

EMC INTEGRATION FOR MICROSOFT PRIVATE CLOUD USING EMC VNX UNIFIED STORAGE

EMC Next-Generation VNX, EMC Storage Integrator for Windows Suite, Microsoft System Center 2012 SP1

- Reduce storage costs by using the VNX flash-optimized unified platform for file and block storage of your Microsoft applications
- Reduce operational costs by streamlining storage management and provisioning for Microsoft Exchange, SharePoint, Lync, and SQL Server with the EMC Storage Integrator for Windows Suite
- Gain agility by orchestrating advanced workflows and processes through ESI integration for Microsoft Private Cloud

EMC Solutions

Abstract

This white paper describes the features and benefits of the EMC® next-generation VNX® storage platform, including performance, support for block and file workloads (SMB3.0), and the complementary software like EMC Storage Integrator (ESI) for Windows, which simplifies and streamlines advanced storage tasks for application and database owners. This solution uses a VNX 7600 with 10,000 Exchange, SharePoint, Lync, and SQL Server users.

November 2013



Copyright © 2013 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided “as is.” EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

All trademarks used herein are the property of their respective owners.

Part Number H12234

Table of contents

Executive summary	6
Business case.....	6
Solution overview	6
Key results	7
Introduction	9
Purpose	9
Scope	9
Audience	10
Terminology.....	10
Technology overview	12
Overview.....	12
EMC next-generation VNX.....	12
next-generation VNX Intel MCx Code Path Optimization.....	12
Flash-optimized hybrid array	13
next-generation VNX performance	13
Unisphere Management Suite	14
EMC FAST Suite	15
FAST VP	15
FAST Cache.....	15
ESI for Windows Suite	15
ESI for Windows	15
ESI SCOM Management Packs.....	16
ESI SCO Integration Pack.....	16
ESI PowerShell Toolkit.....	16
EMC SMI-S Provider.....	16
Microsoft System Center 2012 SP1	16
SCVMM 2012	16
SCO 2012.....	16
SCOM 2012.....	17
Microsoft Windows Server 2012 with Hyper-V	17
Microsoft SMB 3.0	17
Microsoft Collaboration.....	17
Microsoft Exchange Server 2013	17
Microsoft SharePoint Server 2013	17
Microsoft Lync Server 2013	18
Microsoft SQL Server 2012.....	18
Failover Clustering	18

Solution architecture and configuration	19
Overview.....	19
Physical environment.....	19
Hardware resources	21
Software resources	21
Workload profile	22
Storage design.....	23
EMC FAST Suite design.....	26
EMC FAST Cache.....	26
EMC FAST VP	26
Network design.....	27
File storage best practices.....	28
Block storage best practices	28
 System management showcase and configuration.....	 29
Overview.....	29
Interoperability with Microsoft components	30
ESI Management Console	31
Storage monitoring and reporting	32
ESI Management Packs	32
How to enable storage monitoring in SCOM.....	33
Storage monitoring with ESI Management Packs	33
Cloud management.....	36
Hyper-V management.....	36
Storage management	37
Workflow management	41
SCO preparation steps	41
ESI SCO Integration Pack activities	41
Runbook Designer.....	42
 Application design and configuration.....	 44
Overview.....	44
Co-location of SQL Server 2012, Lync Server 2013, and SharePoint Server 2013 volumes and segregation of Exchange Server 2013 volumes.....	44
SQL Server 2012	44
Design and requirements	44
Configuration	45
SharePoint Server 2013	45
Design.....	45
Configuration	46
SharePoint Server 2013 user requirements	46
Exchange Server 2013	47

Exchange Server 2013 user requirements.....	47
Exchange Server 2013 design	48
Microsoft Exchange Server Jetstress 2013	48
Lync Server 2013	49
Configurations.....	49
Hyper-V Cluster	50
Validation.....	51
Overview.....	51
Functional tests	51
Functional testing of System Center 2012 SP1 with EMC components	51
Test methodology	54
Application performance across the environment.....	54
SQL Server 2012	54
SharePoint Server 2013.....	54
Microsoft Jetstress 2013.....	55
Lync Stress and Performance Tool	55
Performance test results	56
Combined application test results	56
Conclusion	59
Summary	59
Findings.....	59
Reference	60
White papers	60
Product documentation.....	60
Other websites.....	60

Executive summary

Business case

More and more organizations are looking to transform to a fully virtualized cloud infrastructure to both improve their service delivery and build a foundation to deliver IT as a service. This journey can be both daunting and confusing because administrators face the challenge of managing the wide range of ever-changing technologies across the modern data centers.

EMC can take you step by step to your cloud and make it simpler to manage and more efficient to operate.

A typical scenario is for an organization to run a Microsoft Windows environment on a Hyper-V platform, and use a unified storage array to support its data centers.

EMC simplifies the management of such a unified storage architecture with the close integration of EMC Storage Integrator (ESI) for Windows Suite with Microsoft System Center 2012 Service Pack 1 (SP1). The integration of these EMC and Microsoft products provides IT departments with a centralized platform to manage the application and storage infrastructure for the whole data center.

Administrators can now take advantage of ESI for Windows Suite to manage, monitor, and automatically deploy EMC storage; in addition, the EMC Fully Automated Storage Tiering (FAST™) Suite of technologies provides intelligent storage self-optimization for EMC® next-generation VNX® storage arrays.

Solution overview

This solution showcases an efficient virtual infrastructure with a cost-effective storage design that enhances a customer's database, collaboration, and unified communications experience.

The solution provides the following configurations:

- A variety of Microsoft application workloads running on a unified storage infrastructure that is hosted on a Windows Server 2012 Hyper-V cluster. Microsoft System Center 2012 SP1 with ESI for Windows Suite manages the environment.
- The Microsoft Server Message Block (SMB) 3.0 protocol-based infrastructure for Microsoft SQL Server 2012, SharePoint Server 2013, and Lync Server 2013 for up to 10,000 users.
- An iSCSI-based infrastructure for Microsoft Exchange Server 2013 with up to 10,000 users.
- The EMC next-generation VNX series storage platform, used for both block (iSCSI) and file (SMB 3.0) access.

Key results

The key findings of this solution are as follows:

- Management, orchestration, and automation
 - The ESI Management Packs for System Center Operations Manager (SCOM) allow SCOM 2012 SP1 to monitor the health and performance of EMC storage systems.
 - The EMC SMI-S Provider allows System Center Virtual Machine Manager (SCVMM) 2012 SP1 to easily manage and provision EMC storage.
 - ESI Integration pack for System Center Orchestrator (SCO) 2012 SP1 allows automated provisioning and allocation of EMC storage with SCO 2012 SP1.
- Performance, availability, and efficiency
 - The next-generation VNX series arrays, together with SMB 3.0, iSCSI block, and Microsoft Windows Server 2012, support a 10,000-user Microsoft application environment for Exchange Server 2013, SQL Server 2012, SharePoint Server 2013, and Lync Server 2013.
 - FAST VP policies allow for the co-location of multiple workloads in the same storage pool, which simplifies storage design without compromising performance. Microsoft applications, which generate complex and mixed workloads, such as SQL Server, Lync, and SharePoint, continue to receive optimal storage performance, while achieving maximum storage efficiency and easy deployment.
 - FAST Cache improves the performance of applications on both the block level and file level. FAST Cache is easily deployed and the performance benefits are quickly evident.
 - The average latency improvement with FAST Suite for the SMB 3.0 applications is more than 91%.

Figure 1 shows the key findings of this solution.

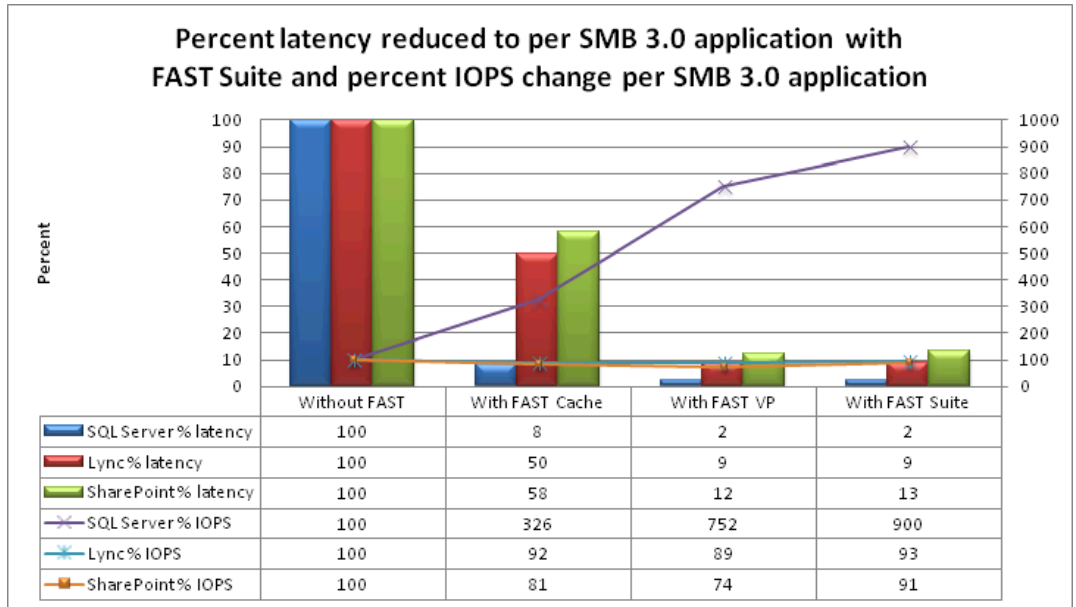


Figure 1. Key findings of this solution

Table 1 lists the latency reduction in percentage.

Table 1. Percent latency reduction

Percent latency reduction	FAST Cache	FAST VP
SQL Server	91%	97%
Lync	50%	90%
SharePoint	42%	87%

Introduction

Purpose

The objective of this solution is to demonstrate EMC's integration with Microsoft System Center 2012 SP1, which allows the management of a Microsoft private cloud environment with Windows Server 2012. In addition, this solution showcases how performance is affected when a variety of Microsoft applications run on the unified storage system (next-generation VNX) with FAST Suite technology.

Scope

A Microsoft-based cloud environment supports critical database, collaboration, and unified communications business applications. The scope of this solution is to test and validate a Microsoft-based cloud environment that uses EMC technologies to enhance the overall experience for cloud architects, system administrators, and end users.

The solution demonstrates how EMC products integrate with Microsoft System Center 2012 SP1 to:

- Simplify management of EMC storage infrastructures natively from within SCVMM by using EMC's SMI-S support.
- Create automated workflows with SCO to support fully-automated storage provisioning by using the EMC's ESI Integration Pack.
- Simplify monitoring and event-based alerts natively within SCOM for EMC's storage platforms by using EMC's Management Pack.

We¹ validated the following storage architecture that supports Microsoft's latest Windows Server 2012 and application offerings:

- On block-based storage—iSCSI
 - Exchange Server 2013
 - For predictable performance based on the building block approach
- On file-based storage—SMB 3.0
 - SharePoint Server 2013
 - Lync Server 2013
 - SQL Server 2012
 - Acceptable performance in a mixed tier for Lync and SharePoint
 - SQL Server added to show a latency-sensitive application with unacceptable performance and how FAST Suite addresses and improves the performance to significantly below the required latency

We used the following EMC products and technologies:

- EMC VNX7600

The solution features are available on all next-generation VNX models.

¹ In this white paper, "we" and "our" refer to the EMC Solutions engineering team that validated the solution.

- EMC Storage Integrator for Windows Suite
 - ESI PowerShell Toolkit
 - ESI SCOM Management Packs
 - ESI SCO Integration Pack
 - EMC SMI-S Provider
- EMC FAST Suite technologies
 - EMC FAST Cache
 - EMC FAST VP
 - EMC Virtual Provisioning (thin LUNs)

Audience

This white paper is intended for EMC employees, including business developers and solution architects, EMC partners, and EMC customers such as product managers and service providers' technologists and architects.

Terminology

This white paper includes the terminology listed in Table 2.

Table 2. Terminology

Term	Definition
ANSI	American National Standards Institute
AVM	Automatic Volume Management
BDM	Background Database Maintenance
CIFS	Common Internet File System
DIR	Deleted items retention
ESE	Extensible Storage Engine
ESI	EMC Storage Integrator
FC	Fibre Channel
GB	Gigabyte
GbE	Gigabit Ethernet. A term describing various technologies for transmitting Ethernet frames at a rate of a gigabit per second.
IM	Instant messaging
IOPS	Input/output operations per second
NIC	Network interface card
OE	Operating Environment
OS	Operating System
PSTN	Public Switch Telephone Network

Term	Definition
QoS	Quality of service
SAN	Storage area network
SAS	Serial-attached SCSI
SCO	Microsoft System Center Orchestrator
SCOM	Microsoft System Center Operations Manager
SCVMM	Microsoft System Center Virtual Machine Manager
SMB	Microsoft Server Message Block
SMI	Storage Management Initiative
SNIA	Storage Networking Industry Association
SP	Storage processor
TB	Terabyte
VHDX	Hyper-V virtual hard disk format—an enhanced format available in Microsoft Windows Server 2012
VSTS	Microsoft Visual Studio Team System

Technology overview

Overview

This section provides an overview of the key components of this solution, as listed in Table 3.

Table 3. Solution components

System	Components
Storage platform	EMC next-generation VNX
System software	<ul style="list-style-type: none">• EMC FAST Suite• EMC Unisphere® Management Suite• ESI for Windows Suite• EMC SMI-S Provider• Microsoft System Center 2012 SP1
Virtualization platform	Microsoft Windows Server 2012 with Hyper-V
Connectivity	Microsoft SMB 3.0 and iSCSI
Applications	<ul style="list-style-type: none">• Microsoft SharePoint Server 2013• Microsoft Lync Server 2013• Microsoft SQL Server 2012• Microsoft Exchange Server 2013

EMC next-generation VNX

EMC next-generation VNX is a flash-optimized, unified storage platform that delivers innovation and enterprise capabilities for file, block, and object storage in a single, scalable, and easy-to-use solution. Ideal for mixed workloads in physical or virtual environments, the next-generation VNX combines powerful and flexible hardware with advanced efficiency, management, and protection software to meet the demanding needs of today's virtualized application environments.

next-generation VNX includes many features and enhancements designed and built upon the first generation's success. These features and enhancements include:

- More capacity with multicore optimization from multicore cache, multicore RAID, and multicore FAST Cache (MCx)
- Greater efficiency with a flash-optimized hybrid array
- Better protection by increasing application availability with an active/active array service processor
- Easier administration and deployment by increasing productivity with the new Unisphere Management Suite

next-generation VNX Intel MCx Code Path Optimization

The flash technology has totally changed the requirements of midrange storage systems. EMC redesigned the midrange storage platform to provide the highest performing storage system at the lowest cost in the market by efficiently optimizing multicore CPUs.

MCx distributes all VNX data services across all cores (up to 32), as shown in Figure 2. The next-generation VNX series with MCx has dramatically improved the file performance for transactional applications like databases or virtual machines over network-attached storage (NAS).

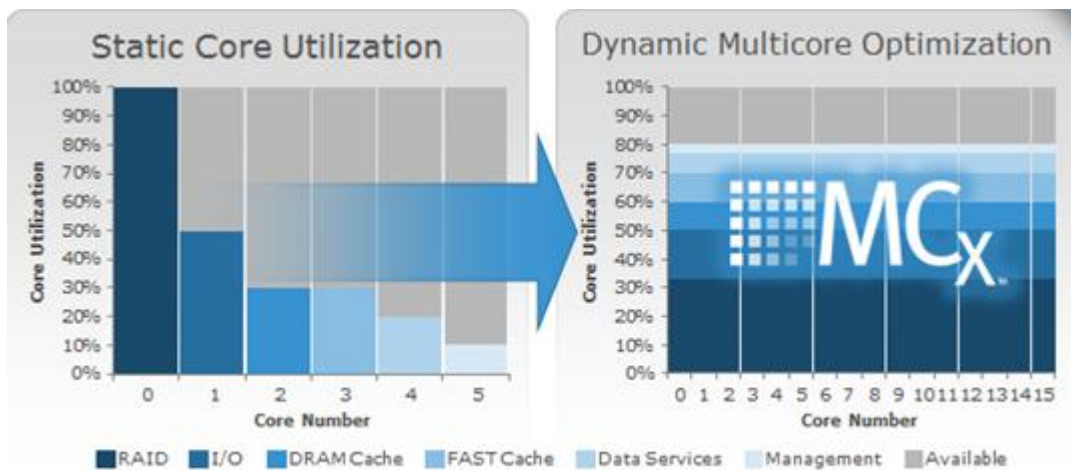


Figure 2. next-generation VNX with multicore optimization

Flash-optimized hybrid array

next-generation VNX is a flash-optimized hybrid array that provides automated tiering to deliver the best performance to your critical data, while intelligently moving less frequently accessed data to lower cost disks.

