




# Cisco UCS Solution for EMC VSPEX Microsoft Private Cloud Fast Track 4.0 Design Guide

Last Updated: October 8, 2014



Building Architectures to Solve Business Problems



## About the Authors

### **Tim Cerling, Technical Marketing Engineer, Cisco**

Tim Cerling is a Technical Marketing Engineer with Cisco's Datacenter Group, focusing on delivering customer-driven solutions on Microsoft Hyper-V and System Center products. Tim has been in the IT business since 1979. He started working with Windows NT 3.5 on the DEC Alpha product line during his 19 year tenure with DEC, and he has continued working with Windows Server technologies since then with Compaq, Microsoft, and now Cisco. During his twelve years as a Windows Server specialist at Microsoft, he co-authored a book on Microsoft virtualization technologies - Mastering Microsoft Virtualization. Tim holds a BA in Computer Science from the University of Iowa.

### **Mike McGhee, Consulting Solutions Engineer, EMC**

Mike is a consulting solutions engineer with EMC Corporation based outside of Boston, Massachusetts. Mike has been with EMC for 15 years focusing on enterprise storage and Microsoft platforms. Currently Mike is focusing on System Center, Hyper-V and Windows Server 2012 integration with storage technologies, and how such technologies help to enable cloud infrastructures.

### **David Feisthammel, Consulting Solutions Engineer, EMC**

Dave is a consulting solutions engineer with EMC Corporation based in Bellevue, Washington, just blocks from the Microsoft headquarters campus. As a member of EMC's Microsoft Partner Engineering team, he focuses on Microsoft's enterprise hybrid cloud technologies, including Windows Server, Hyper-V, and System Center. Dave is an accomplished IT professional with progressive international and domestic experience in the development, implementation, and market launch of IT solutions and products. With nearly three decades of experience in Information Technology, he has presented, lectured, taught, and written on various topics related to systems management.

## Acknowledgments

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, we would like to thank:

Mike Mankovsky, Technical Leader Engineering, Cisco

## 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)

© 2014 Cisco Systems, Inc. All rights reserved



# Cisco UCS Solution for EMC VSPEX Microsoft Private Cloud Fast Track 4.0 Design Guide

---

## Overview

The Microsoft Private Cloud Fast Track program is a joint effort between Microsoft and its hardware partners such as Cisco and EMC. The goal of the program is to help organizations develop and implement private clouds quickly 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 compute, 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 utilizes the core capabilities of the Windows Server (OS), Hyper-V, and System Center to deliver a private cloud infrastructure as a service offering. These are also key software components that are used for every reference implementation.

## Private Cloud Fast Track Program Description

The Infrastructure as a Service Product Line Architecture (PLA) is focused on deploying virtualization fabric and fabric management technologies in Windows Server and System Center to support private cloud scenarios. This PLA includes reference architectures, best practices, and processes for streamlining deployment of these platforms to support private cloud scenarios.

This component of the IaaS PLA focuses on delivering core foundational virtualization fabric infrastructure guidance that aligns to the defined architectural patterns within this and other Windows Server 2012 R2 private cloud programs. The resulting Hyper-V infrastructure in Windows Server 2012 R2 can be leveraged to host advanced workloads, and subsequent releases will contain fabric management scenarios using System Center components. Scenarios relevant to this release include:



- Resilient infrastructure—Maximize the availability of IT infrastructure through cost-effective redundant systems that prevent downtime, whether planned or unplanned.
- Centralized IT—Create pooled resources with a highly virtualized infrastructure that supports maintaining individual tenant rights and service levels.
- Consolidation and migration—Remove legacy systems and move workloads to a scalable high-performance infrastructure.
- Preparation for the cloud—Create the foundational infrastructure to begin transition to a private cloud solution.

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.

Each reference architecture in the Fast Track program combines concise guidance with validated configurations for the compute, network, storage, and virtualization layers. Each architecture presents multiple design patterns for enabling the architecture, and each design pattern describes the minimum requirements for validating each Fast Track solution.

The Cisco and EMC Fast Track Solution presented here is an Open solution. The Cisco and EMC with Microsoft Private Cloud Fast Track solution utilizes the core capabilities of Windows Server 2012 R2, Hyper-V, and System Center 2012 R2 to deliver a Private Cloud - Infrastructure as a Service offering. The key software components of every Reference Implementation are Windows Server 2012 R2, Hyper-V, and System Center 2012 R2. The solution also includes software from Cisco and EMC to form a complete solution that is ready for your enterprise.

## Business Value

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

### Reduced risk

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

### Lower cost of ownership

- A cost-optimized, platform and software-independent solution for rack system integration
- High performance and scalability with Windows Server 2012 R2 operating system and Hyper-V
- Minimized backup times and fulfilled recovery time objectives for each business critical environment
- Highly scalable to address changing business needs
  - Eight Cisco UCS B200 M3 servers as the initial configuration. Expand by adding another chassis with eight more Cisco UCS B200 M3 servers, or other B-series blades, depending on specific needs

- EMC VNX 5400 with 75 x 300GB disks as the initial configuration. Expand by adding more disks, storage array, or different VNX family storage array, depending upon specific needs

