Lenovo Cloud Configuration Guide for Microsoft Hyper-V on Flex System

Last Update: 23 May 2016

Based on Flex System x240 M5 Compute Nodes and NetApp FAS8040 Storage

Uses Windows Server 2012 R2 and Microsoft System Center 2012 R2

Provides architecture and implementation details

Includes a complete Bill of Materials

Scott Smith

David Ye

LENOVO PRESS

Note: Before using this information and the product it supports, read the information in “Notices” on page 139.

Last update on 23 May 2016

This edition applies Microsoft Hyper-V Fast Track running on Flex System x240 M5 Compute Nodes

Contents

Preface

Hyper-V Fast Track (HVFT) for Windows Server 2012 R2 on Lenovo Flex System™ is a pre-tested solution from Lenovo® that improves time-to-value for your private cloud virtualization needs. This paper describes a validated HVFT solution that uses Flex System compute nodes and switching environment with NetApp FAS storage to support the virtualization needs of small-to-medium sized businesses. This configuration provides a highly available fault tolerant solution to support up to 500 virtual machines with the ability to scale up as future demands require.

The design focuses on removing the complexity and uncertainty of deploying and managing a Hyper-V based virtualization environment. The use of a converged Flex System configuration with FCoE helps minimize the number of switches and cables that are needed in a typical environment. Microsoft System Center 2012 R2 with the Lenovo XClarity Upward Integration Modules by Lenovo provide the management environment to monitor and maintain the virtual infrastructure.

This reference configuration architecture and implementation guide targets organizations that are implementing Hyper-V and IT engineers who are familiar with the hardware and software that make up the Lenovo Private Cloud Architecture with Microsoft Hyper-V. Comprehensive experience with the Lenovo Flex System and NetApp FAS technologies is recommended. Also, the Lenovo System x® sales teams and their customers evaluating or pursuing Hyper-V virtualization solutions can benefit from this previously validated configuration.

The team who wrote this paper

This paper s produced by the following subject matter experts:

Scott Smith is a Senior Solutions Architect working at the Lenovo Center for Microsoft Technology. Over the past 15 years, Scott worked to optimize the performance of Lenovo x86 based servers that are running the Microsoft Windows Server operating system and Microsoft application software. Recently, his focus is on Microsoft Hyper-V based solutions with Lenovo System x servers, storage, and networking. He has extensive experience in helping Lenovo customers understand the issues that they are facing and developing solutions that address those issues.

David Ye is a Senior Solutions Architect and has been working at Lenovo Center for Microsoft Technologies for 15 years. He started his career at IBM as a Worldwide Windows Level 3 Support Engineer. In this role, he helped customers solve complex problems and was involved in many critical customer support cases. He is now a Senior Solutions Architect in the System x Enterprise Solutions Technical Services group, where he works with customers on Proof of Concepts, solution sizing, performance optimization, and solution reviews. His area of expertise are Windows Server, SAN Storage, Virtualization, and Microsoft Exchange Server.

Thanks to the following people for their contributions to this project:

- Dan Ghidali, Senior Development Engineer
- Mike Miller, Advisory Software Engineer
- Vinay Kulkarni, Senior Performance Engineer
- David Watts, Lenovo Press
- Karen Lawrence, IBM Redbooks

Thanks to the following members of the NetApp Team:

- Craig Thompson, Enterprise Engineer
- Pavel Lobanov, Reference Architect
- Jacky Ben-Bassat, Market Development Cloud Architect
- Glenn Sizemore, Technical Marketing Engineer

Comments welcome

Your comments are important to us!

We want our papers to be as helpful as possible. Send us your comments about this paper or in one of the following ways:

- Use the online feedback form found at the web page for this document:

  http://lenovopress.com/redp5173

- Send your comments in an email to:

  comments@lenovopress.com

Summary of changes

This section describes the technical changes made in this edition of the Interoperability Guide and in previous editions. This edition might also include minor corrections and editorial changes that are not identified.

23 May 2016

- Updated the bill of materials to include x240 M5 Compute Nodes with Intel Xeon E5-2600 v4 series processors.

22 September 2015

New information

- Reference Architecture is now based on Flex System x240 M5 Compute Nodes

1

Lenovo Flex System solution for Microsoft Hyper-V

Hyper-V Fast Track (HVFT) for Windows Server 2012 R2 on Flex System is a pre-tested solution from Lenovo that improves time-to-value for your private cloud virtualization needs. This paper describes a validated HVFT solution that uses Flex System compute nodes and switching environment with NetApp FAS storage to support the virtualization needs of small- to medium-sized businesses. This configuration provides a highly available, fault tolerant solution to support up to 500 virtual machines (VMs) with scaling capability future demands change.

The design focuses on removing the complexity and uncertainty of deploying and managing a Hyper-V-based virtualization environment. The use of a converged Flex System configuration with Fibre Channel over Ethernet (FCoE) helps minimize the number of switches and cables that are needed in a typical environment. Microsoft System Center 2012 R2, with Lenovo XClarity Upward Integration Modules by Lenovo, provide the management environment to monitor and maintain the virtual infrastructure.

This reference configuration architecture and implementation guide targets organizations that are implementing Hyper-V and IT engineers who are familiar with the hardware and software that make up the Lenovo Private Cloud Architecture with Microsoft Hyper-V. Comprehensive experience with the Lenovo Flex System and NetApp FAS technologies is recommended. Also, the System x sales teams and their customers who are evaluating or pursuing Hyper-V virtualization solutions can benefit from this previously validated configuration.

This chapter includes the following topics:

1.1 Business problem

Today's IT managers are looking for efficient ways to manage and grow their IT infrastructure with confidence. Good IT practices recognize the need for high availability, simplified management, and containing costs through maximum resource utilization. CIOs must rapidly respond to changing business needs with simple, easily deployed configurations with the ability to scale on-demand. Natural disasters, malicious attacks, and even simple software upgrade patches can cripple services and applications until administrators resolve the problems and restore any backed up data.

The challenge of maintaining healthy systems and services becomes more critical as businesses consolidate physical servers into a virtual server infrastructure as part of their journey towards cloud-based solutions. These solutions feature a premise of a flexible consumption model that reduces data center costs and keeps the infrastructure flexible to accommodate those ever-changing business needs.

1.2 Business value

This private cloud reference architecture for Microsoft Hyper-V is based on converged infrastructure technologies from Lenovo and NetApp that are integrated to form an Infrastructure-as-a-Service (IaaS) solution that is cloud-ready. The infrastructure stack includes the Flex System Chassis that is populated with the latest Compute Nodes, Converged Switches, and a NetApp unified storage that is based on Clustered Data ONTAP.

The solution provides a flexible, affordable, and reliable industry-leading cloud-ready platform that is simple to deploy and manage The integration between the hardware components and Microsoft Hyper-V and System Center takes the complexity out of the solution to allow IT functions to focus on optimizing the delivery and the consumption models of the services that best meet the needs of their customers.

This reference architecture is validated under the Microsoft Private Cloud Fast Track program and covers all the components, including the Microsoft software. Another key aspect that this paper addresses is redundancy and fault tolerance across the servers, storage, and networking for the Windows Servers to help ensure a defined level of fault tolerance while managing pooled resources. This document provides consolidated guidance for all the configurations and the authors added step-by-step implementation instructions for the major components.

By pooling computing, networking, and storage capacity with Microsoft Hyper-V in a Windows Failover Cluster helps to eliminate single points of failure so that users have near-continuous access to important server-based, business productivity applications. An independent cluster that is hosting the management fabric that is based on Microsoft System Center 2012 R2 with Lenovo upward integration components provides an environment to deploy, maintain, and monitor the production private cloud. Similar integration, operational efficiency, and high availability are also part of the NetApp storage layer.

IT administration can be improved by simplifying the hardware configuration to a corporate standard with automated deployment and maintenance practices. Templates of pre-configured VMs can be saved and deployed rapidly through self-service portals to the end customers. VMs can be migrated among clustered host servers to support resource balancing, scheduled maintenance, and if there is an unplanned physical or logical outage, VMs can automatically be restarted on the remaining cluster nodes.

It is easy to add compute resources and storage controllers and capacity in a nondisruptive way, and even perform many of the routine maintenance tasks without disrupting the business. As a result, clients minimize downtime, which makes this seamless operation attractive to organizations that are trying to create business and maintain healthy service level agreements.

1.3 Lenovo solution for Microsoft Hyper-V on Flex System

Microsoft Hyper-V technology continues to gain competitive traction as a key component in many customer virtualization environments. Hyper-V is a standard component in Windows Server 2012 R2 Standard and Datacenter editions. Windows 2012 R2 Hyper-V VMs support up to 64 virtual processors (vSPUs) and up to 1 TB of memory.

Individual VMs have their own operating system instance and are isolated from the host operating system and other VMs. VM isolation helps promote higher business-critical application availability while the Microsoft failover clustering feature (which is available in the Windows Server 2012 R2) can dramatically improve production system uptimes.

This Hyper-V Reference Architecture and Implementation guide provides ordering, setup, and configuration information for the highly available virtualization compute environment that is validated as a Microsoft Hyper-V Fast Track Medium configuration. The Microsoft Hyper-V Fast Track Medium configuration provides a validated two-node clustered management fabric that is built around Microsoft System Center 2012 R2, and an eight-node clustered compute fabric for deployment of production resources. This configuration is ideal for large organizations that are ready to take their virtualization to the next level. The design consists of ten Flex System x240 M5 Compute Nodes that are attached to NetApp FAS8040 storage with Clustered Data ONTAP (cDOT).

Networking uses the Flex System CN4093 converged switches. This fault tolerant hardware configuration is clustered by using Microsoft's Windows Server 2012 R2. This configuration can be expanded to multiple chassis for more compute capacity or storage.

In this section, a summary of the Lenovo Hyper-V Reference Architecture software and hardware components and best practice implementation guidelines are provided.

The Lenovo Hyper-V Reference Configuration is constructed with the following enterprise-class components:

- One Flex System Enterprise System Chassis
- A total of 10 Flex System x240 M5 Compute Nodes in a Windows Failover Cluster running Hyper-V:
  - Two clustered Flex System Compute Nodes are used to build the highly available management fabric.
  - Eight clustered Flex System Compute Nodes are used to build a highly available virtualization cluster.
- Two Flex System CN4093 switches that provide fully converged and redundant networking for data and storage (FCoE).
- NetApp FAS8040 Storage System (dual controllers) in a switchless configuration:
  - Clustered Data ONTAP (cDOT) storage OS version 8.2
  - Four DS2246 shelves with SAS disks
  - One DS2246 shelf with SSD

These software and hardware components form a high-performance solution that supports Microsoft Hyper-V environments for most business-critical applications and many custom third-party solutions. Equally important, these components meet the criteria set by Microsoft Private Cloud Fast Track program that promotes robust virtualization environments to help satisfy even the most demanding virtualization requirements. The overall configuration is shown in Figure 1-1.

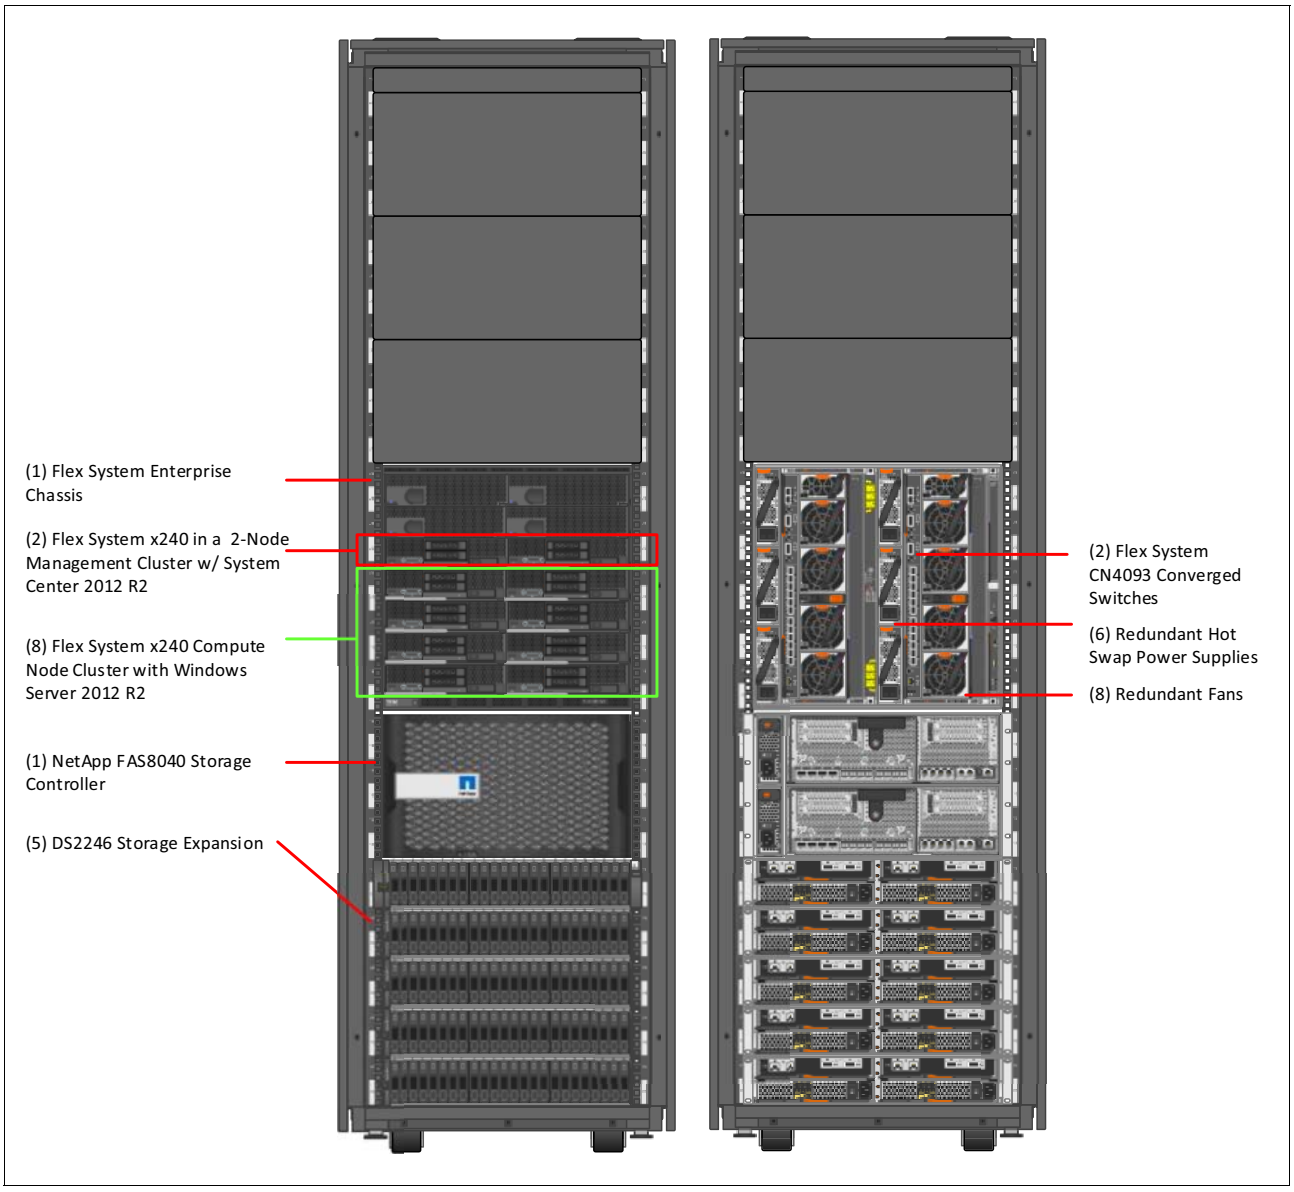

Figure 1-1 Lenovo Hyper-V Fast Track reference configuration

Successful Microsoft Hyper-V deployment and operation is significantly attributed to a set of test-proven planning and deployment techniques. Proper planning includes the sizing of needed server resources (CPU and memory), storage (space and IOPS), and network bandwidth that is needed to support the infrastructure. This information can then be implemented by using industry standard best practices to achieve optimal performance and reserve capacity necessary for the solution.

Note: For more information about this System x solution or to purchase this solution, contact your Lenovo Business Partner and refer to Solution Reference Number CLDHYPFM051.

2

Racking and power distribution considerations

The installation of power distribution units (PDUs) and associated cables must be completed before any system is racked. The prerequisites that must be met before you begin cabling PDUs are described in this chapter.

2.1 Tasks before you begin cabling PDUs

Before cabling the PDUs, consider the following points:

- Ensure that there are separate electrical circuits and receptacles providing enough power to support the required PDUs.
- Redundant electrical circuits to support power to the PDUs are recommended to minimize the possibility of a single electrical circuit failure that can affect the configuration.
- Plan for individual electrical cords from separate PDUs for devices that have redundant power supplies.
- Maintain appropriate shielding and surge suppression practices, and appropriate battery back-up techniques

For more information, see *Flex System Enterprise Chassis & PureFlex Power Requirements Guide*, which is available at this website:

http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102111

For more information about NetApp storage power-related and environmental issues, see this website:

http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102111

3

Components

This highly available virtualization architecture consists of the Lenovo Flex System Enterprise chassis with Flex System CN4093 converged switches, Flex System x240 M5 Compute Nodes that are running Microsoft's Windows Server 2012 R2 operating system, and the NetApp FAS8040 Storage System. Each component provides a key element to the overall solution.

This chapter includes the following topics:

3.1 Flex System Enterprise chassis

The Flex System Enterprise Chassis is a simple, integrated infrastructure platform that supports a combination of compute and networking resources. Working with the NetApp FAS8040 storage, they meet the demands of your application workloads. More chassis can be added as the workloads increases.

The 14-node, 10U chassis delivers high-speed performance that is complete with integrated servers and networking. This flexible chassis is designed for a simple deployment now and for scaling to meet future needs. In addition, the optional Lenovo XClarity Upward Integration Modules for Microsoft System Center provides the integration of the management features of the Flex System into a Microsoft System Center environment. These upward integration modules enhance Microsoft System Center server management capabilities by integrating Lenovo hardware management functionality, which provides affordable, basic management of physical and virtual environments and reduces the time and effort that is required for routine system administration. The Upward Integration Module (UIM) provides discovery, deployment, configuration, monitoring, event management, and power monitoring that is needed to reduce cost and complexity through server consolidation and simplified management.

3.2 Flex System Chassis Management Module

The Flex System Chassis Management Module (CMM) is a hot-swap module that configures and manages all installed chassis components. The CMM provides resource discovery, inventory, monitoring, and alerts for all compute nodes, switches, power supplies, and fans in a single chassis. The CMM provides a communication link with each component's management processor to support power control and out of band remote connectivity, as shown in Figure 3-1.

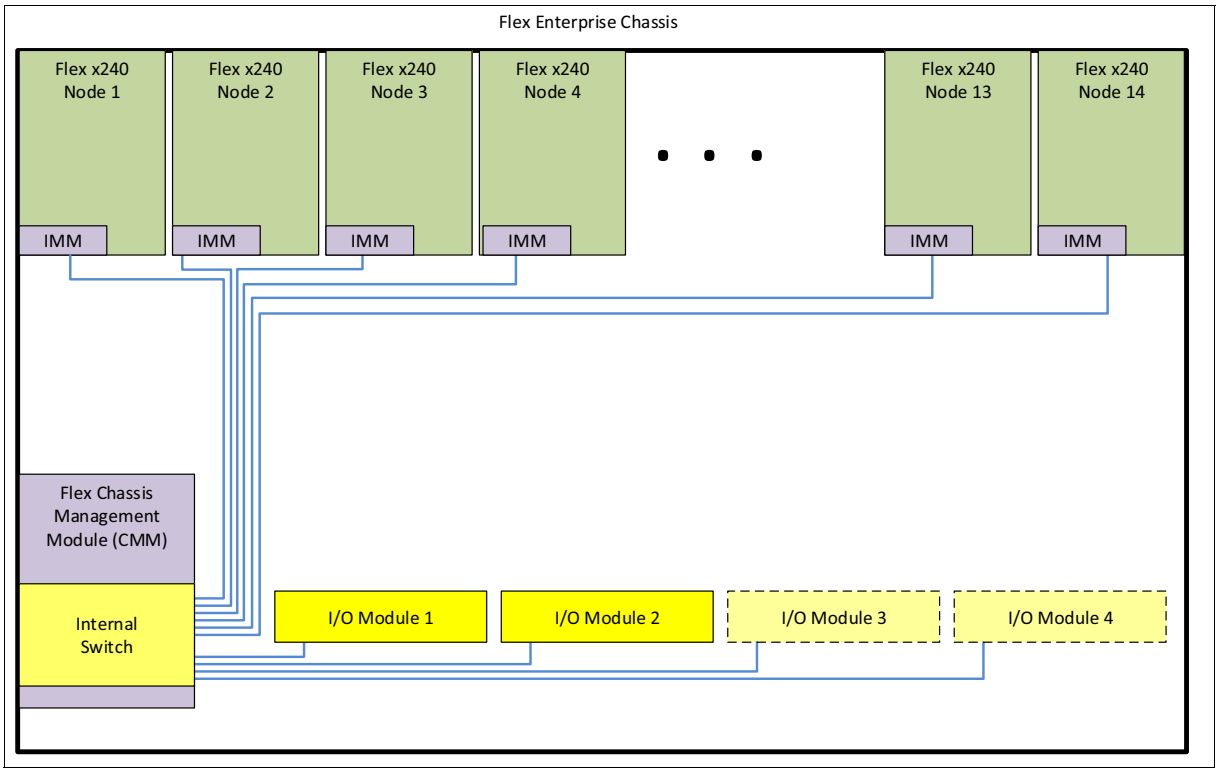

Figure 3-1 CMM management network

Note: The default IP address for the CMM is 192.168.70.100. The default user ID is USERID and the password is PASSW0RD (with a zero).

3.3 Flex System x240 M5 Compute Node

At the core of this reference configuration for Hyper-V, the x240 M5 Compute Nodes (Figure 3-2) deliver the performance and reliability that is required for virtualizing business-critical applications in Hyper-V environments. To provide the expected virtualization performance for handling any Microsoft production environment, the x240 M5 can be equipped with up to two 22-core E5-2600 v4 series processors and up to 1.5 TB of memory. The x240 M5 includes an on-board RAID controller and the choice of hot swap SAS or SATA disks and small form factor (SFF) hot swap solid-state drives (SSDs). Two I/O slots provide ports for data and storage connections though the Flex Enterprise Chassis switches.

The x240 M5 also supports remote management via the Lenovo Integrated Management Module, which enables continuous out of band management capabilities. All of these key features and many others help solidify the dependability that Lenovo customers are accustomed to with System x servers.

By virtualizing with Microsoft Hyper-V technology on Flex System x240 M5 Compute Nodes, businesses reduce physical server space, power consumption, and the total cost of ownership (TCO). Virtualizing the server environment can also result in lower server administration overhead, which gives administrators the ability to manage more systems than in a physical server environment. Highly available critical applications that are on clustered host servers can be managed with greater flexibility and minimal downtime with Microsoft Hyper-V Live Migration capabilities.

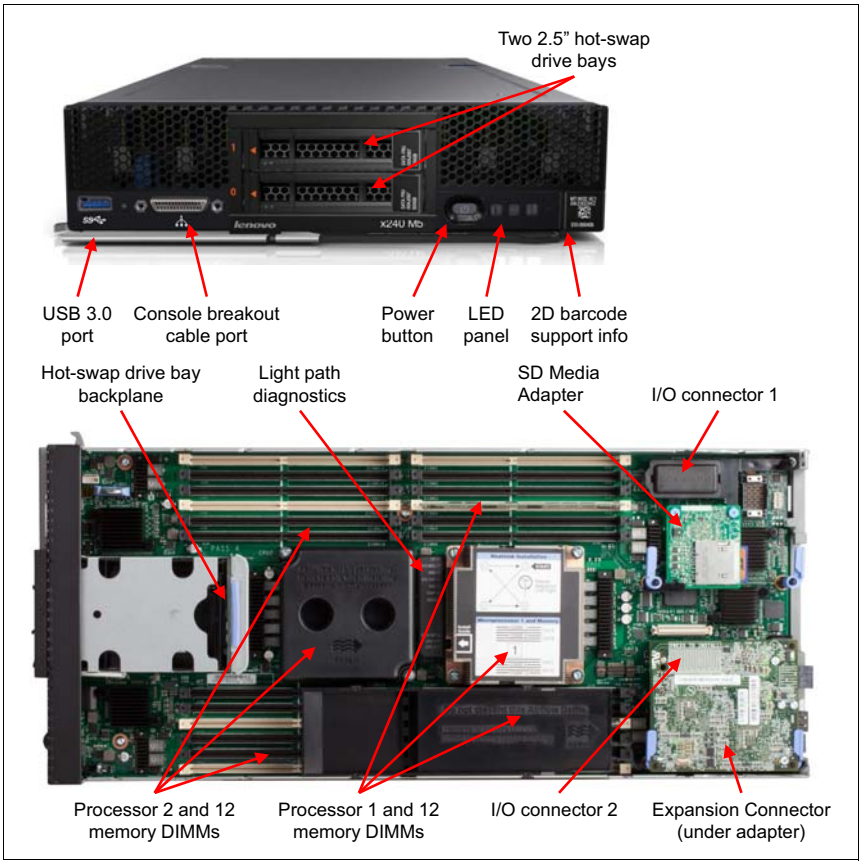

Figure 3-2 Flex System x240 M5

3.4 NetApp FAS Storage running clustered data ONTAP

The storage for this solution is based on the NetApp FAS8040, a mid-range enterprise-class storage system that is part of the FAS8000 family that was introduced by NetApp in 2014. With the latest storage operating system (Clustered Data ONTAP [cDOT] version 8.2), the system provides the foundation for flexible storage infrastructure that meets the demanding characteristics of future cloud solutions that many organizations are seeking today regarding their data.

Non-Disruptive Operations (NDO), High Performance Connectivity, Unified Architecture, Storage Efficiencies, and Integration with Microsoft Hyper-V and System Center are the key features of the NetApp storage system that are described in this solution. The reference architecture focuses on Fibre Channel over Ethernet (FCoE)-based storage, serving network-attached storage (NAS), and SAN protocols.

3.4.1 Overview of the FAS8040 Storage System

The storage system includes two FAS8040 storage system controllers, which are referred to as *nodes* in a storage cluster. The two controllers are configured in an active-active mode so that each is the primary owner of its own disks. However, if there is a controller failure, the other controller assumes ownership of the entire storage system for high availability (HA).

In the FAS storage and Clustered Data ONTAP, the controllers are paired for HA, called *HA Pair*. The FAS8040 is a single chassis enclosure that host both controllers. In other solutions that use a single controller (so there is no HA), half of the chassis is empty and a second controller can be added non-disruptively in the future.

The FAS8040 supports up to 4.3 PB of raw capacity per HA pair. Up to 24 storage nodes (or 12 HA pairs) in pure NAS implementations, or up to eight storage nodes for implementations with block (SAN) protocols to form a single scale-out storage cluster. SSD, SAS, and SATA disks can be used separately to optimize the performance of different applications, and they can be combined to form a pool of automated storage tiering (known as FlashPool). Flash Cache PCIe cards can be added to further enhance performance of certain type of workloads. Flash Pool and Flash Cache are used in this solution. For more information about NetApp's flash technology, see "Related publications" on page 141.

Figure 3-3 and Figure 3-4 on page 14 show the FAS8040 storage system in this solution.

Figure 3-3 NetApp FAS8040

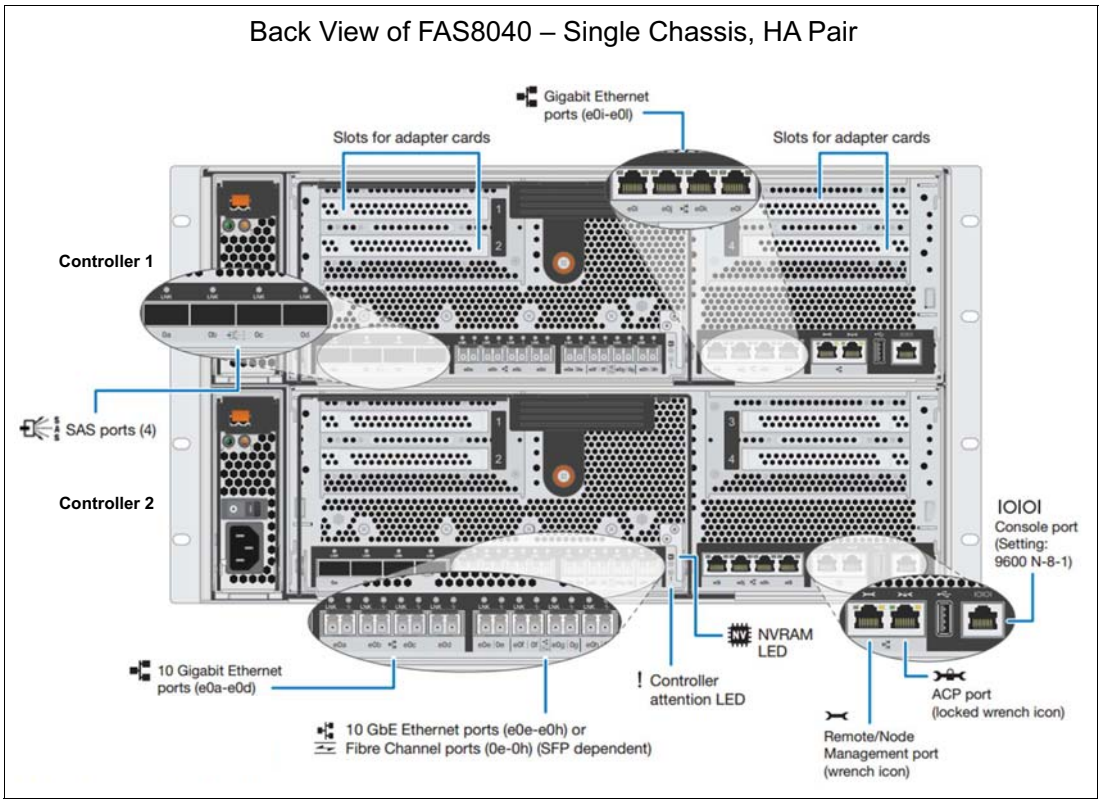

Figure 3-4 FAS8040 HA Pair I/O ports and expansion options

3.4.2 Key capabilities and benefits

As a storage system, the NetApp FAS8040 with cDOT offers the following benefits:

- Support SAN and NAS protocols (FC, FCoE, iSCSI, SMB, and NFS) from a single platform because of its unified architecture
- Integrate with Microsoft System Center to provide automation of storage provisioning and cloning that speeds up the deployment of the various storage services
- Simplify management with an intuitive graphical user interface for performing common and advanced tasks
- Improve storage usage by enabling data deduplication, compression, cloning, and thin provisioning
- Improve availability of cloud-based services because of the Non-Disruptive Operations (NDO) capabilities and the built-in quality of service (QoS) features, which ensures maximum alignment with SLAs
- Optimize cost and performance for mixed workloads by using automated tiering with SAS, SSD, and other Flash technology
- Use the power of cDOT to scale the storage quickly and easily, which addresses capacity, performance, and operations aspects
- Protect data by using built-in features (including point-in-time backup and DR) and extend these features to enable integration at the Hyper-V level

3.5 Flex System CN4093 switches

The Flex System Fabric CN4093 10Gb Converged Scalable Switch (as shown in Figure 3-5) provides unmatched scalability, performance, convergence, and network virtualization, while delivering innovations to address a number of networking concerns today and providing capabilities that help prepare for the future. The switch offers full Layer 2/3 switching and FCoE Full Fabric and Fibre Channel NPV Gateway operations to deliver a truly converged integrated solution. It installs within the I/O module bays of the Flex System Enterprise Chassis and can help clients migrate to a 10 Gb or 40 Gb converged Ethernet infrastructure.

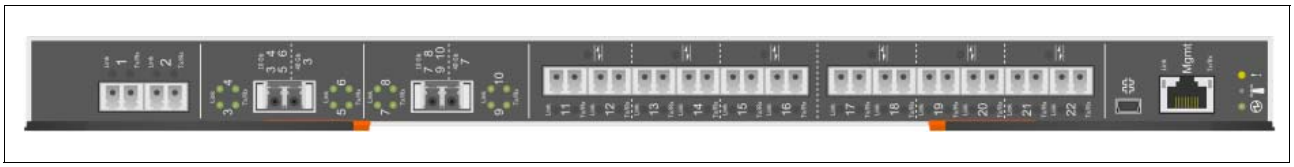

Figure 3-5 Flex System CN4093 Switch

3.6 Microsoft Windows Server 2012 R2

Windows Server 2012 R2 with Hyper-V provides the enterprise with a scalable and highly dynamic platform to support virtualization of most environments with support for up to 4 TB of RAM, 320 logical processors, 1024 virtual machines (VMs) per host, and 64 host nodes per cluster. IT organizations can simplify virtualization resource pools by using key features, such as HA clustering, simultaneous Live Migration (with compression), in-box network teaming, and improved network QoS features. VMs that are running Windows Server 2012 R2 with Hyper-V also increased their resource usage with support for up to 64 vCPU, 1 TB of RAM, shared VHDX, and virtual HBA (vHBA) support.

3.7 Microsoft System Center 2012 R2

Microsoft System Center 2012 R2 with Lenovo Upward Integration Modules enables you to create a comprehensive management environment for the Flex System for Microsoft Hyper-V environment with the following features:

- Platform monitoring and management with System Center Operations Manager
- Bare Metal and Virtualization deployment and management with System Center Virtual Machine Manager
- Self-Service Portal and incident tracking with System Center Service Manager
- Automation management with System Center Orchestrator

3.7.1 System Center Operations Manager Upward Integration Module

The Lenovo Hardware Management Pack for Microsoft System Center Operations Manager enables customers to use the enhanced features of Microsoft System Center Operations Manager for managing the health state of System x servers, Blades, Lenovo BladeCenter® Chassis, Compute Nodes, and Flex System Chassis.

The Hardware Management Pack includes the following key features:

- Rich monitoring of the health for BladeCenter Chassis, Flex System Chassis, and modules by using the SNMP protocol.
- Extensive monitoring of the health of hardware components for the System x servers, BladeCenter x86/x64 blades, and Flex System x86/x64 compute nodes that are running Windows.
- Comprehensive monitoring of the health of software stacks for managing Lenovo hardware.
- Easy determination of overall system health by the aggregation of hardware health monitors.

3.7.2 System Center Virtual Machine Manager Upward Integration Module

The Hardware PRO Pack for Microsoft System Center Virtual Machine Manager (SCVMM) provides performance and resource-optimization (PRO) tips that describe the hardware health state, alerting information, and VM migration recommendations for the physical systems hosting Hyper-V.

Some of the key features are:

- Detailed system health information for physical servers hosting virtual machines while managing Hyper-V or Virtual Server environments.
- Comprehensive health state monitoring of physical processors, memory, hard disks, fans, power supplies, temperature, and voltage.
- Advisory PRO tips about existing or potential hardware problems that warrant Virtual Machine Manager (VMM) administrative operations.

4

Networking and VLANs

In this chapter, we describe the procedures that are used to configure the network and VLANs.

This chapter includes the following topics:

4.1 Flex System switch positions and network connections

The Flex System chassis contains up to four switches. The numbering of these switches is interleaved (as shown in Figure 4-1) and must be kept in mind when work is performed on the switches or cable connections are added to the external ports.

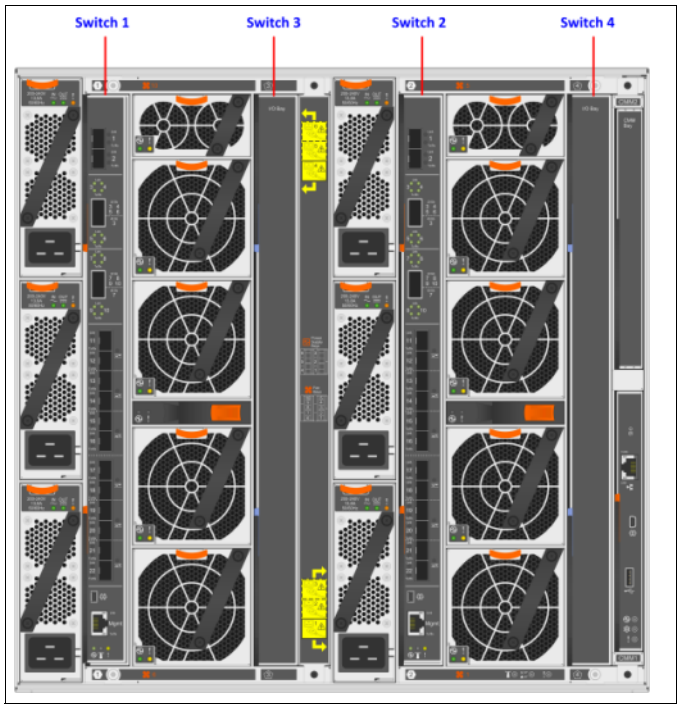

Figure 4-1 Flex System chassis switch position

Each compute node features a single four-port 10Gb CN4054R adapter. Each CN4054R has two ASIC chips and each supports two of the four 10 Gb ports, as shown in Figure 4-2. Each compute node maintains two 10 Gbps connections to each switch. Storage connections and network teams must be distributed across both ASICs to maintain fault tolerance.

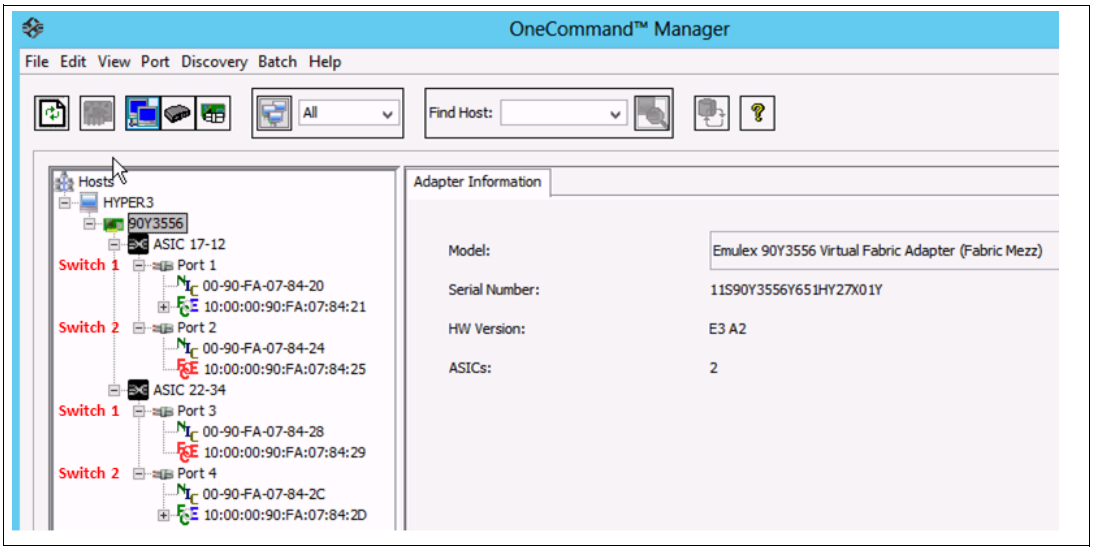

Figure 4-2 View of CN4054R network and storage connections

The connections between the CN4054R converged adapters and the CN4093 converged switches is shown in Figure 4-3.

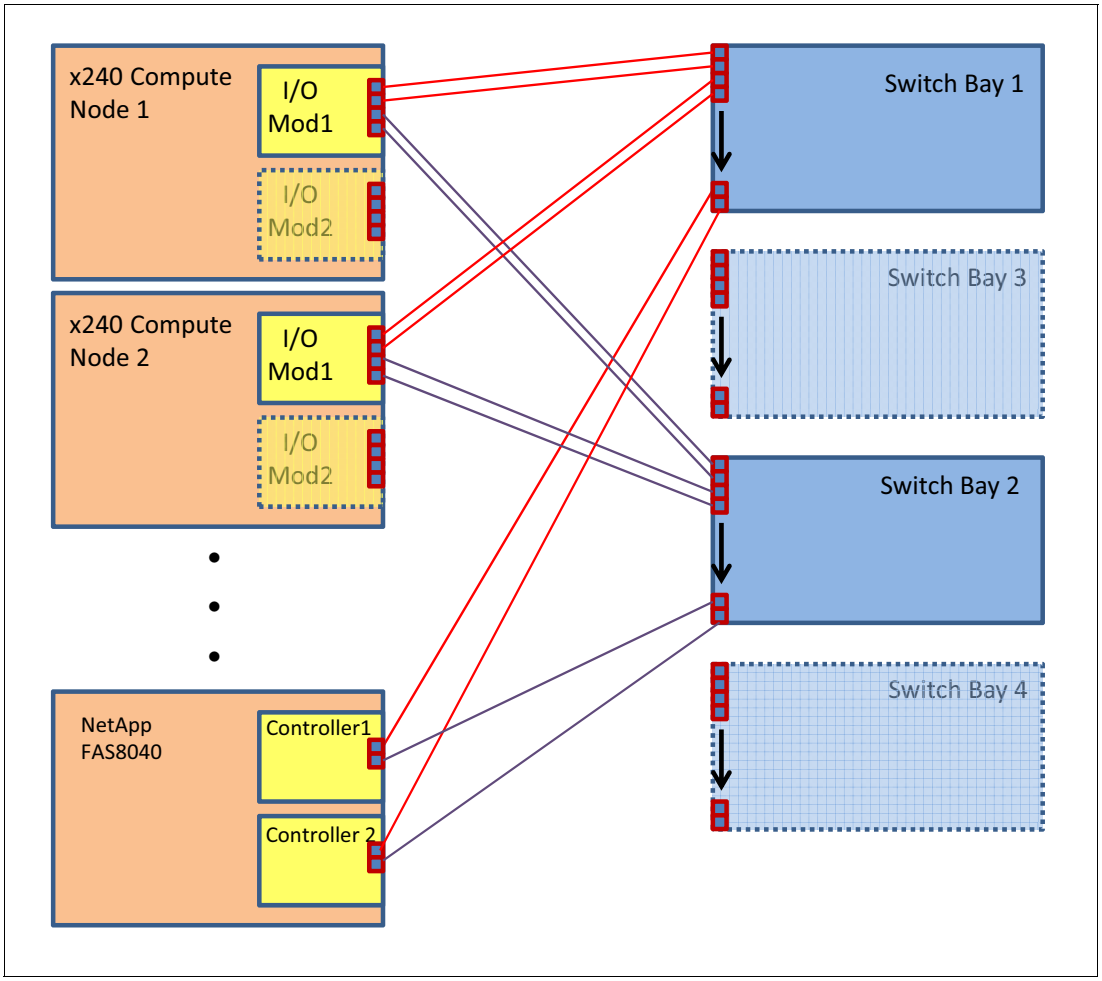

Figure 4-3 Converged storage and server connections

Combinations of physical and virtual isolated networks are configured at the host, switch, and storage layers to satisfy isolation requirements. At the physical host layer, there is a four-port 10GbE Virtual Fabric Adapter for each Hyper-V server (one Flex System CN4054R four-port 10 Gb VFA module). At the physical switch layer, there are two redundant Flex System CN4093 modules with up to 42 internal 10 GbE ports, 2 external 10 GbE SFP+ ports, 12 external SFP+ Omni ports, and 2 external 40 GbE QSFP+ ports (which can also be converted to 8 10 GbE ports) for storage and host connectivity.

