



Deploying Flex System in a BladeCenter Environment

Introduces Flex System and its features and capabilities

Describes network, storage, and management integration aspects

Provides planning and deployment considerations

Discusses implementation details for Microsoft Hyper-V and VMware vSphere infrastructures

Ilya Krutov

Aram Avetisyan

Dusan Tekeljak





Deploying Flex System in a BladeCenter Environment

April 2015

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

Last update on April 2015

This edition applies to:

Flex System

BladeCenter

FastSetup version 3.2

Switch Center version 7.2.1

Upward Integration Modules for Microsoft System Center version 5.5 and 5.6

Upward Integration Modules for VMware vSphere version 3.5 and 3.6

© Copyright Lenovo 2015. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract

Contents

Notices	v
Trademarks	vi
Preface	vii
Authors	vii
Chapter 1. Introduction to Flex System	1
1.1 Flex System overview	2
1.1.1 Flex System Enterprise Chassis	2
1.1.2 Compute nodes	3
1.1.3 Expansion nodes	4
1.1.4 I/O modules	4
1.2 Flex System I/O architecture	5
Chapter 2. Planning for Flex System integration	9
2.1 Server integration	10
2.2 Network integration	10
2.2.1 I/O modules	10
2.2.2 Network adapters	12
2.3 Storage integration	13
2.3.1 Fibre Channel	13
2.3.2 FCoE	16
2.3.3 iSCSI	17
2.4 Management integration	18
2.4.1 Managing hardware	18
2.4.2 Managing network switches	20
2.4.3 Managing FC SAN fabric	21
2.4.4 Managing physical and virtualized environments	21
Chapter 3. Adding Flex System to a BladeCenter environment	23
3.1 Hardware deployment	24
3.1.1 Initial Flex System chassis configuration	24
3.1.2 Firmware updates and basic configuration by using FastSetup	26
3.1.3 Configuring Active Directory Integration for CMM	33
3.2 Networking and storage deployment	34
3.2.1 Configuring the network	34
3.2.2 Enabling UFP in server's UEFI	41
3.2.3 Configuring storage	42
3.3 Integrating VMware vSphere	48
3.3.1 USB memory key with ESXi	48
3.3.2 Configure management network for VMware ESXi	49
3.3.3 Joining hypervisors to the existing environment	51
3.4 Integrating Microsoft Hyper-V	58
3.4.1 BladeCenter environment	58
3.4.2 Deployment considerations	59
3.4.3 Connecting to a Flex System compute node using Remote Console	60
3.4.4 Installing the operating system	61
3.4.5 Management network configuration	65
3.4.6 Joining the node to System Center Virtual Machine Manager	68

3.4.7 VM live migration compatibility	75
Chapter 4. Managing a combined Flex System and BladeCenter environment	77
4.1 Managing a vSphere environment with UIM	78
4.1.1 Enabling UIMs for a newly added ESXi host.	78
4.1.2 Collecting system inventory with UIM	79
4.1.3 Monitoring hardware status.	84
4.2 Using PFA alert to move VMs to another ESXi host	85
4.2.1 Rolling firmware upgrades	89
4.2.2 Changing IMM and UEFI configuration.	91
4.3 Managing a Windows Server environment with UIM.	95
4.3.1 Enabling Hardware Monitoring on the newly deployed Flex System	96
4.3.2 Deploying System Center agents for hardware monitoring and inventory	102
4.3.3 Monitoring hardware status in SCOM.	113
4.3.4 Lenovo Hardware Performance and Resource Optimization Pack for VMM	117
4.3.5 Rolling firmware upgrades by using UIM for System Center VMM.	119
4.3.6 Publishing System Firmware to SCCM server	123
4.3.7 Inventory collection	126
Abbreviations and acronyms	129
Related publications	131
Lenovo Press publications	131
Online resources	131

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 your area. 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®	RackSwitch™	ToolsCenter™
Flex System™	Lenovo(logo)®	UpdateXpress System Packs™
Lenovo®	ServerGuide™	vNIC™
Omni Ports™	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.

Preface

Flex System™ is a next-generation blade platform that is ideally suited to data center environments that require flexible, cost-effective, secure, and energy-efficient server infrastructure.

This Lenovo® Press paper highlights the key features and capabilities of Flex System and provides planning and deployment considerations on how Flex System can be integrated into an existing BladeCenter® environment. Server, network, storage, and management integration aspects also are described.

This paper is intended for those IT professionals who want to learn more about how to integrate Flex System into an existing BladeCenter environment.

Authors

This paper was produced in collaboration with Lenovo Press by a team of subject matter experts from around the world.



Ilya Krutov is a Project Leader at Lenovo Press. He manages and produces pre-sale and post-sale technical publications for various IT topics, including x86 rack and blade servers, server operating systems and software, virtualization and cloud, and datacenter networking. Ilya has more than 15 years of experience in the IT industry, performing various roles, including Team Leader, Portfolio Manager, Brand Manager, IT Specialist, and Certified Instructor. He has written more than 200 books, papers, and other technical documents. He has a Bachelor's degree in Computer Engineering from the Moscow Engineering and Physics Institute (Technical University).



Aram Avetisyan is an IT Specialist with more than 14 years experience. He has expertise in various IT-related fields, including virtualization, operating systems administration, and disaster recovery. He joined IBM in 2011 and became one of the leading virtualization specialists. He holds several industry-level certifications, including VMware Certified Professional - Data Center Design (VCAP5-DCD) and VMware Certified Professional - Data Center Administration (VCP5-DCA). Aram is a VMware Certified Instructor (VCI) and as a VCI, he delivers VMware-authorized courses to IBM employees. Aram is a member of VMware Center of Excellence and was awarded vEXPERT 2014 accreditation.



Dusan Tekelj is an IT Specialist at IBM VMware Center of Excellence in Brno, Czech Republic. Dusan has over 5 years of experience in virtualization. He has expertise in server operating systems, networking, and storage. He works on design and integration projects, including various VMware vSphere on Flex System implementations. Before joining the VMware Center of Excellence at IBM in 2012, Dusan worked with Microsoft cloud technologies. He has a master degree in engineering from Slovak University of Technology and holds several IT industry-leading certifications, including VCAP-DCD, VCAP-DCA, and MCITP.

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

Mark Cadiz
Bruce Corregan
David Tareen
David Watts
Lenovo

Bryan Davis
Karen Lawrence
IBM ITSO



Introduction to Flex System

Lenovo Flex System is a next-generation platform that is ideally suited to data center environments that require flexible, cost-effective, secure, and energy-efficient server infrastructure.

This chapter introduces Flex System and describes its key features, components, and I/O architecture. This chapter includes the following topics:

- ▶ 1.1, “Flex System overview” on page 2
- ▶ 1.2, “Flex System I/O architecture” on page 5

1.1 Flex System overview

The innovative design features of the Flex System products make it possible to configure integrated, customized, highly secure solutions. These solutions meet data center needs and provide for flexible expansion capabilities. The scalable hardware features and the unprecedented power and cooling capabilities of the Flex System components help optimize hardware and power usage, minimize operational costs, and simplify the overall management of the data center.

The following Flex System components are described in this section:

- ▶ 1.1.1, “Flex System Enterprise Chassis”
- ▶ 1.1.2, “Compute nodes” on page 3
- ▶ 1.1.3, “Expansion nodes” on page 4
- ▶ 1.1.4, “I/O modules” on page 4

1.1.1 Flex System Enterprise Chassis

The Flex System Enterprise Chassis is the foundation of the Flex System offering, which features 14 standard (half-width) Flex System form factor compute node bays in a 10U chassis. This offering delivers high-performance connectivity for your integrated compute, storage, networking, and management resources.

Up to a total of 28 independent servers can be accommodated in each Enterprise Chassis, if high-density x222 compute nodes are deployed.

The chassis is designed to support multiple generations of technology. It also offers independently scalable resource pools for higher usage and lower cost per workload.

With the ability to handle up to 14 standard form factor nodes, the Enterprise Chassis provides flexibility and tremendous compute capacity in a 10U package. Also, the rear of the chassis accommodates four high-speed I/O bays that can accommodate up to 40 GbE high-speed networking, 16 Gb Fibre Channel, or 56 Gb InfiniBand. With interconnecting compute nodes, networking, and storage through a high performance and scalable mid-plane, the Enterprise Chassis can support the latest high speed networking technologies.

The ability to support the demands of tomorrow’s workloads is built in with a new I/O architecture, which provides choice and flexibility in fabric and speed. With the ability to use Ethernet, InfiniBand, Fibre Channel (FC), Fibre Channel over Ethernet (FCoE), and iSCSI, the Enterprise Chassis is uniquely positioned to meet the growing and future I/O needs of large and small businesses.

Power and cooling resources are integrated into a chassis, and they are shared by compute, storage, I/O, and management components, which increases power and cooling efficiency.

Figure 1-1 shows Lenovo Flex System Enterprise Chassis.



Figure 1-1 Lenovo Flex System Enterprise Chassis

1.1.2 Compute nodes

Flex System offers compute nodes that vary in architecture, dimension, and capabilities.

Optimized for efficiency, density, performance, reliability, and security, the portfolio includes a range of Intel Xeon based nodes that are designed to make full use of the full capabilities of these processors that can be mixed within the same Enterprise Chassis.

The following Intel based nodes compute nodes are available, which range from two-socket to eight-socket Intel processor families:

- ▶ Intel Xeon processor E5-2400 product family (x220 and x222 Compute Nodes)
- ▶ Intel Xeon processor E5-2600 v2 product families (x240 Compute Nodes)
- ▶ Intel Xeon processor E5-2600 v3 product family (x240 M5 Compute Nodes)
- ▶ Intel Xeon processor E5-2600 v2 product families (x440 Compute Nodes)
- ▶ Intel Xeon processor E7-8800 v2, E7-4800 v2, and E7-2800 v2 product families (x880 X6, x480 X6, and x280 X6 Compute Nodes)

A standard form-factor Flex System x240 Compute Node is shown on the left in Figure 1-2. A full-wide Flex System x440 Compute Node is shown on the right in Figure 1-2.



Figure 1-2 Flex System x240 (left) and x440 (right) Compute Nodes

The nodes are complemented with leadership I/O capabilities of up to 16 channels of high-speed I/O lanes per standard node bay and 32 lanes per full wide node. Various I/O adapters and matching I/O modules are available.

1.1.3 Expansion nodes

Expansion nodes can be attached to certain standard form factor (half-width) Flex System compute nodes, which allows the expansion of the nodes' capabilities with locally attached storage or PCIe adapters.

The Flex System Storage Expansion Node provides locally attached disk expansion to the x240 and x220. SAS and SATA disk are supported.

With the attachment of the Flex System PCIe Expansion Node, an x220 or x240 can have up to four PCIe adapters attached. High performance GPUs can also be installed within the PCIe Expansion Node from companies, such as Intel and NVIDIA.

Figure 1-3 shows the x240 Compute Node with the PCIe Expansion Unit (left) and the Storage Expansion Unit (right) attached.




Figure 1-3 x240 with PCIe Expansion Unit (left) or Storage Expansion Unit (right)

1.1.4 I/O modules

The range of available modules and switches to support key network protocols can be used to configure Flex System to fit in your infrastructure. However, you can do so without sacrificing the ability to be ready for the future. The networking resources in Flex System are standards-based, flexible, and fully integrated into the system. This combination gives you no-compromise networking for your solution. Network resources are virtualized and managed by workload. These capabilities are automated and optimized to make your network more reliable and simpler to manage.

Flex System gives you the following key networking capabilities:

- ▶ Supports the networking infrastructure that you have today, including Ethernet, FC, FCoE, iSCSI, and InfiniBand.
- ▶ Offers industry-leading performance with 1 Gb, 10 Gb, and 40 Gb Ethernet, 8 Gb and 16 Gb Fibre Channel, and QDR and FDR InfiniBand.
- ▶ Provides pay-as-you-grow scalability so you can add ports and bandwidth, when needed.

Networking in data centers is undergoing a transition from a discrete traditional model to a more flexible, optimized model. The network architecture in Flex System addresses the key challenges that customers are facing today in their data centers. The key focus areas of the network architecture on this platform are unified network management, optimized and automated network virtualization, and simplified network infrastructure.

Providing innovation, leadership, and choice in the I/O module portfolio uniquely positions Flex System to provide meaningful solutions to address customer needs.

As an example of the I/O module, the Flex System Fabric EN4093R 10Gb Scalable Switch is shown in Figure 1-4.



Figure 1-4 Lenovo Flex System Fabric EN4093R 10Gb Scalable Switch

1.2 Flex System I/O architecture

The I/O architecture of the Flex System Enterprise Chassis includes an array of connectivity options for server nodes that are installed in the enclosure. Users can decide to use a local switching model that provides superior performance, cable reduction, and a rich feature set. The pass-through technology also can be used, which allows all network switching decisions to be made external to the Enterprise Chassis.

By far, the most versatile option is to use modules that provide local switching capabilities and advanced features that are fully integrated into the operation and management of the Enterprise Chassis.

From a physical I/O module bay perspective, the Enterprise Chassis has four I/O bays in the rear of the chassis. The physical layout of these I/O module bays is shown in Figure 1-5.




Figure 1-5 Rear view of the Enterprise Chassis showing I/O module bays

If a node has a two-port integrated LAN on Motherboard (LOM) as standard, modules 1 and 2 are connected to this LOM. If an I/O adapter is installed in the node's I/O expansion slot 1, modules 1 and 2 are connected to this adapter.