## Technical Benefits

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

### **Faster deployment:**

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

## Program Requirements and Validation

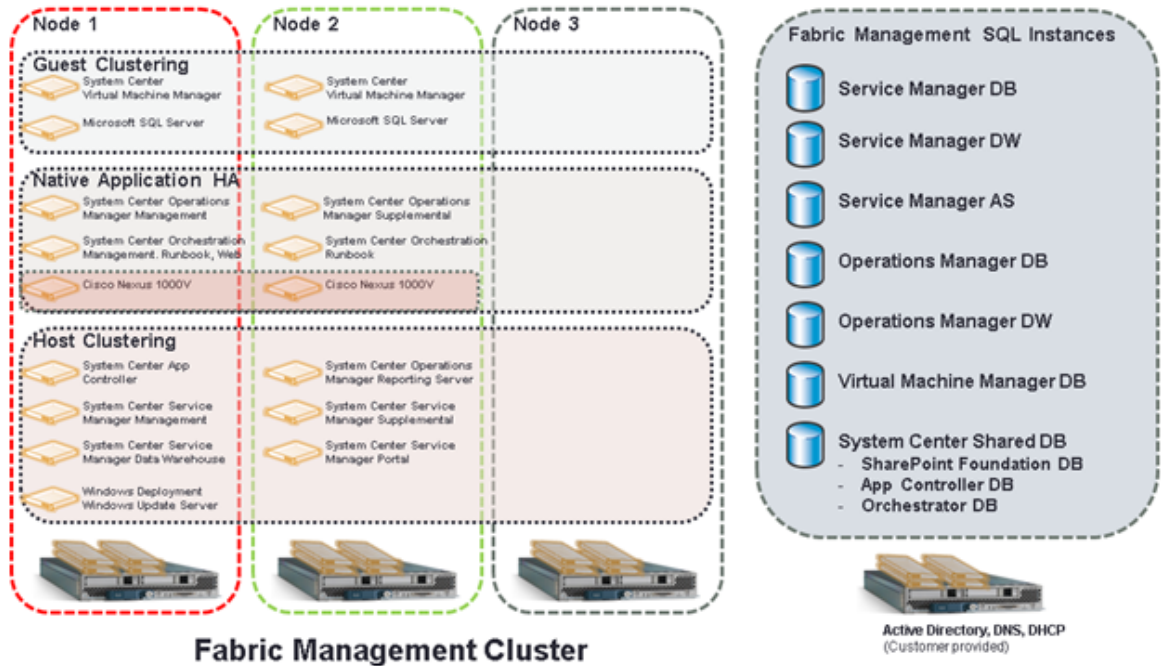
The Microsoft Private Cloud Fast Track program is comprised of three pillars; Engineering, Marketing and Enablement. These three pillars drive the creation of Reference Implementations, making them public and finally making them available for customers to purchase. This Reference Architecture is one step in the "Engineering" phase of the program and towards the validation of a Reference Implementation.

## Core Fast Track Infrastructure

The Cisco and EMC solution is based on a converged infrastructure. Fabric Management VMs are hosted directly on a compute fabric cluster. Other workload VMs are deployed on a second compute cluster. Additional failover clusters can be configured for additional VMs, or the second cluster can be expanded up to 64 nodes. This solution starts with the minimal number of System Center component servers recommended in order to provide full functionality in a production environment. System Center components can be scaled-out by adding additional servers running the particular management component, if needed.

This design pattern is introduced for Fabric Management to include a dedicated three-to-four node Hyper-V failover cluster to host the fabric management virtual machines. This design pattern utilizes both scaled-out and highly available deployments of the System Center components to provide full functionality in a production environment. In addition to the System Center components running as virtual machines, Cisco deploys a pair of Cisco Nexus 1000V virtual machines to handle network management for the VMs.

Figure 1 Private Cloud Fabric Management Infrastructure



The Fabric Management cluster is configured in such a manner to ensure maximum availability of all components of the environment. Each Cisco UCS B200 M3 server is configured with sufficient memory to support the running of all the listed virtual machines illustrated above. By provisioning a third node, the environment retains its highly available capability even during those periods of time when the host nodes are individually taken down for maintenance. For example, in the above figure, even if Node 3 was down for maintenance, a catastrophic failure of Node 2 would not prevent all the virtual machines from continuing to run on Node 1.

## Architecture

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

The Cisco UCS solution validated with Microsoft Private Cloud includes EMC VNX5400 storage, Cisco Nexus 5500 Series network switches, the Cisco Unified Computing Systems (Cisco UCS) platforms, and Microsoft virtualization software in a single package. The computing and storage can fit in one data center rack with networking residing in a separate rack or deployed according to a customer's data center design. Due to port density, the networking components can accommodate multiple configurations of this kind.



Figure 2 Implementation Overview




Figure 2 contains the following components:

- Cisco UCS 5108 chassis each with eight Cisco UCS B200 M3 Blade servers, dual Intel E5-2660 V2 2.50 GHz processors, 256 GB memory, 1240 Virtual Interface Card
- Two Cisco UCS 2108 fabric extenders per chassis
- Two Cisco UCS 6248UP Fabric Interconnects
- Two Cisco Nexus 5548UP Switches
- 10 GE and 8 Gb FC connections
- EMC VNX5400 Unified Platform
- 75 x 300 GB SAS disks

Storage is provided by an EMC VNX5400 storage array with accompanying disk shelves. All systems and fabric links feature redundancy, providing for end-to-end high availability (HA configuration within a single chassis). For server virtualization, the deployment includes Microsoft Hyper-V. While this is the default base design, each of the components can be scaled flexibly to support the specific business requirements in question. For example, more (or different) blades and chassis could be deployed to increase compute capacity, additional disk shelves or SSDs could be deployed to improve I/O capacity and throughput, or special hardware or software features could be added to introduce new features.

## Cisco Unified Computing System

The Cisco Unified Computing System (Cisco UCS) combines Cisco UCS B-Series Blade Servers and C-Series Rack Servers with networking and storage access in a single converged system that simplifies management and delivers greater cost efficiency and agility with increased visibility and control. The Cisco UCS B200 M3 Blade Server delivers performance, versatility, and density without compromise. It addresses the broadest set of workloads, from IT and web infrastructure, through distributed database.