To support all four 10 GbE connections from each server, the CN4093 switches require the Upgrade 1 Feature on Demand (FoD) option. The servers and storage maintain connectivity through two Fibre Channel over Ethernet (FCoE) connections by using Multi-Path I/O (MPIO). The two 10 GbE ports for FCoE also are shared with host management and cluster private and public data networks because these networks often are not bandwidth intensive.

Priority-based flow control at the switch level is set to prioritize up to 50% of the bandwidth to FCoE traffic and Windows quality of service (QoS) settings are applied on the data networks to shape and limit network bandwidth.

On the data network side, Windows Server 2012 R2 NIC teaming is used to provide fault tolerance and load balancing to the following communication networks:

- Host management
- Cluster private
- Cluster public
- Live migration
- Virtual machine (VM)

This setup enables the most efficient use of network resources with a highly optimized configuration for network and storage connectivity.

At the physical switch layer, VLANs are used to provide logical isolation between the various storage and data traffic. A key element is to properly configure the switches to maximize the available bandwidth and reduce congestion; however, based on individual environment preferences, there is flexibility regarding how many VLANs are created and what type of role-based traffic they handle. After a final selection is made, ensure that the switch configurations are saved and backed up.

Except for the NetApp FAS8040 ports, all switch ports must be configured as tagged and the VLAN definitions that are specified on each port as needed. Non-FCoE networks must have VLAN assignments that are made in Windows Server or Hyper-V.

Link Aggregation Control Protocol (LACP) is used to combine two 10 GbE switch ports into a single entity, which is then connected to a similar number of ports on the second switch. LACP teams provide higher bandwidth connections and error correction between LACP team members. LACP teams must be used for the stacking links that are created between the two CN4093 switches. An LACP team can also be used to support the uplink connections to a corporate network. An example of stacked switch configuration that uses LACP configuration is shown in Figure 4-4.

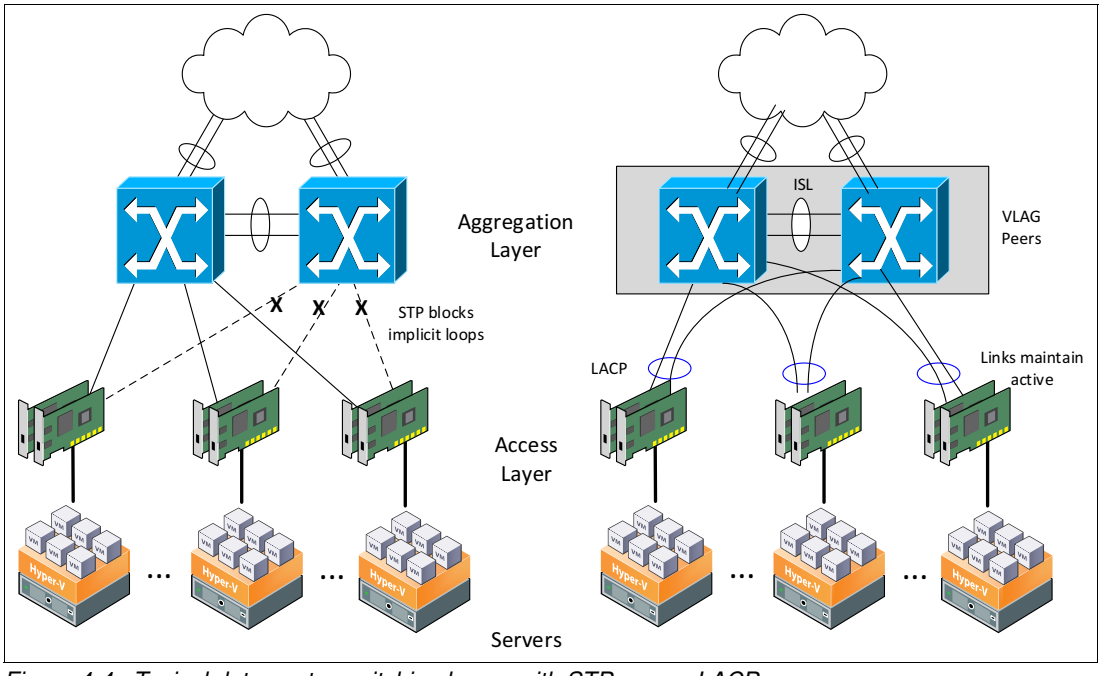

Figure 4-4 Typical data center switching layers with STP versus LACP

A high-level network VLAN overview is shown in Figure 4-5.

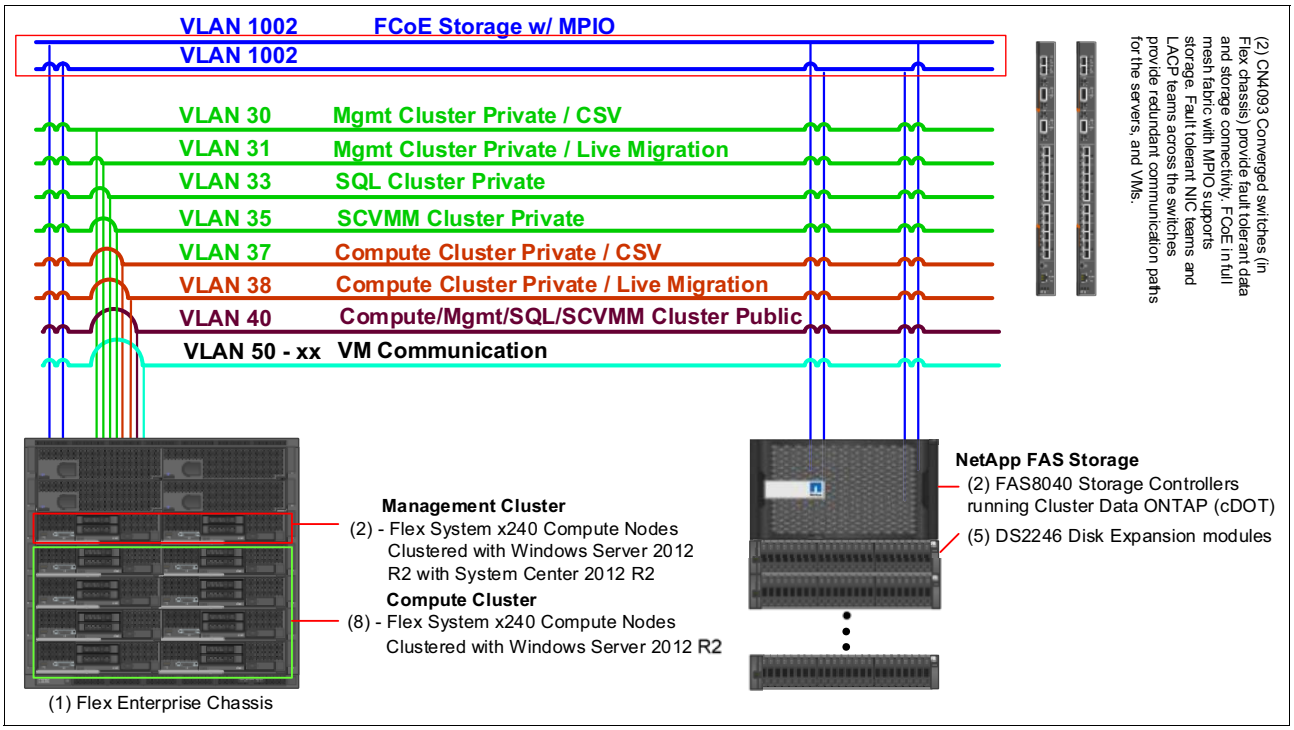

Figure 4-5 Flex System Solution for Microsoft Hyper-V architecture

4.2 VLANs

The validated configuration uses the VLANs that are described in Table 4-1.

Table 4-1 VLAN definitions

Network	Name	Description
VLAN 30	Management Cluster Private Network	Private cluster communication for 2-node management cluster
VLAN 31	Management Cluster Live Migration Network	VM Live Migration traffic for 2-node management cluster
VLAN 33	SQL Cluster Private	Private Cluster communication for SQL Cluster
VLAN 35	SCVVM Cluster Private	Private Cluster communication for SCVMM Cluster
VLAN 37	Production Cluster Private Network	Private cluster communication for 8-node production cluster
VLAN 38	Production Cluster Live Migration Network	VM Live Migration traffic for 8-node production cluster
VLAN 40	VM Communication Network	VM communication
VLAN 50	Cluster Public Network	Used for host management and cluster public network

VLAN 70	Out of band management network	Used for out of band connections to Chassis Management Module (CMM) and IMM devices
VLAN 1002	FCoE Storage Traffic	FCoE Storage Traffic
VLAN 4094	Inter Switch Link (ISL)	Inter Switch Link

4.2.1 Management cluster private and CSV networks (VLAN 30)

This configuration requires that the appropriate switch ports (see Table 4-2 on page 27) for each management x240 M5 Compute Node to set as “tagged”, and the VLAN definitions must include these ports for each switch. The network interfaces that are used to communicate on these VLANs must specify VLAN 30 in VM settings in Hyper-V Manager. There cannot be default gateway settings for cluster private network.

4.2.2 Management cluster live migration network (VLAN 31)

This configuration requires that the appropriate switch ports (see Table 4-2 on page 27) for each management x240 M5 Compute Node to set as “tagged”, and the VLAN definitions must include these ports for each switch. The network interfaces that are used to communicate on these VLANs must specify VLAN 31 in VM settings in Hyper-V Manager. There cannot be any default gateway settings for Live Migration network.

4.2.3 SQL guest cluster private network (VLAN 33)

This configuration requires that the appropriate switch ports (see Table 4-2 on page 27) for each management x240 M5 Compute Node to set as “tagged”, and the VLAN definitions must include these ports for each switch. The network interfaces that are used to communicate on these VLANs must specify VLAN 33 in VM settings in Hyper-V Manager. There cannot be any default gateway settings for SQL cluster private network.

4.2.4 SCVMM guest cluster private network (VLAN 35)

This configuration requires that the appropriate switch ports (see Table 4-2 on page 27) for each management x240 M5 Compute Node to set as “tagged”, and the VLAN definitions must include these ports for each switch. The network interfaces that are used to communicate on these VLANs must specify VLAN 35 in VM settings in Hyper-V Manager. There cannot be any default gateway settings for SCVMM cluster private network.

4.2.5 Production cluster private and CSV networks (VLANs 37)

This configuration requires that the appropriate switch ports (see Table 4-2 on page 27) for each management x240 M5 Compute Node to set as “tagged”, and the VLAN definitions must include these ports for each switch. The network interfaces that are used to communicate on these VLANs must specify VLAN 37 in Windows 2012 R2 Hyper-V vSwitch settings in Hyper-V Manager. There cannot be any default gateway settings for cluster private network.

4.2.6 Production cluster live migration network (VLAN 38)

This configuration requires that the appropriate switch ports (see Table 4-2 on page 27) for each management x240 M5 Compute Node to set as “tagged”, and the VLAN definitions must include these ports for each switch. The network interfaces used to communicate on these VLANs must specify VLAN 38 in Windows 2012 R2 Hyper-V vSwitch settings in Hyper-V Manager. There cannot be any default gateway settings for Live Migration network.

4.2.7 Cluster public network (VLAN 40)

This network supports communication for the host management servers for the management cluster, System Center components, and compute cluster. One team over two 10 GbE ports that is created by using the Windows Server 2012 R2 in-box NIC teaming feature is used to provide fault tolerance, load balancing for host management (cluster public), and cluster private networks.

This NIC team shares the bandwidth with the two FCoE 10 Gb ports. QoS is applied from Windows Server to limit its bandwidth usage. The management cluster also supports VLAN 40 on the VM communications network so that the System Center components can manage the host servers. VLAN identification must be set in Windows Server or Hyper-V.

4.2.8 Production VM communication network (VLAN 50)

This network supports communication for the VMs. One LACP team over two 10 GbE ports that is created by using the Windows Server 2012 in-box NIC teaming feature is used to provide fault tolerance and load balancing for communication for live migration and VM communication. This configuration requires that the appropriate switch ports (see Table 4-2 on page 27) for each production x240 M5 Compute Node be set as ‘tagged’ and the VLAN definitions must include these ports for each switch. Network settings for proper VLAN identification must be performed in each VM’s network interface.

If more segregation between VM networks is required, the VM Team network switch ports can have more VLAN IDs assigned as needed. Each VM can then set the necessary VLAN ID as part of its network settings in Hyper-V.

4.2.9 Out -of-band management network (VLAN 70)

This network supports communication for the out-of-band management network. As shown in Figure 3-1 on page 11, the Chassis Management Module (CMM) provides the communication entry point for the Flex System x240 M5 Integrated Management Module (IMM) and I/O modules. The management ports on the NetApp FAS8040 storage are connected to this VLAN.

Note: It is assumed that an isolated management network is available and connections can be established for each managed device to this network. Routing must be enabled between the management network and the VLAN 40 network for System Center to communicate to these devices.

4.2.10 FCoE storage network (VLAN 1002)

FCoE storage traffic on the physical switches between the NetApp FAS8040 and Flex System ITEs must be isolated on VLAN 1002.

4.2.11 Inter switch link network (VLAN 4094)

A dedicated VLAN to support the stacking between the two switches must be implemented.

4.3 x240 M5 Compute Node network ports

Each host server has one CN4054R 4-port 10 Gb device that is used for network connections and FCoE storage connectivity, public and private cluster communication, and VM communication. The FCoE connections to storage use Multipath I/O drives to ensure fault tolerance and load balancing.

Windows Server 2012 R2 NIC teaming is used for all but the FCoE networks to provide fault tolerance and spread the workload across the network communication interfaces. The NIC teams follow best practices by ensuring the team members are from each of the two ASICs on the CN4054R CNA adapter, so no single ASIC failure can take down the team. For more information, see Figure 4-3 on page 19.

4.3.1 Physical host data access

Consider the following points regarding physical host data access:

- Each compute node uses four connections to the Ethernet network (or networks). Figure 4-3 on page 19 shows four active data connections (two connections from each ASIC).
- By default, the CN4093 switches are set as untagged ports. This designation must be changed to “tagged” and VLAN IDs assigned according to Table 4-3 on page 28. The default VLAN ID remains with a PVID equal to “1”.
- Windows Server 2012 R2 NIC teaming is used to form high-bandwidth, fault-tolerant teams.

4.3.2 Physical host FCoE storage access

Consider the following points regarding physical host FCoE storage access:

- Each compute node uses two connections to the FCoE network. Figure 4-3 on page 19 shows two active FCoE connections (one connection from each ASIC).
- Because the host servers use switch ports for data and storage traffic, each CN4093 switch port is changed from the default “untagged” mode to “tagged”. The default VLAN ID remains with a PVID equal to “1”. Correct VLAN IDs for storage and data must be specified for each switch port, as shown in Table 4-3 on page 28.

4.3.3 Storage controller FCoE access

At the physical storage layer, the NetApp FAS8040 uses 10 Gb Converged Network Adapters (CNAs) for the FCoE storage connectivity. These adapters are called Universal Target Adapters (UTA2). Each controller uses two of its four UTA2 converged ports for FCoE traffic.

The use of the NetApp MPIO Device Specific Module (DSM) manages the multiple I/O paths between the host servers and storage system and optimizes the storage paths for maximum performance. VLANs are used to isolate storage traffic from other data traffic that occurs on the switches. FCoE traffic is prioritized by the CN4093 switches to maximize storage traffic throughput.

Consider the following points:

- Each of the two FAS8040 storage controllers maintains two FCoE connections to help balance storage workloads and redundancy.
- One connection is provided to each controller to each switch (see Figure 4-2 on page 18).
- The CN4093 switches are set as untagged ports by default. The ports for the storage controller must be assigned to the FCoE VLAN ID 1002 for both switches.
- FCoE on the NetApp storage requires FC protocol license.

4.4 Flex System CN4093 converged Ethernet configuration

The Lenovo Hyper-V Virtualization Reference Architecture uses two Flex System CN4093 switches containing up to 64 10 GbE ports each. The CN4093 provides primary storage access and data communication services. Redundancy across the switches is achieved by creating a stacking link between switches 1 and 2. The stacking link is created by using two external 10 GbE links from each switch to form into a LACP team. Corporate uplink connections also can use LACP (as shown in Figure 4-6), depending on the customer configuration. Each of the CN4093 switches requires Upgrade 1 to activate the extra ports that are needed to fully support all the CN4054 ports on each x240 M5 Compute Node.

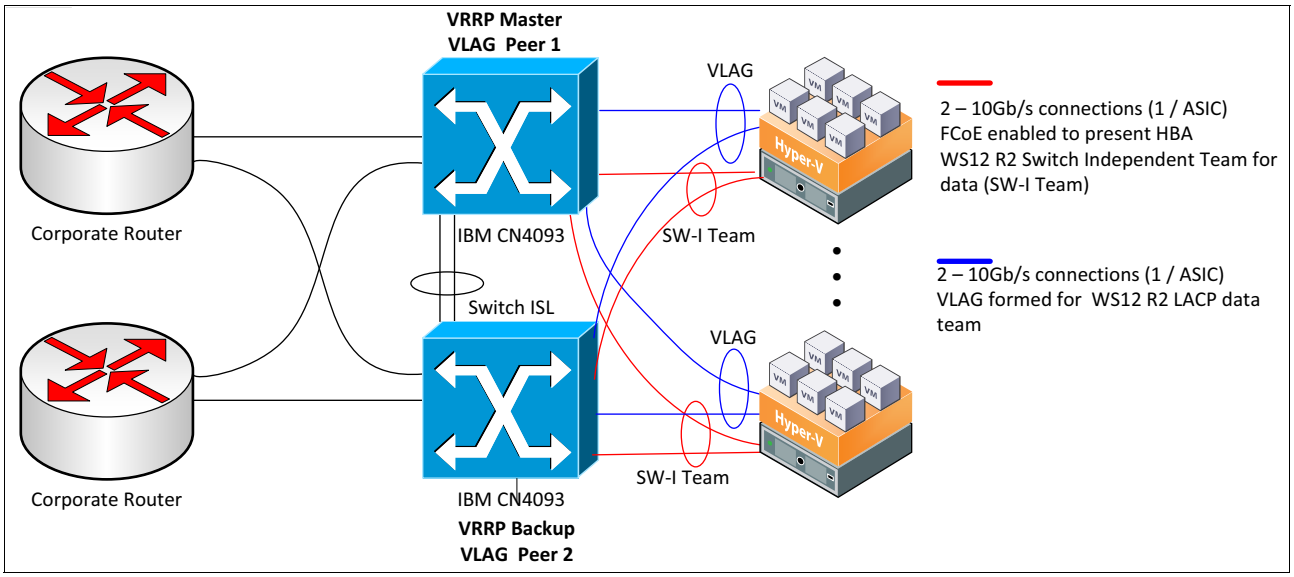

Figure 4-6 Active-Standby uplink configuration that uses LACP

Note: Switch ports that are used for FCoE do not support LACP. Windows 2012 R2 Switch Independent NIC Teaming is used for those ports to support TCP/IP data traffic.

Note: All routing is performed in upstream switches.

Management of the CN4093 switches can be performed through a command-line interface (CLI) or a web-based user interface, as shown in Figure 4-7. The default user name for the CN4093 switches is `admin` and the password for the switches is `admin`. These default values must be changed to a unique password to meet your security requirements.

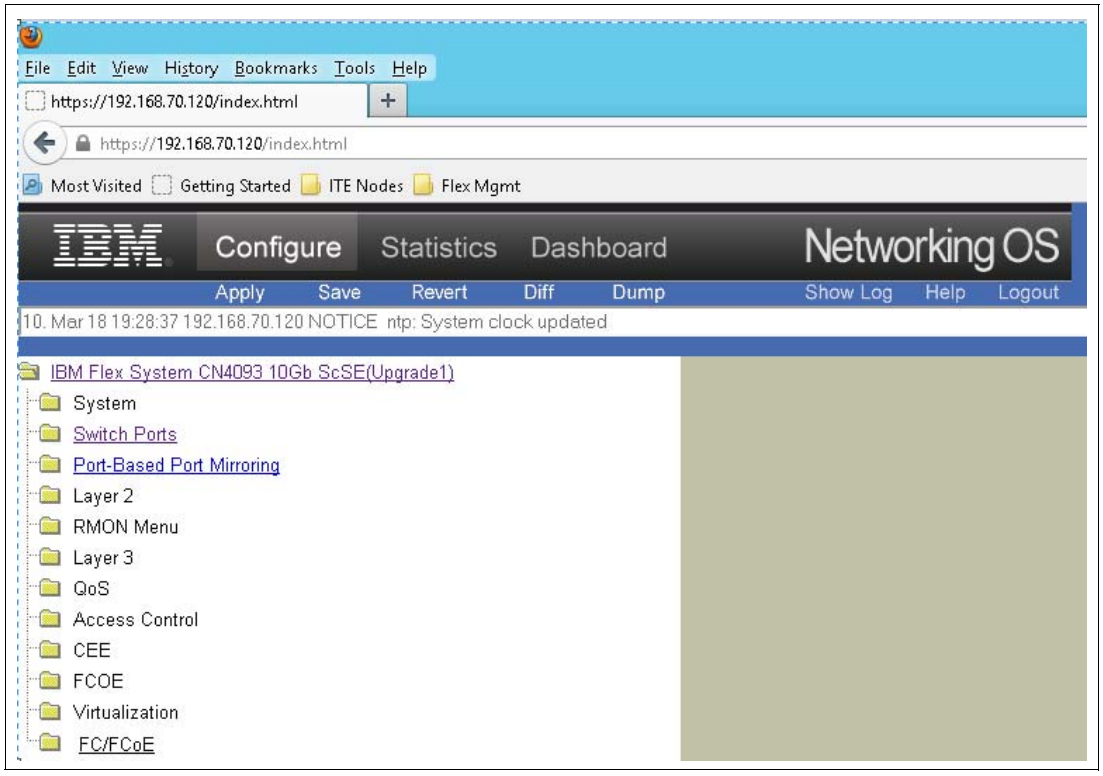

Figure 4-7 CN4093 Administration interface

Note: Spanning Tree must be enabled on all switches according to your organizational requirements.

With the switches stacked, all management is performed from Switch 1 (Master Switch).

By default, the switches are assigned the following management IP address:

- Switch 1: 192.168.70.120
- Switch 2: 192.168.70.121

Table 4-2 lists the roles of each switch port for the two CN4093 switches in the configuration.

Table 4-2 CN4093 Switch port roles

Port	Switch1	Switch2
Int Port A1	FCoE and Cluster Public Team	LACP (Live Migration and VM Communication)
Int Port B1	LACP (Live Migration and VM Communication)	FCoE and Cluster Public Team
Int Port A2	FCoE and Cluster Public Team	LACP (Live Migration and VM Communication)
Int Port A2	LACP (Live Migration and VM Communication)	FCoE and Cluster Public Team
Int Port A3	FCoE and Cluster Public Team	LACP (Live Migration and VM Communication)
Int Port B3	LACP (Live Migration and VM Communication)	FCoE and Cluster Public Team
Int Port A4	FCoE and Cluster Public Team	LACP (Live Migration and VM Communication)
Int Port B4	LACP (Live Migration and VM Communication)	FCoE and Cluster Public Team
Int Port A5	FCoE and Cluster Public Team	LACP (Live Migration and VM Communication)
Int Port B5	LACP (Live Migration and VM Communication)	FCoE and Cluster Public Team
Int Port A6	FCoE and Cluster Public Team	LACP (Live Migration and VM Communication)
Int Port B6	LACP (Live Migration and VM Communication)	FCoE and Cluster Public Team
Int Port A7	FCoE and Cluster Public Team	LACP (Live Migration and VM Communication)
Int Port B7	LACP (Live Migration and VM Communication)	FCoE and Cluster Public Team
Int Port A8	FCoE and Cluster Public Team	LACP (Live Migration and VM Communication)
Int Port B8	LACP (Live Migration and VM Communication)	FCoE and Cluster Public Team
Int Port A9	FCoE and Cluster Public Team	LACP (Live Migration and VM Communication)
Int Port B9	LACP (Live Migration and VM Communication)	FCoE and Cluster Public Team
Int Port A10	FCoE and Cluster Public Team	LACP (Live Migration and VM Communication)
Int Port B10	LACP (Live Migration and VM Communication)	FCoE and Cluster Public Team
Ext Port EXT1	Interswitch Link (LACP)	Interswitch Link (LACP)
Ext Port EXT2	Interswitch Link (LACP)	Interswitch Link (LACP)
Ext Port EXT11	Type FC (used for FCoE configuration)	Type FC (used for FCoE configuration)
Ext Port EXT12	Type FC (used for FCoE configuration)	Type FC (used for FCoE configuration)
Ext Port EXT13	NetApp FAS8040 (FCoE) (Ctrl-1)	NetApp FAS8040 (FCoE) (Ctrl-1)
Ext Port EXT14	NetApp FAS8040 (FCoE) (Ctrl-2)	NetApp FAS8040 (FCoE) (Ctrl-2)
Ext Port EXT15	NetApp FAS8040 (SMB)	NetApp FAS8040 (SMB)
Ext Port EXT16	Unused	Unused
Ext Port EXT17	Uplink (LACP)	Uplink (LACP)
Ext Port EXT18	Uplink (LACP)	Unused (LACP)

Table 4-3 lists the VLAN configuration of the ports for each of the two CN4093 switches in the configuration.

Table 4-3 CN4093 switch port VLAN roles

Port	Device	Tagging	PVID	Switch1 VLANs	Switch2 VLANs
Int Port A1	Compute Node1	Yes	1	37, 40, 1002	38, 50
Int Port B1		Yes	1	38, 50	37, 40, 1002
Int Port A2	Compute Node2	Yes	1	37, 40, 1002	38, 50
Int Port A2		Yes	1	38, 50	37, 40, 1002
Int Port A3	Compute Node3	Yes	1	37, 40, 1002	38, 50
Int Port B3		Yes	1	38, 50	37, 40, 1002
Int Port A4	Compute Node4	Yes	1	37, 40, 1002	38, 50
Int Port B4		Yes	1	38, 50	37, 40, 1002
Int Port A5	Compute Node5	Yes	1	37, 40, 1002	38, 50
Int Port B5		Yes	1	38, 50	37, 40, 1002
Int Port A6	Compute Node6	Yes	1	37, 40, 1002	38, 50
Int Port B6		Yes	1	38, 50	37, 40, 1002
Int Port A7	Compute Node7	Yes	1	37, 40, 1002	38, 50
Int Port B7		Yes	1	38, 50	37, 40, 1002
Int Port A8	Compute Node8	Yes	1	37, 40, 1002	38, 50
Int Port B8		Yes	1	38, 50	37, 40, 1002
Int Port A9	Management Node1	Yes	1	30, 40, 1002	31, 33, 35, 40
Int Port B9		Yes	1	31, 33, 35, 40	30, 40, 1002
Int Port A10	Management Node2	Yes	1	30, 40, 1002	30, 33, 35, 40
Int Port B10		Yes	1	31, 33, 35, 40	30, 40, 1002
Ext Port EXT1	Inter switch Link	Yes	4094	30, 31, 33, 35, 37, 38, 40, 50, 4094	30, 31, 33, 35, 37, 38, 40, 50, 4094
Ext Port EXT2		Yes	4094	30, 31, 33, 35, 37, 38, 40, 50, 4094	30, 31, 33, 35, 37, 38, 40, 50, 4094
Ext Port EXT11	N/A	N/A	N/A	Used for FCoE Configuration	Used for FCoE Configuration
Ext Port EXT12		N/A	N/A	(set to type fc)	(set to type fc)
Ext Port EXT13	NetApp FAS8040 (Ctrl-1)	Yes	1	1002	1002
Ext Port EXT14	NetApp FAS8040 (Ctrl-2)	Yes	1	1002	1002
Ext Port EXT15	NetApp FAS8040 (SMB)	Yes	1	40	40
Ext Port EXT16		Yes	1	Unused	Unused

Port	Device	Tagging	PVID	Switch1 VLANs	Switch2 VLANs
Ext Port EXT17	CorpNet Uplink	Yes	1	Uplink (LACP) VLAN as needed	Uplink (LACP) VLAN as needed
Ext Port EXT18	CorpNet Uplink	Yes	1	Uplink (LACP) VLAN as needed	Uplink (LACP) VLAN as needed

Ports are set as “untagged” by default. All Flex x240 M5 ports are set to “tagged” in the configuration. A preferred VLAN ID (PVID) must remain set to “1”. This setting can be made from the switch GUI in the Configuration window (as shown Figure 4-8) or ISCLI, as described in 4.4.1, “Use of ISCLI to configure CN4093 switches” on page 31.

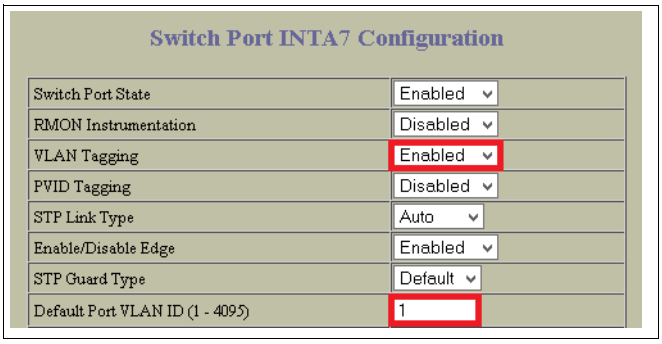

Figure 4-8 Setting VLAN tagging and Default VLAN ID

VLAN assignments for the CN4093 switch ports can be made in the GUI (as shown in Figure 4-9) or by using the ISCLI, as described in 4.4.1, “Use of ISCLI to configure CN4093 switches” on page 31.

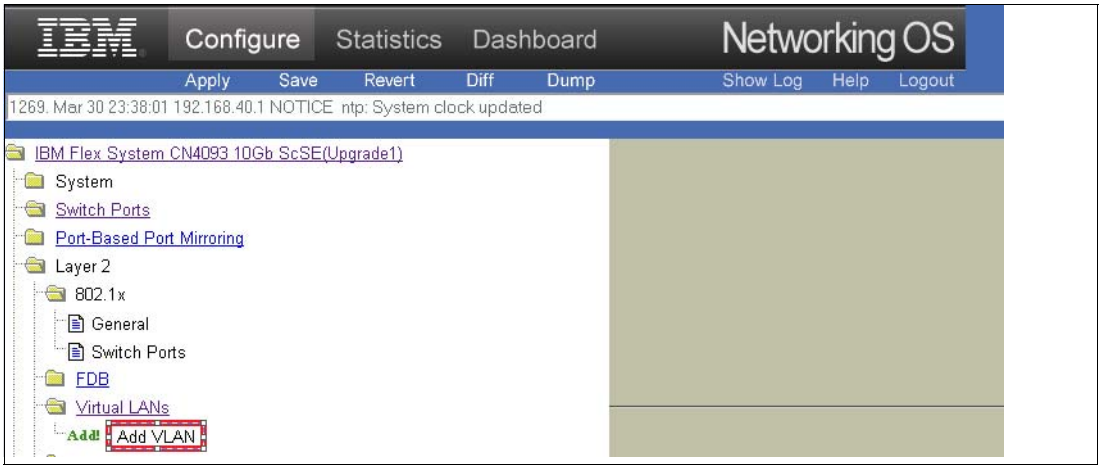

Figure 4-9 Adding ports to VLAN interface

Note: Each LACP team has its own unique LACP ID with each port that is a member of that team. For more information about Switch Stacking, VLANs, and LACP configuration steps, see 4.4.1, “Use of ISCLI to configure CN4093 switches” on page 31

Figure 4-10 shows the concept of the use of LACP teams from the NIC interfaces.

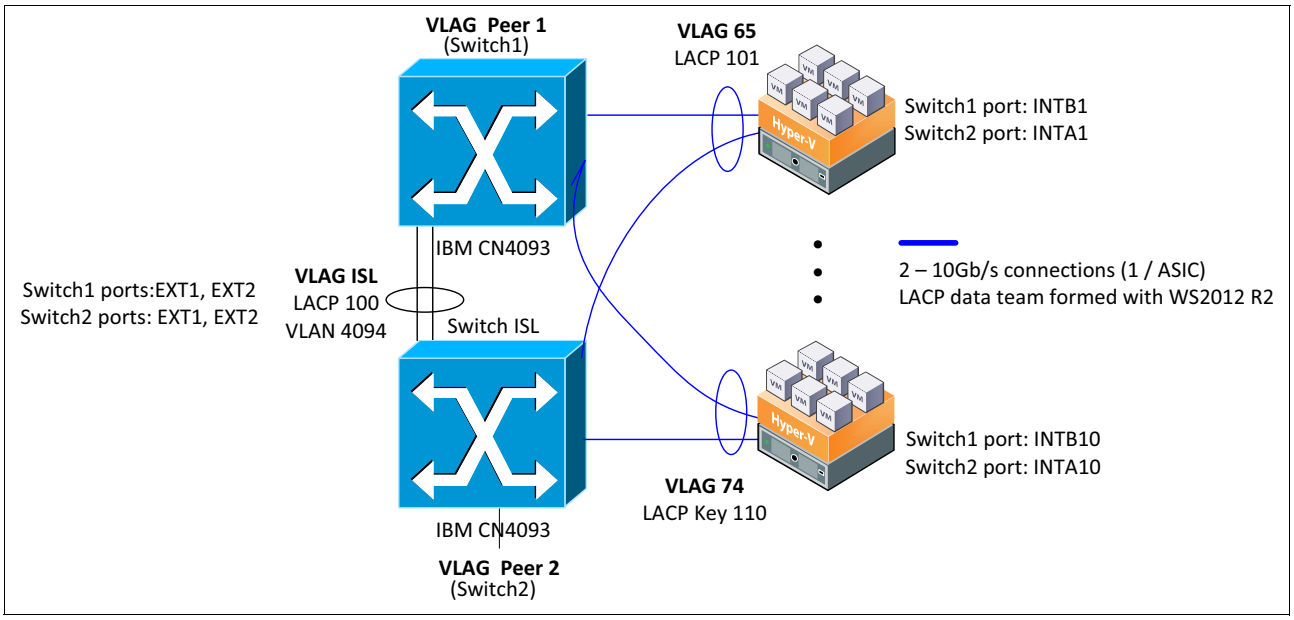

Figure 4-10 LACP configuration

An LACP team is created only between the two ports that are not used for FCoE. Each server's LACP team must consist of one port from each CNA ASIC and span the two CN4093 switches.

Table 4-4 lists the LACP configurations for the switch stack.

Table 4-4 Summary of LACP teams for Switch 1 and 2

VLAG ID (Server)	Switch1 Port	Switch2 Port	LACP Key
65 (Server 1)	IntB1	IntA1	101
66 (Server 2)	IntB2	IntA2	102
67 (Server 3)	IntB3	IntA3	103
68 (Server 4)	IntB4	IntA4	104
69 (Server 5)	IntB5	IntA5	105
70 (Server 6)	IntB6	IntA6	106
71 (Server 7)	IntB7	IntA7	107
72 (Server 8)	IntB8	IntA8	108
73 (Server 9)	IntB9	IntA9	109
74 (Server 10)	IntB10	IntA10	110
ISL	EXT1	EXT1	100
ISL	EXT2	EXT2	100

4.4.1 Use of ISCLI to configure CN4093 switches

This section provides guidance about switch configuration by using the ISCLI command-line environment. Although this information is not an exhaustive step-by-step set of instructions, it does provide information about each major component of the switch configuration, such as Stacking, VLAN, Port Configuration, and FCoE configuration. To access the ISCLI, see the Flex System Fabric CN4093 Industry Standard CLI Command Reference.

Granting privilege and enter configuration mode

Complete the following steps to grant privilege and enter configuration mode:

1. Grant Privilege Mode by using the following command:

```
CN 4093# enable
```

2. Enter Configuration Mode by using the following command:

```
CN 4093# configure terminal
```

Activating the extra 10 Gb external ports

Complete the following steps to activate the extra 10 Gb external ports:

1. Disable one 40 Gb port so port license can be reapplied by using the following command:

```
CN 4093# no boot port-map EXT3
```

2. Enable more 10 Gb ports by using the following commands:

```
CN 4093# boot port-map EXT17
CN 4093# boot port-map EXT18
CN 4093# boot port-map EXT19
CN 4093# boot port-map EXT20
```

Configuring the ISL and VLAG Peer Relationship

Complete the following steps to configure the ISL and VLAG Peer Relationship:

1. Enable VLAG Globally by using the following command:

```
CN 4093(config)# vlag enable
```

2. Configure the ISL ports of each switch and place them into a port trunk group by using the following commands:

```
CN 4093(config)# interface port ext1-ext2
CN 4093(config-if)# tagging
CN 4093(config-if)# pvid 4094
CN 4093(config-if)# lacp mode active
CN 4093(config-if)# lacp key 100
CN 4093(config-if)# exit
```

3. Place the ISL into a dedicated VLAN (VLAN 4094 is recommended) by using the following commands:

```
CN 4093(config)# vlan 4094
CN 4093(config-vlan)# enable
CN 4093(config-vlan)# member ext1-ext2
CN 4093(config-vlan)# exit
```

4. If STP is used on the switch, turn off STP for the ISL by using the following commands:

```
CN 4093(config)# spanning-tree stp 20 vlan 4094
CN 4093(config)# no spanning-tree stp 20 enable
```

5. Configure VLAG Tier ID by using the following command. This ID is used to identify the VLAG switch in a multitier environment:

```
CN 4093(config)# vlag tier-id 10
```

6. Define VLAG peer relationship by using the following commands:

```
CN 4093(config)# vlag isl vlan 4094
CN 4093(config)# vlag isl adminkey 100
CN 4093(config)# exit
```

7. Save the configuration changes by using the following command:

```
CN 4093# write
```

8. Configure the ISL and VLAG peer relationship for the second switch. Ensure the VLAG peer (VLAG Peer 2) is configured by using the same ISL trunk type (dynamic or static), VLAN, STP mode, and tier ID used on VLAG peer 1.

Configuring the Host VLAG

Complete the following steps to configure the Host VLAG:

1. Make each port from the ITEs an active LACP team member by using the following commands. This process must be done for each of the two ports per ITE (once on each switch). For more information, see Table 4-4 on page 30:

```
CN 4093(config)# interface port intb1 (Switch1)
CN 4093(config)# interface port inta1 (Switch2)
CN 4093(config-if)# lacp mode active
CN 4093(config-if)# lacp key 101 (For ITE1)
CN 4093(config-if)# exit
```

2. Enable the VLAG trunk on each switch by using the following command. This change enables LACP teams to be formed across the two CN4093 switches and must be done for each LACP key that is listed in Table 4-4 on page 30 on each switch:

```
CN 4093(config)# vlag adminkey 101 enable
```

3. Configure all required VLAGs on VLAG Peer 1 (Switch 1), and then repeat the configuration for VLAG Peer 2. For each corresponding VLAG on the peer, the port trunk type (dynamic or static), VLAN, STP mode, and ID must be the same as on VLAG Peer 1.
4. Verify the completed configuration by using the following command:

```
CN 4093(config)# show vlag
```

Note: The LACP teams do not show up as active on the switches until they also are formed in Windows Server 2012 with NIC teaming.

Configuring the VLANs

Complete the following steps to configure the VLANs:

1. From the ISCLI of Switch1, run the following commands:

```
CN 4093# enable
CN 4093# configure terminal
CN 4093(config)# interface port inta1-inta10,intb1-intb10
CN 4093(config-if)# tagging
CN 4093(config-if)# exit
CN 4093(config)# interface port ext1-ext2,ext15
CN 4093(config-if)# tagging
```

```
CN 4093(config-if)# exit
CN 4093(config)# vlan 30
CN 4093(config-vlan)# enable
CN 4093(config-vlan)# member inta9-inta10,ext1-ext2
CN 4093(config-vlan)# exit
CN 4093(config)# vlan 31
CN 4093(config-vlan)# enable
CN 4093(config-vlan)# member intb9-intb10,ext1-ext2
CN 4093(config-vlan)# exit
CN 4093(config)# vlan 33
CN 4093(config-vlan)# enable
CN 4093(config-vlan)# member intb9-intb10,ext1-ext2
CN 4093(config-vlan)# exit
CN 4093(config)# vlan 35
CN 4093(config-vlan)# enable
CN 4093(config-vlan)# member intb9-intb10,ext1-ext2
CN 4093(config-vlan)# exit
CN 4093(config)# vlan 37
CN 4093(config-vlan)# enable
CN 4093(config-vlan)# member inta1-inta8,ext1-ext2
CN 4093(config-vlan)# exit
CN 4093(config)# vlan 38
CN 4093(config-vlan)# enable
CN 4093(config-vlan)# member intb1-intb8,ext1-ext2
CN 4093(config-vlan)# exit
CN 4093(config)# vlan 40
CN 4093(config-vlan)# enable
CN 4093(config-vlan)# member inta1-inta10,intb9-intb10
CN 4093(config-vlan)# member ext1-ext2,ext15,ext17-ext18
CN 4093(config-vlan)# exit
CN 4093(config)# vlan 50
CN 4093(config-vlan)# enable
CN 4093(config-vlan)# member intb1-intb10,ext1-ext2,ext17-ext18
CN 4093(config-vlan)# exit
CN 4093(config-vlan)# show vlan
CN 4093(config-vlan)# show interface status
CN 4093(config-vlan)# write
```

2. From the ISCLI of Switch2: Set up the same configuration as Switch1 with the corresponding ports.
3. Back up the configuration to TFTP server (xx.xx.xx.yy is the IP address of the TFTP server) by using the following command:

```
CN 4093# copy running-config tftp filename "file.cfg" address xx.xx.xx.yy
mgt-port
```

Note: Switch ports EXT17 and EXT18 must be configured as a fault tolerant team and used as uplink connections for Routing and CorpNet access. Routing must be configured for VLANs 40, 50, and 70.

