



FlexPod Data Center with Microsoft Fast Track Private Cloud 3

Design Guide


Last Updated: November 20, 2013



Cisco
Validated
Design



Building Architectures to Solve Business Problems



About Cisco Validated Design (CVD) Program

The CVD program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information visit

<http://www.cisco.com/go/designzone>.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0809R)

© 2013 Cisco Systems, Inc. All rights reserved

About the Authors

Mike Mankovsky, Cisco Systems

Mike Mankovsky is a Cisco Unified Computing System architect, focusing on Microsoft solutions with extensive experience in Hyper-V, storage systems, and Microsoft Exchange Server. He has expert product knowledge of Microsoft Windows storage technologies and data protection technologies.

Chris O'Brien, Technical Marketing Manager, Server Access Virtualization Business Unit, Cisco Systems

Chris O'Brien is currently focused on developing infrastructure best practices and solutions that are designed, tested, and documented to facilitate and improve customer deployments. Previously, O'Brien was an application developer and has worked in the IT industry for more than 15 years.

John George, Reference Architect, Infrastructure and Cloud Engineering, NetApp

John George is a Reference Architect in the NetApp Infrastructure and Cloud Engineering team and is focused on developing, validating, and supporting cloud infrastructure solutions that include NetApp products. Before his current role, he supported and administered Nortel's worldwide training network and VPN infrastructure. John holds a Master's degree in computer engineering from Clemson University.

Lindsey Street, Solutions Architect, Infrastructure and Cloud Engineering, NetApp

Lindsey Street is a Solutions Architect in the NetApp Infrastructure and Cloud Engineering team. She focuses on the architecture, implementation, compatibility, and security of innovative vendor technologies to develop competitive and high-performance end-to-end cloud solutions for customers. Lindsey started her career in 2006 at Nortel as an interoperability test engineer, testing customer equipment interoperability for certification. Lindsey has her Bachelors of Science degree in Computer Networking and her Master's of Science in Information Security from East Carolina University.

Chris Reno, Reference Architect, Infrastructure and Cloud Engineering, NetApp

Chris Reno is a reference architect in the NetApp Infrastructure and Cloud Enablement group and is focused on creating, validating, supporting, and evangelizing solutions based on NetApp products. Before being employed in his current role, he worked with NetApp product engineers designing and developing innovative ways to perform Q&A for NetApp products, including enablement of a large grid infrastructure using physical and virtualized compute resources. In these roles, Chris gained expertise in stateless computing, netboot architectures, and virtualization.



Corporate Headquarters:

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

Copyright © 2013 Cisco Systems, Inc. All rights reserved.

Glenn Sizemore, NetApp

Glenn Sizemore is a private cloud reference architect in the Microsoft Solutions Group at NetApp, where he specializes in cloud and automation. Since joining NetApp, Glenn has delivered a variety of Microsoft-based solutions, ranging from general best practices guidance to co-authoring the NetApp Hyper-V Cloud Fast Track with Cisco reference architecture.

FlexPod Data Center with Microsoft Fast Track Private Cloud 3 Design Guide

Goal of This Document

Cisco® Validated Designs consist of systems and solutions that are designed, tested, and documented to facilitate and improve customer deployments. These designs incorporate a wide range of technologies and products into a portfolio of solutions that have been developed to address the business needs of our customers.

The purpose of this document is to describe the Cisco and NetApp® FlexPod® solution, which is a validated approach for deploying Cisco and NetApp technologies as a shared cloud infrastructure.

The Microsoft Private Cloud Fast Track program is a joint effort between Microsoft and its hardware partners. The goal of the program is to help organizations develop and quickly implement private clouds while reducing both complexity and risk. The program provides a reference architecture that combines Microsoft software, consolidated guidance, and validated configurations with partner technology, such as computing, network, and storage architectures, in addition to 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 that are achieved through dedicated private resources. With Private Cloud Fast Track, Microsoft and its hardware partners can help provide organizations both the control and the flexibility that are required to reap the potential benefits of the private cloud.

Private Cloud Fast Track uses the core capabilities of Windows Server, Hyper-V, and System Center to deliver a private cloud infrastructure as a service offering. These key software components are used for every reference implementation.

Audience

The intended audience of this document includes sales engineers, field consultants, professional services, IT managers, partner engineering, and customers who want to take advantage of an infrastructure built to deliver IT efficiency and enable IT innovation.

Changes in FlexPod

The following design elements distinguish this version of FlexPod from previous models:

- End-to-end Fibre Channel over Ethernet (FCoE), delivering a unified Ethernet fabric
- Single-wire Cisco Unified Computing System™ (Cisco UCS®) Manager for Cisco UCS C-Series M3 Rack Servers and the Cisco UCS Virtual Interface Card (VIC) 1225. These features effectively double the server density per I/O module while reducing cabling costs.
- Cisco Nexus® 1000V Switch for Hyper-V is a distributed virtual switching platform that provides advanced networking features, integrated virtual services, and a consistent operational model across physical and virtual environments. Customers can rely on the robust Cisco NX-OS Software command-line interface (CLI) and feature set and Cisco's innovative network services architecture for their virtual environments.

- Server Message Block (SMB) 3.0 file shares provide shared storage for Hyper-V 2012. System Center 2012 SP1 allows you to assign SMB file shares to Hyper-V stand-alone hosts and clusters for storing virtual machine files.
- Virtual Fibre Channel provides access to Fibre Channel storage directly from Hyper-V virtual machines. SQL Server instances running in virtual machines use Virtual Fibre Channel to access database LUNs
- Clustered Data ONTAP is highly scalable. Additional storage controllers and disks can be easily added to existing clusters to scale capacity and performance to meet rising demands.

Deployment Guide

The implementation of this solution is documented in a separate document called FlexPod with Microsoft Private Cloud Fast Track 3.0 Enterprise Deployment Guide. This document can be found at the following link:

http://www.cisco.com/en/US/docs/unified_computing/ucs/UCS_CVDs/flexpod_mspc30_cmode.pdf


FlexPod Technology Overview

Industry trends indicate a vast data center transformation toward shared infrastructure and cloud computing. Enterprise customers are moving away from isolated centers of IT operation toward more cost-effective virtualized environments.

The objective of the move toward virtualization, and eventually to cloud computing, is to increase agility and reduce costs.

Since companies must address resistance to change in both their organizational and technical IT models, achieving this transformation can seem daunting and complex. To accelerate the process and simplify the evolution to a shared cloud infrastructure, Cisco and NetApp have developed a solution called FlexPod with Microsoft Private Cloud Fast Track v3.

FlexPod is a predesigned, best practice data center architecture that is built on Cisco Unified Computing System, the Cisco Nexus family of switches, and NetApp fabric-attached storage (FAS) or V-Series systems (Figure 1). FlexPod is a suitable platform for running a variety of virtualization hypervisors as well as bare metal operating systems and enterprise workloads. FlexPod delivers not only a baseline configuration, but also the flexibility to be sized and optimized to accommodate many different use cases and requirements.

Figure 1 *FlexPod Component Families*

This document describes FlexPod with Microsoft Private Cloud Fast Track v3 from Cisco and NetApp and discusses design choices and deployment best practices using this shared infrastructure platform.

Customer Challenges

As customers transition toward shared infrastructure or cloud computing, they face a number of questions, such as the following:

- How do I start the transition?
- What will be my return on investment?
- How do I build an infrastructure that is ready for the future?
- How do I transition from my current infrastructure cost-effectively?
- Will my applications run properly in a shared infrastructure?
- How do I manage the infrastructure?

The FlexPod architecture is designed to help you answer these questions by providing proven guidance and measurable value. By introducing standardization, FlexPod helps customers mitigate the risk and uncertainty involved in planning, designing, and implementing a new data center infrastructure. The result is a more predictive and adaptable architecture capable of meeting and exceeding customers' IT demands.

FlexPod Program Benefits

Cisco and NetApp have thoroughly tested and verified the FlexPod solution architecture and its many use cases while creating a portfolio of detailed documentation, information, and references to assist customers in transforming their data centers to this shared infrastructure model. This portfolio includes the following items:

- Best practice architectural design

- Workload sizing and scaling guidance
- Implementation and deployment instructions
- Technical specifications (rules for what is, and what is not, a FlexPod configuration)
- Frequently asked questions (FAQs)
- Cisco Validated Designs and NetApp Validated Architectures (NVAs) focused on a variety of use cases

Cisco and NetApp have also built an experienced support team focused on FlexPod solutions, from customer account and technical sales representatives to professional services and technical support engineers. The support alliance formed by NetApp and Cisco provides customers and channel services partners with direct access to technical experts who collaborate with multiple vendors and have access to shared lab resources to resolve potential issues.

FlexPod supports tight integration with virtualized and cloud infrastructures, making it the logical choice for long-term investment. The following IT initiatives are addressed by the FlexPod solution:

Integrated Systems

FlexPod is a prevalidated infrastructure that brings together computing, storage, and network to simplify and accelerate data center builds and application rollouts while reducing the risks. These integrated systems provide a standardized approach in the data center that supports staff expertise, application onboarding, and automation, as well as operational efficiencies that are important for compliance and certification.

Fabric Infrastructure Resilience

FlexPod is a highly available and scalable infrastructure that IT can evolve over time to support multiple physical and virtual application workloads. FlexPod contains no single point of failure at any level, from the server through the network to the storage. The fabric is fully redundant and scalable, providing smooth traffic failover should any individual component fail at the physical or virtual layer.

Fabric Convergence

Cisco Unified Fabric is a data center network that supports both traditional LAN traffic and all types of storage traffic, including the lossless requirements for block-level storage transport over Fibre Channel. Cisco Unified Fabric creates high-performance, low-latency, and highly available networks serving a diverse set of data center needs.

FlexPod Gen-II uses Cisco Unified Fabric to offer a wire-once environment that accelerates application deployment. Cisco Unified Fabric also offers the efficiencies associated with infrastructure consolidation, including:

- Cost savings from the reduction in switches (LAN/SAN switch ports), associated cabling, and rack space, all of which reduce capital expenditures (CapEx)
- Cost savings on power and cooling, which reduce operating expenses (OpEx)
- Migration to the faster 10 Gigabit Ethernet network and in the future, to 40 Gigabit Ethernet and 100 Gigabit Ethernet
- Evolution to a converged network with little disruption to operations: FlexPod Gen-II with Cisco Unified Fabric helps you preserve investments in existing infrastructure, management tools, and staff training and expertise
- Simplified cabling, provisioning, and network maintenance to improve productivity and operational models

Network Virtualization

FlexPod delivers the capability to securely separate and connect virtual machines into the network. Using technologies such as VLANs, quality of service (QoS), Cisco Nexus 1000V Switch for Hyper-V, and VM-FEX, this solution allows network policies and services to be uniformly applied within the integrated computing stack. This capability enables the full utilization of FlexPod while maintaining consistent application and security policy enforcement across the stack, even with workload mobility.

FlexPod provides a uniform approach to IT architecture, offering a well-characterized and documented shared pool of resources for application workloads. FlexPod delivers operational efficiency and consistency with versatility to meet a variety of service-level agreements (SLAs) and IT initiatives, including:


- Application rollouts or application migrations
- Business continuity and disaster recovery
- Desktop virtualization
- Cloud delivery models (public, private, hybrid) and service models (infrastructure as a service [IaaS], platform as a service [PaaS], and software as a service [SaaS])
- Asset consolidation and virtualization

Fast Track Program Overview

The Microsoft Private Cloud Fast Track program outlines the high-level architectural vision that is intended to help partners rapidly develop end-to-end, integrated, and tested virtualization or private cloud solutions for small- and medium-size businesses, as well as for the enterprise and data center that meet or exceed Microsoft validation standards.

The Fast Track program has three main branches, as shown in [Figure 2](#). This guide will focus exclusively on the Enterprise Solutions branch.


Figure 2 *Branches of the Microsoft Private Cloud Fast Track Program*



Fast Track Reference Architectures

Each branch in the Fast Track program uses a reference architecture that defines the requirements that are necessary to design, build, and deliver virtualization and private cloud solutions for small-, medium-, and large-size enterprise implementations. [Figure 3](#) shows examples of these Fast Track reference architectures.


Figure 3 Examples of Fast Track Reference Architectures



Each reference architecture in the Fast Track program combines concise guidance with validated configurations for the computing, network, storage, and virtualization and management layers. Each architecture presents multiple design patterns for using the architecture and describes the minimum requirements for validating each Fast Track solution. This document describes the Enterprise Medium configuration and the Continuous Availability over SMB with SMB 3.0 Enabled Storage design pattern.

The Continuous Availability over SMB with SMB 3.0 Enabled Storage design pattern in this context of Microsoft Private Cloud provides SMB 3.0 file protocol access to VHDX files that are stored on SMB 3.0 file shares that resided on 3rd party storage arrays.(Figure 4).

Figure 4 *Continuous Availability Over SMB with SMB 3.0 Enabled Storage Design Pattern*



FlexPod

This section provides an overview on the FlexPod design.

System Overview

FlexPod is a best practice data center architecture that is built with three components:

- Cisco Unified Computing System
- Cisco Nexus switches
- NetApp fabric-attached storage (FAS) systems

These components are connected and configured according to the best practices of both Cisco and NetApp to provide the ideal platform for running a variety of enterprise workloads with confidence. FlexPod can scale up for greater performance and capacity (adding computing, network, or storage resources individually as needed), or it can scale out for environments that need multiple consistent deployments (rolling out additional FlexPod stacks). FlexPod delivers not only a baseline configuration but also the flexibility to be sized and optimized to accommodate many different use cases.

Typically, the more scalable and flexible a solution is, the more difficult it becomes to maintain a single unified architecture capable of offering the same features and functionality across each implementation. This is one of the key benefits of FlexPod. Each of the component families shown in [Figure 1](#) (Cisco

UCS, Cisco Nexus, and NetApp FAS) offers platform and resource options to scale the infrastructure up or down while supporting the same features and functionality that are required under the configuration and connectivity best practices of FlexPod.

Design Principles

FlexPod addresses four primary design principles: scalability, elasticity, availability, and manageability. These architecture goals are as follows:

- Application availability: Ensure accessible and ready-to-use services
- Scalability: Address increasing demands with appropriate resources
- Flexibility: Provide new services or recovered resources without requiring infrastructure modification
- Manageability: Facilitate efficient infrastructure operations through open standards and APIs




Note

Performance and security are key design criteria that were not directly addressed in this project but have been addressed in other collateral, benchmarking and solution testing efforts. Functionality and basic security elements were validated.

FlexPod Discrete Uplink Design

[Figure 5](#) details FlexPod distinct uplink design with clustered Data ONTAP. As the illustration shows, the design is fully redundant in the compute, network, and storage layers. There is no single point of failure from a device or traffic path perspective.

Figure 5 *FlexPod Discrete Uplink Design with Cluster-Mode Data ONTAP*



The FlexPod Discrete Uplink Design is an end-to-end Ethernet transport solution supporting multiple LAN and SAN protocols, most notably FCoE. The solution provides a unified 10 Gigabit Ethernet enabled fabric defined by dedicated FCoE uplinks and dedicated Ethernet uplinks between the Cisco UCS fabric interconnects and the Cisco Nexus switches, as well as converged connectivity between the NetApp storage devices and the same multipurpose Cisco Nexus platforms.

The Discrete Uplink Design does not employ a dedicated SAN switching environment and requires no dedicated Fibre Channel connectivity. The Cisco Nexus 5500 Series Switches are configured in N port ID virtualization (NPIV) mode, providing storage services for the FCoE-based traffic traversing its fabric.

As illustrated in [Figure 5](#), link aggregation technology plays an important role, providing improved aggregate bandwidth and link resiliency across the solution stack. The NetApp storage controllers, Cisco UCS, and Nexus 5500 platforms all support active port channeling using 802.3ad standard Link Aggregation Control Protocol (LACP). Port channeling is a link aggregation technique offering link fault tolerance and traffic distribution (load balancing) for improved aggregate bandwidth across member ports. In addition, the Cisco Nexus 5000 Series features virtual port channel (vPC) capabilities. vPCs allow links that are physically connected to two different Cisco Nexus 5500 Series devices to appear as a single "logical" port channel to a third device, essentially offering device fault tolerance. vPCs address aggregate bandwidth and link and device resiliency. The Cisco UCS fabric interconnects and NetApp FAS controllers benefit from the Nexus vPC abstraction, gaining link and device resiliency, as well as full utilization of a nonblocking Ethernet fabric.