## Cisco UCS B200 M3 Servers

The enterprise-class Cisco UCS B200 M3 Blade Server further extends the capabilities of the Cisco UCS portfolio in a half-width blade form factor. The Cisco UCS B200 M3 harnesses the power of the latest Intel® Xeon® processor E5-2600 and E5-2600 v2 product families, with up to 768 GB of RAM (using 32-GB DIMMs), two disk drives, and up to dual 4x 10 Gigabit Ethernet throughput. In addition, Cisco UCS has the architectural advantage of not having to power and cool excess switches in each blade chassis. With a larger power budget per blade server, Cisco can design uncompromisable expandability and capabilities in its blade servers, as evidenced by the new Cisco UCS B200 M3, with its leading memory slot and drive capacity.

*Figure 3 Cisco UCS B200 M3 Blade Server*



The Cisco UCS B200 M3 provides the following:

- One or two, multi-core, Intel® Xeon® processor E5-2600 and E5-2600 v2 product families CPUs, for up to 24 processing cores
- 24 DIMM slots for industry-standard double-data-rate 3 (DDR3) memory running up to 1866 MHz and up to 768 GB of total memory (using 32-GB DIMMs)
- Two optional, hot-pluggable SAS or SATA hard disk drives (HDDs) or solid-state drives (SSDs)
- Industry-leading 80 Gbps throughput bandwidth
- Remote management through a Cisco Integrated Management Controller (CIMC) that implements policy established in Cisco UCS Manager
- Out-of-band access by remote keyboard, video, and mouse (KVM) device, Secure Shell (SSH) Protocol, and virtual media (vMedia) as well as the Intelligent Platform Management Interface (IPMI)

In addition, the Cisco UCS B200 M3 is a half-width blade ([Figure 3](#)). Up to eight of these high-density, two-socket blade servers can reside in the 6RU Cisco UCS 5108 Blade Server Chassis, offering one of the highest densities of servers per rack unit in the industry.

## Cisco UCS 5108 Chassis

The Cisco UCS 5100 Series Blade Server Chassis is a crucial building block of the Cisco Unified Computing System, delivering a scalable and flexible blade server chassis for today's and tomorrow's data center while helping reduce TCO. The Cisco UCS 5108 Blade Server Chassis, is six rack units (6RU) high and can mount in an industry-standard 19-inch rack. A chassis can house up to eight half-width Cisco UCS B-Series Blade Servers and can accommodate both half- and full-width blade form factors.

Four single-phase, hot-swappable power supplies are accessible from the front of the chassis. These power supplies are 92 percent efficient and can be configured to support non-redundant, N+1 redundant, and grid-redundant configurations. The rear of the chassis contains eight hot-swappable fans, four power

connectors (one per power supply), and two I/O bays for Cisco UCS 2208XP Fabric Extenders. A passive midplane provides up to 80 Gbps of I/O bandwidth per server slot and up to 160 Gbps of I/O bandwidth for two slots. The chassis is capable of supporting future 40 Gigabit Ethernet standards.

**Figure 4** Cisco UCS 5108 Blade Server Chassis with Blade Servers Front and Back



The Cisco UCS 5108 Blade Server Chassis revolutionizes the use and deployment of blade-based systems. By incorporating unified fabric, integrated, embedded management, and fabric extender technology, Cisco Unified Computing System allows the chassis to use fewer physical components, has no need for independent management, and enables greater energy efficiency than traditional blade server chassis. This simplicity eliminates the need for dedicated chassis management and blade switches, reduces cabling, and enables Cisco Unified Computing System to scale to 20 chassis without adding complexity. The Cisco UCS 5108 chassis is a critical component in delivering the Cisco UCS benefits of data center simplicity and IT responsiveness.


In addition, the Cisco UCS 5108 chassis has the architectural advantage of not having to power and cool excess switches in each chassis. With a larger power budget per blade server, Cisco can design uncompromisable expandability and capabilities in its blade servers, as evidenced by the new Cisco UCS B200 M3 Blade Server, with its leading memory slot and drive capacity, and the Cisco UCS B230 M2 Blade Server, with its industry-leading 32-DIMM memory density in a half-width form factor.

## Cisco UCS 1240 Virtual Interface Card

A Cisco innovation, the Cisco UCS Virtual Interface Card (VIC) 1240 is a 4-port 10 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE)-capable modular LAN on motherboard (mLOM) designed exclusively for the M3 generation of Cisco UCS B-Series Blade Servers. When used in combination with an optional Port Expander, the Cisco UCS VIC 1240 capabilities can be expanded to eight ports of 10 Gigabit Ethernet.

The Cisco UCS VIC 1240 enables a policy-based, stateless, agile server infrastructure that can present up to 256 PCIe standards-compliant interfaces to the host that can be dynamically configured as either network interface cards (NICs) or host bus adapters (HBAs). In addition, the Cisco UCS VIC 1240 supports Cisco Data Center Virtual Machine Fabric Extender (VM-FEX) technology, which extends the Cisco UCS fabric interconnect ports to virtual machines, simplifying server virtualization deployment.

**Figure 5** Cisco UCS Virtual Interface Card (VIC) 1240



The personality of the card is determined dynamically at boot time using the service profile associated with the server. The number, type (NIC or HBA), identity (MAC address and World Wide Name [WWN]), failover policy, bandwidth, and quality-of-service (QoS) policies of the PCIe interfaces are all determined using the service profile. The capability to define, create, and use interfaces on demand provides a stateless and agile server infrastructure.

## Cisco UCS 6248UP Fabric Interconnect

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

**Figure 6** Cisco UCS 6248UP 48-Port Fabric Interconnect



The Cisco UCS 6248UP 48-Port Fabric Interconnect is a one-rack-unit (1RU) 10 Gigabit Ethernet, FCoE and Fiber Channel switch offering up to 960-Gbps throughput and up to 48 ports. The switch has 32 1/10-Gbps fixed Ethernet, FCoE and FC ports and one expansion slot (not required for this Fast Track solution).

