adapters: synthetic and emulated. Synthetic adapters make use of the Hyper-V VMBUS architecture and are the high-performance, native devices. Synthetic devices require the Hyper-V Integration Services be installed within the guest. Emulated adapters are available to all guests even if Integration Services are not available. They are much slower performing and should be used only if synthetic is unavailable.

You can create many virtual networks on the server running Hyper-V to provide a variety of communications channels. For example, you can create networks to provide the following:

- **Private network:** Permits communications between virtual machines only
- **Internal network:** Permits communications between the host server and virtual machines
- **External network:** Permits communications between a virtual machine and a physical network by creating an association to a physical network adapter on the host server

Virtual Processors

Table 4 lists the number of supported virtual processors in a Hyper-V guest. Please note that this information does change; improvements to the Integration Services for Hyper-V are periodically released, adding support for additional operating systems. Please see the following article to access the most current information: http://technet.microsoft.com/en-us/library/cc794868(WS.10).aspx

Table 4. Supported Virtual Processors in Hyper-V Guests

Server Guest Operating System	Editions	Virtual Processors
Windows Server 2008 R2 with SP 1	Standard, Enterprise, Datacenter, and Web editions	1, 2, 3, or 4
Windows Server 2008 R2	Standard, Enterprise, Datacenter, and Windows Web Server 2008 R2	1, 2, 3, or 4
Windows Server 2008	Standard, Standard without Hyper-V, Enterprise, Enterprise without Hyper-V, Datacenter, Datacenter without Hyper-V, Windows Web Server 2008, and HPC Edition	1, 2, 3, or 4
Windows Server 2003 R2 with SP 2	Standard, Enterprise, Datacenter, and Web	1 or 2
Windows Home Server 2011	Standard	1, 2 or 4
Windows Storage Server 2008 R2	Essentials	1, 2 or 4
Windows Small Business Server 2011	Essentials	1 or 2
Windows Small Business Server 2011	Standard	1, 2, or 4
Windows Server 2003 R2 x64 Edition with SP 2	Standard, Enterprise, and Datacenter	1 or 2
Windows Server 2003 with SP 2	Standard, Enterprise, Datacenter, and Web	1 or 2
Windows Server 2003 x64 Edition with SP 2	Standard, Enterprise, and Datacenter	1 or 2
Windows 2000 Server with SP 4 IMPORTANT: Support for this operating system ended on July 13, 2010. For more information, see http://technet.microsoft.com/en-us/library/cc794868(WS.10).aspx	Server, Advanced Server	1
CentOS 6.0 and 6.1	x86 edition and x64 edition	1, 2, or 4
CentOS 5.2-5.7	x86 edition and x64 edition	1, 2, or 4
Red Hat Enterprise Linux 6.0 and 6.1	x86 edition and x64 edition	1, 2, or 4
Red Hat Enterprise Linux 5.7	x86 edition and x64 edition	1, 2, or 4
Red Hat Enterprise Linux 5.6	x86 edition and x64 edition	1, 2, or 4
Red Hat Enterprise Linux 5.5	x86 edition and x64 edition	1, 2, or 4
Red Hat Enterprise Linux 5.4	x86 edition and x64 edition	1, 2, or 4
Red Hat Enterprise Linux 5.3	x86 edition and x64 edition	1, 2, or 4
Red Hat Enterprise Linux 5.2	x86 edition and x64 edition	1, 2, or 4
SUSE Linux Enterprise Server 11 with SP 1	x86 edition and x64 edition	1, 2, or 4
SUSE Linux Enterprise Server 10 with SP 4	x86 edition and x64 edition	1, 2, or 4

Client Guest Operating System	Editions	Virtual Processors
Windows 7 with SP 1	Enterprise, Ultimate, and Professional. This applies to both 32-bit and 64-bit editions, as well as N and KN editions.	1, 2, 3, or 4
Windows 7	Enterprise, Ultimate, and Professional. This applies to both 32-bit and 64-bit editions, as well as N and KN editions.	1, 2, 3, or 4
Windows Vista	Business, Enterprise, and Ultimate, including N and KN editions	1 or 2

Client Guest Operating System	Editions	Virtual Processors
Windows XP with SP 3 (SP3) Important: Performance might be degraded on Windows XP with SP3 when the server running Hyper-V uses an AMD processor. For more information, see Degraded I/O Performance Using a Windows XP Virtual Machine with Windows Server 2008 Hyper-V (http://go.microsoft.com/fwlink/?LinkId=208750).	Professional	1 or 2
Windows XP with SP 2 IMPORTANT: Support for this operating system ended on July 13, 2010.	Professional	1
Windows XP x64 Edition with SP2	Professional	1 or 2

Hyper-V supports a maximum ratio of eight virtual processors (VPs) per one logical processor (LP) for server workloads, and 12 VPs per one LP for VDI workloads. A logical processor is defined as a processing core seen by the host operating system or parent partition. In the case of Intel Hyper-Threading, each thread is considered an LP.

Therefore, a 16-LP server supports a maximum of 128 VPs. That would in turn equate to 128 single-processor VMs, 64 dual-processor VMs, or 32 quad-processor VMs. The 8:1 or 12:1 VP/LP ratios are maximum supported limits. It is recommended that you use lower limits than the maximum.

Management Architecture

This section discusses the details of the underlying management architecture that is included in the solution.

Management Logical Architecture

Design Pattern 2

Figure 26 depicts the management logical architecture if running the management systems directly on the fabric cluster.

Figure 25. Combined Fabric and Management Infrastructure

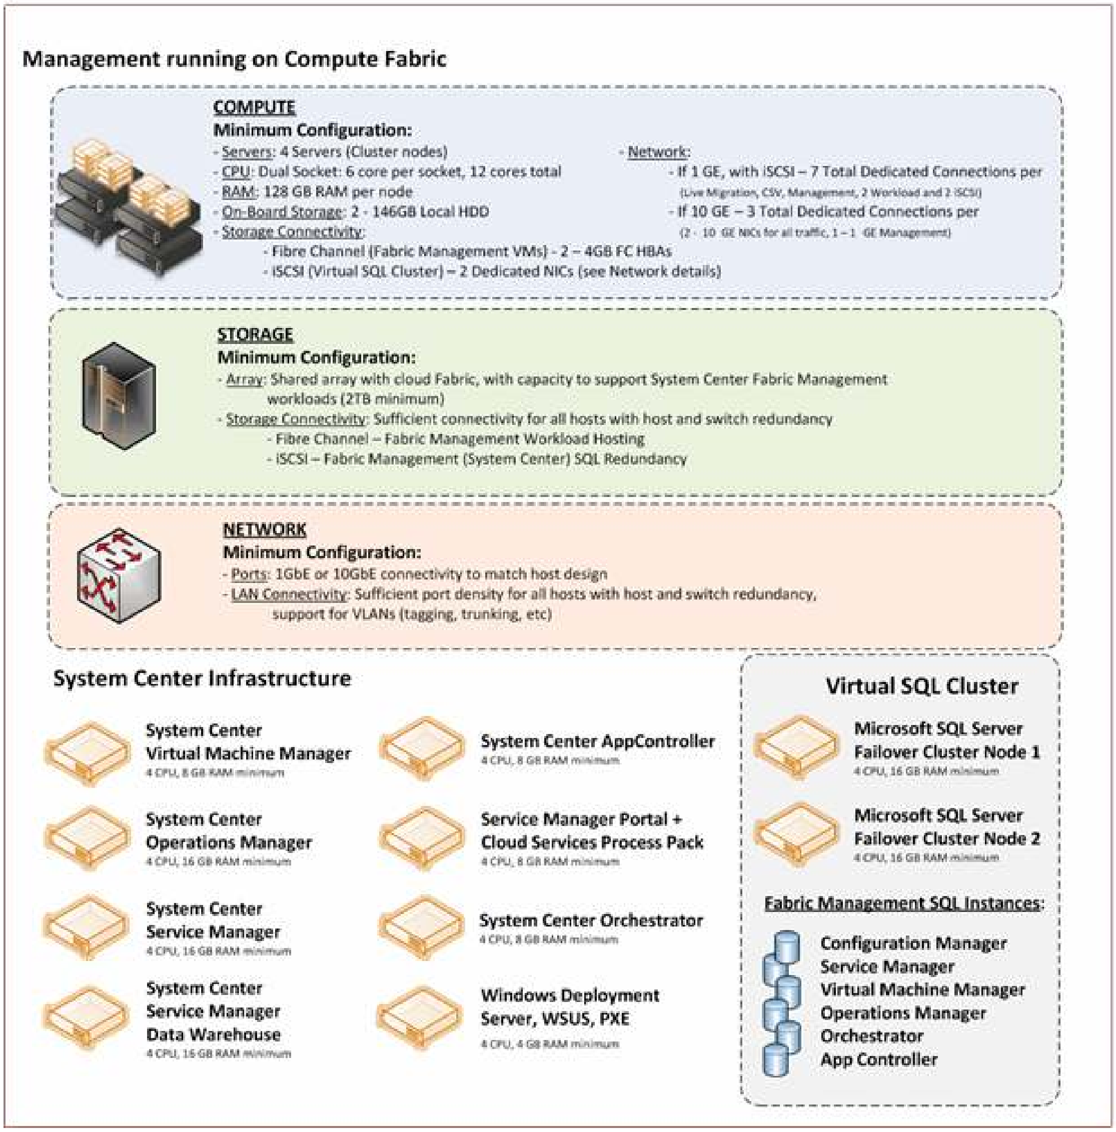
Management running on Compute Fabric
COMPUTE
Minimum Configuration:
- Servers: 4 Servers (Cluster nodes)
- CPU: Dual Socket: 6 core per socket, 12 cores total
- RAM: 128 GB RAM per node
- On-Board Storage: 2 - 146GB Local HDD
- Storage Connectivity:
- Fibre Channel (Fabric Management VMs) - 2 – 4GB FC HBAs
- iSCSI (Virtual SQL Cluster) – 2 Dedicated NICs (see Network details)
- Network:
- If 1 GE, with iSCSI – 7 Total Dedicated Connections per
(Live Migration, CSV, Management, 2 Workload and 2 iSCSI)
- If 10 GE – 3 Total Dedicated Connections per
(2 - 10 GE NICs for all traffic, 1 – 1 GE Management)
STORAGE
Minimum Configuration:
- Array: Shared array with cloud Fabric, with capacity to support System Center Fabric Management workloads (2TB minimum)
- Storage Connectivity: Sufficient connectivity for all hosts with host and switch redundancy
- Fibre Channel – Fabric Management Workload Hosting
- iSCSI – Fabric Management (System Center) SQL Redundancy
NETWORK
Minimum Configuration:
- Ports: 1GbE or 10GbE connectivity to match host design
- LAN Connectivity: Sufficient port density for all hosts with host and switch redundancy, support for VLANs (tagging, trunking, etc)
System Center Infrastructure
System Center Virtual Machine Manager
4 CPU, 8 GB RAM minimum
System Center Operations Manager
4 CPU, 16 GB RAM minimum
System Center Service Manager
4 CPU, 16 GB RAM minimum
System Center Service Manager Data Warehouse
4 CPU, 16 GB RAM minimum
System Center AppController
4 CPU, 8 GB RAM minimum
Service Manager Portal + Cloud Services Process Pack
4 CPU, 8 GB RAM minimum
System Center Orchestrator
4 CPU, 8 GB RAM minimum
Windows Deployment Server, WSUS, PXE
4 CPU, 4 GB RAM minimum
Virtual SQL Cluster
Microsoft SQL Server Failover Cluster Node 1
4 CPU, 16 GB RAM minimum
Microsoft SQL Server Failover Cluster Node 2
4 CPU, 16 GB RAM minimum
Fabric Management SQL Instances:
Configuration Manager
Service Manager
Virtual Machine Manager
Operations Manager
Orchestrator
App Controller