**Note**

The Spanning Tree protocol does not actively block redundant physical links in a properly configured vPC-enabled environment, so all ports should forward on vPC member ports.

This dedicated uplink design leverages FCoE-capable NetApp FAS controllers. From a storage traffic perspective, both standard LACP and Cisco's vPC link aggregation technologies play an important role in FlexPod distinct uplink design. [Figure 5](#) shows the use of dedicated FCoE uplinks between the Cisco UCS Fabric Interconnects and Cisco Nexus 5500 Unified Switches. The Cisco UCS Fabric Interconnects operate in the N-Port Virtualization (NPV) mode, meaning the servers' FC traffic is either manually or automatically pinned to a specific FCoE uplink, in this case either of the two FCoE port channels are pinned. The use of discrete FCoE port channels with distinct VSANs allows an organization to maintain traditional SAN A/B fabric separation best practices, including separate zone databases. The vPC links between the Cisco Nexus 5500 switches' and NetApp storage controllers' Unified Target Adapters (UTAs) are converged, supporting both FCoE and traditional Ethernet traffic at 10 Gigabit providing a robust "last mile" connection between the initiator and target.

The initial storage configuration of this solution is a two-node HA pair with clustered NetApp Data ONTAP®. An HA pair consists of like storage nodes such as FAS 32xx or 62xx Series. Scalability is achieved by adding storage capacity (disk/shelves) to an existing HA pair, or by adding HA pairs into the cluster or storage domain. For SAN environments, the clustered Data ONTAP offering allows up to three HA pairs that include six clustered nodes to form a single logical entity and large resource pool of storage that can be easily managed, logically carved, and efficiently consumed. For NAS environments, up to 24 nodes can be configured. In both the scenarios, the HA interconnect allows each HA node pair to assume control of its partner's storage (disk/shelves) directly. The local physical high-availability storage failover capability does not extend beyond the HA pair. Furthermore, a cluster of nodes does not have to include similar hardware. Rather, individual nodes in an HA pair are configured alike, allowing customers to scale as needed, as they bring additional HA pairs into the larger cluster.

**Note**

It is always advisable to seek advice from experts. Please consider reaching out to your NetApp account team or partner for further guidance.

The "Logical Build" section provides more details regarding the design of the physical components and virtual environment consisting of Windows Server 2012 with Hyper-V, Cisco UCS, and NetApp storage controllers.

Integrated System Components

The following components are required to deploy the Discrete Uplink design:

- Cisco Unified Computing System
- Cisco Nexus 5500 Series Switches
- Cisco Nexus 5596 Series Switches
- Cisco Nexus 1000V Switch for Hyper-V
- NetApp FAS and Data ONTAP
- Windows Server 2012 with Hyper-V Role
- System Center 2012 SP1

Cisco Unified Computing System

The Cisco Unified Computing System is a next-generation solution for blade and rack server computing. It is an innovative data center platform that unites computing, network, storage access, and virtualization into a cohesive system designed to reduce total cost of ownership (TCO) and increase business agility. The system integrates a low-latency, lossless 10 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multi-chassis platform in which all resources participate in a unified management domain. Managed as a single system, whether it has one server or 160 servers with thousands of virtual machines, Cisco Unified Computing System decouples scale from complexity. It accelerates the delivery of new services simply, reliably, and securely through end-to-end provisioning and migration support for both virtualized and nonvirtualized systems.

Cisco Unified Computing System consists of the following components:

- Cisco UCS 6200 Series Fabric Interconnects (www.cisco.com/en/US/products/ps11544/index.html) is a series of line-rate, low-latency, lossless, 10-Gbps Ethernet and FCoE interconnect switches providing the management and communication backbone for Cisco Unified Computing System. Cisco Unified Computing System supports VM-FEX technology.
- Cisco UCS 5100 Series Blade Server Chassis (www.cisco.com/en/US/products/ps10279/index.html) supports up to eight blade servers and up to two fabric extenders in a 6-rack unit (RU) enclosure.
- Cisco UCS B-Series Blade Servers (www.cisco.com/en/US/partner/products/ps10280/index.html): Increase performance, efficiency, versatility, and productivity with these Intel-based blade servers.
- Cisco UCS adapters (www.cisco.com/en/US/products/ps10277/prod_module_series_home.html): Wire-once architecture offers a range of options to converge the fabric, optimize virtualization, and simplify management. Cisco adapters support VM-FEX technology.
- Cisco UCS C-Series Rack Servers (www.cisco.com/en/US/products/ps10493/index.html) deliver unified computing in an industry-standard form factor to reduce TCO and increase agility.
- Cisco UCS Manager (www.cisco.com/en/US/products/ps10281/index.html) provides unified, embedded management of all software and hardware components in the Cisco UCS.

For more information, see www.cisco.com/en/US/products/ps10265/index.html.

Cisco Nexus 5500 Series Switch

The Cisco Nexus 5000 Series is designed for data center environments, with cut-through technology that enables consistent low-latency Ethernet solutions, front-to-back or back-to-front cooling, and data ports in the rear, bringing switching into close proximity with servers and making cable runs short and simple. The switch series is highly serviceable, with redundant, hot-pluggable power supplies and fan modules. It uses data center-class Cisco NX-OS Software for high reliability and ease of management.

The Cisco Nexus 5500 platform extends the industry-leading versatility of the Cisco Nexus 5000 Series 10 Gigabit Ethernet data center-class switches and provides innovative advances toward higher density, lower latency, and multilayer services. The Cisco Nexus 5500 platform is well suited for enterprise-class data center server access-layer deployments across a diverse set of physical, virtual, storage-access, and high-performance computing (HPC) data center environments.

- The switch used in this FlexPod architecture is the Cisco Nexus 5548UP. The following specifications describe the Nexus 5548UP switch:
- A 1RU Gigabit Ethernet and 10 Gigabit Ethernet switch
- 32 fixed unified ports on the base chassis and one expansion slot, totaling 48 ports
- The slot can support any of the three modules: unified ports, 1/2/4/8 native Fibre Channel, and Ethernet or FCoE

- Throughput of up to 960Gbps

For more information, see www.cisco.com/en/US/products/ps9670/index.html.

Cisco Nexus 2232PP 10GE Fabric Extender

The Cisco Nexus 2232PP provides 32 10 Gigabit Ethernet and FCoE Small Form-Factor Pluggable Plus (SFP+) server ports and eight 10 Gigabit Ethernet and FCoE SFP+ uplink ports in a compact 1RU form factor.

The built-in standalone software, Cisco Integrated Management Controller, manages Cisco UCS C-Series Rack Servers. When a Cisco UCS C-Series rack server is integrated with Cisco UCS Manager through the Cisco Nexus 2232 platform, the management controller does not manage the server anymore. Instead, it is managed by the Cisco UCS Manager software, using the Cisco UCS Manager GUI or command-line interface (CLI). The Cisco Nexus 2232 provides data and control traffic support for the integrated Cisco UCS C-Series server.

Cisco Nexus 1000V Switch for Hyper-V

Cisco Nexus 1000V Series Switches provide a comprehensive and extensible architectural platform for virtual machine and cloud networking. These switches are designed to accelerate server virtualization and multitenant cloud deployments in a secure and operationally transparent manner. Integrated into the Windows Server 2012 Hyper-V hypervisor and Virtual Machine Manager 2012 SP1, the Cisco Nexus 1000V Series provides:

- Advanced virtual machine networking, based on the Cisco NX-OS operating system and IEEE 802.1Q switching technology
- Policy-based virtual machine connectivity
- Mobile virtual machine security and network policy
- Nondisruptive operational model for server virtualization and networking teams
- Virtualized network services, with Cisco vPath providing a single architecture for Layer 4 through 7 network services such as load balancing, firewall, and WAN acceleration

These capabilities help ensure that the virtual machine is a basic building block of the data center, with full switching capabilities and a variety of Layer 4 through 7 services in both dedicated and multitenant cloud environments. With the introduction of VXLAN on the Nexus 1000V Series, network isolation among virtual machines can scale beyond traditional VLANs for cloud-scale networking.

For more information on Cisco Nexus 1000V Series Switches and Cisco Nexus 1010 Virtual Services Appliance, see:

- www.cisco.com/en/US/products/ps9902/index.html
- www.cisco.com/en/US/products/ps10785/index.html

NetApp FAS and Clustered Data ONTAP

NetApp solutions are user friendly, easy to manage, and quick to deploy, offering increased availability while consuming fewer IT resources. This means that they dramatically lower the lifetime TCO. Where other solutions manage complexity, NetApp eliminates it. A NetApp solution consists of hardware in the form of controllers and disk storage and the NetApp Data ONTAP operating system, the #1 branded storage OS.



Note

Source: IDC Worldwide Quarterly Disk Storage Systems Tracker Q4 2012, March 2013 (Open Networked Disk Storage Systems revenue)

NetApp offers a unified storage architecture. The term "unified" refers to a family of storage systems that simultaneously support SAN, network-attached storage (NAS), and iSCSI across many operating environments, such as VMware, Windows, and UNIX. This single architecture provides access to data by using industry-standard protocols, including NFS, CIFS, iSCSI, FCP, SCSI, FTP, and HTTP. Connectivity options include standard Ethernet (10/100/1000 megabit or 10 Gigabit Ethernet) and Fibre Channel (1, 2, 4, or 8 Gbps). In addition, all systems can be configured with high-performance solid-state drives (SSDs) or serial ATA (SAS) disks for primary storage applications; low-cost SATA disks for secondary applications (backup, archive, and so on); or a mix of the different disk types.

A storage system running Data ONTAP has a main unit, also known as the chassis or controller, which is the hardware device that receives and sends data. This unit detects and gathers information about the hardware configuration, the storage system components, the operational status, hardware failures, and other error conditions.

A storage system uses storage on disk shelves. The disk shelves are the containers or device carriers that hold disks and associated hardware such as power supplies, connectivity interfaces, and cabling.

If storage requirements change over time, NetApp storage offers the flexibility to change quickly as needed without expensive and disruptive "forklift" upgrades. For example, a logical unit number (LUN) can be changed from FC access to iSCSI access without moving or copying the data. Only a simple dismount of the FC LUN and a mount of the same LUN using iSCSI would be required. In addition, a single copy of data can be shared between Windows and UNIX systems while allowing each environment to access the data through native protocols and applications. If a system was originally purchased with all SATA disks for backup applications, high-performance SAS disks could be added to support primary storage applications such as Oracle, Microsoft Exchange, or Microsoft SQL Server.

NetApp storage solutions provide redundancy and fault tolerance through clustered storage controllers, hot-swappable redundant components (such as cooling fans, power supplies, disk drives, and shelves), and multiple network interfaces. This highly available and flexible architecture enables customers to manage all data under one common infrastructure while achieving mission requirements. The NetApp Unified Storage Architecture allows data storage with high availability and performance, dynamic expansion, and ease of management.

The storage efficiency built into Data ONTAP offers substantial space savings, allowing more data to be stored at lower cost. Data protection provides replication services, making sure that valuable data is backed up and recoverable. The following features provide storage efficiency and data protection:

- **Thin provisioning:** Volumes are created using "virtual" sizing. They appear to be provisioned at their full capacity, but they are actually created much smaller and use additional space only when it is actually needed. Extra unused storage is shared across all volumes, and the volumes can grow and shrink on demand.
- **NetApp Snapshot™ copies:** Automatically scheduled point-in-time copies that write only changed blocks, with no performance penalty. The Snapshot copies consume minimal storage space, because only changes to the active file system are written. Individual files and directories can easily be recovered from any Snapshot copy, and the entire volume can be restored back to any Snapshot state in seconds.
- **NetApp FlexClone® volumes:** Instant "virtual" copies of datasets that use near-zero space. The clones are writable, but only changes to the original are stored, so they provide rapid, space-efficient creation of additional data copies that are ideally suited for test and development environments.
- **Deduplication:** Removes redundant data blocks in primary and secondary storage with flexible policies to determine when the deduplication process is run.
- **Compression:** Compresses data blocks. Compression can be run whether or not deduplication is enabled and can provide additional space savings, whether run alone or together with deduplication.

- **NetApp SnapMirror®**: Volumes can be asynchronously replicated either within the cluster or to another cluster.

All of these capabilities are then exposed through the logical management construct of a NetApp Storage Virtual Machine.

Storage Virtual Machine

The secure logical storage partition through which data is accessed in clustered Data ONTAP is known as a Storage Virtual Machine (SVM), formerly known as Vserver. A cluster serves data through at least one and possibly multiple SVMs. An SVM is a logical abstraction that represents a set of physical resources of the cluster. Data volumes and logical network interfaces (LIFs) are created and assigned to an SVM and can reside on any node in the cluster to which the SVM has been given access. An SVM can own resources on multiple nodes concurrently, and those resources can be moved nondisruptively from one node to another. For example, a flexible volume can be nondisruptively moved to a new node and an aggregate, or a data LIF can be transparently reassigned to a different physical network port. The SVM abstracts the cluster hardware and is not tied to specific physical hardware.

An SVM is capable of supporting multiple data protocols concurrently. Volumes within the SVM can be junctioned together to form a single network-attached storage (NAS) namespace, which makes all of an SVM's data available to NFS and Common Internet File System (CIFS) clients through a single share or mount point. For example, a 24-node cluster licensed for UNIX and Windows file services that has a single SVM configured with thousands of volumes can be accessed from a single network interface on one of the nodes. SVMs also support block-based protocols, and LUNs can be created and exported using iSCSI, Fibre Channel (FC), or Fibre Channel over Ethernet (FCoE). Any or all of these data protocols can be configured for use in a given SVM.

An SVM is a secure entity; therefore, it is only aware of the resources that have been assigned to it and has no knowledge of other SVMs and their resources. Each SVM operates as a separate and distinct entity with its own security domain. Tenants can manage the resources allocated to them through a delegated SVM administration account. Each SVM can connect to unique authentication zones such as Active Directory (AD), Lightweight Directory Access Protocol (LDAP), or network internet service (NIS).

An SVM is effectively isolated from other SVMs that share the same physical hardware.

From a performance perspective, maximum- input/output operations per second (IOPS) and throughput levels can be set per SVM using quality of service (QoS) policy groups. This allows the cluster administrator to quantify the performance capabilities allocated to each SVM.

Clustered Data ONTAP is highly scalable, and additional storage controllers and disks can be easily added to existing clusters to scale capacity and performance to meet rising demands. Because there are virtual storage servers within the cluster, SVMs are also highly scalable. As new nodes or aggregates are added to the cluster, the SVM can be nondisruptively configured to use them. New disk, cache, and network resources can be made available to the SVM to create new data volumes or to migrate existing workloads to these new resources to balance performance.

This scalability also enables the SVM to be highly resilient. SVMs are no longer tied to the lifecycle of a given storage controller. As new replacement hardware is introduced, SVM resources can be nondisruptively moved from the old controllers to the new controllers and the old controllers can be retired from service while the SVM is still online and available to serve data.

Internal to Data ONTAP, an SVM is also referred to as a Vserver. In the command examples presented in this document, the `vserver` command is used to perform actions relating to SVMs.

SVM Components

Logical Interfaces

All SVM networking is done through LIFs created within the SVM. As logical constructs, LIFs are abstracted from the physical networking ports on which they reside.

Flexible Volumes

A flexible volume is the basic unit of storage for an SVM. An SVM has a root volume and can have one or more data volumes. Data volumes can be created in any aggregate that has been delegated by the cluster administrator for use by the SVM. Depending on the data protocols used by the SVM, volumes can contain either LUNs for use with block protocols, files for use with NAS protocols, or both concurrently. For access using NAS protocols, the volume must be added to the SVM namespace through the creation of a client-visible directory called a -junction.

Namespace

Each SVM has a distinct namespace through which all of the NAS data shared from that SVM can be accessed. This namespace can be thought of as a map to all of the junctioned volumes for the SVM, irrespective of the node or the aggregate on which they might physically reside. Volumes can be junctioned at the root of the namespace or beneath other volumes that are part of the namespace hierarchy. For more information about namespaces, refer to [TR-4129: Namespaces in Clustered Data ONTAP](#).