In this hybrid approach, a small percentage of flash drives in the overall system can provide a high percentage of the overall IOPS. The flash-optimized VNX takes full advantage of the low latency of flash to deliver cost-saving optimization and high-performance scalability. The EMC FAST Suite (FAST Cache and FAST VP) tiers both block and file data across heterogeneous drives and boosts the most active data to flash drives, ensuring that customers never have to make cost or performance concessions.

Because data generally is accessed most frequently at the time it is created, new data is first stored on flash drives to provide the best performance. As that data ages and becomes less active, FAST VP tiers the data from high-performance to high-capacity drives automatically, based on customer-defined policies. EMC has enhanced this functionality with four times better granularity and with new FAST VP solid state disks (SSDs) based on enterprise multilevel cell (eMLC) technology to lower the cost per gigabyte. FAST Cache dynamically absorbs unpredicted spikes in system workloads.

next-generation VNX performance

Performance enhancements

VNX storage, enabled with the MCx architecture, is optimized for FLASH 1st and provides unprecedented overall performance, such as transaction performance (cost per IOPS), bandwidth performance (cost per GB/s) with low latency, and optimal capacity efficiency (cost per GB).

next-generation VNX provides the following performance improvements:

- Up to four times more file transactions when compared with dual controller arrays
- Up to three times better file performance for transactional applications (for example, Microsoft Exchange on VMware over NFS) with a 60 percent better response time
- Up to four times more Oracle and Microsoft SQL Server OLTP transactions
- Up to six times more virtual machines

Active/active array service processors

The next-generation VNX architecture provides active/active array service processors, as shown in Figure 3. This eliminates application timeouts during path failover, since both paths are actively serving I/O.

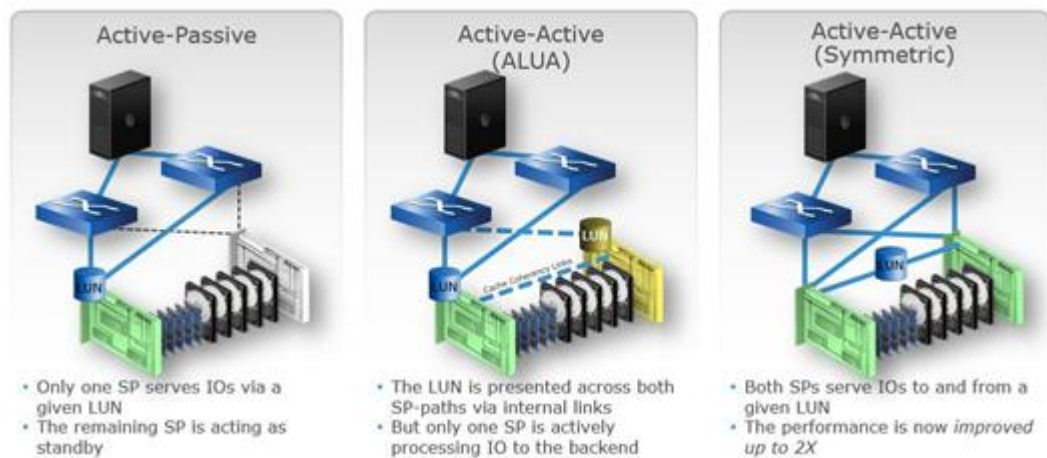


Figure 3. Active/active processors increase performance, resiliency, and efficiency

Load balancing has also been improved and applications can achieve up to two times improvement in performance. Active/active for Block is ideal for applications that require the highest levels of availability and performance, but do not require tiering or efficiency services like compression, deduplication, or snapshot.

Note: The active/active processors are available only for classic logical unit numbers (LUNs), not for pool LUNs.

Unisphere Management Suite

The Unisphere Management Suite extends Unisphere's easy-to-use interface to include VNX Monitoring and Reporting for validating performance and anticipating capacity requirements. As shown in Figure 4, the suite also includes Unisphere Remote for centrally managing up to thousands of EMC VNX and EMC VNXe® systems with support for EMC XtremSW™ Cache.

Unisphere Management Suite

Improve Management Productivity



Figure 4. Unisphere Management Suite

EMC FAST Suite

FAST VP

EMC FAST VP, a part of the EMC FAST Suite, is a simple and elegant solution for dynamically matching storage requirements with changes in the frequency of data access. It is a feature designed to aggressively reduce Total Cost of Ownership (TCO) and maintain performance. By mixing flash, SAS, and NL-SAS drives, you can maintain performance requirements while reducing the drive count. In some cases, you can achieve a nearly two-thirds reduction in drive count, while in other cases, you can double performance throughput by simply adding less than 10 percent of a pool's total capacity to flash drives.

FAST Cache

EMC FAST Cache, a part of the EMC FAST Suite, enables flash drives to function as an expanded cache layer for the array. FAST Cache is an array-wide, non-disruptive cache, available for both file and block storage. Frequently accessed data is copied to the FAST Cache in 64 KB increments and subsequent reads and/or writes to the data chunk are serviced by FAST Cache. This enables immediate promotion of highly active data to flash drives. This dramatically improves the response time for the active data and reduces data hot spots that can occur within a logical unit number (LUN).

ESI for Windows Suite

ESI for Windows Suite is a set of tools for Microsoft Windows and Microsoft applications administrators. The suite includes ESI for Windows, ESI SCOM Management Packs, ESI PowerShell Toolkit (PSToolkit), and ESI SCO Integration Pack.

ESI for Windows

ESI for Windows enables you to view, provision, monitor, and manage block and file storage for Microsoft Windows and Microsoft SharePoint. ESI supports the EMC Symmetrix VMAX series, EMC next-generation VNX and VNXe series, and CLARiiON® CX™ fourth-generation (CX4) series of storage arrays. ESI also supports storage provisioning and discovery for Windows virtual machines running on Microsoft Hyper-V, Citrix XenServer, and VMware vSphere.

ESI SCOM Management Packs

ESI SCOM Management Packs provide storage infrastructure management in a single user interface. With SCOM integration, Windows administrators can discover storage assets, map physical and logical objects, and manage alerts and roll-up health states, with configurable parameter thresholds. The supported platforms are VMAX series and next-generation VNX series for both block and file.

ESI SCO Integration Pack

Microsoft System Center 2012 Orchestrator is a workflow management system that defines, creates, and manages workflows. The ESI SCO Integration Pack provides workflow automation for ESI storage provisioning tasks in runbooks.

The ESI SCO Integration Pack enables you to manage and provision EMC storage for interoperable storage management and process consistency across a data center. This integration pack supports the same EMC block storage systems that ESI for Windows supports.

ESI PowerShell Toolkit

ESI PowerShell Toolkit is a utility that assists administrators and users of Windows with storage system management. PSToolkit cmdlets enable system administrators to obtain storage system information; create and delete storage pools, storage groups, and storage volumes; and map and mask the pools, groups, and volumes to available host servers. This toolkit also enables administrators to create automated scripts to dynamically create and delete virtual machines as needed.

EMC SMI-S Provider

EMC SMI-S Provider V4.4 supports the Storage Networking Industry Association (SNIA) Storage Management Initiative (SMI), which provides an American National Standards Institute (ANSI) standard for storage management. The SMI resulted in a standard management interface with a comprehensive specification, SMI-S. The specification defines the open storage management interface that enables the interoperability of multi-vendor storage-management technologies. These technologies monitor and control storage resources in multi-vendor or storage area network (SAN) topologies.

Microsoft System Center 2012 SP1

Microsoft System Center 2012 SP1 is a comprehensive, unified management platform that enables easy and efficient management of IT environments, including virtual/physical server infrastructures and client devices, in both traditional data centers and private clouds.

SCVMM 2012

SCVMM 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 across the virtual infrastructure.

SCO 2012

SCO provides a workflow management solution for the data center. It enables administrators to automate the creation, monitoring, and deployment of resources in the environment.

SCOM 2012

SCOM provides organizations with a flexible and cost-effective solution for monitoring their storage, computers, networks, and applications in a single console view. Operators can gain rapid insight into the state of their environment as well as receive alerts about availability, performance, configuration, and security situations.

Microsoft Windows Server 2012 with Hyper-V

Microsoft Windows Server 2012 with Hyper-V provides a complete virtualization platform, which offers increased scalability and performance with a flexible solution from the data center to the cloud. It makes it easier for organizations to realize the cost savings from virtualization and to optimize server hardware investments.

Windows Server 2012 Hyper-V high-availability options include incremental backup support, enhancements in clustered environments to support virtual adapters within the virtual machine, and inbox network interface card (NIC) teaming. In Hyper-V, *shared nothing* live migration enables the migration of a virtual machine from a server running Hyper-V to another one without the need for both servers to be in the same cluster or to share storage.

Microsoft SMB 3.0

Microsoft SMB 3.0 is an application-layer network protocol that provides shared access between nodes on a network. It also provides an authenticated interprocess communication mechanism. The SMB 3.0 protocol is supported on VNX systems running VNX File Operating Environment (OE) version 7.1.65 and later. VNX also supports SMB 1.0, SMB 2.x, SMB 3.0 protocol communications between client and server, and SMB 3.0 enhancements that include Continuous Availability (CA), Multichannel, Copy Offload, SMB Encryption.

Microsoft Collaboration

Microsoft Exchange Server 2013

Microsoft Exchange Server 2013 is an enterprise email and communication system that allows businesses and customers to collaborate and share information. EMC enhances Exchange Server 2013 with the industry's broadest choice of storage platforms, software, and services.

Exchange Server 2013 is a redesign of Exchange Server 2010 and adds simplicity of scale, better hardware utilization, and failure isolation. Exchange Server 2013 uses database availability groups (DAGs) and mailbox database copies, along with other features such as single item recovery, retention policies, and lagged database copies. These features provide high availability, site resilience, and Exchange native data protection. Microsoft enhanced the high-availability platform, the Exchange Information Store service, and the Extensible Storage Engine (ESE) to provide greater availability, easier management, and reduced costs.

Microsoft SharePoint Server 2013

Microsoft SharePoint Server 2013 provides a business-collaboration platform for enterprise and commercial organizations. SharePoint enables organizations to share content and information through websites, blogs, wikis, and document libraries, and to manage this content and information collectively from start to finish.

Microsoft Lync Server 2013

Microsoft Lync Server 2013 is an enterprise real-time communications server that provides the infrastructure for enterprise instant messaging, presence, and structured conferences connectivity.

Microsoft SQL Server 2012

Microsoft SQL Server 2012 is Microsoft's database management and analysis system for e-commerce, line-of-business, and data warehousing solutions. SQL Server is widely used to store, retrieve, and manage application data. Because it is used with a range of applications, and each application has different requirements for performance, sizing, availability, recoverability, manageability, and so on, it is important to fully understand these factors and plan accordingly when deploying SQL Server.

Failover Clustering

This solution deploys Windows Server 2012 Failover Clustering with four nodes at the production site. We used the Node and File Share Majority quorum configuration.

When determining where to place the virtual machines, be sure to consider load balancing and failure protection. Always distribute virtual machines with the same application roles across different nodes. For example, this solution separates SharePoint Server virtual machines between different nodes, so if a node fails, only one of the front-end servers is affected.

We configured a fixed size VHDX disk for each virtual machine OS. The virtual machine configuration file and the virtual machine memory swap file (which is the same size as the memory configured for the virtual machine) are located in the same folder as the VHDX.

No cluster disks are located in the cluster resources. All virtual disks are located in a CIFS share, which is used as the storage location.

Solution architecture and configuration

Overview

This solution architecture validates a private cloud infrastructure for companies looking for enterprise consolidation with simplified management. This solution uses ESI for Windows Suite and System Center 2012 SP1 to manage compute, network, and storage resources in a hypervisor. A next-generation VNX array runs mixed workloads of SQL Server 2012, SharePoint Server 2013, and Lync 2013 on the SMB 3.0 file system, and Exchange Server 2013 runs on the iSCSI block system.

Physical environment

Figure 5 displays the physical architecture of the solution.

In this solution, we designed a typical customer environment running multiple Microsoft applications on an EMC next-generation VNX array. We built the entire infrastructure on Windows Server 2012 Hyper-V.

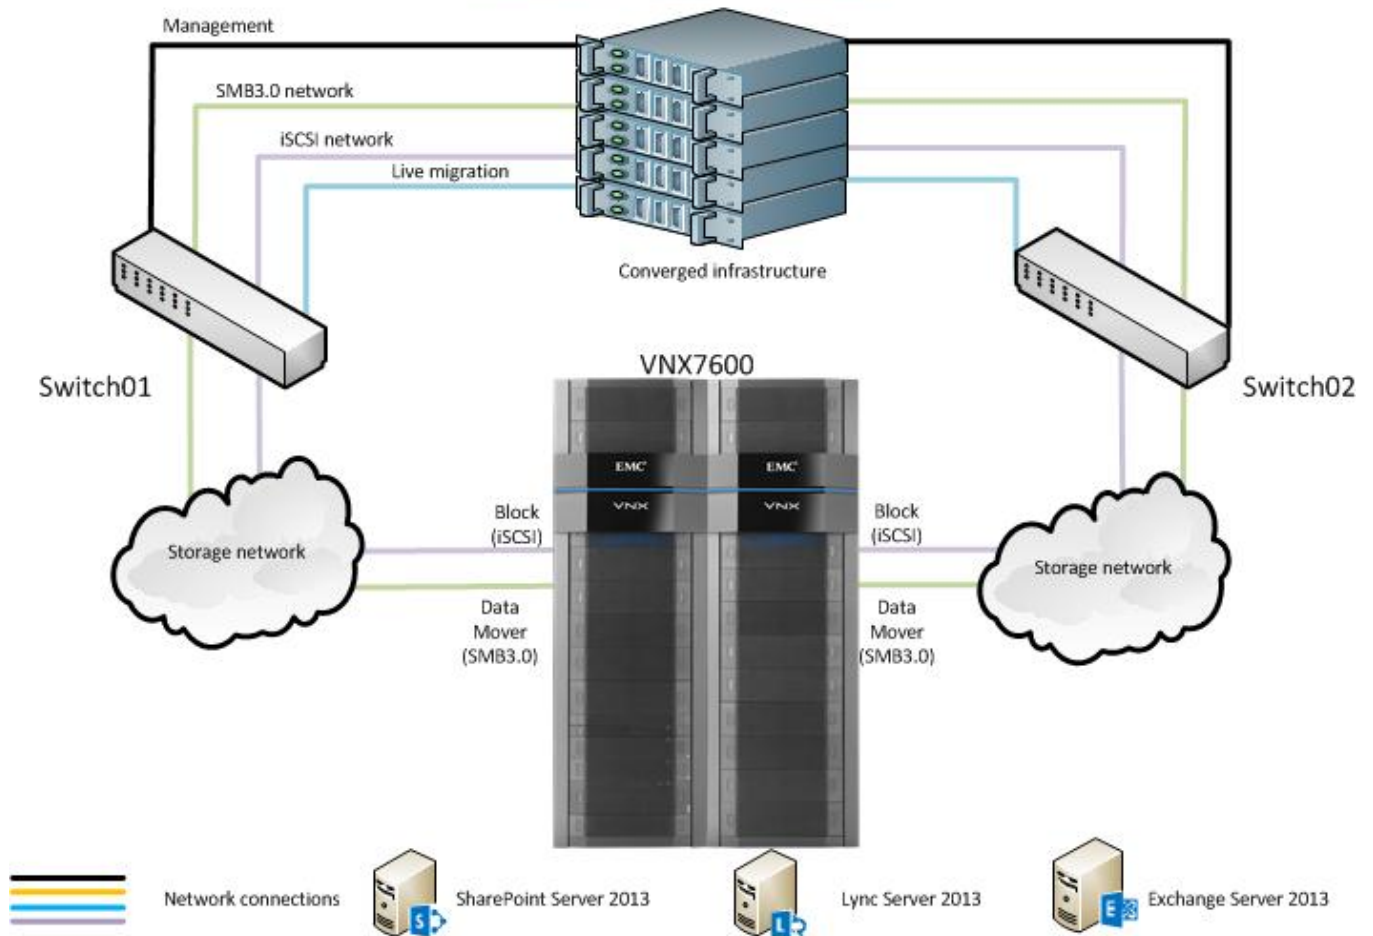
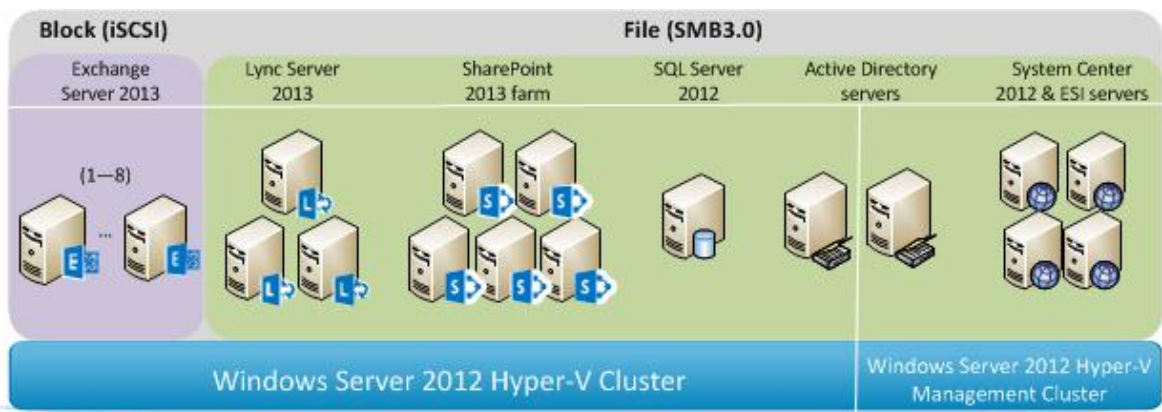