Configuring FCoE

Complete the following steps to configure each switch to support FCoE:

Note: It is easiest to perform the switch FCoE configuration after servers are enabled and configured for FCoE, the operating system is installed, and WWPN is recorded for each server.

1. Enable FCoE on each switch from the ISCLI by using the following commands:

```
CN 4093(config)# cee enable
CN 4093(config)# fcoe fips enable
```

2. Enable FCoE Fibre Channel forwarding on each Switch for VLAN 1002 from the ISCLI by using the following commands:

```
CN 4093# enable
CN 4093# configure terminal
CN 4093(config)# system port ext11-ext12 type fc
CN 4093(config)# vlan 1002
CN 4093(config)# enable
(Switch1)
CN 4093(config-vlan)# member inta1-inta10,ext13-ext14
(Switch2)
CN 4093(config-vlan)# member intb1-intb10,ext13-ext14
CN 4093(config-vlan)# fcf enable
```

3. An FC alias is assigned to each HBA WWPN that is used for easier name identification. An example of FC aliases assignments are listed in Table 4-5 and Table 4-6 on page 35. One port is activated at the switch from each of the ASICs on the CN4054 CNA adapter. WWPN can be viewed from the OneCommand Manager tool for each server, as shown in Figure 4-2 on page 18. NetApp System Manager can be used to view the WWPN for each Storage virtual machine (SVM). For more information, see Chapter 5, “NetApp Storage System” on page 37.

Table 4-5 Switch 1 Fibre Channel alias example

FC Alias	WWPN
ITE1_Port_IntA1	10:00:00:90:fa:07:84:21
ITE2_Port_IntA2	10:00:00:90:fa:0d:4d:93
ITE3_Port_IntA3	10:00:00:90:fa:0d:2a:27
ITE4_Port_IntA4	10:00:00:90:fa:0d:33:75
NTAP_SVM1_Port1	20:0b:00:a0:98:5e:5b:b4
NTAP_SVM1_Port2	20:0d:00:a0:98:5e:5b:b4
NTAP_SVM2_Port1	20:07:00:a0:98:5e:5b:b4
NTAP_SVM2_Port2	20:09:00:a0:98:5e:5b:b4

Table 4-6 Switch 2 Fibre Channel alias example

FC Alias	WWPN
ITE1_Port_IntB1	10:00:00:90:fa:07:84:2d
ITE1_Port_IntB2	10:00:00:90:fa:0d:4d:9f
ITE1_Port_IntB3	10:00:00:90:fa:0d:2a:33
ITE1_Port_IntB4	10:00:00:90:fa:0d:33:81
NTAP_SVM1_Port1	20:0b:00:a0:98:5e:5b:b4
NTAP_SVM1_Port2	20:0b:00:a0:98:5e:5b:b4
NTAP_SVM2_Port1	20:07:00:a0:98:5e:5b:b4
NTAP_SVM2_Port2	20:09:00:a0:98:5e:5b:b4

4. Create FCAliases for each worldwide port name (WWPN) on each switch by using the following commands. Each Flex System ITE presents one WWPN, and each storage virtual machine presents two WWPN per switch for storage connections. FCAliases for the storage VMs must to wait until the WWPNs are set up on the NetApp storage:

```
CN 4093(config)# show fcns database
CN 4093(config)#
CN 4093(config)#fcalias ITE9_Port_IntA1 wwn  10:00:00:90:fa:0d:2a:27
CN 4093(config)#fcalias ITE10_Port_IntA2 wwn  10:00:00:90:fa:0d:33:75
CN 4093(config)#fcalias NTAP_SVM1_Port1 wwn  20:07:00:a0:98:5e:5b:b4
CN 4093(config)#fcalias NTAP_SVM1_Port2 wwn  20:09:00:a0:98:5e:5b:b4
CN 4093(config)#fcalias NTAP_SVM2_Port1 wwn  20:0b:00:a0:98:5e:5b:b4
CN 4093(config)#fcalias NTAP_SVM2_Port2 wwn  20:0b:00:a0:98:5e:5b:b4
CN 4093(config-zone)# exit
CN 4093(config)# show fcalias
CN 4093(config)# write
```

5. Create and populate the following FC zones on each switch. These zones must contain all of the FC aliases that were created:

   - Zone1 must include the SVM1 storage and management servers for switch 1
   - Zone2 must include the SVM2 storage and management servers for switch 1
   - Zone3 must include the SVM1 storage and compute servers for switch 1
   - Zone4 must include the SVM2 storage and compute servers for switch 1
   - Zone5 must include the SVM1 storage and management servers for switch 2
   - Zone6 must include the SVM2 storage and management servers for switch 2
   - Zone7 must include the SVM1 storage and compute servers for switch 2
   - Zone8 must include the SVM2 storage and compute servers for switch 2

   Windows 2012 R2 show supports shared VHDX disks. Shared VHDX disk is used for guest clustering in this configuration. NPIV or iSCSI are alternative methods available if needed:

```
(SVM1 on Switch1)
CN 4093(config-zone)# zone name SW1Zone_SVM1_MgmtSvrs
CN 4093(config-zone)# member fcalias ITE9_PortA9
CN 4093(config-zone)# member fcalias ITE10_PortA10
CN 4093(config-zone)# member fcalias NTAP_SVM1_Port1
CN 4093(config-zone)# member fcalias NTAP_SVM1_Port2
CN 4093(config-zone)# exit
(SVM2 on Switch1)
```

```
CN 4093(config-zone)# zone name SW2Zone_SVM2_MgmtSvrs
CN 4093(config-zone)# member fcalias ITE9_PortB9
CN 4093(config-zone)# member fcalias ITE10_PortB10
CN 4093(config-zone)# member fcalias NTAP_SVM2_Port1
CN 4093(config-zone)# member fcalias NTAP_SVM2_Port2
CN 4093(config-zone)# exit
CN 4093(config)#
CN 4093(config)# show zone
CN 4093(config)# write
```

6. An FC zone set can contain multiple FC zones. Create and activate an FC zone set for each switch that contains the FC zone (or zones) that were created by using the following commands:

```
(Switch1)
CN 4093(config)# zone set name SW1_ZoneSet
CN 4093(config-zoneset)# member SW1Zone_SVM1_MgmtSvrs
CN 4093(config-zoneset)# member SW1Zone_SVM2_MgmtSvrs
CN 4093(config-zoneset)# member SW1Zone_SVM1_ComputeSvrs
CN 4093(config-zoneset)# member SW1Zone_SVM2_ComputeSvrs
CN 4093(config)# exit
CN 4093(config)#
CN 4093(config)# show zone set
CN 4093(config)# zone set activate name SW1_ZoneSet
(Switch2)
CN 4093(config)# zone set name SW2_ZoneSet
CN 4093(config-zoneset)# member SW2Zone_SVM1_MgmtSvrs
CN 4093(config-zoneset)# member SW2Zone_SVM2_MgmtSvrs
CN 4093(config-zoneset)# member SW2Zone_SVM1_ComputeSvrs
CN 4093(config-zoneset)# member SW2Zone_SVM2_ComputeSvrs
CN 4093(config)# exit
CN 4093(config)#
CN 4093(config)# show zone set
CN 4093(config)# zone set activate name SW2_ZoneSet
CN 4093(config)# write
```

7. Back up the configuration to TFTP server (`xx.xx.xx.yy` is the IP address of the TFTP server) by using the following commands:

```
CN 4093# copy running-config tftp mgt-port
Enter the IP address of the TFTP Server: xx.xx.xx.yy
Enter the filename: SW1-Oct-24-2014.cfg
```

4.5 Active Director services

The Lenovo Hyper-V Fast Track reference configuration must be part of an Active Directory (AD) domain. This domain is required to form the Microsoft Windows Server 2012 R2 clusters. For this configuration an AD server must be created and is reachable to this configuration.

5

NetApp Storage System

This chapter describes the basic concepts and configuration steps for understanding and setting up the storage system and some of its supporting tools. Networking information, such as IP addresses for the different interfaces, must be available in advance. NetApp recommends the use of the configuration worksheets that are available in Appendix A, “Appendices” on page 113 as part of the preparation for setting up the storage system. Certain parameters that are used in our configuration (as built) are provided in this chapter for reference.

This chapter includes the following topics:

5.1 Concepts and features of NetApp Storage

This section provides a brief overview of the specific storage elements that are required for our implementation of this solution. It is outside the scope of this document to describe the details of all the storage elements. For more information about references for all storage components and elements, see Appendix A, “Appendices” on page 113.

5.1.1 Storage virtual machine

The secure logical storage partition through which data is accessed in clustered Data ONTAP is known as a storage virtual machine (SVM), which is formerly known as vServer. A storage cluster serves data through at least one and possibly multiple SVMs. An SVM is a logical abstraction that represents a set of physical resources of the storage cluster.

5.1.2 Aggregates

Aggregates provide storage to volumes. Aggregates can be comprised one or more RAID groups.

RAID-DP (known as dual parity RAID) is NetApp’s default RAID configuration and the recommended RAID option for the native FAS storage.

Flash Pool

A Flash Pool is an aggregate that combines solid-state drives (SSDs) and hard disk drives (HDDs) to provide a high performance aggregate more economically than an SSD-only aggregate. Clustered Data ONTAP automatically uses the available drive types in the most optimal method for the required workload. The SSDs provide a high-performance cache for the active data set of the data volumes that are provisioned on the Flash Pool aggregate. This configuration offloads random read operations and repetitive random write operations to improve response times and overall throughput for disk I/O-bound data access operations.

Flexible Volume

Flexible Volume (FlexVol) is the default type of volume that is used in this solution and recommended by NetApp for best practices. These virtual volumes can be thick or thin provisioned. Thin provisioned volumes is the best practice and are used in this solution.

Logical unit numbers

You can use a logical unit number (LUN) the same way you use local disks on the host. LUNs are created in and are associated with volumes. After you create the LUN, you must make it visible to the host. The LUN then appears on the Windows host as a disk.

Offloaded data transfer

Although the offload data transfer (ODX) copy offload feature is not required to run a Hyper-V workload over SMB 3.0, this feature can drastically improve VM deployment time for typical deployments in which the customer must provision multiple VMs. The main advantages of this feature are that it is not apparent to client machines and no data is sent over the network during file copy operations.

NetApp data ONTAP DSM for Windows MPIO

Microsoft MPIO is a protocol-independent feature that supports multiple data paths to a storage device with iSCSI, FC, or SAS. Providing multiple paths that can handle failover increases the availability from a host to the storage system. NetApp Data ONTAP DSMs for Windows MPIO help NetApp storage systems to integrate with Microsoft MPIO on Windows Server 2012 R2 servers and provide high availability (HA) to applications that use path-failover methods.

Data protection and storage efficiency

NetApp storage solutions provide redundancy and fault tolerance through clustered storage controllers and hot-swappable redundant components, such as cooling fans, power supplies, disk drives, and shelves. This highly available and flexible architecture enables customers to manage all data under one common infrastructure while meeting mission-critical uptime requirements.

In addition to the capabilities of the hardware, there are several standard and optional software features for improving data protection and storage utilization. These features include Snapshots, SnapVault, SnapMirror, SnapManager for Hyper-V, Thin Provisioning, FlexClone, Compression, and Deduplication.

OnCommand System Manager

NetApp OnCommand System Manager enables administrators to manage individual storage systems or clusters of NetApp storage systems through an easy-to-use, browser-based interface. System Manager includes wizards and workflows, which simplifies common storage tasks, such as creating volumes, LUNs, qtrees, shares, and exports, which saves time and prevents errors. System Manager works across all NetApp storage resources, including NetApp FAS8040.

5.2 Storage management and integration with Microsoft

One objective of the Microsoft Private Cloud solution is to enable a highly automated environment. Doing so on a large scale and expecting consistent results requires tight integration between the storage technologies and Microsoft System Center, and the Windows Servers. NetApp supports the Microsoft management tools and APIs (as shown in Figure 5-1), which is an overall integration approach that enables automation of rapidly provisioning and de-provisioning of VMs while maximizing storage efficiencies. It also significantly improves data protection of the VMs, enables various DR scenarios, and allows the complete end-to-end automation of a cloud-based solution, including self-service. This section provides an overview of these management and integration tools.

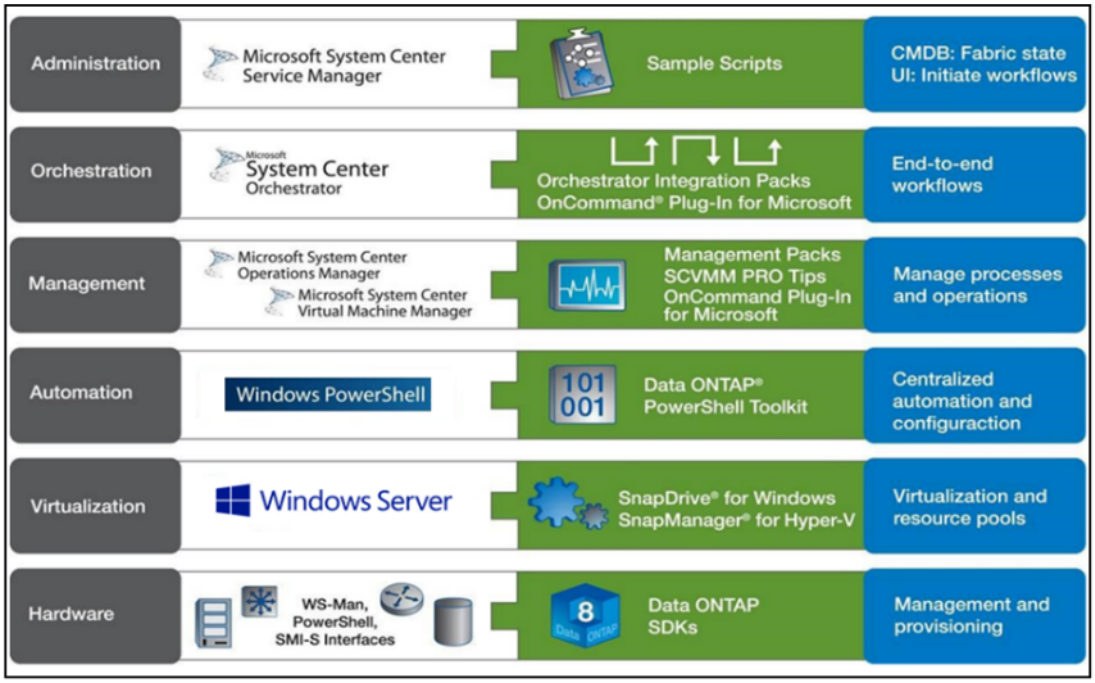

Figure 5-1 NetApp integration with Microsoft

5.2.1 OnCommand plug-in for Microsoft

NetApp OnCommand plug-in for Microsoft (OCPM) enables and simplifies the management of servers and storage systems in Microsoft System Center Operations Manager 2012 R2. OCPM offers native integration with System Center Operations Manager and SCVMM. These integrations provide the intelligence to make System Center fully storage aware, which simplifies the day-to-day administration of NetApp storage and amplifies the effectiveness of System Center monitoring and alerting.

5.2.2 NetApp SMI-S agent

The NetApp SMI-S agent provides a unified storage management interface that can be used to discover, monitor, and manage NetApp storage systems. It provides transparent integration of NetApp storage into Windows Server 2012 R2 and Microsoft SCVMM.

5.2.3 PowerShell toolkit

The standards-based management does not cover all possible contingencies because it is a subset of the capabilities of all vendors. Therefore, to facilitate more advanced deployments, Microsoft System Center Orchestrator, and the NetApp Data ONTAP PowerShell toolkit enable complete end-to-end orchestration and automation workflows. When these solutions are combined with System Center Service Manager, the workflows can be extended as integrated service offerings without the need for complex customer-built portals. They can be extended further through integration with Microsoft Service Management Automation and the Windows Azure Management Pack. Most of the basic functions of managing the storage can be achieved because of the SMI-S integration with System Center. The PowerShell toolkit offers more flexibility if special operation and customization are required.

5.2.4 SnapDrive for Windows

Several components are integrated into the SnapDrive for Windows software and are automatically installed. By using these components, you can manage LUNs, Windows volumes, or SMB shares. You can use these components together to enable SnapDrive for Windows workflows, including provisioning, Snapshot copy management, and backup, restore, and mounting operations.

SnapDrive for Windows “snap-in”. This software module integrates with Microsoft Management Console (MMC) to provide you a graphical interface for managing LUNs on the storage system. The module includes the following features:

- It is in the Windows Server computer management storage tree
- Provides a native MMC snap-in user interface for configuring and managing LUNs
- Supports remote administration so that you can manage SnapDrive on multiple hosts
- Provides SnapMirror integration
- Provides AutoSupport integration, including event notification

5.3 Setup and configuration of NetApp Storage

It is recommended to complete the Configuration Worksheet before working on the storage setup. The Configuration Worksheet helps you with all the preparation aspects and is available in Appendix A, “Appendices” on page 113.

5.3.1 Design considerations

It is important to design the storage layout in a way that effectively uses the various physical and logical elements of the storage system. Capacity and performance requirements are obvious aspects to consider when the storage is designed, but data protection and other operational aspects are also key and should be accounted for.

Proper design of the storage system maximizes the benefits of the use of the various storage features and provides certain flexibility to accommodate inevitable changes in business requirements. Although it is not in the scope of this document to describe storage design in detail, we indicate important aspects to consider and more content in Appendix A, “Appendices” on page 113.

5.3.2 Overview of FAS8040 I/O ports

The FAS8040 offers tremendous flexibility regarding connectivity because it includes multiple I/O ports that are standard and built-in. Also, the controllers include expansion slots for extra I/O cards for configurations that require larger number of connected ports. For our specific solution, all connections are based on the built-in ports of the controllers, as shown in Figure 5-2 and Figure 5-3 on page 43.

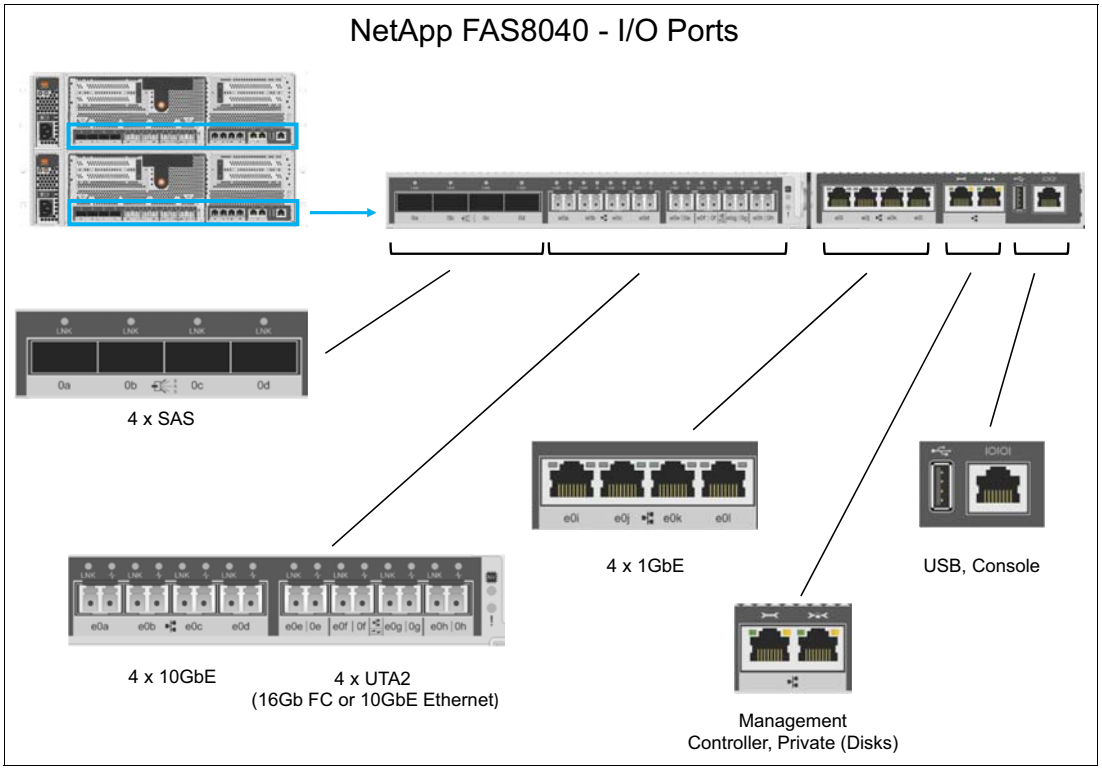

Figure 5-2 Built-in I/O Ports of the FAS8040

Figure 5-3 shows which ports are used for management and data traffic as built for our specific reference architecture.

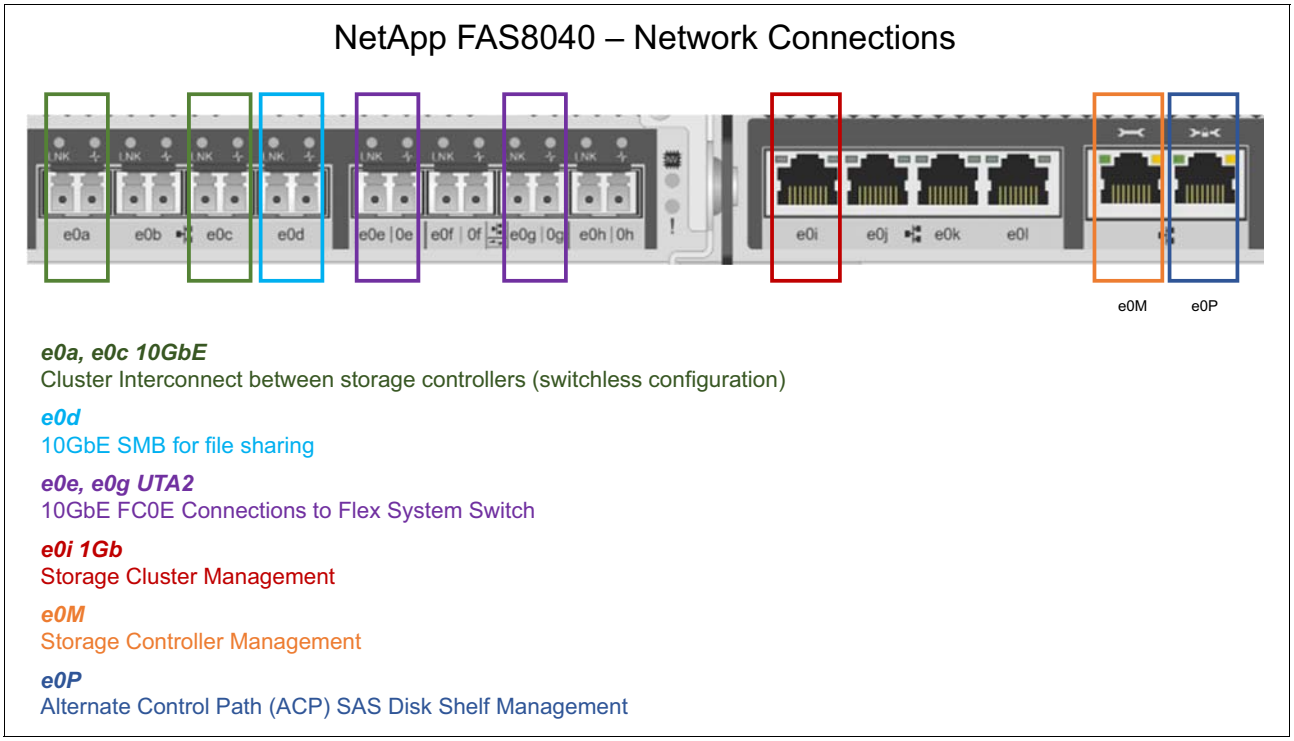

Figure 5-3 I/O ports

Figure 5-4 shows theFAS8040 controllers that are connected with redundant paths to the DS2246 disk shelves that provide the storage capacity. Each DS2246 disk shelf can be populated with up to 24 disks (SAS or SSD). In our specific configuration, four of the DS2246 shelves are fully populated with 900 GB 10 K SAS disks and one shelf is populated with 12 SSDs. Instructions about how to connect the disk shelves in a redundant way to the storage controllers are provided as part of the document that includes the storage system and can also be found at the NetApp support website (http://mysupport.netapp.com). For more information about for cabling the FAS8040 to the CN4093 switches, see Table 4-2 on page 27.

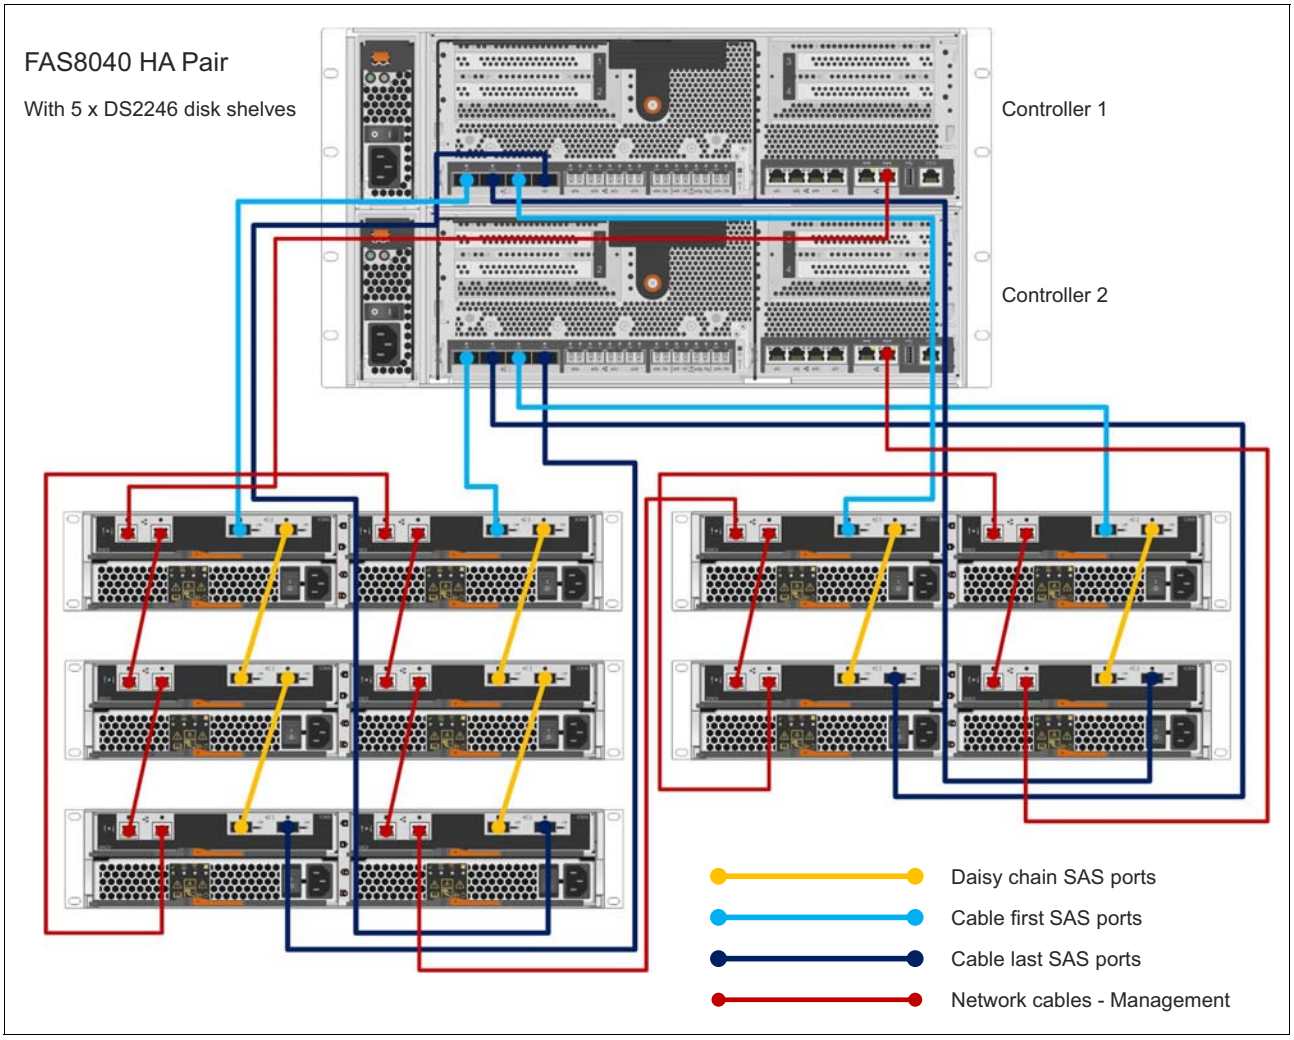

Figure 5-4 FAS8040 and disk shelves cabling

5.4 Flex chassis connections

The NetApp FAS8040 is connected to the CN4093 switches through the external ports on the CN4093 switches in the Flex Chassis. Each FAS8040 controller has one 10 Gb FCoE connection from the four available UTA2 ports to each switch to create a mesh topology for fault tolerance and performance. One other 10 Gb connection from each controller is made to each switch to be used for SMB traffic. Table 5-1 lists the specific ports on the FAS8040 and their mapping to the Flex System switches.

Table 5-1 FAS8040 mapping to Flex System

Storage controller	Local port	Interface	In use	Function/service	Remote device	Remote port
Controller 1						
NTAP_HVFT_Cluster-01	e0e	UTA2 - 10 GbE	Yes	FCoE to Host	Flex System Switch1	EXT13
NTAP_HVFT_Cluster-01	e0f	UTA2 - 10 GbE	No	None	None	None
NTAP_HVFT_Cluster-01	e0g	UTA2 - 10 GbE	Yes	FCoE to Host	Flex System Switch2	EXT13
NTAP_HVFT_Cluster-01	e0h	UTA2 - 10 GbE	No	None	None	None
Controller 2						
NTAP_HVFT_Cluster-02	e0e	UTA2 - 10 GbE	Yes	FCoE	Flex System Switch2	EXT14
NTAP_HVFT_Cluster-02	e0f	UTA2 - 10 GbE	No	None	None	None
NTAP_HVFT_Cluster-02	e0g	UTA2 - 10 GbE	Yes	FCoE	Flex System Switch1	EXT14
NTAP_HVFT_Cluster-02	e0h	UTA2 - 10 GbE	No	None	None	None

5.4.1 Initial storage setup

This section describes the initial setup of the storage cluster system and the layout of the disks. It is assumed that the disks are properly connected to the controller (for the redundant path, see Figure 5-4 on page 44) and that the system is connected to the network and accessible.

The Storage Cluster setup is performed by using the System Setup tool, which is one of the options that is available for setting up the storage that is easy to follow. Manual setup by using the command line is also an option; however, System Setup is a recommended path for users who are less familiar with the FAS storage. System Setup can be downloaded from the following NetApp support website:

http://mysupport.netapp.com

A high-level overview of the disk structure that is used is shown in Figure 5-5 on page 46.

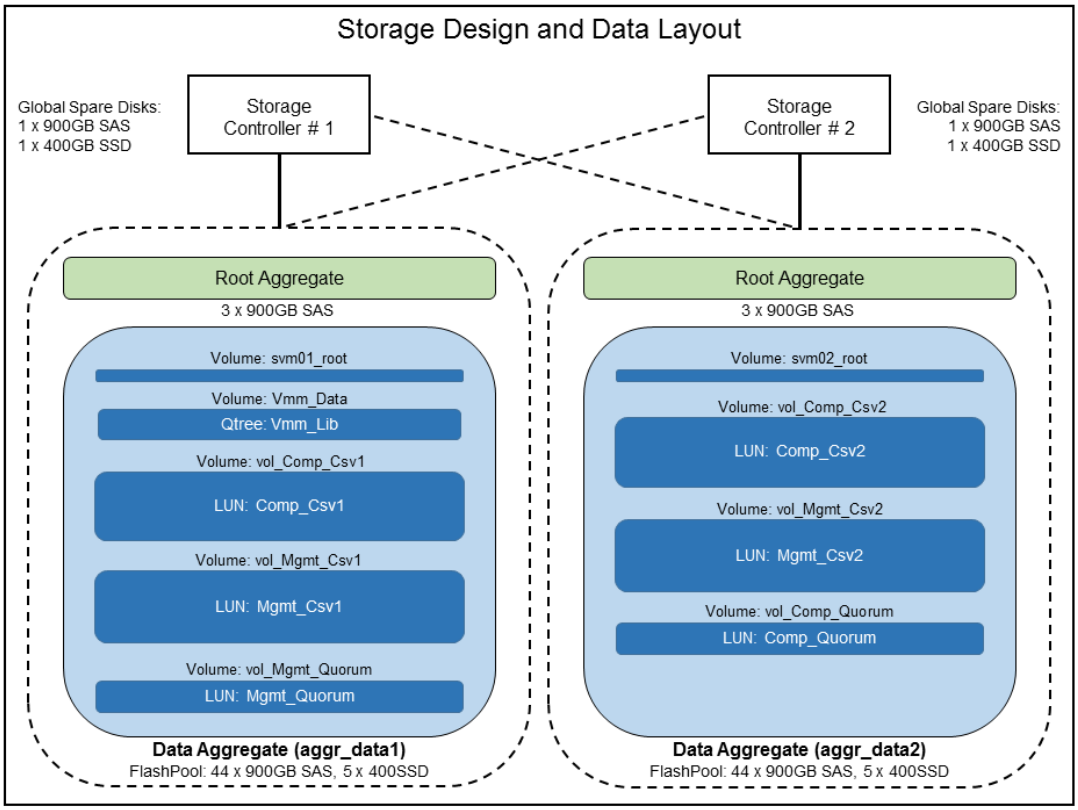

Figure 5-5 Storage Layout of FAS8040

System Setup should be run from your management console with access to the management network and automatically discovers your networked storage. As such, the storage controllers must be connected to the network. A list of all network ports is available in the storage section of Appendix A, “Appendices” on page 113.

Figure 5-6 shows System Setup that is listening for DHCP requests as the controllers boot and assigns a temporary address to use for configuration.

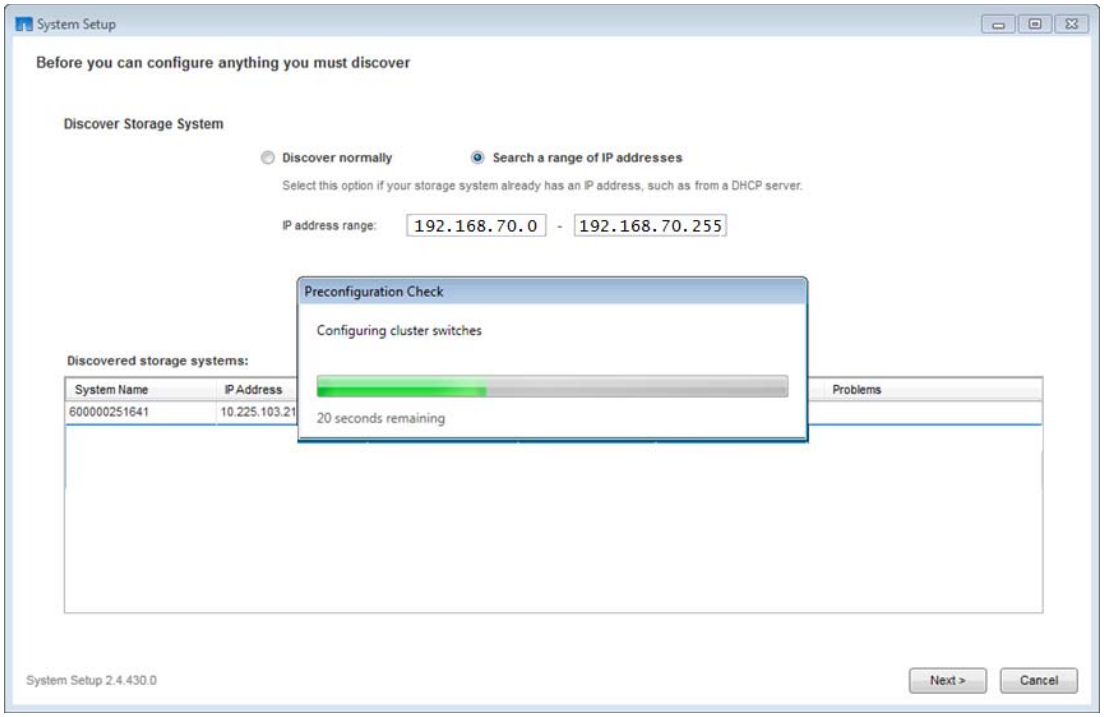

Figure 5-6 Storage discovery for setup

Complete the following steps to set up the storage cluster system:

1. As seen in Figure 5-7, enter the configuration information for the cluster, including the Cluster name and the IP address that is used to manage the cluster. (System Setup completes some default values for node management addresses for the individual nodes.) Complete the information about the Mask, gateway, DNS setup, and the Cluster's password.

   Click **Next** and System Setup creates the inter-cluster network interfaces and the cluster, assigns a management logical interface (LIF), connects to the other node in the cluster, joins the cluster, and then starts configuring.

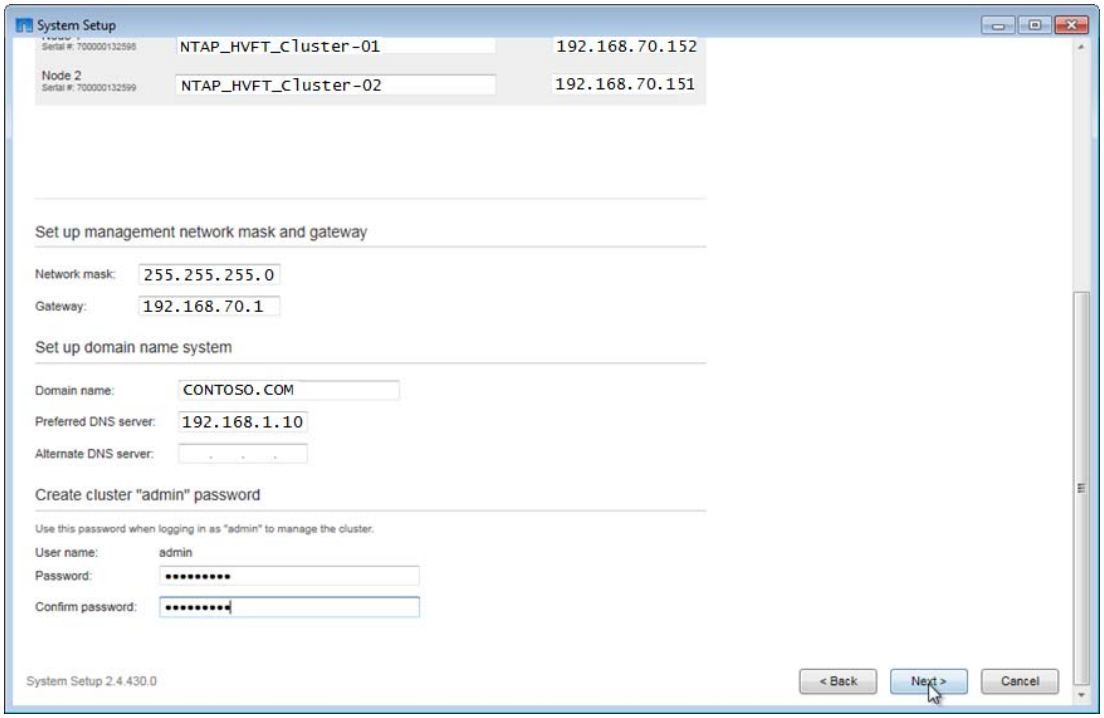

Figure 5-7 Assigning names and network addresses for storage

2. Apply all relevant license keys, as seen in Figure 5-8. For this solution, FC license is required to enable FCoE and CIFS license is used as an option. In our lab system, we enabled more protocols than what is required for this solution. In addition, NetApp recommends that AutoSupport is Enabled so that weekly reports about the health of the customer's system are sent to NetApp Support.

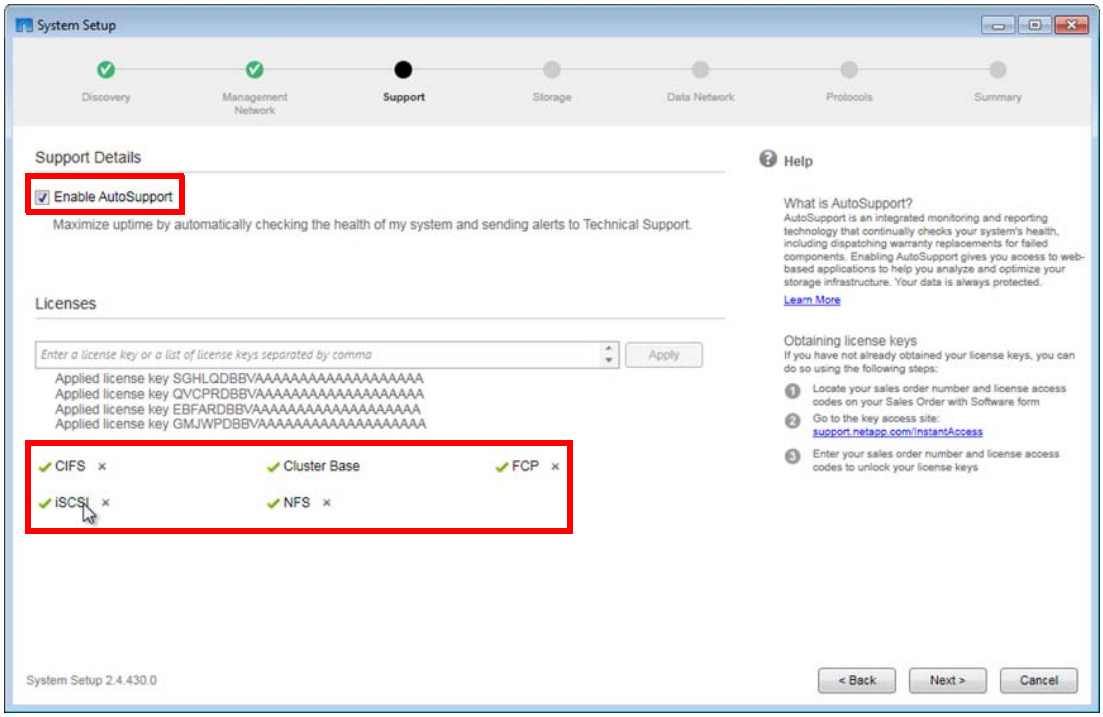

Figure 5-8 License key configuration

3. As seen in Figure 5-9, System Setup automatically arranges disks into the largest possible, recommended configuration for optimal performance and efficiency. You have three choices in this setup, but we create the disk pools (aggregates) later so here by selecting **Do Not Configure**.