Modules 3 and 4 connect to the I/O adapter that is installed in the I/O expansion slot 2 on the node.

These I/O modules provide external connectivity and connect internally to each of the nodes within the chassis. They can be Switch or Pass-thru modules, with a potential to support other types in the future.

From a midplane wiring perspective, the Enterprise Chassis provides 16 lanes out of each half-wide node bay (toward the rear I/O bays) with each lane capable of 16 Gbps or higher speeds. How these lanes are used is a function of which adapters are installed in a node, which I/O module is installed in the rear, and which port licenses are enabled on the I/O module.

How the midplane lanes connect between the node bays upfront and the I/O bays in the rear is shown in Figure 1-6.

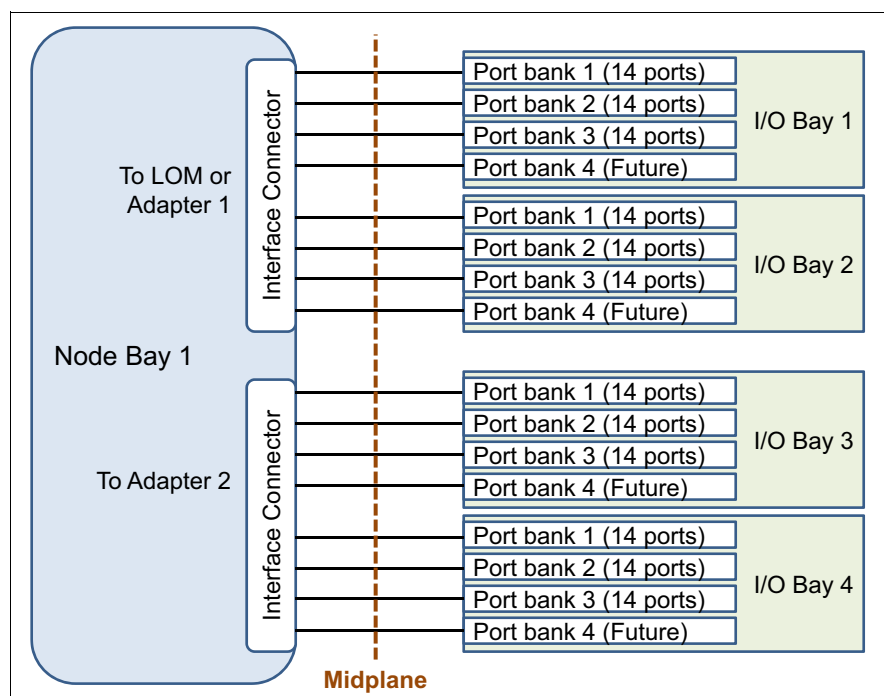


Figure 1-6 A total of 16 lanes of a single half-wide node bay toward the I/O bays

Currently available I/O modules offer one, two, or three banks of 14 internal compute node-facing ports. Each port bank corresponds to a certain port of the adapter that is installed in the compute node. For example, if a node were installed with only the dual port LOM adapter, only two of the 16 lanes are used (one to each of the port banks 1 of the I/O modules that are installed in I/O bays 1 and 2), as shown in Figure 1-7.

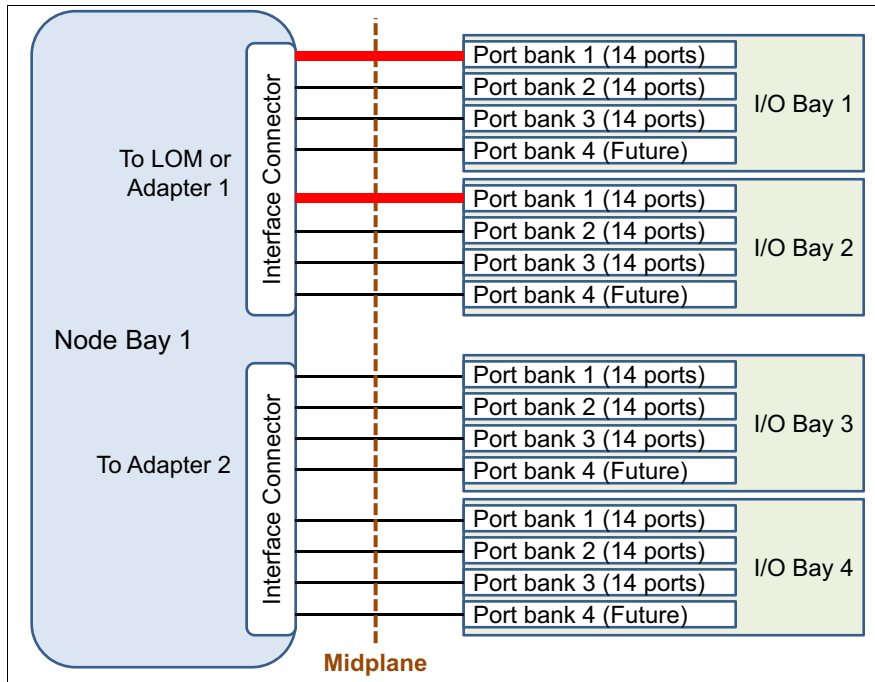


Figure 1-7 Dual port LOM connecting to ports on I/O bays 1 and 2 (all other lanes unused)

If a node was installed with a quad port adapter, four of the 16 lanes are used (one to each of the port banks 1 and 2 of the I/O modules that are installed in I/O bays 1 and 2), as shown in Figure 1-8.

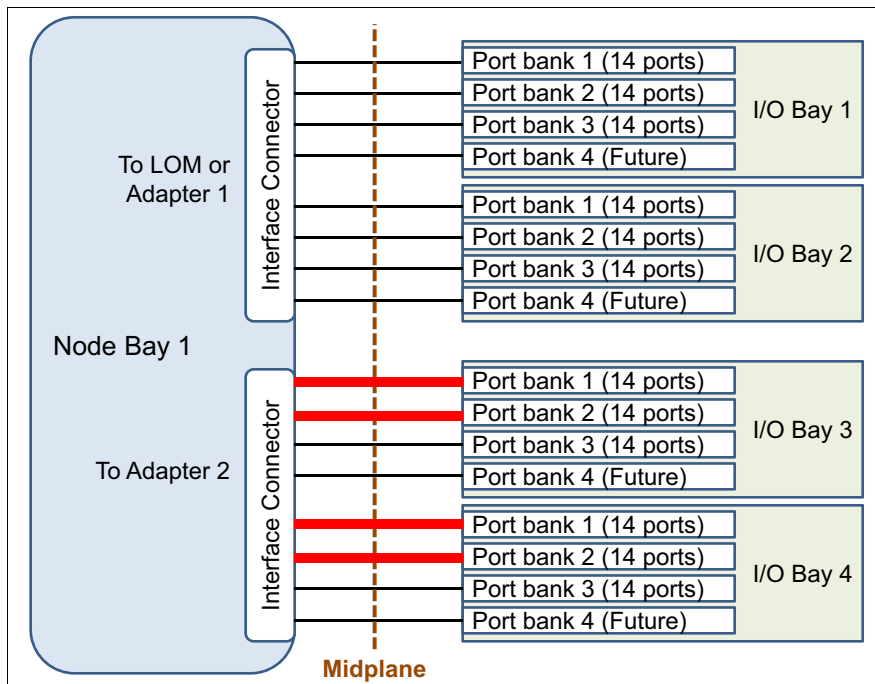


Figure 1-8 Quad-port adapter connecting to ports on I/O bays 1 and 2 (all other lanes unused)

All I/O modules include a base port bank 1 of 14 downstream ports, with the pass-through module supporting only the single set of 14 server facing ports. The Ethernet switching and interconnect I/O modules support more than the base set of ports (up to three port banks,

depending on the I/O module), and the ports in the banks are enabled by the port licenses. FC SAN switches support up to two internal port banks (depending on the I/O module) and the ports in the banks are enabled by the port licenses.

By default, each I/O module ships with certain number of port licenses (depending on the I/O module type). More port licenses can be purchased via Features on Demand (FoD) upgrades. Depending on the I/O module and the software version that is running on the module, the port licenses can be tied to the port bank (for example, port bank 1 in the EN2092, EN4093R, and CN4093 switches and SI4093 interconnect modules that are running Networking OS 7.7 or earlier). They also can be dynamically or statically assigned by the user to the ports in any internal bank and external ports (such as Flexible Port Mapping in the EN2092, EN4093R, and CN4093 switches and SI4093 interconnect modules that are running Networking OS 7.8 or later or Dynamic Port on Demand in the FC5022 switches).

As of this writing, there are limits on the port density of the current I/O modules in that only the first three lanes potentially are available from the I/O module.



Planning for Flex System integration

One of the most important goals of integrating Flex System into an existing BladeCenter environment is to make sure that the existing hardware and software technologies, tools, and applications can be seamlessly deployed and used in the combined environment with no or minimal disruption to the existing processes and services.

Specifically, server, network, storage, and management integration aspects should be addressed. These aspects are described in this chapter.

This chapter includes the following topics:

- ▶ 2.1, “Server integration” on page 10
- ▶ 2.2, “Network integration” on page 10
- ▶ 2.3, “Storage integration” on page 13
- ▶ 2.4, “Management integration” on page 18

2.1 Server integration

This section describes the following Flex System Compute Node selection considerations that are based on your existing blade servers:

- ▶ If you use dual-socket general-purpose HS22, HS22V, or HS23 blade servers, consider adding Flex System x240 or x240 M5 Compute Nodes.

Also, if certain existing applications require scalability that goes beyond two sockets but they do not need enterprise-class hardware features, the four-socket Flex System x440 Compute Node can be used.

- ▶ If you use two- or four-socket HX5 enterprise blade servers, consider the use of Flex System x280, x480, or x880 Compute Node that can be scaled up to eight sockets.
- ▶ If you use entry-level HS12 or HS23E blade servers, consider the use of Flex System x220 or high-density x222 Compute Nodes.

2.2 Network integration

This section describes network integration aspects and provides I/O module and network adapter selection considerations. The following topics are covered:

- ▶ 2.2.1, “I/O modules”
- ▶ 2.2.2, “Network adapters” on page 12

2.2.1 I/O modules

Selecting the Flex System Ethernet I/O module that is a best fit for your existing BladeCenter environment is a process that is unique to each client. The following factors should be considered when you are deciding which Ethernet module is best for a specific environment:

- ▶ If your BladeCenter network is built with Layer 2/3 1 Gb or 1/10 Gb Uplink switches and you plan to continue to use 1 GbE server connectivity, consider the use of Flex System EN2092 1Gb Ethernet Switch for your Flex System solution. EN2092 supports 1 GbE connectivity to the compute nodes and 1 GbE or 10 GbE connections to an upstream network. In addition, EN2092 offers Layer 2/3 network features that are identical to those features in L2/3 GbE switches for BladeCenter. It also can be managed by using the Switch Center or a third-party SNMP-based network management application.

If you plan to use 10 GbE or mixed 1 GbE and 10 GbE server connectivity in your Flex System chassis, consider the use of Flex System Fabric EN4093R or CN4093 10Gb Scalable Switches or SI4093 or SI4091 System Interconnect Module.

- ▶ If your BladeCenter network is built with Virtual Fabric 10Gb Switch Modules, consider the use of Flex System Fabric EN4093R or CN4093 10Gb Scalable Switches or SI4093 or SI4091 System Interconnect Module.
- ▶ If you use pass-through network connectivity with Intelligent Pass-Thru Module or 10Gb Ethernet Pass-Thru Module, consider selecting Flex System EN4091 10Gb Ethernet Pass-Thru Module that supports 1 GbE and 10 GbE speeds.
- ▶ If you use Server Connectivity Module, consider selecting Flex System SI4093 or SI4091 System Interconnect Module. Alternatively, if you plan to use advanced Layer 2/3 features in the future, consider selecting EN4093R 10Gb Scalable Switch that is running in the “easy connect” mode.

The following considerations are important when you are selecting between the EN4093R 10Gb Scalable Switch, the CN4093 10Gb Converged Scalable Switch, and the SI4093 or SI4091 System Interconnect Modules:

- ▶ If you require Fibre Channel Forwarder (FCF) services within the Enterprise Chassis or native Fibre Channel uplinks from the 10 Gb switch, the CN4093 10Gb Converged Scalable Switch is the correct choice.
- ▶ If you do not require FCF services or native Fibre Channel ports on the 10 Gb switch, but need the maximum number of 10 Gb uplinks without purchasing an extra license, support for FCoE transit capabilities, and the most feature-rich solution, the EN4093R 10Gb Scalable Switch is a good choice.
- ▶ If you require ready for use not apparent operation (minimal to no configuration on the switch) and do not need any L3 support or other advanced features (and know that there is no need for more advanced functions), the SI4093 or SI4091 System Interconnect Modules are a potential choice.

When you are selecting switches, there are often many criteria that are involved because each environment has its own unique attributes. The criteria that are listed in Table 2-1 are a good starting point in the decision-making process.