For more information, see

www.netapp.com/us/products/platform-os/data-ontap-8/index.aspx.

Windows Server 2012 Hyper-V

Windows Server 2012 Hyper-V provides significant increase in scalability and expands support for host processors and memory. New features include support for up to 64 processors and 1TB of memory for Hyper V virtual machines, in many cases supporting 4 to 16 times the density of processors, memory, cluster nodes, and running virtual machines. In addition, Windows Server 2012 Hyper-V supports innovative server features, including the ability to project a virtual nonuniform memory access (NUMA) topology onto a VM to provide optimal performance and workload scalability in large VM configurations. Windows Server 2012 Hyper-V also provides improvements to Dynamic Memory, including minimum memory and Hyper-V smart paging. Minimum memory allows Hyper V to reclaim the unused memory from VMs to allow higher VM consolidation numbers. Smart paging is used to bridge the memory gap between minimum and startup memory by allowing VMs to start reliably when the minimum memory setting has indirectly led to an insufficient amount of available physical memory during restart.

In addition to these memory enhancements, Windows Server 2012 Hyper-V allows runtime configuration of memory settings, including increasing the maximum memory and decreasing the minimum memory of running virtual machines. These updated features help make sure that your virtualization infrastructure can support the configuration of large, high-performance VMs to maintain demanding workloads.

Windows Server 2012 Hyper-V includes an update to the virtual hard disk format called VHDX. VHDX provides a higher capacity (up to 64 terabytes of storage), helps provide additional protection from corruption due to power failures, and prevents performance degradation on large-sector physical disks by optimizing structure alignment.

Windows Server 2012 Hyper-V also includes virtual FC support, allowing virtual machines to have unmediated access to SAN LUNs. Virtual FC enables scenarios such as running the Windows Failover Cluster Management feature inside the guest operating system of a VM connected to shared FC storage. Virtual FC supports multipath I/O (MPIO), NPIV for one-to-many mappings, and up to four virtual FC adapters per virtual machine.

Windows Server 2012 introduces several networking enhancements, including support for SR-IOV, third-party extensions to the Hyper-V extensible switch, QoS minimum bandwidth, network virtualization, and data center bridging (DCB).

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 easier management and distribution of workloads across the physical infrastructure. The ability of the virtualization layer to migrate running virtual machines from one server to another without downtime, as well as many other features that are provided by hypervisor-based virtualization technologies, enables a rich set of solution capabilities. These capabilities can be used by the automation, management, and orchestration layers to maintain desired states and proactively address decaying hardware and other issues that would otherwise cause faults or service disruptions.

Similar to the hardware layer, the automation, management, and orchestration layers must be able to manage the virtualization layer. Virtualization provides an abstraction of software from hardware that moves the majority of management and automation to software instead of requiring people to perform manual operations on physical hardware.

[Table 1](#) summarizes the scale improvements and feature enhancements in Windows Server 2012 Hyper-V compared to Windows Server 2008.

Table 1 *Comparison of Hyper-V Capabilities in Windows Server 2008 and Windows Server 2012 Hyper-V*

	Windows Server 2008	Windows Server 2008 R2	Windows Server 2012
Scale			
HW logical processor (LP) support	16 LPs	64 LPs	320 LPs
Physical memory support	1TB	1TB	0. 14TB
Cluster scale	16 nodes up to 1000 VMs	16 nodes up to 1000 VMs	64 nodes up to 4000 VMs
Virtual processor (VP) support	Up to 4 VPs	Up to 4 VPs	Up to 64 VPs
VM memory	Up to 64GB	Up to 64GB	Up to 1TB
Live migration	Yes, one at a time	Yes, one at a time	Yes, with no limits. As many as hardware will allow.
Live storage migration	No. Quick storage migration via Virtual Machine Manager (VMM)	0. 2No. Quick storage migration via VMM.	0. 3Yes, with no limits. As many as hardware will allow.
Servers in a cluster	16	16	64
VP-to-LP ratio	8:1	8:1 for server	No limits. As many as hardware will allow.
		12:1 for client (virtual desktop infrastructure (VDI))	
Storage			
Live storage migration	No. Quick storage migration via VMM.	No. Quick storage migration via VMM.	0. 4Yes, with no limits. As many as hardware will allow.
VMs on file storage	No	No	0. 5Yes, via Server Message Block (SMB) 3

Guest Fibre Channel	No	No	0. 6Yes
Virtual disk format	Virtual hard disk (VHD) up to 2 TB	VHD up to 2 TB	0. 7VHD up to 2 TB
			0. 8VHDX up to 64 TB
VM guest clustering	Yes, via iSCSI	Yes, via iSCSI	0. 9Yes, via iSCSI, Fibre Channel, and SMB 3
Native 4k disk support	No	No	Yes
Live VHD merge	No, offline.	No, offline.	Yes
Live new parent	No	No	Yes
Secure offloaded data transfer (ODX)	No	0. 10No	0. 11Yes
Networking		0. 12	0. 13
NIC teaming	Yes, via partners.	Yes, via partners.	Windows NIC Teaming inbox
VLAN tagging	Yes	Yes	Yes
MAC spoofing protection	No	Yes, with R2 SP1	Yes
Address Resolution Protocol (ARP) spoofing Protection	No	Yes, with R2 SP1	Yes
SR-IOV networking	No	No	Yes
Network QoS	No	No	Yes
Network metering	No	No	Yes
Network monitor modes	No	No	0. 14Yes
IP Security (IPsec) task offload	No	No	0. 15Yes
VM trunk mode	No	No	Yes
Manageability			
Hyper-V PowerShell	No	No	Yes
Network PowerShell	No	No	Yes
Storage PowerShell	No	No	Yes
SCONFIG	No	Yes	Yes
Enable/disable shell	No	No	Yes, MinShell
	(Server core at OS setup)	0. 16(Se rver core at OS setup)	
VMConnect support for Microsoft RemoteFX	–	No	0. 17Yes

The Hyper-V host cluster requires different types of network access, as described in [Table 2](#).

Table 2 *Host Cluster Networks*

Network Access Type	Purpose of Network Access Type	Network Traffic Requirements	Recommended Network Access
Virtual machine access	Workloads running on virtual machines usually require external network connectivity to service client requests.	Varies	Public access that can be teamed for link aggregation or to fail over the cluster.
Cluster and cluster shared volumes (CSV)	Preferred network used by the cluster for communications to maintain cluster health. Also used by CSV to send data between owner and nonowner nodes. If storage access is interrupted, this network is used to access CSV or to maintain and back up CSV. The cluster should have access to more than one network for communication to make sure that the cluster is highly available.	Usually low bandwidth and low latency. Occasionally, high bandwidth.	Private access
Live migration	Transfer VM memory and state.	High bandwidth and low latency during migrations	Private access
SMB	Access to VHDX and virtual machines files that are stored on SMB 3.0 file shares.	High bandwidth and low latency	Usually dedicated and private access.
Storage	Access storage through iSCSI .	High bandwidth and low latency	Usually dedicated and private access.
Management	Managing the Hyper-V management operating system. This network is used by Hyper-V Manager.	Low bandwidth	Public access that can be teamed to fail over the cluster.

Highly available host servers are a critical component of a dynamic, virtual infrastructure. 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 (nodes) are connected physically. If one of the cluster nodes fails, another node begins to provide service. In the case of a planned live migration, users experience no perceptible service interruption.

Microsoft System Center 2012 SP1

Microsoft System Center 2012 SP1 helps organizations deliver flexible and cost-effective private cloud infrastructure in a self-service model, while using existing data center hardware and software investments. It provides a common management experience across data centers and private or partner hosted clouds. To deliver the best experience for modern applications, System Center 2012 SP1 offers deep insight into applications, right down to client script performance. System Center 2012 SP1 delivers the tools and capabilities that organizations need to scale their capacity and, where necessary, use cloud resources as well.

Microsoft System Center 2012 offers unique application management capabilities that can enable you to deliver agile, predictable application services. Using the App Controller, Operations Manager, and Virtual Machine Manager components of System Center 2012, you can deliver "applications as a service," where a "service" is a deployed instance of a cloud-style application, along with its associated configuration and virtual infrastructure. The following application management capabilities are included:

Standardized Application Provisioning

- Virtual Machine Manager offers service templates to help you define standardized application blueprints. A service template would typically include specifications for the hardware, operating system, and application packages that compose the service.
- Supports multiple package types for Microsoft .NET applications, including MS Deploy for the web tier (IIS), Microsoft Server Application Virtualization (Server App-V) for the application tier, and SQL Server DAC for the data tier.
- Specifies application configuration requirements such as topology, elasticity and scale-out rules, health thresholds, and upgrade rules.
- Server App-V, a unique technology in Virtual Machine Manager, optimizes applications for private cloud deployments by abstracting the application from the underlying OS and virtual infrastructure. By enabling image-based management, Server App-V simplifies application upgrades and maintenance.

Comprehensive Hybrid Application Management

- App Controller offers application owners a single view to manage application services and virtual machines, whether they are on-premises, at service providers, or on Windows Azure.
- App Controller provides the ability to deploy and migrate virtual machines to the Windows Azure Virtual Machine service. You can migrate core applications such as Microsoft SQL Server, Active Directory, and Microsoft SharePoint Server from on-premises environments to Windows Azure with just a few mouse clicks.

360-Degree Application Monitoring, Diagnosis, and Dev-Ops

- Operations Manager offers deep application and transaction monitoring insight for .NET applications (and J2EE application servers) and helps you efficiently isolate the root cause of application performance issues down to the offending line of code.
- Outside-in monitoring with Global Service Monitor (GSM) and Operations Manager provides real time visibility into application performance as experienced by end users.
- Operations Manager and GSM integrate with Microsoft Visual Studio to facilitate dev-ops collaboration, thereby helping you remediate application issues faster.
- Operations Manager offers easy-to-use reporting and custom dashboarding.

Using the Service Manager and Orchestrator components of System Center 2012, you can automate core organizational process workflows such as incident management, problem management, change management, and release management. You can also integrate and extend your existing toolsets and build flexible workflows (or runbooks) to automate processes across your IT assets and organizations. The following service delivery and automation capabilities are provided:

Standardize IT Services

- Define standardized service offerings by using dependencies in a centralized configuration management database (CMDB).
- Publish standardized service offerings through the Service Catalog offered by Service Manager.
- Provision and allocate pooled infrastructure resources to internal business unit ITs (BUIs) using the Cloud Services Process Pack (CSPP) that's natively integrated into Service Manager.
- Chargeback (or showback) storage, network, and computing costs to BUIs; specify pricing for BUIs at different levels of granularity.
- Make sure compliance with pertinent industry regulations and business needs with the IT GRC Process Pack.

Enable IT Service Consumers to Identify, Access, and Request Services

- Enable self-service infrastructure with the self-service portal offered by Service Manager.
- Set access and resource quota levels on a per-user or per-BUIT basis.
- Capture and track required service request information.

Automate Processes and Systems Necessary to Fulfill Service Requests

- Integrate and extend automation across System Center and third-party management toolsets (including BMC, HP, IBM, and VMware) with Orchestrator Integration Packs; extend automation to Windows Azure virtual machine workflows without the need for coding or scripting.
- Orchestrate automated workflows across multiple processes, departments, and systems.
- Automate provisioning of service requests for end-to-end request fulfillment.

Microsoft System Center 2012 SP1 provides a common management toolset to help you configure, provision, monitor, and operate your IT infrastructure. If your infrastructure is like that of most organizations, you have physical and virtual resources running heterogeneous operating systems. The integrated physical, virtual, private, and public cloud management capabilities in System Center 2012 can help you ensure efficient IT management and optimized ROI of those resources. The following infrastructure management capabilities are provided:

Provision your Physical and Virtual Infrastructure

- Support deployment and configuration of virtual servers and Hyper-V with Virtual Machine Manager.
- Manage VMware vSphere and Citrix XenServer using one interface.
- Automatically deploy Hyper-V to bare metal servers and create Hyper-V clusters.
- Provision everything from operating systems to physical servers, patches, and endpoint protection with Configuration Manager.

Provision Private Clouds

- Use "create cloud" functionality in Virtual Machine Manager to aggregate virtual resources running on Hyper-V, vSphere, and XenServer into a unified private cloud fabric.
- Customize and assign private cloud resources to suit your organization's needs.
- Deliver self-service capability for application owners to request and automate provisioning of new private cloud resources.

Operate Your Infrastructure

- Use a single console and customizable dashboards in Operations Manager to monitor and manage your physical, virtual, networking, application, and cloud resources.
- Dynamically optimize virtual resources for load balancing and power efficiency.
- Protect your physical and virtual resources with Endpoint Protection and Data Protection Manager.
- Automatically patch your physical and virtual resources with Configuration Manager and Virtual Machine Manager.
- Automatically track and create custom reports for hardware inventory, software inventory, and software usage metering.

Domain and Element Management

This section provides general descriptions of the domain and element Managers used during the validation effort. The following managers are used:

- Cisco UCS Manager
- Cisco UCS Power Tool
- Cisco UCS System Center Operations Manager 2012 Management Pack
- Cisco UCS System Center Orchestrator 2012 Integration Pack
- Cisco UCS System Center Virtual Machine Manager 2012 User Interface Extensions
- Cisco Nexus 1000V for Hyper-V VSM
- NetApp OnCommand® System Manager
- NetApp SnapDrive® for Windows
- NetApp SnapManager® for Hyper-V
- Microsoft System Center 2012 SP1
 - App Controller
 - Operations Manager
 - Orchestrator
 - Service Manager
 - Virtual Machine Manager

Cisco UCS Manager

Cisco UCS Manager provides unified, centralized, embedded management of all Cisco Unified Computing System software and hardware components across multiple chassis and thousands of virtual machines. Administrators use this software to manage the entire Cisco Unified Computing System as a single logical entity through an intuitive GUI, a CLI, or an XML API.

Cisco UCS Manager resides on a pair of Cisco UCS 6200 Series Fabric Interconnects using a clustered, active-standby configuration for high availability. The software gives administrators a single interface for performing server provisioning, device discovery, inventory, configuration, diagnostics, monitoring, fault detection, auditing, and statistics collection. Cisco UCS Manager service profiles and templates support versatile role- and policy-based management, and system configuration information can be exported to CMDBs to facilitate processes based on ITIL® concepts.

Service profiles let server, network, and storage administrators treat Cisco UCS servers as raw computing capacity to be allocated and reallocated as needed. The profiles define server I/O properties and are stored in the Cisco UCS 6200 Series Fabric Interconnects. Using service profiles, administrators can provision infrastructure resources in minutes instead of days, creating a more dynamic environment and more efficient use of server capacity.

Each service profile consists of a server software definition and the server's LAN and SAN connectivity requirements. When a service profile is deployed to a server, Cisco UCS Manager automatically configures the server, adapters, fabric extenders, and fabric interconnects to match the configuration specified in the profile. The automatic configuration of servers, NICs, host bus adapters (HBAs), and LAN and SAN switches lowers the risk of human error, improves consistency, and decreases server deployment times.

Service profiles benefit both virtualized and nonvirtualized environments. The profiles increase the mobility of nonvirtualized servers, such as when moving workloads from server to server or taking a server offline for service or an upgrade. Profiles can also be used in conjunction with virtualization clusters to bring new resources online easily, complementing existing virtual machine mobility.

For more information on Cisco UCS Manager, visit www.cisco.com/en/US/products/ps10281/index.html.

Cisco UCS PowerTool

Cisco UCS PowerTool is a PowerShell module that helps automate all aspects of Cisco UCS Manager, including server, network, storage, and hypervisor management. PowerTool enables easy integration with existing IT management processes and tools.

Cisco UCS PowerTool is a flexible and powerful command-line toolkit that includes more than 1500 PowerShell cmdlets, providing customers with an efficient, cost-effective, and easy-to-use interface to integrate and automate Cisco UCS management with Microsoft products and many third-party products. It lets you take advantage of the flexible and powerful scripting environment offered by Microsoft PowerShell.

Cisco UCS System Center Operations Manager 2012 Management Pack

The Cisco UCS Management Pack for Microsoft System Center Operations Manager provides visibility into the health, performance, and availability of a Cisco UCS domain through a single, familiar, and easy-to-use interface. The management pack contains rules which monitor, for example, chassis, blade servers, rack servers, and service profiles across multiple Cisco UCS domains.