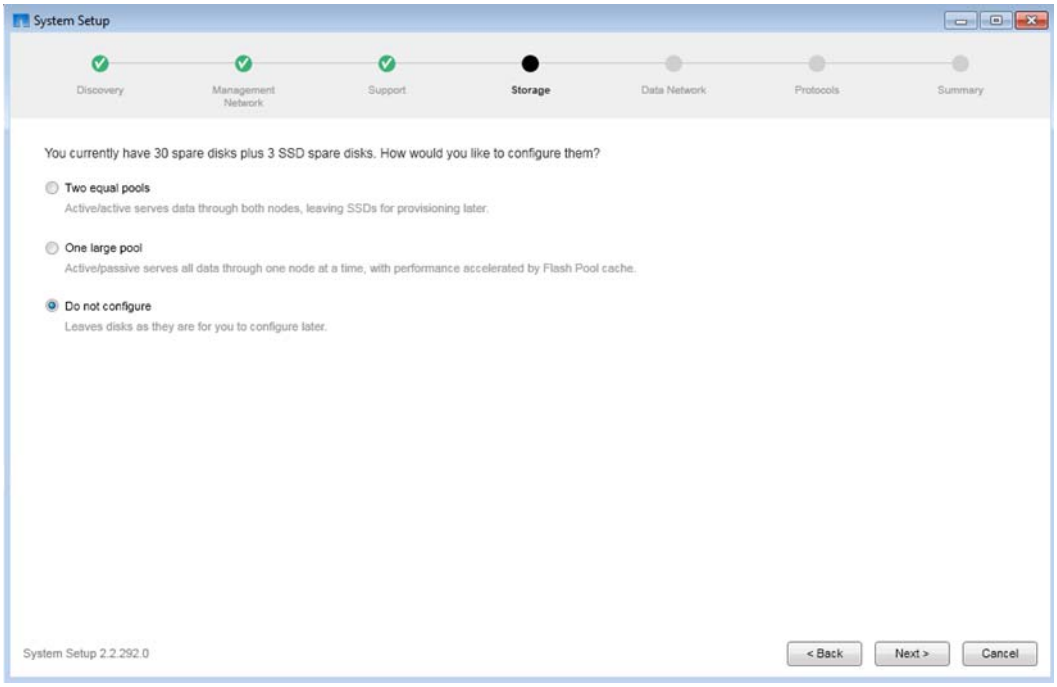

Figure 5-9 Disk configuration

4. Figure 5-10 shows the creation of the virtual storage server (SVMs). The configuration includes two SVMs (one on each of the two storage controllers): `svm01` and `svm02`.

   We use System Manager for creating the SVM and skip this step in the System Setup tool. It is common to create multiple SVMs and other components; therefore, working within the same tool often is more productive. The System Setup tool is designed for initial setup only and supports creating the first aggregate.

   Ensure that the **Create first storage VM and set up data network option** is not selected, as shown in Figure 5-10. Click **Next**.

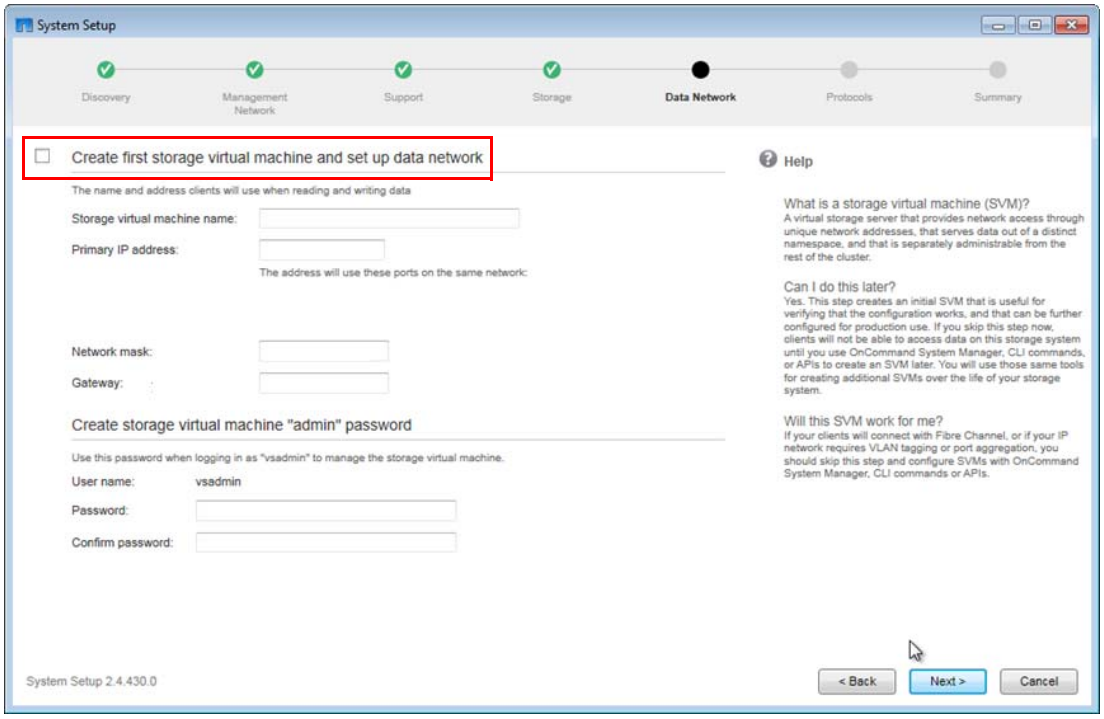

Figure 5-10 Creating SVMs

5. Follow the System setup wizard to the next step to enable the protocols for the storage system (licenses must be applied in the previous steps) and complete the setup.

5.4.2 Creating aggregates

In this section, we use OnCommand System Manager to create more aggregates that are used for data. OnCommand System Manager must be installed and the Clustered Storage must be added before these steps are started. For more information about how to install System Manager, see Appendix A, “Appendices” on page 113.

Complete the following steps to create aggregates:

1. As shown in Figure 5-11, use the OnCommand System Manger and expand Storage in the left pane. Click **Aggregates** and then click **Create**. Select the first controller (node NTAP_HVFT_Cluster-01).

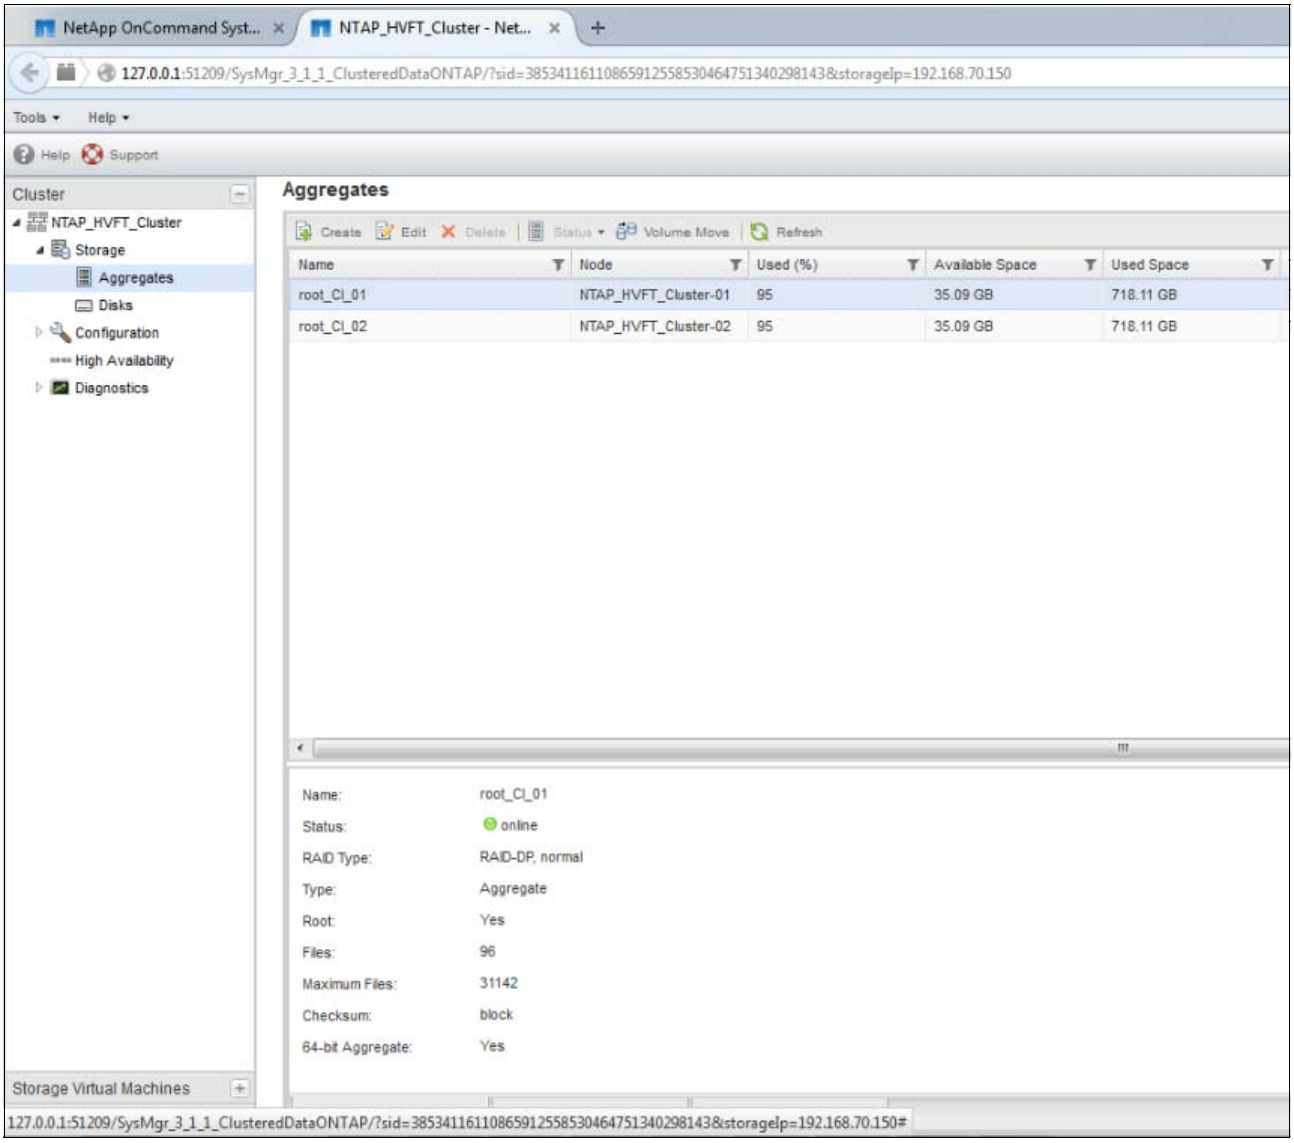

Figure 5-11 Creating aggregates

2. Click **Next** to proceed with initial storage configuration of the new system. Then, click **Next** in the Wizard welcome window.

3. Choose the appropriate option for your wanted configuration of aggregates, as shown in Figure 5-12. For this example, we choose the RAID-DP and Flash Pool options. We name the aggregate `aggr_data1`.

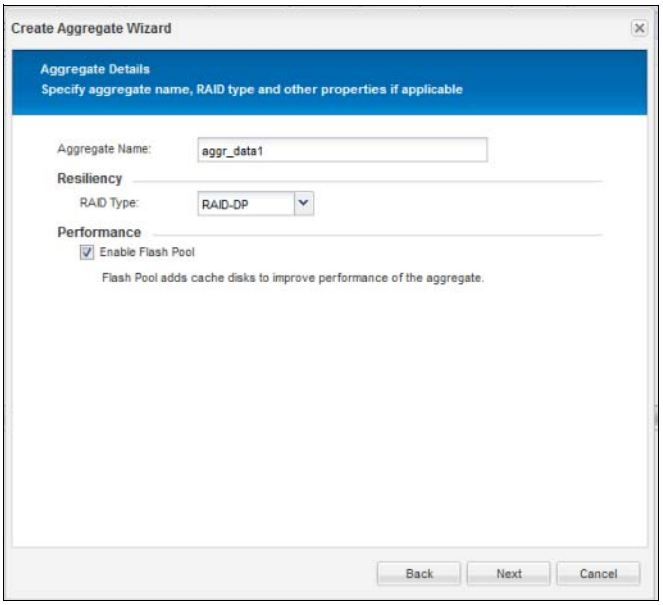

Figure 5-12 Selecting RAID option

4. Specify the number of disks that are used to build the aggregate, as shown in Figure 5-13.

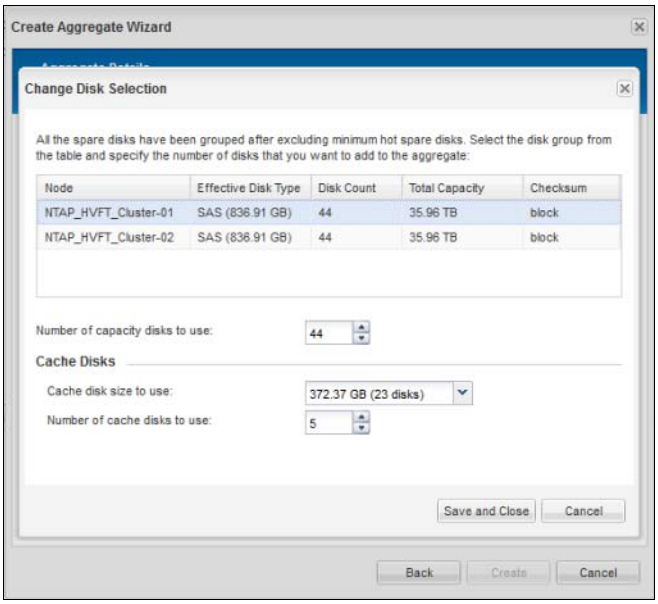

Figure 5-13 Disk selection

The current configuration of disk assignment to the aggregate is specific to this reference architecture and “as built” in our lab.

It is important to carefully design the storage and calculate the required number of disks from each type (SAS and SSD) to properly address performance that is required and benefit from the Flash Pool feature. Also, consider data protection aspects, what workloads are served on the aggregate, and expected growth.

Aggregates with Flash Pool have lower limited capacity than standard aggregates that are not Flash Pool. For more information about those limitations, see the following NetApp support website before you design your storage layout:

http://mysupport.netapp.com

For our specific system, each storage controller has one aggregate for data (`aggr_data1`) that has 44 data disks (900 GB 10 K SAS) and five 400 GB SSDs. We do not configure parity disks or spare disks as part of this process. Parity disks are configured automatically later and spare disks are global per controller.

5. After selecting the disk, select the size of the RAID group.

   In our example, the RAID group size is 15 disks, as shown in Figure 5-14. Although the storage system can accommodate different-sized RAID groups, best practices should be followed because there is trade-off balance to consider and maintain to optimize the performance, minimize exposure to data loss risks, and support growth. For more information about best practices and other references to the NetApp storage system, see Appendix A, “Appendices” on page 113.

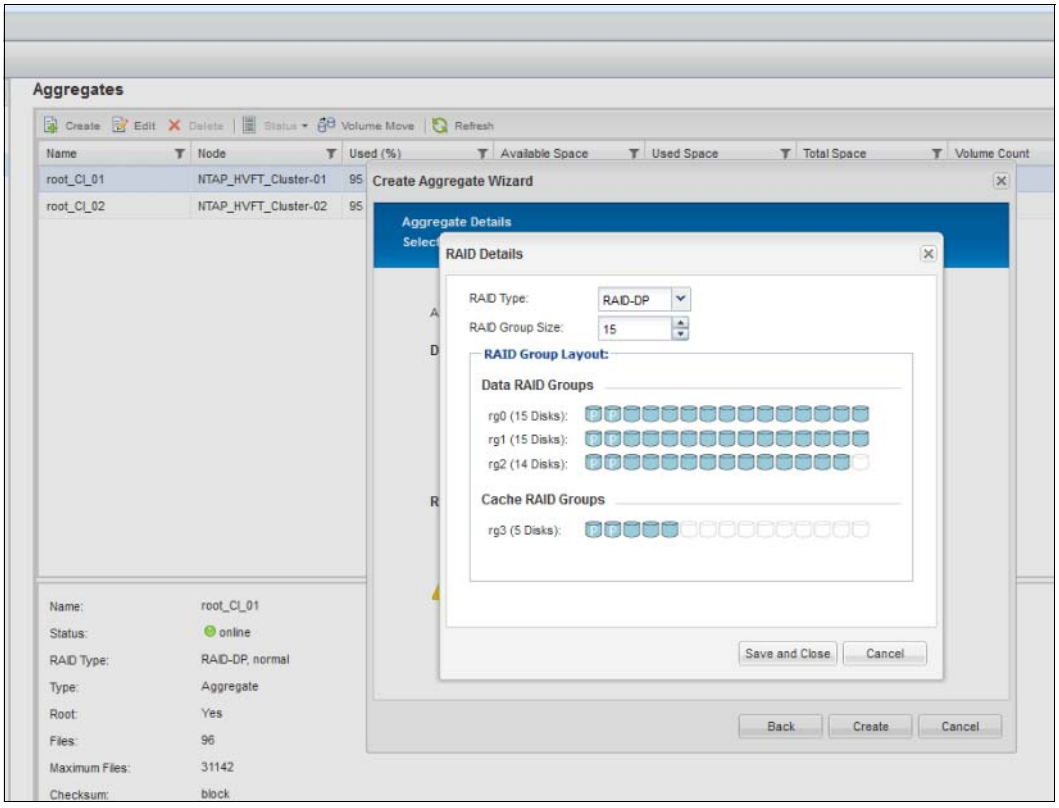

Figure 5-14 Selecting RAID group size

6. After selecting the RAID group size, save this option and close. Then, click **Finish** for the system to create the aggregate. Take note of the comment about Flash Pool and how it can affect the timing for creating the aggregate. Flash Pool is used in our lab system for this reference document as shown in Step 4 of this process.

Repeat this process for the second controller. We follow the same steps to create `aggre_data2` for storage node `NTAP_HVFT_Cluster-02`.

You can use System Manager to see the two created aggregates, as shown in Figure 5-15.

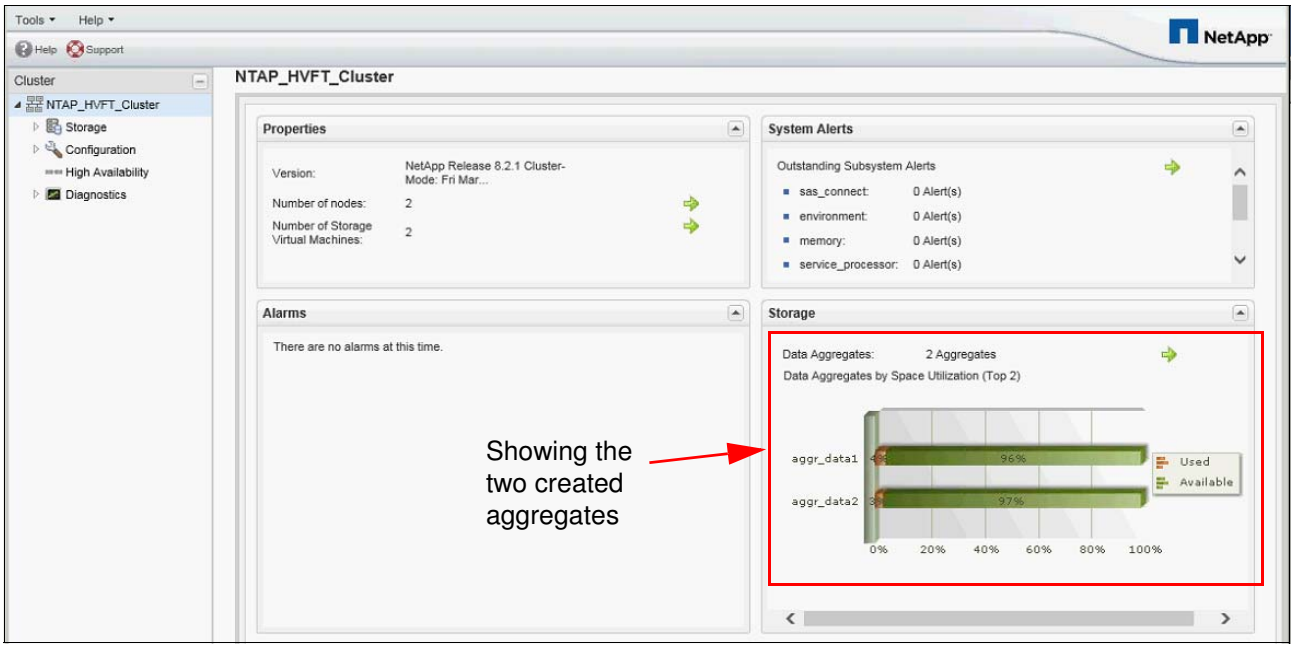

Figure 5-15 Aggregates view within System Manager

5.4.3 Creating storage virtual machines

After creating the aggregates, we can configure the SVM. We use System Manager to set up SVM and configure the protocols that are used in our configuration, CIFS, and FC/FCoE.

Note: CIFS and FCoE protocols must be licensed on the storage cluster.

As part of this workflow, we perform the following overall tasks:

- Set up an SVM with FlexVol volumes.
- Configure CIFS and FCoE protocols on the SVM.
- Set up an SVM administrator account.

Complete the following steps:

1. From the Home tab, double-click the appropriate storage system.
2. Expand the SVM's hierarchy in the left navigation pane.
3. In the navigation pane, select the cluster.
4. Click **Create**.
5. In the SVM Setup window, specify the SVM details, such as the SVM name, select **FlexVol volumes** as the volume type, select **CIFS** and **NFS** as the protocols that are allowed, and specify the SVM language, root volume security style, and an aggregate for the root volume.

   The default language setting for an SVM is UTF-8 (EN_US Unicode encoding).

   By default, the aggregate with the maximum amount of free space is selected as the container for the root volume of the SVM. Based on the protocols that are selected, the default security style and the root aggregate also are selected.

The security style is set to NTFS if you select CIFS protocol or a combination of CIFS protocol and the other protocols. The security style is set to UNIX if you select NFS, iSCSI, or FC/FCoE or a combination of these protocols.

6. Specify the DNS domain names and the name server IP addresses to configure the DNS services. The default values are selected from the SVM configurations.
7. Click **Submit & Continue**.

The SVM is created with the specified configuration as shown in Figure 5-16. The SVM that you created starts automatically. The root volume name is automatically generated as `SVM name_root`. By default, the `vsadmin` user account is created and is in the locked state.

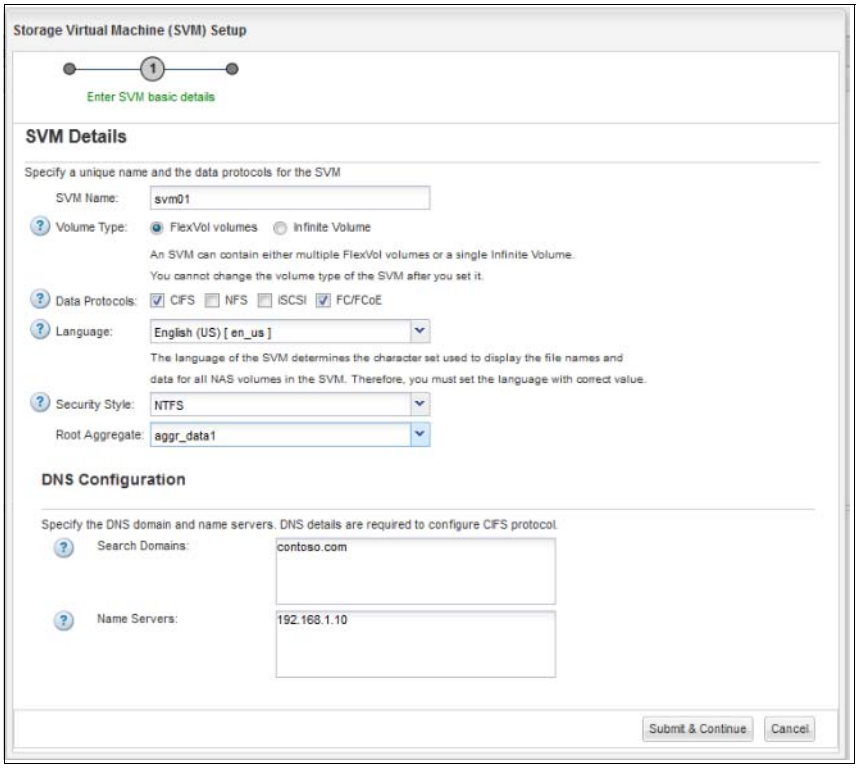

Figure 5-16 SVM configuration

Configuring CIFS and NFS protocols on an SVM

You can configure CIFS and NFS protocols on the SVM to provide file-level data access for NAS clients. To enable the CIFS protocol, you must create the data LIFs and the CIFS server. To enable the NFS protocol, you can specify the NIS details and the data LIFs. NFS is not used in this solution and CIFS protocol is used, but optional.

Complete the following steps:

1. In the Data LIF Configuration section, specify the network details to create data LIFs. You can retain the same data LIF configuration for CIFS and NFS, or configure a new LIF for each protocol.
2. Specify the following information to create a CIFS server:
   - CIFS server name.
   - Active Directory to associate with the CIFS server.
   - Organizational unit (OU) within the Active Directory domain to associate with the CIFS server; by default, this parameter is set to CN=Computers.

– Credentials of an administrative account that has sufficient privileges to add the CIFS server to the OU.
– Optional: You can also specify the IP addresses of the NIS servers and NIS domain name to configure NIS services on the SVM.

3. Click **Submit & Close**, as shown in Figure 5-17.

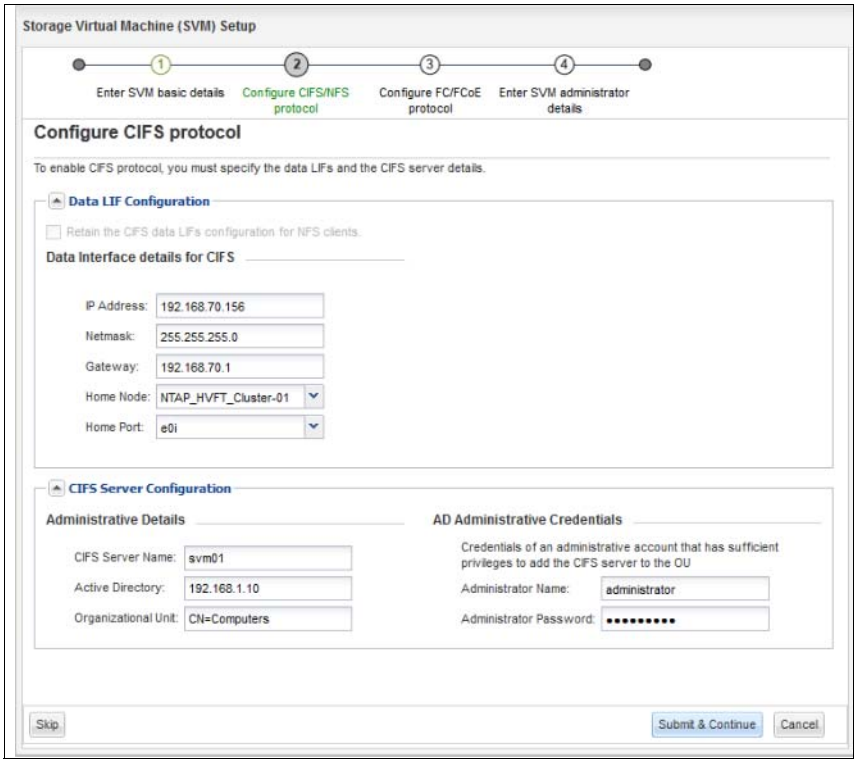

Figure 5-17 CIFS setup

Configuring FC and FCoE

You can configure FC and FCoE protocols on the SVM to provide block-level data access for SAN clients. To enable the FCoE protocol, you must create the data LIFs. To enable the FCoE protocol, you can specify the FCoE details and the data LIFs.

Complete the following steps:

1. In the Data LIF Configuration section, specify the network details to create data LIFs. You can retain the same data LIF configuration for CIFS and NFS or configure a new LIF for each protocol.

2. Click **Submit & Close**, as shown in Figure 5-18.

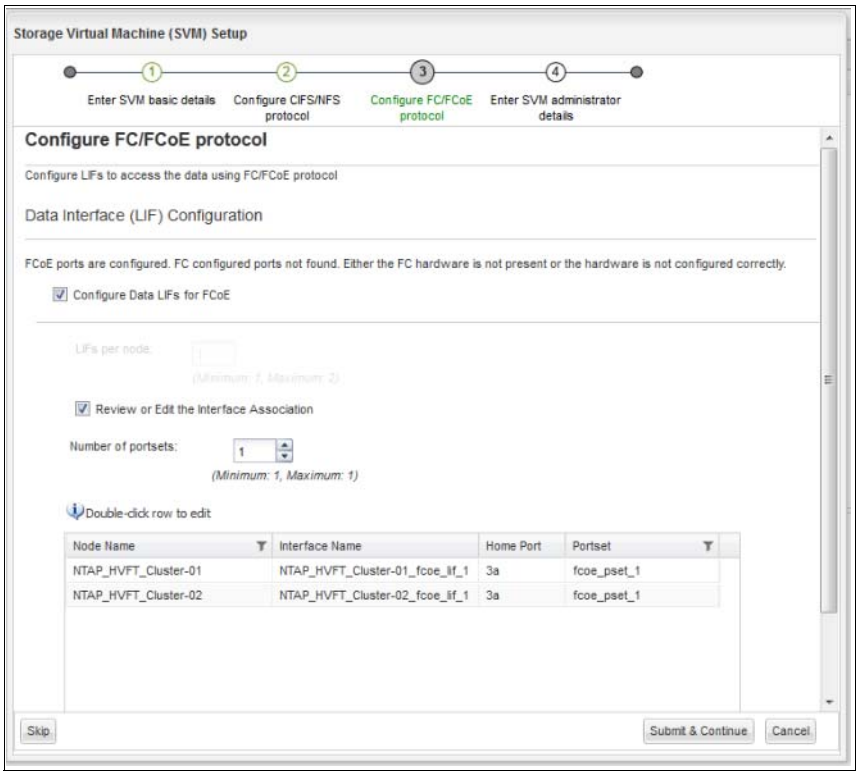

Figure 5-18 SVM protocol setup for FCoE

Delegating administration to an SVM administrator

After setting up a functional SVM or an SVM with basic network configuration, you can optionally delegate the administration of the SVM to an SVM administrator.

In the Administrator Details section, set up a password for the vsadmin user account.

By default, the data LIFs have management access, as shown in Figure 5-19 on page 59.

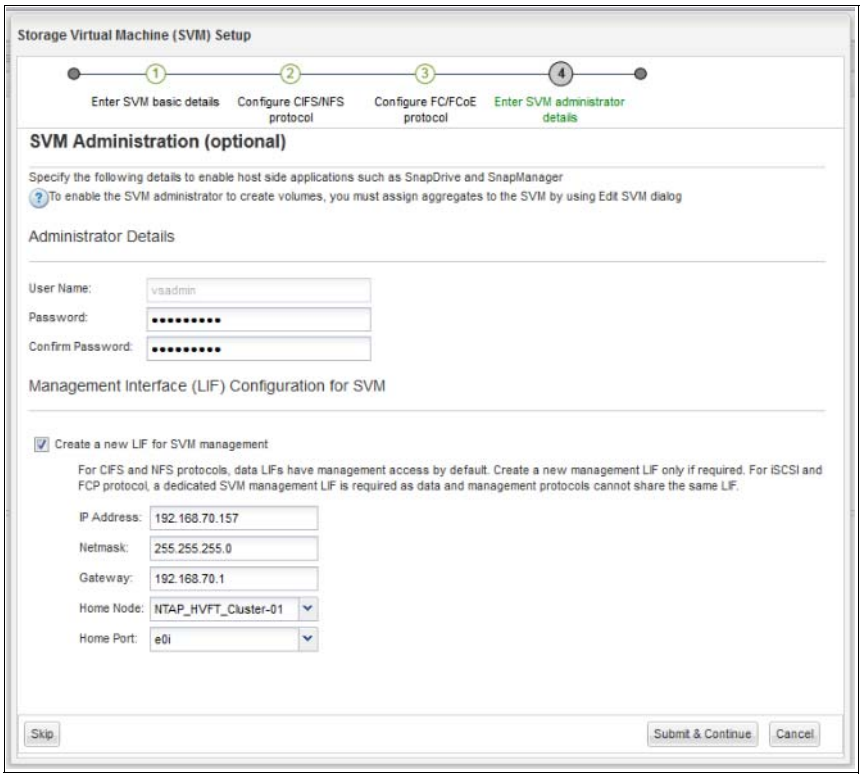

Figure 5-19 SVM Administrator setup

The SVM (or SVMs) can be seen in System Manager, as shown in Figure 5-20.

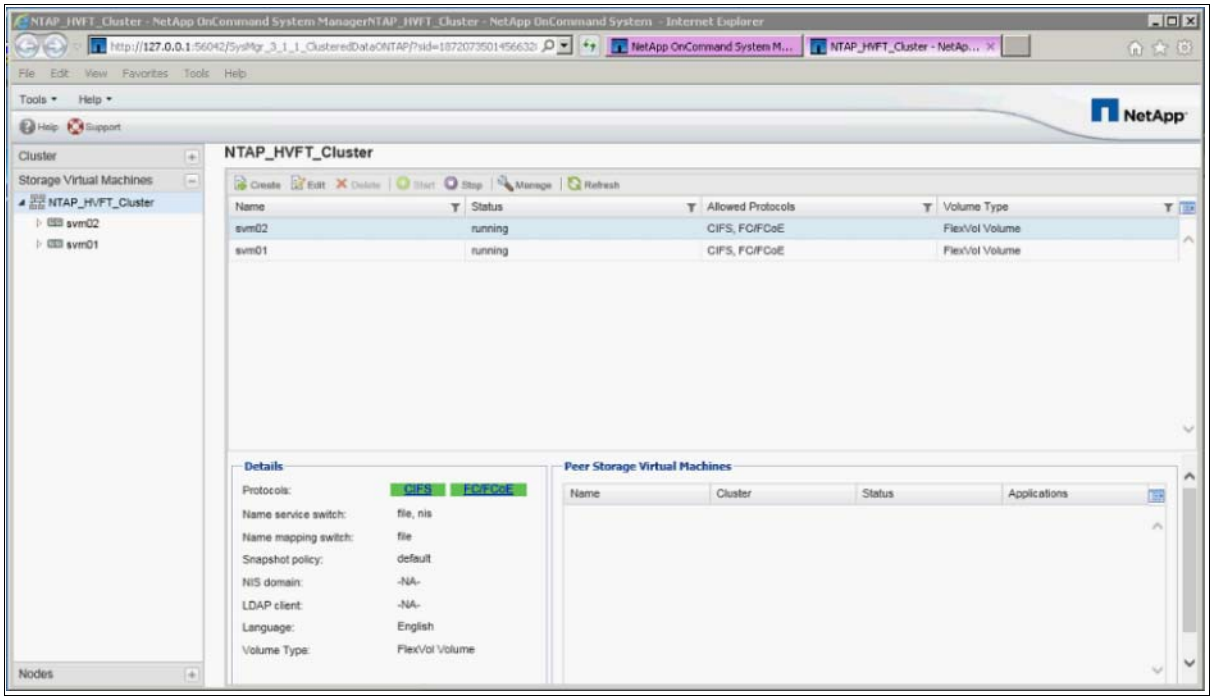

Figure 5-20 System Manager view of the SVMs

5.4.4 Configuring storage volumes

This section describes the process of creating storage volumes (FlexVol) that are required step before the LUNs that connect to the servers are created. Volumes easily can be created by using System Manager. Figure 5-5 on page 46 shows the volumes that are used in this configuration and the capacity that is allocated to each volume as a reference.

Volume sizes vary depending on the size of the data and the usage of storage efficiency features. Volumes sizes that are described in this section are for reference only and do not represent best practices or minimum required capacity for which to plan.

5.4.5 NetApp FAS and Windows cluster storage considerations

Microsoft Windows Failover Clustering supports Cluster Shared Volumes (CSVs). Cluster Shared Volumes provide shared primary storage for VM configuration files and virtual hard disks consistently across the entire cluster. All CSVs are visible to all cluster nodes and allow simultaneous access from each node. For optimum performance, CSV volumes should be evenly distributed between the two FAS8040 controllers as their preferred path.

These dedicated storage LUNs are assigned across the storage VMs to help distribute the workload across the controllers. Host mapping uses a host's worldwide port names to limit host access to those LUNs they are authorized to access. Thin provisioning is standard and is considered part of the best practices. With thin provisioning, not all of the required capacity is allocated during the creation of the LUN; capacity is added automatically from a shared pool only if needed. The host is not aware of the actual allocated size.

For more information about best practices for setting up the storage and following the steps that are described in this solution, see NetApp Technical Report (TR) 4226 *NetApp SnapManager 2.0 for Hyper-V on Clustered Data ONTAP 8.2*, which is available at this website:

`http://www.netapp.com/us/media/tr-4226.pdf`

Other procedures that are describes in this paper assure that Windows failover cluster works as expected.

5.4.6 Defining the host server and mapping the LUN

The next step is to define the host servers to the FAS storage system and map LUNs to them. This process limits access to the LUNs to authorized host servers only.

Before setting up the LUNs and mapping to the hosts, we must ensure that SnapDrive and NetApp MPIO drivers are installed on the servers. We also want to ensure that the volumes are created because the LUNs are components that are in a volume on the NetApp storage and that the storage controllers are connected to the Flex System Switches. For more information about the LUNs in this configuration, see Figure 5-5 on page 46.

Note: Several steps are included in this process and they must be followed in order. For more information, see *Installation and Administration Guide* for Data ONTAP DSM 4.1, which is available at this website:

`https://library.netapp.com/ecm/ecm_download_file/ECMP1307349`

For Windows MPIO setup, the following high-level steps are used:

1. Ensure that each Windows Server includes Microsoft DSM.
2. Install NetApp SnapDrive for Windows (version 7 or later).
3. Install the NetApp DSM.
4. Record the FC initiator identifiers.
5. Create LUNs and map them to the respective host servers.

6

Flex System x240 M5 Management Fabric Setup

The Flex System x240 M5 Management Fabric setup procedure is explained in this chapter.

This chapter includes the following topics:

6.1 Cluster management servers

The management fabric for this environment is built around a two-node Windows Server cluster that consists of two dual-socket Flex System x240 M5 Compute Nodes with 256 GB of RAM, and one CN4054R 4-port converged network adapter for each node. With redundant CN4093 converged switches, and fault tolerant FAS8040 storage system and the clustered Flex System x240 M5 servers eliminates single points of failure and allows near continuous important management environment.

The clustered management servers run Windows Server 2012 R2 with Hyper-V, and Failover Clustering to host highly available virtual machines (VMs) that are running various System Center 2012 R2 applications. Lenovo XClarity Upward Integration Modules are installed to provide improved granularity and visibility of Lenovo hardware components in the management fabric. The high availability (HA) configuration that is used for each management component varies depending on the needs and capabilities of the individual component. A high-level overview of the management fabric is shown in Figure 6-1 on page 65.

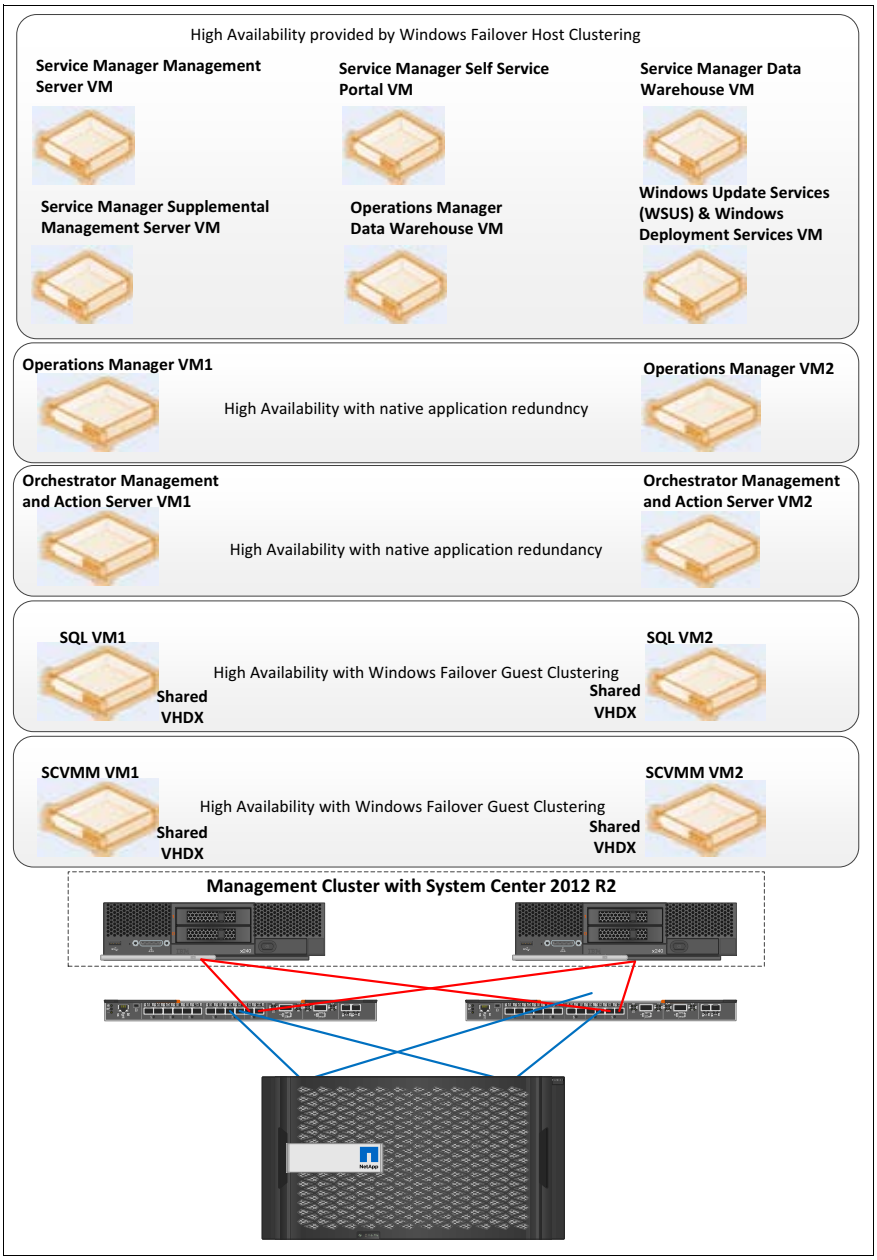

Figure 6-1 HA provided by Windows Failover Hosting Clustering

Setup involves the installation and configuration of Windows Server 2012 R2 Datacenter edition, networking, and storage on each node. Highly available VMs can then be created to perform the various management tasks for the management framework.