Figure 5. Solution architecture

Hardware resources

Table 4 details the hardware resources used in this solution.

Table 4. Hardware resources

Equipment	Quantity	Configuration
Storage platform	1	EMC VNX7600 with: <ul style="list-style-type: none"> • Two Data Mover Six Core Intel 2.80 GHz Westmere CPU modules with 24 GB memory • Two storage processors (SPs) • 44 x 3 TB 7.2 K rpm NL-SAS • 25 x 900 GB 10k rpm SAS • 5 x FAST VP SSD • 6 x FAST Cache SSD
Networking	1	Two Cisco UCS 6140XP switch with 40*10-GbE ports 4 virtual NICs mapping to 4 Ethernet NICs on each of blades of Cisco UCS
Server infrastructure	6 (blades)	Single chassis, 6 x blades. Each blade contains: <ul style="list-style-type: none"> • 2 x processors (20 cores in total) • 256 GB DDR3 RAM

Software resources Table 5 details the software resources used in this solution.

Table 5. Software resources

Resource	Quantity	Version	Purpose	
EMC software	EMC VNX OE for Block	1	RTM	VNX operating environment
	EMC VNX OE for File	1	RTM	VNX operating environment
	ESI for Windows Suite	1	3.0	Provisioning new storage and enhanced storage views and reporting
	EMC SMI-S Provider	1	4.4	Storage management support for EMC arrays
Microsoft software	Microsoft Windows Server	19	2012 Datacenter RTM	Hyper-V nodes and virtual machine operating system (OS)
	Microsoft SQL Server	3	2012 Enterprise Edition	Database servers for: <ul style="list-style-type: none"> • SQL Server OLTP • SharePoint • Database server for Microsoft System Center

Resource	Quantity	Version	Purpose
Microsoft SharePoint Server	4	2013 Enterprise	SharePoint farm
Microsoft Lync Server	1	2013	Conference and instant messaging
Microsoft Exchange Jetstress	4	2013	Jetstress testing
SCVMM	1	2012 SP1	Virtual infrastructure management
SCOM	1	2012 SP1	Environment health monitor
SCO	1	2012 SP1	Workflow management

Workload profile

The following tables detail the environment workload profiles used in validating this solution.

Table 6 details the Microsoft SQL Server 2012 workload profile.

Table 6. Microsoft SQL Server 2012 workload profile

Profile characteristic	Quantity/Type/Size
SQL Server instances	1
OLTP database (OLTP_DB)	10,000 users/100 GB
Workload type	OLTP-like (90:10 read/write ratio)

Table 7 details the Microsoft SharePoint Server 2013 workload profile.

Table 7. Microsoft SharePoint Server 2013 workload profile

Profile characteristic	Quantity/Type/Size
Total user count	10,000 users
Total data	4.3 TB
Content database sizes	4 TB 200 GB 100 GB
Number of content databases	3
Document size range	200 KB to 2 MB, average 335 KB
Total site collection count	3
Size per site	10 GB to 400 GB
Sites per site collection	10 to 100
Total site count	120

Table 8 details the Microsoft Exchange Server 2013 workload profile.

Table 8. Microsoft Exchange Server 2013 workload profile

Profile characteristic	Quantity/Type/Size
Number of Exchange Server 2013 users	10,000
Number of Mailbox Server virtual machines	8 (4 tested)
Number of DAGs and database copies	1 DAG with 2 copies
Number of users per Mailbox Server	2,500 total mailboxes (1,250 active and 1,250 passive during normal operating conditions)
User profile (in DAG configuration)	150 messages/user/day (0.10 IOPS)
Read/write ratio	3:2 in a DAG configuration
Mailbox size	2 GB
Target average message size	75 KB
Deleted items retention window	14 days
Logs protection buffer	3 days
24 x 7 background database maintenance configuration	Enabled

Storage design

We designed the storage to accommodate a customer with 10,000 users running Microsoft applications.

We designed the exchange storage according to best practices as per our building blocks. We separated the Exchange data pool and log pools because of their different I/O requirements. We also created a small pool for the Exchange OS LUNs. The Exchange LUNs were all thin LUNs as per next-generation VNX best practices.

SharePoint, Lync, and SQL Server shared a mixed pool with performance set to the Lowest Available Tier (mostly on NL-SAS disks). As per next-generation VNX best practices all file LUNs were thick LUNs.

Because SQL Server was on the lowest tier and shared the pool with other applications, its performance suffered. Introducing FAST Suite alleviated this issue and guaranteed optimal performance for all the applications sharing this pool.

We set up FAST Cache according to the best practices with all FAST Cache SSDs on enclosure 0_0. FAST VP SSDs were also on enclosure 0_0 but the best practice is to spread these SSDs across all buses if possible.

Table 9 details the VNX7600 block level storage configuration used for this solution.

Table 9. EMC VNX7600 block level storage configuration

Application	Storage pool name	Disk type	RAID type	Number of disks	Number of LUNs	Protocol
Exchange	ExchangePool	3 TB 7.2 K rpm NL-SAS	RAID 10	32	16	iSCSI
	ExchangeLogPool	3 TB 7.2 K rpm NL-SAS	RAID 10	4	16	iSCSI
	OSPool for exchange	900 GB 10k rpm SAS	RAID 5	5	1	iSCSI

Table 10 details the VNX7600 file level storage configuration used for this solution.

Table 10. EMC VNX7600 file level storage configuration

Application	Storage pool name	Disk type	RAID type	Number of disks	Number of LUNs	File system name	CIFS share name
Lync	MixedPool	900 GB 10 K rpm SAS	RAID 5	15	10	MixedFS	MixedShare
SQL Server		3 TB 7.2 K rpm NL-SAS	RAID 6	8			
SharePoint							

Table 11 shows the FAST Cache configuration.

Table 11. FAST Cache configuration

Number of disks	RAID type	Disk type	Purpose
6	RAID 10	FAST Cache SSD	FAST Cache

Table 12 shows the FAST VP configuration.

Table 12. FAST VP configuration

Number of disks	RAID type	Disk type	Purpose
5	RAID 5	FAST VP SSD	Mixed pool additional FAST VP SSD added for tiering

Figure 6 shows the disk layout in this solution.

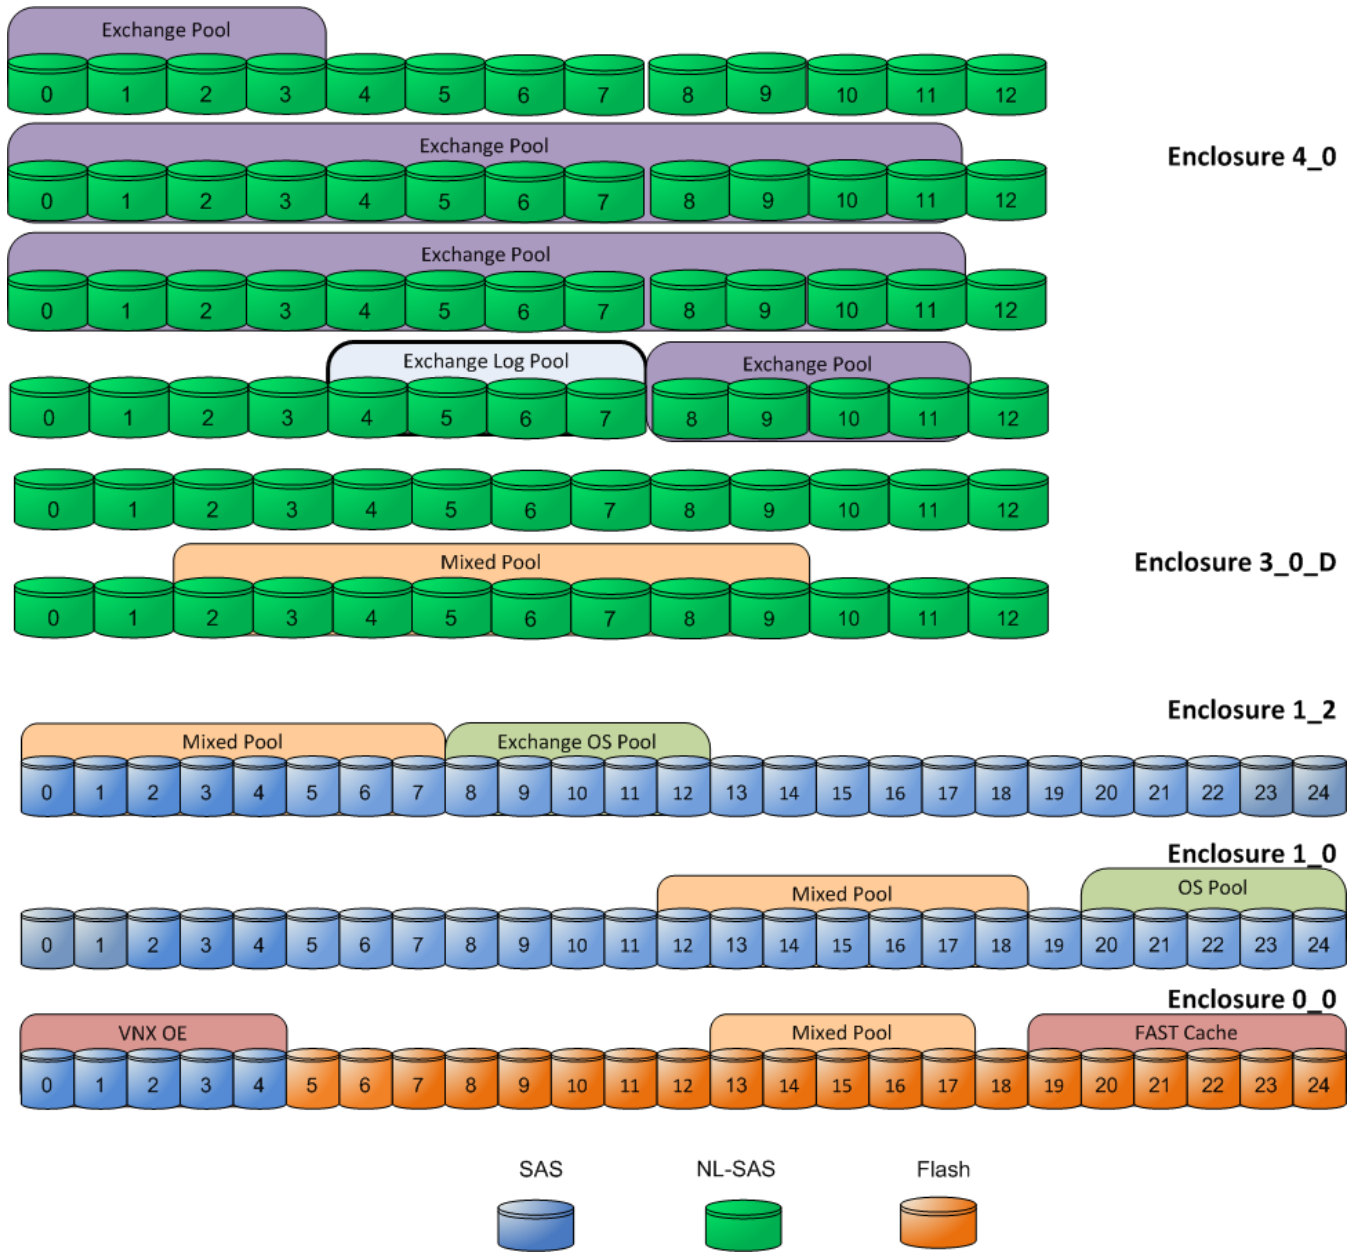


Figure 6. Disk layout of the solution

EMC FAST Suite design

EMC FAST Cache

FAST Cache consisted of six disks in a RAID 10 (3+3) configuration on flash 100 GB disks. During testing we enabled FAST Cache on the mixed pool LUNs and the Exchange Pool LUNs but not on the Exchange Log Pool LUNs. This is because the Log Pool I/O pattern is large sequential writes and this type of I/O does not suit FAST Cache.

EMC FAST VP

We configured FAST VP for the mixed pool. Initially the mixed pool had only SAS 900 GB disks and NL-SAS 3 TB disks, and the FAST VP policy was set to Lowest Available Tier. We then added five 100 GB flash disks in a RAID 5 (4+1) configuration to the mixed pool. We changed the FAST VP policy to Auto-Tier and ran relocations over several hours.

Table 13 details the FAST VP design.

Table 13. FAST VP design

Tiering policy	Storage pool description	Number of disks	RAID type	Disk type
FAST VP Lowest Available Tier	Mixed pool	15	RAID 5	900 GB 10 K rpm SAS
		8	RAID 6	3 TB 7.2 K rpm NL-SAS
FAST VP Auto-Tier	Mixed pool with additional FAST VP SSD added for tiering	15	RAID 5	900 GB 10 K rpm SAS
		8	RAID 6	3 TB 7.2 K rpm NL-SAS
		5	RAID 5	100 GB FAST VP SSD

Network design

The network design included a 10 GbE NIC on a Data Mover, a 10 GbE iSCSI I/O module on a storage processor, Cisco Fabric Interconnect, and a 10 GbE NIC card on the UCS server. We set up redundancy for each layer.

Figure 7 shows the network topology.




Figure 7. Network topology

This solution uses virtualization technology at the host level. We configured each host with four 10 GbE NICs and mapped each one to a virtual NIC in Hyper-V and Cisco UCS respectively. We separated our networking into four NICs so that network traffic for different applications or functions would not affect each other. We configured each subnetwork with its own vLAN and subnet segment.

Figure 8 displays the network configuration view with the following four network types from host side:

- Cluster—For live migration and cluster network load
- iSCSI—For Exchange storage network
- Management—For Hyper-V hosts and virtual machine management
- SMB—For SharePoint, SQL and Lync storage network

Networks	
Name	Status
Cluster	Up
iSCSI	Up
Management	Up
SMB	Up

Figure 8. Network configuration view from host side

In this solution we have two kinds of network connectivity between the host and VNX7600 storage. The storage processor connects with the host through an iSCSI network while Data Movers connect with the host through a CIFS share using the SMB 3.0 protocol. Exchange Server runs on the iSCSI protocol, while all other applications run on the SMB 3.0 protocol.