## Cisco Nexus 5548UP Switch

The Cisco Nexus 5548UP switch, using cut-through architecture, supports line-rate 10 Gigabit Ethernet on all ports while maintaining consistently low latency independent of packet size and services enabled. It supports a set of network technologies known collectively as Data Center Bridging (DCB) that increases the reliability, efficiency, and scalability of Ethernet networks. These features allow the switches to support multiple traffic classes over a lossless Ethernet fabric, thus enabling consolidation of LAN, SAN, and cluster environments. Its ability to connect Fibre Channel over Ethernet (FCoE) to native Fibre Channel protects existing storage system investments while dramatically simplifying in-rack cabling.

In addition to supporting standard 10 Gigabit Ethernet network interface cards (NICs) on servers, the Cisco Nexus 5000 Series integrates with multifunction adapters called converged network adapters (CNAs) that combine the functions of Ethernet NICs and Fibre Channel host bus adapters (HBAs), making the transition to a single, unified network fabric transparent and consistent with existing practices, management software, and OS drivers. The switch series is compatible with integrated transceivers and Twinax cabling solutions that deliver cost-effective connectivity for 10 Gigabit Ethernet to servers at the rack level, eliminating the need for expensive optical transceivers. The Cisco Nexus 5000 Series portfolio also provides the flexibility to connect directly to servers using 10GBASE-T connections or fiber with Enhanced Small Form-Factor Pluggable (SFP+) transceivers.

**Figure 7** Cisco Nexus 5548UP Switch



The Cisco Nexus 5548UP is a 1RU 10 Gigabit Ethernet, Fibre Channel, and FCoE switch offering up to 960 Gbps of throughput and up to 48 ports. The switch has 32 unified ports and one expansion slot (not required for this Fast Track solution).

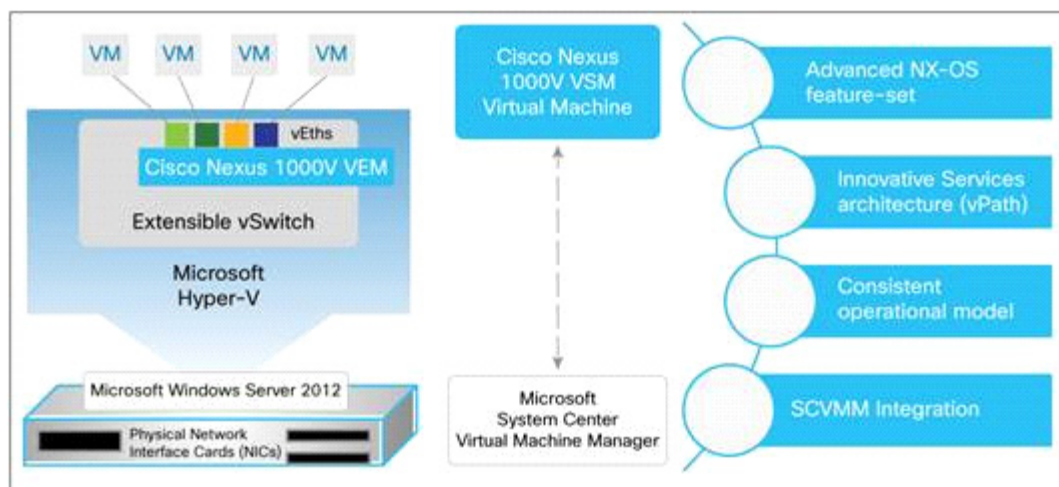
## Cisco Nexus 1000V

Cisco Nexus1000V Switch provides a comprehensive and extensible architectural platform for virtual machine and cloud networking. The switch is designed to accelerate server virtualization and multitenant cloud deployments in a secure and operationally transparent manner.

The Cisco Nexus 1000V Switch for Microsoft Hyper-V is a distributed software switching platform for Microsoft Windows Server environments. It provides:

- Advanced Cisco NX-OS Software feature set and associated partner ecosystem
- Innovative network services architecture to support scalable, multitenant environments
- Consistent operating model across physical and virtual environments and across hypervisors
- Tight integration with Microsoft System Center Virtual Machine Manager

**Figure 8** Cisco Nexus 1000V Switch for Microsoft Hyper-V Components



The Cisco Nexus 1000V Switch brings the robust architecture associated with traditional Cisco physical modular switches to Microsoft Hyper-V environments. The solution has two main components.

- The Cisco Nexus 1000V virtual Ethernet module (VEM) is a software component deployed on each Microsoft Hyper-V host as a forwarding extension. Each virtual machine on the host is connected to the VEM through virtual Ethernet (vEth) ports.
- The Cisco Nexus 1000V virtual supervisor module (VSM) is the management component that controls multiple VEMs and helps in the definition of virtual machine-focused network policies. It is a virtual machine running Cisco NX-OS on a Microsoft Hyper-V host and is similar to the supervisor module on a physical modular switch.

In addition to the VEM and VSM, Cisco Nexus 1000V Switches include Cisco vPath technology and provide a scalable, multitenant network services infrastructure for Microsoft Hyper-V environments.

The Cisco Nexus 1000V uses the extensible switch framework offered by Microsoft Windows Server 2012 with Hyper-V and the management ecosystem offered by Microsoft SCVMM and thus provides a transparent operating experience for Microsoft Hyper-V environments.

## EMC

### VNX5400

EMC VNX storage arrays implement a modular architecture that integrates hardware components for block, file, and object with concurrent support for NAS, iSCSI, Fibre Channel, and FCoE protocols. The series delivers file (NAS) functionality via X-blade data movers and block (iSCSI, FCoE, and FC) storage via dual storage processors leveraging full 6 Gb SAS disk drive topology. The system leverages the patented MCx™ multicore storage software operating environment that delivers unparalleled performance efficiency.