6.2 Pre-operating system installation

Complete the following steps before an operating system is installed:

1. Confirm that the CN4054R 4-port Ethernet devices are installed in each compute node.

   The FCoE Feature on Demand (FoD) must be imported on each node through an IMM connection, as shown in Figure 6-2 on page 66.

Note: The default IMM address for each x240 M5 Compute Node is 192.168.70.1xx, where xx is equal to the two-digit slot number that the compute node is installed in (Slot 1 = 01)

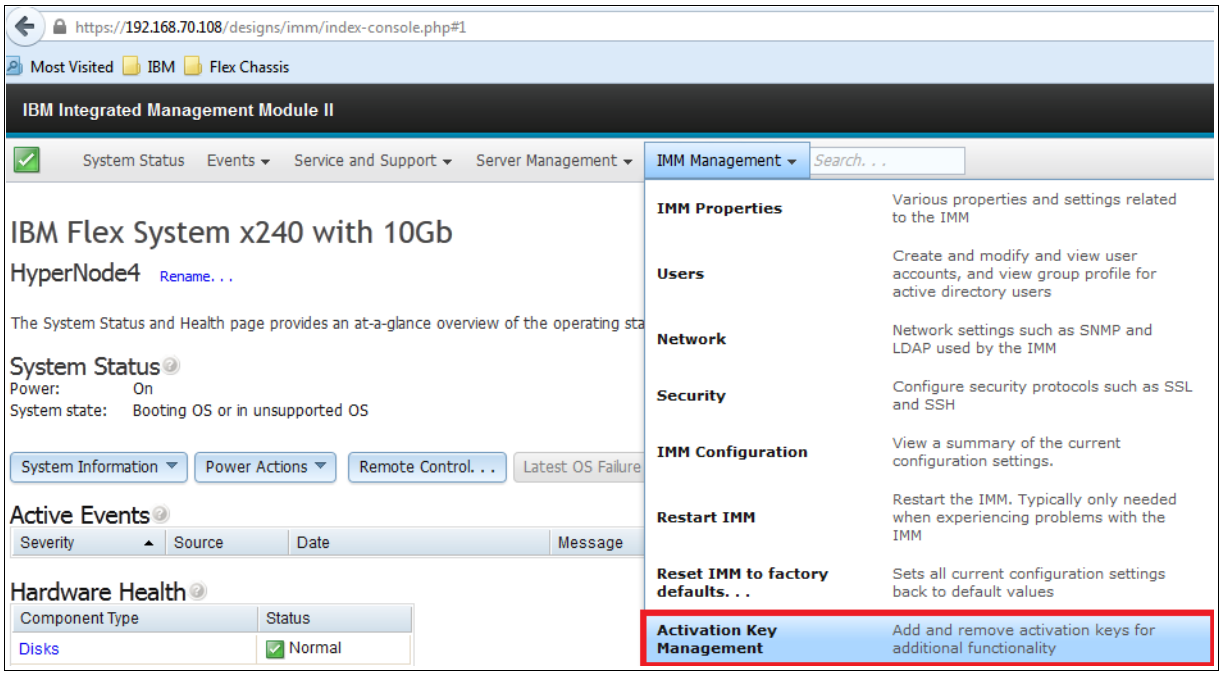

Figure 6-2 FoD key activation

2. It is best to include any FoD keys that are needed with the initial order so that they can be installed at the factory. If FoD keys are ordered separately, contact your Lenovo Sales representative for assistance.

   For more information about FoD, see the Lenovo Press paper, *Using Features on Demand*, which is available at this website:

   `http://lenovopress.com/redp4895`

3. If the FoD Unique Identifier (FUI) is needed, it can be found in the Unified Extensible Firmware Interface (UEFI) by selecting **System Settings** → **Network ... Emulex Device** → **Feature on Demand**. There should be two unique FUI per CN4054 card, one for each ASIC.

   Note: After Windows and the Emulex OneCommand tool is installed, the `HBACmd` utility can be used to retrieve the FUI by using the following format:

   ```
   C:\Program Files\Emulex\Util\OCManager> HbaCmd getfodinfo xx-xx-xx-xx-xx-xx
   ```

   Where `xx-xx-xx-xx-xx-xx` is the MAC address of the first NIC port on each ASIC.

4. The Flex x240 M5 firmware for the UEFI, IMM, DSA, SAS, and CNA devices must be evaluated and flashed to the latest firmware, if necessary:

   – For out-of-band updates, Bootable Media Creator creates a bootable image of the latest Flex x240 M5 updates (download previously) and is available at this website:

   `http://www.ibm.com/support/entry/portal/docdisplay?lndocid=TOOL-BOMC`

An external DVD device is required or mounted to the server by using the virtual media capabilities of the IMM.

For more information, see *ToolsCenter Bootable Media Creator Installation and Users Guide 9.41*, which is available at this website:

http://download.boulder.ibm.com/ibmdl/pub/systems/support/system_x_pdf/bomc_userguide.pdf

– Fast Setup is an optional tool that can be downloaded and used to configure multiple System x, BladeCenter, or Flex System servers simultaneously and is available at this website:

http://www-947.ibm.com/support/entry/portal/docdisplay?lndocid=TOOL-FASTSET

– In-band updates can be applied by using the UpdateXpress tool to download and install recommended updates on the Flex System platform.

For more information about UpdateXpress, see this website:

http://ibm.com/support/entry/portal/docdisplay?lndocid=serv-xpress

For more information, see *UpdateXpress System Pack Installation and Users Guide*, which is available at this website:

http://ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5085892

5. By default, the x240 M5 Compute Node settings are set to balance power consumption and performance. To change these settings, boot to UEFI mode and select **System Settings** → **Operating Mode**, as shown in Figure 6-3. Change the settings to fit your organization's needs.

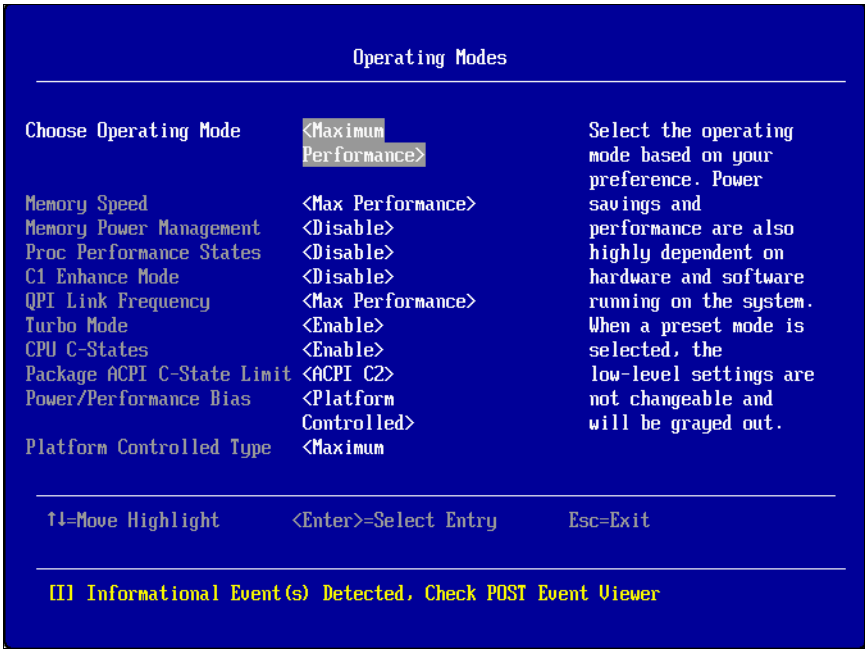

Figure 6-3 Operating Mode settings in UEFI

6. The Multichannel settings subdivide the network interface into four virtual network interfaces. In this configuration, we use native features of Windows Server 2012 R2 to present virtual interfaces to the operating system and VMs. The Multichannel mode can be disabled on the CN4054R in the UEFI by selecting **System Settings** → **Network … Emulex Device**, as shown in Figure 6-4 on page 68.

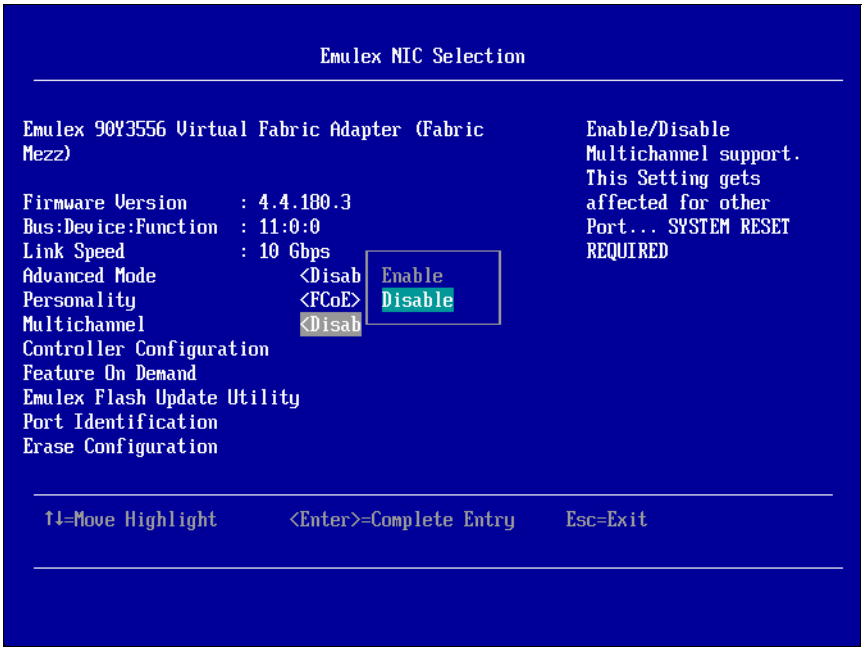

Figure 6-4 Disabling Multichannel Mode in CN4054R in UEFI

7. Enable FCoE on the CN4054R converged network adapter through the UEFI be selecting **System Settings** → **Network** → **Network Device List**. Select the top device on each bus, as shown in Figure 6-5.

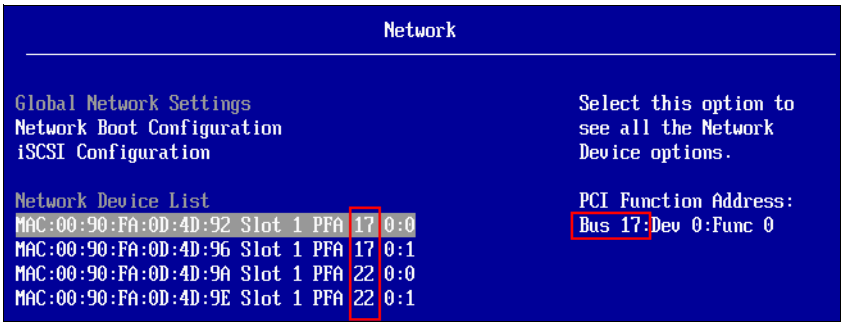

Figure 6-5 Identifying CNA devices to be configured

8. Change the setting from `NIC` to `FCoE`, as shown in Figure 6-6.

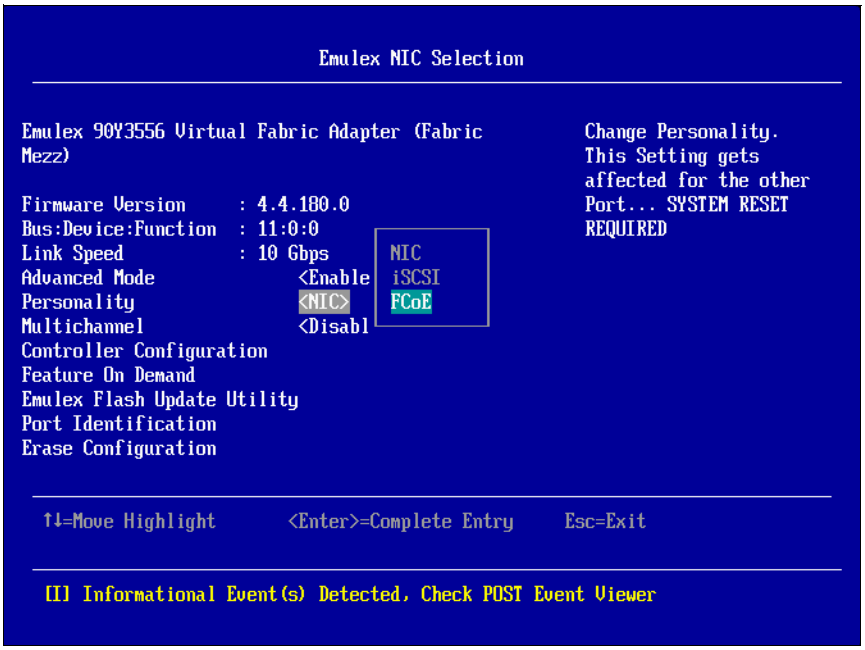

Figure 6-6 Changing CNA property to support FCoE

Repeat this step for the second device on the other bus (second ASIC).

9. Set the IMM VLAN ID to the predetermined VLAN. VLAN settings are disabled by default in the IMM settings by selecting **UEFI** → **Systems Settings** → **Integrated Management Module** → **Network Configuration** → **VLAN Support**.
10. The two local disks should be configured as a RAID 1 array.
11. Confirm that CN4093 switches were configured. Consider the following points:
    - Stacking, LACP ports defined, and VLAN configuration must be created and assigned as described in 4.2, “VLANs” on page 21.
    - Switch FCoE configuration is easier after the operating system is installed and WWPN can be collected for each server.
12. The volumes and LUNs on the FAS8040 must be configured as described in Chapter 5, “NetApp Storage System” on page 37 to be ready for host mapping assignments.

    LUN mapping is easiest after the operating system is installed and WWPN can be collected for each server.

6.3 Installing and configuring the operating system

Windows Server 2012 R2 Datacenter allows unlimited Windows Server VM instances or licenses on the host servers and is the preferred version for deploying Hyper-V compute configurations. Windows Server 2012 R2 Standard also supports clustering, but provides licensing rights only for up to two Windows Server VM instances (more licenses are needed for extra VMs). Windows Server 2012 R2 Standard Edition is intended for physical servers that have few or no VMs that are running on it.

Complete the following steps to install and configure the operating system:

1. Install Windows Server 2012 R2 Datacenter Edition.
2. Set your server name and join the domain.
3. Install the Hyper-V role and Failover Clustering feature.
4. Run Windows Update to ensure that any new patches are installed.
5. The latest CN4054R NIC and FCoE firmware and drivers must be downloaded from the Fix Central and installed. A complete package of all platform-related updates by operating system is available from Fix Central as an UpdateXpress System Pack, which is available at this website:

   http://www.ibm.com/support/fixcentral/

   The associated Windows Installer also is download separately on Fix Central.
6. Install the Emulex OneCommand Manager utility to provide more information for the CN4054R converged network adapter. The utility is available at this website:

   http://www.emulex.com/downloads/oem-qualified-downloads/ibm/vfafc-software-kits

   Multipath I/O is used to provide balanced and fault tolerant paths to the FAS8040 storage. This multipath I/O requires another NetApp MPIO DSM driver to be installed on the host servers before attaching the storage. The driver is available at this website:

   http://mysupport.netapp.com/NOW/cgi-bin/software/?product=Multipath+I%2FO&platform=Windows

   For more information about the NetApp MPIO DSM driver, see Appendix A, “Appendices” on page 113.
7. Install the IBM Systems Director platform agent 6.3.5, which is available at this website:

   http://www-03.ibm.com/systems/director/downloads/agents.html

6.4 Network configuration

One key new feature of Windows Server 2012 is in-box NIC teaming. In-box teaming can provide fault tolerance, link aggregation, and can be tailored to host or VM connectivity. Two separate Windows Server 2012 R2 teams are created for each host in following configuration. One team is used to support host server management and cluster private traffic; another team is used to support Live Migration and VM Communication.

Note: Be careful when you are identifying and enumerating the network interfaces in each host to ensure that teams are spread across the two network interfaces and properly routed to the correct switches. Use Emulex OneCommand Manager Utility to review each network interface and MAC address.

Note: For more information about mapping server ports to the two CN4093 switch ports, see Figure 4-2 on page 18 and Figure 4-3 on page 19. Each ASIC has two network interfaces: the top one is for switch 1, the bottom for switch 2.

The following PowerShell commands also can be useful (for more information about other PowerShell scripts, see “PowerShell scripts” on page 132). Each team must contain a member from each “bus”:

- `Get-NetAdapterHardwareInfo | Sort-Object –Property Bus,Function`
- `Get-NetAdapter –InterfaceDescription “IBM*”`
- `Rename-NetAdapter –name “Ethernet” –NewName “Port_IntA9 SW1”`

Figure 6-7 shows an NIC interface naming example.

```
PS C:\Users\Administrator.CONTOSO> Get-NetAdapter -InterfaceDescription "IBM O*"

Name                      InterfaceDescription                    ifIndex Status       MacAddress
----                      --------------------                    ------- ------       ----------
Port_IntB9 SW2            IBM OCm11104-N-X Emulex Virtual Fa...#2      13 Up           00-90-FA-6A-4E-42
Port_IntB9 SW1 VLAG       IBM OCm11104-N-X Emulex Virtual Fabr...      12 Up           00-90-FA-6A-4E-3E
Port_IntA9 SW2 VLAG       IBM OCm11104-N-X Emulex Virtual Fa...#4      15 Up           00-90-FA-6A-4E-3A
Port_IntA9 SW1            IBM OCm11104-N-X Emulex Virtual Fa...#3      14 Up           00-90-FA-6A-4E-36
```

Figure 6-7 Displaying all CN4093R network adapters with PowerShell

Windows Server 2012 R2 in-box NIC teaming can be found in the Server Manager console, as shown in Figure 6-8. The basic NIC teaming tools are available in the Server Manager GUI however it is better to use PowerShell for the added options and flexibility.

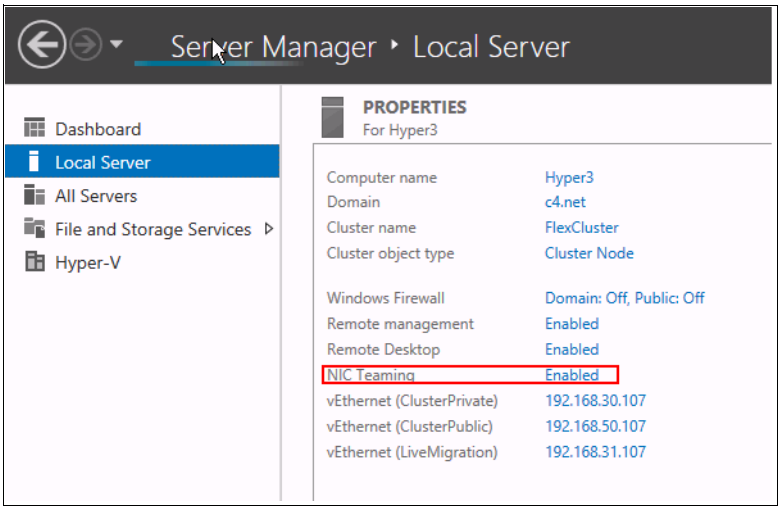

Figure 6-8 NIC teaming in Server Manager

Note: Run `LBFOAdmin.exe` from the command line to see the Network Teaming graphical user interface.

The team that is created to support cluster public and private communication between the host servers share ports that are used for FCoE traffic. The CN4093 switches prioritize FCoE traffic over Ethernet data traffic. To further reduce the potential of bandwidth over subscription, quality of service (QoS) limits are placed on these interfaces in Windows Server 2012 R2. This team must be created by using the two NIC ports, as described in Table 4-2 on page 27.

Complete the following steps:

1. Create the “ClusterTeam” by using the default “Switch Independent” mode and “Address Hash” mode with the following PowerShell commands:

   ```
   New-NetLbfoTeam –name “ClusterTeam” –TeamMembers “Port_IntA9 SW1”, “Port_IntB9 SW2” –TeamingMode SwitchIndependent
   ```

   Note: The team member names can vary by host server.

2. The second team is created off the NIC interfaces that are configured to support LACP. The use of LACP teaming and is then backed by the full aggregated bandwidth of the two ports, as shown in the following commands:

   ```
   New-NetLbfoTeam –name “VLAGTeam” –TeamMembers “Port_IntA9 SW2 VLAG”, “Port_IntB9 SW1 VLAG” –TeamingMode Lacp
   ```

   When Windows Server 2012 R2 NIC teaming is complete, there should be two teams displayed when queried in PowerShell, as shown in Figure 6-9.

   ```
   PS C:\Users\Administrator.CONTOSO> get-netlbfoteam

   Name                   : VLAGTeam
   Members                : {Port_IntA9 SW2 VLAG, Port_IntB9 SW1 VLAG}
   TeamNics               : VLAGTeam
   TeamingMode            : Lacp
   LoadBalancingAlgorithm : Dynamic
   Status                 : Up

   Name                   : ClusterTeam
   Members                : {Port_IntB9 SW2, Port_IntA9 SW1}
   TeamNics               : ClusterTeam
   TeamingMode            : SwitchIndependent
   LoadBalancingAlgorithm : Dynamic
   Status                 : Up
   ```

   Figure 6-9 Windows Server NIC teaming

3. Virtual switches are created on each of these teams. Each Hyper-V virtual switch often can provide one interface to the management operating system. A second virtual NIC for use by the host operating system is created on the “ClusterTeam” to provide a segregated network path for the Cluster Private/CSV network. The following PowerShell commands are used to create the virtual switches and second virtual NIC:

   ```
   New-VMSwitch –Name “MgmtClusterPublic” –NetAdapterName “ClusterTeam” –MinimumBandwidthMode Absolute –AllowManagementOS $true

   Add-VMNetworkAdapter –ManagementOS -Name “MgmtClusterPrivate” –SwitchName “MgmtClusterPublic”

   New-VMSwitch –Name “VM_Communication” –NetAdapterName “VLAGTeam” –MinimumBandwidthMode Weight –AllowManagementOS $true
   ```

4. Rename the management-facing network interface on the VM_Communication team to reflect the role that it is fulfilling by using the following commands:

   ```
   Rename-VMNetworkAdapter -NewName “MgmtLiveMigration” –ManagmentOS –name “VM_Communication”
   ```

5. Confirm that the network interfaces are available to the management operating systems by using the following Powershell command and as shown in Figure 6-10:

```
Get-VMNetworkAdapter –ManagementOS
```

```
PS C:\Users\Administrator.CONTOSO> Get-VMNetworkAdapter -ManagementOS

Name              IsManagementOs VMName SwitchName        MacAddress   Status
----              -------------- ------ ----------        ----------   ------
MgmtLiveMigration  True                  VM_Communication  0090FA1ECD64 {Ok}
MgmtClusterPrivate True                  MgmtClusterPublic 00155DFAAC02 {Ok}
MgmtClusterPublic  True                  MgmtClusterPublic 0090FA1ECD6C {Ok}
```

Figure 6-10 Displaying all network interfaces created for host partition from vSwitches

6. Assign VLAN IDs to each of these interfaces by using the following PowerShell commands:

```
Set-VMNetworkAdapterVlan –ManagementOS –VMNetworkAdapterName
“MgmtClusterPublic”     –Access –Vlanid 40

Set-VMNetworkAdapterVlan –ManagementOS –VMNetworkAdapterName
“MgmtClusterPrivate”     –Access –Vlanid 30

Set-VMNetworkAdapterVlan –ManagementOS –VMNetworkAdapterName
“MgmtLiveMigration”     –Access –Vlanid 31
```

7. Confirm your Management operating system network adapter names and VLAN assignments by using the following command and as shown in Figure 6-11:

```
Get-VMNetworkAdapterVlan –ManagementOS
```

```
PS C:\Users\Administrator.CONTOSO> Get-VMNetworkAdapterVlan -ManagementOS

VMName VMNetworkAdapterName Mode   VlanList
------ -------------------- ----   --------
       MgmtLiveMigration    Access 31
       MgmtClusterPrivate   Access 30
       MgmtClusterPublic    Access 40
```

Figure 6-11 VLAN assignments on management operating system NIC interfaces

8. Bandwidth limits must be placed on these network interfaces. The “ClusterPublic” virtual switch was created with “Absolute” weighting, which allows a maximum bandwidth cap to be placed on these network interfaces by using the following PowerShell commands (maximum bandwidth is defined in bits per second):

```
Set–VMNetworkAdpater –ManagementOS  -Name MgmtClusterPublic  -MaximumBandwidth
2GB

Set–VMNetworkAdpater –ManagementOS  -Name MgmtClusterPrivate  -MaximumBandwidth
1GB
```

   The network interface that is used for Hyper-V Live Migration uses the team that was created by using “Weighted” mode. A minimum bandwidth setting of 30% is set for the “LiveMigration” network, as shown in the following example:

```
Set–VMNetworkAdpater –ManagementOS  -Name MgmtLiveMigration
-MinimumBandwidthWeight  30
```

9. Assign TCP/IP addresses and confirm network connectivity for all of the network connections on each VLAN, as shown in the following example:

```
New-NetIPAddress –InterfaceAlias “vEthernet (MgmtClusterPublic)” –IPAddress
192.168.40.21 –PrefixLength 24

New-NetIPAddress –InterfaceAlias “vEthernet (MgmtClusterPrivate)” –IPAddress
192.168.30.21 –PrefixLength 24
```

```
New-NetIPAddress –InterfaceAlias “vEthernet (MgmtLiveMigration)” –IPAddress
192.168.31.21 –PrefixLength 24
```

10. Confirm the cluster public network (VLAN 40) is at the top of the network binding order.

    The Cluster Private and Live Migration networks must not have any default gateway defined.

6.5 Host storage connections

The following volumes are used to support the Management host cluster:

- Two 4 TB volumes to be used as a Cluster Shared Volumes
- 1 GB Volume to be used as the Management Cluster Quorum

After the switch FCoE configuration and the FAS8040 LUN mapping is complete, storage logical unit numbers (LUNs) should be visible to the host servers.

Complete the following steps:

1. Confirm that the disks are visible in Windows Disk Manager (a disk rescan might be required).
2. From one server, bring each disk online and format it as a GUID Partition Table (GPT) disk for use by the cluster. Assigning drive letters is optional because they are used for specific clustering roles, such as CSV, and Quorum.

   Validate that each potential host server can see the disks and bring the disks online.

Note: Only one server can have the disks online at a time until they are added to Cluster Shared Volumes.

6.6 Creating management host clusters

Microsoft Windows clustering is used to join the two host servers in a highly available configuration with which both servers can run VMs to support a production environment. VM workloads must be balanced across all hosts and careful attention must be paid to ensure that the combined resources of all VMs does not exceed those resources that are available on N-1 cluster nodes. Staying below this threshold allows any single server to be taken out of the cluster while minimizing the effect on your management servers.

A policy of monitoring resource utilization, such as CPU, Memory, and Disk (space and I/O), helps keep the cluster running at optimal levels and allows for proper planning for more resources as needed.

Complete the following steps:

1. Temporarily disable the default USB Remote NDIS Network Device on all cluster nodes. If the device is not disabled, the validation issues a warning during network detection because all nodes are sharing an IP address. These devices can be re-enabled after validation.

2. By using the Failover Cluster Manager on one of the two management nodes, run the Cluster Validation Wizard to assess the two physical host servers as potential cluster candidates and address any errors.

   The cluster validation wizard checks for available cluster compatible host servers, storage, and networking (see Figure 6-12).

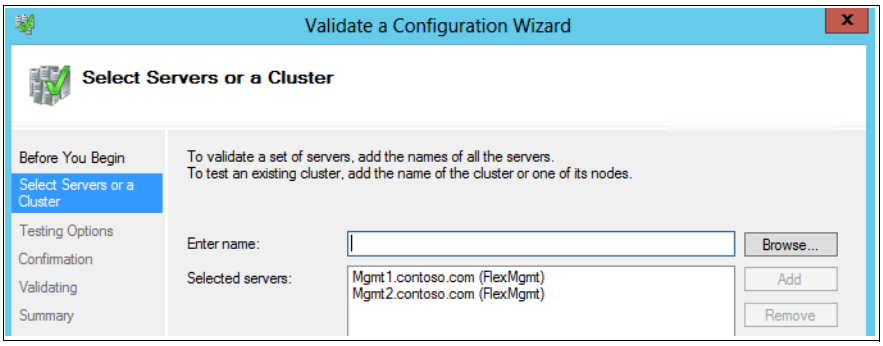

Figure 6-12 Cluster Validation Wizard

   Ensure that the intended cluster storage is online in only one of the cluster nodes.

   Address any issues that are flagged during the validation.

   As shown in Figure 6-13, use the Failover Cluster Manager to create a cluster with the two physical host servers that are to be used for the management cluster. Step through the cluster creation wizard. (you need a cluster name and IP address).

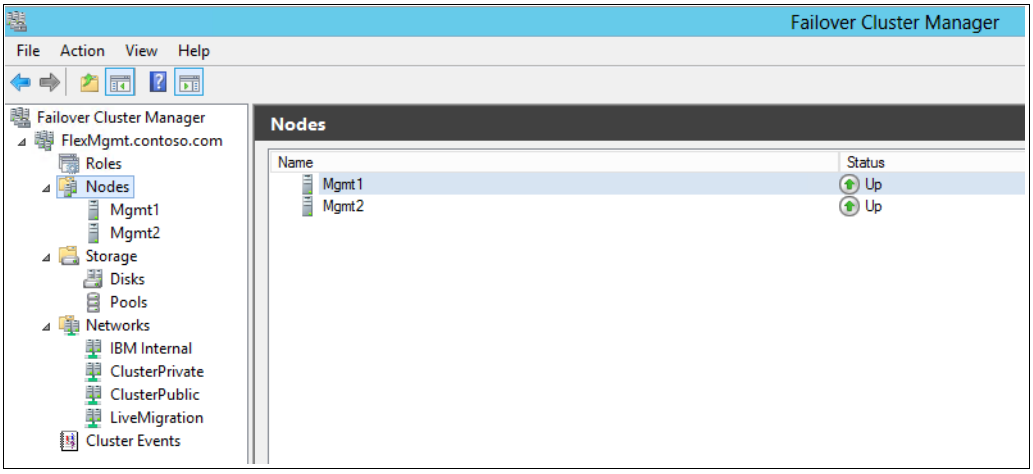

Figure 6-13 Management Nodes in Failover Cluster Manager

   By using Hyper-V Manager, set the default paths for VM creating to use the Cluster Shared Volumes.

3. Add the created storage volumes as a Cluster Shared Volumes (CSVs), as shown in Figure 6-14.

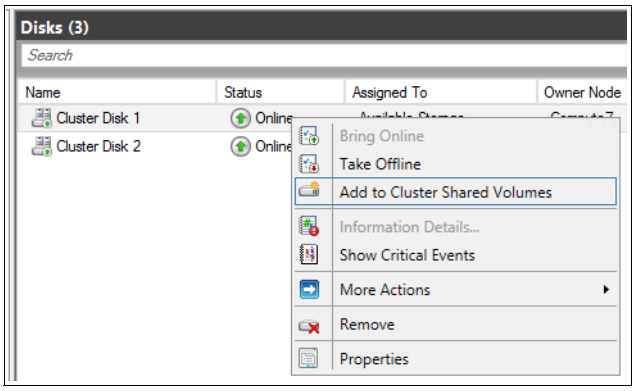

Figure 6-14 Changing cluster storage volume to a Cluster Shared Volume

Note: The cluster quorum volume should not be configured as a CSV.

4. By using Hyper-V Manager, set the default paths for VM creation to use the CSVs.
5. Configure the Live Migration settings to route this network traffic over the dedicated link, as shown in Figure 6-15.

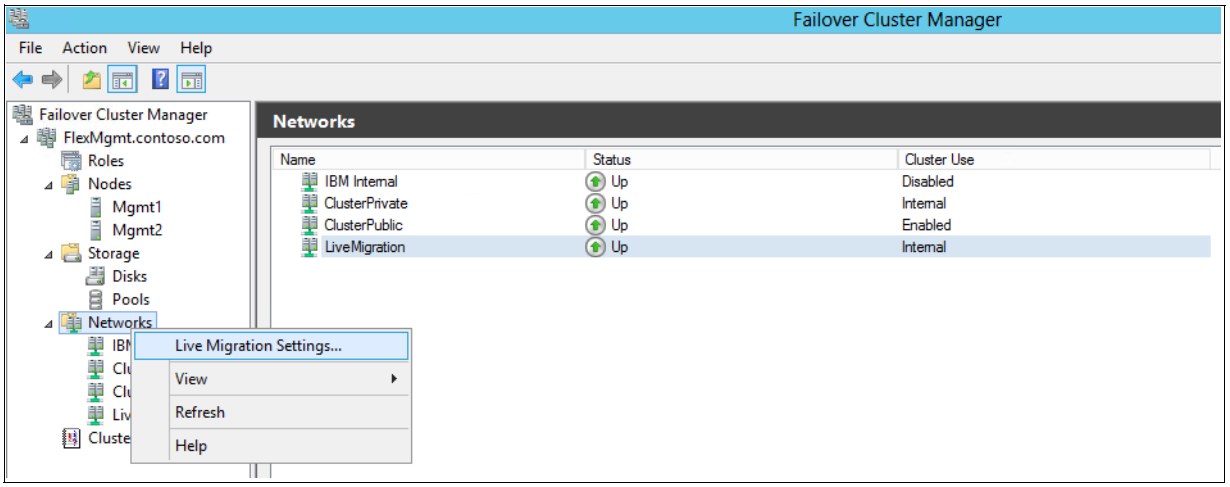

Figure 6-15 Changing Live Migration settings

6.7 VM Fibre Channel storage connections

Most of the storage that is used in a virtual environment consists of virtual hard disk drives in the form of VHDX files, which are used by Hyper-V. The VHDX files often are on storage that is provided by the Cluster Shared Volume and managed by the host server.

A new feature of Windows Server 2012 R2 is the ability to share VHDX files (hosted on a CSV) between VMs. Guest clustering between VMs is an example of why this option is used. In some cases, you must create direct Fibre Channel (FC) connections from the VM to the storage. These situations are increasingly uncommon; however, NPIV can be enabled to provide “virtual” HBA adapters in the VM if needed.

Complete the following to set up and configure direct FC connections to a VM:

1. Install the NetApp MPIO driver in the VM and then shut down the VM.
2. By using the Emulex OneCommand tool, enable NPIV on each of the Fibre Channel HBA ports that is used for storage access on the host operating system, as shown in Figure 6-16. The default setting for NPIV is “disabled.” Restart the host system.

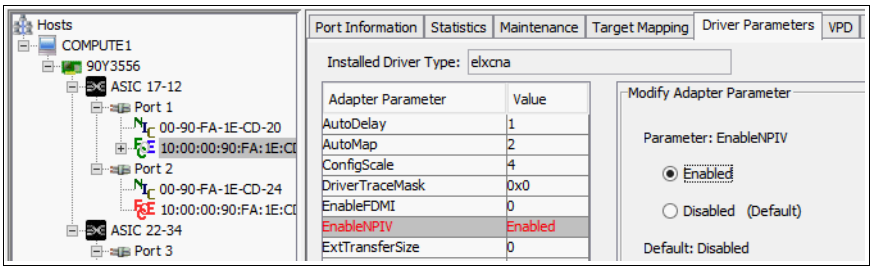

Figure 6-16 Enabling NPIV on FC ports

3. From the Hyper-V manager console, open the settings for the VM that has FC connectivity and add two vHBAs to the VM under “Add Hardware.” Associate each vHBA with one of the two vSANs, as shown in Figure 6-17.

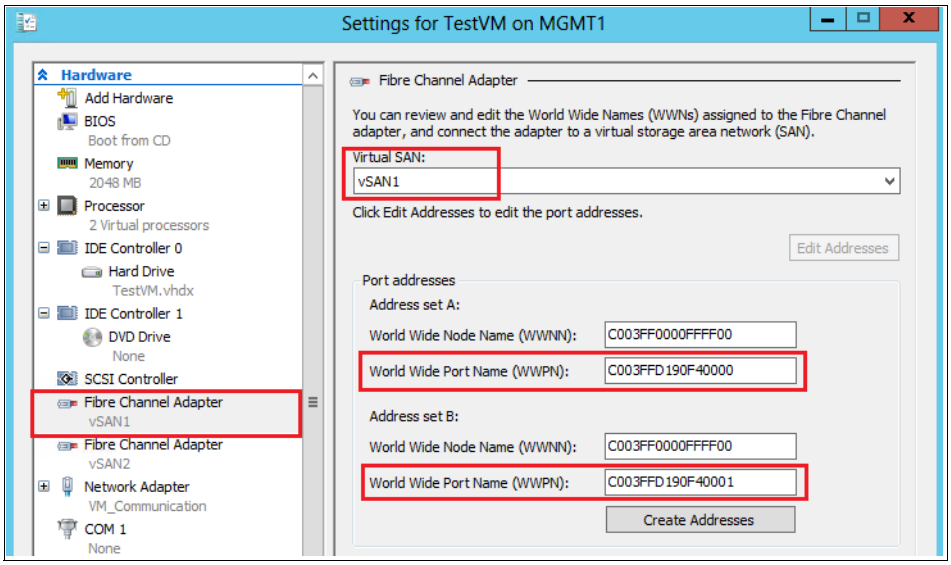

Figure 6-17 Adding vHBA to Hyper-V VM

4. By using Hyper-V Manager Virtual SAN Manager, create two virtual SAN switches with one HBA port in each virtual SAN, as shown in Figure 6-18. Previous worldwide port name (WWPN) information can be reviewed to determine proper port-to-switch correlation.

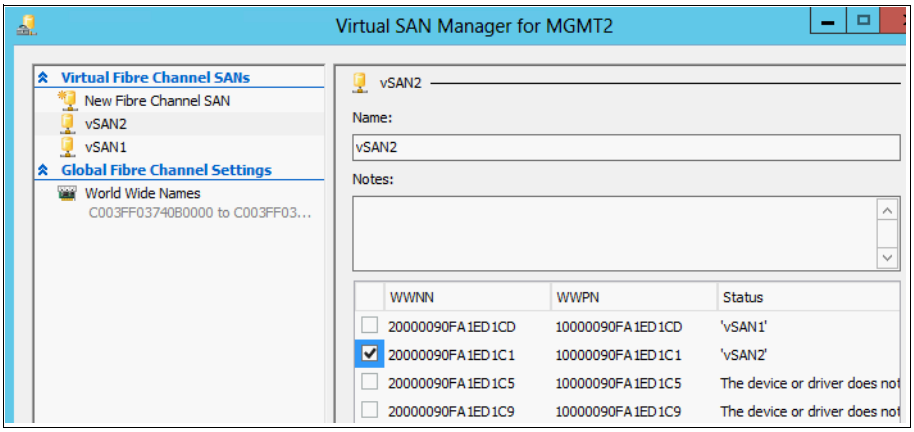

Figure 6-18 Creating vSAN with Hyper-V Manager

5. Record the WWPN that is provided to the new vHBA under “Address Set A.” If this VM is Live Migrated, also record the WWPN that is under “Address Set B.”

   Hyper-V Live Migration alternates the WWPN port names between “Address Set A” and “Address Set B” during migration between host servers.

6. Start the VM.

7. The new active WWPN from the VM now must be seen by each Flex System CN4093 switch. A new fcalias must be created for this new WWPN on each switch. The second WWPN does not become visible until it becomes active; however, it should be added to the fcalias list, zone, and zone set now by using the following commands:

```
CN 4093(config)# fcalias vmSQL1_vSAN1_Addr_A wwn  c0:03:ff:d1:90:f4:00:00
CN 4093(config)# fcalias vmSQL1_vSAN1_Addr_B wwn  c0:03:ff:d1:90:f4:00:01
CN 4093(config)# fcalias vmSQL2_vSAN1_Addr_A wwn  c0:03:ff:d1:90:f4:00:00
CN 4093(config)# fcalias vmSQL2_vSAN1_Addr_B wwn  c0:03:ff:d1:90:f4:00:01
CN 4093(config)# write
```

   A zone must be created on each switch that contains the fcaliases of the two storage ports that are presented by the NetApp storage VM and new VM WWPNs by using the following commands:

```
CN 4093(config-zone)# zone name SW1Zone_SQL_Cluster
CN 4093(config-zone)# member fcalias vmSQL1_vSAN1_Addr_A
CN 4093(config-zone)# member fcalias vmSQL1_vSAN1_Addr_B
CN 4093(config-zone)# member fcalias vmSQL2_vSAN1_Addr_A
CN 4093(config-zone)# member fcalias vmSQL2_vSAN1_Addr_B
CN 4093(config-zone)# member fcalias NTAP_SVM1_Port1
CN 4093(config-zone)# member  fcalias NTAP_SVM1_Port2
CN 4093(config-zone)# exit
CN 4093(config)# show zone
CN 4093(config)# write
```

8. Add the zone (or zones) that was created to the current zone set and activate by using the following commands:

```
CN 4093(config)# zoneset name SW1_ZoneSet
CN 4093(config-zoneset)# member SW1Zone_SQL_Cluster
```

```
CN 4093(config)# exit
CN 4093(config)# show zoneset
CN 4093(config)# zoneset activate name SW1_ZoneSet
CN 4093(config)# write
```

9. Repeat the process for the second vHBA (vSAN2) on the other CN4093 switch.

Note: SAN connectivity must be configured for four of the management fabric virtual machines (`vmSQL1`, `vmSQL2`, `vmVMM1`, and `vmVMM2`) that is using “Guest Clustering”.