File storage best practices

The best practices for VNX file storage are:

- Spread input/output (I/O) evenly across all the disks in a storage pool by dividing the number of disks in the pool by 4 and rounding up to the nearest 10. The result is the number of LUNs the pool should contain.
- Use Automatic Volume Management (AVM) to achieve balance between performance and capacity on the file system.
- Do not use Direct Writes.

Note: When the number of disks is set, the number of storage pools does not affect file system performance.

Block storage best practices

The best practices for VNX block storage are:

- Use thin LUNs first for all LUNs.
- Use RAID 5 for SAS disk tiers and RAID 6 for NL-SAS disk tiers.
- Maintain a multiple of the preferred drive count for each tier selected. For example, in a SAS disk tier of the mixed pool we configured 15 disks as RAID 5 (4+1).

System management showcase and configuration

Overview

This solution uses the following EMC and Microsoft components to provide a centralized management platform for an enterprise-level private cloud using EMC storage infrastructure. The EMC ESI and SMI-S components serve as the enablers for the Microsoft components. This section also explains the relationship between these EMC components and Microsoft components.

Table 14 shows the integration functions of EMC components and Microsoft components.

Table 14. Integration functions of EMC components and Microsoft components

EMC components	Microsoft components	Integration function
ESI Management Packs	SCOM	Storage discovery and health monitoring
ESI Integration Pack	SCO	Automatic storage provision
SMI-S	SCVMM	Storage management

Note: These components are free to EMC customers.

Interoperability with Microsoft components

Figure 9 is a visual representation of the relationship between the EMC and Microsoft components.

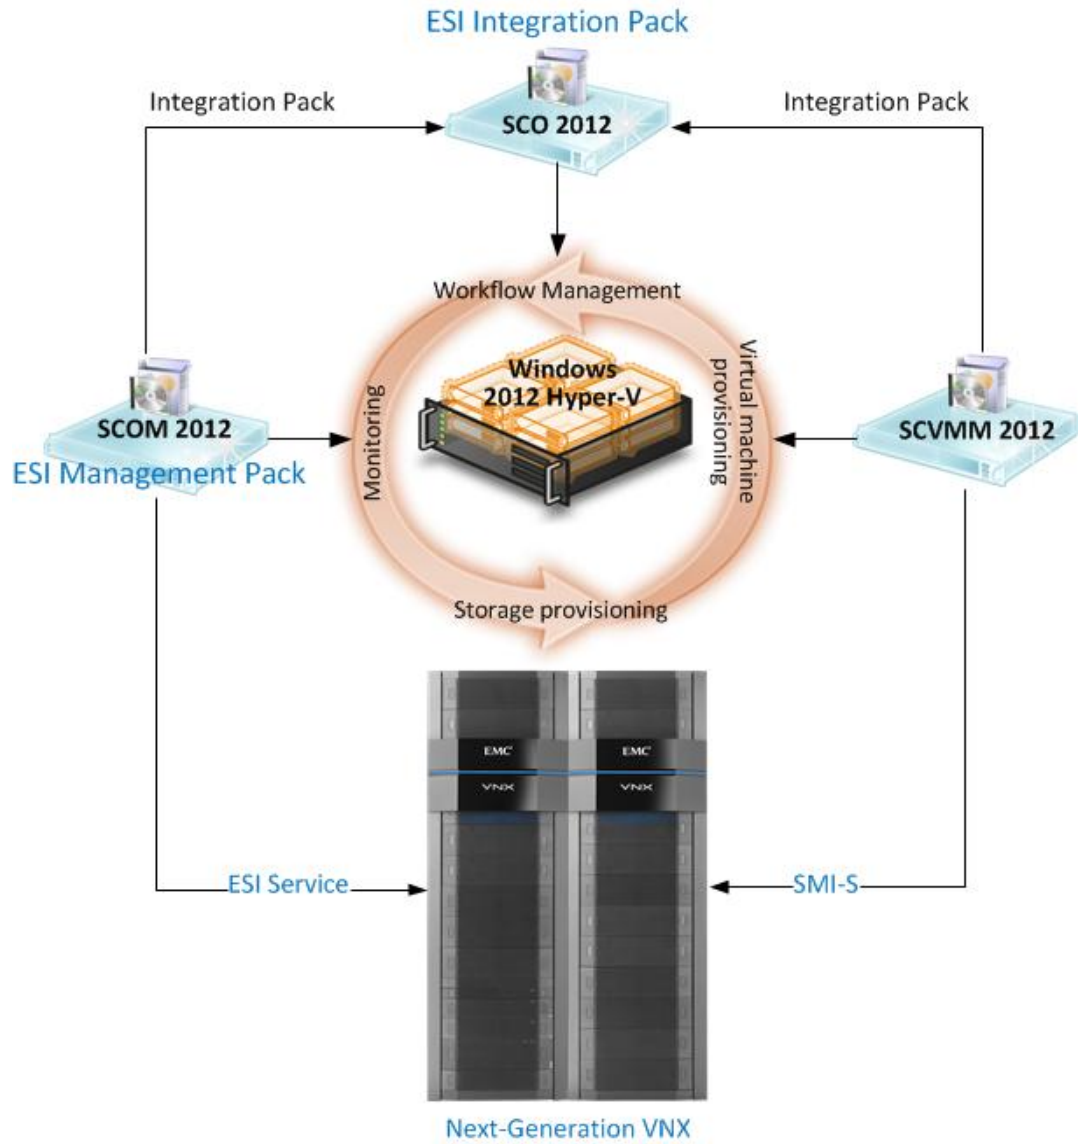


Figure 9. System management design for the solution

ESI Management Console

The technical overview section provides general information about ESI for Windows Suite. This section introduces the block/file view switch, a new feature of the ESI Management Console.

Besides providing storage management and hypervisor management, the ESI Management Console can also perform management on both block level and file level VNX arrays. You can switch between block management and file management simply by choosing **Block View** or **File View**.

Figure 10 shows the Block View status.

EMC Storage Integrator (x64)

File Action View Help

EMC Storage Integrator

- Storage Systems
 - 7600 - FNM0013C
 - Hosts
 - SSE-UCS1-B7 - 10
 - Host Clusters
 - Hypervisors
 - Applications
 - Replication

Friendly Name: 7600
 Array Name:
 Serial Number:
 System Type: VNX
 Model: VNX7600
 Software Revision:
 Switch View: Block View

Storage Pools	LUNs	Registered Hosts	Storage Groups	Service Nodes			
Name	User Capacity	Available Capacity	Subscribed Capacity	RAID Type	Provision Type	More Information	
Mixed Pool	26.090 TB	1.103 TB	24.987 TB	Mixed	Thick/Thin	Pool	
ExchangeLo...	5.374 TB	5.325 TB	1.802 TB	RAID1/0	Thick/Thin	Pool	
SCVM_jJCSI_...	10.746 TB	10.718 TB	61.119 GB	RAID5	Thick/Thin	Pool	
ExchangePool	5.374 TB	5.318 TB	1.808 TB	RAID1/0	Thick/Thin	Pool	
ExchangePool	42.989 TB	11.385 TB	39.358 TB	RAID1/0	Thick/Thin	Pool	
ExchangePo...	42.989 TB	12.245 TB	38.999 TB	RAID1/0	Thick/Thin	Pool	
ExchangePo...	42.989 TB	11.484 TB	39.358 TB	RAID1/0	Thick/Thin	Pool	
SCVM_CIFS_...	10.746 TB	759.230 GB	10.105 TB	RAID5	Thick/Thin	Pool	
OSPool	3.204 TB	93.683 GB	3.113 TB	RAID5	Thick/Thin	Pool	
OSPool for E...	3.204 TB	3.115 TB	3.250 TB	RAID5	Thick/Thin	Pool	
ExchangeLo...	5.374 TB	5.325 TB	1.802 TB	RAID1/0	Thick/Thin	Pool	
MetaLuns	0	0	0	N/A	N/A	Meta Luns	

Figure 10. Block View of VNX7600

Figure 11 show the File View status.

EMC Storage Integrator (x64)

File Action View Help

EMC Storage Integrator

- Storage Systems
 - 7600 - FNM0013C
 - Hosts
 - SSE-UCS1-B7 - 10
 - Host Clusters
 - Hypervisors
 - Applications
 - Replication

Friendly Name: 7600
 Array Name:
 Serial Number:
 System Type: VNX
 Model: VNX7600
 Software Revision:
 Switch View: File View

Shared Folder Pools	Shared Folders			
Name	Service Protocol	Total Capacity	Export Path	Parent Pool
OSFS	CIFS	3,000 TB	\\192.168.121.247\OSFS	OSFS
MixFS	CIFS	15,900 TB	\\192.168.121.247\MixFS	MixFS
test	CIFS	5,000 GB	\\192.168.121.247\test	OSFS
SCVM_1T	CIFS	1,000 TB	\\192.168.121.247\SCVM_1T	SCVM_1T
SCVM_500G	CIFS	500,000 GB	\\192.168.121.247\SCVM_500G	SCVM_500G

Figure 11. File View of VNX7600

This section highlights the storage monitoring features available in SCOM and lists key steps for discovering and monitoring EMC storage using SCOM 2012 SP1 with ESI Management Packs. This solution also uses SCOM 2012 SP1 to discover and monitor the health, performance, and availability of the entire virtual infrastructure across Exchange, Lync, SQL Server, and SharePoint applications as well as the operating systems and hypervisors.

ESI Management Packs

The following six ESI Management Pack files are provided free of charge to EMC customers:

- EMC.ESI.Library.mp: Defines the storage system monitors that ESI supports and the system components in SCOM.
- EMC.ESI.Discovery.mp: Contains all discoveries for all storage systems and system components defined in the library.
- EMC.ESI.Monitoring.mp: Contains the monitors, rules, and diagnostics for the storage systems and system components.
- EMC.ESI.Presentation.mp: Presents the storage system and system component folders and views in SCOM.
- EMC.ESI.Reporting.mp: Reports the health and capacity of the storage systems and system components.
- EMC.SI.Customization.xml: Contains all overrides and customizations for your specific storage environment. Import this management pack only during the initial installation. If you import it again, you will lose all of your previous overrides and customizations.

Figure 12 shows the ESI Management Packs loaded and available in SCOM.

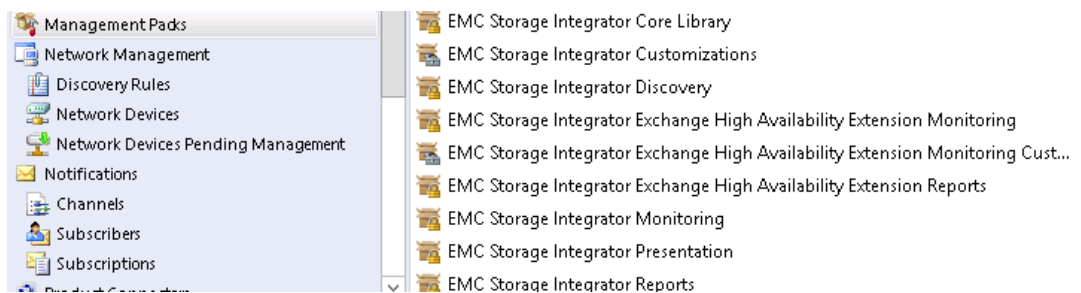


Figure 12. ESI Management Packs in SCOM

For example, use the ESI Monitoring Management Pack to discover EMC storage system components. SCOM agents discover and monitor these components. The monitoring agent retrieves data from the ESI Service by using a RESTful HTTPS connection, which in turn retrieves the data from the storage systems. The monitoring agent then inserts the monitoring data into the SCOM database.

Figure 13 shows the overall topology of storage systems.

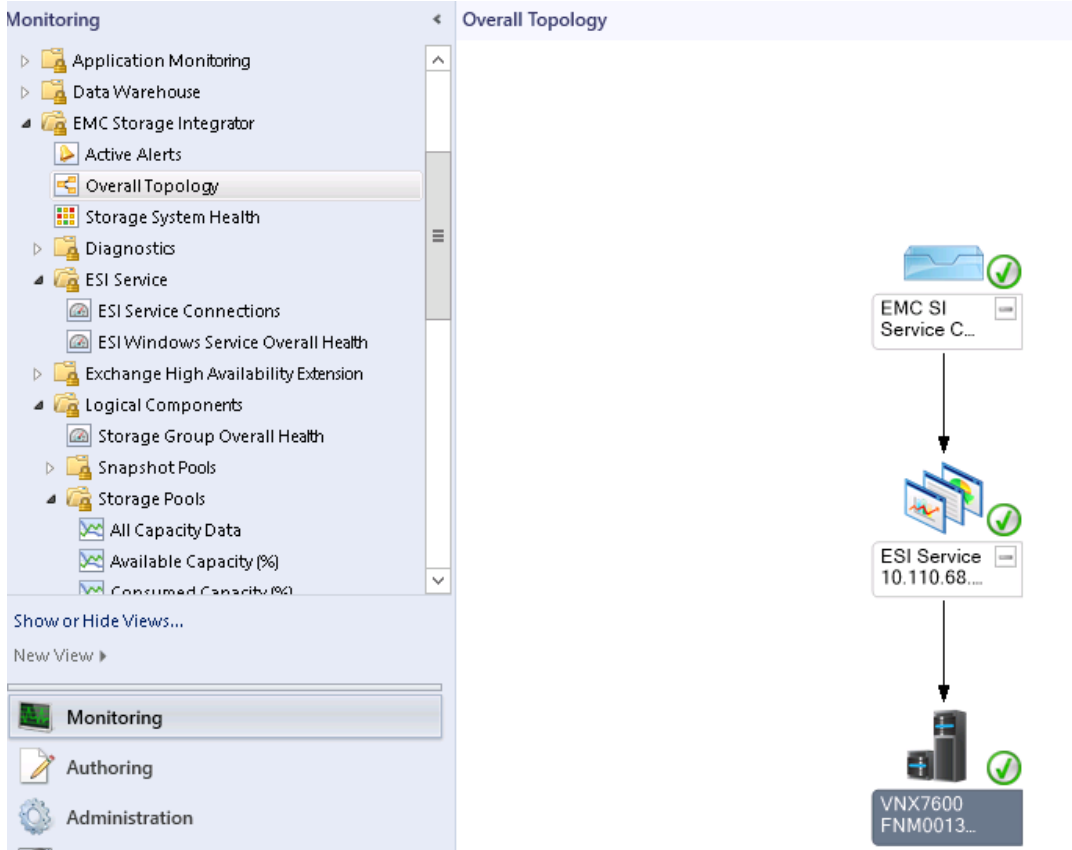


Figure 13. Overall topology view of storage system in SCOM

How to enable storage monitoring in SCOM

The ESI management packs are installed and imported on the SCOM server, with the ESI Service providing the communication link between ESI and the SCOM Management Group.

To enable storage monitoring in SCOM:

1. Install ESI Service on SCOM.
2. Install the ESI SCOM Management Packs.
3. Add the storage system into ESI Service.
4. Import ESI Management Packs into SCOM.

Note: For detailed ESI setup and configuration instructions, refer to the *EMC Storage Integrator for Windows Suite Product Guide*.

Storage monitoring with ESI Management Packs

After storage monitoring was enabled in SCOM, the ESI Service discovered the VNX7600. SCOM sends alerts for any issues with EMC storage platform components, such as disk devices, power supplies, and storage pools.

Figure 14 shows an example of the overall health monitoring view for storage pools on both block level and file level. SCOM was intentionally configured to trigger alerts when storage pools exceeded 90 percent capacity utilization. These alerts, which show a critical error state, notify you that the storage pool capacity needs to be expanded.