The new VNX series is EMC's next generation of midrange-to-enterprise products. It utilizes EMC's VNX Operating Environment for Block and File that you can manage with one easy to use GUI. This VNX software environment offers significant advancements in efficiency, simplicity, and performance.

The VNX series is designed specifically to unleash the power of Flash - designed from the ground up to maximize the performance of Flash, offering the lowest latency as well as the highest performance with the ability to add HDDs, offering the overall lowest \$/GB.

The new MCx architecture takes full advantage of the Intel Xeon E5-2600 (Sandy Bridge) multicore CPUs and allows cache and back-end processing software to scale in a linear fashion. In addition, other hardware improvements, such as the 6 Gb/s x 4 lanes SAS back end with an x 8 wide (6 Gb/s x 8 lanes) high-bandwidth option, PCIe Gen 3 I/O modules, and increased memory integrate to deliver:

- Up to 4X more file transactions
- Up to 4X more Oracle and SQL OLTP transactions
- Up to 4X more virtual machines
- Up to 3X more bandwidth for Oracle and SQL data warehousing

MCx is a combination of Multi-Core Cache (MCC), Multi-Core RAID (MCR), and Multi-Core FAST Cache, which enables the array to fully leverage the new Intel multi-core CPU architecture. MCx is a re-architecture of the core Block OE stack within the new VNX to ensure optimum performance at high scale with little to no customer intervention in an adaptive architecture. MCx delivers core functionality improvements that make the new VNX platform more robust, more reliable, more predictable, and easier to use. The MCx architecture is a true enterprise-ready storage platform.

MCx is designed to scale performance across all cores and sockets. It can scale down to 2 cores and all the way up to 40 cores and beyond. Along with the ability to scale across processor cores, significant improvements have been made in IO path latency.

Symmetric Active/Active allows clients to access a Classic LUN (not supported on pool LUNs) simultaneously through both SPs for improved reliability, ease-of management and improved performance. Since all paths are active, there is no need for the storage processor "trespass" to gain access to the LUN "owned" by the other storage processor on a path failure, eliminating application timeouts. The same is true for an SP failure, as the SP merely picks up all of the I/Os from the host through the alternate "optimized" path.

The EMC FAST™ Suite is an unrivaled set of software that tiers data across heterogeneous drives and boosts the most active data to cache. Highly active data is served from up to 4.2 terabytes of Flash drives with FAST Cache, which dynamically absorbs unpredicted spikes in system workloads. As that data ages and becomes less active over time, FAST VP (Fully Automated Storage Tiering for Virtual Pools) automatically tiers the data from high-performance to high-capacity drives in 256 megabyte increments, resulting in overall lower costs, independent of application type or data age.

The VNX series also offers file-level and block-level deduplication and compression. VNX block-based deduplication can dramatically lower the costs of the flash tier and is especially ideal for storage hosting Virtual machines and other environments with redundant data across multiple sources. VNX block-based compression, intended for relatively inactive LUNs such as backup copies and static data repositories, automatically compresses data, enabling customers to recapture capacity and reduce the data footprint by up to 50 percent. VNX file-level deduplication and compressions also reduces disk space by up to 50 percent by selectively compressing and deduplicating inactive files.

The new VNX series supports the SMB 3.0 protocol. SMB 3.0 protocol support is available with Microsoft Windows 8 and Microsoft Windows Server 2012 systems and has significant improvements over the previous SMB versions. While updating the SMB protocol, Microsoft focused on features that would have the most impact on server and cloud environments. Some of these features include:

- **SMB Multichannel:** Increases network performance in Hyper-V environments through network link aggregation, and provides fault tolerance to Hyper-V hosts by routing data over multiple network paths.
- **SMB Transparent Failover:** Provides applications a continuously available connection to storage by letting administrators configure clustered Windows file shares.
- **Branch Cache:** SMB 3.0 features Branch Cache improvements that better optimize bandwidth over wide-area network (WAN) connections between Branch Cache content servers and remote clients.
- **SMB Encryption:** Protects data on unsecured networks by encrypting in-flight data between the client and server.
- **Windows PowerShell Cmdlets:** Ease administration tasks by providing end-to-end command-line management of file shares.

These enhancements bring added flexibility and reliability to Windows Server 2012 and Hyper-V deployments that use SMB-based storage on EMC storage arrays. EMC fully supports SMB 3.0 within its unified storage platforms, such as the EMC VNX5400.

In addition to SMB 3.0 support, the new VNX series also supports the following:

- **Storage Management Initiative Specification (SMI-S):** SMI-S is the primary storage management interface for Windows Server 2012 and SCVMM. It can be used to provision storage directly from Windows Server 2012 and supports heterogeneous storage environments. It is based on the WBEM, CIM, and CIM/XML open standards.

- **Offloaded Data Transfer (ODX):** When a host performs a large data transfer within the same array, the data needs to travel from the array to the host and then back to the array. This is a waste of bandwidth and host resources. ODX is a SCSI token-based copy method that offloads data copies to the array, saving both bandwidth and host resources.
- **Thin Provisioning Support:** When a storage pool high water mark is met with a thin LUN, Server 2012 will be alerted and the event log will be updated. Server 2012 supports space reclamation when files are deleted from an NTFS formatted Thin LUN. Space is returned to the LUN in 256MB slice intervals.
- **N\_Port ID Virtualization (NPIV):** NPIV support has been added to the Unisphere Host Agent on Windows Server 2012. This allows multiple virtual FC initiators to use a single physical port. This enables Hyper-V VMs to have a dedicated path to the array, so push registrations from VMs are now allowed through the virtual port.

EMC virtual provisioning gives storage administrators flexibility in deploying storage to Windows Server 2012 and Hyper-V hosts. Virtual provisioning, which is also known as thin provisioning, lets storage administrators allocate storage on demand by presenting more storage to hosts than is physically available on the EMC storage array. As the operating system writes data to a thin-provisioned logical unit number (LUN), the storage array automatically allocates storage on-demand to handle the write requests.