Consider the following points:

- The FAS8040 storage now must see the new WWPNs and can be mapped to a new host with storage assigned as needed.
- The alternative ‘B’ addresses must be included on the host mapping. The addresses show as unverified until they are brought online.
- LUNs can then be assigned to the host VM in the FAS8040 storage (which is shown with System Manager, see Figure 6-7 on page 71).
- The new LUNs must be visible in Windows Disk Manager. A disk rescan might be required.

6.8 Setting up and configuring VMs

Setting up the management fabric VMs uses several methods of Windows clustering to create a robust and fault-tolerant management environment.

The operating system can be installed on a VM by using various methods. One approach is to modify the VM DVD drive settings to specify an image file that points to the Windows installation ISO image, then start the VM to begin the installation. Other deployment methods, such as the use of a VHD file with a Sysprep’d image, WDS server, or SCCM, also can be used.

After the operating system is installed and while the VM is running, the following steps must be complete before the application software is installed:

1. Run Windows Update.
2. Install the integration services on the VM. Although the latest Windows Server builds include integration services built-in, it is important to make sure the Hyper-V child and parent run the same version of integration components.
3. Windows Server that is running a VM must automatically activate if the host and VM are running Windows Server 2012 R2.

Note: Hyper-V supports Dynamic Memory in VMs, which allows some flexibility in the assignment of memory resources to VMs. However, some applications might experience performance-related issues if the VM’s memory settings are not configured correctly. It is suggested that the management fabric uses static memory assignments to ensure optimal performance.

7

Setting up and configuring System Center 2012 R2

The procedures for setting up and configuring System Center 2012 R2 are provided in this chapter, which included the following topics:

7.1 Microsoft System Center 2012 R2

Microsoft System Center 2012 R2 with Lenovo XClarity Upward Integration Modules (UIM) provides an environment to configure, monitor, and manage a Windows Server with Hyper-V based configuration.

System Center 2012 R2 is a collection of several server applications that are integrated to provide a management fabric. Although the components have scale points that are well within the present configuration, they are listed in Table 7-1 to provide a reference point as your virtual environment scales up.

Table 7-1 Scale points for System Center 2012 components

Component	Scalability Reference	Notes
Virtual Machine Manager	A total of 800 hosts/25,000 virtual machine (VMs) per instance.	A Virtual Machine Manager instance is defined as a stand-alone or cluster installation. Although not required, scalability is limited to 5,000 VMs when Service Provider Foundation (SPF) is installed. A single SPF installation can support up to five Virtual Machine Manager instances.
App Controller	Scalability is proportional to Virtual Machine Manager.	Supports 250 VMs per Virtual Machine Manager user role.
Operations Manager	A total of 3,000 agents per management server, 15,000 agents per management group, 50,000 agentless managed devices per management group.	N/A
Orchestrator	Simultaneous execution of 50 runbooks per Runbook server.	N/A
Service Manager	Large deployment supports up to 20,000 computers.	Topology dependent. In Fast Track, Service Manager is used solely for VM management. An advanced deployment topology can support up to 50,000 computers.

The Fast Track fabric management architecture uses System Center 2012 R2 Datacenter edition. For more information about System Center 2012, see *What's New in System Center 2012 R2*, which is available at this website:

http://technet.microsoft.com/en-us/library/dn249519.aspx

7.2 SQL Server 2012 Setup and Configuration

Microsoft SQL Server 2012 performs the data repository functions for all the key components of System Center management fabric. Each System Center component stores its data on its own clustered SQL Server database instance.

Following the Microsoft Hyper-V Private Cloud Fast Track best practices, our solution consists of two SQL Server 2012 virtual machines (VMs) that are configured as a failover cluster with guest clustering. There are seven instances of SQL installed on each of these two machines with each instance supporting a specific role in the management fabric. Figure 7-1 shows the high-level architecture of a SQL Server guest failover cluster.

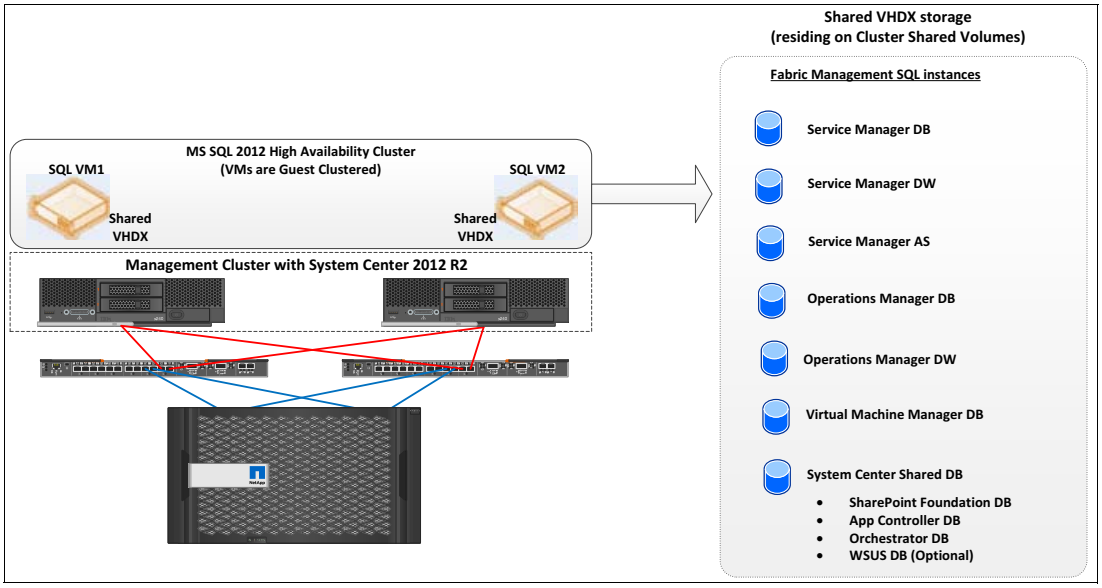

Figure 7-1 Microsoft SQL Server 2012 Cluster

These VMs feature the following base configuration:

- Eight vCPU
- 16 GB RAM
- Two virtual network interfaces:
  - One interface that is configured for the Host Management network (ClusterPublic/VLAN40) that uses VM_Communications vSwitch to connect to Lenovo vNIC™
  - One interface that is configured for SQL Cluster Private (VLAN33) that uses VM_Communications vSwitch to connect to vNIC
- One 60 GB virtual hard disk for operating system installation (VHDX format)
- One 1 GB Shared VHDX Disk (for Guest Cluster Quorum)
- A total of 14 Shared VHDX Disks (for Database Instances)
- Windows Server 2012 R2 Datacenter edition that is installed and updated
- High availability that is achieved by using guest clustering between VMs

7.2.1 SQL clustered instances

A total of seven Clustered SQL instances are created to service the management fabric. A high-level overview of each instance is listed in Table 7-2 on page 84.

Table 7-2 Database Instance Requirements

Fabric management component	Instance name (Suggested)	Components	Collation[a]	Storage volumes[b]
Virtual Machine Manager	SCVMMDB	Database Engine	SQL_Latin1_General_CP1_CI_AS	2 LUNs
Windows Server Update Services (optional)	SCVMMDB	Database Engine	SQL_Latin1_General_CP1_CI_AS	N/A – Shared instance with Virtual Machine Manager
Operations Manager	SCOMDB	Database Engine, Full-Text Search	SQL_Latin1_General_CP1_CI_AS	2 LUNs
Operations Manager Data Warehouse	SCOMDW	Database Engine, Full-Text Search	SQL_Latin1_General_CP1_CI_AS	2 LUNs
Service Manager	SCSMDB	Database Engine, Full-Text Search	Latin1_General_100_CI_AS	2 LUNs
Service Manager Data Warehouse	SCSMDW	Database Engine, Full-Text Search	Latin1_General_100_CI_AS	2 LUNs
	SCSMAS	Analysis Services	Latin1_General_100_CI_AS	2 LUNs
Service Manager Web Parts and Portal	SCDB	Database Engine	SQL_Latin1_General_CP1_CI_AS	N/A – Shared instance with Orchestrator and App Controller
Orchestrator	SCDB	Database Engine	SQL_Latin1_General_CP1_CI_AS	2 LUNs
App Controller	SCDB	Database Engine	SQL_Latin1_General_CP1_CI_AS	N/A – Shared instance with Orchestrator and Service Manager Portal

a. The default SQL collation settings are not supported for multi-lingual installations of the Service Manager component. Use only the default SQL collation if multiple languages are not required. The same collation must be used for all Service Manager databases (management, DW, and reporting services).

b. Extra LUNs might be required for TempDB management in larger scale configurations.

7.2.2 SQL cluster storage

A total of 15 storage volumes (VHDXs) must be created and mapped to the two SQL Server VMs on the Cluster Shared Volumes. These shared VHDX files must be evenly distributed across the two CSVs to better balance the I/O: one for the cluster quorum and 14 for the database instances (two per instance).

A VHDX can be marked as shared by examining the advance properties when it is attached to a VM. The SQL VMs must be created and configured for networking and storage before the SQL cluster is created. Table 7-3 lists suggested SQL Server instance names, disk sizes, assignments, and IP addresses.

Table 7-3 Database Instance guidance

Component	Service manager manage-ment server	Service manager data ware-house server	Service manager analysis server	App controller, orchestrator, Microsoft SharePoint services farm and WSUS	VM manager	Operations manager	Operations manager data warehouse
SQL Server Instance Name	SCSMDB	SCSMDW	SCSMAS	SCDB	SCVMMDB	SCOMDB	SCOMDW
SQL Server Instance Failover Cluster Network Name	SCSMDB	SCSMDW	SCSMAS	SCDB	SCVMMDB	SCOMDB	SCOMDW
SQL Server Instance DATA Cluster Disk Resource	Cluster Disk2 (145 GB)	Cluster Disk4 (1 TB)	Cluster Disk6 (8 GB)	Cluster Disk8 (10 GB)	Cluster Disk 10 (6 GB)	Cluster Disk 12 (130 GB)	Cluster Disk 14 (1 TB)
SQL Server Instance LOG Cluster Disk Resource	Cluster Disk3 (70 GB)	Cluster Disk5 (500 GB)	Cluster Disk7 (4 GB)	Cluster Disk9 (5 GB)	Cluster Disk 11 (3 GB)	Cluster Disk 13 (65 GB)	Cluster Disk 15 (500 GB)
SQL Server Instance Install Drive	E:	G:	I:	K:	M:	O:	Q:
SQL Server Instance DATA Drive	E:	G:	I:	K:	M:	O:	Q:
SQL Server Instance LOG Drive	F:	H:	J:	L:	N:	P:	R:
SQL Server Instance TEMPDB Drive	F:	H:	J:	L:	N:	P:	R:
Cluster Service Name	SQL Server (SCSMDB)	SQL Server (SCSMDW)	SQL Server (SCSMAS)	SQL Server (SCDB)	SQL Server (SCVMMDB)	SQL Server (SCOMDB)	SQL Server (SCOMDW)
Clustered SQL Server Instance IP address	192.168.40 .91	192.168.40. 92	192.168.4 0.93	192.168.40.94	192.168.40. 95	192.168.40. 96	192.168.40. 97
Host Cluster Public Network Interface Subnet Mask	255.255.25 5.0	255.255.25 5.0	255.255.2 55.0	255.255.255.0	255.255.25 5.0	255.255.255 .0	255.255.25 5.0
Host vSwitch for Network Interface Name	VM_Comm	VM_Comm	VM_Comm	VM_Comm	VM_Comm	VM_Comm	VM_Comm
SQL Server Instance Listening TCP/IP Port	10437	10438	10439	1433[a]	10434	10435	10436

Component	Service manager manage-ment server	Service manager data ware-house server	Service manager analysis server	App controller, orchestrator, Microsoft SharePoint services farm and WSUS	VM manager	Operations manager	Operations manager data warehouse
SQL Server Instance Preferred Owners	Node2, Node1	Node2, Node1	Node2, Node1	Node1, Node2	Node1, Node2	Node1, Node2	Node1, Node2

a. The SCDB instance must be configured to port 1433 if the Cloud Services Process Pack is used.

7.2.3 SQL server guest clustering

The two VMs that are hosting SQL Server 2012 use guest clustering to create a highly available fault tolerant configuration.

Complete the following steps to set up a 2-node VM guest cluster:

1. Validate that the SQL Quorum VHDX can be attached and is visible to both SQL nodes.
2. Validate the SQL data volumes can be attached and are visible to both SQL nodes.
3. Map the storage volumes to the two VMs and validate that the volume name and drive letter are consistent across both SQL nodes.
4. Validate the network connections.
5. Validate and create the Windows cluster with the two SQL nodes.

The SQL nodes can also be highly available on the management host cluster, but must be assigned a preferred node as shown in these steps. Both SQL nodes can be migrated to one management host (if needed) to support maintenance operations.

For more information about installation instructions for the SQL Server 2012 configuration, see the Microsoft Infrastructure as a Service (IaaS) Deployment Guide, which is available at this website:

http://gallery.technet.microsoft.com/Infrastructure-as-a-ecf1cc0b

7.3 Setting up and configuring the file server

It is recommended to store all necessary files that are needed for host setup and configuration on a file server. The NetApp FAS8040 can provide a dedicated CIFS/SMB file share for this purpose. CIFS license is required and must be added to support this feature.

A highly available VM that us running as a Windows File Server is needed to present the FAS8040 NAS share to SCVMM.

The File Server VM must feature the following base VM configuration:

- One vCPU
- 2 GB RAM
- One virtual network interface that is configured for the Host Management network (ClusterPublic/VLAN40), which uses VM_Communications vSwitch to connect to vNIC
- One 50 GB virtual hard disk for operating system installation (VHDX format)

- One 50 GB virtual hard disk for data volume (VHDX format)
- Windows Server 2012 R2 Datacenter edition that is installed and updated
- Highly available through host clustering

An initial share that is created on the Windows File Server is needed to setup SCVMM. After SCVMM is configured, another Virtual Machine Manager (VMM) Library share can be added from within SCVMM to this server as an unmanaged share. This file server then redirects I/O to the FAS8040 share.

To integrate NetApp shares for use by SCVMM as a VMM library, the FAS8040 shares must be created under a qtree.

7.4 NetApp FAS8040 SMI-S Provider for storage automation

The NetApp Storage Management Initiative Specification (SMI-S) Agent easily can be run within a VM because only IP connectivity to the SCVMM/SCOM server and NetApp storage controllers is needed.

Note: Installing SMI-S Agent on the server in which SCVMM is installed is not supported.

The SMI-S Provider VM must feature the following base VM configuration:

- Two vCPU
- 4 GB RAM
- One virtual network interface that is configured for the Host Management network (ClusterPublic/VLAN40), which uses VM_Communications vSwitch to connect to vNIC
- One 50 GB virtual hard disk for OS installation (VHDX format)
- Windows Server 2012 R2 Datacenter edition that is installed and updated.
- Highly available through Windows Failover host clustering.

The NetApp SMI-S provider can be downloaded from the following NetApp SMI-S 5.1 Provider website:

http://mysupport.netapp.com/NOW/cgi-bin/software/?product=Data+ONTAP+SMI-S+Agent&platform=Windows

Installing SMI-S Provider

Complete the following steps to install SMI-S Provider:

1. Right-click **smisagent-5-1** and select **Install**.
2. On the Welcome page of the Data ONTAP SMI-S Agent Setup wizard, click **Next**.
3. On the Ready to Install Data ONTAP SMI-S Agent page, click **Install**.
4. On the Completed the Data ONTAP SMI-S Agent Setup Wizard page, click **Finish** to complete the installation.

Configuring SMI-S Provider

To configure the NetApp SMI-S provider, complete the following steps:

1. Open the application window, right-click **Data ONTAP SMI-S Agent**, and select **Run as Administrator**.

2. Change the directory into the SMI-S program files (`%ProgramFiles(x86)%\ONTAP\smis\pegasus\bin`).
3. Add the Storage Virtual Machine (SVM) SVM credentials to the SMI-S Agent (as shown in Figure 7-2) by selecting **Smis addsecure → VserverIpAddress> → <VserverAdmin> → <VserverAdminPassword>**.

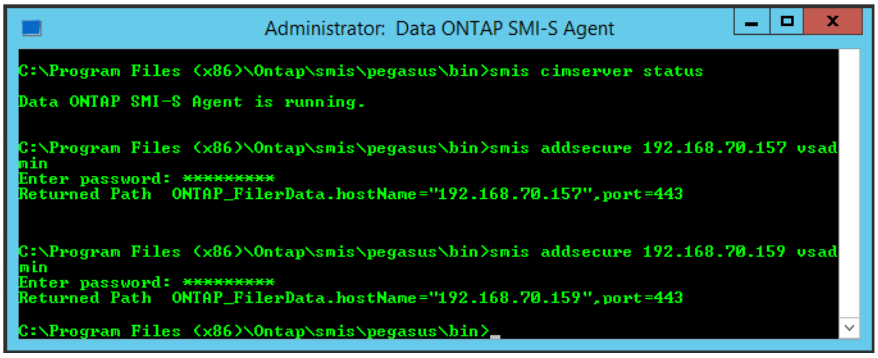

Figure 7-2 Adding SVMs to SMI-S provider

4. Enable user authentication by using the following command:

   `Cimconfig -p -s enableAuthentication=true`
5. Restart the agent/cimserver by using the following command:

   `smis cimserver restart`
6. Add SMI-S Run As account to the SMIS configuration by using the following command:

   `cimuser -a -u FT-SMIS-User -w <password>`

For more information about the NetApp SMI-S provider, see Appendix A, “Appendices” on page 113.

7.5 Setting up and configuring System Center Virtual Machine Manager 2012 SP1

The Systems Center 2012 Virtual Machine Manager component of this configuration consists of two VMs that are using guest clustering and running Microsoft System Center 2012 SP1 VMM to provide virtualization management services to this environment. Figure 7-3 shows the high-level architecture of a VMM high-availability cluster.

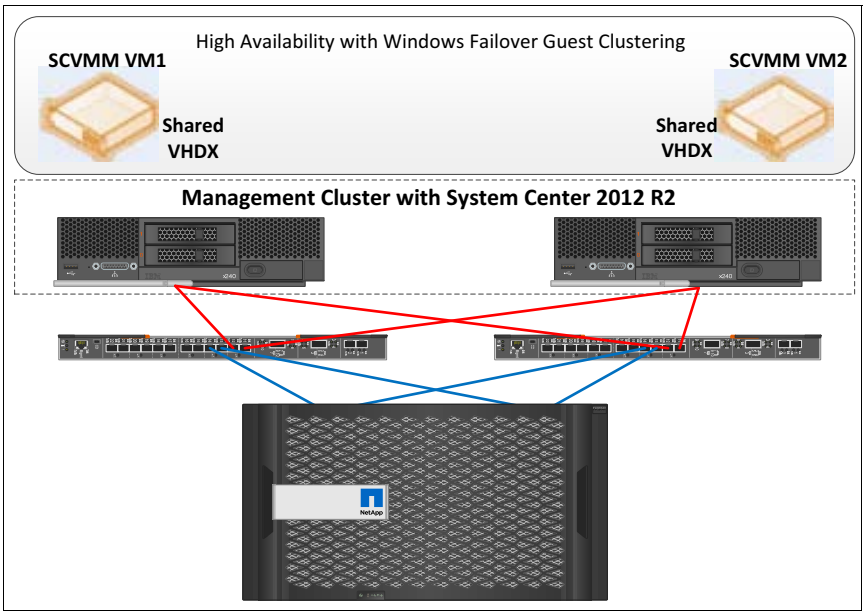

Figure 7-3 VMM high-availability cluster

These VMs must feature the following VM base configuration:

- Four vCPU
- 8 GB RAM
- Two virtual network interfaces:
  - One instance that is configured for the Host Management network (`ClusterPublic/VLAN40`), which uses VM_Communications vSwitch to connect to vNIC
  - One instance that is configured for SCVMM Cluster Private, which uses VM_Communications vSwitch to connect to vNIC
- One 60 GB virtual hard disk for OS and application installation (VHDX format)
- One 1 GB Shared VHDX Disk (for Guest Cluster Quorum)
- Access to a Windows File Server as the location of VMM Library
- Windows Server 2012 R2 Datacenter edition that is installed and updated
- High availability that is achieved by using guest clustering between VMs

Note: A dedicated SQL Server 2012 instance that is running on a separate, configured SQL Server cluster also is required to support this workload.

Note: The Microsoft IaaS Deployment Guide recommends installing the Virtual Machine Manager Library as a Clustered file resource by using the two SQL servers. For more information, see this website:

http://gallery.technet.microsoft.com/Infrastructure-as-a-ecf1cc0b

7.5.1 SCVMM guest clustering for Virtual Machine Manager

The two VMs that are hosting SCVMM use guest clustering to create a highly available fault tolerant configuration. This configuration requires a shared VHDX disk to be used as the cluster quorum.

Complete the following steps to set up a 2-node VM guest cluster:

1. Validate that the Quorum VHDX can be attached and is visible to both SCVMM nodes.
2. Validate network connections.
3. Validate and create the Windows cluster.

For more information about installation instructions for the System Center Virtual Machine Manager 2012 R2 configuration, see *Microsoft Infrastructure as a Service (IaaS) Deployment Guide*, which is available at this website:

http://gallery.technet.microsoft.com/Infrastructure-as-a-ecf1cc0b

7.5.2 Pro Pack for Microsoft System Center Virtual Machine Manager

The Lenovo Hardware Pro Pack Modules for System Center Virtual Machine Manager provides performance, health alerting, and VM migration recommendations for the Flex Systems that are used in this configuration. This package is installed on the System Center Operations Manager by using the instructions that are featured in Lenovo Hardware Performance and Resource Optimization Pack for Microsoft System Virtual Machine Manager, which is available at this website:

http://ibm.com/support/entry/portal/docdisplay?lndocid=LNVO-MANAGE

7.5.3 Setting up and configuring NetApp FAS8040 SMI-S for Microsoft System Center Virtual Machine Manager

The NetApp FAS and cDOT supports storage automation through the implementation of the SMI-S. Core storage management tasks, such as discovery, classify, provisioning and de-provisioning, and mapping can be completed from the System Center Virtual Machine Manager 2012 R2 console.

Complete the following steps to set up and configure NetApp FAS8040 SMI-S for Microsoft System Center Virtual Machine Manager:

1. Select the **Fabric** panel at the bottom, left side of the SCVMM console and choose **SAN and NAS devices discovered and managed by a SMI-S provider**, as shown in Figure 7-4.

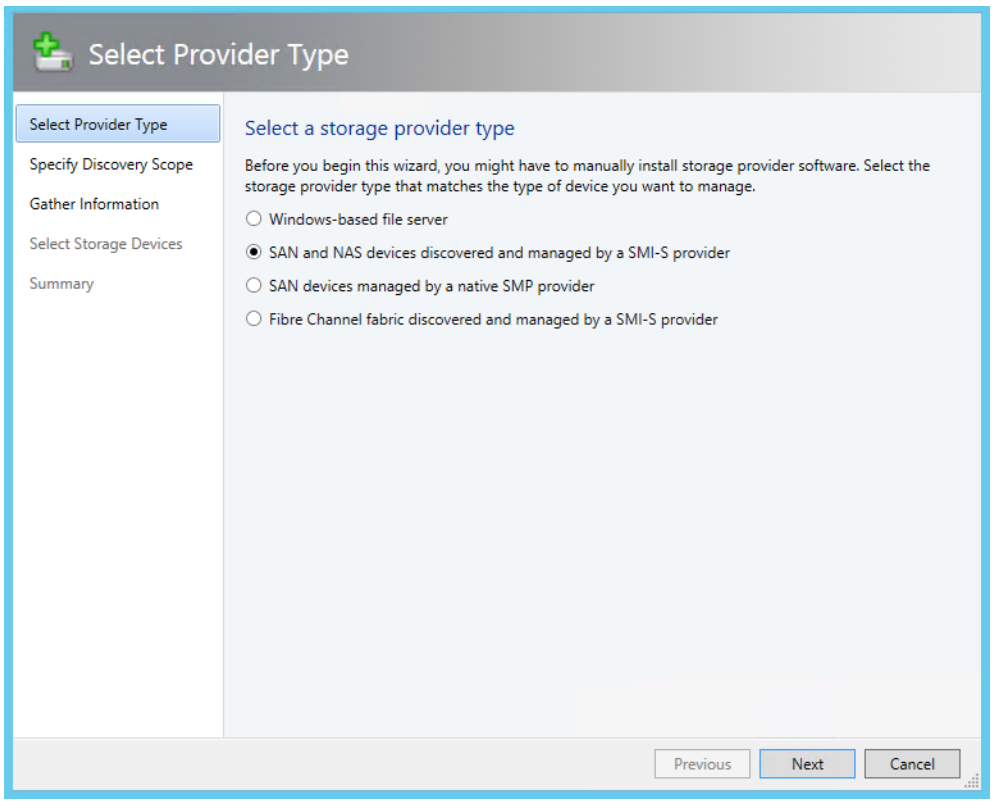

Figure 7-4 SCVMM SMI-S provider setup

2. Click **Specify Discovery Scope** on the next step of the SCVMM Provider setup wizard. Ensure that SSL is selected for communication with the FAS8040, as shown in Figure 7-5 on page 92.

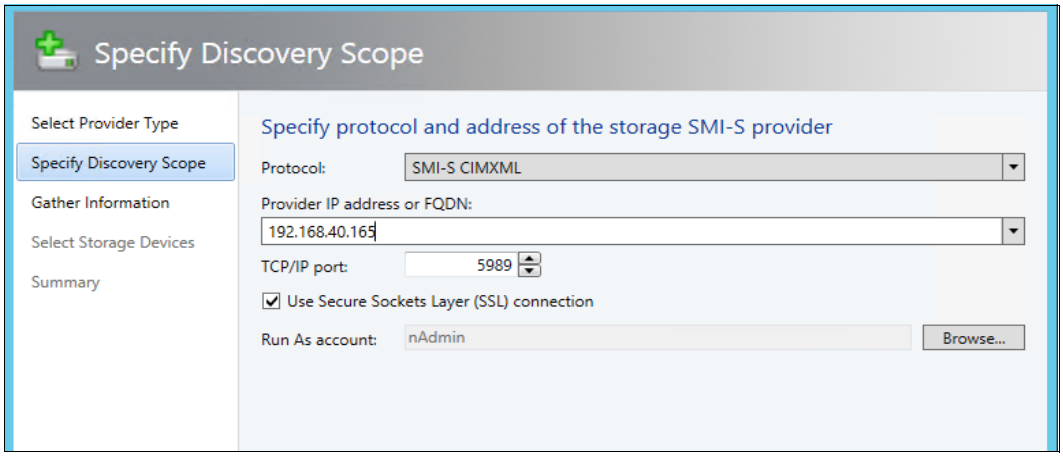

Figure 7-5 SCVMM Provider Discovery Scope

SCVMM must query the SMI-S provider and discover the SVMs that are running on the FAS8040, as shown in Figure 7-6.

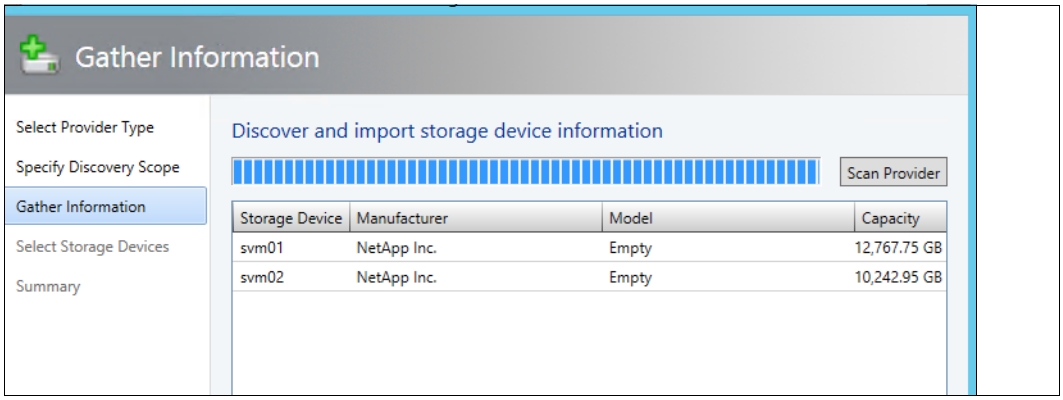

Figure 7-6 SCVMM discovery for FAS8040 Storage

3. Select the FAS8040 volumes to manage with SMI-S, classify as needed, and click **Finish** (see Figure 7-7).

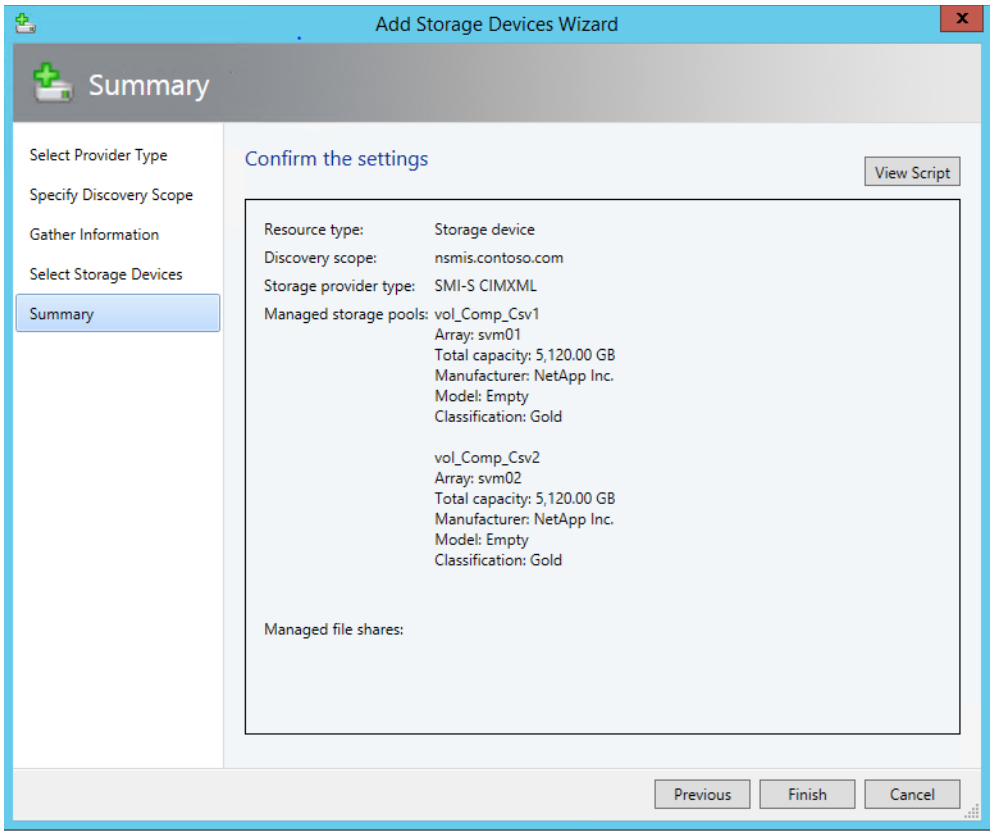

Figure 7-7 NetApp FAS8040 Managed Volumes

7.5.4 Bare metal provisioning

Bare metal provisioning compute nodes can be accomplished by using several methods. Windows Deployment Server (which is included in Windows Server) provides a straightforward, bare metal deployment framework. A more comprehensive deployment solution is found in Microsoft System Center Virtual Machine Manager (SCVMM) 2012 R2.

7.6 Setting up and configuring System Center Operations Manager 2012 R2

The System Center Operations Manager 2012 R2 component of this configuration consists of two VMs that use native application failover and redundancy features and operating as a single management group. A third VM is configured as highly available under the host cluster and used as the Operations Manager Reporting Server. Figure 7-8 shows the high-level architecture of the Operations Manager components in the management cluster.

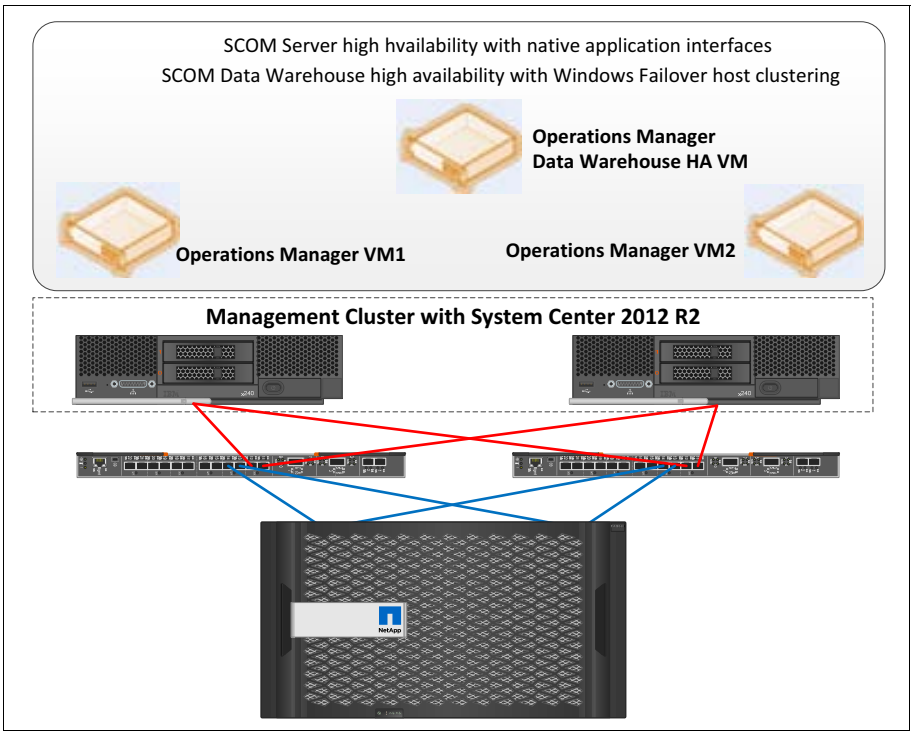

Figure 7-8 System Center Operations Manager VMs

These VMs must feature the following base VM configuration:

- Eight vCPU
- 16 GB RAM
- One virtual network interface that is configured for the Host Management network (ClusterPublic/VLAN40), which uses VM_Communications vSwitch to connect to vNIC

  Routing must be provided between the Out Of Band (OOB) Network and the Host Management Network (MgmtClusterPublic) for Operations Manager to communicate to the Flex CMM.
- One 60 GB virtual hard disk for operating system and application installation (VHDX format).
- Windows Server 2012 R2 Datacenter edition installed and updated.
- High availability is achieved through native application interfaces for the SCOM servers.
- High availability is achieved through host clustering for the SCOM Data Warehouse.

Note: Two dedicated SQL Server 2012 instances that are running on the separate previously configured SQL Server cluster is required to support these workloads.

- One instance for the Operations Manager Database
- One instance for the Operations Manager Data Warehouse

For more information about installation instructions for the System Center Operations Manager 2012 R2 configuration, see *Microsoft Infrastructure as a Service (IaaS) Fabric Management Deployment Guide*, which is available at this website:

http://gallery.technet.microsoft.com/Infrastructure-as-a-ecf1cc0b

7.6.1 Lenovo XClarity Integrator offerings for Microsoft System Center

After the System Center Operations Manager is set up and configured, the Lenovo XClarity Integrator package can be installed and configured.

The Lenovo XClarity Integrator expands the server management capability by integrating Lenovo hardware management functionality and reduces the time and effort that is needed for routine system administration. It also provides discovery, deployment, configuration, monitoring, event, and power management. The software is available from this website:

https://ibm.com/support/entry/portal/docdisplay?lndocid=LNVO-MANAGE

Complete the following steps:

1. Install the SCOM package on the Operations Management Server.
2. Install the optional Power CIM provider.
3. Install and activate the License tool on the SCOM server.
4. Discover and manage the Flex System x240 M5 servers.
5. Create the connection to Flex Chassis through the Flex System CMM by using SNMP V3.

7.6.2 Pro Pack for Microsoft System Center Virtual Machine Manager

Complete the following tasks to set up Lenovo Pro Pack for SCVMM:

- Verify that the SCOM and SCVMM servers are integrated.
- Install the Lenovo Hardware Performance and Resource Optimization Pack for Microsoft System Center Virtual Machine Manager on the Operations Management Server, which is available from this website:

  https://ibm.com/support/entry/portal/docdisplay?lndocid=LNVO-PROPACK

7.6.3 NetApp Storage Management Monitoring with the OnCommand Module for Microsoft System Center Operations Manager

The NetApp OnCommand plug-in for Microsoft integrates with System Center Operations Manager to monitor storage health and provide alerts.

The NetApp OnCommand Plug-In for Microsoft 4.0.1 x64 is available for download at this website:

http://mysupport.netapp.com/NOW/cgi-bin/software/?product=OnCommand+Plug-in+for+Microsoft&platform=Microsoft+SCOM

Complete the following steps to install the SCOM portion of the OnCommand module on the SCOM server:

1. Install the SCOM Management Packs and SCOM Console Integration.
2. Enter the required AD credentials or SCOM RunAs account.

3. Enter the SQL Server that can be used for the OCPM database. Use the SCOM database instance with appropriate AD credentials.
4. Under the Clustered Data ONTAP folder, select management servers, and configure storage.
5. Tune SCOM as needed for your environment.

For more information, see *NetApp OnCommand 4.0 for Microsoft Best Practices*, which is available at this website:

http://www.netapp.com/us/media/tr-4244.pdf

7.7 Setting up and configuring System Center Orchestrator 2012 R2

Microsoft System Center Orchestrator 2012 is a workflow management solution for the datacenter. Orchestrator enables the automation of resources creation, monitoring, and deployment in your environment.

System Center Orchestrator 2012 is configured as a highly available platform by clustering two VMs with native applications services that are built into the product. Figure 7-9 shows the high-level architecture of the Orchestrator Native Application high availability that is running on the Management Cluster.

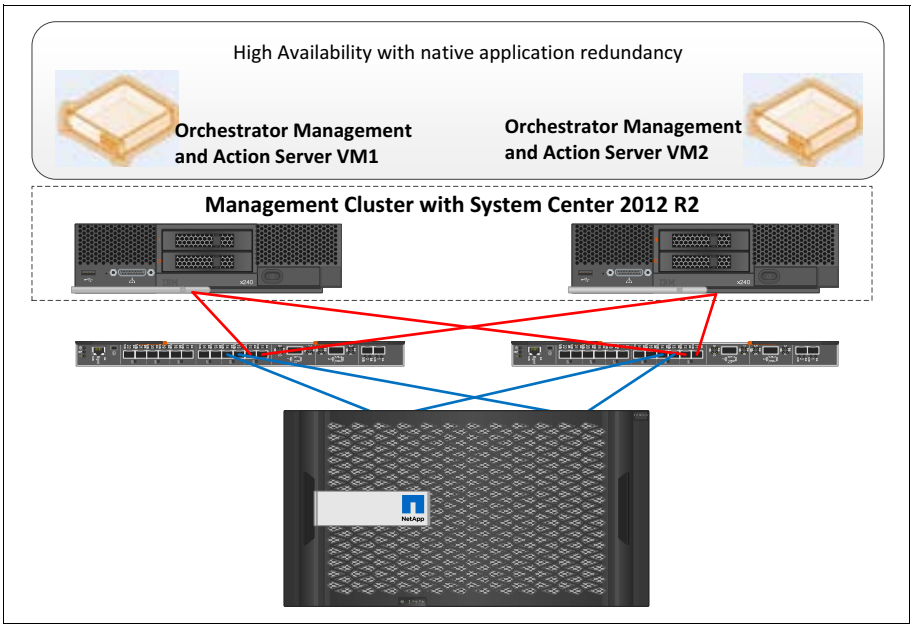

Figure 7-9 System Center Orchestrator highly available Cluster

These VMs must feature the following base VM configuration:

- Four vCPU
- 8 GB RAM
- One virtual network interface that is configured for the Host Management network (ClusterPublic/VLAN40), which uses VM_Communications vSwitch to connect to vNIC
- One 60 GB virtual hard disk for operating system and application installation (VHDX format)
- Windows Server 2012 R2 Datacenter edition installed and updated
- Native application interface used for high availability

Note: Consider the following points:

- Do not use host clustering.
- A dedicated SQL Server 2012 instance that is running on a separate, configured SQL Server cluster is required to support this workload.

For more information about installation instructions for the System Center Orchestrator 2012 R2 configuration, see *Microsoft Infrastructure as a Service (IaaS) Fabric Management Deployment Guide*, which is available at this website:

`http://gallery.technet.microsoft.com/Infrastructure-as-a-ecf1cc0b`

7.8 Setting up and configuring System Center Service Manager 2012 R2

Microsoft System Center Service Manager 2012 provides an integrated management platform with a robust set of capabilities that provides built-in processes for incident and problem resolution, change control, and asset lifecycle management.

The Microsoft Hyper-V Private Cloud Fast Track solution consists of four Service Manager components. Each component runs on its own VM instance, as shown in Figure 7-10 on page 98. High availability for each of these VMs is achieved through Windows Failover Clustering on the hosts with the ability to rapidly live migrate or failover the VM between hosts.