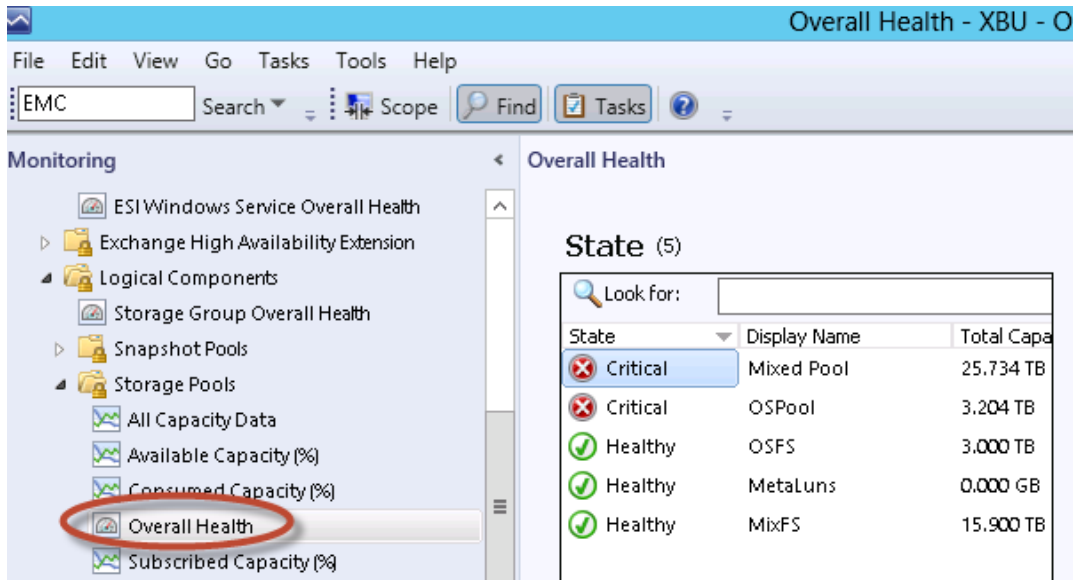


Figure 14. Storage Pools Overall Health monitoring view

Figure 15 shows the Detail View of the EMC SI Storage Pool. You can see from the details about this pool that the ESI for SCOM management pack discovered a CIFS share.

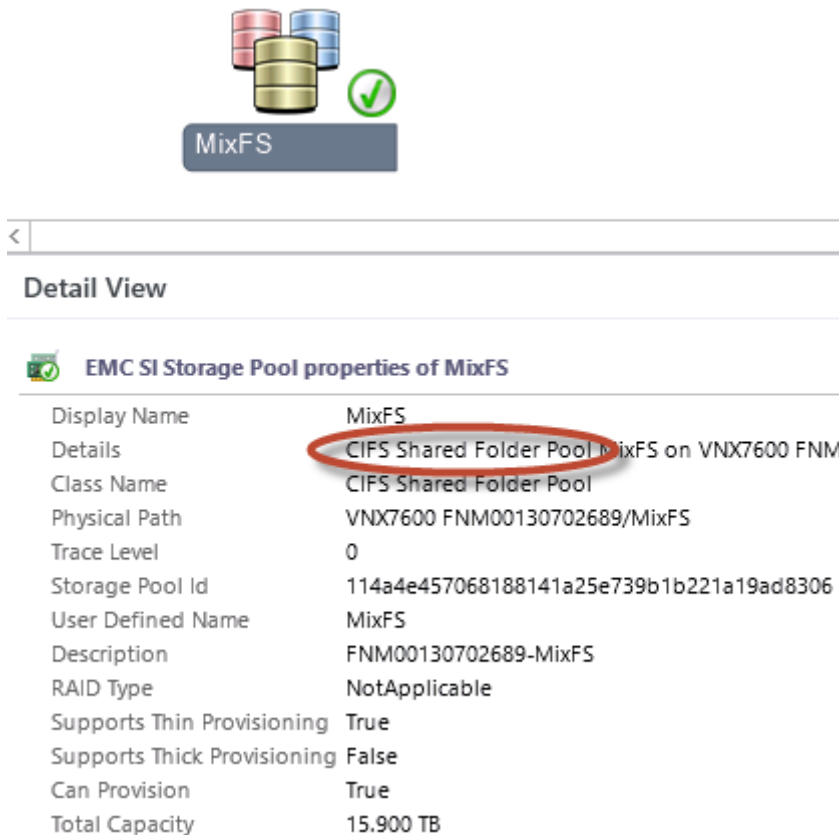


Figure 15. File system pool properties

Figure 16 shows the overall health of a disk drive in SCOM. From this console, you can see all disk drive related and status information.

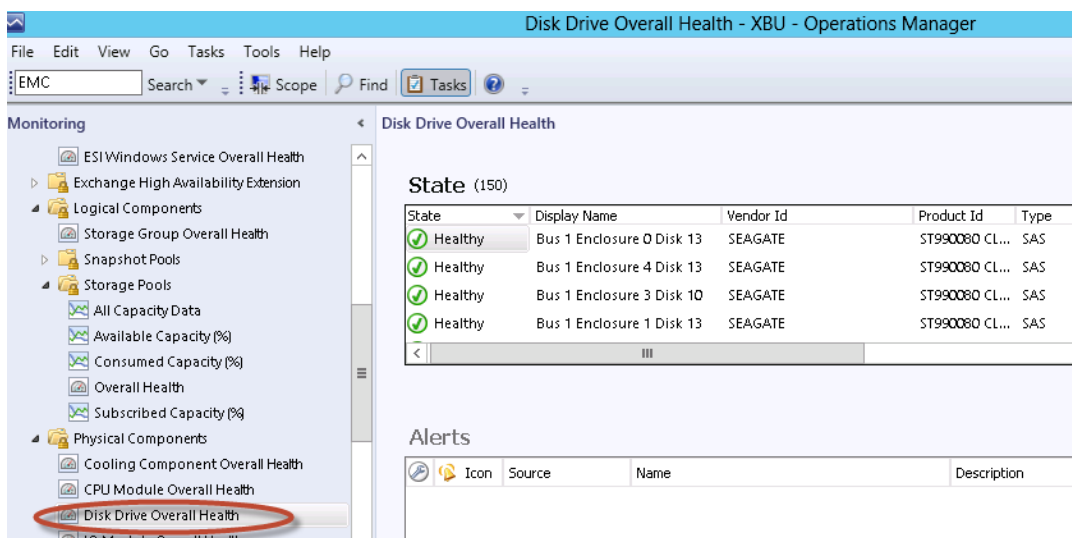


Figure 16. Disk Drive Overall Health monitoring view

For detailed information about the error state, right-click the error state, select **Open**, and choose **Health Explorer**, as shown in Figure 17.

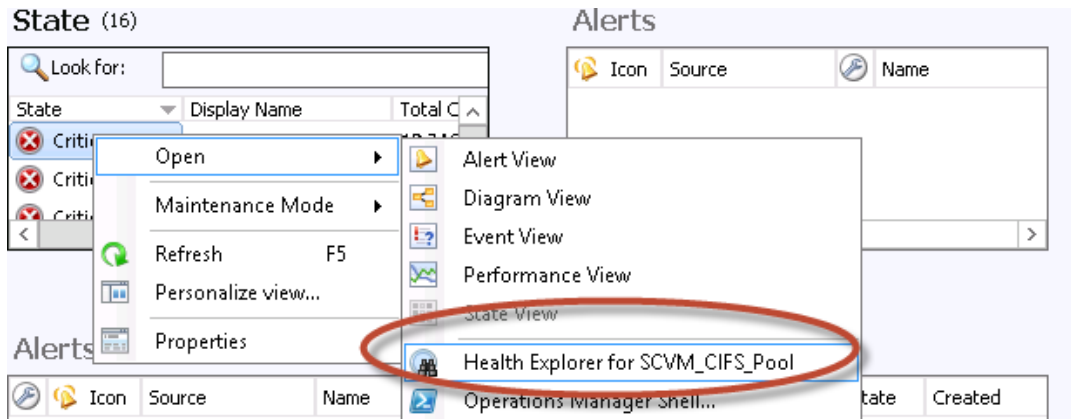


Figure 17. How to view detailed information about error states

Cloud management

This solution includes a cloud with both a Hyper-V infrastructure and an EMC storage system. To fulfill the purpose of cloud management, this solution uses SCVMM 2012 SP1 as an interface to manage Hyper-V and EMC unified storage arrays through EMC's close integration with System Center.

Note: SCVMM 2012 SP1 must be installed on Windows Server 2012.

Hyper-V management

This solution includes one Hyper-V cluster containing four nodes and another two-node management cluster, all running Windows Server 2012. We added the nodes to SCVMM 2012 SP1 and configured properties, allowing for configuration of virtual machines and other functions such as live migration and storage migration.

Note: Refer to the article [Technical Documentation Download for System Center 2012—Virtual Machine Manager](#) on Microsoft's website.

Figure 18 shows SCVMM 2012 SP1 managing the Windows Server 2012 features on the four Hyper-V nodes.

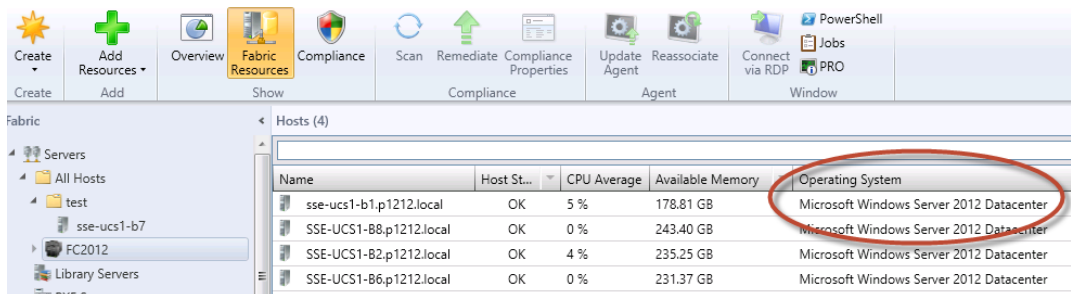


Figure 18. SCVMM 2012 managing Hyper-V features of Windows Server 2012

Storage management

Besides managing and monitoring virtual environments, SCVMM 2012 SP1 also provides new storage discovery, storage classification, and storage allocation features. You can use the SCVMM console to discover, classify, and provision storage on EMC arrays. SCVMM 2012 SP1 fully automates the assignment of storage to a Hyper-V host or Hyper-V host cluster, and tracks storage that it manages.

The installation of SCVMM 2012 SP1 includes the Microsoft Storage Management Service. This new Microsoft Storage Management Service can communicate with external arrays through an SMI-S provider.

EMC storage discovery

EMC provides its own SMI-S provider for SCVMM 2012 SP1. To configure SCVMM to discover the EMC storage systems, perform the following steps:

1. Set up the SMI-S provider server.
2. Add the storage system into the SMI-S provider server.
3. Configure the SMI-S provider in SCVMM 2012 SP1 by expanding the storage pane. Right-click **Providers**, select **Add storage devices**, input the SMI-S provider server IP address, and assign **Run As account** for this provider.

Figure 19 shows the status of the newly added SMI-S provider.

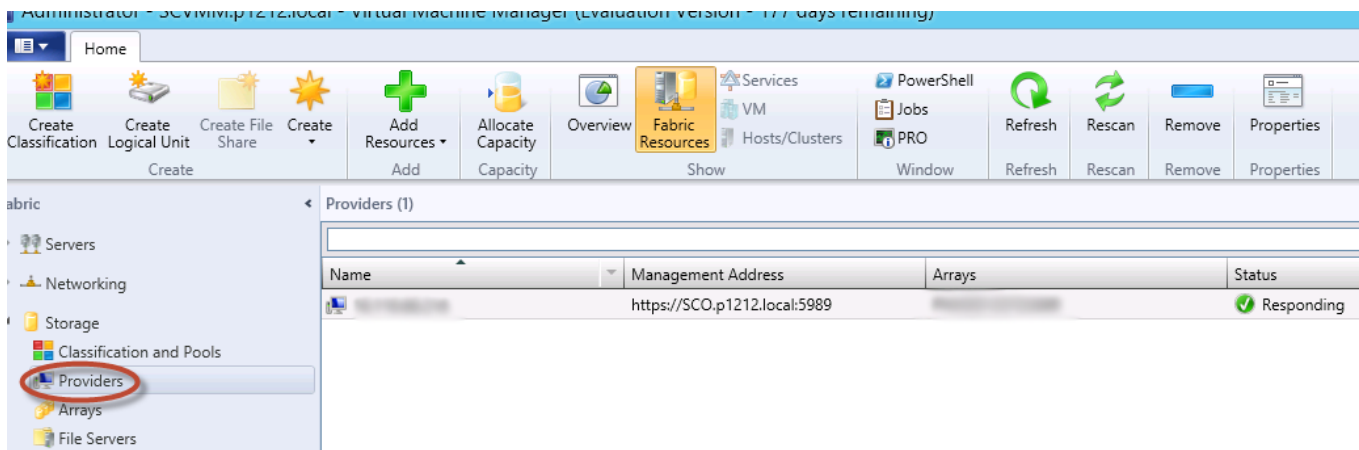


Figure 19. SMI-S provider status

Note: For detailed information about SMI-S provider setup and configuration, refer to the SMI-S provider release notes.

EMC SMI-S Provider is installed on the SCO server in this solution. Figure 20 shows the storage overview in the SCVMM storage management console.

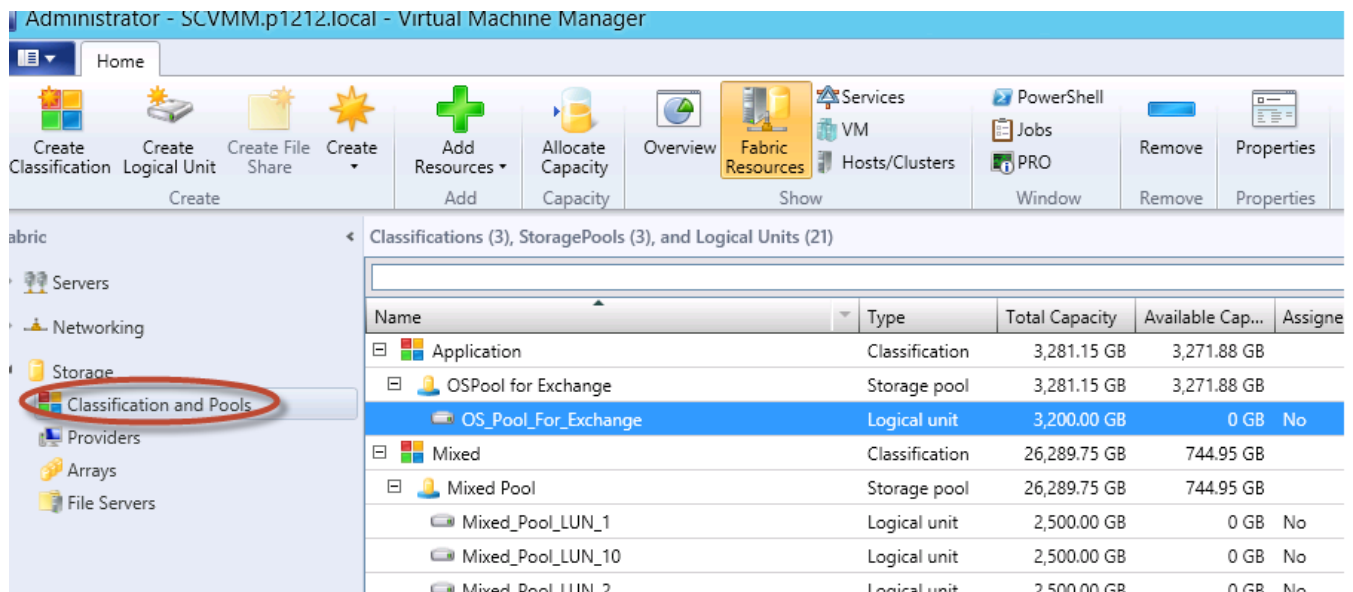


Figure 20. Storage overview of SCVMM 2012 SP1

Components

Figure 20 shows the four subcategories under the Storage section:

- **Classification and Pools:** Lets you assign user-defined storage classifications to discovered storage pools, typically by quality of service (QoS). For example, you can assign a classification of GOLD to storage pools that have the highest performance and availability.
- **Providers:** Lists the SMI-S provider server.
- **Arrays:** Lists the storage arrays that are discovered.