The management architecture consists of a minimum of four physical nodes in a failover cluster with SAN-attached storage that supports iSCSI and redundant network connections. This is to provide a highly available platform for the management systems and capacity for workloads. In this scenario, we have scaled down the high-availability options for management in order to facilitate a smaller management footprint on the fabric.

The management systems include:

- 2 x SQL servers in a guest cluster configuration
- 1 x Virtual Machine Manager server
- 1 x Operations Manager server
- 1 x Orchestrator server

- 1 x Service Manager servers
- 1 x Service Manager Data Warehouse
- 1 x Service Manager Self-Service Portal w/ Cloud Service Process Pack
- 1 x App Controller server
- 1 x deployment server providing Windows Deployment Services (WDS), preboot execution environment (PXE), and Windows Server Update Services (WSUS)

The Cisco and EMC design uses a management layer that implements:

- System Center Virtual Machine Manager 2012
- System Center Operations Manager 2012
- System Center Orchestrator 2012
- System Center Virtual Machine Manager 2012 Self-Service Portal
- System Center Service Manager 2012
- System Center App Controller 2012
- (Optional) Windows Server 2008 R2 SP providing WDS, PXE, and WSUS

Management Systems Architecture

Prerequisite Infrastructure

The following section outlines this architecture and its dependencies within a customer environment.

Active Directory Domain Services

Active Directory Domain Services (AD DS) is a required foundational component. Fast Track provides support for Windows Server 2008 and Windows Server 2008 R2 SP1 AD DS customer deployments. Previous versions are not directly supported for all workflow provisioning and deprovisioning automation. It is assumed that AD DS deployments exist at the customer site and deployment of these services is not in scope for the typical deployment.

- **Forests and domains:** The preferred approach is to integrate into an existing AD DS forest and domain. This is not a hard requirement; a dedicated resource forest or domain can also be employed as an additional part of the deployment. Fast Track does support multiple domains and/or multiple forests in a trusted environment using two-way forest trusts.
- **Trusts:** Fast Track enables multidomain support within a single forest where two-way forest (Kerberos) trusts exist between all domains. This is referred to as multidomain or interforest support. Also supported are interforest or multiforest scenarios as well as intraforest environments.

DNS

DNS name resolution is a required element for System Center 2012 components and the process automation solution. Active Directory-integrated DNS is required for automated provisioning and deprovisioning components within the System Center Orchestrator runbook as part of the solution. We provide full support and automation for Windows Server 2008 and Windows Server 2008 R2 SP1 DNS Active Directory-integrated DNS deployments.

Use of non-Microsoft or non-Active Directory-integrated DNS solutions may be possible but would not provide for automated creation and removal of DNS records related to virtual machine provisioning and deprovisioning processes. Use of solutions outside of Active Directory-integrated DNS would either require manual intervention for these scenarios or require modifications to Cloud Services Process Pack Orchestrator runbooks.

DHCP

To support dynamic provisioning and management of physical and virtual compute capacity within the IaaS infrastructure, you should use DHCP for all physical and virtual machines by default to support runbook automation. For physical hosts such as the fabric management cluster nodes and the scale-unit cluster nodes, DHCP reservations are recommended so that physical servers and NICs always have known IP addresses while providing centralized management of those addresses via DHCP.

Windows DHCP is required for automated provisioning and deprovisioning components within System Center Orchestrator runbooks as part of the solution. This is used to support host cluster provisioning, DHCP reservations, and other areas supporting dynamic provisioning of compute within the infrastructure. We provide full support and automation for Windows Server 2008 and Windows Server 2008 R2 SP1 versions of the DHCP server role. Use of solutions outside of the Windows DHCP Server role require additional testing and validation activities.

The Cisco and EMC design implements the use fixed IP address assignment for core infrastructure requirements. Customer configuration can utilize DHCP services as necessary, and are fully supported within the solution design.

SQL Server

Two SQL servers are deployed to support the solution. Two of the SQL servers are configured as a failover cluster, containing all the databases for each System Center product. By default, a two-node SQL guest VM cluster is utilized. Each SQL VM is configured with four vCPUs, at least 16 GB of RAM (32 GB recommended for large scale configurations), and four vNICs. The SQL VMs access iSCSI-based shared storage with two LUNs configured for each hosted database. A guest cluster is utilized for several reasons: namely, to maintain HA of the System Center databases during both host and guest OS patching and during host or guest failures. Should the CPU, RAM, and I/O needs of the solution exceed what two VMs are able to provide, additional VMs can be added to the virtual SQL cluster and each SQL instance moved to its own VM in the cluster. This requires SQL 2008 Enterprise Edition, which is the recommendation. In addition, where organizations can support it, solid-state drive (SSD) storage should be used to provide the necessary I/O for these databases.

Management SQL Server Configuration

- 2 non-HA VMs on different Hyper-V hosts
- Windows Server 2008 R2 SP1 Enterprise or Datacenter
- 4 vCPU
- 16-GB memory (do not use dynamic memory)
- 4 vNICs (1 client connections, 1 cluster communications, 2 iSCSI)
- Storage: 1 witness iSCSI LUN, 1 MSDTC iSCSI LUN, 20 x dedicated iSCSI LUNs

Databases

Table 5 shows the required Microsoft SQL Server database instances required for the installation of the System Center components. Two LUNs are recommended for each instance. One LUN would contain the data and the other LUN the log file.

Table 5. SQL Server Database Instances

Fabric Management Component	Instance Name (Suggested)	Components	Collation [4]	Storage Requirements
Virtual Machine Manager	SCVMMDB	Database Engine	SQL_Latin1_General_CP1_CI_AS	2 LUNs
Operations Manager	SCOMDB	Database Engine, Full-Text Search	SQL_Latin1_General_CP1_CI_AS	2 LUNs
Operations Manager Data Warehouse	SCOMDW	Database Engine, Full-Text Search	SQL_Latin1_General_CP1_CI_AS	2 LUNs
Service Manager	SCSMDB	Database Engine, Full-Text Search	SQL_Latin1_General_CP1_CI_AS	2 LUNs
Service Manager Data Warehouse	SCSMDW	Database Engine, Full-Text Search	SQL_Latin1_General_CP1_CI_AS	2 LUNs
	SCSMAS	Analysis Services	SQL_Latin1_General_CP1_CI_AS	2 LUNs
	SCSPFarm	Database Engine	SQL_Latin1_General_CP1_CI_AS	2 LUNs
Orchestrator	SCODB	Database Engine	SQL_Latin1_General_CP1_CI_AS	2 LUNs
App Controller	SCACDB	Database Engine	SQL_Latin1_General_CP1_CI_AS	2 LUNs
Windows Server Update Services (optional)	SCWSUSDB	Database Engine	SQL_Latin1_General_CP1_CI_AS	2 LUNs

Figure 26 shows Microsoft's recommended placement of the various databases and log files.

Figure 26. Suggested Storage Placement of Databases and Log Files

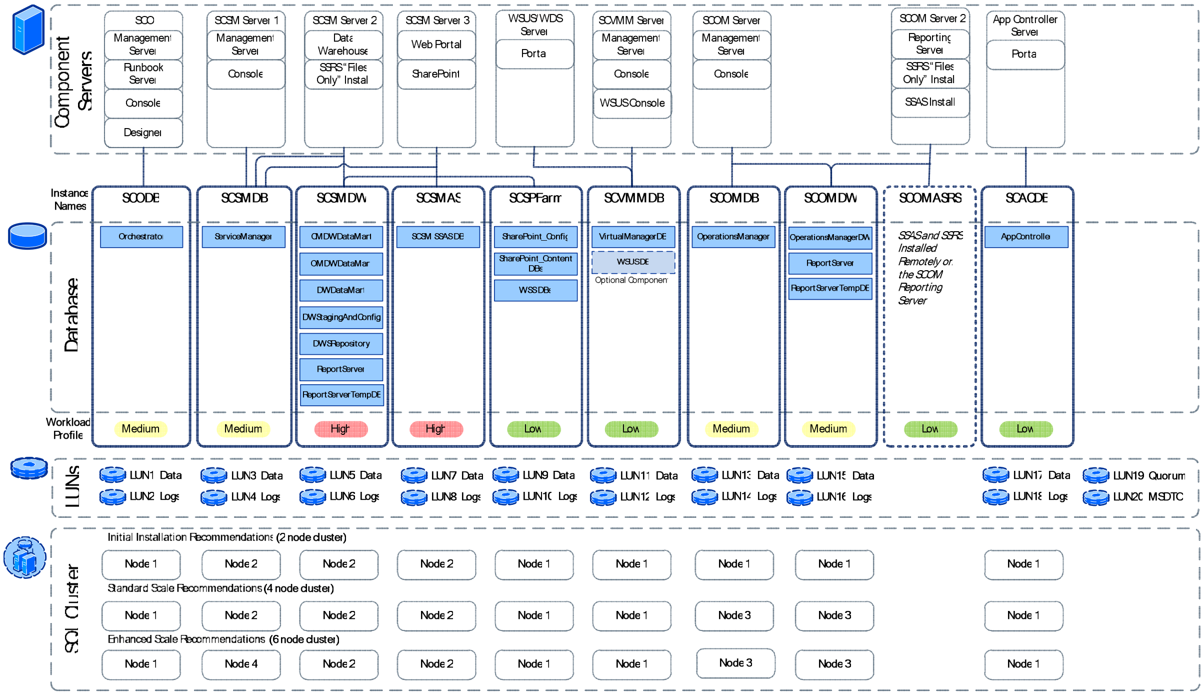

[4] The default SQL collation settings are not supported for multilingual installations of the Service Manager component. Only use the default SQL collation if multiple languages are not required. Note that the same collation must be used for all Service Manager databases (Management, Data Warehouse, and Reporting Services).