One weakness of thin-provisioned storage is that once the storage array allocates storage to a thin-provisioned LUN, the storage remains allocated even if the host deletes data from the LUN. As time passes, LUN file-system utilization might remain constant even though storage allocation on the array continues to increase.

Windows Server 2012 can detect thin-provisioned storage on EMC storage arrays and reclaim unused space, including when Windows Server 2012 is deployed within a Hyper-V virtual machine. For example, suppose a storage administrator creates and assigns a 100 GB LUN to a host that is unaware of the LUNs thin-provisioned status. If the host writes 50 GB of data to the LUN, the storage array allocates 50 GB of storage to the LUN and removes that same amount from the storage pool. But if the host eventually deletes 10 GB of data from the LUN, the storage array will still keep 50 GB allocated to the LUN, even though 10 GB is now unused. With Windows Server 2012, the unused 10 GB of storage would automatically be reclaimed and returned to the pool, where it could be used by other applications.

With virtual provisioning detection, storage administrators can confidently deploy thin-provisioned LUNs knowing that unused storage will be returned to the storage pool, which keeps storage allocation in equilibrium with file-system utilization. This feature is especially important in large Hyper-V private cloud deployments or self-service Hyper-V private cloud environments where users continually create and delete virtual machines.

## Microsoft

### Windows Server 2012 R2 - Hyper-V

Microsoft includes its type-1 hypervisor, Hyper-V, as a role of the Windows Server operating system or as a no-cost down-load from their web site. If running Windows Server operating system environments, the physical host is licensed for the operating system and no additional cost is incurred for the hypervisor. The Hyper-V role enables you to create and manage a virtualized computing environment by using virtualization technology that is built in to Windows Server. 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 WMI provider, and other virtualization components such as the virtual machine bus (VMBus), virtualization service provider (VSP) and virtual infrastructure driver (VID).

The management tools for the Hyper-V role included with the system for no additional cost consist of:

- GUI-based management tools: Hyper-V Manager, a Microsoft Management Console (MMC) snap-in, and Virtual Machine Connection, which provides access to the video output of a virtual machine so you can interact with the virtual machine.
- 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 functionality not available through the GUI.

With each release of the Windows Server operating system, Microsoft has been enhancing existing features and adding new capabilities. The following table shows new and updated features available in the last two releases.

Feature	Description
Authorization	The Hyper-V Administrators group is introduced and is implemented as a local security group.
Automatic Virtual Machine Activation	VMs installed on a computer where Windows Server 2012 R2 is properly activated automatically activate, even in disconnected environments.
Client Hyper-V	Hyper-V available in a desktop operating system version of Windows.
Dynamic Memory	Support for configuring minimum memory and Smart Paging. Smart Paging provides a reliable restart experience for VMs configured with less minimum memory than startup memory.
Enhanced Session Mode	Redirection of local resources (e.g. display, audio, printers, clipboard, USB devices) to a Virtual Machine Connection session.
Export	Export VM checkpoint while VM is running.
Failover Clustering	Detect physical storage failures on storage devices not managed by Failover Clustering (SMB 3.0 file shares). Detect network connectivity for VMs.
Import	Import a virtual machine after copying the files manually, rather than exporting the virtual machine first.
Integration Services	Copy files to VM without network connection.
Linux Support	More built-in Linux Integration Services for newer distributions and more Hyper-V features are supported for Linux virtual machines. Features are part of the official Linux distribution.
Live Migration	Perform a live migration in a non-clustered environment, as well as perform more than one live migration at the same time and use higher network bandwidths. Improved performance with compression and SMB 3.0. Cross-version.
Networking	Many new features to support and enhance NVGRE, teaming, diagnostics, and extensions.
PowerShell	More than 160 cmdlets to manage Hyper-V, virtual machines, and virtual hard disks.
Quality of Service	Storage QoS enables management of storage throughput for virtual hard disks that are accessed by virtual machines.

Replica	Replicate virtual machines between storage systems, clusters, and data centers in two sites to provide business continuity and disaster recovery.
Resize Virtual Hard Disk	Resize virtual hard disks while the virtual machine is running.
Resource Metering	Track and gather data about physical processor, memory, storage, and network usage by specific virtual machines.
Scale and Resiliency	Larger compute and storage resources than was previously possible and improved handling of hardware errors.
Second Generation VM	Generation 2 VMs provide secure boot, boot from SCSI VHD or DVD, PXE boot from standard NIC, UEFI firmware support.
Shared Virtual Disk	Enables clustering virtual machines by using shared virtual hard disk (VHDX) files.
SMB Storage	In addition to storing on local disks, iSCSI, and Fibre Channel, Hyper-V supports use of SMB 3.0 file shares to provide storage for virtual machines.
Snapshots	After a virtual machine snapshot is deleted, the storage space that the snapshot consumed before being deleted is now made available while the virtual machine is running.
Storage Migration	Move the virtual hard disks used by a virtual machine to different physical storage while the virtual machine remains running.
SR-IOV	Assign a network adapter that supports single-root I/O virtualization (SR-IOV) directly to a virtual machine.
Virtual Fibre Channel	Allows connection directly to Fibre Channel storage from within the guest operating system that runs in a virtual machine.
Virtual Hard Disk Format	New format has been introduced to meet evolving requirements and take advantage of innovations in storage hardware.
Virtual NUMA	A virtual NUMA topology is made available to the guest operating system in a virtual machine.
Virtual Switch	The architecture of the virtual switch has been updated to provide an open framework that allows third parties to add new functionality to the virtual switch.

## System Center 2012 R2

Microsoft System Center is a complete management suite to help you capture and aggregate knowledge about your infrastructure, policies, processes, and best practices so your IT staff can build manageable systems and automate repeatable operations.