Cisco UCS System Center Orchestrator 2012 Integration Pack

The Cisco UCS Integration Pack is an add-on for Microsoft System Center 2012 and 2012 SP1 Orchestrator (SCO) that enables you to automate Cisco UCS Manager tasks. You can use Cisco UCS Integration Pack to create workflows that interact with and transfer information to other Microsoft System Center products such as Microsoft System Center 2012 Operations Manager.


Cisco UCS System Center Virtual Machine Manager 2012 User Interface Extensions

Cisco UCS Add-in for Virtual Machine Manager provides an extension to the virtual machine manager user interface. The extended interface enables you to manage the Cisco UCS servers (blade servers and rack-mount servers).

Cisco Nexus 1000V for Hyper-V

The Cisco Nexus 1000V is a logical switch that fully integrates into Windows Server 2012 Hyper-V and Virtual Machine Manager 2012 SP1. The Cisco Nexus 1000V operationally emulates a physical modular switch, with a Virtual Supervisor Module (VSM) providing control and management functionality to multiple line cards. In the case of the Nexus 1000V, the Cisco Virtual Ethernet Module (VEM) is a forwarding extension for the Hyper-V logical switch when installed on the Hyper-V host. [Figure 6](#) describes the Cisco Nexus 1000V architecture.


Figure 6 *Cisco Nexus 1000V for Hyper-V Architecture*



NetApp OnCommand Plug-in for Microsoft

NetApp OnCommand Plug-in for Microsoft (OCPM) (Figure 7) enables and simplifies the management of servers and storage systems from within Microsoft® System Center SP1. OCPM offers native integration with System Center Operations Manager and System Center Virtual Machine Manager. These integrations provide the intelligence to make System Center fully storage aware, thus simplifying day-to-day administration of NetApp storage and amplifying the effectiveness of monitoring and alerting in Microsoft System Center.


Figure 7 *NetApp OnCommand Plug-in for Microsoft*



NetApp OnCommand System Manager

NetApp OnCommand System Manager makes it possible for administrators to manage individual or clusters of NetApp storage systems through an easy-to-use browser-based interface. System Manager comes with wizards and workflows, simplifying common storage tasks such as creating volumes, LUNs, qtrees, shares, and exports, which saves time and prevents errors. System Manager works across all NetApp storage: FAS2000, FAS3000, and FAS6000 series as well as V-Series systems. Figure 8 shows a sample screen in NetApp OnCommand System Manager.

Figure 8 Sample NetApp OnCommand System Manager Screen



NetApp SnapDrive for Windows

NetApp SnapDrive for Windows (SDW) is an enterprise-class storage and data management application that simplifies storage management and increases availability of application data. The key functionality includes storage provisioning, file system-consistent data Snapshot copies, rapid application recovery, and the ability to manage data easily. SDW complements the native file system and volume manager and integrates seamlessly with the clustering technology supported by the host OS.

NetApp SnapManager for Hyper-V

NetApp SnapManager for Hyper-V (SMHV) automates and simplifies backup and restore operations for virtual machines running in Microsoft Windows Server 2012 Hyper-V environments hosted on Data ONTAP storage systems. SMHV enables application-consistent dataset backups according to protection policies set by the storage administrator. VM backups can also be restored from those application-consistent backups.

SnapManager for Hyper-V makes it possible to back up and restore multiple VMs across multiple hosts. Policies can be applied to the datasets to automate backup tasks such as scheduling, retention, and replication.

System Center 2012 SP1 App Controller

App Controller is a member of the Microsoft System Center suite. It provides a common self-service experience that can help administrators easily configure, deploy, and manage virtual machines and services across private clouds. App Controller provides the user interface for connecting and managing workloads post-provisioning.

System Center 2012 SP1 Operations Manager

Operations Manager is a member of the Microsoft System Center suite. It provides infrastructure monitoring that is flexible and cost-effective, helps ensure the predictable performance and availability of vital applications, and offers comprehensive monitoring for your data center and private cloud.

System Center 2012 SP1 Service Manager

Service Manager is a member of the Microsoft System Center suite. It provides an integrated platform for automating and adapting your organization's IT service management best practices, such as those found in Microsoft Operations Framework (MOF) and ITIL. It provides built-in processes for incident and problem resolution, change control, and asset lifecycle management.

System Center 2012 SP1 Virtual Machine Manager

Virtual Machine Manager (VMM) is a member of the Microsoft System Center suite. It is a management solution for the virtualized data center, enabling you to configure and manage your virtualization host, networking, and storage resources in order to create and deploy virtual machines and services to private clouds that you have created.

Microsoft SQL Server 2012 SP1

Microsoft SQL Server is a highly available database management and analysis system for e-commerce, line-of-business, and data warehousing solutions. It stores data and provides reporting services for the System Center components.

A Closer Look at FlexPod Discrete Uplink Design

Physical Build: Hardware and Software Revisions

Table 3 describes the hardware and software versions used during validation. It is important to note that Cisco, NetApp, and Microsoft have interoperability matrices that should be referenced to determine support for a specific implementation of FlexPod. Please refer to the following links:

- NetApp Interoperability Matrix Tool: <http://support.netapp.com/matrix/>
- Cisco UCS Hardware and Software Interoperability Tool:
www.cisco.com/web/techdoc/ucs/interoperability/matrix/matrix.html
- Windows Server 2012 Catalog:
www.windowsservercatalog.com/results.aspx?&chtext=&cstext=&csstext=&chbtext=&bCatID=1283&cpID=0&avc=38&ava=0&avq=0&OR=1&PGS=25&ready=0

Table 3 **Validated Software and Firmware Versions**

Layer	Device	Image	Comments
Computing	Cisco UCS 6200 Series Fabric Interconnects	2.1(3a)	Includes UCSM
	Cisco UCS 5108 Chassis	2.1(3a)	Includes the UCS-IOM 2208XP
	Cisco Nexus 2232 Fabric Extender	2.1(3a)	
	Cisco UCS B200 M3 and C220 M3	2.1(3a)	B200 M3 using Cisco UCS VIC 1240 C220 M3 using Cisco VIC 1225
	Cisco E-NIC	2.3.0.4	Ethernet NIC driver
	Cisco F-NIC	2.3.0.12	HBA driver
Network	Cisco Nexus 5000 NX-OS	6.0(2)N1(2a)	
Network	Cisco Nexus 1000V NX-OS	5.2(1)SM1(5.1)	Nexus 1000V for Microsoft Hyper-V
Storage	NetApp FAS Model 3250-AE	Data ONTAP 8.2	
Software	Windows Server 2012	6.02.9200	Updated with current updates
	Cisco UCS PowerTool	1.01.0	
	Cisco UCS Management Pack for System Center Operations Manager 2012	2.6.1	
	Cisco UCS Integration Pack for System Center Orchestrator 2012	1.0	

	Cisco UCS User Interface Extensions for System Center Virtual Machine Manager 2012	1.0.1	
	NetApp OnCommand	4.0	
	NetApp SnapDrive for Windows	7.0	
	NetApp SnapManager for Hyper-V	2.0	
	Microsoft System Center App Controller		Updated with current updates
	Microsoft System Center Operations Manager	7.0.9538.0	Updated with current updates
	Microsoft System Center Orchestrator	7.1.3002.0	Updated with current updates
	Microsoft System Center Service Manager	7.5.2905.0	Updated with current updates
	Microsoft System Center Virtual Machine Manager	3.1.6018.0	Must include Rollup Update 2 and the latest updates
	Microsoft SQL Server 2012 SP1	11.0.3321.0	Updated with current updates

Logical Build

Figure 5 illustrates the FlexPod Discrete Uplink Design. The design is physically redundant across the stack, addressing Layer 1 high availability requirements, but there are additional Cisco and NetApp technologies and features that make for an even more effective solution. This section discusses the logical configuration validated for FlexPod.


FlexPod allows organizations to adjust the individual components of the system to meet their particular scale or performance requirements. One key design decision in the Cisco UCS domain is the selection of I/O components. Numerous combinations of I/O adapter, I/O module (IOM), and fabric interconnect are available, so it is important to understand the impact of these selections on the overall flexibility, scalability and resiliency of the fabric.

Figure 9 illustrates the available backplane connections in the Cisco UCS 5100 series chassis. As the illustration shows, each of the two fabric extenders (IOMs) has four 10GBASE KR (802.3ap) standardized Ethernet backplane paths available for connection to the half-width blade slot. This means that each half-width slot has the potential to support up to 80Gb of aggregate traffic. What is realized depends on several factors, namely:

- Fabric extender model (2204 or 2208)
- Modular LAN on Motherboard (mLOM) card
- Mezzanine slot card

The Cisco UCS 2208XP Fabric Extender has eight 10 Gigabit Ethernet, FCoE-capable, SFP+ ports that connect the blade chassis to the fabric interconnect. The Cisco UCS 2204XP has four external ports with identical characteristics to connect to the fabric interconnect. Each Cisco UCS 2208XP has 32 10 Gigabit Ethernet ports connected through the midplane KR lanes to each half-width slot in the chassis, while the 2204XP has 16. This means the 2204XP enables two KR lanes per half-width blade slot while the 2208XP enables all four. The number of KR lanes indicates the potential I/O available to the chassis and therefore to the blades.


Figure 9 Cisco UCS B-Series M3 Server Chassis Backplane Connections



Port aggregation is supported by the second-generation Cisco UCS 6200 Series Fabric Interconnects, 2200 Series Fabric Extenders, and 1200 Series VICs. This capability allows for workload rebalancing between these devices, providing link fault tolerance in addition to increased aggregate bandwidth within the fabric. It should be noted that in the presence of second-generation VICs and fabric extenders, fabric port channels will automatically be created in the fabric. Fabric port channels between the fabric extenders and fabric interconnects are controlled via the Chassis/FEX discovery policy. Figure 10 illustrates the two modes of operation for this policy. In Discrete Mode each FEX KR connection and therefore server connection is tied or pinned to a network fabric connection homed to a port on the fabric interconnect. In the presence of a failure on the external "link," all KR connections are disabled within

the FEX I/O module. In the case of a fabric port channel discovery policy, the failure of a network fabric link allows for redistribution of flows across the remaining port channel members. This is less disruptive to the fabric.

Figure 10 *Example of Discrete Mode Versus Port Channel Mode*



Note

First-generation Cisco UCS hardware is compatible with the second-generation gear, but it will operate only in discrete mode.

Figure 11 represents one of the Cisco UCS B200 M3 backplane connections validated for the FlexPod. The B200 M3 uses a VIC 1240 in the mLOM slot with an empty mezzanine slot. The FEX 2204XP enables 2 KR lanes to the half-width blade, while the global discovery policy dictates the formation of a fabric port channel.

Figure 11 *Validated Cisco UCS Backplane Configurations: VIC 1240 Only*





Figure 12 illustrates another Cisco UCS B200 M3 instance in the test bed. In this instance the mezzanine slot is populated with the port expander option. This passive device provides connectivity for the unused ports on the UCS VIC 1240, essentially enabling the 40-Gb potential of the mLOM card. Beyond the raw capacity improvements is the creation of two more automatic port channels between the fabric extender and the server. This provides link resiliency at the adapter level and double the bandwidth available to the system. (dual 2x10 Gb).

Figure 12 Validated Cisco UCS Backplane Configurations: VIC 1240 with Port Extender**Note**

See Appendix B for additional combinations of UCS second-generation hardware and the connectivity options they afford.

The FlexPod defines two FCoE port channels (Po1 and Po2) and two LAN port channels (Po13 and Po14). The FCoE port channels carry only Fibre Channel traffic that is associated with a VSAN/VLAN set, with the set in turn supported only on one side of the fabric A or B. As in this example, the virtual HBA (vHBA) "FABRIC-A" is defined in the service profile. The vHBA uses a virtual circuit, VC 737, to traverse the Cisco UCS unified fabric to port channel Po1, where FCoE traffic egresses the Cisco UCS domain and enters the Cisco Nexus 5500 platform. Fabric A supports a distinct VSAN, which is not present on Fabric B, thus maintaining fabric isolation.

It has been said that design is the art of compromise; however, with the FlexPod architecture there is very little sacrifice. Availability and performance are present. The question becomes what combination meets the application and business requirements of the organization. [Table 4](#) describes the availability and performance aspects of the second-generation Cisco UCS I/O gear.

Table 4 *Cisco UCS B-Series M3 FEX 2204XP and 2208XP Options*

Reliability Technique	Fabric failover, adapter redundancy, and port channel				VIC 1240 and VIC 1280
	Fabric failover and adapter redundancy		VIC 1240 and VIC 1280		
	Fabric failover and port channel		VIC 1240 with Port expander		VIC 1240 with port expander
	Fabric failover	VIC 1240	VIC 1240		
		20 Gb	40 Gb	60 Gb	80 Gb
Aggregate Bandwidth (Performance)					

*Orange shading indicates that the FEX 2208XP is in use. All other values are based on the FEX 2204XP model.

**Note**

Table 4 assumes the presence of Cisco UCS 6200 Series Fabric Interconnects

**Note**

Third-party generation 3 PCIe adapters were not validated.

A balanced fabric is critical within any data center environment. Given the myriad traffic types (live migration, CSV, FCoE, public, control traffic, etc.) the FlexPod must be able to provide for specific traffic requirements while simultaneously being able to absorb traffic spikes and protect against traffic loss. To address these requirements, the Cisco UCS QoS system classes and Cisco Nexus QoS policies should be configured. In this validation effort the FlexPod was configured to support jumbo frames with a maximum transmission unit (MTU) size of 9000. This class was assigned to the Best-Effort class. With regard to jumbo frames, it is important to make sure the MTU settings are applied uniformly across the stack to prevent fragmentation and the negative performance implications inconsistent MTUs may introduce.

Cisco UCS C-Series Server Design

Cisco UCS Manager 2.1 provides two connectivity modes for Cisco UCS C-Series Rack Server management:

- Dual-wire management (shared LAN on motherboard [LOM]): This management mode is supported in Cisco UCS Manager releases earlier than 2.1. Shared LOM ports on the rack server are used exclusively for carrying management traffic. A separate cable connected to one of the ports on the PCIe card carries the data traffic. Using two separate cables for managing data traffic and management traffic is also referred to as dual-wire management.
- Single-wire management (Sideband): Cisco UCS Manager version 2.1 introduces an additional rack server management mode using Network Controller Sideband Interface (NC-SI). Cisco UCS Virtual Interface Card 1225 uses the NC-SI, which can carry both data traffic and management traffic on the same cable. This new feature is referred to as single-wire management and will allow for denser server to FEX deployments.


From a functional perspective, the 1 RU Nexus FEX 2232PP replaces the Cisco UCS 2204 or 2208 IOMs that are located with the UCS 5108 blade chassis. Each 10 Gigabit Ethernet VIC port connects to Fabric A or B via the FEX. The FEX and fabric interconnects form port channels automatically based on the chassis discovery policy, providing a link resiliency to the C-Series server. This is identical to the behavior of the IOM to fabric interconnect connectivity. From a logical perspective, the virtual circuits formed within the UCS domain are consistent between the B-Series and C-Series deployment models and the virtual constructs formed at the Hyper-V.

Cisco Nexus 5500 Series Switch

As [Figure 13](#) shows, the Cisco Nexus 5500 Series Switch provides Ethernet and, in particular, FCoE connectivity for the Cisco UCS domain as well as for the NetApp storage controllers. From an Ethernet perspective, the Nexus 5500 uses virtual PortChannel (vPC) allowing links that are physically connected to two different Cisco Nexus 5000 Series devices to appear as a single port channel to a third device, in this case the UCS fabric interconnects and NetApp controllers. vPCs provide the following benefits:

- Allow a single device to use a port channel across two upstream devices
- Eliminate Spanning Tree Protocol blocked ports
- Provide a loop-free topology
- Use all available uplink bandwidth
- Provide fast convergence if either the link or a device fails
- Provide link-level resiliency
- Help ensure high availability

Figure 13 Discrete Uplink Design Cisco Nexu 5500 and NetApp Storage Focus



vPCs requires a "peer link," which is documented as port channel 10 in [Figure 13](#). It is important to note that the VLAN associated with the FCoE traffic does not traverse this peer link. Remember that the FCoE VLAN is associated or mapped to a VSAN, typically using the same numeric ID. It is crucial that the fabrics do not mix, maintaining SAN A/B isolation best practices. In addition, the vPC links facing the Cisco UCS fabric interconnects, vPC13 and vPC14, do not carry any FCoE traffic. Do not define any FCoE VLANs on these links. However, the vPCs connected to the NetApp UTAs are converged, supporting both FCoE and all other VLANs associated with LAN protocols.