Virtual Machine Manager

System Center Virtual Machine Manager 2012 is required. One library share on the VMM server is utilized. Additional library servers can be added as needed (for instance, one per physical location). VMM and Operations Manager integration is configured during the installation process. The following hardware configuration is used for the VMM virtual machine:

- 1 HA VMs, guest clustered
- Windows Server 2008 R2 SP1 Enterprise or Datacenter
- 4 vCPU
- 8 GB Memory
- 1 vNIC
- Storage: 1 OS VHD, 1 x Data VHD

Operations Manager

System Center Operations Manager 2012 is required. One OM server is deployed in a single management group using a dedicated SQL instance on the virtualized SQL cluster. An OM agent gets installed on every guest VM as well as every management host and scale unit cluster node to support health monitoring functionality. Note that Operations Manager Gateway servers and additional management servers are supported for custom solutions; however, for the base reference implementation these additional roles are not implemented.

The Operations Manager installation uses a dedicated SQL instance on the virtualized SQL cluster. The installation follows a “split SQL configuration”: SQL Server Reporting Services (SSRS) and OpsMgr components reside on the OpsMgr VM, while the SSRS and OpsMgr databases utilize a dedicated instance on the virtualized SQL cluster.

- The following hardware configuration is used for the OM virtual machine: 1HA VM, guest clustered
- Windows Server 2008 R2 SP1 Enterprise or Datacenter
- 4 vCPU
- 16-GB memory
- 1 vNIC
- Storage: 1 OS VHD

The following Operations Manager Management Packs are required:

- Virtual Machine Manager 2012
- Windows Server Base Operating System
- Windows Server Failover Clustering
- Windows Server 2008 Hyper-V
- Microsoft SQL Server Management Pack
- Microsoft Windows Server Internet Information Services (IIS) 2000/2003/2008
- System Center MPs
- Server OEM third-party MPs

System Center Operations Manager is enabled with management packs for management of Microsoft infrastructure and products, and is extended with Cisco and EMC management packs to support hardware and software implemented within the solution.

Service Manager

System Center Service Manager 2012 is optional. The Service Manager Management server is installed on two virtual machines. A third virtual machine hosts the Service Manager Data Warehouse server. Both the Service Manager database and the Data Warehouse database, use a dedicated SQL instance on the virtualized SQL cluster. The Service Manager Portal is hosted on a fourth VM with the portal. The following VM configurations are used.

- The following hardware configuration is used for the SM virtual machine: 1 HA VM
- Windows Server 2008 R2 SP1 Enterprise or Datacenter
- 4 vCPU
- 16-GB memory
- 1 vNIC
- Storage: 1 OS VHD

The following hardware configuration is used for the SM Data Warehouse virtual machine:

- 1 HA VM
- Windows Server 2008 R2 SP1 Enterprise or Datacenter
- 4 vCPU
- 16-GB memory
- 1 vNIC
- Storage: 1 OS VHD

The following hardware configuration is used for the SM Portal Server virtual machine:

- 1 HA VM
- Windows Server 2008 R2 SP1 Enterprise or Datacenter
- 4 vCPU
- 8-GB Memory
- 1 vNICs
- Storage: 1 OS VHD

Orchestrator

The Orchestrator installation uses a dedicated SQL instance on the virtualized SQL cluster.

We use a single Orchestrator Runbook server. For HA and scale purposes additional Ocrchestrator Runbook servers can be deployed. Orchestrator provides built-in failover capability (it does not use failover clustering). By default, if an Orchestrator server fails, any workflows that were running on that server will be started (not restarted) on the other Orchestrator server. The difference between starting and restarting is that restarting implies saving or maintaining the state and enabling an instance of a workflow to keep running. Orchestrator only guarantees that it will start any workflows that were started on the failed server. The state may (likely will) be lost, meaning a request may fail. Most workflows have some degree of state management built in which helps mitigate this risk.

The other reason two Orchestrator servers may be deployed is for scalability. By default, each Orchestrator Runbook server can run a maximum of 50 simultaneous workflows. This limit can be increased depending on server resources, but an additional server is needed to accommodate larger scale environments.

The following hardware configuration is used for the SM Portal Server virtual machine:

- 1 HA VM
- Windows Server 2008 R2 SP1 Enterprise or Datacenter
- 4 vCPU
- 8-GB memory
- 1 vNIC
- Storage: 1 OS VHD

App Controller

System Center App Controller is optional; however, if the Service Manager Portal is utilized, App Controller must also be installer. App Controller uses a dedicated SQL instance on the virtualized SQL cluster. A single App Controller server is installed on the host cluster.

Service Manager provides the Service Catalog and Service Request mechanism, Orchestrator provides the automated provisioning, and App Controller provides the end-user interface for connecting to and managing workloads post-provisioning.

The following hardware configuration is used for the SM Portal Server virtual machine:

- 1 HA VM
- Windows Server 2008 R2 SP1 Enterprise or Datacenter
- 4 vCPU
- 8-GB Memory
- 1 vNIC
- Storage: 1 OS VHD

Management Scenarios Architecture

The primary management scenarios addressed in Fast Track are as follows, although the management layer can provide many more capabilities.

- Fabric management
- Fabric provisioning
- VM provisioning and deprovisioning
- IT service provisioning (including platform and application provisioning)
- Resource optimization
- Fabric and IT service maintenance
- Fabric and IT service monitoring
- Reporting (used by chargeback, capacity, service management, health, performance) Service management
- User Self-Service

- Backup and disaster recovery
- Security

Fabric Management

Fabric management is the act of pooling multiple disparate computing resources together and being able to subdivide, allocate, and manage them as a single fabric. Hardware integration and the various methods next make this possible.

Storage Integration

In Virtual Machine Manager, you can discover, classify, and provision remote storage on supported storage arrays through the VMM console. VMM fully automates the assignment of storage to a Hyper-V host or Hyper-V host cluster, and tracks the storage that is managed by VMM.

To enable the new storage features, VMM uses the new Microsoft Storage Management Service to communicate with external arrays through an Storage Management Initiative Specification (SMI-S) provider. The Storage Management Service is installed by default during the installation of VMM. You must install a supported SMI-S provider on an available server, and then add the provider to VMM management.

The Cisco and EMC design implements a combination of EMC Storage Integrator (ESI) PowerShell and SMI-S functionality in compliance of this requirement. ESI/PowerShell is utilized for initial configuration requirements. SMI-S support for the SCVMM 2012 installation is provided to allow for ongoing storage management.

Network Integration

Networking in Virtual Machine Manager includes several enhancements that enable administrators to efficiently provision network resources for a virtualized environment.

The networking enhancements include the ability to create and define logical networks. A logical network, together with one or more associated network sites, is a user-defined named grouping of IP subnets, VLANs, or IP subnet/VLAN pairs that is used to organize and simplify network assignments. Some possible examples include BACKEND, FRONTEND, LAB, MANAGEMENT and BACKUP. Logical networks represent an abstraction of the underlying physical network infrastructure which enables you to model the network based on business needs and connectivity properties. After a logical network is created, it can be used to specify the network on which a host or a virtual machine (standalone or part of a service) is deployed. Users can assign logical networks as part of virtual machine and service creation without having to understand the network details.

Static IP Address and MAC Address Pool Assignment

If you associate one or more IP subnets with a network site, you can create static IP address pools from those subnets. Static IP address pools enable VMM to automatically allocate static IP addresses to Windows-based virtual machines that are running on any managed Hyper-V, VMware ESX, or Citrix XenServer host. VMM can automatically assign static IP addresses from the pool to standalone virtual machines, to virtual machines that are deployed as part of a service, and to physical computers when you use VMM to deploy them as Hyper-V hosts. Additionally, when you create a static IP address pool, you can define a reserved range of IP addresses for the virtual IP (VIP) addresses of load balancers. VMM automatically assigns a virtual IP address to a load balancer during the deployment of a load-balanced service tier.

Load Balancer Integration

You can discover and add hardware load balancers to System Center 2012 Virtual Machine Manager (VMM). By adding load balancers to VMM management and by creating associated virtual IP templates (VIP templates), users who create services can automatically provision load balancers when they create and deploy a service.

Fabric Provisioning

In accordance with the principle of standardization and automation, creating the fabric and adding capacity should always be an automated process. In Virtual Machine Manager, this is achieved through a multistep process.

1. Provisioning Hyper-V hosts
2. Configuring host properties, networking, and storage
3. Creating Hyper-V host clusters

Each step in this process has dependencies:

1. Provisioning Hyper-V hosts
    a. A PXE boot server
    b. Dynamic DNS registration
    c. A standard base image to be used for Hyper-V Hosts
    d. Hardware Driver files in the VMM Library
    e. A Host Profile in the VMM Library
    f. Baseboard Management Controller (BMC) on the physical server
2. Configuring host properties, networking, and storage
    a. Host property settings
    b. Storage integration from above plus the addition MPIO and/or iSCSI configuration
    c. Network: You must have already configured the logical networks that you want to associate with the physical network adapter. If the logical network has associated network sites, one or more of the network sites must be scoped to the host group where the host resides.
3. Creating Hyper-V hHost clusters
    a. The hosts must meet all requirements for Windows Server Failover Clustering.
    b. The hosts must be managed by VMM.