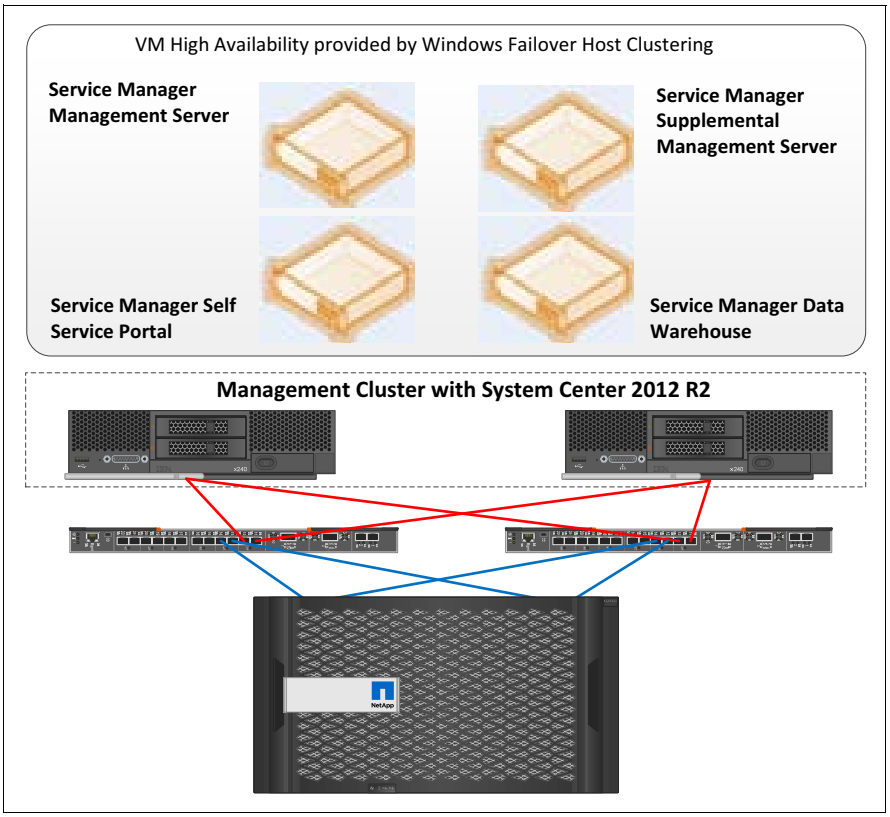

Figure 7-10 Service Manager architecture

The two Service Manager VMs must feature the following base VM configuration:

- Four vCPU
- 16 GB RAM
- One virtual network interface that is configured for the Host Management network (ClusterPublic/VLAN40), which uses VM_Communications vSwitch to connect to vNIC
- One 60 GB virtual hard disk for operating system and application installation (VHDX format)
- Windows Server 2012 Datacenter edition that is installed and updated
- Highly available through Windows Failover host clustering

The Service Manager Data Warehouse VMs must feature the following base VM configuration:

- Eight vCPU
- 16 GB RAM
- One virtual network interface that is configured for the Host Management network (ClusterPublic/VLAN40), which uses VM_Communications vSwitch to connect to vNIC
- One 60 GB virtual hard disk for operating system and application installation (VHDX format)
- Windows Server 2012 Datacenter edition that is installed and updated
- Highly available through Windows Failover host clustering

The two Service Manager Self Service Portal VM must feature the following VM base configuration:

- Eight vCPU
- 16 GB RAM
- One virtual network interface that is configured for the Host Management network (ClusterPublic/VLAN40), which uses VM_Communications vSwitch to connect to vNIC
- One 60 GB virtual hard disk for operating system and application installation (VHDX format)
- Windows Server 2008 R2 SP1 Datacenter edition that is installed and updated
- Highly available through Windows Failover host clustering

> **Note:** Two dedicated SQL Server 2012 instances that are running on a separate, configured SQL Server cluster are required to support these workloads.

- One instance for the Service Manager Database
- One instance for the Service Manager Data Warehouse

For more information about installation instructions for the System Center Service Manager 2012 R2 configuration, see *Microsoft Infrastructure as a Service (IaaS) Fabric Management Deployment Guide*, which is available at this website:

`http://gallery.technet.microsoft.com/Infrastructure-as-a-ecf1cc0b`

7.9 Setting up and configuring WSUS and Windows Deployment Services server

Windows Server Update Services (WSUS) enables you to deploy the latest Microsoft critical updates to servers that are running the Windows operating system. By using WSUS with Cluster Aware Updating, you can control which specific updates are applied consistently across the clustered Windows Server environment.

Windows Deployment Services (WDS) enables you to install an operating system across the network with PXE enabled network interfaces. System Center Virtual Machine Manager requires a WDS server to support the bare metal provisioning services in SCVMM 2012 R2.

Figure 7-11 shows the high-level architecture of WSUS highly available VM that is running on Windows Failover Cluster.

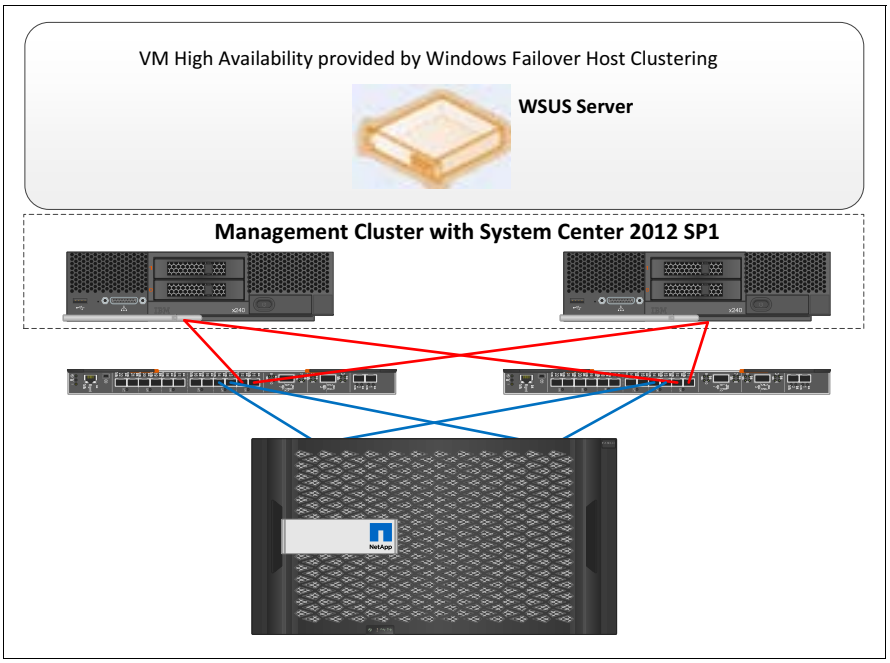

Figure 7-11 WSUS highly available VM

The WSUS/WDS VM must feature the following base VM configuration:

- Two vCPU
- 4 GB RAM
- One virtual network interface that is configured for the Host Management network (ClusterPublic/VLAN40), which uses VM_Communications vSwitch to connect to vNIC
- One 60 GB virtual hard disk for operating system and application installation (VHDX format)
- Windows Server 2012 Datacenter edition that is installed and updated
- Highly available through Windows Failover host clustering

Note: WSUS does require a small database. It is suggested to host the WSUS database on the SCDB instance on the previously configured SQL Server cluster to support this workload.

For more information about how to set up and configure WSUS, see the Microsoft TechNet document *Deploy Windows Update Services in Your Organization*, which is available at this website:

`http://technet.microsoft.com/en-us/library/hh852340.aspx`

7.10 Setting up and configuring Cluster Aware Updating

Cluster Aware Updating (CAU) is a feature in Windows Server 2012 R2 that simplifies the management and updating of cluster nodes. With CAU managing the Windows Server cluster updates, one node is paused at a time, and the VMs are migrated off the node. Updates are then applied, and the cluster node is rebooted and brought back online. This process is repeated until the entire cluster is updated.

The Cluster Aware Updating feature is installed as part of the cluster management administrative tools along with clustering or as part of the Remote Server Administration tools if installed on a separate management console.

From within the CAU console (see Figure 7-12), select the cluster to configure with self-updating options.

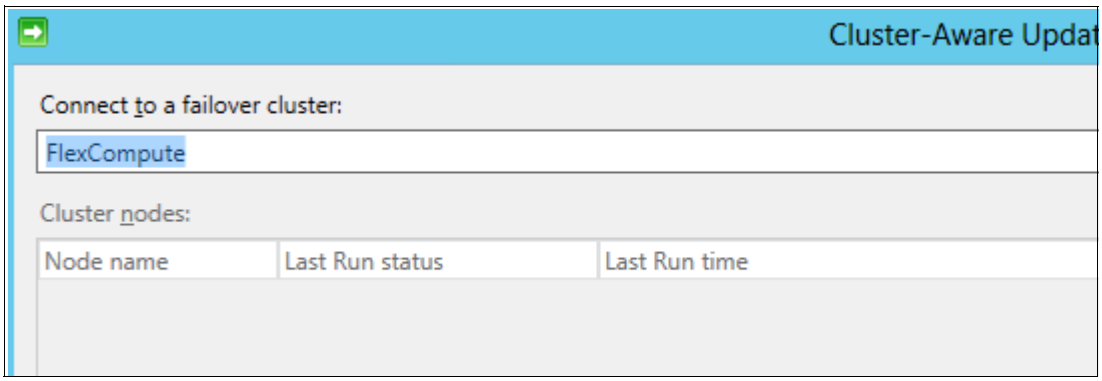

Figure 7-12 CAU console cluster selection

Install the CAU role on the cluster nodes, as shown in Figure 7-13.

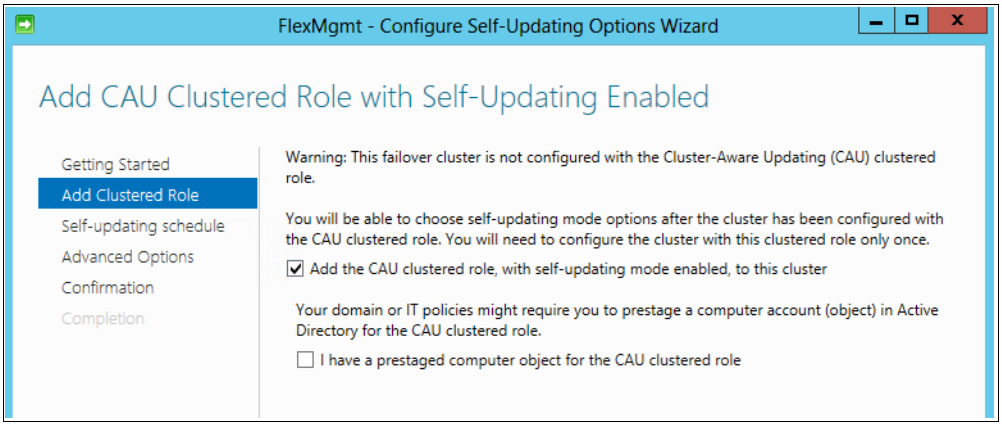

Figure 7-13 Installing CAU on cluster

Complete the Wizard by selecting the self-updating mode, Self-Updating schedule, any optional parameters that are needed, and confirm (see Figure 7-14).

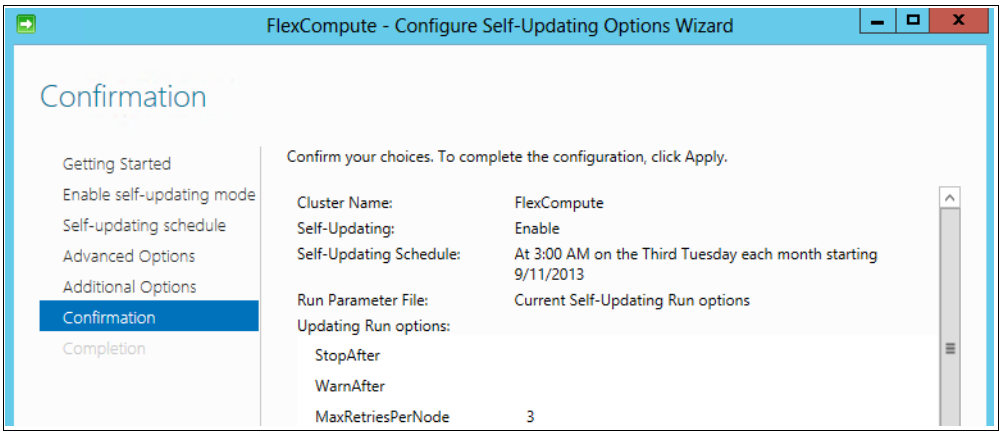

Figure 7-14 Complete CAU wizard

8

Flex System x240 M5 compute node setup

This chapter describes how to set up and configure the Flex System x240 M5 compute node and includes the following topics:

8.1 Beginning the setup

The production fabric for this configuration consists of eight dual socket Flex System x240 M5 Compute Nodes with 128 GB of RAM, and one CN4054R 4-port converged network adapter each. These Flex System compute nodes are configured into a highly available Windows Server cluster to support rapid deployment, resource balancing, and failover by the management fabric.

Setup involves the installation and configuration of Windows Server 2012 R2 Datacenter edition, networking, and storage on each node. After setup and discovery by the previously configured management fabric, production virtual machines can be deployed on this compute cluster.

8.2 Pre-operating system installation

Complete the following steps:

1. Confirm the CN4054R 4-port Ethernet devices are installed in each compute node. The FCoE Feature on Demand (FoD) must be imported on each node via an IMM connection.
2. Validate that firmware levels are consistent across all compute nodes and update if needed.
3. Verify Unified Extensible Firmware Interface (UEFI) performance settings are consistent with corporate operation procedures.
4. Verify that the IMM VLAN settings are consistent with design plan.
5. Configure the CN4054R card to disable multichannel and enable FCoE.
6. Confirm that CN4093 switches were configured. Stacking, LACP teams, and VLANs must be created and assigned.
7. The FAS8040 must be configured as defined and be ready for host mapping assignments to map the volumes for the volumes when the worldwide port name (WWPN) becomes available.
8. The two local disks must be configured as a RAID 1 array.

8.3 Installing and configuring the operating system

Windows Server 2012 R2 Datacenter allows unlimited Windows virtual machine rights on the host servers and is the preferred version for building Hyper-V compute configurations.

Complete the following steps to install and configure the operating system:

1. Install Windows Server 2012 R2 Datacenter Edition.
2. Set your server name and join the domain.
3. Install the Hyper-V role and Failover Clustering feature.
4. Update Windows Server to ensure that any new patches are installed.
5. Download the latest CN4054R NIC and FCoE drivers from this website:

   http://www.ibm.com/support/fixcentral/

6. Install the Emulex OneCommand Manager utility to provide more information for the CN4054R converged network adapter. The utility is available from this website:

   http://www.emulex.com/downloads/oem-qualified-downloads/ibm/vfafc-software-kits

7. Multipath I/O is used to provide balanced and fault tolerant paths to the NetApp FAS8040 storage. This Multipath I/O requires another NetApp MPIO DSM-specific driver to be installed on the host servers before attaching the storage. Download the driver from this website:

   http://mysupport.netapp.com/NOW/cgi-bin/software

8. Install the IBM Systems Director platform agent 6.3.5, which is available from this website:

   http://www-03.ibm.com/systems/director/downloads/agents.html

8.4 Network configuration

Two separate Windows Server 2012 R2 NIC teams are created in the configuration that is described in this section. One team is used to support host server management and cluster private traffic; a second team is used to support Live Migration and VM Communication.

Pay careful attention when you are identifying and enumerating the network interfaces in each host to ensure that teams are spread across the two network interfaces and are properly routed to the correct switch ports. Use Emulex OneCommand Manager Utility to make available each network interface and MAC address.

Note: For more information about mapping the server ports to the two CN4093 switch ports, see Figure 4-2 on page 18 and Figure 4-3 on page 19. Each ASIC has two network interfaces: the top interface is for switch 1, and the bottom interface is for switch 2.

The following PowerShell commands also can be useful (for more information about PowerShell scripts, see Appendix A, “Appendices” on page 113):

```
Get-NetAdapterHardwareInfo | Sort-Object –Property Bus,Function
Get-NetAdapter –InterfaceDescription “IBM O*”
Rename-NetAdapter –name “Ethernet” –NewName “Port_IntA1_SW1”
```

Figure 8-1 shows a Network Interface Card (NIC) interface naming example.

```
PS C:\Users\Administrator.CONTOSO> Get-NetAdapter -InterfaceDescription "IBM O*"

Name                      InterfaceDescription                    ifIndex Status       MacAddress
----                      --------------------                    ------- ------       ----------
Port_IntB1 SW2            IBM OCm11104-N-X Emulex Virtual Fa...#2      13 Up           00-90-FA-67-FE-A2
Port_IntB1 SW1 VLAG       IBM OCm11104-N-X Emulex Virtual Fa...#4      15 Up           00-90-FA-67-FE-9E
Port_IntA1 SW2 VLAG       IBM OCm11104-N-X Emulex Virtual Fabr...      12 Up           00-90-FA-67-FE-9A
Port_IntA1 SW1            IBM OCm11104-N-X Emulex Virtual Fa...#3      14 Up           00-90-FA-67-FE-96


PS C:\Users\Administrator.CONTOSO> _
```

Figure 8-1 Displaying all Emulex network adapters with PowerShell

The team that was created to support cluster public and private communication with the host servers shares ports that are used for FCoE traffic. The CN4093 switches prioritize FCoE traffic over Ethernet data traffic. To further reduce the potential of bandwidth over subscription, quality of service (QoS) limits are placed on these interfaces in Windows Server 2012. This team is created by using the two NIC ports, as described in Table 4-2 on page 27.

Complete the following steps to create the NIC teaming:

1. Create the ClusterTeam team by using the default Switch Independent mode and Address Hash mode by using the following PowerShell commands:

   ```
   New-NetLbfoTeam –name “ClusterTeam” –TeamMembers “Port_IntA1 SW1”, “Port_IntB1
   SW2” –TeamingMode SwitchIndependent
   ```

   Note: The team member names vary by host server.

2. Create the second team off the NIC interfaces that are configured to support LACP. The use of LACP teaming and is then backed by the full aggregated bandwidth of the two ports, as shown in the following commands:

   ```
   New-NetLbfoTeam –name “VLAGTeam” –TeamMembers “Port_IntA1 SW2 VLAG”,
   “Port_IntB1 SW1 VLAG” –TeamingMode Lacp
   ```

   When Windows Server 2012 R2 NIC teaming is complete, there are two teams that are displayed when queried in PowerShell, as shown in Figure 8-2.

```
PS C:\Users\Administrator.CONTOSO> Get-NetLbfoTeam


Name                   : ClusterTeam
Members                : {Port_IntA1 SW1, Port_IntB1 SW2}
TeamNics               : ClusterTeam
TeamingMode            : SwitchIndependent
LoadBalancingAlgorithm : Dynamic
Status                 : Up

Name                   : VLAGTeam
Members                : {Port_IntA1 SW2 VLAG, Port_IntB1 SW1 VLAG}
TeamNics               : VLAGTeam
TeamingMode            : Lacp
LoadBalancingAlgorithm : Dynamic
Status                 : Up
```

Figure 8-2 Windows Server NIC teaming.

3. Virtual switches are created on each of these teams. Each Hyper-V virtual switch can provide one interface to the management operating system. Create a second virtual NIC for use by the host operating system on the ClusterTeam to provide a segregated network path for the Cluster Private/Cluster Shared Volume (CSV) network. The following PowerShell commands are used to create the virtual switches and second virtual NIC:

   ```
   New-VMSwitch –Name “ComputeClusterPublic” –NetAdapterName “ClusterTeam”
   –MinimumBandwidthMode Absolute –AllowManagementOS $true

   Add-VMNetworkAdapter –ManagementOS -Name “ComputeClusterPrivate” –SwitchName
   “ComputeClusterPublic”

   New-VMSwitch –Name “VM_Communication” –NetAdapterName “VLAGTeam”
   –MinimumBandwidthMode Weight –AllowManagementOS $true
   ```

4. Rename the management-facing network interface on the VM_Communication team to reflect the role that it is fulfilling by using the following command:

   ```
   Rename-VMNetworkAdapter -NewName “LiveMigration” –ManagmentOS –name
   “VM_Communication”
   ```

5. Confirm that the network interfaces are available to the management operating systems by using the following Powershell command, as shown in Figure 8-3:

```
Get-VMNetworkAdapter –ManagementOS
```

```
PS C:\Users\Administrator.CONTOSO> Get-VMNetworkAdapter -ManagementOS

Name                IsManagementOs VMName SwitchName          MacAddress   Status
----                -------------- ------ ----------          ----------   ------
ComputeClusterPrivate True                 ComputeClusterPublic 00155D363305 {Ok}
ComputeClusterPublic  True                 ComputeClusterPublic 0090FA1ECD2C {Ok}
ComputeLiveMigration  True                 VM_Communication     0090FA1ECD24 {Ok}
```

Figure 8-3 Displaying all network interfaces created for host partition from vSwitches

6. Assign VLAN IDs to each of these interfaces and verify by using the following PowerShell commands:

```
Set-VMNetworkAdapterVlan –ManagementOS –VMNetworkAdapterName
“MgmtClusterPublic”     –Access –Vlanid 40
Set-VMNetworkAdapterVlan –ManagementOS –VMNetworkAdapterName
“MgmtClusterPrivate”     –Access –Vlanid 30
Set-VMNetworkAdapterVlan –ManagementOS –VMNetworkAdapterName
“MgmtLiveMigration”     –Access –Vlanid 31
```

7. Confirm that the network interfaces VLAN IDs are available to the management operating systems by using the following PowerShell command, as shown in Figure 8-4:

```
Get-VMNetworkAdapterVlan –ManagementOS
```

```
PS C:\Users\Administrator.CONTOSO> Get-VMNetworkAdapterVlan -ManagementOS

VMName VMNetworkAdapterName  Mode   VlanList
------ --------------------  ----   --------
       ComputeClusterPrivate Access 37
       ComputeClusterPublic  Access 40
       ComputeLiveMigration  Access 38
```

Figure 8-4 VLAN Assignments on Management OS NIC interfaces

8. Place bandwidth limits on these network interfaces. The ClusterPublic virtual switch was created with Absolute weighting. By using this configuration, a maximum bandwidth cap can be placed on these network interfaces by using the following PowerShell commands (maximum bandwidth is defined in bits per second):

```
Set–VMNetworkAdpater –ManagementOS  -Name ClusterPublic  -MaximumBandwidth 2GB
Set–VMNetworkAdpater –ManagementOS  -Name ClusterPrivate  -MaximumBandwidth 1GB
```

9. The network interface that is used for Hyper-V Live Migration uses the team that was created by using Weighted mode. Set a minimum bandwidth of 30% for the LiveMigration network by using the following commands:

```
Set–VMNetworkAdpater –ManagementOS  -Name ClusterPrivate
-MinimumBandwidthWeight  30
```

To confirm your Management operating system network adapter names and VLAN, assign TCP/IP addresses and confirm network connectivity for all network connections on each VLAN by using the following commands:

```
New-NetIPAddress –InterfaceAlias “vEthernet (ClusterPublic)” –IPAddress
192.168.40.21 –PrefixLength 24
```

The cluster public network (VLAN 40) should be at the top of the network binding order. The Cluster Private and Live Migration network cannot have any defined default gateways.

8.5 Host storage connections

The following volumes are required to support the Compute host cluster:

- Two 4 TB volume to be used as a Cluster Shared Volumes
- 1 GB volume to be used as the Compute Cluster Quorum

After the switch FCoE configuration and the storage LUN mapping are completed, complete the following steps to make the storage LUNs visible to the host servers:

1. Confirm that the disks are visible in Windows Disk Manager. A disk rescan might be required.
2. From one server, bring each disk online and format it as a GUID Partition Table (GPT) disk for use by the cluster. Assigning drive letters is optional because they are used for specific clustering roles, such as CSV, and Quorum. Validate that each potential host server can see the disks and bring them online.

Note: Only one server can have the disks online at a time until they must be added to Cluster Shared Volumes.

8.6 Creating compute host clusters

Microsoft Windows clustering is used to join the eight host servers in a highly available configuration with which eight servers run virtual machines to support a production environment. Virtual machine workloads must be balanced across all hosts with careful attention to ensure that the combined resources of all virtual machines do not exceed those resources that are available on N-1 cluster nodes. Remaining below this threshold minimizes the effect to your production servers if any single server is taken out of the cluster.

A policy of monitoring resource usage, such as CPU, Memory, and Disk (space, and I/O), keeps the cluster running at optimal levels, and helps proper planning to add resources as needed.

Complete the following steps:

1. Temporarily disable the default USB Remote NDIS Network Device on all cluster nodes because it causes the validation to issue a warning during network detection because of all nodes sharing the same IP address. These devices can be re-enabled after validation.
2. Open the Failover Cluster Manager on one of the eight compute nodes, run the cluster validation wizard to assess the eight physical host servers as potential cluster candidates, and address any errors.

   The cluster validation wizard checks for available cluster compatible host servers, storage, and networking, as shown in Figure 8-5 on page 109.

   Ensure that the intended cluster storage is online in only one of the cluster nodes.

   Address any issues that are flagged during the validation.

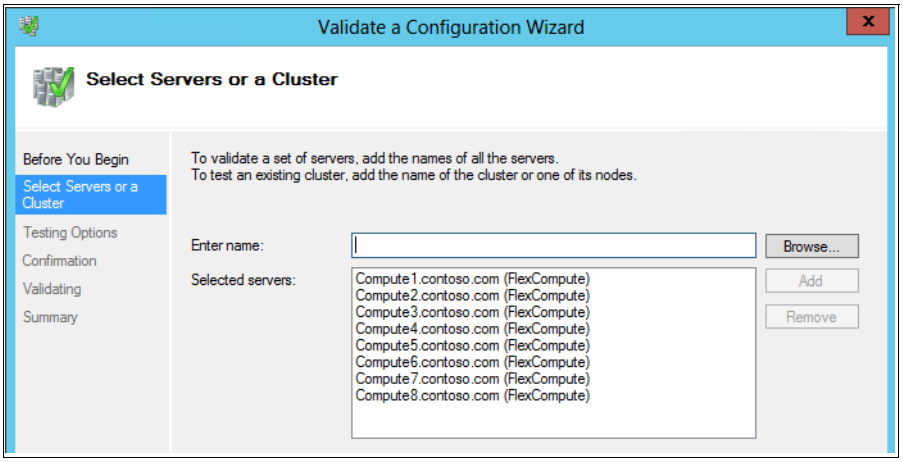

Figure 8-5 Cluster Validation Wizard

3. Use the Failover Cluster Manager to create a cluster with the eight physical host servers to be used for the compute cluster. Follow the steps that are found in the cluster wizard.

   You need a cluster name and IP address. Figure 8-6 shows the Failover Cluster Manager with the 8-node compute cluster visible.

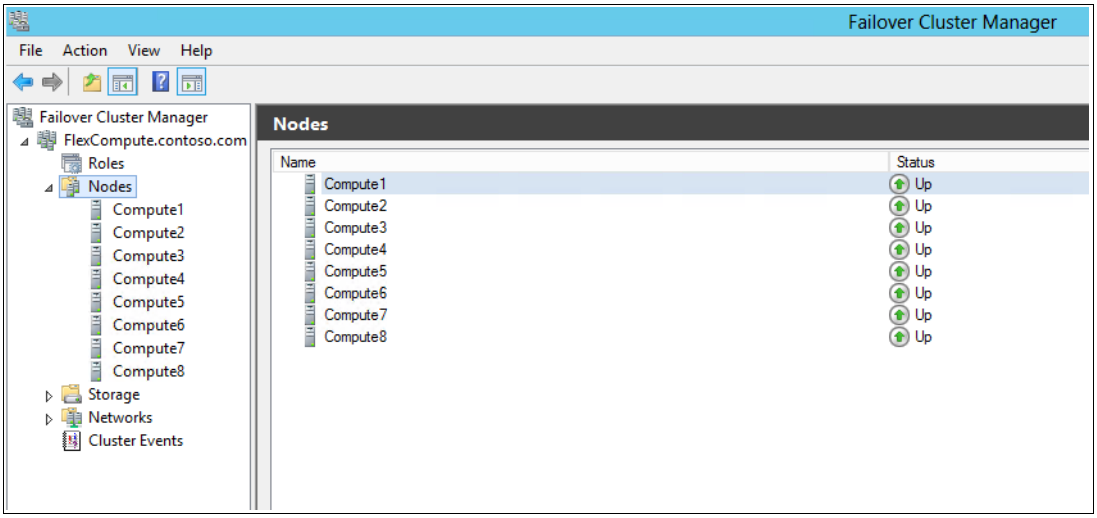

Figure 8-6 Compute Cluster Failover Cluster Manager

4. Add the storage volume that is dedicated to Cluster Shared Volumes.
5. By using Hyper-V Manager, set the default paths for creating the VM to use the Cluster Shared Volumes on each host.

6. Configure the Live Migration settings to route this network traffic over the dedicated link, as shown in Figure 8-7.

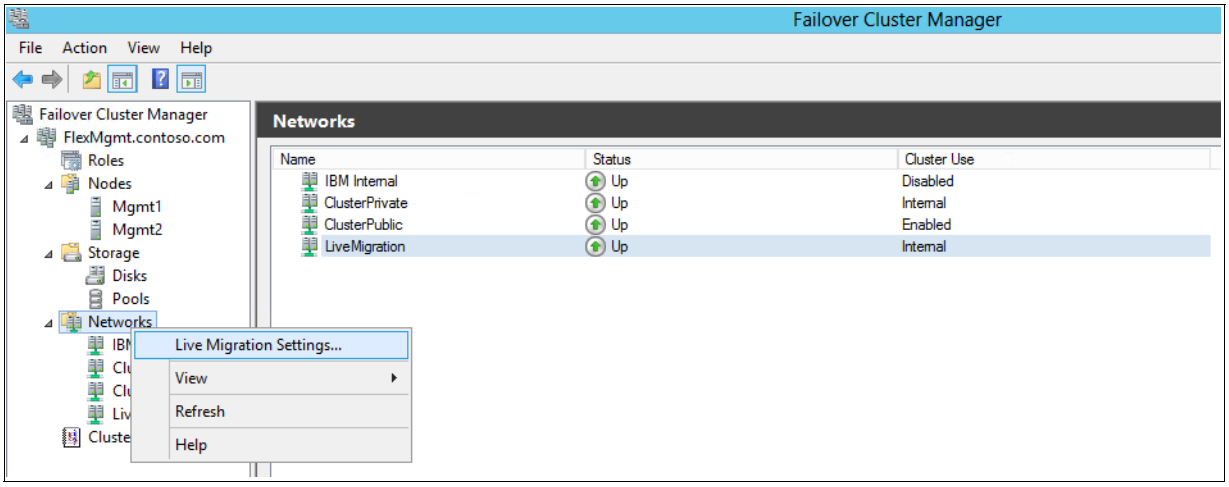

Figure 8-7 Changing Live Migration settings

9

Summary

Upon completing implementation steps, a highly available, elastic, and flexible virtualization environment is created that is based on Flex System, NetApp clustered Data ONTAP Storage, and Microsoft Hyper-V. This environment consists of a two-node cluster that forms the management fabric that is based on Microsoft System Center 2012 R2 and Lenovo XClarity Upward Integration Modules. Also created is an eight-node compute cluster that is used for deployment of production resources.

With enterprise-class multi-level software and hardware, fault tolerance is achieved by configuring a robust collection of industry-leading Flex System Compute Nodes, networking components, and NetApp FAS storage to meet Microsoft's Private Cloud Fast Track program guidelines. The program's unique framework promotes standardized and highly manageable and automated virtualization environments, which can help satisfy even the most challenging business critical virtualization demands, such as high availability, rapid deployment with self-service provisioning, and load balancing. Also promoted is the ability to easily scale the compute cluster and the supporting storage elements for data capacity and throughput performance as business needs grow or change.

Integration between the various components that are followed by best practices from Lenovo, NetApp, and Microsoft that are described in this document lead to optimized infrastructure with predicted performance that is essential to all cloud-based solutions and for meeting the Service Level Objectives that are set by the business.

Following the recommended best practices also ensures minimal experience of downtime and better data protection without compromising performance.

A

Appendices

This appendix includes the following topics:

Bill of Materials

Table A-1 lists the Bill of Material lists for the Microsoft Hyper-V Private Cloud solution in this paper.

Table A-1 Bill of materials

Part number or feature code	Description	Quantity
Flex System Enterprise Chassis		
8721HC1	Flex System Enterprise Chassis Base Model	1
A0TA	Flex System Enterprise Chassis	1
A1NF	Flex System Console Breakout Cable	1
A0UE	Flex System Chassis Management Module	1
A3HH	Flex System Fabric CN4093 10Gb Scalable Switch	2
5053	SFP+ SR Transceiver	8
3701	5m LC-LC Fibre Cable (networking)	8
3268	SFP RJ45 Transceiver	2
3802	1.5m Blue Cat5e Cable	2
A1PH	1m Passive DAC SFP+ Cable	2
A3HL	Flex System Fabric CN4093 10Gb Scalable Switch (Upgrade 1)	2
A0TM	Flex System Chassis Management Module	1
A0UA	Flex System Enterprise Chassis 80mm Fan Module	4
A0UC	Flex System Enterprise Chassis 2500W Power Module Standard	2
A0UD	Flex System Enterprise Chassis 2500W Power Module	4
6252	2.5m, 16A/100-240V, C19 to IEC 320-C20 Rack Power Cable	6
A10X	Flex System Enterprise Chassis Rack Kit	1
A0TW	System Documentation and Software - US English	1
Flex System x240 M5 Management Node		
9532AC1	Flex System node x240 M5 Base Model	2
A1BA	Lenovo Flex System x240 Compute Node	2
ATCG	Intel Xeon Processor E5-2690 v4 14C 2.6GHz 35MB 2400MHz 135W	2
ATD7	Addl Intel Xeon Processor E5-2690 v4 14C 2.6GHz 35MB 2400MHz 135W	2
ATCA	16GB TruDDR4 Memory (2Rx4, 1.2V) PC4-19200 CL17 2400MHz LP RDIMM	32
A4K2	Lenovo Flex System CN4054R 10Gb Virtual Fabric Adapter	2
A1R0	Lenovo Flex System Virtual Fabric Adapter Upgrade (SW Upgrade)	2
AT89	Lenovo 300GB 10K 12Gbps SAS 2.5" G3HS HDD	4

Part number or feature code	Description	Quantity
A5SG	Lenovo Flex System x240 M5 2.5" HDD Backplane	2
Flex System x240 M5 Compute Node		
9532AC1	Flex System node x240 M5 Base Model	8
A1BA	Lenovo Flex System x240 M5 Compute Node	8
ATCG	Intel Xeon Processor E5-2690 v4 14C 2.6GHz 35MB 2400MHz 135W	8
ATD7	Addl Intel Xeon Processor E5-2690 v4 14C 2.6GHz 35MB 2400MHz 135W	8
ATCA	16GB TruDDR4 Memory (2Rx4, 1.2V) PC4-19200 CL17 2400MHz LP RDIMM	64
A4K2	Lenovo Flex System CN4054R 10Gb Virtual Fabric Adapter	8
A1R0	Lenovo Flex System Virtual Fabric Adapter Upgrade (SW Upgrade)	8
AT89	Lenovo 300GB 10K 12Gbps SAS 2.5" G3HS HDD	16
A5SG	Lenovo Flex System x240 M5 2.5" HDD Backplane	8
Rack		
9363RC4	42U 1100mm Enterprise V2 Dynamic Rack	1
A2EV	RFID Tag, AG/AP: 902-928Mhz	1
	DPI Single phase 60A/208V C19 Enterprise PDU (US)	2
4275	5U black plastic filler panel	6
4271	1U black plastic filler panel	2
2304	Rack Assembly - 42U Rack	1
NetApp FAS8040 Storage System Hardware		
FAS8040A-001-R6	FAS8040 High Availability System	2
X6227-R6-C	Chassis, FAS8040/60/80 W/CNTRL Slots, AC PS, -C	1
X5515A-R6-C	Rackmount Kit, 4N2,DS14-Middle, -C, R6	1
X1974A-R6-C	Flash Cache 1TB PCIe Module 2,-C	2
DS2246-21.6TB-0P-R6-C	DSK SHLF, 24x900GB, 6G, 0P, -C	4
DS2246-SL048-12S-0P-R6-C	SSD SHLF, 12x400GB, 0P, -C	1
X5526A-R6-C	Rackmount Kit, 4-Post,Universal, -C, R6	5
DOC-80XX-C	Documents, 80xx, -C	1
X800E-R6-C	Power Cable North America, -C, R6	12
X6585-R6-C	Cable, Ethernet, 3m RJ45 CAT6, -C	1
X6536-R6-C	Cable, Cntlr-Shelf/Switch, 5m, LC/LC, Op, -C	12
X6553-R6-C	Cable, Cntlr-Shelf/Switch, 2m, LC/LC, Op, -C	2
X6557-R6-C	Cable, SAS Cntlr-Shelf/Shelf-Shelf/HA, 0.5m, -C	6

Part number or feature code	Description	Quantity
X6559-R6-C	Cable, SAS Cntlr-Shelf/Shelf-Shelf/HA, 5m, -C	8
X6560-R6-C	Cable, Ethernet, 0.5m RJ45 CAT6, -C	6
X6562-R6-C	Cable, Ethernet, 5m RJ45 CAT6, -C	4
X6599A-R6-C	SFP+ Optical 10Gb Shortwave, FAS80X0, -C	8
X6596-R6-C	SFP+ FC Optical 16Gb, -C	8
NetApp FAS8040 Storage System Software		
SW-2-CL-BASE	SW-2, Base, CL, Node	1
SW-2-8040A-FCP-C	SW-2, FCP, 8040A, -C	2
SW-2-8040A-CIFS-C	SW-2, CIFS, 8040A, -C	2
SW-SDR-WIN	SnapDrive Software, Windows	4
SW-NOC-CORE	SW, NetApp OnCommand core	1
SW-2-8040A-SMIRROR-C	SW-2, SnapMirror, 8040A, -C (Optional)	2
SW-2-8040A-FLEXCLN-C	SW-2, Flexclone, 8040A, -C (Optional)	2
SW-2-8040A-SRESTORE-C	SW-2, SnapRestore, 8040A, -C (Optional)	2
SW-SMHV-MSFT	SnapManager Hyper-V SW (Optional)	2

Networking worksheets

Table A-2 lists the mapping of hosts to switch ports and VLAN assignments.

Table A-2 Networking worksheet

Port	Device	Tagging	PVID	Switch1 VLANs	Switch2 VLANs
Int Port A1	Compute Node1	Yes	1	37, 40, 1002	38, 50
Int Port B1		Yes	1	38, 50	37, 40, 1002
Int Port A2	Compute Node2	Yes	1	37, 40, 1002	38, 50
Int Port A2		Yes	1	38, 50	37, 40, 1002
Int Port A3	Compute Node3	Yes	1	37, 40, 1002	38, 50
Int Port B3		Yes	1	38, 50	37, 40, 1002
Int Port A4	Compute Node4	Yes	1	37, 40, 1002	38, 50
Int Port B4		Yes	1	38, 50	37, 40, 1002
Int Port A5	Compute Node5	Yes	1	37, 40, 1002	38, 50
Int Port B5		Yes	1	38, 50	37, 40, 1002
Int Port A6	Compute Node6	Yes	1	37, 40, 1002	38, 50
Int Port B6		Yes	1	38, 50	37, 40, 1002

Port	Device	Tagging	PVID	Switch1 VLANs	Switch2 VLANs
Int Port A7	Compute Node7	Yes	1	37, 40, 1002	38, 50
Int Port B7		Yes	1	38, 50	37, 40, 1002
Int Port A8	Compute Node8	Yes	1	37, 40, 1002	38, 50
Int Port B8		Yes	1	38, 50	37, 40, 1002
Int Port A9	Management Node1	Yes	1	30, 40, 1002	31, 33, 35, 40
Int Port B9		Yes	1	31, 33, 35, 40	30, 40, 1002
Int Port A10	Management Node2	Yes	1	30, 40, 1002	30, 33, 35, 40
Int Port B10		Yes	1	31, 33, 35, 40	30, 40, 1002
Ext Port EXT1	Inter switch Link	Yes	4094	30, 31, 33, 35, 37, 38, 40, 50, 4094	30, 31, 33, 35, 37, 38, 40, 50, 4094
Ext Port EXT2		Yes	4094	30, 31, 33, 35, 37, 38, 40, 50, 4094	30, 31, 33, 35, 37, 38, 40, 50, 4094
Ext Port EXT11	N/A	N/A	N/A	Used for FCoE Configuration	Used for FCoE Configuration
Ext Port EXT12		N/A	N/A	(set to type fc)	(set to type fc)
Ext Port EXT13	NetApp FAS8040 (Ctrl-1)	Yes	1	1002	1002
Ext Port EXT14	NetApp FAS8040 (Ctrl-2)	Yes	1	1002	1002
Ext Port EXT15	NetApp FAS8040 (SMB)	Yes	1	40	40
Ext Port EXT16		Yes	1	Unused	Unused
Ext Port EXT17	CorpNet Uplink	Yes	1	Uplink (LACP) VLAN as needed	Uplink (LACP) VLAN as needed
Ext Port EXT18	CorpNet Uplink	Yes	1	Uplink (LACP) VLAN as needed	Uplink (LACP) VLAN as needed

Switch configuration

This section describes the configuration files for Switch1 and Switch2.