Component	Description
App Controller	Provides a unified console that helps you manage public clouds and private clouds, as well as cloud-based virtual machines and services.
Operations Manager	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 datacenter and cloud, both private and public.
Orchestrator	Provides orchestration, integration, and automation of IT processes through the creation of runbooks, enabling you to define and standardize best practices and improve operational efficiency.

Service Manager	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 Information Technology Infrastructure Library (ITIL). It provides built-in processes for incident and problem resolution, change control, and asset lifecycle management.
Virtual Machine Manager	A management solution for the virtualized datacenter, 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.
Configuration Manager, Data Protection Manager, Endpoint Protection	Additional components included in the System Center license but not included as part of this solution.

## Integration Components

Both Cisco and EMC have integration components to extend the capabilities of Microsoft's System Center environment.

### Windows PowerShell

Windows PowerShell is a task-based command-line shell and scripting language designed especially for system administration. Built on the .NET Framework, Windows PowerShell helps IT professionals and power users control and automate the administration of the Windows operating system and applications that run on Windows.

Windows PowerShell commands, called cmdlets, let you manage the computers from the command line. Windows PowerShell providers let you access data stores, such as the registry and certificate store, as easily as you access the file system. In addition, Windows PowerShell has a rich expression parser and a fully developed scripting language.

Cisco provides two different PowerShell modules:

- **PowerTool**—Cisco UCS PowerTool is a PowerShell module that helps automate all aspects of Cisco UCS Manager including server, network, storage, and hypervisor management. Cisco UCS PowerTool enables easy integration with existing IT management processes and tools. The PowerTool cmdlets work on the Cisco UCS Manager's Management Information Tree (MIT). These cmdlets allow you to create, modify, or delete actions on the Managed Objects (MOs) in the tree.
- **Cisco Nexus 1000V**—Microsoft's Hyper-V implements an extensible virtual switch, allowing third parties the ability to add functions to the basic switch. Cisco Nexus 1000V Series Switches provide a comprehensive and extensible architectural platform for virtual machine (VM) and cloud networking. The switches are designed to accelerate server virtualization and multitenant cloud deployments in a secure and operationally transparent manner. Cisco provides a PowerShell module that enables managing the Cisco Nexus 1000V.

EMC provides a suite of integration tools and services in the EMC Storage Integrator (ESI) for Windows Suite. ESI enables administrators to manage EMC storage systems such as VMAX and VNX, EMC Solutions Enabler with SMI-S Provider, and EMC RecoverPoint in Microsoft cloud environments including Exchange, SharePoint, SQL Server, System Center, and PowerShell.

Included with the ESI Suite are the ESI PowerShell Toolkit and ESI Service PowerShell Toolkit. The ESI PowerShell Toolkit enables you to provision and manage storage to Microsoft Windows hosts that use EMC storage. This toolkit includes a set of PowerShell cmdlets to manage EMC storage systems

from the PowerShell command line. The ESI PowerShell Toolkit provides access to most of the provisioning functionality offered by the ESI Microsoft Management Console (MMC) application and shares a common configuration set with the MMC application. The ESI PowerShell Toolkit provides cmdlets to manage:

- Connections to host and storage systems and to provision block storage
- Disk devices in hypervisor environments, including Microsoft Hyper-V
- Block device snapshots for EMC VNX storage arrays

The ESI Service PowerShell Toolkit provides cmdlets for setting up the ESI Service for use with the ESI System Center Operations Manager (SCOM) Management Packs. The ESI Service PowerShell Toolkit provides cmdlets for setting up the ESI Service to:

- Get the storage system entity data and entity relationships for the entity graph
- Configure service security
- Register EMC storage systems

## System Center Operations Manager

System Center Operations Manager (SCOM) is a cross-platform data center management system for operating systems and hypervisors. It uses a single interface that shows state, health and performance information of computer systems. It also provides alerts generated according to some availability, performance, configuration or security situation being identified. It works with Microsoft Windows Server and Unix-based hosts.

SCOM uses the term 'management pack' to refer to a set of filtering rules specific to some monitored application. While Microsoft and other software vendors make management packages available for their products, SCOM also provides for authoring custom management packs. While an administrator role is needed to install agents, configure monitored computers and create management packs, rights to simply view the list of recent alerts can be given to any valid user account.

The Cisco UCS management pack provides deep visibility into the health, performance, and availability of Cisco UCS through a single, familiar and easy to use interface. The management pack contains rules which monitor chassis, blades, rack servers, service profiles etc. across multiple Cisco Unified Computing Systems.

The EMC SCOM Management Packs enable you to manage EMC storage systems with SCOM by providing consolidated and simplified dashboard views of storage entities. These views enable you to:

- Discover and monitor the health status and events of your EMC storage systems and system components in SCOM
- Receive alerts for any possible problems with disk drives, power supplies, storage pools, and other types of physical and logical components in SCOM

## System Center Orchestrator

System Center Orchestrator is an automation platform for orchestrating and integrating IT tools to decrease the cost of datacenter operations while improving the reliability of IT processes. It enables IT organizations to automate best practices, such as those found in Microsoft Operations Framework (MOF) and Information Technology Infrastructure Library (ITIL). System Center Orchestrator operates through workflow processes that coordinate System Center and other management tools to automate incident response, change and compliance, and service-lifecycle management processes.

The Cisco UCS Integration Pack is an add-on for Microsoft System Center Orchestrator (SCO) that enables automation of Cisco UCS Manager tasks. The Cisco UCS Integration Pack is used to create workflows that interact with and transfer information to other Microsoft System Center products such as Microsoft System Center Operations Manager.

## System Center Virtual Machine Manager

System Center Virtual Machine Manager (VMM) is a management solution for the virtualized datacenter, 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.