The Cisco and EMC design implements a solution based on System Center Virtual Machine Manager 2012 and its integration with the EMC SMI-S provider. The actual implementation makes use of both Cisco's and EMC's respective PowerShell modules, which can be executed from the VMM library.

VMM Private Clouds

Once you have configured the fabric resources (such as storage, networking, library servers and shares, host groups, and hosts), you can sub-divide and allocate them for self-service consumption via the creation of VMM private clouds. During private cloud creation, you select the underlying fabric resources that will be available in the private cloud, configure library paths for private cloud users, and set the capacity for the private cloud. For example, you may want to create a cloud for use by the Finance Department. You will be able to:

- Name the cloud—for example, Finance.
- Scope it to one or more host groups.

- Select which logical networks, load balancers, and VIP templates are available to the cloud.
- Specify which storage classifications are available to the cloud.
- Select which library shares are available to the cloud for VM storage.
- Specify granular capacity limits to the cloud (for example, virtual CPUs, memory, and so on).
- Select which capability profiles are available to the cloud. (Capability profiles match the type of hypervisor platforms that are running in the selected host groups. The built-in capability profiles represent the minimum and maximum values that can be configured for a virtual machine for each supported hypervisor platform.)

VM Provisioning and Deprovisioning

One of the primary cloud attributes is user self-service—that is, giving the consumer of a service the ability to request that service and have it be automatically provisioned for them. In the Microsoft private cloud solution, this refers to the ability for the user to request one or more virtual machines or to delete one or more of their existing virtual machines.

The infrastructure scenario supporting this capability is the VM provisioning and deprovisioning process. This process is initiated from the self-service portal or tenant user interface and triggers and automated process or workflow in the infrastructure through System Center Virtual Machine Manager to either create or delete a virtual machine based on the authorized settings input by the user or tenant. Provisioning can be template-based, such as requesting a small, medium, or large VM template, or a series of selections can be made by the user (vCPUs, RAM, and so on). If authorized, the provisioning process should create a new VM per the user’s request, add the VM to any relevant management products in the Microsoft private cloud (such as System Center), and enable access to the VM by the requestor.

IT Service Provisioning

In VMM, a service is a set of virtual machines that are configured and deployed together and are managed as a single entity—for example, a deployment of a multitier line-of-business application.

In the VMM console, you use the Service Template Designer to create a service template, which defines the configuration of the service. The service template includes information about the virtual machines that are deployed as part of the service, which applications to install on the virtual machines, and the networking configuration needed for the service (including the use of a load balancer).

Resource Optimization

Elasticity, perception of infinite capacity, and perception of continuous availability are Microsoft private cloud architecture principles that relate to resource optimization. This management scenario deals with optimizing resources by dynamically moving workloads around the infrastructure based on performance, capacity, and availability metrics. Examples include the option to distribute workloads across the infrastructure for maximum performance or consolidating as many workloads as possible to the smallest number of hosts for a higher consolidation ratio.

VMM Dynamic Optimization migrates virtual machines to perform resource balancing within host clusters that support Live Migration, according to settings you enter.

Dynamic Optimization attempts to correct three possible scenarios, in priority order:

1. VMs that have configuration problems on their current host

2. VMs that are causing their host to exceed configured performance thresholds
3. Unbalanced resource consumption on hosts

VMM Power Optimization is an optional feature of Dynamic Optimization; it is available only when a host group is configured to migrate virtual machines through Dynamic Optimization. Through Power Optimization, VMM helps to save energy by turning off hosts that are not needed to meet resource requirements within a host cluster, and turns the hosts back on when they are needed again.

By default, VMM performs Power Optimization all of the time when the feature is turned on. However, you can schedule the hours and days during the week when Power Optimization is performed. For example, you might initially schedule power optimization only on weekends, when you anticipate low resource usage on your hosts. After observing the effects of Power Optimization in your environment, you might increase the hours.

For Power Optimization, the computers must have a baseboard management controller (BMC) that enables out-of-band management.

Fabric and IT Service Maintenance

The Microsoft private cloud must make it possible to perform maintenance on any component of the solution without impacting the availability of the solution. Examples include the need to update or patch a host server, add additional storage to the SAN, and so on. During maintenance, the system should ensure that unnecessary alerts or events are not generated in the management systems during planned maintenance.

VMM 2012 includes the built-in ability to maintain the fabric servers in a controlled, orchestrated manner.

Fabric servers include the following physical computers managed by VMM: Hyper-V hosts and Hyper-V clusters, library servers, preboot execution environment (PXE) servers, the Windows Server Update Management (WSUS) server, and the VMM management server.

VMM supports on-demand compliance scanning and remediation of the fabric. Administrators can monitor the update status of the servers. They can scan for compliance and remediate updates for selected servers. Administrators also can exempt resources from installation of an update.

VMM supports orchestrated updates of Hyper-V host clusters. When a VMM administrator performs update remediation on a host cluster, VMM places one cluster node at a time in maintenance mode and then installs updates. If the cluster supports live migration, intelligent placement is used to migrate virtual machines off the cluster node. If the cluster does not support live migration, VMM saves the state for the virtual machines.

The feature requires the use of a Windows Server Update Management (WSUS) server.

Fabric and IT Service Monitoring

The Microsoft private cloud must make it possible to monitor every major component of the solution and generate alerts based on performance, capacity, and availability metrics. Examples include monitoring server availability, CPU, and storage utilization.

Monitoring of the fabric is performed via the integration of Operations Manager and Virtual Machine Manager. Enabling this integration allows Operations Manager to automatically discover, monitor, and report on essential performance and health characteristics of any object managed by VMM, including:

- Health and performance of all VMM-managed Hosts & VMs

- Diagram views in Operations Manager reflecting all VMM deployed hosts, services, VMs, private clouds, IP address pools, storage pools, and more
- Performance and resource optimization (PRO), which can now be configured at a very granular level and delegated to specific self-service users
- Monitoring and automated remediation of physical servers, storage, and network devices

Note: For additional in-guest workload and application-specific monitoring, simply deploy an Operations Manger agent within the VM operating system and enable the desired Management Pack. Although this is not considered “fabric” monitoring, is important to be aware of.

Cisco and EMC both provide additional Management Packs to support the design infrastructure. This includes Management Packs for EMC VNX storage, EMC S/W stack as appropriate, Cisco Nexus switches, and Cisco UCS infrastructure.

Reporting

The cloud solution must provide a centralized reporting capability. The reporting capability should provide standard reports detailing capacity, utilization, and other system metrics. The reporting functionality serves as the foundation for capacity or utilization-based billing and chargeback to tenants.

In a service-oriented IT model, reporting serves the following purposes:

- Systems performance and health
- Capacity metering and planning
- Service level availability
- Usage-based metering and chargeback
- Incident and problem reports which help IT focus efforts

As a result of Virtual Machine Manager and Operations Manager integration, several reports are created and available by default. However, for metering and chargeback reports and incident and problem reports, you must use Service Manager and the Cloud Services Process Pack.

Service Management

The goal of Service Manager 2012 is to support IT service management in a broad sense. This includes implementing Information Technology Infrastructure Library (ITIL) processes, such as change management and incident management, and it can also include processes for other things, such as allocating resources from a private cloud.

Service Manager 2012 maintains a configuration management database (CMDB). The CMDB is the repository for nearly all configuration and management-related information in the System Center 2012 environment. With the System Center Cloud Services Process Pack, this information includes Virtual Machine Manager (VMM) 2012 resources like virtual machine templates, virtual machine service templates, and so on, which are copied regularly from the VMM 2012 library into the CMDB.

This allows objects such as VMs and users to be tied to Orchestrator runbooks for automated request fulfillment, metering, chargeback, and more.

Tailoring the Service Management component to meet customer needs is an additional effort.

User Self-Service

The Microsoft User Self-Service solution consists of three elements:

- Service Manager Self-Service Portal
- Cloud Services Process Pack
- App Controller

Service Manager 2012 provides its own Self-Service Portal. Using the information in the CMDB, Service Manager 2012 can create a service catalog that shows the services available to a particular user.

For example, suppose a user wants to create a virtual machine in the group's cloud. Instead of passing the request directly on to VMM 2012 as System Center App Controller 2012 does, Service Manager 2012 starts a workflow to handle the request. The workflow contacts the user's manager to get an approval for this request. If the request is approved, the workflow then starts a System Center Orchestrator 2012 Runbook.

The Service Manager Self-Service Portal consists of two parts and has the prerequisite of a Service Manager Server and database:

- Web content server
- SharePoint Web Part

These roles must be colocated on a single, dedicated server.

Cloud Services Process Pack is an add-on component that enables IaaS capabilities through the Service Manager Self-Service Portal and Orchestrator Runbooks. It provides:

- Standardized and well-defined processes for requesting and managing cloud services, including the ability to define projects, capacity pools, and virtual machines.
- Natively supported request, approval, and notification to enable businesses to effectively manage their own allocated infrastructure capacity pools.

After a request has been fulfilled, App Controller is the portal a self-service user would use in order to connect to and manage their virtual machines and services. App Controller connects directly to Virtual Machine Manager utilizing the credentials of the authenticated user to display his or her VMs, and services, and to provide a configurable set of actions.

The Cisco and EMC design implements the System Center management stack, supporting the user self-service functionality through VMM 2012. System Center Orchestrator 2012 and System Center Service Manager 2012 are implemented within the design. Customizing these components to fit the customer requirements is an additional effort.

Backup and Disaster Recovery

In a virtualized data center, there are three commonly used backup types: host-based, guest-based, and SAN-snapshot-based. Table 6 shows a comparison of these different types of backup.