Table 2-1 Switch module selection criteria

Suitable switch module	I/O modules				
Requirement	EN2092	SI4091	SI4093	EN4093R	CN4093
Gigabit Ethernet to nodes	Yes	Yes	Yes	Yes	Yes
10 Gb Ethernet to nodes	No	Yes	Yes	Yes	Yes
10 Gb Ethernet uplinks	Yes	Yes	Yes	Yes	Yes
40 Gb Ethernet uplinks	No	No	Yes	Yes	Yes
Basic Layer 2 switching	Yes	Yes	Yes	Yes	Yes
Advanced Layer 2 switching: IEEE features (STP, QoS)	Yes	No	No	Yes	Yes
Layer 3 switching (forwarding, routing, ACL filtering)	Yes	No	No	Yes	Yes
10 Gb Ethernet CEE	No	Yes	Yes	Yes	Yes
FCoE FIP Snooping Bridge support	No	Yes	Yes	Yes	Yes
FCF support	No	No	No	No	Yes
Native FC port support	No	No	No	No	Yes
Switch stacking	No	No	No	Yes	Yes
802.1Qbg Edge Virtual Bridge support	No	No	Yes	Yes	Yes
vLAG support	No	No	No	Yes	Yes
UFP support	No	No	Yes	Yes	Yes
Virtual Fabric mode vNIC™ support	No	No	No	Yes	Yes
Switch independent mode vNIC support	No	Yes	Yes	Yes	Yes
SPAR support	No	No	Yes	Yes	Yes
Openflow support	No	No	No	Yes	No

2.2.2 Network adapters

This section describes the selection considerations for the network adapters based on your existing BladeCenter server configurations.

The following topics are covered:

- ▶ “Broadcom technology-based network adapters”
- ▶ “Mellanox technology-based network adapters”
- ▶ “Emulex technology-based network adapters” on page 12

Compatibility: This section describes general considerations for selecting an adapter that are based on technology requirements. For more information about the compatibility information between adapters and compute nodes, see the Flex System Interoperability Guide, which is available at this website:

<http://lenovopress.com/fsig>

Broadcom technology-based network adapters

If you are standardized on Broadcom technology for your network adapters, the following integration choices are available:

- ▶ 1 GbE connectivity

If you plan to use up to two 1 GbE ports and you selected the x220 Compute Nodes, the x220s have the embedded Broadcom based dual-port 1 GbE controller.

If the compute node you selected does not have embedded Broadcom NICs (such as x240, x440, and X6 compute nodes), consider the use of EN2024 4-port 1Gb Ethernet Adapter. This option gives you up to four 1 GbE ports per compute node.

If you used up to eight 1 GbE ports on your blade servers, consider installing two EN2024 adapters into the compute node for a total of eight 1 GbE ports.

Alternatively, consider the use of CN4022 2-port 10Gb Converged Adapter on which each 10 Gb port can be divided into four virtual NICs with configurable bandwidth.

- ▶ 10 GbE connectivity

If you used Broadcom 10Gb Virtual Fabric Adapters in your BladeCenter infrastructure, consider the use of CN4022 2-port Converged Adapter.

Mellanox technology-based network adapters

If you are standardized on Mellanox technology for your 10 GbE network connectivity with Mellanox 10Gb Ethernet Adapters for BladeCenter, consider the use of the EN4132 2-port 10Gb Ethernet network adapters in your Flex System solution.

Emulex technology-based network adapters

If you are standardized on Emulex technology for your network adapters, the following integration choices are available:

- ▶ 1 GbE connectivity

If you used integrated 1 GbE ports on the HS23, you can select x240 models with integrated dual-port LOM that supports 1 GbE and 10 GbE speeds.

- ▶ 10 GbE connectivity

If your existing blades have dual-port 10 Gb Emulex Virtual Fabric Adapters, consider selecting corresponding Flex System Compute Node offerings that have integrated dual-port 10 GbE LOM. If there are no models with integrated LOM, consider installing CN4052 10Gb Virtual Fabric Adapter in the compute node.

If your existing blade servers use four 10 GbE ports (HS23 with VFAs installed), consider the use of CN4054 or CN4054R 10Gb Virtual Fabric Adapter. By using this option, four 10 GbE ports are available.

If you need more than four 10 GbE ports, consider the use of CN4058S 10Gb Virtual Fabric Adapter which gives you eight 10 GbE ports. Currently, only six of eight ports can be used.

2.3 Storage integration

This section describes storage integration aspects that must be considered. The following topics are covered:

- ▶ 2.3.1, “Fibre Channel” on page 13
- ▶ 2.3.2, “FCoE” on page 16
- ▶ 2.3.3, “iSCSI” on page 17

2.3.1 Fibre Channel

Fibre Channel (FC) is a proven and reliable network for storage interconnect. The Flex System Enterprise Chassis FC portfolio offers various choices to meet your needs and interoperate with existing SAN infrastructure.

FC requirements

If Enterprise Chassis is integrated into FC storage fabric, ensure that the following requirements are met. Check the compatibility guides from your storage system vendor for confirmation:

- ▶ Enterprise Chassis server hardware and host bus adapter (HBA) are supported by the storage system. For more information, see the System Storage Interoperation Center (SSIC) or the third-party storage system vendors support matrixes.
- ▶ The FC fabric that is used or proposed for use is supported by the storage system.
- ▶ The operating systems that are deployed are supported by the compute nodes and storage system.
- ▶ Multipath drivers exist and are supported by the operating system and storage system (in case you plan for redundancy).
- ▶ Clustering software is supported by the storage system (in case you plan to implement clustering technologies).

If any of these requirements are not met, consider another solution that is supported.

Almost every vendor of storage systems or storage fabrics has extensive compatibility matrixes that include supported HBAs, SAN switches, and operating systems. For more information about IBM System Storage compatibility, see the System Storage Interoperability Center at this website:

<http://www.ibm.com/systems/support/storage/config/ssic>

FC switch selection and fabric interoperability rules

Flex System Enterprise Chassis provides integrated FC switching functions by using the following switch options:

- ▶ Flex System FC3171 8Gb SAN Switch
- ▶ Flex System FC3171 8Gb SAN Pass-thru
- ▶ Flex System FC5022 16Gb SAN Scalable Switch

If your existing storage network is standardized on the Brocade technology, consider selecting FC5022 16Gb SAN Scalable Switches.

If you use QLogic technology-based switches or pass-through or both in your existing BladeCenter infrastructure, consider selecting FC3171 8Gb SAN Switch or pass-through modules.

Considerations for the FC5022 16Gb SAN Scalable Switch

The module can function in Fabric OS Native mode or Brocade Access Gateway mode. The switch ships with Fabric OS mode as the default. The mode can be changed by using operating system commands or web tools.

Access Gateway simplifies SAN deployment by using N_Port ID Virtualization (NPIV). NPIV provides FC switch functions that improve switch scalability, manageability, and interoperability.

Considerations for the FC3171 modules

These I/O modules provide seamless integration of Flex System Enterprise Chassis into existing Fibre Channel fabric. They avoid any multivendor interoperability issues by using NPIV technology.

All ports are licensed on both of these switches (there are no port licensing requirements). The I/O module has 14 internal ports and 6 external ports that are presented at the rear of the chassis.

Attention: If you need Full Fabric capabilities at any time in the future, purchase the Full Fabric Switch Module (FC3171 8Gb SAN Switch) instead of the pass-through module (FC3171 8Gb SAN Pass-thru). The pass-through module never can be upgraded.

You can reconfigure the FC3171 8Gb SAN Switch to become a pass-through module by using the switch GUI or command-line interface (CLI). The module can be converted back to a full function SAN switch at any time. The switch requires a reset when you turn on or off transparent mode.

Select a SAN module that can provide the required functionality with seamless integration into the existing storage infrastructure, as shown in Table 2-2 on page 15. There are no strict rules to follow during integration planning. However, several considerations must be taken into account.

Almost all switches support interoperability standards, which means that almost any switch can be integrated into existing fabric by using interoperability mode. Interoperability mode is a special mode that is used for integration of different vendors' FC fabrics into one. However, only standards-based functionality is available in the interoperability mode. Advanced features of a storage fabric's vendor might not be available.

Brocade, QLogic, and Cisco have interoperability modes on their fabric switches. Check the compatibility matrixes for a list of supported and unsupported features in the interoperability mode.

Table 2-2 provides a high-level overview of standard and advanced functions that are available for particular Enterprise Chassis SAN switches. It lists how these switches might be used for designing new storage networks or integrating with existing storage networks.

Table 2-2 SAN module feature comparison and interoperability

	FC5022 16Gb SAN Scalable Switch	FC3171 8Gb SAN Switch	FC5022 16Gb SAN Scalable Switch in Brocade Access Gateway mode	FC3171 8Gb SAN Pass-thru (and FC3171 8Gb SAN Switch in pass-through mode)
Basic FC connectivity				
FC-SW-2 interoperability	Yes ^a	Yes	Not applicable	Not applicable
Zoning	Yes	Yes	Not applicable	Not applicable
Maximum number of Domain IDs	239	239	Not applicable	Not applicable
Advanced FC connectivity				
Port Aggregation	Yes	No ^b	Not applicable	Not applicable
Advanced fabric security	Yes	Yes	Not applicable	Not applicable
Interoperability (existing fabric)				
Brocade fabric interoperability	Yes	No	Yes	Yes
QLogic fabric interoperability	No	No	No	No
Cisco fabric interoperability	No	No	Yes	Yes

a. Indicates that a feature is supported without any restrictions for existing fabric, but with restrictions for added fabric, and vice versa.

b. Does not necessarily mean that a feature is not supported. Instead, it means that severe restrictions apply to the existing fabric. Some functions of the existing fabric potentially must be disabled (if used).

Remember: Advanced (proprietary) FC connectivity features from different vendors might be incompatible with each other, even those features that provide almost the same function. For example, Brocade and Cisco support port aggregation. However, Brocade uses ISL trunking and Cisco uses PortChannels, which are incompatible with each other.

For example, if you integrate FC3052 2-port 8Gb FC Adapter (Brocade) into QLogic fabric, you cannot use Brocade proprietary features, such as ISL trunking. However, QLogic fabric does not lose functionality. Conversely, if you integrate QLogic fabric into existing Brocade fabric, placing all Brocade switches in interoperability mode loses Advanced Fabric Services functions.

If you plan to integrate Enterprise Chassis into an FC fabric that is not listed here, QLogic might be a good choice. However, this configuration is possible with interoperability mode only, so extended functions are not supported. A better way is to use the FC3171 8Gb SAN Pass-thru or Brocade Access Gateway.

Switch selection and interoperability have the following rules:

- FC3171 8Gb SAN Switch is used when Enterprise Chassis is integrated into existing QLogic fabric or when basic FC functionality is required; that is, with one Enterprise Chassis with a direct-connected storage server.

- FC5022 16Gb SAN Scalable Switch is used when Enterprise Chassis is integrated into existing Brocade fabric or when advanced FC connectivity is required. You might use this switch when several Enterprise Chassis are connected to high-performance storage systems.

If you plan to use advanced features, such as ISL trunking, you might need to acquire specific licenses for these features.

Tip: The use of FC storage fabric from the same vendor often avoids possible operational, management, and troubleshooting issues.

If Enterprise Chassis is attached to a storage system, support is provided by the storage system's vendor. For more information about supported configurations, see the vendor's specific compatibility information.

For more information about IBM storage compatibility, see the System Storage Interoperation Center at this website:

<http://ibm.com/systems/support/storage/ssic>

2.3.2 FCoE

One common way to reduce administration costs is by converging technologies that are implemented on separate infrastructures. FCoE removes the need for separate Ethernet and FC HBAs on the servers. Instead, a Converged Network Adapter (CNA) is installed in the server.

Although Lenovo does not mandate the use of FCoE, the choice of using separate Ethernet and SAN switches inside the chassis or choosing a converged FCoE solution is left up to the client. Flex System offers both connectivity solutions.

A CNA presents what appears to be an NIC and an HBA to the operating system, but the output out of the node is 10 Gb Ethernet. The adapter can be the integrated 10Gb LOM with FCoE upgrade applied, or it can be a converged adapter 10 Gb, such as the CN4054R 10Gb Virtual Fabric Adapter that includes FCoE.

The CNA is then connected via the chassis midplane to an internal switch that passes these FCoE packets onwards to an external switch that contains a Fibre Channel Forwarder (where the FC is "broken out", such as the EN4093R), or by using a switch that is integrated inside the chassis that includes an FC Forwarder. Such a switch is the CN4093 10Gb Converged Scalable Switch, which can break out FC and Ethernet to the rear of the Flex System chassis. The CN4093 10Gb Converged Scalable Switch has external Omni Ports™ that can be configured as FC or Ethernet.

If your BladeCenter uses FCoE storage connectivity, you have the following options for your prospective Flex System solution:

- If you use Virtual Fabric 10Gb Switch Module as a transit FCoE switch that is connected to the upstream FCF switch, you can use EN4093R as an FCoE transit switch that is connected to the same upstream network. For more information about the compatibility of the solution, consult with your storage system vendor.
- If you use Virtual Fabric 10Gb Switch Module as a transit FCoE switch that is connected to the QLogic FC Gateway Module that is installed in the BladeCenter H chassis, you can use CN4093 as a Full Fabric FCoE switch with native FC ports that can be connected to the existing storage system. For more information about the compatibility of the solution, consult with your storage system vendor.

- If you use 10Gb Ethernet pass-through Module for BladeCenter, you can select Flex System EN4091 10Gb Ethernet pass-through and connect your Flex System solution to the existing upstream network if you have the sufficient number of network ports. For more information about the compatibility of the solution, consult with your storage system vendor.