Refer to the article [Configuring Storage in VMM Overview](#) on the [Microsoft TechNet](#) website for supported arrays.

- **File Servers:** Shows the designated network file shares on Windows Server 2012 computers that SCVMM 2012 SP1 supports. The network file shares appear as the storage locations for virtual machine files, such as configuration, virtual hard disk files (.vhd/.vhdx) and checkpoints. This functionality uses the new SMB 3.0 protocol that was introduced in Windows Server 2012.

SMB 3.0 file system management

Besides block-based storage, SCVMM also discovers and manages the file system. In this solution, we created a CIFS share using the SMB 3.0 protocol on VNX7600. SCVMM can discover and manage this share through SMI-S Provider.

Figure 21 shows the file share registered in SCVMM. The storage file system exported this file share using the SMB 3.0 protocol.

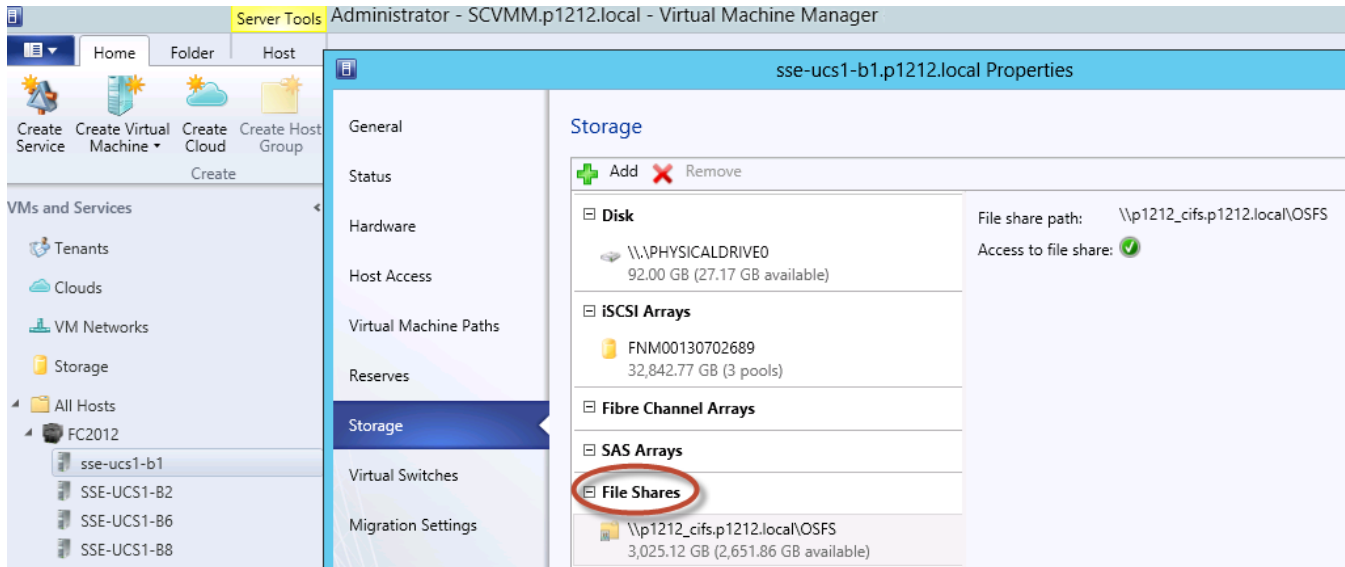


Figure 21. File share status in SCVMM

Storage pool management

You cannot use SCVMM to create VNX storage pools, but you can use it to manage existing storage pools and create LUNs within pools. To configure storage pools of EMC storage:

1. Expand the **Arrays** node under **Storage**.
2. Right-click one of the EMC arrays and choose **Properties**.
3. Select the storage pools that SCVMM should manage, and then click **OK**.

Figure 22 shows the array properties and managed storage pools on VNX7600.

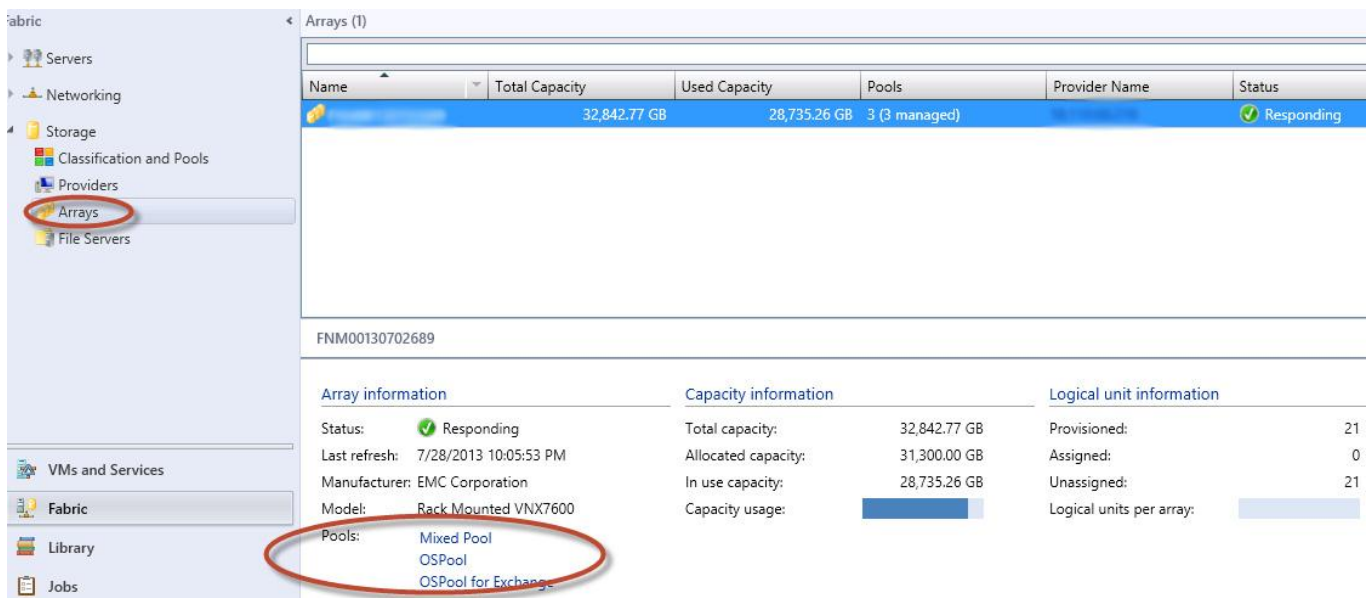


Figure 22. Managing storage pools on EMC storage

LUN creation with SCVMM

After you have selected the storage pools that SCVMM can manage, you can create LUNs. To create a thin LUN for a sample storage pool for the file system called **SCVM_CIFS_Pool**, follow these steps:

1. Right-click the storage pane and choose **Create Logical Unit**.
2. Define the storage pool in which you want to create this LUN.
3. Type the name of the LUN, choose the size, and select thin or fixed size storage logical unit.

Figure 23 shows the Create Logical Unit dialog box.

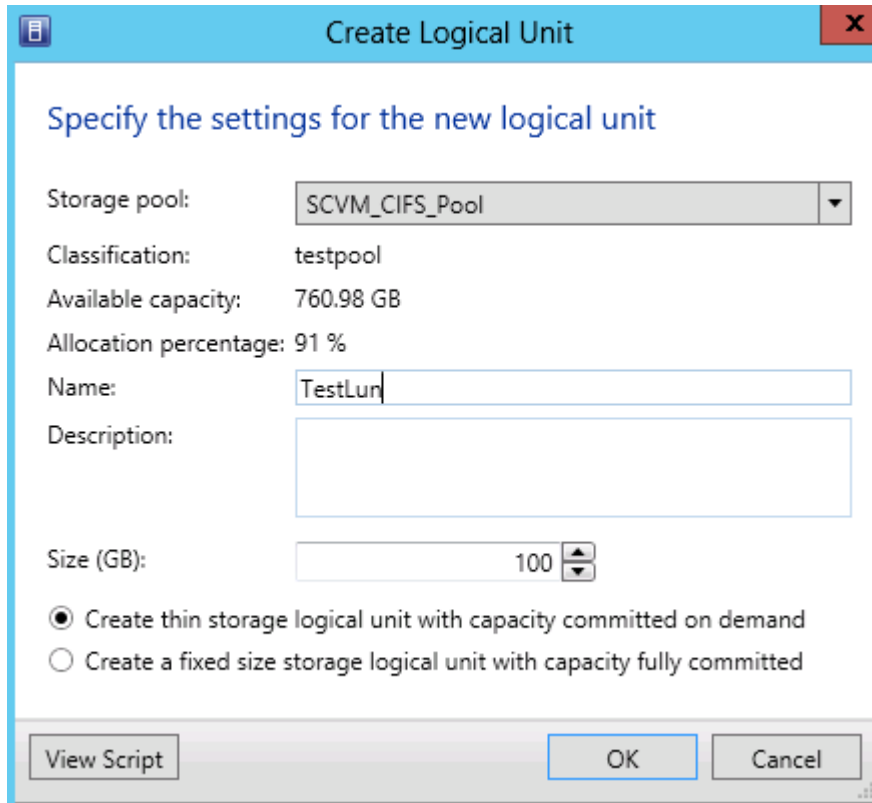


Figure 23. Creating logical unit with SCVMM

After you have successfully created a thin LUN, you can find it under the storage pool, as shown in Figure 24.

Name	Type	Total Capacity	Provisioning Type
SCVM_LUN_6	Logical unit	1,000.00 GB	Fixed
SCVM_LUN_7	Logical unit	1,000.00 GB	Fixed
SCVM_LUN_8	Logical unit	1,000.00 GB	Fixed
SCVM_LUN_9	Logical unit	1,000.00 GB	Fixed
TestLun	Logical unit	100.00 GB	Thin
SCVM_iSCSI_Pool	Storage pool	10,985.92 GB	

Figure 24. Thin LUN under storage pool created by SCVMM

Workflow management

This section shows how to use SCO 2012 SP1 and ESI Integration Pack together to automatically provision virtual machines and concurrently provision the required storage.

In SCO 2012 SP1 use a runbook to automatically create and deploy virtual machines, provision storage, and perform numerous other activities. Use the SCO 2012 SP1 Runbook Designer to design and manage runbooks.

SCO preparation steps

Before designing a runbook, complete the following steps on the SCO server:

1. Install and import the SCVMM Integration Pack.
2. Install and import the SCOM Integration Pack.
3. Install ESI and the ESI Integration Pack.
4. Import the ESI Integration Pack.

ESI SCO Integration Pack activities

The ESI SCO Integration Pack includes 16 activities, as shown in Figure 25.

ESI SCO Integration Pack	
Connect To Host Or Cluster System	Disconnect Host, Cluster, or Storage System
Connect To VMAX Storage System	Initialize Cluster Disk
Connect To VNX-Block Storage System	Initialize Host Disk
Connect To VNXe Storage System	Present LUN to Cluster
Create Cluster Filesystem	Present LUN to Host
Create Filesystem	Remove a LUN from Cluster
Create New LUN	Remove a LUN from Host
Delete LUN	Resize Host Volume

Figure 25. ESI SCO Integration Pack activities

Note: For SCO installation and IP installation and importing, refer to the [Microsoft TechNet](#) website.

Figure 26 shows an example of the “Initialize Host Disk” activity.

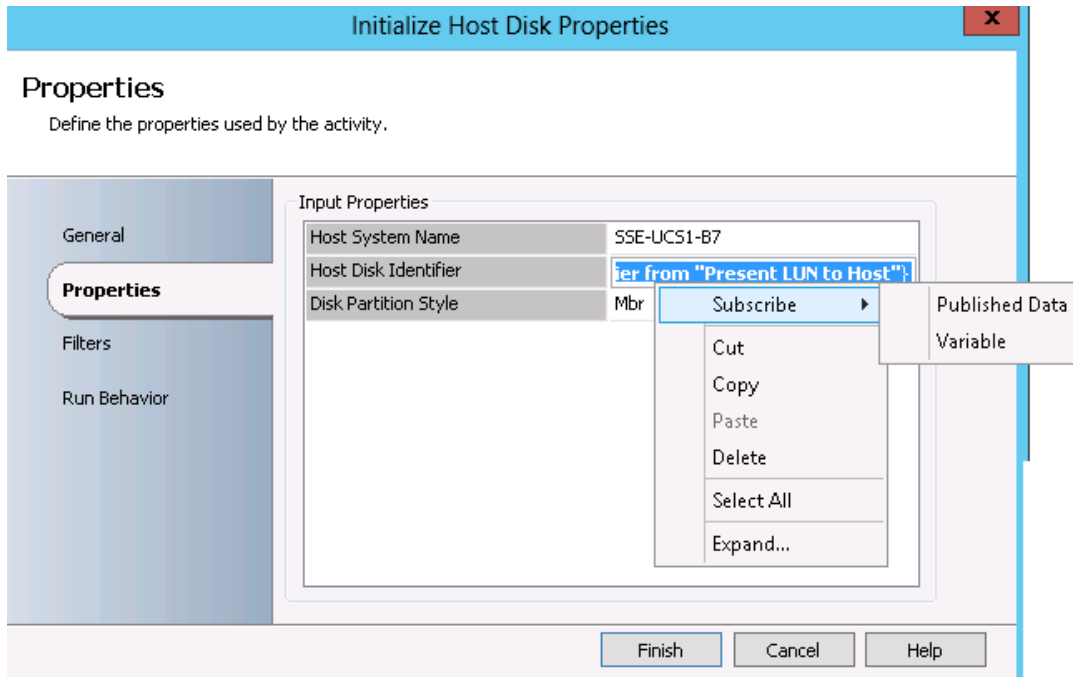


Figure 26. Initialize Host Disk Properties page

As you can see from the properties of this activity, you can use either **Published Data** or **Variable** for each input’s properties, which can save you time. For example, you can simply use data that is published from the “**Present LUN to Host**” activity as input for the **Host Disk Identifier** property.

Runbook Designer

In this section, we design a runbook to complete the following steps:

1. Create a virtual machine and deploy it on the server.
2. Provide storage provisioning.
3. Create a status report of key steps and show them in SCOM.

Figure 27 shows the Runbook Designer GUI. The 10 activities in this runbook come from ESI IP, SCVMM IP, and SCOM IP. Those activities include storage connection, storage provisioning, virtual machine provisioning, and status report.

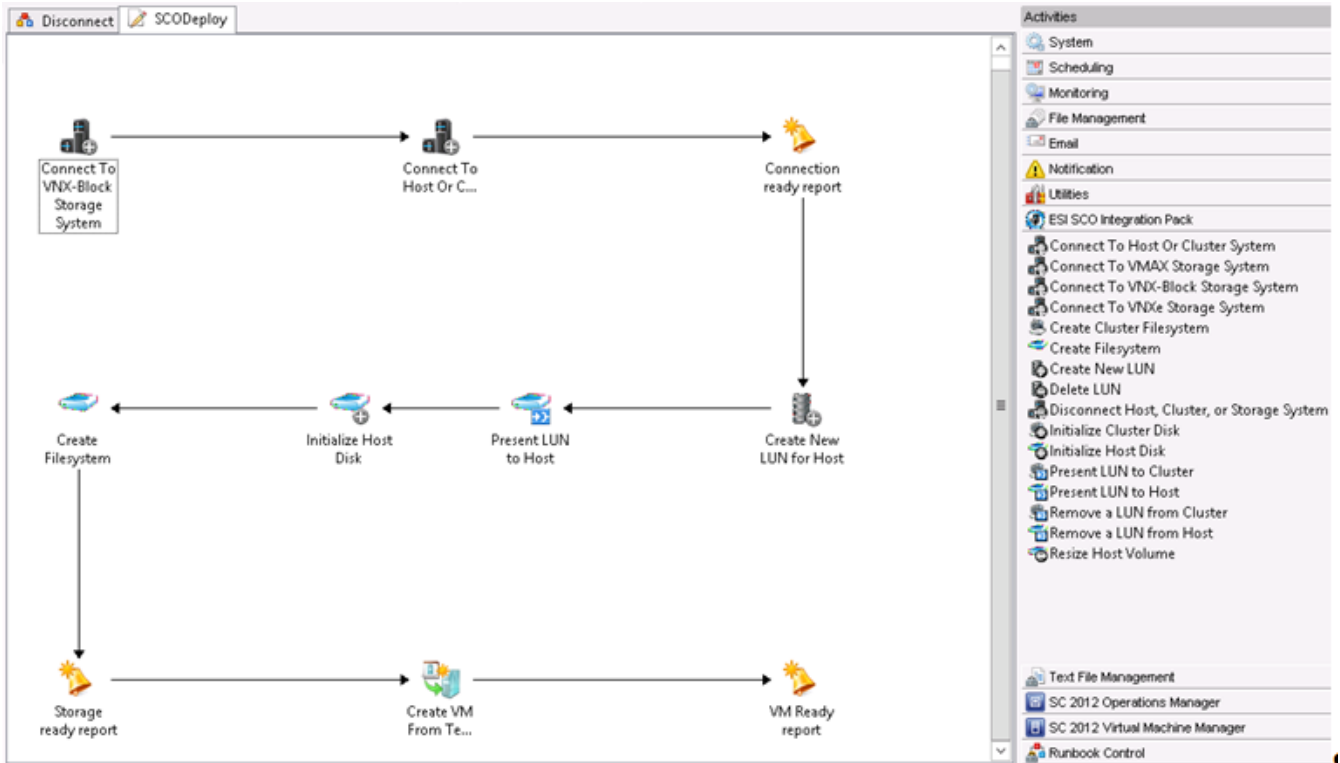


Figure 27. Runbook Designer GUI

For the test results of this runbook, refer to [Validation](#).