Table 6. Types of Backup for Virtualized Data Centers

Capability	Host-Based	Guest- Based	SAN Snapshot
Protection of VM configuration	X		X*
Protection of Host & Cluster configuration	X		X*
Protection of virtualization-specific data such as VM snapshots	X		X
Protection of data inside the VM	X	X	X
Protection of data inside the VM stored on pass-through disks		X	X
Support for VSS-based backups for supported operating systems and applications	X	X	X*
Support for continuous data protection	X	X	
Ability to granularly recover specific files or applications inside the VM		X	

*Depends on storage vendor's level of Hyper-V Integration

Security

The three pillars of IT security are confidentiality, integrity, and availability (CIA).

IT infrastructure threat modeling is the practice of considering what attacks might be attempted against the different components in an IT infrastructure. Generally, threat modeling assumes the following conditions:

- Organizations have resources (in this case, IT components) that they wish to protect.
- All resources are likely to exhibit some vulnerabilities.
- People might exploit these vulnerabilities to cause damage or gain unauthorized access to information.
- Properly applied security countermeasures help mitigate threats that exist because of vulnerabilities.

The IT infrastructure threat modeling process is a systematic analysis of IT components that compiles component information into profiles. The goal of the process is to develop a threat model portfolio, which is a collection of component profiles.

One way to establish these pillars as a basis for threat modeling IT infrastructure is through Microsoft Operations Framework (MOF) 4.0, which provides practical guidance for managing IT practices and activities throughout the entire IT lifecycle.

The Reliability Service Management Function (SMF) in the Plan Phase of MOF addresses creating plans for confidentiality, integrity, availability, continuity, and capacity. The Policy SMF in the Plan Phase provides context to help understand the reasons for policies, their creation, validation, and enforcement, and includes processes to communicate policy, incorporate feedback, and help IT maintain compliance with directives. The Deliver Phase contains several SMFs that help ensure that project planning, solution building, and the final release of the solution are accomplished in ways that fulfill requirements and create a solution that is fully supportable and maintainable when operating in production.

Figure 28 shows a decision tree covering various questions that must be addressed to assign a threat priority to a potential security issue.

Figure 27. Threat Priority Decision Tree

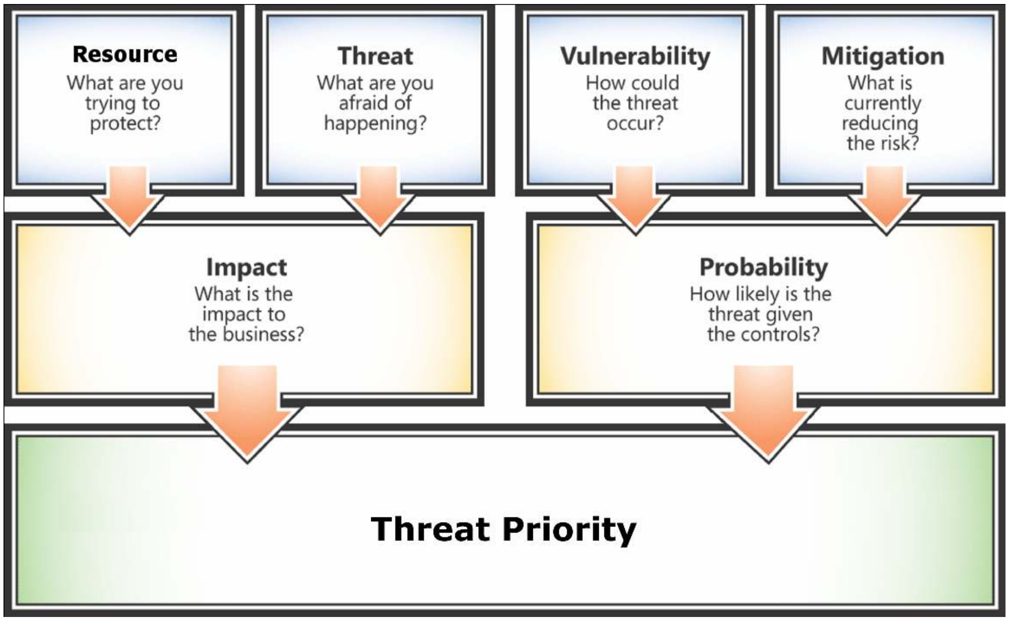
Resource
What are you trying to protect?
Threat
What are you afraid of happening?
Vulnerability
How could the threat occur?
Mitigation
What is currently reducing the risk?
Impact
What is the impact to the business?
Probability
How likely is the threat given the controls?
Threat Priority

For more information, refer to the IT Infrastructure Threat Modeling Guide.

Security Risk Management Guide

Security for the Microsoft private cloud is founded on three pillars: protected infrastructure, application access, and network access. For more information, refer to the Security Risk Management Guide.

Protected Infrastructure

A defense-in-depth strategy is utilized at each layer of the Microsoft private cloud architecture. Security technologies and controls must be implemented in a coordinated fashion.

An entry point represents data or process flow that traverses a trust boundary. Any portions of an IT infrastructure in which data or processes traverse from a less-trusted zone into a more-trusted zone should have a higher review priority.

Users, processes, and IT components all operate at specific trust levels that vary between fully trusted and fully untrusted. Typically, parity exists between the level of trust assigned to a user, process, or IT component and the level of trust associated with the zone in which the user, process, or component resides.

Malicious software poses numerous threats to organizations, from intercepting a user's login credentials with a keystroke logger to achieving complete control over a computer or an entire network by using a rootkit. Malicious software can cause websites to become inaccessible, destroy or corrupt data, and reformat hard disks. Effects can include additional costs such as to disinfect computers, restore files, reenter or recreate lost data. Virus attacks can also cause project teams to miss deadlines, leading to breach of contract or loss of customer confidence. Organizations that are subject to regulatory compliance can be prosecuted and fined.

A defense-in-depth strategy, with overlapping layers of security, is the best way to counter these threats, and the least-privileged user account (LUA) approach is an important part of that defensive strategy. The LUA approach ensures that users follow the principle of least privilege and always log in with limited user accounts. This strategy also aims to limit the use of administrative credentials to administrators, and then only for administrative tasks.

Cisco and EMC support the standard Microsoft security capabilities built into the Windows operating system and the System Center products. Configuration of security to meet customer requirements is an additional effort.

Application Access

Active Directory provides the means to manage the identities and relationships that make up the Microsoft private cloud. Integrated with Windows Server 2008 R2, Active Directory provides out-of-the-box functionality needed to centrally configure and administer system, user, and application settings.

Windows Identity Foundation enables .NET developers to externalize identity logic from their applications, improving developer productivity, enhancing application security, and enabling interoperability. Enjoy greater productivity, applying the same tools and programming model to build on-premises software as well as cloud services. Create more secure applications by reducing custom implementations and using a single simplified identity model based on claims.

The Cisco and EMC design includes support for Windows Active Directory services. Implementation of claims-based access is dependent upon customer input and is an additional effort.

Network Access

Windows Firewall with Advanced Security combines a host firewall and Internet Protocol security (IPsec). Unlike a perimeter firewall, Windows Firewall with Advanced Security runs on each computer running this version of Windows and provides local protection from network attacks that might pass through your perimeter network or originate inside your organization. It also provides computer-to-computer connection security by allowing you to require authentication and data protection for communications.

Network Access Protection (NAP) is a platform that allows network administrators to define specific levels of network access based on a client's identity, the groups to which the client belongs, and the degree to which the client complies with corporate governance policy. If a client is not compliant, NAP provides a mechanism for automatically bringing the client into compliance (a process known as remediation) and then dynamically increasing its level of network access. NAP is supported by Windows Server 2008 R2, Windows Server 2008, Windows 7, Windows Vista®, and Windows® XP with Service Pack 3 (SP3). NAP includes an application programming interface that developers and vendors can use to integrate their products and use this health state validation, access enforcement, and ongoing compliance evaluation.

You can logically isolate server and domain resources to limit access to authenticated and authorized computers. You can create a logical network inside an existing physical network, where computers share a common set of requirements for secure communications. In order to establish connectivity, each computer in the logically isolated network must provide authentication credentials to other computers in the isolated network to prevent unauthorized computers and programs from gaining access to resources inappropriately. Requests from computers that are not part of the isolated network will be ignored.

End-point Protection (Antivirus and Antimalware)

Cisco and EMC support the running of various antivirus and antimalware products. Customer usage and implementation is an additional effort.

System Center Endpoint Protection

Desktop management and security have traditionally existed as two separate disciplines, yet both play central roles in keeping users safe and productive. Management ensures proper system configuration, deploys patches against vulnerabilities, and delivers necessary security updates. Security provides critical threat detection, incident response, and remediation of system infection.

System Center 2012 Endpoint Protection (formerly known as Forefront Endpoint Protection 2012) aligns these two work streams into a single infrastructure. System Center 2012 Endpoint Protection makes it easier to protect critical desktop and server operating systems against viruses, spyware, rootkits, and other threats. Key features include:

- **Single console for endpoint management and security:** Configuration Manager provides a single interface for managing and securing desktops that reduces complexity and improves troubleshooting and reporting insights.
- **Central policy creation:** Administrators have a central location for creating and applying all client-related policies.
- **Enterprise scalability:** Use of the Configuration Manager infrastructure in System Center 2012 Endpoint Protection makes it possible to efficiently deploy clients and policies in the largest organizations around the globe. By using Configuration Manager distribution points and an automatic software deployment model, organizations can quickly deploy updates without relying on Windows Server Update Service (WSUS).
- **Highly accurate and efficient threat detection:** The antimalware engine in System Center 2012 Endpoint Protection protects against the latest malware and rootkits with a low false-positive rate, and keeps employees productive with scanning that has a low impact on performance.
- **Behavioral threat detection:** System Center 2012 Endpoint Protection uses system behavior and file reputation data to identify and block attacks on client systems from previously unknown threats. Detection methods include behavior monitoring, the cloud-based Dynamic Signature Service, and dynamic translation.
- **Vulnerability shielding:** System Center 2012 Endpoint Protection blocks exploitation of endpoint vulnerabilities with deep protocol analysis of network traffic.
- **Automated agent replacement:** System Center 2012 Endpoint Protection automatically detects and removes the most common endpoint security agents, dramatically lowering the time and effort needed to deploy new protection.
- **Windows Firewall management:** System Center 2012 Endpoint Protection ensures that Windows Firewall is active and working properly to protect against network-layer threats. It also enables administrators to more easily manage these protections across the enterprise.