Switch1

Switch1 features the following file content:

```
version "7.8.6"
switch-type "IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch(Upgrade1)"
iscli-new
!
!
!
```

```
system port EXT11,EXT12 type fc
interface fc EXT11
   switchport trunk allowed vlan 1,1002
interface fc EXT12
   switchport trunk allowed vlan 1,1002
!
system idle 60
!
!
interface port INTA1
   switchport mode trunk
   switchport trunk allowed vlan 1,37,40,1002
   exit
!
interface port INTA2
   switchport mode trunk
   switchport trunk allowed vlan 1,37,40,1002
   exit
!
interface port INTA3
   switchport mode trunk
   switchport trunk allowed vlan 1,37,40,1002
   exit
!
interface port INTA4
   switchport mode trunk
   switchport trunk allowed vlan 1,37,40,1002
   exit
!
interface port INTA5
   switchport mode trunk
   switchport trunk allowed vlan 1,37,40,1002
   exit
!
interface port INTA6
   switchport mode trunk
   switchport trunk allowed vlan 1,37,40,1002
   exit
!
interface port INTA7
   switchport mode trunk
   switchport trunk allowed vlan 1,37,40,1002
   exit
!
interface port INTA8
   switchport mode trunk
   switchport trunk allowed vlan 1,37,40,1002
   exit
!
interface port INTA9
   switchport mode trunk
   switchport trunk allowed vlan 1,30,40,1002
   exit
!
interface port INTA10
```

```
   switchport mode trunk
   switchport trunk allowed vlan 1,30,40,1002
   exit
!
interface port INTB1
   switchport mode trunk
   switchport trunk allowed vlan 1,38,50
   exit
!
interface port INTB2
   switchport mode trunk
   switchport trunk allowed vlan 1,38,50
   exit
!
interface port INTB3
   switchport mode trunk
   switchport trunk allowed vlan 1,38,50
   exit
!
interface port INTB4
   switchport mode trunk
   switchport trunk allowed vlan 1,38,50
   exit
!
interface port INTB5
   switchport mode trunk
   switchport trunk allowed vlan 1,38,50
   exit
!
interface port INTB6
   switchport mode trunk
   switchport trunk allowed vlan 1,38,50
   exit
!
interface port INTB7
   switchport mode trunk
   switchport trunk allowed vlan 1,38,50
   exit
!
interface port INTB8
   switchport mode trunk
   switchport trunk allowed vlan 1,38,50
   exit
!
interface port INTB9
   switchport mode trunk
   switchport trunk allowed vlan 1,31,33,35,40
   exit
!
interface port INTB10
   switchport mode trunk
   switchport trunk allowed vlan 1,31,33,35,40
   exit
!
interface port EXT1
```

```
   switchport mode trunk
   switchport trunk allowed vlan 30-31,33,35,37-38,40,50,4094
   switchport trunk native vlan 4094
   exit
!
interface port EXT2
   switchport mode trunk
   switchport trunk allowed vlan 30-31,33,35,37-38,40,50,4094
   switchport trunk native vlan 4094
   exit
!
interface port EXT13
   description "NetApp Cntrl-1"
   switchport mode trunk
   switchport trunk allowed vlan 1,1002
   exit
!
interface port EXT14
   description "NetApp Cntrl-2"
   switchport mode trunk
   switchport trunk allowed vlan 1,1002
   exit
!
interface port EXT15
   description “NetApp SMB”
switchport mode trunk
   switchport trunk allowed vlan 1,40
   exit
!
interface port EXT17
   description “Uplink1”
switchport mode trunk
   switchport trunk allowed vlan 1,40,50
   exit
!
interface port EXT18
   description “Uplink2”
   switchport mode trunk
   switchport trunk allowed vlan 1,40,50
   exit
vlan 30
   name "VLAN 30"
!
vlan 31
   name "VLAN 31"
!
vlan 33
   name "VLAN 33"
!
vlan 35
   name "VLAN 35"
!
vlan 37
   name "VLAN 37"
!
```

```
vlan 38
   name "VLAN 38"
!
vlan 40
   name "VLAN 40"
!
vlan 50
   name "VLAN 50"
!
vlan 1002
   name "VLAN 1002"
   fcf enable
!
vlan 4094
   name "VLAN 4094"
!
!
no spanning-tree stp 20 enable
spanning-tree stp 20 vlan 4094
!
spanning-tree stp 30 vlan 30
!
spanning-tree stp 31 vlan 31
!
spanning-tree stp 33 vlan 33
!
spanning-tree stp 35 vlan 35
!
spanning-tree stp 37 vlan 37
!
spanning-tree stp 38 vlan 38
!
spanning-tree stp 40 vlan 40
!
spanning-tree stp 50 vlan 50
!
spanning-tree stp 113 vlan 1002
!
!
no logging console
!
interface port INTB1
   lacp mode active
   lacp key 101
!
interface port INTB2
   lacp mode active
   lacp key 102
!
interface port INTB3
   lacp mode active
   lacp key 103
!
interface port INTB4
   lacp mode active
```

```
   lacp key 104
!
interface port INTB5
   lacp mode active
   lacp key 105
!
interface port INTB6
   lacp mode active
   lacp key 106
!
interface port INTB7
   lacp mode active
   lacp key 107
!
interface port INTB8
   lacp mode active
   lacp key 108
!
interface port INTB9
   lacp mode active
   lacp key 109
!
interface port INTB10
   lacp mode active
   lacp key 110
!
interface port EXT1
   lacp mode active
   lacp key 100
!
interface port EXT2
   lacp mode active
   lacp key 100
!
!
!
vlag enable
vlag tier-id 10
vlag isl adminkey 100
vlag adminkey 101 enable
vlag adminkey 102 enable
vlag adminkey 103 enable
vlag adminkey 104 enable
vlag adminkey 105 enable
vlag adminkey 106 enable
vlag adminkey 107 enable
vlag adminkey 108 enable
vlag adminkey 109 enable
vlag adminkey 110 enable
!
!
!
fcoe fips enable
!
cee enable
```

```
!
!
!
fcalias ITE1_Port_IntA1 wwn 10:00:00:90:fa:67:fe:97
fcalias ITE2_Port_IntA2 wwn 10:00:00:90:fa:6a:3c:cd
fcalias ITE3_Port_IntA3 wwn 10:00:00:90:fa:5d:8b:07
fcalias ITE4_Port_IntA4 wwn 10:00:00:90:fa:6a:0e:61
fcalias ITE5_Port_IntA5 wwn 10:00:00:90:fa:67:ea:d3
fcalias ITE6_Port_IntA6 wwn 10:00:00:90:fa:67:f7:7f
fcalias ITE7_Port_IntA7 wwn 10:00:00:90:fa:6a:ff:9f
fcalias ITE8_Port_IntA8 wwn 10:00:00:90:fa:6a:47:6b
fcalias ITE9_Port_IntA9 wwn 10:00:00:90:fa:6a:4e:37
fcalias ITE10_Port_IntA10 wwn 10:00:00:90:fa:6a:42:17
fcalias NTAP_SVM1_Port1 wwn 20:0b:00:a0:98:5e:5b:b4
fcalias NTAP_SVM1_Port2 wwn 20:0d:00:a0:98:5e:5b:b4
fcalias NTAP_SVM2_Port1 wwn 20:07:00:a0:98:5e:5b:b4
fcalias NTAP_SVM2_Port2 wwn 20:09:00:a0:98:5e:5b:b4
!
zone name SW1Zone_SVM1_ComputeSvrs
   member fcalias ITE1_Port_IntA1
   member fcalias ITE2_Port_IntA2
   member fcalias ITE3_Port_IntA3
   member fcalias ITE4_Port_IntA4
   member fcalias ITE5_Port_IntA5
   member fcalias ITE6_Port_IntA6
   member fcalias ITE7_Port_IntA7
   member fcalias ITE8_Port_IntA8
   member fcalias NTAP_SVM1_Port1
   member fcalias NTAP_SVM1_Port2
zone name SW1Zone_SVM2_ComputeSvrs
   member fcalias ITE1_Port_IntA1
   member fcalias ITE2_Port_IntA2
   member fcalias ITE3_Port_IntA3
   member fcalias ITE4_Port_IntA4
   member fcalias ITE5_Port_IntA5
   member fcalias ITE6_Port_IntA6
   member fcalias ITE7_Port_IntA7
   member fcalias ITE8_Port_IntA8
   member fcalias NTAP_SVM2_Port1
   member fcalias NTAP_SVM2_Port2
zone name SW1Zone_SVM1_MgmtSvrs
   member fcalias ITE9_Port_IntA9
   member fcalias ITE10_Port_IntA10
   member fcalias NTAP_SVM1_Port1
   member fcalias NTAP_SVM1_Port2
zone name SW1Zone_SVM2_MgmtSvrs
   member fcalias ITE9_Port_IntA9
   member fcalias ITE10_Port_IntA10
   member fcalias NTAP_SVM2_Port1
   member fcalias NTAP_SVM2_Port2
zoneset name SW1_ZoneSet
   member SW1Zone_SVM1_ComputeSvrs
   member SW1Zone_SVM2_ComputeSvrs
   member SW1Zone_SVM1_MgmtSvrs
   member SW1Zone_SVM2_MgmtSvrs
```

```
zoneset activate name SW1_ZoneSet
!
!
!
ntp enable
ntp ipv6 primary-server fe80::211:25ff:fec3:7ed2 MGT
ntp interval 4095
ntp authenticate
ntp primary-key 53694
!
ntp message-digest-key 11217 md5-ekey
"c137f0b68022a0a2acbeb3f7439203506bb455055e742f7ec1568d9f98eea8829246054379d99a140239bcbec953d67
94d03fc52c668451fab51c0e20b413073"
!
ntp message-digest-key 25226 md5-ekey
"d436e5b79022a0a2bcbea6f653920350a603dc938d55e89ba7cff4f7fe083bc804d15cd027c575a48f460cd54d615be
7b29c5d08fe415a5ed4a78974e95374ac"
!
ntp message-digest-key 38739 md5-ekey
"e444d5c580008080ac9cb6a643b0237278f4a3eeabde13c32752a5c2509496a6a96d41db81526c68ae670e4468710cf
c099a10959e3945b56e4929660b9e8796"
!
ntp message-digest-key 53694 md5-ekey
"f2dcc35d92088208be94a2b651b821fad31366b85c96bec47caef66010ea327015bd333709329350bfebc6729673e33
e7a87a765c06e7f0fdd7d825dd8e2ec16"
!
ntp trusted-key 11217,25226,38739,53694
!
end
```

Switch2

Switch2 features the following file content:

```
version "7.8.6"
switch-type "IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch(Upgrade1)"
iscli-new
!
!
system port EXT11,EXT12 type fc
interface fc EXT11
   switchport trunk allowed vlan 1,1002
interface fc EXT12
   switchport trunk allowed vlan 1,1002
!
system idle 60
!
!
interface port INTA1
   switchport mode trunk
   switchport trunk allowed vlan 1,38,50
   exit
!
```

```
interface port INTA2
   switchport mode trunk
   switchport trunk allowed vlan 1,38,50
   exit
!
interface port INTA3
   switchport mode trunk
   switchport trunk allowed vlan 1,38,50
   exit
!
interface port INTA4
   switchport mode trunk
   switchport trunk allowed vlan 1,38,50
   exit
!
interface port INTA5
   switchport mode trunk
   switchport trunk allowed vlan 1,38,50
   exit
!
interface port INTA6
   switchport mode trunk
   switchport trunk allowed vlan 1,38,50
   exit
!
interface port INTA7
   switchport mode trunk
   switchport trunk allowed vlan 1,38,50
   exit
!
interface port INTA8
   switchport mode trunk
   switchport trunk allowed vlan 1,38,50
   exit
!
interface port INTA9
   switchport mode trunk
   switchport trunk allowed vlan 1,31,33,35,40
   exit
!
interface port INTA10
   switchport mode trunk
   switchport trunk allowed vlan 1,31,33,35,40
   exit
!
interface port INTB1
   switchport mode trunk
   switchport trunk allowed vlan 1,37,40,1002
   exit
!
interface port INTB2
   switchport mode trunk
   switchport trunk allowed vlan 1,37,40,1002
   exit
!
```

```
interface port INTB3
   switchport mode trunk
   switchport trunk allowed vlan 1,37,40,1002
   exit
!
interface port INTB4
   switchport mode trunk
   switchport trunk allowed vlan 1,37,40,1002
   exit
!
interface port INTB5
   switchport mode trunk
   switchport trunk allowed vlan 1,37,40,1002
   exit
!
interface port INTB6
   switchport mode trunk
   switchport trunk allowed vlan 1,37,40,1002
   exit
!
interface port INTB7
   switchport mode trunk
   switchport trunk allowed vlan 1,37,40,1002
   exit
!
interface port INTB8
   switchport mode trunk
   switchport trunk allowed vlan 1,37,40,1002
   exit
!
interface port INTB9
   switchport mode trunk
   switchport trunk allowed vlan 1-2,30,40,1002
   exit
!
interface port INTB10
   switchport mode trunk
   switchport trunk allowed vlan 1-2,30,40,1002
   switchport trunk native vlan 2
   exit
!
interface port EXT1
   switchport mode trunk
   switchport trunk allowed vlan 30-31,33,35,37-38,40,50,4094
   switchport trunk native vlan 4094
   exit
!
interface port EXT2
   switchport mode trunk
   switchport trunk allowed vlan 30-31,33,35,37-38,40,50,4094
   switchport trunk native vlan 4094
   exit
!
interface port EXT13
   description "NetApp Cntrl-1"
```

```
   switchport mode trunk
   switchport trunk allowed vlan 1,1002
   exit
!
interface port EXT14
   description "NetApp Cntrl-2"
   switchport mode trunk
   switchport trunk allowed vlan 1,1002
   exit
!
interface port EXT15
   description “NetApp SMB”
switchport mode trunk
   switchport trunk allowed vlan 1,40
   exit
!
interface port EXT17
   description “Uplink1”
switchport mode trunk
   switchport trunk allowed vlan 1,40,50
   exit
!
interface port EXT18
   description “Uplink2”
   switchport mode trunk
   switchport trunk allowed vlan 1,40,50
   exit
!
vlan 30
   name "VLAN 30"
!
vlan 31
   name "VLAN 31"
!
vlan 33
   name "VLAN 33"
!
vlan 35
   name "VLAN 35"
!
vlan 37
   name "VLAN 37"
!
vlan 38
   name "VLAN 38"
!
vlan 40
   name "VLAN 40"
!
vlan 50
   name "VLAN 50"
!
vlan 1002
   name "VLAN 1002"
   fcf enable
```

```
!
vlan 4094
   name "VLAN 4094"
!
!
no spanning-tree stp 20 enable
spanning-tree stp 20 vlan 4094
!
spanning-tree stp 30 vlan 30
!
spanning-tree stp 31 vlan 31
!
spanning-tree stp 33 vlan 33
!
spanning-tree stp 35 vlan 35
!
spanning-tree stp 37 vlan 37
!
spanning-tree stp 38 vlan 38
!
spanning-tree stp 40 vlan 40
!
spanning-tree stp 50 vlan 50
!
spanning-tree stp 113 vlan 1002
!
!
no logging console
!
interface port INTA1
   lacp mode active
   lacp key 101
!
interface port INTA2
   lacp mode active
   lacp key 102
!
interface port INTA3
   lacp mode active
   lacp key 103
!
interface port INTA4
   lacp mode active
   lacp key 104
!
interface port INTA5
   lacp mode active
   lacp key 105
!
interface port INTA6
   lacp mode active
   lacp key 106
!
interface port INTA7
   lacp mode active
```

```
   lacp key 107
!
interface port INTA8
   lacp mode active
   lacp key 108
!
interface port INTA9
   lacp mode active
   lacp key 109
!
interface port INTA10
   lacp mode active
   lacp key 110
!
interface port EXT1
   lacp mode active
   lacp key 100
!
interface port EXT2
   lacp mode active
   lacp key 100
!
!
!
vlag enable
vlag tier-id 10
vlag isl adminkey 100
vlag adminkey 101 enable
vlag adminkey 102 enable
vlag adminkey 103 enable
vlag adminkey 104 enable
vlag adminkey 105 enable
vlag adminkey 106 enable
vlag adminkey 107 enable
vlag adminkey 108 enable
vlag adminkey 109 enable
vlag adminkey 110 enable
!
!
!
fcoe fips enable
!
cee enable
!
!
!
fcalias ITE1_Port_IntB1 wwn 10:00:00:90:fa:67:fe:a3
fcalias ITE2_Port_IntB2 wwn 10:00:00:90:fa:6a:3c:d9
fcalias ITE3_Port_IntB3 wwn 10:00:00:90:fa:5d:8b:13
fcalias ITE4_Port_IntB4 wwn 10:00:00:90:fa:6a:0e:6d
fcalias ITE5_Port_IntB5 wwn 10:00:00:90:fa:67:ea:df
fcalias ITE6_Port_IntB6 wwn 10:00:00:90:fa:67:f7:8b
fcalias ITE7_Port_IntB7 wwn 10:00:00:90:fa:6a:ff:ab
fcalias ITE8_Port_IntB8 wwn 10:00:00:90:fa:6a:47:77
fcalias ITE9_Port_IntB9 wwn 10:00:00:90:fa:6a:4e:43
```

```
fcalias ITE10_Port_IntB10 wwn 10:00:00:90:fa:6a:42:23
fcalias NTAP_SVM1_Port1 wwn 20:0e:00:a0:98:5e:5b:b4
fcalias NTAP_SVM1_Port2 wwn 20:0c:00:a0:98:5e:5b:b4
fcalias NTAP_SVM2_Port1 wwn 20:08:00:a0:98:5e:5b:b4
fcalias NTAP_SVM2_Port2 wwn 20:06:00:a0:98:5e:5b:b4
zone name SW2Zone_ComputeSvrs
   member fcalias ITE3_Port_IntB3
   member pwwn 10:00:00:90:fa:6a:3c:d9
   member pwwn 10:00:00:90:fa:5d:8b:13
   member pwwn 10:00:00:90:fa:6a:0e:6d
   member pwwn 10:00:00:90:fa:67:ea:df
   member pwwn 10:00:00:90:fa:67:f7:8b
   member pwwn 10:00:00:90:fa:6a:ff:ab
   member pwwn 10:00:00:90:fa:6a:47:77
   member pwwn 50:05:07:68:02:50:02:c1
   member pwwn 50:05:07:68:02:50:02:c2
zone name SW2Zone_MgmtSvrs
   member pwwn 10:00:00:90:fa:6a:4e:43
   member pwwn 10:00:00:90:fa:6a:42:23
   member pwwn 50:05:07:68:02:50:02:c1
   member pwwn 50:05:07:68:02:50:02:c2
zone name SW2Zone_SVM1_ComputeSvrs
   member fcalias ITE1_Port_IntB1
   member fcalias ITE2_Port_IntB2
   member fcalias ITE3_Port_IntB3
   member fcalias ITE4_Port_IntB4
   member fcalias ITE5_Port_IntB5
   member fcalias ITE6_Port_IntB6
   member fcalias ITE7_Port_IntB7
   member fcalias ITE8_Port_IntB8
   member fcalias NTAP_SVM1_Port1
   member fcalias NTAP_SVM1_Port2
zone name SW2Zone_SVM2_ComputeSvrs
   member fcalias ITE1_Port_IntB1
   member fcalias ITE2_Port_IntB2
   member fcalias ITE3_Port_IntB3
   member fcalias ITE4_Port_IntB4
   member fcalias ITE5_Port_IntB5
   member fcalias ITE6_Port_IntB6
   member fcalias ITE7_Port_IntB7
   member fcalias ITE8_Port_IntB8
   member fcalias NTAP_SVM2_Port1
   member fcalias NTAP_SVM2_Port2
zone name SW2Zone_SVM1_MgmtSvrs
   member fcalias ITE9Port_IntB9
   member fcalias ITE10_Port_IntB10
   member fcalias NTAP_SVM1_Port1
   member fcalias NTAP_SVM1_Port2
zone name SW2Zone_SVM2_MgmtSvrs
   member fcalias ITE9Port_IntB9
   member fcalias ITE10_Port_IntB10
   member fcalias NTAP_SVM2_Port1
   member fcalias NTAP_SVM2_Port2
zoneset name SW2_ZoneSet
   member SW2Zone_SVM1_ComputeSvrs
```

```
   member SW2Zone_SVM2_ComputeSvrs
   member SW2Zone_SVM1_MgmtSvrs
   member SW2Zone_SVM2_MgmtSvrs
zoneset activate name SW2_ZoneSet
!
!
!
ntp enable
ntp ipv6 primary-server fe80::211:25ff:fec3:7ed2 MGT
ntp interval 4095
ntp authenticate
ntp primary-key 53694
!
ntp message-digest-key 11217 md5-ekey
"fa66db0d8a408a08b636b2e649f029fac7b1bddfd37cb3f873efe6b862109d9ec2d03cb541cc8eb3dbe61d220ad4b74
d22282589f5d3c7f955d1eab310a99b99"
!
ntp message-digest-key 25226 md5-ekey
"b37f921482408200be36f3f741f021f24d412452c14a51d27bf21fe4e77932c5396ccfcb374751bcd388375e187cb61
e03e6f9e5170fd1d9ee4230a8cc4bb1ef"
!
ntp message-digest-key 26424 md5-ekey
"ad518c3a8840882ab436e7f34bf02bd89f95048923706b6b55565610b8e2fbfe9a6a821dbd889ace6509da40257e647
55e9de1127adaa106e1102f00913f2fc8"
!
ntp message-digest-key 38739 md5-ekey
"5e057f6e0a002a2a3676b6a7c9b089d8b8eaf22617904ab22143d324f04fee2639bc172764aff3aee54a719c7302401
cf0cf0b9c078b74b919121ba269756132"
!
ntp message-digest-key 52895 md5-ekey
"498968e2088028a234f6a3a3cb308b5000244ee03bd9652a2efa428125d541ca1968e544e102ff812f90cb3a4dea396
8a95c749ec878614fadec42ba188b548b"
!
ntp message-digest-key 53694 md5-ekey
"48dc69b708c028a234b6a2f6cb708b507555445b916f0961cbe83a8163c5d4fba80ca5b2fa71459941e41342e6a3c32
a3fe1e78a3a028b7cdc5e17f141de62e1"
!
ntp trusted-key 11217,25226,26424,38739,52895,53694
!
end
```

PowerShell scripts

In this section, some sample Powershell scripts are provided that can be useful to automate some of the repetitive tasks when the Microsoft Hyper-V Private Cloud is set up.

Management node network configuration

This section includes the following management node network configuration scripts:

```
#Build the string list for the Physical NIC Interface Names. This script needs to be modified
per ITE.
# The port names would vary by ITE
$NewNameList = "PortA9_SW1", "PortA9_SW2 VLAG", "PortB9_SW1 VLAG", "PortB9_SW2"
$NIC_List=Get-NetAdapterHardwareInfo | Sort-Object -Property Bus,Function
#Assumes that the list returned will be in the same order since sorted by Bus & Function.
Rename-NetAdapter $NIC_List[0].name -NewName $NewNameList[0]
Rename-NetAdapter $NIC_List[1].name -NewName $NewNameList[1]
Rename-NetAdapter $NIC_List[2].name -NewName $NewNameList[2]
Rename-NetAdapter $NIC_List[3].name -NewName $NewNameList[3]
#Create the Managment Team, vSwitch, and vNICs
New-NetLbfoTeam -name "ClusterTeam" -TeamMembers $NIC_List[0].name, $NIC_List[3].name
-TeamingMode SwitchIndependent
New-VMSwitch "MgmtClusterPublic" -NetAdapterName "ClusterTeam" -MinimumBandwidthMode Absolute
-AllowManagementOS $true
Add-VMNetworkAdapter -ManagementOS -Name "MgmtClusterPrivate" -SwitchName "MgmtClusterPublic"
#Create the LACP Team, vSwitch, and vNIC
#New-NetLbfoTeam -name "LACPTeam" -TeamMembers $NIC_List[1].name, $NIC_List[2].name -TeamingMode
Lacp
New-VMSwitch "VM_Communication" -NetAdapterName "LACPTeam" -MinimumBandwidthMode Weight
-AllowManagementOS $true
Rename-VMNetworkAdapter -NewName "MgmtLiveMigration" -ManagementOS -Name "VM_Communication"
#Set the Bandwidth restrictions for each VMNetworkAdapter facing the ManagementOS
Set-VMNetworkAdapter -ManagementOS -Name "MgmtClusterPublic" -MaximumBandwidth 2GB
Set-VMNetworkAdapter -ManagementOS -Name "MgmtClusterPrivate" -MaximumBandwidth 1GB
Set-VMNetworkAdapter -ManagementOS -Name "MgmtLiveMigration" -MinimumBandwidthWeight 30
#Set the VLANs for the above devices (Compute will be different than Mgmt)
Set-VMNetworkAdapterVLAN -ManagementOS -VMNetworkAdapterName "MgmtClusterPublic" -Access -VlanId
40
Set-VMNetworkAdapterVLAN -ManagementOS -VMNetworkAdapterName "MgmtClusterPrivate" -Access
-VlanId 30
Set-VMNetworkAdapterVLAN -ManagementOS -VMNetworkAdapterName "MgmtLiveMigration" -Access -VlanId
31
#Get-NetAdapterHardwareInfo | Sort-Object -Property bus,Function
Get-VMNetworkAdapter -ManagementOS
#Set the IP Addresses for the Management Interfaces (each ITE should have a unique address)
New-NetIPAddress -InterfaceAlias "vEthernet (MgmtClusterPublic)" -IPAddress 192.168.40.21
-PrefixLength 24
New-NetIPAddress -InterfaceAlias "vEthernet (MgmtClusterPrivate)" -IPAddress 192.168.30.21
-PrefixLength 24
New-NetIPAddress -InterfaceAlias "vEthernet (MgmtLiveMigration)" -IPAddress 192.168.31.21
-PrefixLength 24
#Set the DNS Address
Set-DnsClientServerAddress -Interfacealias "vEthernet (MgmtClusterPublic)" -ServerAddress
192.168.40.10
```

Compute node network configuration

This section includes the following node network configuration scripts:

```
#Build the string list for the Physical NIC Interface Names. This script needs to be modified
per ITE.
$NewNameList = "PortA1_SW1", "PortA1_SW2 VLAG", "PortB1_SW1 VLAG", "PortB1_SW2"
$NIC_List=Get-NetAdapterHardwareInfo | Sort-Object -Property Bus,Function
#Assumes that the list returned will be in the same order since sorted by Bus & Function.
Rename-NetAdapter $NIC_List[0].name -NewName $NewNameList[0]
Rename-NetAdapter $NIC_List[1].name -NewName $NewNameList[1]
Rename-NetAdapter $NIC_List[2].name -NewName $NewNameList[2]
Rename-NetAdapter $NIC_List[3].name -NewName $NewNameList[3]
#Create the Managment Team, vSwitch, and vNICs
New-NetLbfoTeam -name "ClusterTeam" -TeamMembers $NIC_List[0].name, $NIC_List[3].name
-TeamingMode SwitchIndependent
New-VMSwitch -Name "ComputeClusterPublic" -NetAdapterName "ClusterTeam" -MinimumBandwidthMode
Absolute -AllowManagementOS $true
Add-VMNetworkAdapter -ManagementOS -Name "ComputeClusterPrivate" -SwitchName
"ComputeClusterPublic"
#Create the LACP Team, vSwitch, and vNIC
New-NetLbfoTeam -name "LACPTeam" -TeamMembers $NIC_List[1].name, $NIC_List[2].name -TeamingMode
Lacp
New-VMSwitch -Name "VM_Communication" -NetAdapterName "LACPTeam" -MinimumBandwidthMode Weight
-AllowManagementOS $true
Rename-VMNetworkAdapter -NewName "ComputeLiveMigration" -ManagementOS -Name "VM_Communication"
#Set the Bandwidth restrictions for each VMNetworkAdapter facing the ManagementOS
Set-VMNetworkAdapter -ManagementOS -Name "ComputeClusterPublic" -MaximumBandwidth 2GB
Set-VMNetworkAdapter -ManagementOS -Name "ComputeClusterPrivate" -MaximumBandwidth 1GB
Set-VMNetworkAdapter -ManagementOS -Name "ComputeLiveMigration" -MinimumBandwidthWeight 30
#Set the VLANs for the above devices (Compute will be different than Mgmt)
Set-VMNetworkAdapterVLAN -ManagementOS -VMNetworkAdapterName "ComputeClusterPublic" -Access
-VlanId 40
Set-VMNetworkAdapterVLAN -ManagementOS -VMNetworkAdapterName "ComputeClusterPrivate" -Access
-VlanId 37
Set-VMNetworkAdapterVLAN -ManagementOS -VMNetworkAdapterName "ComputeLiveMigration" -Access
-VlanId 38
#Set the IP Addresses for the Management Interfaces. (each ITE should have a unique address)
New-NetIPAddress -InterfaceAlias "vEthernet (ComputeClusterPublic)" -IPAddress 192.168.40.31
-PrefixLength 24
New-NetIPAddress -InterfaceAlias "vEthernet (ComputeClusterPrivate)" -IPAddress 192.168.37.31
-PrefixLength 24
New-NetIPAddress -InterfaceAlias "vEthernet (ComputeLiveMigration)" -IPAddress 192.168.38.31
-PrefixLength 24
#Set the DNS Address
Set-DnsClientServerAddress -Interfacealias "vEthernet (ComputeClusterPublic)" -ServerAddress
192.168.40.10
#Get-NetAdapterHardwareInfo | Sort-Object -Property bus,Function
Get-VMNetworkAdapter -ManagementOS
Get-VMNetworkAdapterVLAN –ManagementOS
```

Network addresses

Table A-3 lists sample network address assignments for the Microsoft Hyper-V Private Cloud solution that is described in this paper.

Table A-3 Network addresses

Network	IP address or range
VLAN 30 Addresses (Management Cluster Private)	
Management Server1	192.168.30.21
Management Server2	192.168.30.22
VLAN 31 Addresses (Management Cluster Live Migration)	
Management Server1	192.168.31.21
Management Server2	192.168.31.22
VLAN 33 Addresses (SQL Cluster Private)	
SQL1	192.168.33.41
SQL2	192.168.33.42
VLAN 35 Addresses (SCVMM Cluster Private)	
SCVMM1	192.168.35.51
SCVMM2	192.168.35.52
VLAN 37 Addresses (Production Cluster Private)	
Hyper-V Compute Node1	192.168.37.31
Hyper-V Compute Node2	192.168.37.32
Hyper-V Compute Node3	192.168.37.33
Hyper-V Compute Node4	192.168.37.34
Hyper-V Compute Node5	192.168.37.35
Hyper-V Compute Node6	192.168.37.36
Hyper-V Compute Node7	192.168.37.37
Hyper-V Compute Node8	192.168.37.38
VLAN 38 Addresses (Production Cluster Live Migration)	
Hyper-V Compute Node1	192.168.38.31
Hyper-V Compute Node2	192.168.38.32
Hyper-V Compute Node3	192.168.38.33
Hyper-V Compute Node4	192.168.38.34
Hyper-V Compute Node5	192.168.38.35
Hyper-V Compute Node6	192.168.38.36
Hyper-V Compute Node7	192.168.38.37
Hyper-V Compute Node8	192.168.38.38

Network	IP address or range
VLAN 40 Addresses (Management/Cluster Public)	
AD Server	192.168.40.10
AD Server (ALT)	192.168.40.11
Management Server 1	192.168.40.21
Management Server 2	192.168.40.22
Management Cluster (FlexMgmt1)	192.168.40.29
SQL Server VM1	192.168.40.41
SQL Server VM2	192.168.40.42
SQL Cluster (SQLCluster)	192.168.40.49
SQL Service Manager Cluster (SCSMDB)	192.168.40.91
SQL Service Manager Data Warehouse Cluster (SCSMDW)	192.168.40.92
SQL Service Manager Analysis Cluster (SCSMAS)	192.168.40.93
SQL Service App Controller/Orchestrator/Sharepoint/WSUS Cluster (SCDB)	192.168.40.94
SQL Service Virtual Machine Manager Cluster (SCVMMDB)	192.168.40.95
SQL Service Operations Manager Cluster (SCOMDB)	192.168.40.96
SQL Service Operations Manager Data Warehouse Cluster (SCOMDW)	192.168.40.97
SCVMM Server VM1	192.168.40.51
SCVMM Server VM2	192.168.40.52
SCVMM Cluster (Cluster of the two VMs)	192.168.40.59
HAVMM Cluster (The role under the cluster)	192.168.40.58
System Center Operations Manager VM1	192.168.40.61
System Center Operations Manager VM2	192.168.40.62
System Center Ops Manager Reporting Server	192.168.40.66
System Center Orchestrator (Mgmt & Action Svr)	192.168.40.71
System Center Orchestrator (Supplemental Action Svr)	192.168.40.72
Service Manager (Change Management)	192.168.40.81
Service Manager (Data warehouse)	192.168.40.82
Service Manager (Portal)	192.168.40.83
System Center App Server	192.168.40.84
Optional Windows Deployment Server (& WSUS)	192.168.40.99
Hyper-V Production Server 1	192.168.40.31
Hyper-V Production Server 2	192.168.40.32
Hyper-V Production Server 3	192.168.40.33
Hyper-V Production Server 4	192.168.40.34

Network	IP address or range
Hyper-V Production Server 5	192.168.40.35
Hyper-V Production Server 6	192.168.40.36
Hyper-V Production Server 7	192.168.40.37
Hyper-V Production Server 8	192.168.40.38
Hyper-V Production Cluster (FlexCompute)	192.168.40.39
WSUS Server	192.168.40.101
WDS Server	192.168.40.231
NetApp FAS8040 – Controller1 - CIFS/SMB	192.168.40.89
NetApp FAS8040 – Controller2 – CIFS/SMB	192.168.40.90
NetApp FAS8040 SMI-S Provider	192.168.40.165
VLAN 50 Addresses (Production)	
Production VMs as needed	192.168.50.xx
VLAN 70 Addresses (Out of Band Mgmt Network)	
CMM	192.168.70.100
NetApp Cluster	192.168.70.150
NetApp Controller1 cDOT	192.168.70.151
NetApp Controller2 cDOT	192.168.70.152
NetApp Controller1 Service Processor	192.168.70.161
NetApp Controller2 Service Processor	192.168.70.162
NetApp SVM01	192.168.70.157
NetApp SVM02	192.168.70.159
CN4093 Switch1	192.168.70.120
CN4093 Switch2	192.168.70.121
Stacked CN4093	192.168.70.120
ITE-1 - ITE10	192.168.70.101 - 110
TFTP Server (to backup CN4093 switch)	192.168.70.xx
VLAN 30 Addresses (Management Cluster Private)	
Management Server1	192.168.30.21
Management Server2	192.168.30.22
VLAN 31 Addresses (Management Cluster Live Migration)	
Management Server1	192.168.31.21
Management Server2	192.168.31.22
VLAN 33 Addresses (SQL Cluster Private)	
SQL1	192.168.33.41

Network	IP address or range
SQL2	192.168.33.42
VLAN 35 Addresses (SCVMM Cluster Private)	
SCVMM1	192.168.35.51
SCVMM2	192.168.35.52
VLAN 37 Addresses (Production Cluster Private)	
Hyper-V Compute Node1	192.168.37.31
Hyper-V Compute Node2	192.168.37.32
Hyper-V Compute Node3	192.168.37.33
Hyper-V Compute Node4	192.168.37.34
Hyper-V Compute Node5	192.168.37.35
Hyper-V Compute Node6	192.168.37.36
Hyper-V Compute Node7	192.168.37.37
Hyper-V Compute Node8	192.168.37.38

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consulty our local Lenovo representative for information on the products and services currently available in yourarea. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service.

Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
1009 Think Place - Building One
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Trademarks

Lenovo, the Lenovo logo, and For Those Who Do are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. These and other Lenovo trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by Lenovo at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of Lenovo trademarks is available on the Web at http://www.lenovo.com/legal/copytrade.html.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

BladeCenter®	Lenovo®	vNIC™
Bootable Media Creator™	Lenovo(logo)®	
Flex System™	System x®	

The following terms are trademarks of other companies:

Intel, Intel Xeon, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Related publications

The publications that are listed in this section are considered particularly suitable for a more detailed discussion of the topics that are covered in this paper.

Lenovo Press publications

The following Lenovo Press publications provide additional information about the topic in this document.

- Product Guide on the Lenovo Flex System x240 M5 (E5-2600 v4)

  http://lenovopress.com/lp0093
- *Using System x Features on Demand*, REDP4895

  https://lenovopress.com/redp4895

Other publications and online resources

The following publications and websites are also relevant as further information sources:

- Flex System Updates:

  https://www-304.ibm.com/software/brandcatalog/puresystems/centre/update
- Flex System x240 Compute Node Installation and Service Guide:

  http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.8737.doc/dw1ko_book.pdf
- Flex System Chassis Management Module Installation Guide:

  http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.cmm.doc/dw1ku_cmm_ig_book.pdf
- Flex System Chassis Management Module User's Guide:

  http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.cmm.doc/dw1kt_cmm_ug_pdf.pdf
- Flex System Chassis Management Module Command-Line Interface Reference Guide:

  http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.cmm.doc/dw1ku_cmm_ig_book.pdf
- Flex System Power Guide:

  https://support.lenovo.com/documents/LNVO-POWINF
- Server Guide:

  https://support.lenovo.com/documents/LNVO-GUIDE
- Firmware update and best practices guide:

  https://support.lenovo.com/documents/MIGR-5082923
- Bootable Media Creator:

  https://support.lenovo.com/documents/LVNO-BOMC

- IBM Systems Director Agent Download (Platform Agent):

  http://ibm.com/systems/software/director/downloads/agents.html
- Upward Integration for Microsoft System Center Operations Manager:

  https://support.lenovo.com/documents/MIGR-5082204
- Upward Integration Modules for Microsoft System Center:

  https://support.lenovo.com/documents/MIGR-5087849
- Flex System CN4093 10Gb Converged Scalable Switch Application Guide:

  http://pic.dhe.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.networkdevices.doc/00ay506.pdf
- Flex System CN4093 10Gb Converged Scalable Switch ISCLI Reference Guide:

  http://pic.dhe.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.networkdevices.doc/00ay508.pdf
- Fast Setup:

  http://ibm.com/support/entry/portal/docdisplay?lndocid=TOOL-FASTSET
- Emulex OneCommand Utility:

  http://www.emulex.com/downloads/oem-qualified-downloads/ibm/vfafc-software-kits

Microsoft resources

For more information, see the following resources that are available from Microsoft:

- Microsoft Windows Server 2012 R2 Hyper-V:

  http://technet.microsoft.com/library/hh831531
- Microsoft System Center 2012 R2:

  http://www.microsoft.com/en-us/server-cloud/products/system-center-2012-r2/default.aspx
- IaaS Fabric Management Fast Track Deployment Guide:

  http://gallery.technet.microsoft.com/Infrastructure-as-a-ecf1cc0b
- IaaS Product Line Fabric Management Fast Track Architecture Guide:

  http://gallery.technet.microsoft.com/Infrastructure-as-a-ecf1cc0b

NetApp resources

For more information, see the following NetApp resources:

- NetApp Support Site:

  http://mysupport.netapp.com/
- NetApp Documentation, All products, main page:

  http://mysupport.netapp.com/documentation/productsatoz/index.html?tab=category
- NetApp Software download, main page:

  http://mysupport.netapp.com/NOW/cgi-bin/software/
- NetApp hardware information and specs (Hardware Universe):

  http://hwu.netapp.com/Home/Index
- NetApp Interoperability Matrix Tool (IMT):

  http://support.netapp.com/matrix/mtx/login.do

- NetApp Private Cloud Capacity Planning Guide:

  http://www.netapp.com/us/system/pdf-reader.aspx?m=tr-4329.pdf
- NetApp Sizing Tools (System Performance Modeler):

  https://spm.netapp.com
- NetApp Storage Cluster Configuration Worksheet:

  https://library.netapp.com/ecmdocs/ECMP1654308/html/GUID-425BDF6A-0C22-41E2-AC7E-CD34E3CC502D.html
- NetApp Volume Configuration Worksheet:

  https://library.netapp.com/ecmdocs/ECMP1610207/html/GUID-D0FFEE72-5A9B-4B5A-B4E4-7C018A349956.html
- NetApp SMB share configuration worksheet:

  https://library.netapp.com/ecmdocs/ECMP1610207/html/GUID-B1B02DE1-72BC-4B7B-9945-A6D886B92730.html

Clustered Data ONTAP resources

For more information, see the following Clustered Data ONTOP resources:

- NetApp Clustered Data ONTAP (8.2) documentation:

  http://mysupport.netapp.com/documentation/docweb/index.html?productID=61820&language=en-US
- System Administration Guide for Cluster Administrators:

  https://library.netapp.com/ecm/ecm_download_file/ECMP1368701
- NetApp Clustered Data ONTAP Network Management Guide:

  https://library.netapp.com/ecm/ecm_download_file/ECMP1636021
- NetApp TR-4129: Namespaces in Clustered Data ONTAP:

  http://www.netapp.com/us/media/tr-4129.pdf
- Clustered Data ONTAP File Access Management Guide for CIFS (SMB):

  https://library.netapp.com/ecm/ecm_get_file/ECMP1366834

Storage Design resources:

For more information, see the following Storage Design resources:

- SAN Administration Guide:

  https://library.netapp.com/ecm/ecm_download_file/ECMP1368525
- SAN Configuration Guide:

  https://library.netapp.com/ecm/ecm_download_file/ECMP1368692
- Logical Storage Management Guide:

  https://library.netapp.com/ecm/ecm_download_file/ECMP1368017
- Physical Storage Management Guide:

  https://library.netapp.com/ecm/ecm_download_file/ECMP1368404
- System Administration Guide for SVM Administrators:

  https://library.netapp.com/ecm/ecm_download_file/ECMP1368704

Add on Software and Microsoft Integration resources

For more information, see the following resources:

- NetApp TR-4175: Microsoft Windows Server 2012 Hyper-V Storage Performance:

  http://www.netapp.com/us/media/tr-4175.pdf
- NetApp SnapDrive TR-4228: SnapDrive 7.0 for Windows for Clustered Data ONTAP 8.2:

  http://www.netapp.com/us/media/tr-4228.pdf
- NetApp Best Practices for Microsoft Hyper-V over SMB 3.0 (TR-4172):

  http://www.netapp.com/us/media/tr-4172.pdf
- NetApp TR-4226: NetApp SnapManager 2.0 for Hyper-V on Clustered Data ONTAP 8.2

  http://www.netapp.com/us/media/tr-4226.pdf
- Data ONTAP DSM 4.1 for Windows – Installation and Administration Guide:

  https://library.netapp.com/ecm/ecm_download_file/ECMP1307349
- Download ONTAP MPIO DSM 4.1 Driver for Windows:

  http://mysupport.netapp.com/NOW/cgi-bin/software/?product=Multipath+I%2FO&platform=Windows
- NetApp OnCommand Plug-In 4.0 for Microsoft Best Practices Guide (TR-4244):

  http://www.netapp.com/us/media/tr-4244.pdf
- NetApp Download page for OnCommand Plug-In for Microsoft:

  http://mysupport.netapp.com/NOW/cgi-bin/software/?product=OnCommand+Plug-in+for+Microsoft&platform=Microsoft+SCOM
- NetApp TR-4271: Best Practices and Implementation Guide for NetApp SMI-S Agent 5.1

  http://www.netapp.com/us/media/tr-4271.pdf
- NetApp Download page for SMI-S:

  http://mysupport.netapp.com/NOW/cgi-bin/software/?product=Data+ONTAP+SMI-S+Agent&platform=Windows