For information about IBM storage compatibility, see the System Storage Interoperation Center at this website:

<http://ibm.com/systems/support/storage/ssic>

2.3.3 iSCSI

iSCSI uses a traditional Ethernet network for block I/O between storage system and servers. Servers and storage systems are connected to the LAN and use iSCSI to communicate with each other. Because iSCSI uses a standard TCP/IP stack, you can use iSCSI connections across LAN or wide area network (WAN) connections.

The software iSCSI initiator is specialized software that uses a server's processor for iSCSI protocol processing. A hardware iSCSI initiator exists as microcode that is built in to the LAN on Motherboard (LOM) on the node or on the I/O Adapter if it is supported.

Software and hardware initiator implementations provide iSCSI capabilities for Ethernet NICs. However, an operating system driver can be used only after the locally installed operating system is turned on and running. In contrast, the NIC built-in microcode is used for boot-from-SAN implementations, but cannot be used for storage access when the operating system is already running.

The iSCSI compatibility information normally lists support only for iSCSI storage that is attached by using hardware iSCSI offload adapters in the servers. Flex System compute nodes support any type of iSCSI (1 Gb or 10 Gb) storage if the software iSCSI initiator device drivers meet the storage requirements for operating system and device driver levels.

Software initiators can be obtained from the operating system vendor. For example, Microsoft offers a software iSCSI initiator for download. The initiators also can be obtained as a part of an NIC firmware upgrade (if supported by NIC).

Tip: Consider the use of a separate network segment for iSCSI traffic. That is, isolate NICs, switches or virtual local area networks (VLANs), and storage system ports that participate in iSCSI communications from other traffic.

If you plan for redundancy, you must use multipath drivers. These drivers often are provided by the operating system vendor for iSCSI implementations, even if you plan to use hardware initiators.

HA clustering solutions can be implemented by using iSCSI, but certain restrictions might apply. For more information, see the storage system vendor compatibility guides.

When you plan your iSCSI solution, consider the following points:

- Flex System Enterprise Chassis nodes, the initiators, and the operating system are supported by an iSCSI storage system. For more information, see the compatibility guides from the storage vendor.
- Multipath drivers exist and are supported by the operating system and the storage system (when redundancy is planned). For more information, see the compatibility guides from the operating system vendor and storage vendor.

For more information about selecting the Flex System Ethernet I/O modules for your iSCSI connectivity, see 2.2, “Network integration” on page 10.

For information about IBM storage compatibility, see the System Storage Interoperation Center at this website:

<http://ibm.com/systems/support/storage/ssic>

2.4 Management integration

This section describes management integration aspects and tools that can be used to manage a combined BladeCenter and Flex System environment from a single pane of glass. If you are not using some or any of the tools that are described in this section, consider rolling them out because it might help you perform deployment, configuration, and support tasks more efficiently.

The following topics are covered:

- ▶ 2.4.1, “Managing hardware”
- ▶ 2.4.2, “Managing network switches”
- ▶ 2.4.3, “Managing FC SAN fabric”
- ▶ 2.4.4, “Managing physical and virtualized environments”

2.4.1 Managing hardware

One essential management task is hardware management. This task includes initial hardware inventory and configuration, acquiring and installing updates, operating system deployment, and health status monitoring and reporting.

This section describes the following useful tools that can help unify the management of a combined BladeCenter and Flex System environment:

- ▶ “Chassis Management Module”
- ▶ “ToolsCenter” on page 19
- ▶ “ToolsCenter Suite” on page 19
- ▶ “FastSetup” on page 19
- ▶ “Other deployment tools” on page 20

Chassis Management Module

The Chassis Management Module (CMM) provides multi-chassis management in the Flex System environment and is used to communicate with the management controller in each compute node. As with the Advanced Management Module (AMM) in the BladeCenter environment, it provides system monitoring, event recording, and alerts. It also manages the chassis, its devices, and the compute nodes. The chassis supports up to two CMMs. If one CMM fails, the second CMM can detect its inactivity, self-activate, and take control of the system without any disruption. The CMM is central to the management of the chassis and is required in the Enterprise Chassis.

Through an embedded firmware stack, the CMM implements functions to monitor, control, and provide external user interfaces to manage all chassis resources.

You can use the CMM to perform the following functions:

- ▶ Define login IDs and passwords.
- ▶ Configure security settings, such as data encryption and user account security.
- ▶ Select recipients for alert notification of specific events.
- ▶ Monitor the status of the compute nodes and other components.
- ▶ Find chassis component information.
- ▶ Discover other chassis in the network and enable access to them.
- ▶ Control the chassis, compute nodes, and other components.
- ▶ Access the I/O modules to configure them.
- ▶ Change the start sequence in a compute node.
- ▶ Set the date and time.
- ▶ Use a remote console for the compute nodes.
- ▶ Enable multi-chassis monitoring.
- ▶ Set power policies and view power consumption history for chassis components.

ToolsCenter

The ToolsCenter™ is a collection of server management tools with which you can manage your BladeCenter and Flex System environment. ToolsCenter makes managing your server environment less complicated, more productive, and cost-effective.

For more information about the ToolsCenter, see this website (requires IBM ID):

<https://www-947.ibm.com/support/entry/myportal/docdisplay?ln docid=tool-center>

The ToolsCenter web page contains the links to the tools that are described in this section.

ToolsCenter Suite

ToolsCenter Suite (TCS) is a consolidation of server management tools that helps simplify the management of System x® and BladeCenter servers, and Flex System compute nodes. TCS provides functions to collect system health information, system full information, configure system settings, update system firmware and drivers, and Features on Demand (FoD) mass activation key management.

Among the TCS features, FoD mass activation key management is the inventory, acquisition, and installation FoD keys for multiple end points. All tasks can be done in a simple, unified, web-based, single-user GUI. It is suitable for a user with little knowledge or an advanced IT professional user.

FastSetup

FastSetup is a no-cost software tool that helps simplify the maintenance and deployment of select x86 servers, including System x rack servers, BladeCenter blade servers, and Flex System compute nodes. The intuitive GUI starts all phases of server setup, including discovery, updating, and configuration. Features include templates that enable replication of settings across many servers and automation that reduces hands-on time and user errors. Wizards and other default settings enable flexible customization capabilities. The low-touch, set-once and walk-away feature reduces the hands-on server setup time from days to minutes, particularly for larger deployments.

Other deployment tools

The following deployment tools also are part of TCS that can help simplify deployment of Lenovo System x platform-based environments:

- ▶ **ServerGuide™**

ServerGuide is a System x server installation assistant that simplifies the process of installing and configuring System x, BladeCenter, and Flex System servers. ServerGuide goes beyond hardware configuration by assisting with the installation of your operating system, the latest system device drivers, and other system components with minimal user intervention.

- ▶ **UpdateXPress**

UpdateXpress can help reduce your cost of computing by providing an effective and simple way to update device drivers, server firmware, and the firmware of supported options that are contained within the server on most of your System x, BladeCenter, and Flex System products. If you purchased a System x server, UpdateXpress is available for download at no charge. UpdateXpress System Packs™ (UXSPs) contain an integration-tested bundle of online, updateable firmware and device drivers for your servers.

- ▶ **Scripting toolkit**

The Scripting Toolkit is a collection of system-configuration tools and installation scripts that you can use to deploy software to your System x server in a repeatable, predictable manner. When used with ServerGuide and UpdateXpress, the ServerGuide Scripting Toolkit provides a total solution for deploying System x servers in an unattended mode.

2.4.2 Managing network switches

Ethernet I/O modules also can be managed by the CLI, web interface, Switch Center, or any third-party SNMP-based management tool.

The EN4093R 10Gb Scalable Switch, CN4093 10Gb Converged Scalable Switch, and the EN2092 1Gb Ethernet Switch modules all offer two CLI options (because it is a non-managed device, the pass-through module has no user interface). The default CLI for these Ethernet switch modules is the Networking operating system CLI, which is a menu-driven interface. A user also can enable an optional CLI that is known as industry standard CLI (isCLI) that more closely resembles Cisco IOS CLI. The SI4091 and SI4093 System Interconnect Modules support only the isCLI option for CLI access.

For more information about how to configure various features and the operation of the various user interfaces, see the *Application and Command Reference* guides, which are available at this website:

<http://pic.dhe.ibm.com/infocenter/flexsys/information/index.jsp>

Switch Center

Switch Center provides remote monitoring and management of Ethernet and converged switches from Lenovo. It is designed to simplify and centralize the management of your BladeCenter, Flex System, and RackSwitch™ Ethernet and converged switches.

The Switch Center offers the following features:

- ▶ Improve network visibility and drive availability, reliability, and performance
- ▶ Simplify management of large groups of switches with automatic discovery
- ▶ Automate and integrate management, deployment, and monitoring
- ▶ Simple network management protocol (SNMP) based configuration and management

- ▶ Support of network policies for virtualization
- ▶ Authentication and authorization
- ▶ Fault and performance management
- ▶ Integration with VMware Virtual Center and vSphere clients

For more information about Switch Center, see this website:

<http://www.ibm.com/systems/networking/software/snsc/index.html>

Any third-party management platforms that support SNMP also can be used to configure and manage the modules.

2.4.3 Managing FC SAN fabric

If you are planning to connect your Flex System to the existing FC SAN fabric, consider the use of NPIV mode on the Flex System FC SAN modules. NPIV mode does not add any domain IDs and management points to your existing storage network; therefore, you can continue to use your SAN fabric management tools without any changes.

If you are standardized on a SAN fabric management application from a specific vendor, such as Brocade Network Advisor or QLogic Enterprise Fabric Suite, consider selecting respective FC SAN modules for Flex System, if available (for more information, see “FC switch selection and fabric interoperability rules” on page 14).

2.4.4 Managing physical and virtualized environments

For managing physical and virtualized environments, you can continue use the tools that are deployed in your existing infrastructure. Also, for VMware vSphere and Microsoft Windows Server and Hyper-V environments, Lenovo offers powerful extensions that are called Upward Integration Modules (UIMs) that integrate hardware management features, such as status monitoring, firmware upgrades, and predictive failure alerts (PFA) into a management application (VMware vCenter and Microsoft System Center).

Upward integration for VMware vSphere

UIMs for VMware vSphere provide IT administrators with the ability to integrate the management features of the System x offerings with VMware vCenter. Lenovo expands the virtualization management capabilities of VMware vCenter with Lenovo hardware management functionality, which provides affordable, basic management of physical and virtual environments to reduce the time and effort that is required for routine system administration. It also provides the discovery, configuration, monitoring, event management, and power monitoring that is needed to reduce cost and complexity through server consolidation and simplified management.

Combined with the management features of Lenovo x86 offerings, VMware vCenter enhances and extends VMware's virtualization technologies and hardware service management to help you dramatically reduce complexity and cost. The following key features are included:

- ▶ Provides an overview of the host or cluster status, including information summary and health messages of the managed entities.
- ▶ Collects and analyzes system information to help diagnose system problems.
- ▶ Acquires and applies the latest UpdateXpress System Packs and individual firmware updates to your ESXi system.

- ▶ Provides nondisruptive system updates that automate the update process of the hosts in a cluster environment without any workload interruption.
- ▶ Monitors and provides a summary of power usage, thermal history, and fan speed and a trend chart of the managed host. Enables or disables the Power Metric function on a host and set the power capping for a power-capping capable host to limit the server power usage. Also, it supports power throttling and provide notification if the server power usage exceeds the specific value.
- ▶ Manages the current system settings on the host, including IMM, uEFI, and boot order settings for the host.
- ▶ Monitors the server hardware status and automatically evacuates virtual machines in response to predictive failure alerts to protect your workloads.

For more information, see the System x Upwards Integration Modules for VMware vSphere product page at this website:

<https://www-947.ibm.com/support/entry/myportal/docdisplay?ln docid=migr-vmware>

Upward integration for Microsoft System Center

The System x UIM for Microsoft System Center provides IT administrators with the ability to integrate the management features of the System x, BladeCenter, and Flex System servers with Microsoft System Center. Lenovo expands Microsoft System Center server management capabilities by integrating Lenovo hardware management functionality, which provides affordable, basic management of physical and virtual environments to reduce the time and effort that is required for routine system administration.

It also provides the discovery, configuration, monitoring, event management, and power monitoring that is needed to reduce cost and complexity through server consolidation and simplified management. The UIM for Microsoft System Center provides the following features:

- ▶ Integrated end-to-end management of System x hardware with monitoring of physical and virtual server health
- ▶ Operating system deployment with the latest firmware and driver update management
- ▶ Automated VM migration that is based on server health or power consumption
- ▶ Perform hardware configuration and firmware and driver updates and checks for the latest updates from the support website
- ▶ Collect Lenovo specific hardware inventory of System x or x86 blade servers
- ▶ Power on and off blades via Microsoft System Center console
- ▶ Author configuration packs to perform compliance checking on System x or BladeCenter x86 servers or Flex System compute nodes
- ▶ Manage servers remotely, independent of operating system state
- ▶ One year of software service and maintenance (three years available as an option)

UIMs for Microsoft System Center can be purchased as a one- or three-year software service and maintenance license.

For more information, see the Upward Integration for Microsoft System Center bundle product page at this website:

<https://www-947.ibm.com/support/entry/myportal/docdisplay?ln docid=migr-5087849>



Adding Flex System to a BladeCenter environment

This chapter describes deployment and configuration examples for Flex System that is being integrated into a BladeCenter environment. These examples include the initial configuration of Flex System hardware by using FastSetup, networking configuration with Switch Center, and integration into virtual infrastructures and virtualization management tools.