Storage Management

Cisco and EMC provide a range of services and tools to manage storage provisioning operations.

EMC Unisphere

EMC Unisphere provides a flexible, integrated experience for managing existing EMC CLARiiON storage systems, existing EMC Celerra storage systems, and EMC’s VNX storage systems. With Unisphere, customers can manage file and block data with one simple interface.

The current version of Unisphere provides a cohesive user interface to manage all file and block functionalities. Figure 29 shows the dashboard view provided by Unisphere, representing a unified view of storage assets and utilization.

Figure 28. EMC Unisphere Dashboard

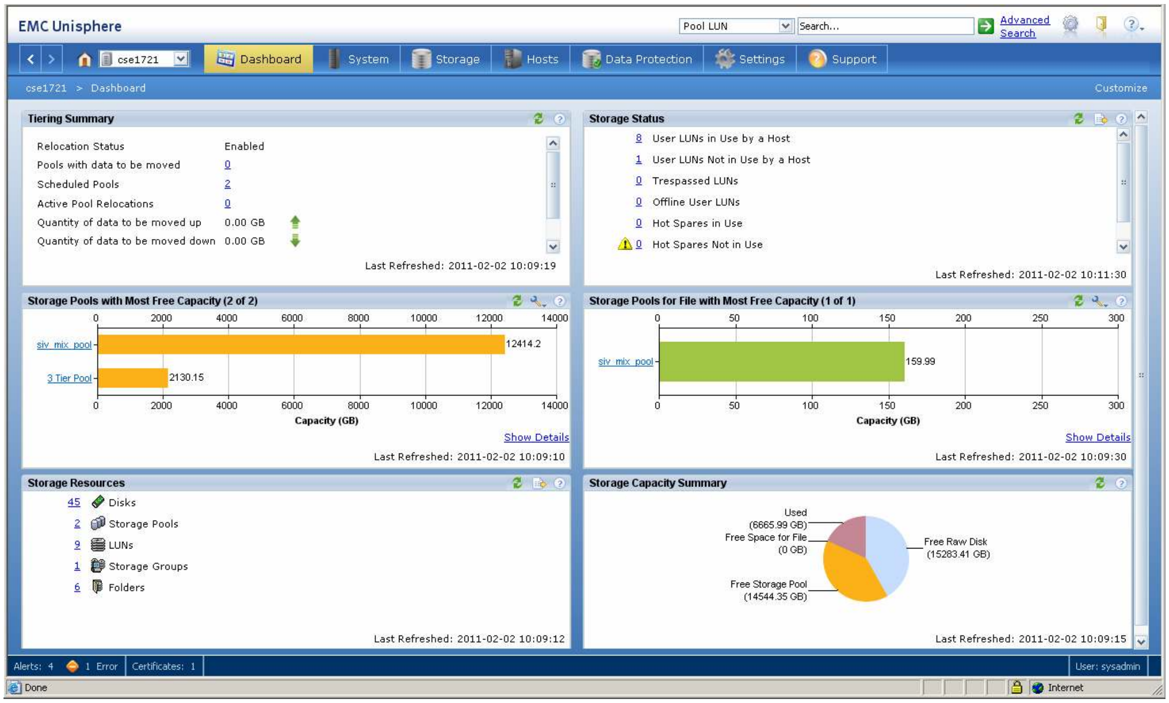

Tiering Summary

Relocation Status	Enabled
Pools with data to be moved	0
Scheduled Pools	2
Active Pool Relocations	0
Quantity of data to be moved up	0.00 GB
Quantity of data to be moved down	0.00 GB

Last Refreshed: 2011-02-02 10:09:19

Storage Status

8	User LUNs in Use by a Host
1	User LUNs Not in Use by a Host
0	Trespassed LUNs
0	Offline User LUNs
0	Hot Spares in Use
0	Hot Spares Not in Use

Last Refreshed: 2011-02-02 10:11:30

Last Refreshed: 2011-02-02 10:09:10

Last Refreshed: 2011-02-02 10:09:30

Storage Resources

45	Disks
2	Storage Pools
9	LUNs
1	Storage Groups
6	Folders

Last Refreshed: 2011-02-02 10:09:12

Storage Capacity Summary
Used (6665.99 GB)
Free Space for File (0 GB)
Free Raw Disk (15283.41 GB)
Free Storage Pool (14544.35 GB)

Last Refreshed: 2011-02-02 10:09:15

Unisphere provides simplicity, flexibility, and automation —all key requirements for optimal storage management. Unisphere’s ease of use is reflected in its intuitive task-based controls, customizable dashboards, and single-click access to "real-time" support tools and online customer communities. Unisphere’s wizards help you provision and manage your storage while automatically implementing best practices for your configuration.

EMC Storage Integrator

EMC Storage Integrator (ESI) for Windows provides capabilities for viewing and provisioning storage. As a part of the viewing capability, ESI depicts storage system configuration as shown in Figure 30, as well as Windows server to storage mapping of resources.

Figure 29. Storage System Configuration in EMC Storage Integrator

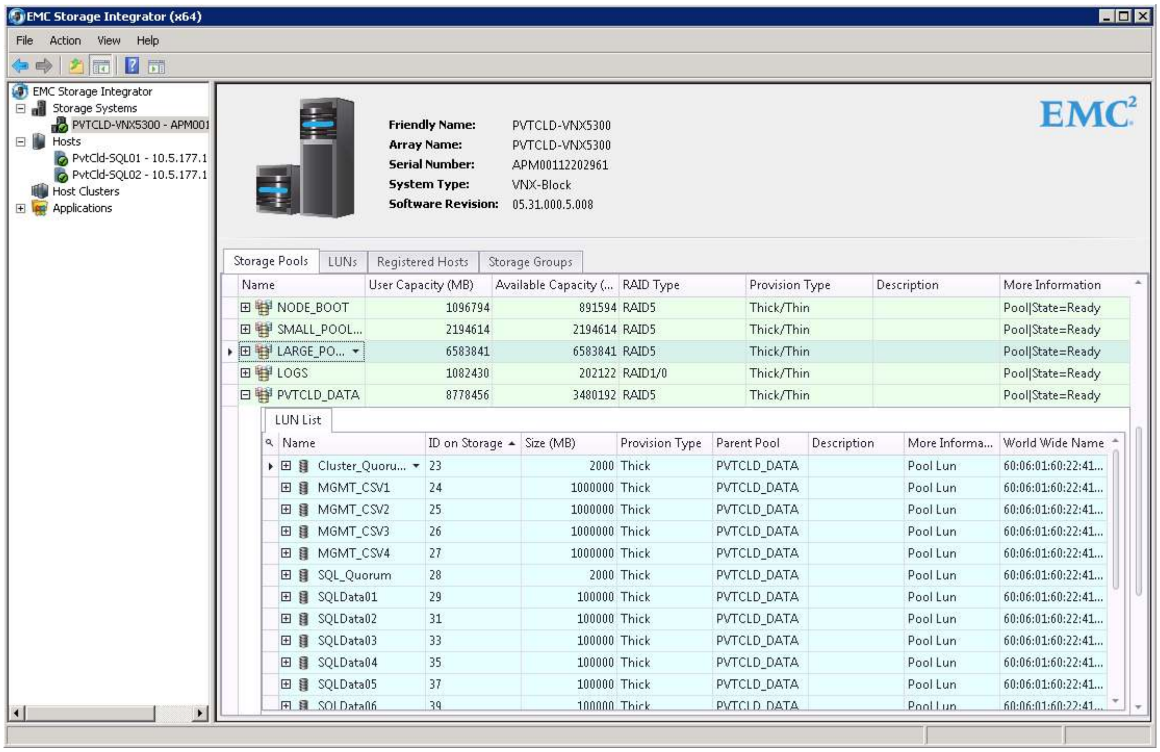

Friendly Name: PVTCLD-VNX5300
Array Name: PVTCLD-VNX5300
Serial Number: APM00112202961
System Type: VNX-Block
Software Revision: 05.31.000.5.008

Name	User Capacity (MB)	Available Capacity (...	RAID Type	Provision Type	Description	More Information
NODE_BOOT	1096794	891594	RAID5	Thick/Thin		Pool|State=Ready
SMALL_POOL...	2194614	2194614	RAID5	Thick/Thin		Pool|State=Ready
LARGE_PO...	6583841	6583841	RAID5	Thick/Thin		Pool|State=Ready
LOGS	1082430	202122	RAID1/0	Thick/Thin		Pool|State=Ready
PVTCLD_DATA	8778456	3480192	RAID5	Thick/Thin		Pool|State=Ready

Name	ID on Storage	Size (MB)	Provision Type	Parent Pool	Description	More Informa...	World Wide Name
Cluster_Quoru...	23	2000	Thick	PVTCLD_DATA		Pool Lun	60:06:01:60:22:41...
MGMT_CSV1	24	1000000	Thick	PVTCLD_DATA		Pool Lun	60:06:01:60:22:41...
MGMT_CSV2	25	1000000	Thick	PVTCLD_DATA		Pool Lun	60:06:01:60:22:41...
MGMT_CSV3	26	1000000	Thick	PVTCLD_DATA		Pool Lun	60:06:01:60:22:41...
MGMT_CSV4	27	1000000	Thick	PVTCLD_DATA		Pool Lun	60:06:01:60:22:41...
SQL_Quorum	28	2000	Thick	PVTCLD_DATA		Pool Lun	60:06:01:60:22:41...
SQLData01	29	100000	Thick	PVTCLD_DATA		Pool Lun	60:06:01:60:22:41...
SQLData02	31	100000	Thick	PVTCLD_DATA		Pool Lun	60:06:01:60:22:41...
SQLData03	33	100000	Thick	PVTCLD_DATA		Pool Lun	60:06:01:60:22:41...
SQLData04	35	100000	Thick	PVTCLD_DATA		Pool Lun	60:06:01:60:22:41...
SQLData05	37	100000	Thick	PVTCLD_DATA		Pool Lun	60:06:01:60:22:41...
SQLData06	39	100000	Thick	PVTCLD_DATA		Pool Lun	60:06:01:60:22:41...