Application design and configuration

Overview

The environment was configured to support an application profile that consists of:

- A single SQL Server 2012 instance
- A medium-sized SharePoint farm
- A medium-sized Lync Server
- A medium-sized Exchange consisting of two DAG copies

In this solution, SQL Server, Lync Server, and SharePoint were enabled by FAST Suite and co-located in a single pool consisting of flash, SAS, and NL-SAS drives, on the VNX7600. The Exchange DAG copies were each hosted in their own pool of NL-SAS drives.

Co-location of SQL Server 2012, Lync Server 2013, and SharePoint Server 2013 volumes and segregation of Exchange Server 2013 volumes

Co-location of application workloads such as SQL Server, Lync, and SharePoint can help simplify storage layout. Pay attention to the type of the I/O workloads when doing this. In this solution, an OLTP SQL Server workload was located with Lync and SharePoint in a mixed storage pool consisting of SAS and NL-SAS disks. SQL Server was co-located with Lync and SharePoint to show the performance improvement that the next-generation VNX can bring even in the case of a single pool design with a workload that is too heavy. The tiering policy for the LUNs was initially set to the Lowest Available Tier of NL-SAS disks.

Exchange workloads should be segregated with each DAG copy located in its own pool. In this configuration, two pools of NL-SAS drives were configured on the VNX.

We used FAST Cache to improve the performance of the applications. FAST Cache was applied at a pool level so we could improve the performance of both the block and file pools at the same time.

After adding five flash disks to the mixed storage pool, we used FAST VP to move the busiest data to the highest tier. Then we changed the LUN policy to “Auto-tier” to allow the movement of data to the appropriate tier.

SQL Server 2012

Design and requirements

The SQL Server 2012 configuration consisted of one instance that hosts a single 100 GB database. A heavy OLTP-like workload of more than 1,000 transactions per second and with a read/write ratio of approximately 90:10 was run against the database that has a total of 10,000 users.

Table 15. SQL Server user requirements

Profile characteristic	Quantity/Type/Size
SQL Server Instances	1
OLTP database (OLTP_DB)	10,000 users/100 GB
Workload type	OLTP-like (90:10 read/write ratio)

Configuration

The OLTP user database consisted of eight data files and a log file. The temp database and system databases were located in the same mixed storage pool. Although this is not an SQL Server storage best practice, we wanted to create a worst scenario and alleviate it using FAST Suite.

Note: The temp database in typical OLTP environments generally produce a small number of semi-random IOPS, while the log file produces minimal small sequential writes.

In this solution co-locating both the user and temp database provided a simple design while still guaranteeing the performance.

Table 16 shows the virtual machine configuration.

Table 16. SQL Server virtual machine configuration

Virtual machine role	Quantity	vCPU	Memory (GB)	VHDX	Disk size (GB)
SQL Server	1	8	32	Boot	60
				SQL database	500
				Temp database	500
				SQL logs	20

SharePoint Server 2013 Design

The SharePoint farm was designed to serve a medium-sized corporate organization of up to 10,000 users at 10 percent concurrency with two web servers, one crawl server, one Microsoft SQL Server, and one application server. The farm was configured as a Publishing Portal that contained several Document Center sites.

The farm, excluding boot volumes for the virtual machines, was stored on a single large mixed storage pool.

The farm consisted of three content databases (4 TB, 100 GB, and 200 GB).

Configuration

Without advanced storage technologies, SharePoint requires a complex and costly storage infrastructure to sustain performance, growth, and recoverability.

In this solution, our three guiding principles for storage design and deployment were simplicity, performance, and capacity for growth. The entire farm was stored in the same mixed storage pool. Individual VHDX volumes were created for the content databases, the search databases, the temp databases, and the configuration databases. Individual VHDX files enable granular recovery at a virtual volume level of elements of the SharePoint farm, such as a content database. The crawl server Index volume was also a separate VHDX on the same mixed pool. Both of the web servers were query servers and therefore had their query volumes as VHDX volumes on the same pool.

SharePoint Server 2013 user requirements

Table 17 lists the SharePoint user requirements in this solution.

Table 17. SharePoint user requirements

Item	Value
Total user count	10,000 with 2 web servers
Usage profile(s) (%browse/ % search/%modify/%download/%upload)	60%/15%/10%/10%/5%
User concurrency	10%

Table 18 details the virtual machine disk configurations for the SharePoint farm.

Table 18. Virtual machine configurations for the SharePoint farm

Virtual machine role	Number of virtual machines	vCPU	Memory (GB)	VHDX disk size	Description
Web server (with Query role)	2	8	16	60 GB	Boot
				300 GB	Query
Application	1	1	4	60 GB	Boot
Index	1	8	16	60 GB	Boot
				500 GB	Index
SQL Server	1	8	64	60 GB	Boot
				7 TB	Content databases
				500 GB	Content database logs
				550 GB	Search database

Virtual machine role	Number of virtual machines	vCPU	Memory (GB)	VHDX disk size	Description
				55 GB	Search database logs
				200 GB	Search temp databases
				40 GB	Search temp database logs
				100 GB	Configuration databases
				10 GB	Configuration database logs

Exchange Server 2013

In a continuing Microsoft tradition, Exchange Server 2013 provides a new architecture, enhanced availability features, and further optimized storage I/O operations. Exchange Server 2013 includes significant changes compared to the Exchange 2010 version and requires more server resources, CPU, and memory. It needs additional memory to support changes to the Store database engine to reduce demands for disk I/Os. The rendering of mailbox information and a new search indexing system on the mailbox role requires an increase in CPU requirements. Specifically, the storage schema was improved and modified to accommodate high availability and larger mailboxes with fewer I/Os per second (IOPS) per mailbox.

Exchange Server 2013 user requirements

Table 19 details the Exchange Server 2013 user requirements in this solution.

Table 19. Exchange user requirements

Item	Value
Total number of users (mailboxes) in Exchange environment	10,000
Storage Infrastructure	SAN
Type of deployment (physical or virtual)	Virtual (Windows Server 2012 Hyper-V)
High availability requirements	One DAG with two database copies
Mailbox size limit	2 GB maximum quota
User profile	150 messages per user per day (0.101 IOPS)
Target average message size	75 KB
Number of mailbox servers	8 (4 tested)
Number of mailboxes per server	2,500 (1,250 active/1,250 passive)

Item	Value
Number of databases per server	4
Number of users per database	625
Deleted items retention (DIR) period	14 days
Log protection buffer (to protect against log truncation failure)	3 days
Background Database Maintenance (BDM) configuration	Enabled 24 x7
Database read/write ratio	3:2 (60%/40%) in a DAG configuration
User concurrency requirements	100%
Disk type	3 TB NL-SAS (7,200 rpm)
Storage platform	EMC next-generation VNX

Exchange Server 2013 design

EMC recently released a best practice for Exchange Server 2013 design, which includes storage design, network design, and computing design. We strongly recommend that you refer to this best practice when designing an Exchange environment. In order to have a predictable performance, we have included 32 NL-SAS disks for the Exchange database and 4 NL-SAS disks for Exchange log to meet I/O and capacity requirements. Exchange OS volumes were placed in a separate pool. All those pools were using iSCSI (Block) protocol.

For more information, refer to *Microsoft Exchange Server Best Practices and Design Guidelines for EMC Storage*.

Microsoft Exchange Server Jetstress 2013

In this solution, we used Jetstress 2013 to test our storage design. We simulated the worst scenario: Half of the Exchange Servers were down, and only four servers were left to support 10,000 users. Each virtual machine was assigned four vCPUs and 4 GB memory to run Jetstress 2013, as shown in Figure 28.




Figure 28. Jetstress 2013 test scenario

Lync Server 2013

To support enterprise-level collaboration requirements, Microsoft Lync Server 2013 must be carefully designed and deployed. The best type of deployment depends on the workloads you want to deploy, and your organization’s geographical and business status.

Based on the requirement of 10,000 users, we deployed one front-end pool that contained two front-end servers and one SQL Server. Figure 29 shows the entire Lync system topology. For test purposes, we deployed only what we needed for sample workloads, that is, an SQL back-end server and front-end servers.

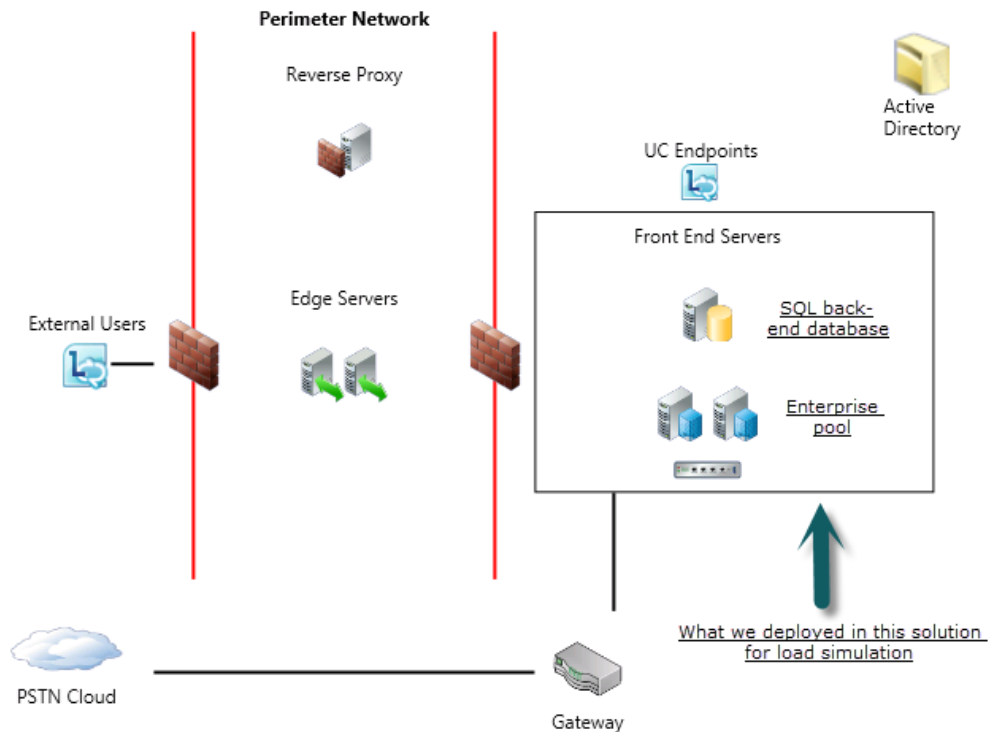


Figure 29. Topology of a Lync system

Configurations

For the Lync Server virtual machines, we set the memory and CPU based on requirements from Microsoft best practices. For more information, see the [Lync Server](#) website.

Table 20. Lync Server virtual machines configuration

Server name	Roles	vCPU	Memory (GB)
LyncFE01	Front-end server	4	32
LyncFE02	Front-end server	4	32

Server name	Roles	vCPU	Memory (GB)
LyncBE	Back-end SQL Server	4	32

We designed the storage for Lync Server so that it shared the same mixed storage pool as SharePoint and SQL Server. This pool contains all Lync Server VHDXs.

Hyper-V Cluster

This solution deploys a Hyper-V cluster consisting of four nodes to increase the availability of virtual machines and applications. When determining where to place the virtual machines, be sure to plan for load balancing and failure protection. You should distribute virtual machines with the same application roles on different Hyper-V root servers, as shown in Table 21.

Table 21. Hyper-V cluster deployment

	Virtual machine role	vCPU	Memory (GB)
Node 1	SharePoint crawl server	8	16
	Lync Server (front-end)	4	32
	Lync Server (back-end)	4	4
	Exchange Server	10	72
	Exchange Server	10	72
Total		36	196
Node 2	SharePoint web server	8	16
	Exchange Server	10	72
	Exchange Server	10	72
	Lync Server (front-end)	4	32
Total		32	192
Node 3	Exchange Server	10	72
	Exchange Server	10	72
	SharePoint SQL Server	10	64
Total		30	208
Node 4	SharePoint App Server	1	4
	SharePoint web server	8	16
	SQL Server	8	32
	Exchange Server	10	72
	Exchange Server	10	72
Total		37	196

Validation

Overview

The Proven Solutions group tested and validated environments from a customer point of view using either industry standard or custom application testing to simulate user loads.

The testing section covers the following areas:

1. Functionality and integration of ESI for Windows Suite
 - a. Storage monitoring
 - b. Automated storage provisioning as part of virtual machine creation
2. Application testing methodology
3. Performance test results

Functional tests

Functional testing of System Center 2012 SP1 with EMC components

To validate the functionality of System Center 2012 SP1 and ESI for Windows Suite, the functional tests we ran covered the following scenarios.

SCOM 2012 SP1 with ESI Management Packs testing

To verify that SCOM 2012 SP1 with ESI Management Packs can provide real-time monitoring of EMC storage systems, we removed the disk and looked for an alert from SCOM.

Figure 30 shows that while we were using a command to remove disk A7 from Bus 4, Enclosure 0, SCOM immediately showed the alert.