The vPC peer keepalive link is a required component of a vPC configuration. The peer keepalive link allows each vPC-enabled switch to monitor the health of its peer. This link accelerates convergence and reduces the occurrence of split-brain scenarios. In this validated solution, the vPC peer keepalive link uses the out-of-band management network. (This link is not shown in [Figure 13](#).)

Each Cisco Nexus 5500 Series Switch defines a port channel dedicated to FCoE and connected to the Cisco UCS fabric interconnects, in this instance Po15 and Po16. Each discrete port channel supports a single VLAN associated with Fabric A or Fabric B. A virtual Fibre Channel interface (vfc) is then bound to the logical port channel interface. This same construct is applied to the vPCs facing the NetApp storage controllers, in this example vfc11 and vfc12. This assures universal accessibility of the fabric to each NetApp storage node in case of failures. To maintain SAN A and B isolation, vfc 11 and 12 are associated with a different VLAN/VSAN pairing, meaning the vPCs facing the NetApp storage systems support all LAN and FCoE traffic but have unique FCoE VLANs defined on each Nexus switch.

**Note**

It is considered a best practice to name your vfc for the port channel it is residing on; for example, vfc15 is on port channel 15.

The Cisco Nexus 5500 in the FlexPod design provides Fibre Channel services to the Cisco UCS and NetApp FAS platforms. Internally the Nexus 5500 platforms are performing FC zoning to enforce access policy between UCS-based initiators and FAS-based targets.

FlexPod is a converged infrastructure platform. This convergence is possible due to the support of Ethernet enhancements across the integrated computing stack with regard to bandwidth allocation and flow control based on traffic classification. As such, it is important to implement these QoS techniques to help ensure quality of service in the FlexPod.

- Priority Flow Control (PFC) 802.1Qbb: Lossless Ethernet using a PAUSE on a per class of service (CoS)
- Enhanced Transmission Selection (ETS) 802.1Qaz: Traffic protection through bandwidth management
- Data Center Bridging Capability Exchange (DCBX): Negotiates Ethernet functionality between devices (PFC, ETS, and CoS values)

The Cisco Nexus 5500 supports these capabilities through QoS policy. QoS is enabled by default and managed using Cisco Modular QoS CLI (MQC), providing class-based traffic control. The Nexus system will instantiate basic QoS classes for Ethernet traffic and a system FCoE class (class-fcoe) when the FCoE feature is enabled. It is important to align the QoS setting (CoS, MTU) within the Nexus 5500 and the UCS fabric interconnects. DCBX signaling can affect the NetApp controller, so be sure to allocate the proper bandwidth, based on the site's application needs, to the appropriate CoS classes and keep MTU settings consistent in the environment to avoid fragmentation issues and improve performance.

The following list summarizes the best practices used in the validation of the FlexPod architecture:

- Cisco Nexus 5500 features enabled
 - FCoE uses PFC, ETS, and DCBX to provide a lossless fabric
 - NPIV, allows the network fabric port (N-port) to be virtualized and support multiple Fibre Channel initiators on a single physical port
 - LACP
 - Cisco vPC for link and device resiliency
 - Link Layer Discovery Protocol (LLDP) allows the Nexus 5000 to share and discover DCBX features and capabilities between neighboring FCoE-capable devices.
 - Enable Cisco Discovery Protocol for infrastructure visibility and troubleshooting
- vPC considerations
 - Define a unique domain ID
 - Set the priority of the intended vPC primary switch lower than the secondary (default priority is 32768)
 - Establish peer keepalive connectivity. It is recommended to use the out-of-band management network (mgmt0) or a dedicated switched virtual interface (SVI)
 - Enable the vPC auto-recovery feature

- Enable IP arp synchronization to optimize convergence across the vPC peer link. Note: Cisco Fabric Services over Ethernet is responsible for synchronization of configuration, Spanning Tree, MAC and VLAN information, which removes the requirement for explicit configuration. The service is enabled by default.
- A minimum of two 10 Gigabit Ethernet connections are required for vPC.
- All port channels should be configured in LACP active mode
- Spanning tree considerations
 - Help ensure that the path cost method is set to long. This setting accounts for 10 Gigabit Ethernet links in the environment.
 - Do not modify the spanning tree priority, the assumption being that this is an access layer deployment.
 - Loopguard is disabled by default.
 - Bridge Protocol Data Unit (BPDU) guard and filtering are enabled by default.
 - Bridge assurance is enabled only on the vPC peer link.
 - Ports facing the NetApp storage controller and Cisco Unified Computing System are defined as "edge" trunk ports.

For configuration details, refer to the Cisco Nexus 5000 Series Switches configuration guides at www.cisco.com/en/US/products/ps9670/products_installation_and_configuration_guides_list.html.

Hyper-V

The Hyper-V role enables you to create and manage a virtualized computing environment by using virtualization technology that is built into Windows Server 2012. Installing the Hyper-V role installs the required components and optionally installs management tools. The required components include Windows hypervisor, Hyper-V Virtual Machine Management Service, the virtualization Windows Management Interface (WMI) provider, and other virtualization components such as the VMBus, virtualization service provider (VSP), and virtual infrastructure driver (VID).

The management tools for the Hyper-V role consist of:

- GUI-based management tools: Hyper-V Manager, an MMC snap-in, and Virtual Machine Connection, which provides access to the video output of a virtual machine so you can interact with the VM.
- Hyper-V-specific cmdlets for Windows PowerShell. Windows Server 2012 includes a Hyper-V module, which provides command-line access to all the functionality available in the GUI, as well as functionality not available through the GUI.

Windows Server 2012 introduced many new features and enhancements for Hyper-V. The following are some of the more notable enhancements that are used in this design.

- **Host Scale-Up:** Greatly expands support for host processors and memory. New features include support for up to 64 virtual processors and 1 TB of memory for Hyper-V guests, a new VHDX virtual hard disk format with a larger disk capacity of up to 64 TB, and additional resiliency. These features help ensure that your virtualization infrastructure can support the configuration of large, high-performance virtual machines to support workloads that might need to scale up significantly.
- **SR-IOV-capable network devices** that let an SR-IOV virtual function of a physical network adapter be assigned directly to a virtual machine. This increases network throughput and reduces network latency while also reducing the host CPU overhead that is required for processing network traffic. Refer back to Figure 6 to see the architecture of SR-IOV support in Hyper-V.

FlexPod Discrete Uplink Design with Clustered Data


Figure 14 shows FlexPod with Clustered Data. Clustered Data ONTAP consists of two or more storage controllers with shared media. The NetApp FAS controllers use redundant 10-Gb converged adapters configured in a two-port interface group (IFGRP). Each port of the IFGRP is connected to one of the upstream switches, allowing multiple active paths by using the Nexus vPC feature. IFGRP is a mechanism that allows the aggregation of a network interface into one logical unit. Combining links aids in network availability and bandwidth. NetApp provides three types of IFGRPs for network port aggregation and redundancy:

- Single mode
- Static multimode
- Dynamic multimode

Dynamic multimode IFGRPs are recommended due to the increased reliability and error reporting and also because of their compatibility with vPCs. A dynamic multimode IFGRP uses LACP to group multiple interfaces together to act as a single logical link. This provides intelligent communication between the storage controller and the Cisco Nexus switches and enables load balancing across physical interfaces as well as failover capabilities.

From a Fibre Channel perspective, SAN A (red in Figure 14) and SAN B (blue in Figure 14) fabric isolation is maintained across the architecture with dedicated FCoE channels and virtual interfaces. The design allocates Fibre Channel interfaces with SAN A and SAN B access for each controller in the HA pair.

Figure 14 Discrete Uplink Design with Clustered Data ONTAP



Private Cloud Architecture Principles

The Fast Track architecture attempts to achieve the principles, patterns, and concepts outlined in the Microsoft TechNet article "[Private Cloud Principles, Patterns, and Concepts](#)." Please refer to this article if you need clarification on why a particular design choice was made in the Fast Track management architecture. The introduction to the article provides a synopsis:


- A key goal is to allow IT organizations to utilize the principles and concepts described in the reference architecture for [Private Cloud](#) content set to offer Infrastructure as a Service (IaaS), allowing any workload hosted on this infrastructure to automatically inherit a set of cloud-like attributes. Fundamentally, the consumer should have the perception of infinite capacity and continuous availability of the services they consume. They should also see a clear correlation between the amount of services they consume and the price they pay for these services.
- Achieving this requires virtualization of all elements of the infrastructure, compute (processing and memory), network, and storage, into a fabric that is presented to the container or the virtual machine. It also requires the IT organization to take a service provider's approach to delivering infrastructure,

necessitating a high degree of IT service management maturity. Moreover, most of the operational functions must be automated to minimize the variance as much as possible while creating a set of predictable models that simplify management.

Private Cloud Reference Model

Infrastructure as a service (IaaS) is the application of private cloud architecture principles to deliver infrastructure. As the cloud ecosystem matures, product features and capabilities broaden and deepen. The reference model described in this section and shown in [Figure 15](#) is used as a guide for delivering a holistic solution that spans all the layers required for mature IaaS. The model acts as a guide to assist architects in their efforts to holistically address the development of a private cloud architecture. This model is a reference only. Some elements are emphasized more than others in the technical reference architecture, and that preference is based on experience operating private clouds in real-world environments.

Figure 15 *Private Cloud Reference Model*



The reference model is split into the following layers:

- The software, platform, and infrastructure layers represent the technology stack. Each layer provides services to the layer above.
- The service operations and management layers represent the process perspective and include the management tools required to implement the process.
- The service delivery layer represents the alignment between business and IT.

This reference model is a deliberate attempt to blend the technology and process perspectives, because cloud computing is as much about service management as it is about the technologies involved in it. For examples, see [ITIL](#) and [MOF](#). For further reading, please see [Private Cloud Reference Model](#).

Private Cloud Management Overview

Fabric Management

As we discuss later in the "Management Architecture" section, fabric management involves treating discrete capacity pools of servers, storage, and networks as a single fabric. Key capabilities of the fabric management system include:

- Hardware integration
- Fabric provisioning
- Virtual machine and application provisioning
- Resource optimization
- Health and performance monitoring
- Maintenance
- Reporting

Process Automation and Orchestration

The orchestration layer managing the automation and management components 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 virtual machine 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 that combines complex workflows with events and activities across multiple management-system components and forms an end-to-end IT business process. The orchestration layer must provide the ability to design, test, implement, and monitor these IT workflows.

Service Delivery

Service Management System

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

User Self-Service

Self-service capability is a characteristic of private cloud computing and must be present in any implementation. The intent is to permit users to approach a self-service capability and be presented with options available for provisioning. The capability may be basic, provisioning of a virtual machine with a predefined configuration; more advanced, allowing configuration options to the base configuration; or complex, when implementing a platform capability or service.

Self-service capability is a critical business driver that allows members of an organization to become more agile in responding to business needs with IT capabilities that align and conform to internal business and IT requirements. The interface between IT and the business should be abstracted to a well-defined, simple, and approved set of service options. The options should be presented as a menu in a portal or available from the command line. The business can select these services from the catalog, start the provisioning process, and be notified upon completion, at which point they are charged only for the services actually used.

Service Catalog

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

- Classes of services that are available
- Requirements to be eligible for each service class
- Service-level attributes and targets included with each service class
- Cost models for each service class

The service catalog might also include specific virtual machine templates 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 underutilization. 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 risks will be managed in a disaster scenario to help make sure minimum service levels are maintained. The principles of resiliency and automation are fundamental here.

Service Level Management

Service-level management (SLM) is the process of negotiating 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. SLM also 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 identifying a business need, then moves to managing a business relationship, and concludes when that service becomes available. Service strategy drives service design. After launch, the service is transitioned to operations and refined through continual service improvement. A service provider's approach is critical to successful service lifecycle management.

Operations

Change Management

Change management controls 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 the services. Change management focuses on understanding and balancing the cost and risk of making the change versus the potential benefit of the change to the business or the service. Providing predictability and minimizing human involvement are the core principles for achieving a mature service management process and making sure changes can be made without affecting the perception of continuous availability.

Incident and Problem Management

Incident management quickly resolves events that affect, or threaten to affect, services with minimal disruption. Problem management identifies and resolves the root causes of incidents. It also tries to prevent or minimize the impact of possible incidents.

Configuration Management

Configuration management involves making sure that the assets required to deliver services are properly controlled. The goal is to have accurate and effective information about those assets available when and where it is needed. This information includes details about asset configuration and the relationships between assets.


Configuration management typically requires a CMDB, which is used to store configuration records throughout their lifecycle. The configuration management system maintains one or more CMDBs, and each CMDB stores attributes of configuration items and relationships to other configuration items.

Management Architecture

Fabric and Fabric Management

At a high level, the Fast Track architectures include the concepts of a computing, storage, and network fabric. This fabric is logically and physical independent of components such as System Center that provide management of the fabric, that is, fabric management ([Figure 16](#)).

Figure 16 **High-Level Diagram of Fast Track Architecture**



Fabric

The fabric is defined as all of the physical and virtual resources under the scope of management within the fabric management infrastructure. The fabric is typically the entire computing, storage, and network infrastructure, usually implemented as Hyper-V host clusters managed by the System Center infrastructure.

For private cloud infrastructures, the fabric constitutes a resource pool that consists of one or more scale units. In a modular architecture, a scale unit is the point to which a module in the architecture can scale before another module is required. For example, an individual server is a scale unit because it can be expanded to a certain point in terms of CPU and RAM; however, once it reaches its maximum scalability, an additional server is required to continue scaling. Each scale unit also has an associated amount of physical installation and configuration labor. With large-scale units, such as a preconfigured full rack of servers, the labor overhead can be minimized.

It is critical to know the scale limits of all components, both hardware and software, when determining the optimum scale units for the overall architecture. Scale units allow the documentation of all the requirements needed for implementation, including space; power; heating, ventilation and air conditioning (HVAC); and connectivity.

Fabric Management

Fabric management involves treating discrete capacity pools of servers, storage, and networks as a single fabric. The fabric is then subdivided into capacity clouds, or resource pools, that carry characteristics such as delegation of access and administration, SLAs, and cost metering. Fabric management allows the centralization and automation of complex management functions that can be carried out in a highly standardized, repeatable fashion to increase availability and lower operational costs.

Fabric Management Host Architecture


In a private cloud infrastructure, it is recommended that the systems that make up the resource pools be physically separate from the systems that provide management. Much like the concept of having a top-of-rack switch, this separation is recommended to provide separate fabric management hosts to manage the underlying services that provide capacity to the private cloud infrastructure. This model helps make sure that the availability of the fabric is separated from fabric management and, regardless of the state of the underlying fabric resource pools, management of the infrastructure and its workloads is maintained at all times.

To support this level of availability and separation, Fast Track private cloud architectures should contain a separate set of hosts, a minimum of two, configured as a failover cluster in which the Hyper-V role is enabled. Furthermore, these hosts should contain highly available virtualized instances of the management infrastructure, System Center, to support fabric management operations that are stored on a dedicated Continuously Available SMB 3.0 share.

All management hosts will use the Windows Server 2012 Datacenter Edition operating system with the Hyper-V role enabled. For the specified scalability, the supporting System Center products and their dependencies will run within Hyper-V virtual machines on the management hosts.

For enterprise implementations, a minimum two-node fabric management cluster is required, with four nodes recommended for scale and availability, to provide high availability of the fabric management workloads. This fabric management cluster is dedicated to the virtual machines running the suite of products providing IaaS management functionality and is not intended to run additional customer workloads outside of those that provide management capabilities over the fabric infrastructure. For additional management scale points, additional management host capacity might be required. The host architecture is illustrated in Figure 17.

Figure 17 Management Fabric Infrastructure



Management Host Computing (CPU)

The management virtual machine workloads are expected to have a fairly high level of utilization. A conservative virtual CPU to logical processor ratio of two or fewer should be used. This ratio implies a minimum of two sockets per fabric management host, with six to eight cores per socket. During maintenance or failure of one of the two nodes, this CPU ratio will be temporarily exceeded.