For storage provisioning operations, ESI simplifies and automates the steps of creating a LUN, presenting the LUN to a Windows host, partitioning and finally formatting and mounting of the Windows volume. Support for Hyper-V also includes the ability to utilize the ESI PowerShell environment to provide all major storage actions in a scriptable manner.

EMC SMI-S Provider

The EMC SMI-S provider fully integrates with the System Center Virtual Machine Manager 2012 storage management framework. This level of integration allows System Center VMM to automate most storage management operations, including storage volume replication, as well as provisioning of new storage system objects for deployment into the managed environment.

This integration with System Center VMM 2012, coupled with end-user self-service portal configurations, allows customers to deploy a fully automated and scalable solution for provisioning of discrete VMs or tiered applications.

Orchestrator Integration Pack

EMC also provides an integration pack for System Center Orchestrator 2012 that provides support for creation of workflows. This implementation provides support for managing operations that may fall outside the scope of SCVMM 2012 or ESI PowerShell automation.

Due to the requirements enforced by the 32-bit implementation of System Center Orchestrator 2012, the Orchestrator 2012 Integration Kit for VNX is currently limited to customers of the Fast Track solution.

Network Management

The Cisco Nexus series of switches provides a unified management layer of Ethernet, IP, and Fibre Channel Protocol combined into a single management platform.

Consistent Management for Cisco Products

The switch platform's network features can be managed using the Cisco command-line interface (CLI), and the Fibre Channel and FCoE features can be managed through the Cisco Fabric Manager suite. Cisco Data Center Network Manager (DCNM) also supports the Cisco Nexus 5500 Switch platform. The capability to manage Ethernet and FCoE features independently with existing Cisco tools preserves existing management models, best practices, and investments in staff training. In addition, Simple Network Management Protocol (SNMP) MIBs, XML, and the Cisco CLI are made available to customers for switch management through third-party and custom-developed tools. The switch platform uses Cisco NX-OS for superior operating efficiency, pervasive security, and continuous operation even through software upgrades.

Cisco Data Center Network Manager

The Cisco Nexus 5000 is supported in Cisco DCNM. Cisco DCNM is designed for hardware platforms enabled for Cisco NX-OS, which are the Cisco Nexus Family of products. Cisco DCNM is a Cisco management solution that increases overall data center infrastructure uptime and reliability, hence improving business continuity. Focused on the management requirements of the data center network, Cisco DCNM provides a robust framework and comprehensive feature set that meets the routing, switching, and storage administration needs of present and future data centers. In particular, Cisco DCNM automates the provisioning process, proactively monitors the LAN by detecting performance degradation, secures the network, and streamlines the diagnosis of dysfunctional network elements.

Server Management Utilities

The Cisco UCS platform allows the solution’s compute resources to be managed manually or automatically.

Server Out-of-Band Management Configuration

Cisco UCS Manager provides unified, embedded management of all software and hardware components of the Cisco UCS across multiple chassis, rack-mount servers, and thousands of virtual machines. Cisco UCS Manager manages Cisco UCS as a single entity through an intuitive GUI, a command-line interface (CLI), PowerShell, or an XML API for comprehensive access to all Cisco UCS Manager functions. Included in UCS Manager is a keyboard, video, and mouse (KVM) capability to access the console of any physical host in order to provide console debugging and configuration.

The Cisco UCS Management Pack for System Center Operations Manager graphically depicts Cisco Unified Computing System hardware, service profiles, host operating systems, and virtual machines. The correlation of events with the blades and service profiles they affect simplifies identification of root causes and accelerates problem resolution.

Cisco UCS PowerTool is a flexible and powerful command line toolkit, a library of PowerShell cmdlets. This provides customers with an efficient, cost-effective, and easy-to-use interface for integrating and automating Cisco UCS management with Microsoft products and many third-party products. This is accomplished by taking advantage of the flexible and powerful scripting environment offered by Microsoft PowerShell.

Flexible, Role-Based Management

Cisco UCS Manager offers role-based management that helps organizations make more efficient use of their limited administrator resources. Server, network, and storage administrators maintain responsibility and accountability for their domain policies within an integrated management environment. Roles and privileges in the system can easily be modified and new roles quickly created.

Administrators can focus on defining policies needed to provision computing infrastructure and network connectivity. They can also collaborate on strategic architectural issues because implementation of basic server configurations is now highly accelerated and automated.

Policy-Based Provisioning of Server, Network, and Storage Access Resources

Cisco UCS Manager uses service profiles to provision and manage Cisco UCS blade servers and rack-mount servers and their I/O properties within a single management domain.

Service profiles are created by server, network, and storage administrators. Infrastructure policies needed to deploy applications are encapsulated in the service profile. The policies coordinate and automate element management at every layer of the hardware stack, including RAID levels, BIOS settings, firmware revisions and settings, adapter identities and settings, VLAN and VSAN network settings, network quality of service (QoS), and data center connectivity.

Service profile templates are used to simplify the creation of service profiles, helping ensure consistent policies within the system for a given service or application. This approach makes it just as easy to configure one server or hundreds of servers with perhaps thousands of virtual machines.

Multiple Interface Options

Cisco UCS Manager has a GUI and a CLI for use by server, network, and storage administrators. Cisco UCS Manager also provides a powerful XML API for integration with existing data center systems management tools. Some examples of additional management interfaces are Intelligent Platform Management Interface (IPMI); keyboard, video, and mouse (KVM); serial-over-LAN (SoL); and Simple Network Management Protocol (SNMP). The XML interface allows the entire system to be monitored or configured externally by higher-level systems management tools from Cisco’s many ecosystem partners.

Figure 31 shows a view of the Cisco UCS 5108 Server Chassis equipment.

Figure 30. Cisco UCS 5108 Server Chassis

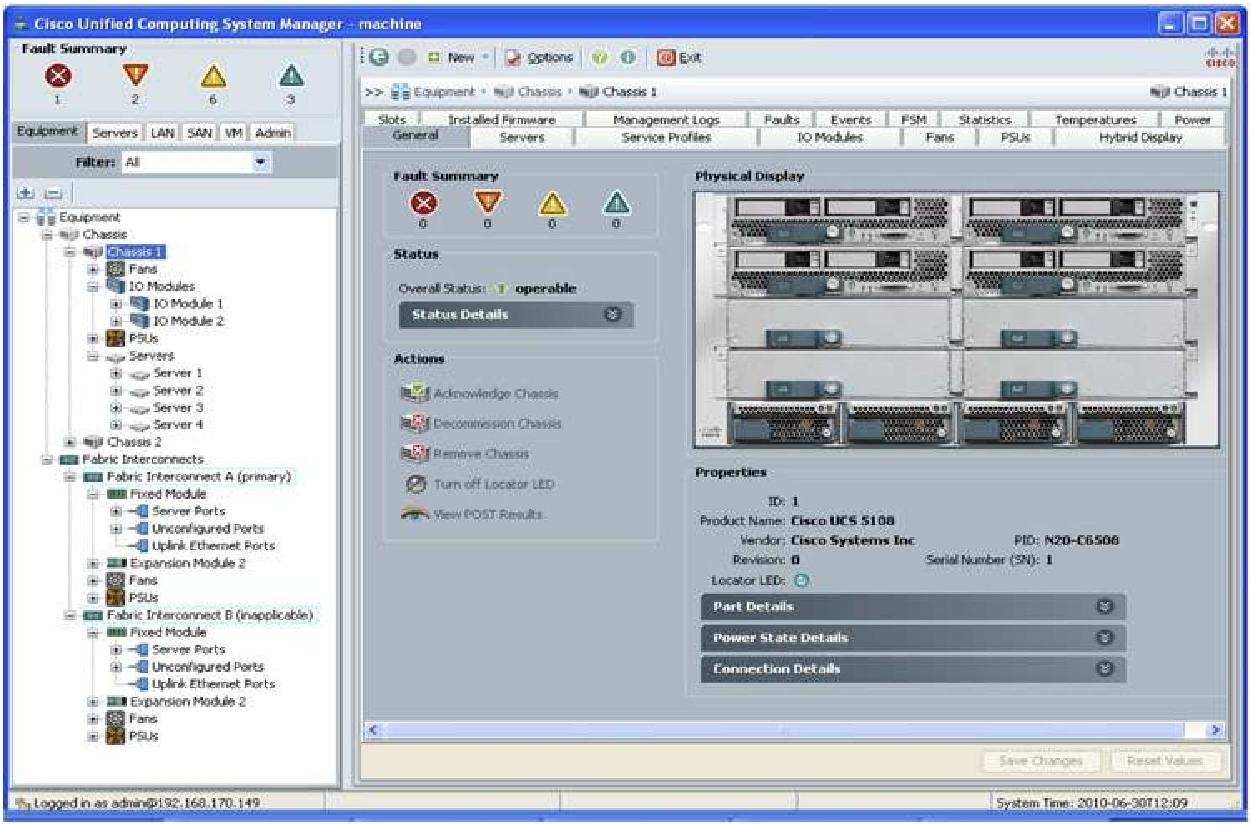

Filter: All

Overall Status: operable

ID: 1
Product Name: Cisco UCS 5108
Vendor: Cisco Systems Inc PID: N20-C6508
Revision: 0 Serial Number (SN): 1
Locator LED:

Logged in as admin@192.168.170.149
System Time: 2010-06-30T12:09

Service Management

The Service Management layer provides the means for automating and adapting IT service management best practices, such as those found in Microsoft Operations Framework and the IT Infrastructure Library (ITIL), to provide built-in processes for incident resolution, problem resolution, and change control.

Microsoft Operations Framework (MOF) 4.0 provides relevant, practical, and accessible guidance for today's IT pros. MOF strives to seamlessly blend business and IT goals while establishing and implementing reliable, cost-effective IT services. MOF is a free, downloadable framework that encompasses the entire service management lifecycle (see Figure 32). Read about MOF online.