The Cisco UCS Manager Add-in for SCVMM 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). Using the add-in, you can perform tasks such as viewing server details, viewing service profiles, launching the host KVM console, associating service profiles to a server or server pool, and many other functions.

SCVMM 2012 introduced standards-based discovery and automation of iSCSI/Fibre Channel (FC) block storage resources in a virtualized datacenter environment. These new capabilities build on the Storage Management Initiative Specification (SMI-S) developed by the Storage Networking Industry Association (SNIA). The SMI-S standardized management interface enables an application such as SCVMM to discover, assign, configure, and automate storage for heterogeneous arrays in a unified way. To take advantage of this new storage capability, EMC updated its SMI-S Provider to support the SCVMM 2012 R2 release.

The EMC SMI-S Provider aligns with the SNIA goal to design a single interface that supports unified management of multiple types of storage arrays. The one-to-many model enabled by the SMI-S standard makes it possible for SCVMM to interoperate, by means of the EMC SMI-S Provider, with multiple disparate storage systems from the same SCVMM Console.

Some of the key value propositions for SCVMM and storage integration can be found in the following table:

Value add	Description
Reduce costs	<ul style="list-style-type: none"> <li data-bbox="378 1331 1409 1520">· <b>On-demand storage:</b> Aligns IT costs with business priorities by synchronizing storage allocation with fluctuating user demand. SCVMM elastic infrastructure supports thin provisioning, that is, SCVMM supports expanding (or contracting) the allocation of storage resources on EMC storage systems in response to waxing or waning demand.</li> <li data-bbox="378 1556 1409 1661">· <b>Ease-of-use:</b> Simplifies consumption of storage capacity, and thus saves time and lowers costs, by enabling the interaction of EMC storage systems with, and the integration of storage automation capabilities within, the SCVMM private cloud.</li> </ul>

Simplify administration	<ul style="list-style-type: none"> <li>• Private cloud GUI: Allows administration of private cloud assets (including storage) through a single management UI, the SCVMM Console, available to SCVMM or cloud administrators.</li> <li>• Private cloud CLI: Enables automation through SCVMM's comprehensive set of Windows PowerShell commands (cmdlets).</li> <li>• <b>Reduce errors:</b> Minimizes errors by providing the SCVMM UI or CLI to view and request storage.</li> <li>• <b>Simpler storage requests:</b> Automates storage requests to eliminate delays of days or weeks.</li> </ul>
Deploy faster	<ul style="list-style-type: none"> <li>• <b>Deploy VMs faster and at scale:</b> Supports rapid provisioning of VMs to Hyper-V hosts or host clusters at scale. SCVMM can communicate directly with SAN arrays to provision storage for VMs. SCVMM 2012 can provision storage for a VM in the following ways:             <ul style="list-style-type: none"> <li>– Create a new logical unit from an available storage pool: You can control the number and size of each logical unit.</li> <li>– <b>Create a writeable snapshot of an existing logical unit:</b> You can provision many VMs quickly by rapidly creating multiple copies of an existing virtual disk. This puts minimal load on hosts and uses space on the array efficiently.</li> <li>– <b>Create a clone of an existing logical unit:</b> You can offload the creating of a full copy of a virtual disk from the host to the array. Typically, clones are not as space-efficient as snapshots and take longer to create.</li> </ul> </li> <li>• Reduce load: Rapid provisioning of VMs using SAN-based storage resources takes full advantage of EMC array capabilities while placing no load on the network.</li> </ul>

In addition to the previously mentioned iSCSI and Fibre capabilities, the new VNX platform contains an embedded SMI-S provider which enables SCVMM to manage the network attached storage (NAS) based capabilities of the array. The VNX NAS provider offers the following capabilities when used in conjunction with SCVMM 2012 R2. Register the VNX within SCVMM to allow the array to be a "managed" file sharing device:

- View VNX based file share information
- Create file shares
- Remove file shares

## Conclusion

The Cisco UCS and EMC VSPEX solution for Microsoft's Private Cloud Fast Track 4.0 is an optimal shared infrastructure for deploying private cloud solution. Cisco and EMC have created and validated a solution that is both flexible and scalable, easily adapted to differing customer environments. Customers can be assured that they can start with the sized solution they need and readily expand it as their needs grow.

# Bill of Materials

Custom Name	SKU	Description	Qty
Blade Server	UCSB-B200-M3	UCS B200 M3 Blades w/o CPU, memory, HDD	8
UCSB-B200-M3		2.50 GHz E5-2660 V2	16
		16GB DDR3-1866-MHz	128
	UCSB-MLOM-40G-01	Cisco UCS VIC 1240 modular LOM for M3 blade server	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	UCS 5108 Blade Svr AC Chassis/0 PSU/8 fans/0 fabric extender	1
N20-C6508	UCS-IOM-220 4XP	UCS 2204XP I/O Module (4 External, 16 Internal 10Gb Ports)	2
	N01-UAC1	Single phase AC power module for UCS 5108	1
	N20-CAK	Access. kit for 5108 Blade Chassis including Railkit, KVM dongle	1
	N20-FAN5	Fan module for UCS 5108	4
	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-6248 UP	UCS 6248UP 1RU Fabric Int/No PSU/32 UP/ 12p LIC
UCS-FI-6248UP	UCS-ACC-62 48UP	UCS 6248UP Chassis Accessory Kit	2
	N10-MGT010	UCS Manager v2.2	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

Nexus 5548UP	N5K-C5548U P-OSM.P	Nexus 5548UP Storage Solutions Bundle, Full Stor Serv Lic, OSM	2
	N55-48PO-SS K9.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-750 W.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-K IT.P	Nexus 5548 Chassis Accessory Ki	2
	N5KUK9-521 N1.1.P	Nexus 5000 Base OS Software Rel 5.2(1)N1(1)	2
Nexus 1000V			2