This chapter includes the following topics:

- ▶ 3.1, “Hardware deployment” on page 24
- ▶ 3.2, “Networking and storage deployment” on page 34
- ▶ 3.3, “Integrating VMware vSphere” on page 48
- ▶ 3.4, “Integrating Microsoft Hyper-V” on page 58

3.1 Hardware deployment

This section describes the following Flex System hardware initial configurations steps:

- ▶ 3.1.1, “Initial Flex System chassis configuration”
- ▶ 3.1.2, “Firmware updates and basic configuration by using FastSetup” on page 26
- ▶ 3.1.3, “Configuring Active Directory Integration for CMM” on page 33

3.1.1 Initial Flex System chassis configuration

After your Flex System in place, you must configure the Chassis Management Module (CMM) and other chassis components. Complete the following steps:

1. By default, the CMM has IP 192.168.70.100; therefore, to configure the CMM, set an IP from subnet 192.168.70.0/24 on your machine and connect it to the CMM if there is no DHCP server that is running in the network. If the DHCP server is running, CMM receives an IP address from the DHCP server.
2. Open a browser on your machine and browse to <https://192.168.70.100> or the dynamic IP address that was received via DHCP.
3. Log in by using the default credentials. Consider changing the password on your first login.

Note: Management Modules use the following default credentials:

- ▶ User name: USERID
- ▶ Password: PASSWORD

The Initial Setup wizard opens.

4. Read the Welcome page and click **Next**.
5. Examine the list of your components on the Inventory and Health page, as shown in Figure 3-1. Click **Next**.

Initial Setup Wizard

☒ Welcome

Inventory and Health

Import Existing Configuration

General Settings

Date and Time

IP Configuration

IO Modules

Security Policy

DNS

Event Recipients

Confirm

Inventory and Health

Shows the currently detected inventory and health of your components

Examine the list of your components below and confirm that all components are present a

Health status

Active events

Device Name	Device Type	Health Status	Bay	M
Standby CMM	Management Module	Normal	1	.
SN#Y011BG25302F	Management Module	Normal	2	.
Node 01 (Discovering)	Compute Node	Discovery	1	.
Node 02 (Discovering)	Compute Node	Discovery	2	.
Node 03 (Discovering)	Compute Node	Discovery	3	.
Node 04 (Discovering)	Compute Node	Discovery	4	.

Figure 3-1 Inventory and Health

- If you saved a configuration, you can import it by using the Import Existing Configuration page. If you do not have anything to import, click **Next**.
- Enter the Management module name, Chassis description, and other information in the General Settings page. Click **Next**.
- On Date and Time page, configure your time zone, current date, and time. Click **Apply**, then click **Next**.
- Set the Hostname of CMM and configure the IP settings on IP Configuration page, as shown in Figure 3-2. Click **Next**.

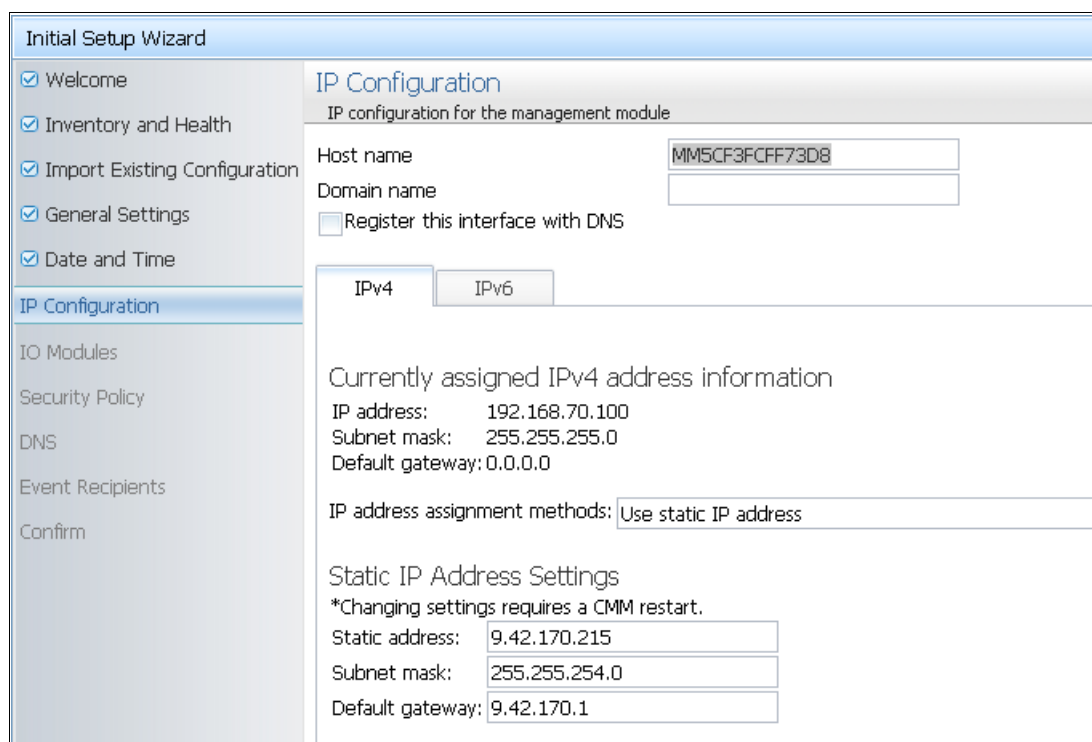
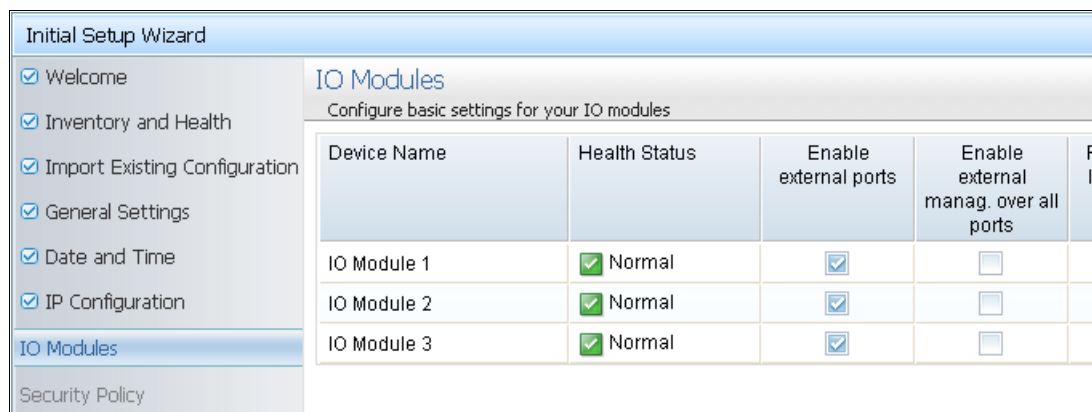


Figure 3-2 CMM IP settings.

- On the IO Modules page, configure the basic settings for your I/O modules, as shown in Figure 3-3. Click **Next**.



Device Name	Health Status	Enable external ports	Enable external manag. over all ports	Port IP
IO Module 1	✓ Normal	<input checked="" type="checkbox"/>	<input type="checkbox"/>	
IO Module 2	✓ Normal	<input checked="" type="checkbox"/>	<input type="checkbox"/>	
IO Module 3	✓ Normal	<input checked="" type="checkbox"/>	<input type="checkbox"/>	

Figure 3-3 I/O Modules basic settings

11. On the Security Policy page, select security policy that best suits your company needs. Click **Next**.
12. Based on your environment, configure the DNS configuration on the DNS page. Click **Next**.
13. On the Event Recipients page, you can configure email address of users that are to be notified about all events. Click **Next**.
14. Review configuration on Summary page. Click **Finish**.

After your CMM is configured and booted, you can connect to it through your management network by using the IP address that was configured in Initial Setup wizard.

You can proceed with configuration of IP addresses for the rest of components from the Component IP Configuration page, which is under the Chassis Management menu of the CMM interface.

3.1.2 Firmware updates and basic configuration by using FastSetup

Complete the following steps to use the FastSetup tool to configure Flex System in the same way as you use FastSetup to configure BladeCenter:

1. Start the FastSetup tool. A welcome opens. Click **Next**.
2. On the Initial Configuration page, configure the proxy server (if needed) and select the Network interface that is used to communicate with Flex System, as shown in Figure 3-4. Click **Next**.

Network Access

Tell IBM FastSetup how your local workstation is connected to the Internet and connected to the LAN. A connection to Internet is required to download firmware updates from ibm.com. The LAN is used to access the resources to be managed. The proxy configuration will be saved when Internet Explorer is allowed to save the cookie from IBM FastSetup.

Proxy Settings(optional)

If your local workstation requires a proxy server to connect to the Internet, enter the information below. The proxy configuration will be saved when Internet Explorer is allowed to save the cookie from IBM FastSetup.

IP/host name:

Port:

User name:

Password:

LAN Access

IBM FastSetup has detected the following network adapters. Select the adapter corresponding to the network IBM FastSetup should use to access the resources you want to manage.


	Ethernet Adapter	Description	IP Address
	Local Area Connection 5	Broadcom BCM5709S NetXtreme II GigE (NDIS VBD Client) #4	9.42.171.21

Figure 3-4 FastSetup: Network Access

3. On the Resource Selection page, select **Flex System (CMM, x86 Compute Nodes and I/O Modules)**. Click **Next**.
4. On the Task Selection page, select a task that is based on your needs. In our example, we use **Full Setup**. Click **Next**.
5. On the System Discovery page, you can allow Fast Setup to scan your network and find all accessible CMMs or you can add your CMM manually. After the CMM is discovered (as shown in Figure 3-5), click **Next**.

System Discovery






Select a method for discovering a system in your environment. You can choose to automatically discover a system, manually enter a system IP address to discover, or select a system from a list of previously discovered systems.

 **EIZUI111V**
 Run successfully.
[Details...](#)

☐ Automatically discover systems in this subnet
☒ Manually enter a system IP address or host name to discover

System IP address or host name:



IP Addresses/Host Name to Discover:

 |  |  |  |  Actions ▼

	IP Address/Host Name
<input checked="" type="checkbox"/>	9.42.170.215

☐ Select from a list of previously discovered systems

Select the system that you would like to configure:

 | 


	Name	Model	Machine Type	URL/IP Address	Status
<input checked="" type="checkbox"/>	LAB_Chassi	Flex Chassis	8721HC1	9.42.170.215	 Valid

Figure 3-5 Discovering CMM in IBM FastSetup

6. It takes some time to discover the details about your Flex System. You can monitor discovery status on the Inventory and Health page. After the discovery is finished, click **Next**.

- On the Device Selection page, select the component that you want to configure and update it by using Fast Setup, as shown in Figure 3-6. Click **Next**.

Device Selection

Select the devices to manage.

Actions ▼

	Name	Description	Power	Status
<input type="checkbox"/>	Slot 2 (Node 02 (x240_02))	IBM Flex System x240 Compute Node with embedded 10Gb Virtual Fabric	ON	Warning
<input checked="" type="checkbox"/>	Slot 3 (Node 03 (x240_03))	IBM Flex System x240 Compute Node with embedded 10Gb Virtual Fabric	ON	OK
<input type="checkbox"/>	Slot 4 (Node 04 (x240_04))	IBM Flex System x240 Compute Node with embedded 10Gb Virtual Fabric	ON	OK
	I/O Modules			
<input checked="" type="checkbox"/>	Slot 1 (IBM Flex System Fabric EN4093 10Gb Scalable Switch)	EN4093 10Gb Ethernet Switch	ON	OK

Figure 3-6 Select components

- On the Temporary IP Settings page, enter the temporary IP for the Flex System compute node, as shown in Figure 3-7. This IP address is used to boot the compute node from the FastSetup machine. Make sure that this IP address has full network access to a machine where Fast Setup is run. Click **Next**.

Temporary IP Settings

In order to manage your device, IBM FastSetup will boot your selected device into maintenance mode. Select an option below for applying a temporary network configuration.

☐ DHCP - Assign addresses using DHCP
☐ Address pools - Assign static IP addresses from pools
☒ Custom - Specify a temporary static IP address for each server

☐ Use the same network mask for all servers:
☐ Use the same gateway address for all servers:

System	Description	IP Address	Network Mask	Gateway
LAB_Chassi	Flex Chassis			
Slot 3	IBM Flex System x240 Compute Node with embedded 10Gb Virtual Fabric	9.42.171.26	255.255.254.0	9.42.170.1

Figure 3-7 Configure temporary IP address for the compute node

9. On the Adapter Port Settings page, select which NIC Fast Setup should use for communication, as shown in Figure 3-8. Click **Next**.

Adapter Port Settings

For each server, specify the adapter port connected to the data network.

☐ All servers use the same adapter port: NIC 1 (I/O Bay 1(A3))

Actions

System	Description	IP Address	Adapter Port - MAC Address (I/O Bay(Internal Port))
LAB_Chassi			
Slot 3	IBM Flex System x240 Compute Node with embedded 10Gb Virtual Fabric	9.42.171.26	NIC 1 - 34:40:b5:be:7d:00 (I/O Bay 1(A3))

Figure 3-8 Select NIC to use

10. The compute node must be rebooted after you click **Next** in the previous step. A pop-up window opens in which you confirm your choice.
11. FastSetup is collecting all possible information about the compute node. After the process completes (as shown in Figure 3-9), click **Next**.