The following recommendation is provided for each fabric management host within the configuration:

- Minimum 12 logical CPUs and 96 virtual CPUs

Management Host Memory (RAM)

Host memory should be sized accordingly to support the System Center products and their dependencies providing IaaS management functionality. [Table 5](#) lists recommendations for each fabric management host within the configuration.

Table 5 *Recommendations for Fabric Management Hosts*

Management Fabric Cluster	Host Memory
2-node	256-GB RAM recommended

Management Host Network

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, provided that two or more physical adapters can be teamed for redundancy and multiple vNICs and VLANs can be presented to the hosts for traffic segmentation and bandwidth control. 10 Gigabit Ethernet or higher network interfaces must be used to reduce bandwidth contention and simplify the network configuration through consolidation.

Management Host Storage Connectivity

The requirement for storage is simply that shared storage is provided with sufficient connectivity, but no particular storage technology is required. The following guidance is provided to assist with storage connectivity choices. For storage attached directly to the host, an internal SATA or SAS controller is required (for boot volumes), unless the design features booting from SAN for the host operating system. Depending on the storage device used, the following adapters are required to allow shared storage access:

- If using SMB 3.0 file shares, two or more 10 Gigabit Ethernet NICs or converged network adapters (CNAs)
- If using FC SAN connections, two or more HBAs
- If using iSCSI, two or more 10 Gigabit Ethernet NICs or HBAs
- If using FCoE, two or more 10 Gigabit Ethernet CNAs

Management Host Storage


The management components require three types of storage:

- FCoE LUNs for booting Windows Server from SAN
- Continuously Available SMB3.0 share for the management virtual machines
- vFCP LUNs for the virtualized SQL Server cluster

Management Logic Architecture

[Figure 18](#) depicts the management logical architecture if using a dedicated three-node management cluster.

Figure 18 Management Architecture Using a Dedicated Three-Node Management Cluster



The management architecture consists of a minimum of two physical nodes in a failover cluster with shared storage and redundant network connections. This architecture provides a highly available platform for the management systems. Some management systems have additional highly available options, and in these cases, the most effective highly available option will be used.

The management systems include:

- Two SQL Servers in a guest cluster configuration
- Two System Center 2012 SP1 VMM servers in a guest cluster configuration
- Two System Center 2012 SP1 Operations Manager management servers using the built-in failover and redundancy features (up to four management servers may be required for agent-managed monitoring of up to 8000 virtual machines)
- One System Center 2012 SP1 Operations Manager Reporting server
- Two System Center 2012 SP1 Orchestrator servers using the built-in failover and redundancy features
- Two System Center 2012 SP1 Service Manager Management servers
- One System Center 2012 SP1 Service Manager Data Warehouse
- One System Center 2012 SP1 Service Manager Self-Service Portal
- One System Center 2012 SP1 App Controller server
- One deployment server providing Windows Deployment Services (WDS) and Windows Server Update Services (WSUS) (optional)
- Two Cisco Nexus 1000V for Hyper-V VSMs

Management Systems Architecture

This section outlines the management systems architecture and its dependencies within a customer environment.

System Center Component Scalability

The System Center 2012 SP1 product comprises several components that have differing scale points. In order to deploy the System Center suite to support a Fast Track private cloud installation, these requirements must be normalized across components.

Table 6 provides guidance on the scalability of each component.

Table 6 **Scalability of System Center 2012 Components**

Component	Scalability Reference	Notes
Virtual Machine Manager (VMM)	800 hosts and 25,000 virtual machines per instance	A VMM instance is defined as a standalone or cluster installation. While not required, scalability is limited to 5000 virtual machines when Service Provider Foundation (SPF) is installed. A single SPF installation can support up to five VMM instances.
App Controller	Scalability is proportional to VMM.	Supports 250 virtual machines per VMM user role.
Operations Manager	3000 agents per management server, 15,000 agents per management group, 50,000 agentless managed devices per management group.	
Orchestrator	Simultaneous execution of 50 runbooks per runbook server.	
Service Manager	Large deployment supports up to 20,000 computers.	Topology dependent. Note that in Fast Track Service Manager is used solely for private cloud virtual machine management. An advanced deployment topology can support up to 50,000 computers.

As shown by the component scalability in the table, the default Fast Track deployment can support the management of up to 8000 virtual machines and associated fabric hosts, based on the deployment of a single 64-node Windows Server 2012 Hyper-V failover cluster. Note that individual components such as Operations Manager can be scaled further to support larger and more complex environments. In these cases, a four-node management cluster would be required to support scale.

Prerequisite Infrastructure

Active Directory Domain Services (AD DS)

AD DS is a required foundational component. Fast Track provides support for Windows Server 2008, Windows Server 2008 R2, and Windows Server 2012 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, but this is not a strict requirement. A dedicated resource forest or domain may also be employed as an additional part of the deployment. Fast Track does support multiple domains or multiple forests in a trusted environment using two-way forest trusts.
- **Trusts:** Fast Track allows multidomain support within a single forest in which two-way forest (Kerberos) trusts exist between all domains. This is referred to as multidomain or interforest support.

Domain Name System (DNS)

DNS name resolution is a required element for System Center 2012 SP1 components and the process automation solution. AD DS integrated DNS is required for automated provisioning and deprovisioning components within Orchestrator Runbook as part of the solution. The solution provides full support and automation for Windows Server 2008, Windows Server 2008 R2, and Windows Server 2012 AD DS integrated DNS deployments.

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

Dynamic Host Configuration Protocol (DHCP)

To support dynamic provisioning and management of physical and virtual computing capacity within the IaaS infrastructure, 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 have known IP addresses while providing centralized management of those addresses through DHCP.

Windows DHCP is required for automated provisioning and deprovisioning components within Orchestrator runbooks as part of the solution. DHCP is used to support host cluster provisioning, DHCP reservations, and other areas supporting dynamic provisioning of computing within the infrastructure. The solution provides full support and automation for Windows Server 2008, Windows Server 2008 R2, and Windows Server 2012 versions of the DHCP server role. Use of solutions outside of the Windows DHCP server role requires additional testing and validation.

Consolidated SQL Server Design

Under System Center 2012 SP1, the support matrix for the various versions of SQL Server has been simplified. System Center 2012 SP1 supports SQL Server 2008 R2 and SQL Server 2012 fully, with limited component support for SQL Server 2008.

[Table 7](#) provides a compatibility matrix.

Table 7 **Compatibility of System Center 2012 SP1 Components with SQL Server**

Component	SQL Server 2008 R2	SQL Server 2012
App Controller	Release to manufacturing (RTM) or later	RTM or later
Operations Manager	SP1 or later	RTM or later
Orchestrator	SP1 or later	RTM or later
Service Manager	SP1 or later	RTM or later
Virtual Machine Manager	SP1 or later	RTM or later

To support advanced availability scenarios and more flexible storage options, SQL Server 2012 is required for Fast Track deployments of fabric management. Two SQL Server 2012 virtual machines must be deployed as a guest failover cluster to support the solution, with an option to scale to a four-node cluster. This multinode SQL Server failover cluster will contain all the databases for each System Center product in discrete instances by product and function. This separation of instances allows for division by unique requirements and scaling over time as the needs of each component grow.

**Note**

Not all features are supported for failover cluster installations. Some features cannot be combined on instances, and some allow configuration only at initial installation.

As a general rule, database engine services and analysis services will be hosted in separate instances within the failover cluster. Because of the support for SQL Server Reporting Services (SSRS) in a failover cluster, SSRS will be installed on the hosting System Center component server, the Operations Manager Reporting Server. This installation, however, will be "files only," and the SSRS configuration will configure remote reporting services databases hosted on the component instance on the SQL cluster. The exception to this is the System Center Operations Manager Analysis Services and Reporting Services configuration. For this instance, Analysis Services and Reporting Services must be installed with the same server and with the same instance to support VMM and Operations Manager integration. All instances are required to be configured with Windows authentication. In System Center 2012 SP1, the App Controller and Orchestrator components can share an instance of SQL Server with the SharePoint farm, providing additional consolidation of the SQL instance requirements.

[Table 8](#) outlines the options required for each instance.

Table 8 Database Instances and Requirements

Fabric Management Component	Instance Name (Suggested)	Components	Collation ¹	Storage Requirements ²
Virtual Machine Manager	SCVMM DB	Database engine	SQL_Latin1_General_CP1_CI_AS	2 LUNs
Windows Server Update Services (optional)	SCVMM DB	Database engine	SQL_Latin1_General_CP1_CI_AS	NA - Shared instance with VMM
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	Latin1_General_100_CI_AS	2 LUNs
Service Manager Data Warehouse	SCSMDW	Database engine, full-text search	Latin1_General_100_CI_AS	2 LUNs
	SCSMAS	Analysis services	Latin1_General_100_CI_AS	2 LUNs
Service Manager Web Parts and Portal	SCDB	Database engine	SQL_Latin1_General_CP1_CI_AS	NA - Shared instance with Orchestrator and App Controller
Orchestrator	SCDB	Database engine	SQL_Latin1_General_CP1_CI_AS	2 LUNs
App Controller	SCDB	Database engine	SQL_Latin1_General_CP1_CI_AS	NA - Shared instance with Orchestrator and Service Manager portal

¹ The default SQL collation settings are not supported for multilingual installations of the Service Manager component. Use the default SQL collation only 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).

² Note that additional LUNs may be required for TempDB management in larger-scale configurations.


Each SQL Server 2012 virtual machine will be configured with four virtual CPUs, at least 16 GB of RAM (32 GB is recommended for large-scale configurations), and three virtual NICs (for LAN, cluster communications, and iSCSI). Each SQL Server VM will access iSCSI-based shared storage with two LUNs configured for each hosted database. Should the needs of the solution exceed what two SQL Server VMs are able to provide, additional VMs can be added to the virtual SQL Server cluster and each SQL Server instance moved to its own VM in the cluster. The Fast Track configuration requires the use of SQL Server 2012 Enterprise Edition using the latest service pack and cumulative update on a dedicated Windows Server 2012 failover cluster. In addition, SSD storage should be used, if possible, to provide the necessary I/O for these databases. The instances and associated recommended node placement are outlined in [Figure 19](#).

**Note**

The default SQL collation settings are not supported for multilingual installations of the Service Manager component. Use the default SQL collation only 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).

**Note**

The additional LUNs may be required for TempDB management in larger-scale configurations.

Figure 19 Configuration of System Center SQL Instances**Note**

For a more detailed version of this diagram, see the [Appendix](#).

SQL Server Configuration

- Two highly available virtual machines (optional third or fourth node for reserve capacity and failover)
- Windows Server 2012 Datacenter
- SQL Server 2012 Enterprise Edition (current service pack and cumulative update)
- One 60-GB VHDX per SQL Server VM on the host cluster's SMB 3.0 file shares
- Eight virtual CPUs per SQL Server VM
- 16 GB memory (32 GB recommended; do not enable Dynamic Memory)
- Two vNICs (one for client connections, one for cluster communications)
- Two vFC HBAs for SQL Server database LUN access
- Storage: One operating system VHDX per SQL Server VM and 15 dedicated cluster LUNs (14 LUNs for System Center and 1 LUN for quorum)

SQL Server Data Locations

[Table 9](#) lists the locations of SQL Server data.

Table 9 SQL Server Data Locations

LUN	Component(s)	Instance Name	Purpose	Size
LUN 1/2	Service Manager Management	SCSMDB	Instance database and logs	145 GB/70 GB
LUN 3/4	Service Manager Data Warehouse	SCSMDW	Instance database and logs	1 TB/ 500 GB
LUN 5/6	Service Manager Analysis Service	SCSMAS	Instance database and logs	8 GB/4 GB
LUN 7/8	Service Manager SharePoint Farm Orchestrator App Controller	SCDB	Instance database and logs	10 GB/5 GB
LUN 9/10	Virtual Machine Manager Windows Server Update Services	SCVMMDB	Instance database and logs	6 GB/3 GB

LUN 11/12	Operations Manager	SCOMDB	Instance database and logs	130 GB/65 GB
LUN 13/14	Operations Manager Data Warehouse	SCOMDW	Instance database and logs	1 TB/ 500 GB
LUN 15	N/A	N/A	SQL Server failover cluster quorum	1 GB

Note that the Operations Manager and Service Manager database sizing assumes a managed infrastructure of 8000 virtual machines. Additional references for sizing are provided in the component sections below.

Virtual Machine Manager

System Center 2012 SP1 VMM is required. Two VMM servers are deployed and configured in a failover cluster, using a dedicated SQL Server instance on the virtualized SQL Server cluster. One library share on the VMM servers will be utilized. Additional library servers can be added as needed. The VMM and Operations Manager integration is configured during the installation process. The following hardware configurations will be used:

Servers

- Two guest clustered virtual machines
- Windows Server 2012
- Four virtual CPUs
- 8 GB memory
- Two vNICs (one for client connections, one for cluster communications)
- Storage: One operating system VHDX, one data VHDX (pass-through volume, iSCSI LUN or virtual Fibre Channel LUN)

Operations Manager

System Center 2012 SP1 Operations Manager is required. A minimum of two Operations Manager servers are deployed in a single management group, using a dedicated SQL Server instance on the virtualized SQL Server cluster. An Operations Manager agent is required to be installed on every management host and scale unit cluster node to support health monitoring functionality. Additionally, agents may be installed on every guest VM to provide guest-level monitoring capabilities. 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 Server instance on the virtualized SQL Server cluster. The installation will follow a split SQL Server configuration: SQL Server Reporting Services and Operations Manager components will reside on the Operations Manager VM, while the SQL Server Reporting Services and Operations Manager databases will use a dedicated instance on the virtualized SQL Server cluster. Note that for the Fast Track implementation the data warehouse is sized for 90-day retention instead of the default retention period.

Estimated SQL Server Database Sizes

- 130-GB Operations Manager database
- 1-TB Operations Manager data warehouse database

Operations Manager Management Servers

- Two highly available virtual machines (if monitoring up to 8000 agent-managed VMs, up to 4 Operations Manager management servers are required)
- Windows Server 2012
- Eight virtual CPUs
- 16-GB memory
- One vNIC
- Storage: One operating system VHDX

Operations Manager Reporting Server

- One highly available virtual machine
- Windows Server 2012
- Eight virtual CPUs
- 16-GB memory (if monitoring up to 8000 agent-managed VMs, up to 32-GB management servers are required)
- One vNIC
- Storage: One operating system VHDX

**Note**

Additional guidance on database and data warehouse sizing for Operations Manager can be found at <http://blogs.technet.com/b/momteam/archive/2012/04/02/operations-manager-2012-sizing-helper-tool.aspx>. Further guidance is provided at <http://blogs.technet.com/b/kevinholman/archive/2010/01/05/understanding-and-modifying-data-warehouse-retention-and-grooming.aspx> and <http://blogs.technet.com/b/momteam/archive/2008/05/14/data-warehouse-data-retention-policy-dwdata-rp-exe.aspx>.

Management Packs

In addition to the management packs required for Virtual Machine Manager and Operations Manager integration, the following management packs should be included as part of the Fast Track design:

- Virtual Machine Manager 2012 SP1 management packs and prerequisites as required for VMM and Operations Manager integration
- Server OEM third-party management packs

Service Manager

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 Server instance on the virtualized SQL Server cluster. The Service Manager portal is hosted on a fourth virtual machine with the portal. Note that for the Fast Track implementation the data warehouse is sized for 90-day retention instead of the default retention period.

**Note**

Additional guidance on database and data warehouse sizing for Service Manager can be found at <http://go.microsoft.com/fwlink/p/?LinkID=232378>. Additional guidance is provided at <http://blogs.technet.com/b/servicemanager/archive/2009/09/18/data-retention-policies-aka-grooming-in-the-servicemanager-database.aspx>.