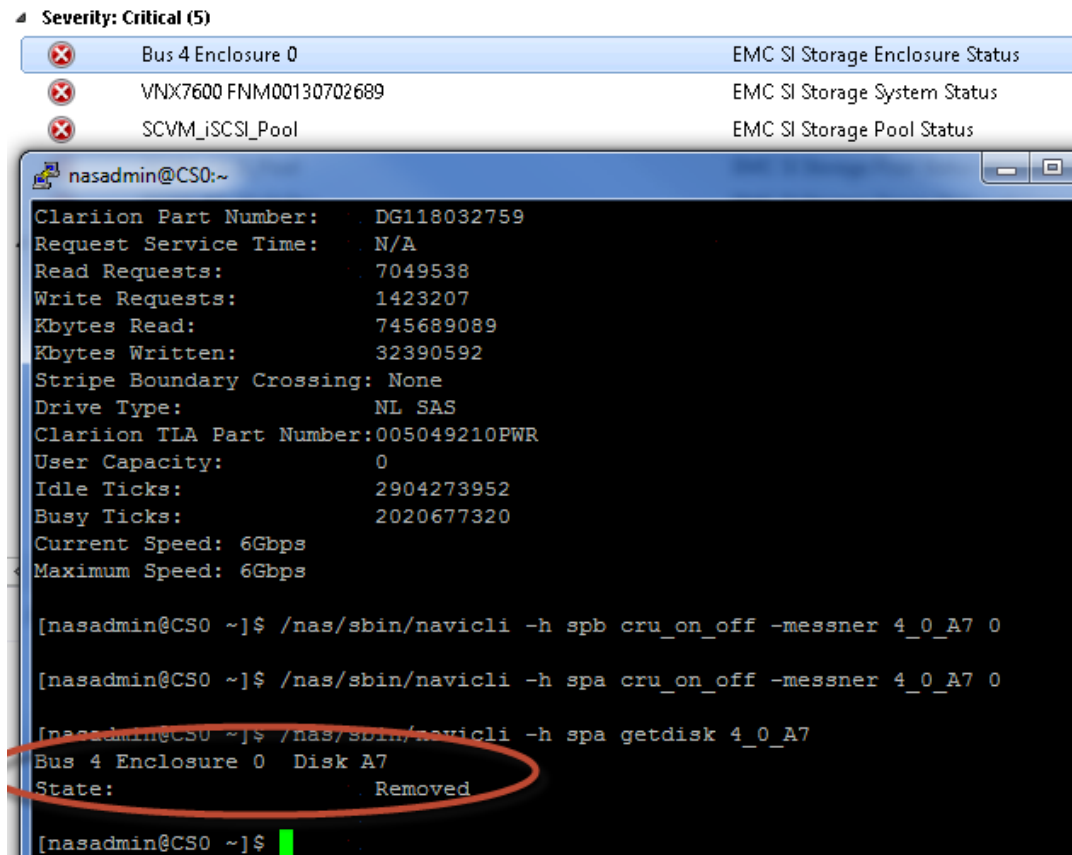


Figure 30. Removing disk A7 from Bus 4, Enclosure 0

SCO 2012 SP1 with ESI Integration Pack

To verify that using ESI Integration Pack for SCO and SCVMM Integration Pack automated storage provisioning, we performed the following steps:

1. We created a new runbook in the SCO Runbook Designer and assigned the following activities in order:
 - a. Connect to storage and create storage for a virtual machine using ESI Integration Pack activities.
 - b. Create a virtual machine from a template using SCVMM Integration Pack activities.
 - c. Create a report using SCOM Integration Pack activities.

ESI Integration Pack enables users to create not only a virtual volume, but also a new EMC storage LUN, and to assign the appropriate virtual volumes to the new storage LUN.

As shown in Figure 31, SCOM creates reports so that users can easily get runbook status from the SCOM console.




Figure 31. Runbook overview

2. After finishing the design of this runbook, we tested it by clicking the **Run** button. Several minutes later, we successfully created the test virtual machine **SCOTEST** with its storage on VNX.

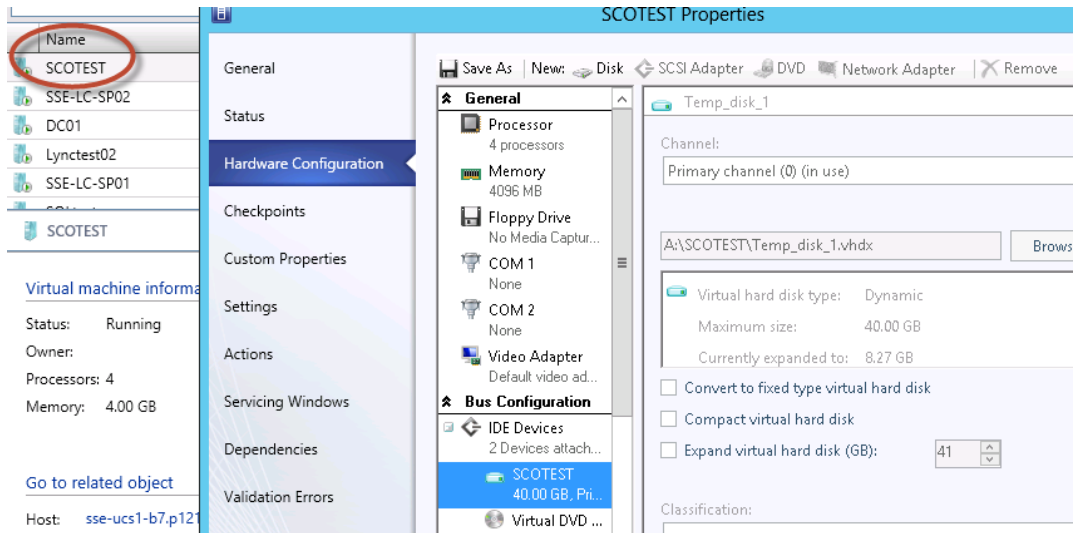


Figure 32. Virtual machine SCOTEST created with its storage on VNX

From the SCOM console, we could also see the status of each major step when it finished.

Figure 33 shows that all three major steps successfully finished.

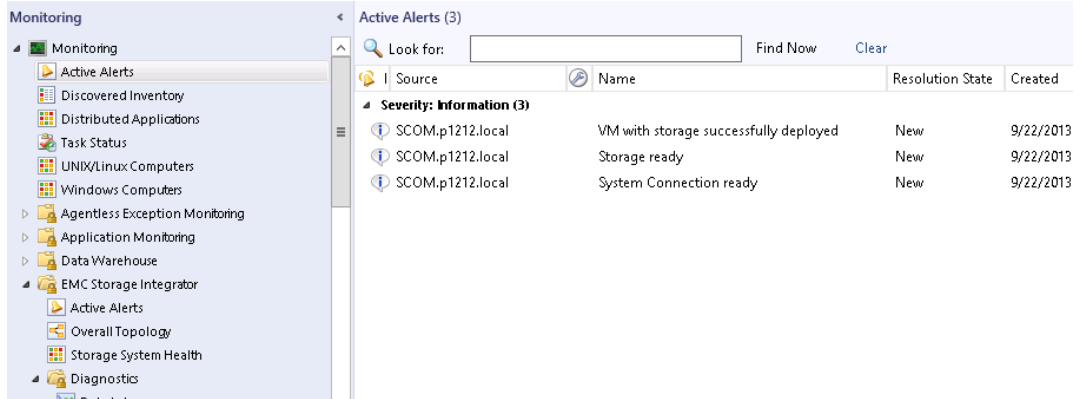


Figure 33. SCOM report of runbook

Test methodology **Application performance across the environment**

This solution also validated the performance of the final configuration across all applications in the environment.

SQL Server 2012

We conducted SQL Server testing by running OLTP-like workloads against a single SQL Server instance. Latency of less than 20 ms is the Microsoft best practice limit for storage latency for transactional databases performance.

SharePoint Server 2013

We used Microsoft Visual Studio Team System (VSTS) to simulate the load on the SharePoint farm. We also used a client load emulation tool to ensure that the SharePoint farm was operating at the optimal performance level.

All users adhered to a Microsoft heavy-user profile, which specifies 60 requests per hour. We applied a think time of zero percent to all tests. “Zero think time” is the elimination of typical user decision-making time when browsing, searching, or modifying data using Microsoft Office SharePoint Server. Every user request is completed from start to end without a pause, which generates a continuous workload on the system.

The maximum user capacity is derived from the following formula:

$$\# = \text{seconds per hour} / \text{RPH} / \text{Concurrency\%} * \text{RPS}$$

Example: $3600 / 60 / 1\% * 34.15 = 204,900$

Example: $3600 / 60 / 10\% * 34.15 = 20,490$ (supported user capacity for 10 percent concurrency)

Table 22 lists the response time limits for each test type.

Table 22. Response time

Test type	Action	Percentage	Response time
Browse	User browse	60	Less than 3 seconds

Test type	Action	Percentage	Response time
Search	Unique value search	15	Less than 3 seconds
Modify	Metadata modify and upload	10	Less than 3 seconds
Download	Download a document	10	Less than 3 seconds
Upload	Upload a document	5	Less than 5 seconds

Microsoft Jetstress 2013

Microsoft Jetstress 2013 was used to simulate the load. Based on Exchange Server 2013 user IOPS, our goal is to achieve at least 1,010 IOPS total (10,000 users plus 0.101 IOPS per user) with average database read latency under 20ms.

Table 23 shows the Jetstress 2013 test profile.

Table 23. Jetstress 2013 test profile

Test tool	Number of users	IOPS per user	Target total IOPS	Average latency	Log write latency
Jetstress 2013	10,000	0.101	1010	Less than 20 ms	Less than 10 ms

Lync Stress and Performance Tool

To simulate Lync client load in this solution, we used the Lync Server 2013 Stress and Performance Tool to generate and simulate user loads. This tool can help you with capacity planning for Lync Server 2013 by:

- Simplifying your hardware planning for Lync Server 2013
- Providing you with increased knowledge and best practices for performance tuning
- Measuring the performance of your intended Lync Server 2013 deployments

Test scenario

We simulated 10,000 users with high load. Each test scenario included:

- Instant messaging (IM)—Provided an instant messaging platform with conversation history, including support for public IM connectivity with users of public IM networks such as MSN/Windows Live, Yahoo!, AOL, and Google Talk.
- Application sharing—Enabled users to share Microsoft Office applications such as Microsoft Word, Microsoft Excel, and Microsoft OneNote.
- Audio conference—Provided a user experience that is familiar to users of traditional audio bridge services including Public Switch Telephone Network (PSTN) dial-in services with touch-tone call control commands.

- Data collaboration—Enabled users to create whiteboards and polls, or share Microsoft PowerPoint slides.

Table 24 shows the number of users in each scenario.

Table 24. 10,000-user scenario

Name	Load level	Number of users
Audio conferences	High	375
Application sharing	High	185
Data collaboration	High	75
Instant messaging	High	9,365

Note: The Lync Stress and Performance Tool includes three load types: Low, Medium, and High. In this solution, we used the High load for simulation; the tool itself defines the number of users in each test scenario.

For detailed instructions about using the Lync Server 2013 Stress and Performance Tool, see the [Microsoft TechNet](#) website.

Performance test results

Combined application test results

The test results show the improvement in performance (either IOPS, latency, or both) for all the applications. SQL Server shows the best improvement in terms of latency and IOPS. We purposely put SQL Server in the mixed pool and on the lowest tier to highlight a customer issue and show the massive improvement that FAST Suite on the next-generation VNX brings to it.

FAST Cache is an easy and quick improvement that customers can apply to their VNX to improve performance for pools. FAST VP will monitor and find which slices of data to relocate and will normally relocate on a fixed schedule so improvement may take a while longer to be evident.

Figure 34 shows the average host latencies for SQL Server, Exchange, Lync, and SharePoint.




Figure 34. Average host latencies

As you can see from Figure 34, the three SMB 3.0 applications (SharePoint, SQL Server, and Lync) have a reduction in host latencies. Exchange shows a small increase as some SMB 3.0 applications also increase their IOPS with the lower latencies. Exchange used the iSCSI protocol to connect to storage and was in a separate pool with only NL-SAS drives. FAST VP was not used for this pool as it was not tiered. Exchange shows a higher latency (still under 20 ms), but the increase of IOPS shows an improvement overall.

SharePoint and Lync show 90% and 87% respective improvement in latencies while still satisfying the user load requirement for 10,000 users. Both FAST Cache and FAST VP lower the latencies for these applications.

SQL Server shows an improvement in latency of 97% demonstrating how much FAST Suite benefits OLTP workloads.

The SQL Server latency in this chart and the chart in Figure 35 was measured as the average transfers per second of the database disk only. The Exchange latency was measured using the average read latency only as this was the slowest latency on the Exchange servers. The Lync latency was measured using average read latency only. The SharePoint latency was measured using the SharePoint SQL Server Content database data disk only.

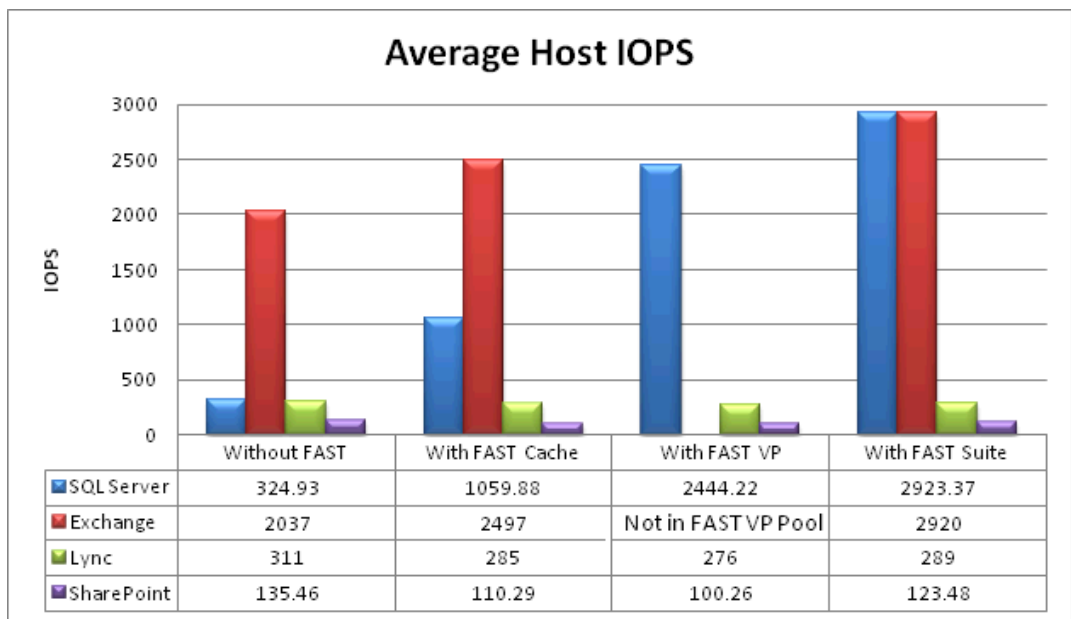


Figure 35. Host IOPS

Figure 35 shows the increase in IOPS as the latency decreases for most applications. SQL Server shows the most dramatic increase in IOPS. Exchange also shows a significant increase in IOPS. Exchange was not relevant to the FAST VP test because it was not in a tiered pool like the other applications.

The increase in IOPS from the SMB 3.0 mixed storage pool, however, did impact Exchange latencies.

Conclusion

Summary

- The environment in this solution validates Windows Server 2012 Hyper-V, and an EMC VNX storage array hosting the SQL Server, Lync Server, and SharePoint Server environments on the SMB 3.0 file system, as well as an Exchange Server 2013 environment on the iSCSI block system.
- This solution highlights the close integration of ESI for Windows Suite with Microsoft System Center 2012 SP1, which allows administrators to control applications and EMC next-generation VNX storage infrastructure throughout the data center.
- Automatically provisioned EMC storage also provides easy deployment of Windows virtual machines.
- This solution showcases the simplified storage design that the EMC FAST Suite technology enables. In addition, it uses a single three-tier EMC storage pool on a VNX to co-locate a SQL Server environment supporting OLTP workloads with a SharePoint Publishing Portal deployment and a Lync Server environment.
- FAST Suites' components, FAST Cache and FAST VP, make it possible for applications to be hosted on VNX File shares over SMB 3.0 with acceptable performance and low latencies.

Findings

In this solution, we found that:

- The ESI Management Packs enable the SCOM 2012 SP1 to monitor the health and performance of next-generation VNX storage systems.
- The EMC SMI-S Provider allows SCVMM 2012 SP1 to easily manage and provision next-generation VNX storage.
- ESI Integration Pack for SCO allows automated provisioning and allocation of EMC storage with SCO 2012 SP1.
- SCVMM 2012 SP1 enables simplified and centralized management of Hyper-V virtualization platforms.
- Unified storage with FAST Suite technology delivers on average more than a 91% latency reduction for applications hosted on SMB 3.0.
- Live Failover Clustering of virtual machines is easy to configure with the virtual machine disks on CIFS shares through SMB 3.0.

Reference

White papers

The following white papers, available from the EMC.com website, provide additional and relevant information. If you do not have access to a document, contact your EMC technical support professional.

- *EMC VNX Series: Introduction to SMB 3.0 Support*
- *EMC VNX Unified Best Practices for Performance*
- *Introduction to the EMC VNX Series—A Detailed Review*
- *EMC VNX Virtual Provisioning—Applied Technology*
- *EMC VNX FAST VP—A Detailed Review*
- *Microsoft Exchange Server Best Practices and Design Guidelines for EMC Storage*

Product documentation

For additional information, see the product document websites listed below:

- [SharePoint 2013](#)
- [Lync Server 2013](#)
- [Windows Server 2012](#)

Other websites

For additional information, see the [SMB 3.0](#) and [Failover Clustering](#) websites.