Device Inventory

A detailed inventory is being collected on the selected devices to obtain the current firmware levels.

Actions

Device Name	Description	Build ID	Release Date	Version	Status
LAB_Chassi	Flex Chassis				✓ Finished
Servers					
slot 3	IBM Flex System x240 Compute Node with embedded 10Gb Virtual Fabric				✓ Finished

Figure 3-9 Gathering details about the compute node

12. In the Management Module Configuration page, you can edit some of CMM network settings, as shown in Figure 3-10. Click **Next**.

Management Module Configuration

Select the Management Module you want to configure and select Configure Settings.

Configure Settings

	Chassis	Bay	IP Address	Host Name	Status
	LAB_Chassi	2	9.42.170.215	flexchassis1	100%

Figure 3-10 Management Module configuration

13. On the Management Module Update page (as shown in Figure 3-11), you can apply Firmware updates to CMM. Select the applicable updates and then click **Next**. If no updates are needed, click **Next** only.

Management Module Updates

Select the type of update you want to apply, select the Management Modules to which it will be applied, and click Apply Updates.

☒ Update using the latest available Management Module firmware.

☐ Select from a list of all available Management Module firmware levels.

	Chassis	Description	Installed Version	Pending Version	Status
	LAB_Chassi		2PET12I	<div>unavailable</div>	N/A

Figure 3-11 Update CMM Firmware

14. On the I/O Module Configuration page (as shown in Figure 3-12), configure the network parameters for all selected I/O modules. Click **Next**.

I/O Module Configuration

Configure the IP address for the I/O Modules in each chassis.

Actions ▾

Chassis	Description	State	Configuration Type	IP Address	Network Mask	Gateway	
LAB_Chassi	Flex Chassis						
	Bay 1	IBM Flex System Fabric EN4093 10Gb Scalable Switch	Enabled ▾	Static ▾	9.42.171.8	255.255.254.0	9.42.170.1

Figure 3-12 Configure network settings for the I/O modules

- On the I/O Module Updates page (as shown in Figure 3-13), you can update the firmware of your I/O module. If there are applicable updates, select those updates, and then click **Next**. If no updates are needed, click **Next**.

I/O Module Updates

Select the type of update you want to apply, select the I/O Modules to which the update will be applied, and click Apply Updates.

☒ Update using the latest available I/O Module firmware
☐ Select from a list of all available I/O Module firmware levels

	System	Description	Installed Version	Pending Version	Status
LAB_Chassi					
<input type="checkbox"/>	Slot 1	IBM Flex System Fabric EN4093 10Gb Scalable Switch	Boot ROM: 7.5.3.0 Main Application 1: 7.5.3.0 Main Application 2: 7.2.2.2		N/A

Figure 3-13 I/O modules firmware update

- Compute node firmware updates can be applied on the Server Updates page of the FastSetup tool, as shown in Figure 3-14. If there are applicable updates, select those updates and then click **Next**. If no updates are need, click **Next**.

Server Updates

Select the update type you want to apply. Then select the servers or components to which it will be applied and click Apply Updates.

☐ Update using the UpdateXpress System Pack (UXSP) - server level only
☐ Update using the latest available component firmware
☒ Select from a list of all available component firmware levels

	System	Description	Installed Version	Pending Version	Status
LAB_Chassi					
<input type="checkbox"/>	Slot 3	IBM Flex System x240 Compute Node with embedded 10Gb Virtual Fabric			Loaded
<input checked="" type="checkbox"/>		Emulex UCNA Adapter Firmware Update	Unrecognized (N/A)	4.6.281.21-1(12/03/2013)	
<input checked="" type="checkbox"/>		IBM Online SAS/SATA Hard Disk Drive Update Program	Unrecognized (N/A)	1.13.04(02/27/2014)	

Figure 3-14 Update firmware on the compute node

17. Complete the following steps in the Server Configuration page:

- a. On the IMM Configuration page, you can configure network settings for the IMM on the compute node, as shown in Figure 3-15.

Integrated Management Module Configuration
Select the Management Module you want to configure, then select Configure Settings.

Configure Settings

	Chassis	Bay	IP Address	Host Name	Status
<input checked="" type="radio"/>	LAB_Chassi	3	9.42.171.18	IMM2-3440b5bf4d71	<input checked="" type="checkbox"/> 100%

Figure 3-15 Network configuration of the IMM

- b. On the RAID Configuration page, configure the RAID level that is to be used on your compute node, as shown in Figure 3-16.

Configure RAID Array

☒ Select the desired RAID level and the drives to be included in the array.
☐ Create a RAID0 array using all available drives

Server: Node 03 (x240_03) **RAID Controller:** SAS2004
Desired RAID level: RAID0 **Volume Size(2048-1907738):** 2,048 MB
Minimum drives: 2 **Added number of drives:** 0

Available Drives					Added Drives	
<input type="checkbox"/>	953	Drive 1	SAS_H DD	Yes		
<input type="checkbox"/>	953	Drive 2	SAS_H DD	Yes		

>>

Figure 3-16 Configure RAID on the compute node

- c. On the UEFI Settings page (as shown in Figure 3-17), configure the Boot Order for UEFI or reset the order to the default settings.

Configure Basic Settings

☒ Set all UEFI settings to default values
☐ Specify boot order and set all other UEFI settings to default values
☐ Specify boot order only

Figure 3-17 Configure UEFI settings in FastSetup

18. You can perform the following tasks on the Summary page:

- Export settings for future use.
- Export Firmware repository. This task can be useful if you must run Fast Setup on a machine that does not have an Internet connection.
- On the Completion State page, you can select the action that you want to perform on the compute nodes after all of the actions are performed. You can shut down all of the servers after all of the tasks are complete or reboot them. Select one of the options, and then click **Next**.

3.1.3 Configuring Active Directory Integration for CMM

The CMM offers possibility to integrate into your centralized user management system by using LDAP or Microsoft Active Directory integration. In this section, we describe how to integrate it with Active Directory by using basic settings.

Note: Make sure you configured DNS settings and your domain is discoverable via DNS_SRV records.

Complete the following steps to configure LDAP authentication on the CMM:

1. Log in to the CMM web console by using your local account.
2. Click **Mgt Module Management** → **Network** → **LDAP Client**.
3. In the LDAP Authentication section, select **Use LDAP Servers for Authentication Only (with local authorization)**.
4. In the LDAP Servers field section, select **Use DNS to find LDAP Servers**.
5. In the Miscellaneous Settings Finding Method section, select **w/ Login credentials**.
6. You can leave the others fields empty and click **Apply**, as shown in Figure 3-18.

Lightweight Directory Access Protocol (LDAP) Client

The CMM contains a LDAP client that can be configured to provide user authentication through one or more LDAP servers. The client can be configured to be discovered dynamically or manually pre-configured. Use the dropdown list to select which of these two methods you want to use.

LDAP Authentication: Use LDAP Servers for Authentication Only (with local authorization) ▼

LDAP Servers: Use DNS to find LDAP Servers ▼

Active Directory Forest Name:

Domain Name:

Active Directory Settings

Use Mgt Module Management > User accounts for user configuration

Miscellaneous Settings

Root DN:

UID search attribute:

Binding method: w/ Login credentials ▼

Figure 3-18 LDAP configuration

Complete the following steps to pair Active Directory Groups to CMM roles:

1. Click **Mgt Module Management** → **User Accounts** → **Group Profiles**, as shown in Figure 3-19.

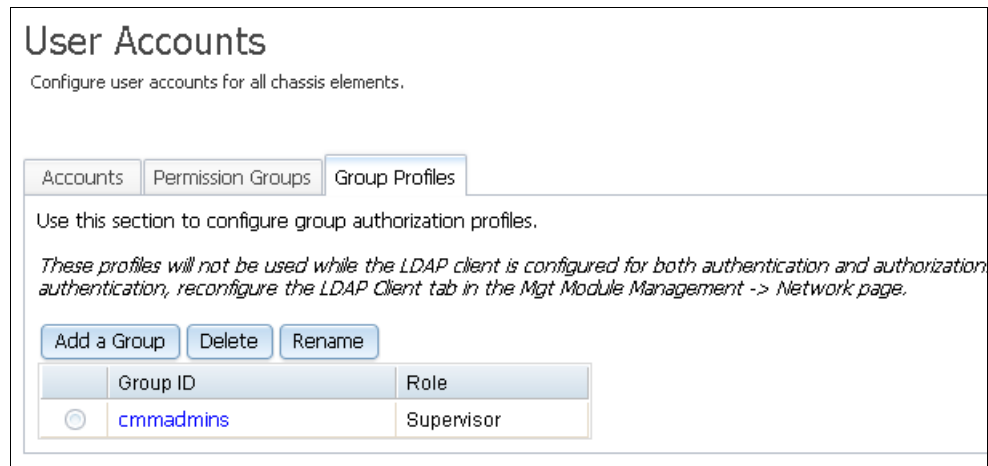


Figure 3-19 Active Directory Group-Role Mapping

2. Click **Add a Group**. Complete the following steps:
 - a. In the Group Profile Name tab, enter Group Name (the Active Directory group name). Click **Next**.
 - b. Define the Role. Click **Next**.
 - c. Define the Authority, click **Next**.
 - d. Define the Access Scope, click **Finish**.
3. Log out and log in by using a domain account.

3.2 Networking and storage deployment

To integrate your physical or virtualized environment that is running on your BladeCenter environment with your new Flex System, you must meet several prerequisites. BladeCenter managed hosts and Flex System compute nodes should have access to same networks and VLANs. They also should have access to the shared storage, if any. Network and storage solutions have various interfaces that can be used to configure network and storage.

In this section, we describe the examples of configuring BladeCenter and Flex System network switches via Switch Center and iSCSI storage connectivity with a web GUI by using V7000 storage as an example. The following topics are covered:

- ▶ 3.2.1, “Configuring the network”
- ▶ 3.2.3, “Configuring storage” on page 42

3.2.1 Configuring the network

To proceed with the configuration of the network, physical interconnections must be done. In the example that is described in this section, we use the Switch Center tool to create VLANs and assign those VLANs to network ports in the network switch, which is installed into the Flex System chassis. We assume that Switch Center is already installed in the network and it is used to configure BladeCenter switches.

Note: Consider deploying Switch Center for the unified management of BladeCenter and Flex System Ethernet switches from a single console.

We also configure Unified Fabric Port (UFP) on the 10 GbE ports of the compute node to create separated network subinterfaces for different types of traffic. The following tasks are performed:

- ▶ “Adding a switch to Switch Center”
- ▶ “Configuring VLANs on the switch with Switch Center” on page 37
- ▶ “Enabling UFP on the EN4093R by using Switch Center” on page 39
- ▶ “Enabling UFP in server’s UEFI” on page 41

Adding a switch to Switch Center

Complete the following steps to add a Flex System I/O module to Switch Center:

1. Run Switch Center on the machine where it is installed. A login window opens. If this is the first time that you logged in to Switch Center, use the default credentials. Consider changing the password on your first login.

Note: Switch Center uses the following default credentials:

- ▶ User name: admin
- ▶ Password: admin

2. On the main page of Switch Center, click **Device List Page**, as shown in Figure 3-20.

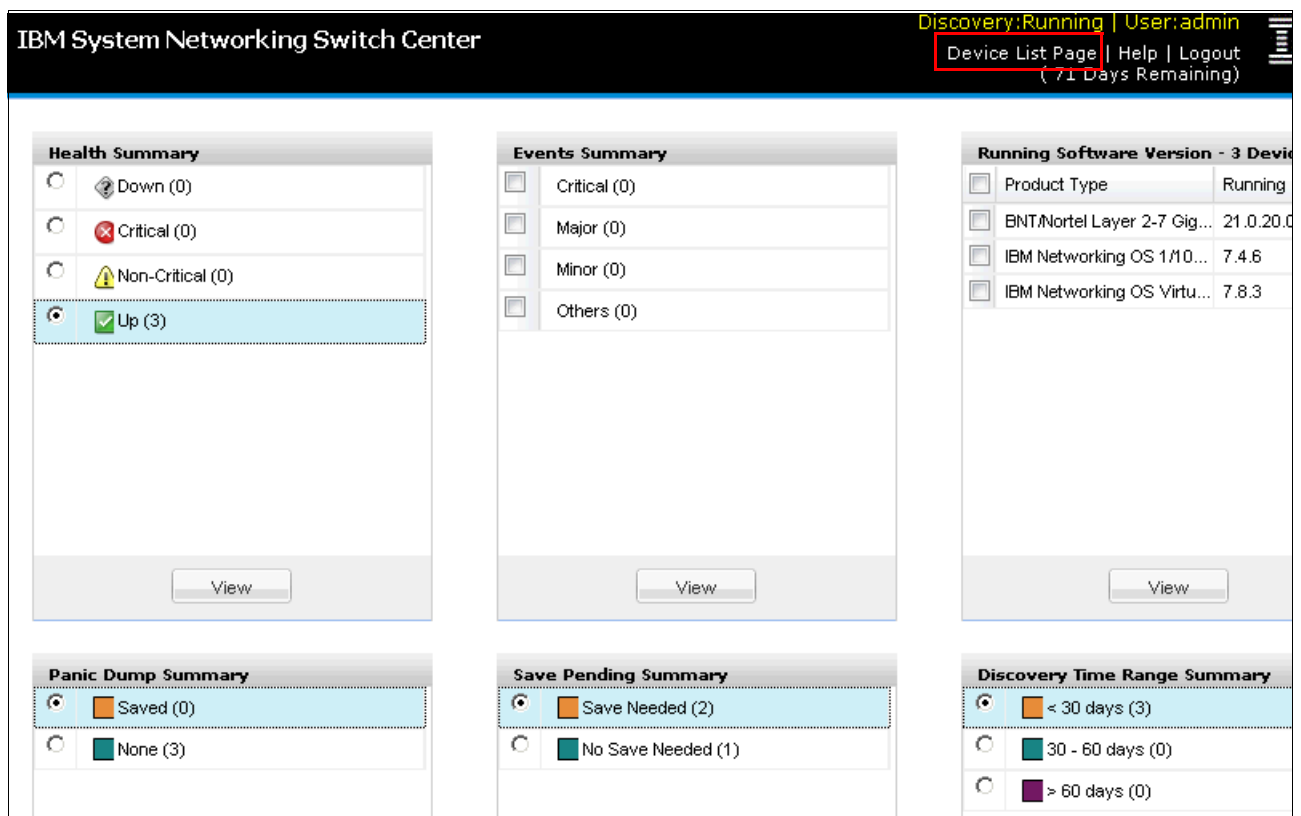


Figure 3-20 Switch Center main page

- On the Device List page, browse to the Flex System category and click **Add a Switch**, as shown in Figure 3-21.