Service Manager Management Servers

- Two highly available virtual machines
- Windows Server 2012
- Four virtual CPUs
- 16-GB memory
- One vNIC
- Storage: One operating system VHDX

Service Manager Data Warehouse Server

- One highly available virtual machine
- Windows Server 2012
- Eight virtual CPUs
- 16-GB memory
- One vNIC
- Storage: One operating system VHDX

Service Manager Portal Server

- One highly available virtual machine
- Windows Server 2008 R2 SP1
- Eight virtual CPUs
- 16-GB memory
- One vNIC
- Storage: One operating system VHDX

Service Manager Estimated SQL Server Database Sizes

- 145-GB Service Manager database
- 1-TB Service Manager data warehouse database

Orchestrator

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

Use two Orchestrator runbook servers for high availability and scale purposes. Orchestrator provides built-in failover capability, but 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 state and allowing an instance of a workflow to keep running. Orchestrator assures only that

it will start any workflows that were started on the failed server. State may, and likely will, be lost, which means a request might fail. Many workflows have some degree of state management built in that helps mitigate this risk.

The other reason two Orchestrator servers are deployed by default 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.

Orchestrator Servers

- Two non-high availability virtual machines
- Windows Server 2012
- Four virtual CPUs
- 8-GB memory
- One vNIC
- Storage: One operating system VHDX

App Controller

System Center 2012 SP1 App Controller is optional. However, if the Service Manager portal is utilized, App Controller must also be installed. App Controller uses a dedicated SQL Server instance on the virtualized SQL Server cluster. A single App Controller uses App Controller server that is installed on the management host cluster.

Service Manager provides the service catalog and service request mechanism. Orchestrator provides the automated provisioning. App Controller provides the user interface for connecting and managing workloads post-provisioning.

Application Controller Server

- One highly available virtual machine
- Windows Server 2012
- Four virtual CPUs
- 8-GB memory
- One vNIC
- Storage: One operating system VHDX

Cisco Nexus 1000V for Hyper-V Virtual Supervisor Module (VSM)

The Cisco Nexus 1000V for Hyper-V VSM provides the user interface for creating and managing the Nexus 1000V for Hyper-V configuration. Once the configuration is created or updated, it is imported by VMM and applied to logical switches that are configured with the Nexus 1000V extension. The Nexus 1000V VSM has built-in high availability capability and runs in Active/Passive mode with its partner VSM.

Cisco Nexus 1000V for Hyper-V VSM

- Two highly available virtual machines
- Cisco NX-OS

- One virtual CPU
- 4-GB memory
- Three vNICs
- Storage: One operating system VHDX

Fabric Management Requirement Summary

Table 10 summarizes the fabric management virtual machine requirements of the System Center component that supports the product or operating system role.

Table 10 Virtual Machine Requirements of System Center Components

Component Roles	Virtual CPU	RAM (GB)	Virtual Hard Disk (GB)
SQL Server cluster node 1	8	16	60
SQL Server cluster node 2	8	16	60
Virtual Machine Manager	4	8	60
Virtual Machine Manager	4	8	60
App Controller	4	8	60
Operations Manager management server	8	16	60
Operations Manager supplemental management server	8	16	60
Operations Manager reporting server	8	16	60
Orchestrator runbook server	4	8	60
Orchestrator supplemental runbook server	4	8	60
Service Manager management server	4	16	60
Service Manager supplemental management server	4	16	60
Service Manager portal	8	16	60
Service Manager data warehouse	8	16	60
Windows Deployment Services/Windows Server Update Services	2	4	60
0. 1 Nexus 1000V for Hyper-V VSM 1	1	4	4
Nexus 1000V for Hyper-V VSM 1	1	4	4
Totals	88	196 GB	908 GB

Management Scenarios

Listed below are the primary management scenarios addressed in Fast Track, although the management layer can provide many more capabilities:

- Fabric management

- Fabric provisioning
- Virtual machine provisioning and deprovisioning
- IT service provisioning (including platform and application provisioning)
- Fabric and IT service maintenance
- Fabric and IT service monitoring
- Resource optimization
- Service management
- Reporting (used by chargeback, capacity, service management, health, and performance)
- 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. The various methods outlined in the sections that follow make fabric management possible.

Hardware Integration

Hardware integration refers to the management system being able to perform deployment or operational tasks directly against the underlying physical infrastructure, such as storage arrays, network devices, and servers.

Storage Integration and Management

In VMM, 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.

SAN Integration

To activate the storage features, VMM uses the Windows Storage Management API (SMAPI) to manage external storage using symmetric multiprocessing (SMP), or uses SMAPI together with the Microsoft standards-based storage management service to communicate with Storage Management Initiative - Specification (SMI-S) compliant storage. The Microsoft standards-based storage management service is an optional server feature that allows communication with SMI-S storage providers. It is activated during installation of System Center 2012 SP1. The Data ONTAP SMI-S Agent is installed on a VM that the VMM management server can access over the network by IP address or by fully qualified domain name (FQDN) to enable the management of the NetApp storage array.

NetApp SMB 3.0 Storage Integration

Clustered Data ONTAP provides support for using SMB 3.0 file shares as shared storage for Hyper-V 2012. System Center 2012 SP1 allows you to assign SMB file shares to Hyper-V standalone hosts and clusters via the included SMI-S extension.

Network Integration and Management

Networking in VMM includes several enhancements that allow administrators to efficiently provision network resources for a virtualized environment. The networking enhancements include the following:

Logical Networks

System Center 2012 SP1 allows you to easily connect virtual machines to a network that serves a particular function in your environment, for example, the "back-end," "front-end," or "backup" network. To do this, associate IP subnets and, if needed, VLANs together into named units called logical networks. You can design your logical networks to fit your environment.

Load Balancer Integration

Networking in VMM includes load-balancing integration so that you can automatically provision load balancers in your virtualized environment. Load-balancing integration works together with other network enhancements in VMM.

By adding a load balancer to VMM, you can load balance requests to the virtual machines that make up a service tier. You can use Microsoft Windows Network Load Balancing (NLB), or you can add supported hardware load balancers through the VMM console. NLB is included when you install VMM. NLB uses round robin as the load-balancing method.

To add supported hardware load balancers, you must install a configuration provider that is available from the load balancer manufacturer. The configuration provider is a plug-in to VMM that translates VMM PowerShell commands to API calls that are specific to a load balancer manufacturer and model.

Switches and Ports

VMM in System Center 2012 SP1 allows you to consistently configure identical capabilities for network adapters across multiple hosts by using port profiles and logical switches. Port profiles and logical switches act as containers for the properties or capabilities that you want your network adapters to have. Instead of configuring individual properties or capabilities for each network adapter, you can specify the capabilities in port profiles and logical switches, which you can then apply to the appropriate adapters. This can simplify the configuration process.

Virtual Machine Networks

Virtual machine networks offer the ability to use network virtualization that extends the concept of server virtualization to make it possible for you to deploy multiple virtual networks (VM networks) on the same physical network. However, VM networks can be configured in multiple ways:

- **Network virtualization (Hyper-V network virtualization):** If you wish to support multiple tenants, also called clients or customers, with their own networks, isolated from the networks of others, use network virtualization. To do this, create a logical network, and on top of that logical network create multiple VM networks, each of which uses the option to isolate using Hyper-V network virtualization. With this isolation, your tenants can use any IP addresses that they want for their virtual machines, regardless of the IP addresses that are used on other VM networks. Also, you can allow your tenants to configure some aspects of their own networks, based on limits that you specify.

**Note**

If using network virtualization and the virtual machines require network communication outside of the private subnet, you will need to provide a gateway. See "How to Add a Gateway in System Center 2012 SP1," <http://technet.microsoft.com/en-us/library/jj614618.aspx>.

- VLAN-based configuration: If you are working with networks that use familiar VLAN technology for network isolation, you can manage those networks as they are, using VMM to simplify the management process.

Virtual Switch Extension Management

VMM 2012 SP1 can use a vendor-provided network-management console and the VMM management server together. You can configure settings or capabilities in the vendor-provided network-management console, also known as the management console for a forwarding extension, and then use the console and the VMM management server in a coordinated way.

To do this, you must first install the provider software that is provided by the vendor on the VMM management server. Then you can add the virtual switch extension manager to VMM, which will cause the VMM management server to connect to the vendor network-management database and import network settings and capabilities from that database. The result is that you can see those settings and capabilities, and all your other settings and capabilities, together in VMM.

Fabric Provisioning

In accordance with the principle of standardization and automation, creating the fabric and adding capacity should be an automated process. There are multiple scenarios for adding fabric resources in VMM. This section is specifically referring to bare-metal provisioning of Hyper-V hosts and host clusters. In VMM, this is achieved through a multistep process:

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

Each step in this process has dependencies:

- Provisioning Hyper-V hosts
 - A Preboot Execution Environment (PXE) boot server
 - Dynamic DNS registration
 - A standard base image to be used for Hyper-V hosts
 - Hardware driver files in the VMM library
 - A host profile in the VMM library
 - A baseboard management controller (BMC) on the physical server
- Configuring host properties, networking, and storage
 - Host property settings
 - The storage integration from above plus addition MPIO and/or iSCSI configuration
 - For the 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

- Create Hyper-V host clusters
 - The hosts must meet all requirements for Windows Server failover clustering
 - The hosts must be managed by VMM

VMM Private Clouds

When you have configured the fabric resources, you can subdivide and allocate them for self-service consumption through 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 might want to create a cloud for use by the finance department. You will be able to:

- Name the cloud
- Scope it to one or more host groups
- Specify which network capabilities are available to the cloud
- Specify which storage classifications are available to the cloud
- Select which library shares are available to the cloud for virtual machine storage
- Specify granular capacity limits to the cloud (virtual CPU, memory, storage, 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.

Virtual Machine Provisioning and Deprovisioning

One of the primary cloud attributes is user self-service capability. In this solution, self-service capability refers to the ability for the user to request one or more VMs or to delete one or more of their existing VMs. 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 an automated process or workflow in the infrastructure through VMM to create or delete a VM based on the input from the user or tenant. Provisioning can be template based, such as requesting a small, medium, or large VM template, or it can be a series of selections made by the user. If authorized, the provisioning process could create a new VM per the user's request, add the VM to any relevant management products in the private cloud, and allow access to the VM by the requestor.

To facilitate this capability, the administrator needs to have preconfigured the following VMM items:

- VMM library resources, such as virtual hard disks
- Networking components, such as logical networks and load balancers
- Hyper-V hosts and host groups
- Private clouds
- Hardware profiles, guest operating system profiles, application profiles, and SQL Server profiles
- VM templates

IT Service Provisioning

In VMM, a service is a set of virtual machines that are configured, deployed, and managed as a single entity. An example would be the deployment of a multitier line-of-business application with front-end, middle, and data tier virtual machines.

In the VMM console, use the service template designer to create a service template that defines the configuration of the service. The service template includes information about the VMs that are deployed as part of the service, including which applications to install on the VMs and the networking configuration needed for the service.

Resource Optimization

Elasticity, perception of infinite capacity, and perception of continuous availability are the 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 onto the smallest number of hosts for a higher consolidation ratio.

Dynamic Optimization

Based on user settings, VMM dynamic optimization migrates virtual machines to perform resource balancing within host clusters that support live migration. Two or more Hyper-V hosts are required in a host cluster to allow dynamic optimization.

Dynamic optimization looks to correct the three possible scenarios listed below in priority order:

1. Virtual machines that have configuration problems on their current host
2. Virtual machines that are causing their host to exceed configured performance thresholds
3. Unbalanced resource consumption on hosts

Power Optimization

VMM power optimization is an optional feature of dynamic optimization, and 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 BMC (Baseboard Management Controller) that allows out-of-band management.

Power optimization makes sure that the cluster maintains a quorum if an active node fails. For clusters created outside of VMM and added to VMM, power optimization requires more than four nodes. For each additional one or two nodes in a cluster, one node can be powered down. For instance:

- One node can be powered down for a cluster of five or six nodes
- Two nodes can be powered down for a cluster of seven or eight nodes

- Three nodes can be powered down for a cluster of nine or ten nodes

When VMM creates a cluster, it creates a quorum disk and uses that disk as part of the quorum model. For clusters created by VMM, power optimization can be set up for clusters of more than three nodes. This means that the number of nodes that can be powered down is as follows:

- One node can be powered down for a cluster of four or five nodes
- Two nodes can be powered down for a cluster of six or seven nodes
- Three nodes can be powered down for a cluster of eight or nine nodes

Fabric and IT Service Maintenance

A private cloud solution must provide the ability to perform maintenance on any component without affecting the availability of the solution. Examples include the need to update or patch a host server or add additional storage to the SAN. The system should not generate unnecessary alerts or events in the management systems during planned maintenance.

VMM supports on-demand compliance scanning and remediation of the fabric. Fabric servers include the following physical computers managed by VMM: Hyper-V hosts and Hyper-V clusters, library servers, PXE servers, the WSUS server, and the VMM management server. Administrators can monitor the update status of the servers. In addition, they can scan for compliance and remediate updates for selected servers and 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 the 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 state for the VMs.

The use of this feature requires either a dedicated WSUS server integrated with VMM or an existing WSUS server from a Configuration Manager environment.

If you use an existing WSUS server from a Configuration Manager environment, changes to configuration settings for the WSUS server—for example, update classifications, languages, and proxy settings—should be made only from the Configuration Manager. The VMM administrator can view the configuration settings from the VMM console, but cannot make changes.

Fabric and IT Service Monitoring

A private cloud solution must provide the ability to monitor every major component of the solution and generate alerts based on performance, capacity, and availability metrics. Examples of availability metrics include monitoring server availability, CPU, and storage utilization. Monitoring of the fabric is performed through the integration of Operations Manager and VMM. Enabling this integration allows Operations Manager to automatically discover, monitor, and report on essential performance and health characteristics of any object managed by VMM:

- Health and performance of all VMM managed hosts and virtual machines
- Diagram views in Operations Manager reflecting all of VMM's deployed hosts, services, virtual machines, private clouds, IP address pools, and storage pools
- Performance and resource optimization, which can 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 Manager agent within the virtual machine operating system and install the desired management pack. Be aware that this scenario would not be considered fabric monitoring.

Reporting

A private 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-based, 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 that help IT focus efforts

As a result of VMM and Operations Manager integration, several reports are created and available by default. However, metering and chargeback reports and incident and problem reports are enabled by the use of Service Manager.

VMM and Operations Manager Integration Default Reports

[Table 11](#) lists default reports created through the integration of VMM and Operations Manager.

Table 11 **Default Reports Available Through VMM and Operations Manager Integration**

Report	Description
Capacity utilization	Details usage for VM hosts and other objects. It provides an overview of how capacity is being used in your data center. This information can inform decisions about how many systems you need to support your VMs.
Chargeback	Provides information to calculate chargeback to cost centers for VMs. You can set up this report by cost center grouping to summarize CPU, memory, disk, and network usage for VMs within your cost centers. NOTE: The cost center is a property of VMs that can also be set on VM templates.
Host group forecasting	Predicts host activity based on disk space, memory, disk I/O, network I/O, and CPU usage history.
Host utilization	Shows the number of VMs running on each host and average usage, along with total or maximum values for host processors, memory, and disk space.
Host utilization growth	Shows the percentage change in resource usage and the number of VMs running on selected hosts during a specified time period.
Power savings	Shows how much power is saved through power optimization. You can view total hours of processor power saved for a date range and host group, as well as detailed information for each host in a host group. For more information about power optimization, see Configuring Dynamic Optimization and Power Optimization in Virtual Machine Manager .
SAN usage forecasting	Predicts SAN usage based on history.
Virtual machine allocation	Provides information about allocation of VMs.
Virtual machine utilization	Provides information about resource utilization by VMs, including average usage and total or maximum values for VM processors, memory, and disk space.
Virtualization candidates	Helps identify physical computers that are good candidates for conversion to VMs. You can use this report to identify little-used servers and display average values for a set of commonly requested performance counters for CPU, memory, and disk usage, along with hardware configurations, including processor speed, number of processors, and total RAM. You can limit the report to computers that meet specified CPU and RAM requirements, and sort the results by selected columns in the report.

**Note**

You can also design your own reports.

Service Management System

The goal of System Center 2012 Service Manager is to support IT service management in a broad sense. This includes implementing ITIL and MOF processes, such as change and incident management, and it can also include processes like allocating resources from a private cloud.

Service Manager maintains a CMDB. The CMDB is the repository for nearly all configuration- and management-related information in the System Center 2012 environment. For the System Center Cloud Services Process Pack, this information includes VMM resources such as VM templates, and VM service templates, which are all copied regularly from the VMM library into the CMDB. This allows objects such as VMs and users to be tied to Orchestrator runbooks for automated tasks such as request fulfillment, metering, and chargeback.