Figure 31. Microsoft Operations Framework

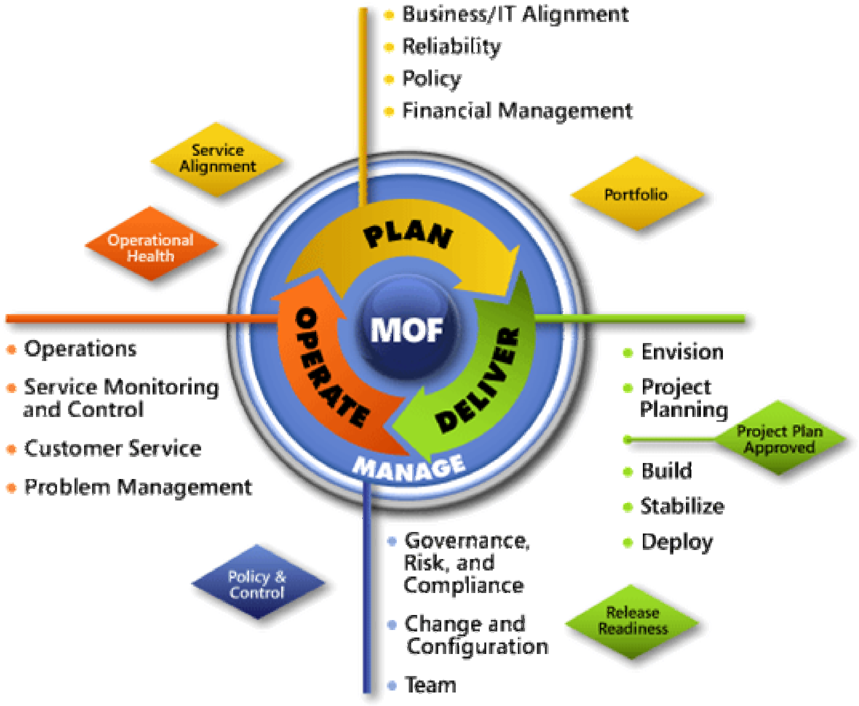

Service Delivery Layer

As the primary interface with the business, the service delivery layer is expected to know or obtain answers the following questions:

- What services does the business want?
- What level of service are business decision-makers willing to pay for?
- How can private cloud move IT from being a cost center to becoming a strategic partner to the business?

With these questions in mind, there are two main problems within the service layer that IT must address:

- How do we provide a cloud-like platform for business services that meets business objectives?
- How do we adopt an easily understood, usage-based cost model that can be used to influence business decisions?

An organization must adopt the Private Cloud Architecture Principles in order to meet the business objectives of a cloud-like service. Figure 33 shows the components of the service delivery layer.

Figure 32. Components of the Service Delivery Layer

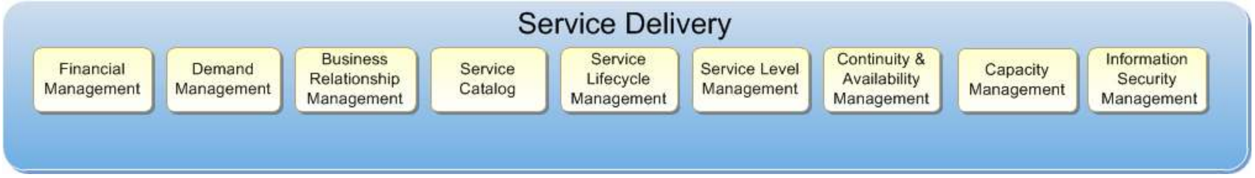
Service Delivery
Financial Management
Demand Management
Business Relationship Management
Service Catalog
Service Lifecycle Management
Service Level Management
Continuity & Availability Management
Capacity Management
Information Security Management

Financial Management

Financial Management incorporates the functions and processes used to meet a service provider's budgeting, accounting, metering, and charging requirements. The primary concerns around financial management in a private cloud are providing cost transparency to the business and structuring a usage-based cost model for the consumer. Achieving these goals is a basic precursor to achieving the principle of encouraging desired consumer behavior.

Demand Management

Demand management involves understanding and influencing customer demands for services, plus the provision of capacity to meet these demands. The principles of perceived infinite capacity and continuous availability are fundamental to stimulating customer demand for cloud-based services. A resilient, predictable environment and predictable capacity management are necessary to adhere to these principles. Cost, quality, and agility factors influence consumer demand for these services.

Business Relationship Management

Business relationship management is the strategic interface between the business and IT. If an IT department is to adhere to the principle that it must act as a service provider, mature business relationship management is critical. The business should define the functionality of required services and partner with IT on solution procurement. The business will also need to work closely with IT to define future capacity requirements to continue adhering to the principle of perceived infinite capacity.

Service Catalog

The output of demand and business relationship management will be a list of services or service classes offered and documented in the service catalog. This catalog describes each service class, eligibility requirements for each service class, service-level attributes, targets included with each service class (such as availability targets), and cost models for each service class. The catalog must be managed over time to reflect changing business needs and objectives.

Service Lifecycle Management

Service lifecycle management takes an end-to-end management view of a service. A typical journey starts with identification of a business need, through business relationship management, to the time when that service becomes available. Service strategy drives service design. After launch, the service is transitioned to operations and refined through continual service improvement. Taking a service provider's approach is critical to successful service lifecycle management.

Service-Level Management

Service-level management is the process of negotiating service-level agreements (SLAs) and making sure the agreements are met. SLAs define target levels for cost, quality, and agility by service class as well as the metrics for measuring actual performance. Managing SLAs is necessary for achieving the perception of infinite capacity and continuous availability. This, too, requires a service provider's approach by IT.

Continuity and Availability Management

Availability management defines processes necessary to achieve the perception of continuous availability. Continuity management defines how risk will be managed in a disaster scenario to make sure minimum service levels are maintained. The principles of resiliency and automation are fundamental here.

Capacity Management

Capacity management defines the processes necessary to achieve the perception of infinite capacity. Capacity must be managed to meet existing and future peak demand while controlling under-utilization. Business relationship and demand management are key inputs into effective capacity management and require a service provider's approach. Predictability and optimization of resource usage are primary principles in achieving capacity management objectives.

Information Security Management

Information security management strives to make sure that all requirements are met for confidentiality, integrity, and availability of the organization's assets, information, data, and services. An organization's particular information security policies will drive the architecture, design, and operations of a private cloud. Resource segmentation and multitenacy requirements are important factors to consider during this process.

Operations Layer

The operations layer defines the operational processes and procedures necessary to deliver IT as a Service. This layer uses IT service management concepts that can be found in prevailing best practice such as ITIL or Microsoft Operations Framework (MOF).

The main focus of the operations layer is to execute the business requirements defined at the service delivery layer. Cloud-like service attributes cannot be achieved through technology alone; mature IT service management is also required.

The operations capabilities are common to all the services: IaaS, PaaS, and SaaS. Figure 34 shows the components of the operations layer.

Figure 33. Components of the Operations Layer

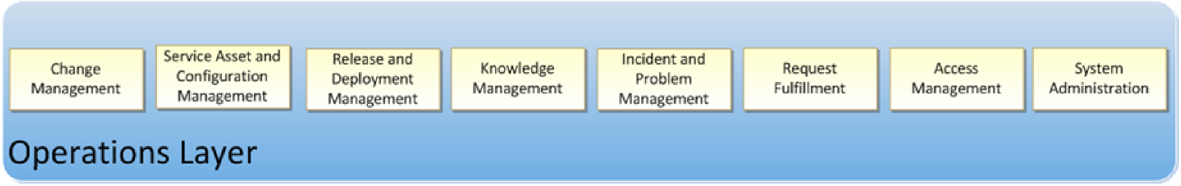

Change Management

Change management is responsible for controlling the lifecycle of all changes. Its primary objective is to implement beneficial changes with minimum disruption to the perception of continuous availability. Change management determines the cost and risk of making changes and balances them against the benefits to the business or service. Driving predictability and minimizing human involvement are the core principles behind a mature change management process.

Service Asset and Configuration Management

Service asset and configuration management maintain information on the assets, components, and infrastructure needed to provide a service. Accurate configuration data for each component, and its relationship to other components, must be captured and maintained. This data should include historical and expected future states in addition to the current state, and be easily available to those who need it. Mature service asset and configuration management processes are necessary for achieving predictability.

Release and Deployment Management

Release and deployment management is responsible for seeing that changes to a service are built, tested, and deployed with minimal disruption to the service or production environment. Change management provides the approval mechanism (determining what will be changed and why), but release and deployment management is the mechanism for determining how changes are implemented. Driving predictability and minimizing human involvement in the release and deployment process are key to achieving cost, quality, and agility goals.

Knowledge Management

Knowledge management is responsible for gathering, analyzing, storing, and sharing information within an organization. Mature knowledge management processes are necessary to achieve a service provider’s approach and a key element of IT service management.

Incident and Problem Management

The goal of incident and problem management is to resolve disruptive, or potentially disruptive, events with maximum speed and minimum disruption. Problem management also identifies root causes of past incidents and seeks to identify and prevent (or minimize the impact of) future ones. In a private cloud, the resiliency of the infrastructure helps make sure that faults, when they occur, have minimal impact on service availability. Resilient design promotes rapid restoration of service continuity. Driving predictability and minimizing human involvement are necessary to achieve this resiliency.

Request Fulfillment

The goal of request fulfillment is to manage user requests for services. As IT adopts a service provider's approach, it should define available services in a service catalog based on business functionality. The catalog should encourage desired user behavior by exposing cost, quality, and agility factors to the user. Self-service portals, when appropriate, can assist the drive towards minimal human involvement.

Access Management

The goal of access management is to deny access to unauthorized users while making sure that authorized users have access to needed services. Access management implements security policies defined by information security management at the service delivery layer. Maintaining smooth access for authorized users is critical to achieving the perception of continuous availability. Adopting a service provider's approach to access management will also make sure that resource segmentation and multitenacy are addressed.

Systems Administration

The goal of systems administration is to perform the daily, weekly, monthly, and as-needed tasks required for system health. A mature approach to systems administration is required for achieving a service provider's approach and for driving predictability. The vast majority of systems administration tasks should be automated.

Conclusion

The Cisco and EMC Microsoft Private Cloud solution provides a highly scalable and reliable platform for a variety of virtualized workloads. The goal of this program is to help you quickly deploy a private cloud environment within your enterprise without the expense or risk associated with designing and building your own custom solution.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA C11-711496-00 07/12