Figure 3-21 Add a switch in Switch Center

- In the Add a Switch window, enter the switch connection details and then, click **Open**, as shown in Figure 3-22.



Figure 3-22 Connection details for a network switch

- After the switch is added, it is shown in the list of available switches (the BladeCenter and Flex System I/O modules are listed). Click the highlighted IP of the switch to see the Device Console, as shown in Figure 3-23.

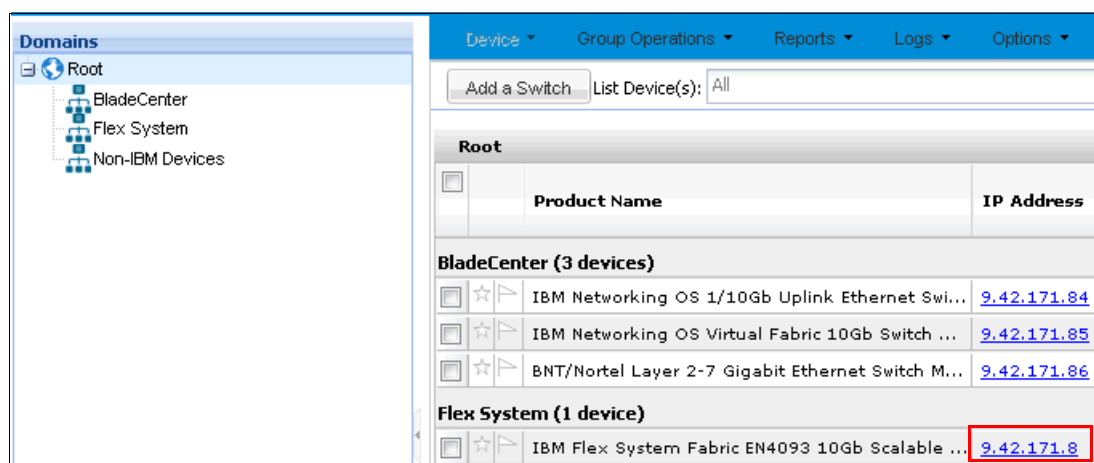


Figure 3-23 Selecting the switch to configure

Configuring VLANs on the switch with Switch Center

Complete the following steps to configure VLANs on the switch by using Switch Center:

1. To add VLANs to the switch, in Device Console, click the **Configure** tab, select the **Layer 2** folder and then, click **Virtual LANs**. Click **Insert** to insert a new VLAN, as shown in Figure 3-24.

IBM System Networking Switch Center - Device Console

EN4093, 9.42.171.8 IBM Flex System Fabric EN4093 10Gb Scalable Switch

Monitor **Configure**

Configuration tree:

- Configure
 - Switch
 - Config/Image/Dump Control
 - Layer 2**
 - General
 - Trunk
 - LACP
 - 802.1x
 - MSTP/RSTP
 - CIST
 - Spanning Tree Protocol
 - Forwarding Database
 - VLAG
 - Hot Links
 - Virtual LANs**
 - LLDP
 - Failover
 - ECP
 - Layer 3
 - Ports
 - QoS
 - Access Control List
 - CEE
 - FCoE
 - Virtualization

Actions: Help

VLAN Memberships Private VLAN Protocol VLAN VMA

VMAP for Server Ports VMAP for All Ports

Virtual LANs - VLAN Memberships

VLAN	Name	Ports	State	Spanning Tree Group
1	Default VLAN	INTA1-INTA	enabled	1
20	Production	INTA3-INTA	enabled	20
30	LiveMigration	INTA3-INTA	enabled	30
42	Management	INTA3-INTA	enabled	42
101	vPort1		enabled	101
102	vPort2	INTA4	enabled	102
103	vPort3		enabled	103
104	vPort4		enabled	104
1002	iSCSI SAN VL...	INTA3-INTA	enabled	2
4095	Mgmt VLAN	EXTM-MGT1	enabled	128

Submit Apply Refresh **Insert** Modify Delete Export Print Help

Figure 3-24 Virtual LANs menu

2. Enter the required details for the VLAN. In our example, we add VLAN 80 and assign that VLAN to ports INTA3 and EXT1, as shown in Figure 3-25. Click **OK**.

Virtual LANs - VLAN Memberships - Insert Form

VLAN: 80 1..4095

Name: Management_VLAN_80 0..32 characters

Ports: INTA3;EXT1 Browse...

State: ☒ enabled ☐ disabled

Spanning Tree Group: 80 0..127 [1-127 for 802.1d; 1 for RSTP; 0-32 for MSTP]

Management State: ☒ enabled ☐ disabled

Virtual Ports: Browse...

OK Cancel

Figure 3-25 VLAN details

3. To apply and save the configuration, click **Submit** at the bottom center of the page. Click **Apply** (which is next to Submit). To preserve this change across the reboot of the switch, click **Actions** → **Save**, as shown in Figure 3-26.

EN4093, 9.42.171.8 IBM Flex System Fabric EN4093 10Gb Scalable Switch

Monitor **Configure**

Configure

- Switch
- Config/Image/Dump Control
- Layer 2
 - General
 - Trunk
 - LACP
 - 802.1x
 - MSTP/RSTP
 - CIST
 - Spanning Tree Protocol
 - Forwarding Database
 - VLAP
 - Hot Links
 - Virtual LANs**
 - LLDP
 - Failover
 - ECP
- Layer 3
- Ports

Actions Help

Apply **Save** Diff Config Diff Flash Config Dump Syslog Dump Revert Revert Apply Reboot Switch Clear Panic Dump Exit

Private VLAN Protocol VLAN VMAP

VMAP for All Ports

Memberships

Ports	State	Spanning Tree Group	M S
INTA1-INTA	enabled	1	c
INTA3-INTA	enabled	20	c
INTA3-INTA	enabled	30	c
INTA3-INTA	enabled	42	c
INTA3;EXT1	enabled	80	e
INTA3-INTA	enabled	101	c
INTA4	enabled	102	c
INTA4	enabled	103	c
104	vPort4	104	c
1002	iSCSI SAN VL...	2	c
4095	Mgmt VLAN	128	e

Figure 3-26 Selecting Actions → Save

4. Complete the following steps:
 - a. Click **Ports** on the left side pane.
 - b. Find the port to which you assigned a created VLAN.
 - c. Double-click **VLAN Tag State** and change it to tagged.
 - d. Double-click **Default VLAN** in front of it and select the VLAN that you want to use as the default, as shown in Figure 3-27.




Figure 3-27 Configure port

- e. Click **Submit** at the bottom center of the page.
- f. Click **Apply** (which is next to Submit).
- g. To preserve this change across reboot of the switch, click **Actions** → **Save**.

The switch is now configured to accept traffic from VLAN 80.

Enabling UFP on the EN4093R by using Switch Center

UFP is an approach to NIC virtualization. It is similar to Virtual Fabric vNIC, but with enhanced flexibility and should be considered the direction for future development in the virtual NIC area for Lenovo switching solutions. With Flex System, UFP is supported today on the EN4093R 10Gb Scalable Switch, CN4093 10Gb Converged Scalable Switch, and SI4093 System Interconnect Module. With BladeCenter, UFP is supported on the Virtual Fabric 10Gb Switch Module.

If your BladeCenter infrastructure is standardized on other type of vNIC (for example, Virtual Fabric vNIC), it can be configured in a similar way. If you use Switch Independent vNIC, no switch configuration is required.

For more information about UFP and other NIC virtualization choices, see *NIC Virtualization in Flex System Fabric Solutions*, SG24-8223:

<http://lenovopress.com/sg248223>

Complete the following steps to enable UFP on the switch and on the port:

1. In Switch Center, open Device Console, then click **Configure** → **Virtualization** → **UFP**. In the General tab, select **enabled** and then, click **Apply**, as shown in Figure 3-28.

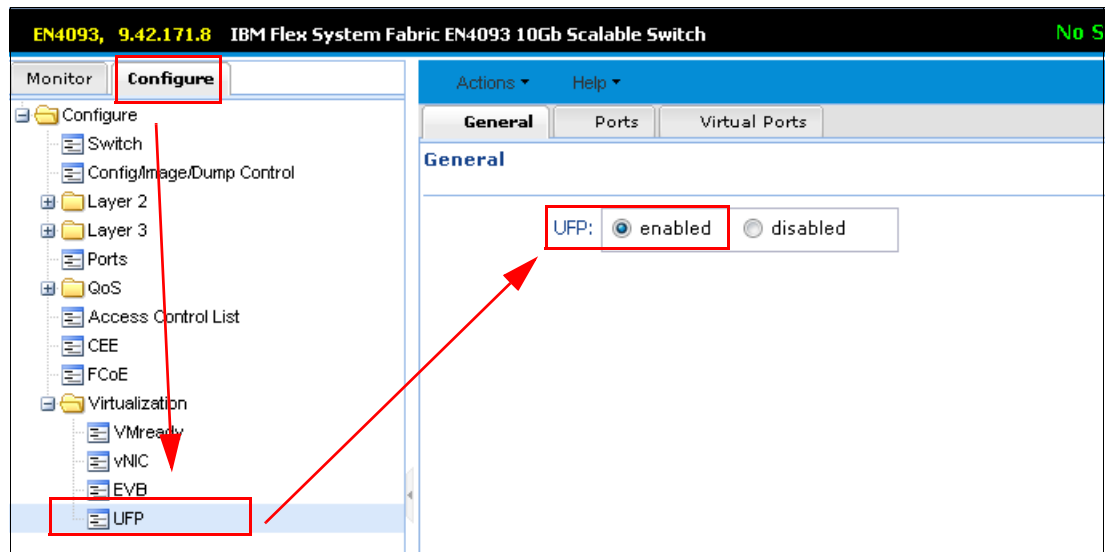


Figure 3-28 Enable UFP

2. Click the Ports tab, enable UFP for the wanted port by selecting **enable** in the state field and then, click **Apply**, as shown in Figure 3-29.

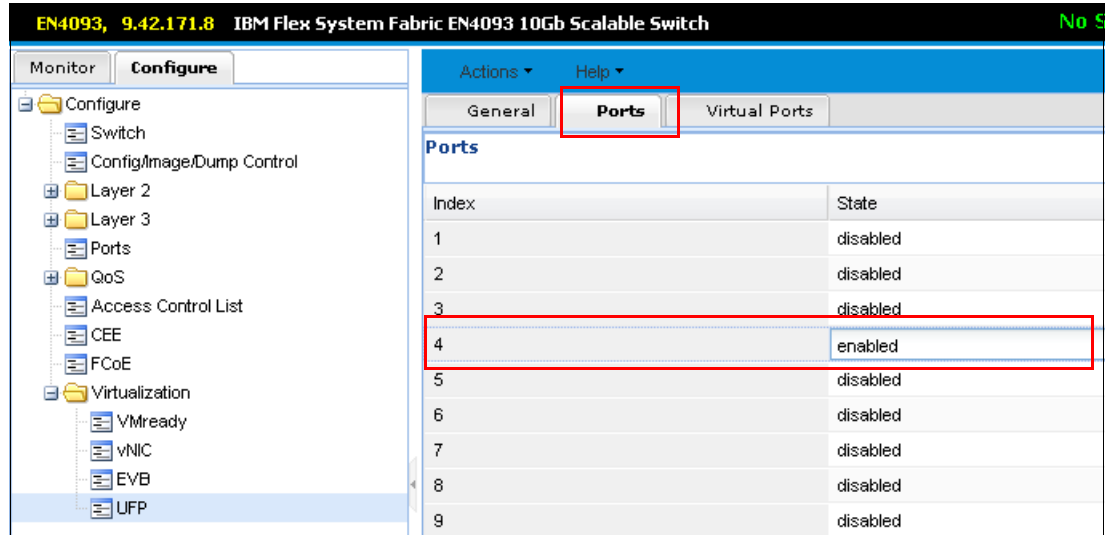


Figure 3-29 Enable UFP for Port

3. In the Virtual Ports tab, you can configure virtual ports, their VLANs membership, Network mode, and bandwidth, as shown in Figure 3-30.