User Self-Service

The Microsoft Private Cloud Self-Service Solution consists of the:

- Service Manager self-service portal with the Cloud Services Process Pack
- App Controller

Service Manager 2012 SP1 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, when a user wants to create a virtual machine in the group's cloud, instead of passing the request directly on to VMM, as the App Controller does, Service Manager starts an Orchestrator workflow to handle the request. The workflow contacts the user's manager to get an approval for the request. If the request is approved, the workflow then starts an Orchestrator 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 co-located on a single dedicated server.

The [Cloud Services Process Pack](#) is an add-on component that allows 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, which includes the ability to define projects, capacity pools, and virtual machines
- Natively supported request, approval, and notification to allow businesses to effectively manage their own allocated infrastructure capacity pools

App Controller is the portal that a self-service user would utilize after a request has been fulfilled in order to connect to and manage his or her virtual machines and services. App Controller connects directly to VMM, using the credentials of the authenticated user to display his or her virtual machines and services, and to provide a configurable set of actions.

Service Management

The service management layer provides the means for automating and adapting IT service management best practices, found in MOF 4.0 and ITIL, to provide built-in processes for incident resolution, problem resolution, and change control.

MOF provides relevant, practical, and accessible guidance for today's IT professionals. It is a downloadable framework that encompasses the entire service management lifecycle ([Figure 20](#)). For more information about MOF, see [MOF 4.0](#).

Figure 20 MOF 4.0 Model



Operations Manager also has the ability to integrate with Visual Studio Team Foundation Server. Streamlining the communications between development and IT operations teams, often called dev-ops, can help decrease the time it takes for the application maintenance and delivery to move into production stage, where your application delivers value to customers. To speed interactions between these teams, it is essential to quickly detect and fix problems that might need assistance from the engineering team. For more information, see <http://technet.microsoft.com/en-us/library/jj614609.aspx>.

Security


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

IT infrastructure threat modeling is the practice of considering what attacks might be attempted against the different components in an IT infrastructure (Figure 21). 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 vulnerability
- 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 MOF, a framework that provides practical guidance for managing IT practices and activities throughout the entire IT lifecycle. The **effective 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 and 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 make sure 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 21 **Security Threat Modeling**

For more information on threat modeling, see the following resources:

- [IT Infrastructure Threat Modeling Guide](#)
- [Security Risk Management Guide](#)

Security for Microsoft private cloud is founded on three pillars: protected infrastructure, application access, and network access, as described in the sections that follow

Protected Infrastructure

A defense-in-depth strategy is used 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 crosses a trust boundary. Any portions of an IT infrastructure in which data or processes cross 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 to disinfect computers, restore files, and reenter or re-create 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 a strong way to counter these threats. The least-privileged user account (LUA) approach is an important part of that defensive strategy. The LUA approach directs users to follow the principle of least privilege and log in with limited user accounts. This strategy also aims to limit the use of administrative credentials to administrators for administrative tasks only.

Application Access

AD DS provides the means to manage the identities and relationships that make up a Microsoft private cloud. Integrated with Windows Server 2008 R2 and Windows Server 2012, AD DS provides the functionality needed to centrally configure and administer system, user, and application settings.

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

Network Access

Windows Firewall with Advanced Security combines a host firewall and IP Security (IPsec). Unlike a perimeter firewall, Windows Firewall with Advanced Security runs on each computer, running a particular version of Windows, and provides local defense from network attacks that might pass through your perimeter network or originate inside your organization. It also contributes to 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 includes an API 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. This involves create a logical network inside an existing physical network in which computers share a common set of requirements for more 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.

Desktop management and security have traditionally existed as two separate disciplines, yet both play central roles in helping to keep users safe and productive. Management provides 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 SP1 Endpoint Protection (formerly known as Forefront Endpoint Protection 2012) aligns these two work streams into a single infrastructure.

Key Features

System Center 2012 SP1 Endpoint Protection makes it easier to help protect critical desktop and server operating systems against viruses, spyware, rootkits, and other threats.

- 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 large organizations around the globe. By using Configuration Manager distribution points and an automatic software deployment model, organizations can quickly deploy updates without relying on WSUS.
- **Highly accurate and efficient threat detection:** The anti-malware engine in System Center 2012 SP1 Endpoint Protection helps protect against the latest malware and rootkits, with a low false-positive rate, and helps to keep employees productive with scanning that has a low impact on performance.
- **Behavioral threat detection:** System Center 2012 SP1 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 SP1 Endpoint Protection helps prevent exploitation of endpoint vulnerabilities with deep protocol analysis of network traffic.
- **Automated agent replacement:** System Center 2012 SP1 Endpoint Protection automatically detects and removes common endpoint security agents, to lower the time and effort needed to deploy new protection.
- **Windows Firewall management:** System Center 2012 SP1 Endpoint Protection makes sure that Windows Firewall is active and working properly to help protect against network-layer threats. It also allows administrators to more easily manage protections across the enterprise.

Service Delivery Layer

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

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

With these questions in mind, IT departments must address two main issues within the service layer:

- How do we provide a cloudlike 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. See the section "Private Cloud Architecture Principles," earlier in this document, for more information on these principles.

Figure 22 *Service Delivery Layer of Dynamic Data Center Model*



The components of the service delivery layer are as follows:

- **Financial management:** Financial management incorporates the functions and processes used to meet a service provider's budgeting, accounting, metering, and charging requirements. The primary financial management concerns 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, and includes the 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 with predictable capacity management is 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 capabilities of the required services and partner with the IT department on solution procurement. The business will also need to work closely with the IT department to define future capacity requirements to continue to adhere 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 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 critical to successful service lifecycle management.
- **Service-level management:** Service-level management is the process of negotiating SLAs and making sure they 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 to achieve the perception of infinite capacity and continuous availability. This, too, requires IT departments to implement a service provider's approach.
- **Continuity and availability management:** Availability management defines processes necessary to achieve the perception of continuous availability. Continuity management defines how risks will be managed in a disaster scenario to help make sure that 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 underutilization. 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 multi-tenancy requirements are important factors to consider during this process.

Operations

The operations layer defines the operational processes and procedures necessary to deliver IT as a service (Figure 23). This layer uses IT service management concepts that can be found in prevailing best practices such as ITIL and MOF.

The main focus of the operations layer is to carry out the business requirements defined at the service delivery layer. Cloudlike service attributes cannot be achieved through technology alone; mature IT service management will be required. The operations capabilities are common to all three services are IaaS, platform as a service (PaaS), and software as a service (SaaS).

Figure 23 **Operations Layer of Dynamic Data Center Model**



The components of the operations layer include the following:

- **Change management:** Change management is responsible for controlling the lifecycle of all changes. The 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 potential benefits to the business or service. Offering 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 maintains 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, current, and expected future states, and it should 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 involves 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. Predictability and minimal human involvement in the release and deployment process are critical 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 are 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 the 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 a minimal impact on service availability. Resilient design promotes rapid restoration of service continuity. Predictability and minimal human involvement are necessary to achieve this resiliency.

- **Request fulfillment:** The goal of request fulfillment is to manage user requests for services. As the IT department 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 toward 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 multi-tenancy 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 promoting predictability. The vast majority of systems administration tasks should be automated.

Conclusion

FlexPod with Microsoft Private Cloud is the optimal shared infrastructure foundation on which to deploy a variety of IT workloads. Cisco and NetApp have created a platform that is both flexible and scalable for multiple use cases and applications. One common use case is to deploy Windows Server 2012 with Hyper-V as the virtualization solution, as described in this document. From virtual desktop infrastructure to Microsoft Exchange Server, Microsoft SharePoint Server, Microsoft SQL Server, and SAP, FlexPod can efficiently and effectively support business-critical applications running simultaneously from the same shared infrastructure. The flexibility and scalability of FlexPod also enable customers to start out with a right-sized infrastructure that can ultimately grow with and adapt to their evolving business requirements.

Appendix

Validated Bill of Materials

The following product information is provided for reference and will require modification depending on specific customer environments. Considerations include optics, cabling preferences, application workload, and performance expectations.



Note

FlexPod architecture can support numerous Cisco and NetApp configurations that are not covered in this Cisco UCS bill of materials. Please work with your Cisco, NetApp, or partner account team if you need assistance.

Table 12 NetApp Platform Bill of Materials

Product	Description	Qty
DS2246-1011-24S-R	DSK SHLF,24x600GB,10K,6Gb SAS,IOM6,-C,R5	4
FAS-V32XX-CHASS	FAS/V32XX,Chassis,AC PS,-C,R6	2
FAS3210AE-IB-BAS	FAS3240 HA System with Controller & IOXM	2
SW-3240A-SNAPM	SW,SnapMgr Application Integration, 3240A,-C	2
SW-3210A-ONTAP8	SW,Data ONTAP Essentials,3240A,-C	2
X1107A-R6-C	NIC 2-Port Bare Cage SFP+ 10GbE PCIe,-C	2
X5515A-R6-C	Rackmount Kit,4N2,DS14-Middle,-C,R6	1
X5526A-R6-C	Rackmount Kit,4-Post,Universal,-C,R6	1
X6536-R6-C	Cable,Cntlr-Shelf/Switch,5m,LC/LC,Op,-C	4
X6558-R6-C	Cable,SAS Cntlr-Shelf/Shelf-Shelf/HA,2m,-C	4
X6559-R6-C	Cable,SAS Cntlr-Shelf/Shelf-Shelf/HA,5m,-C	4
X6560-R6-C	Cable,Ethernet,0.5m RJ45 CAT6,-C	4
X6561-R6-C	Cable,Ethernet,2m RJ45 CAT6,-C	1
X6562-R6-C	Cable,Ethernet,5m RJ45 CAT6,-C	4
X800E-R6-C	Power Cable North America,-C,R6	12
DOC-32XX-C	Documents,32XX,-C	1
CS-O2-NOINSTALL	SupportEdge Premium 4hr Onsite, w/o Install - Mths:36	1
X6557-R6-C	Cable,SAS Cntlr-Shelf/Shelf-Shelf/HA,0.5m,-C	4
SW-FLASH-CACHE	SW, Flash Cache,-C	2
X1970A-R5-C	Flash Cache PCIe 256GB Module,-C	2
X2065A-R6-C	HBA SAS 4-Port Copper 3/6 Gb QSFP PCIe,-C	2
SW-3240A-FCP-C	SW,FCP,3240A,-C	2
SW-3240A-SRESTO	SW,SnapRestore,3240A,-C	2
SW-ISCSI-C	SW,iSCSI,-C	2


Table 13 Cisco Unified Computing Bill of Materials

Custom Name	SKU	Description	Qty
Blade Server UCSB-B200-M3	UCSB-B200-M3	UCS B200 M3 Blade Server w/o CPU, memory, HDD, mLOM/mezz	8
	UCS-CPU-E52660B	2.20 GHz E5-2660/95W 10C/25MB Cache/DDR3 1866MHz	16
	UCS-MR-1X162RZ-A	16GB DDR3-1866-MHz RDIMM/PC3-14900/dual rank/1.5v	128
	UCS-VIC-M82-8P	Cisco UCS VIC 1280 dual 40Gb capable Virtual Interface Card	8
	UCSB-MLOM-40G-01	Cisco UCS VIC 1240 modular LOM for M3 blade servers	8
	N20-BBLKD	UCS 2.5 inch HDD blanking panel	16
	UCSB-HS-01-EP	CPU Heat Sink for UCS B200 M3 and B420 M3	16
Chassis N20-C6508	N20-C6508	UCS 5108 Blade Svr AC Chassis/0 PSU/8 fans/0 fabric extender	1

	CAB-C19-CBN	Cabinet Jumper Power Cord, 0 VAC 16A, C20-C19 Connectors	4
	UCS-IOM-2208XP	UCS 2208XP I/O Module (8 External, 16 Internal 10Gb Ports)	2
	N01-UAC1	Single phase AC power module for UCS 5108	4
	N20-CAK	Access. kit for 5108 Blade Chassis including Railkit, KVM dongle	1
	N20-FAN5	Fan module for UCS 5108	8
	N20-PAC5-00W	00W AC power supply unit for UCS 5108	4
	N20-FW010	UCS 5108 Blade Server Chassis FW package	1
Fabric Interconnect UCS-FI-6248UP	UCS-FI-6248UP	UCS 6248UP 1RU Fabric Int/No PSU/32 UP/ 12p LIC	2
	UCS-ACC-6248UP	UCS 6248UP Chassis Accessory Kit	2
	UCS-PSU-6248UP-AC	UCS 6248UP Power Supply/100-240VAC	4
	N10-MGT010	UCS Manager v2.1	2
	CAB-9K12A-NA	Power Cord, 1VAC 13A NEMA 5-15 Plug, North America	2
	UCS-LIC-10GE	UCS 6200 Series ONLY Fabric Int 1PORT 1/10GE/FC-port license	20
	UCS-FAN-6248UP	UCS 6248UP Fan Module	4
	UCS-FI-DL2	UCS 6248 Layer 2 Daughter Card	2

Name	Catalog Num	Description	Qty
Nexus 5548UP	N5K-C5548UP-OSM.P	Nexus 5548UP Storage Solutions Bundle, Full Stor Serv Lic, OSM	2
	N55-48PO-SSK9.P	Nexus 5500 Storage License, 48 Ports, OSM	2
	N55-DL2.P	Nexus 5548 Layer 2 Daughter Card	2
	N55-M-BLNK.P	Nexus 5500 Module Blank Cover	2
	N55-PAC-750W.P	Nexus 5500 PS, 750W, Front to Back Airflow(Port-Side Outlet)	4
	N5548P-FAN.P	Nexus 5548P and 5548UP Fan Module, Front to Back Airflow	4
	CAB-C13-C14-2M.P	Power Cord Jumper, C13-C14 Connectors, 2 Meter Length	4
	N5548-ACC-KIT.P	Nexus 5548 Chassis Accessory Kit	2
	N5KUK9-602N1.2a.P	Nexus 5000 Base OS Software Rel 6.0(2)N1(2a)	2

Figure 24 Detailed Fast Track SQL Server Database Design Diagram



References

Cisco Unified Computing System:

www.cisco.com/en/US/products/ps10265/index.html

Cisco UCS 6200 Series Fabric Interconnects:

www.cisco.com/en/US/products/ps11544/index.html

Cisco UCS 5100 Series Blade Server Chassis:

www.cisco.com/en/US/products/ps10279/index.html

Cisco UCS B-Series Blade Servers:

www.cisco.com/en/US/partner/products/ps10280/index.html

Cisco UCS Adapters:

www.cisco.com/en/US/products/ps10277/prod_module_series_home.html

Cisco UCS Manager:

www.cisco.com/en/US/products/ps10281/index.html

Cisco Nexus 5000 Series Switches:

www.cisco.com/en/US/products/ps9670/index.html

Cisco Nexus 1000V for Hyper-V:

www.cisco.com/en/US/products/ps13056/index.html

Microsoft Private Cloud Fast Track

www.microsoft.com/en-us/server-cloud/private-cloud/fast-track.aspx

Microsoft System Center 2012 SP1

www.microsoft.com/en-us/server-cloud/system-center/default.aspx

Microsoft Windows Server 2012

www.microsoft.com/en-us/server-cloud/windows-server/default.aspx

Microsoft Windows Server 2012 with Hyper-V

www.microsoft.com/en-us/server-cloud/windows-server/server-virtualization.aspx

NetApp FAS3200 Series Storage System

www.netapp.com/us/products/storage-systems/fas3200/

NetApp OnCommand Management Software

www.netapp.com/us/products/management-software/index.aspx

NetApp SnapDrive for Windows

www.netapp.com/us/products/management-software/snapdrive-windows.aspx

NetApp SnapManager for Hyper-V

www.netapp.com/us/products/management-software/snapmanager-hyperv.aspx

NetApp SnapManager for SQL Server

www.netapp.com/us/products/management-software/snapmanager-hyperv.aspx

Interoperability Matrix

NetApp Interoperability Matrix Tool

<http://support.netapp.com/matix/>

Cisco UCS Hardware and Software Interoperability Tool

www.cisco.com/web/techdoc/ucs/interoperability/matrix/matrix.html