Port Index	vPort Index	State	Network Mode	Network Default VLAN	Network Default Tag	QoS Min Guaranteed Bandwidth	QoS Max Allowed Bandwidth
1	1	disabled	tunnel	0	disabled	25	100
1	2	disabled	tunnel	0	disabled	25	100
1	3	disabled	tunnel	0	disabled	25	100
1	4	disabled	tunnel	0	disabled	25	100
2	1	disabled	tunnel	0	disabled	25	100
2	2	disabled	tunnel	0	disabled	25	100
2	3	disabled	tunnel	0	disabled	25	100
2	4	disabled	tunnel	0	disabled	25	100
3	1	disabled	tunnel	0	disabled	25	100
3	2	disabled	tunnel	0	disabled	25	100
3	3	disabled	tunnel	0	disabled	25	100
3	4	disabled	tunnel	0	disabled	25	100
4	1	enabled	access	42	disabled	25	25
4	2	enabled	access	1002	disabled	25	100
4	3	enabled	trunk	103	disabled	25	100
4	4	disabled	tunnel	104	disabled	25	100

Figure 3-30 Virtual Ports configuration

3.2.2 Enabling UFP in server's UEFI

To enable the UFP function, you must configure Flex System compute node. Complete the following steps:

1. Connect to the console of your compute node and enter the UEFI configuration by pressing F1 during start.
2. Enable Multichannel Mode in UFP Mode personality for the network adapter by clicking **System Settings** → **Network** in UEFI. Complete the following steps:
 - a. From the Network Device List, open the first adapter.
 - b. Press Enter to enter configuration mode.

- c. Change Multichannel Mode to Unified Fabric Protocol Mode, as shown in Figure 3-31.




Figure 3-31 Emulex NIC configuration utility

- d. Exit and apply to the second adapter, if needed
3. Exit and save UEFI configuration.
4. Reboot the server and your UFP configuration is finished.

After you configure network adapters in UFP mode, you are configuring virtual ports settings only in the network switch side; no other configuration in UEFI is needed.

3.2.3 Configuring storage

To connect Flex System to the existing environment, many factors must be considered. One factor is the type of storage that is used, SAN or NAS. If SAN is used, Fibre Channel, FCoE, or iSCSI is used. If you are running virtualization or clustered environment, you should consider that a single LUN should be allocated to several physical hosts and shared across them. In this section, the following topics are covered:

- ▶ “Configuring iSCSI adapter”
- ▶ “Storage LUN provisioning” on page 44

Configuring iSCSI adapter

The Converged Network Adapter that is included with Flex System can work as an iSCSI initiator or as a FCoE HBA. In our example, we describe the scenario of configuring it for use with an existing iSCSI SAN.

Complete the following steps to configure your iSCSI initiator parameters:

1. Connect to the console of your compute node and enter the System Setup utility.
2. Enable iSCSI personality for network adapter by clicking **System Settings** → **Network**. Complete the following steps:
 - a. From Network Device List, open first Adapter.

- b. Press Enter to begin the configuration mode, as shown in Figure 3-31 on page 42.
 - c. Change Personality to iSCSI.
 - d. Exit and configure the second NIC, if needed (the second port on the same NIC is configured with the first port).
 - e. Exit and save the UEFI configuration.
 - f. Reboot the system re-enter the System Setup utility.
3. Click **System Settings** → **Storage**. Complete the following steps:
 - a. Enter the Emulex iSCSI Utility for the particular adapter, as shown in Figure 3-32.

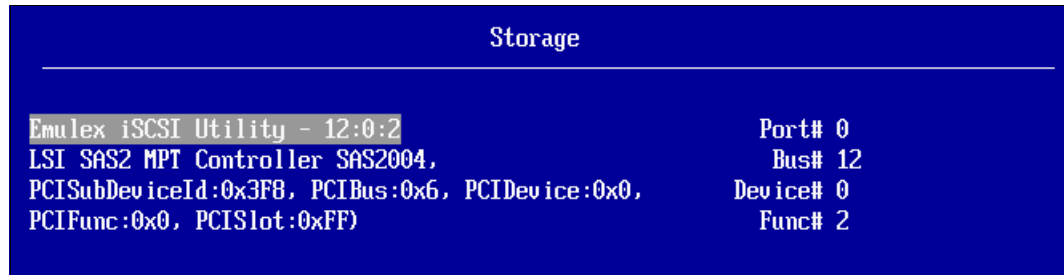


Figure 3-32 UEFI: Emulex iSCSI Utility

- b. Enter the iSCSI Initiator Name and save the changes, as shown in Figure 3-33.



Figure 3-33 uEFI iSCSI Controller Configuration menu

- c. Enter the Network Configuration.
 - d. Enter the iSCSI Target Configuration.
 - e. Exit the configuration window.
 - f. Repeat steps a - e for the second network adapter, if needed.
4. Reboot the computer.

Note: For more information about FCoE and iSCSI configuration, see *Storage and Network Convergence Using FCoE and iSCSI*, SG24-7986:

<http://lenovopress.com/sg247986>

Storage LUN provisioning

Now that the iSCSI initiator is configured, it must be able to see existing LUNs.

In our example, we show you how to allocate a LUN to a new Host in IBM Storwize V7000 web GUI that uses V7000 Storage Node as an example. Complete the following steps:

1. In a web browser, enter the IP address of the V7000 and login.
2. Hover your mouse over the Hosts icon on the left side of the window and click **Hosts** in the drop-down menu, as shown in Figure 3-34.

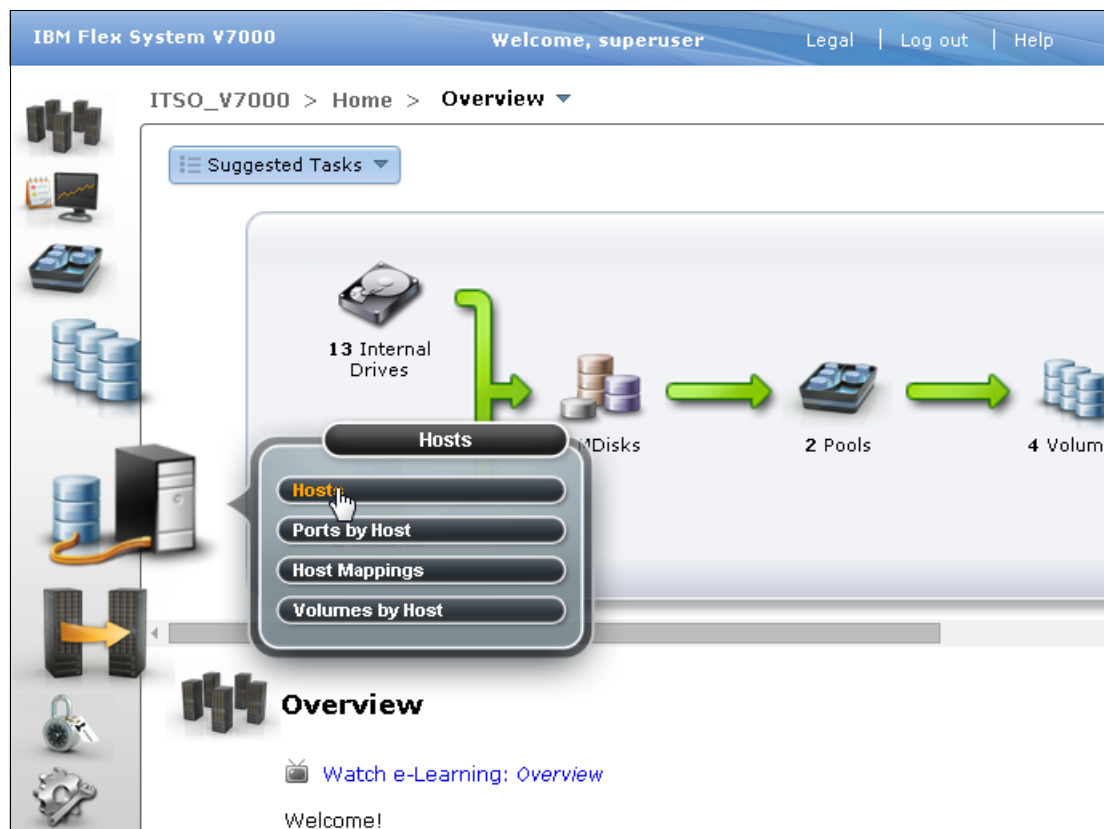


Figure 3-34 Enter Hosts menu

3. In the Hosts menu, click **New Host**, as shown in Figure 3-35.

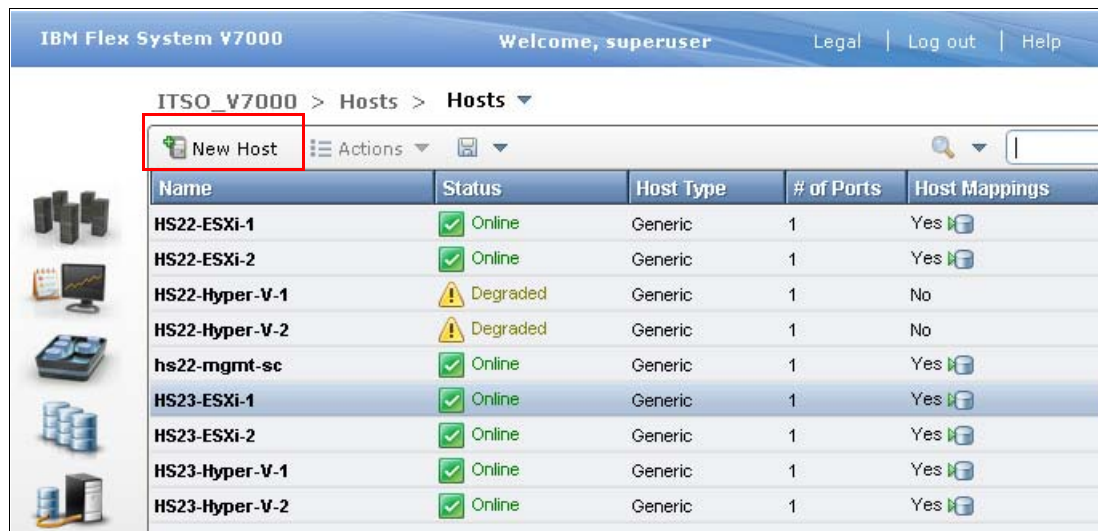


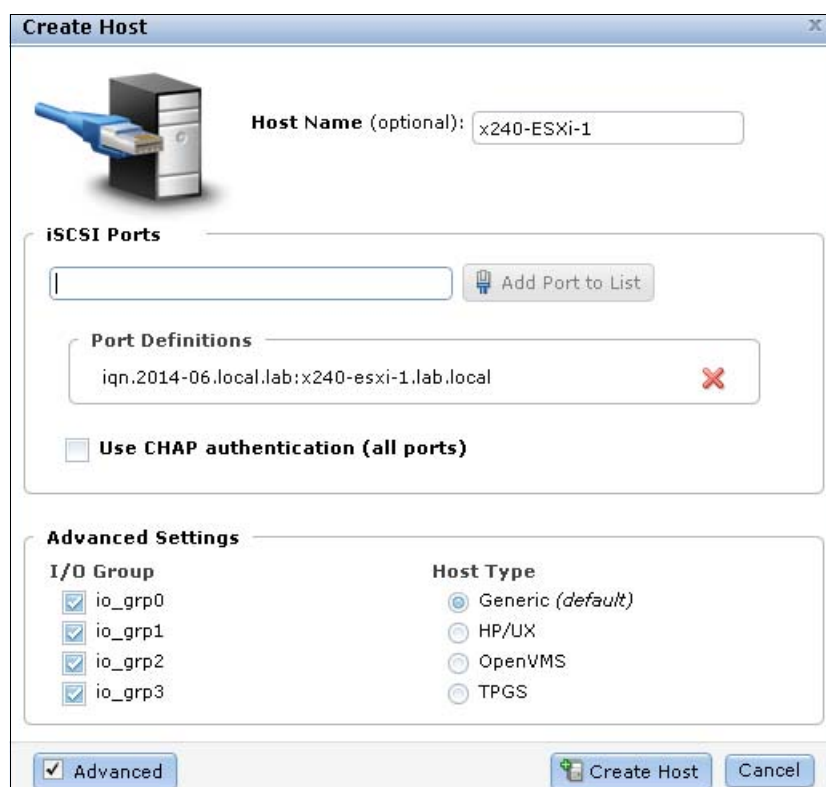
Figure 3-35 Add Host

4. In the window that opens, select **iSCSI Host**, as shown in Figure 3-36.



Figure 3-36 Selecting iSCSI Host

5. Enter the required information, including the Host Name and iSCSI qualified name (IQN) for adding an iSCSI host, as shown in Figure 3-37. Click **Create Host**.




The 'Create Host' dialog box is shown with the following fields and options:

- Host Name (optional):** x240-ESXi-1
- iSCSI Ports:** A text input field is empty, with an 'Add Port to List' button to its right.
- Port Definitions:** A text input field contains 'iqn.2014-06.local.lab:x240-esxi-1.lab.local'.
- Use CHAP authentication (all ports):** An unchecked checkbox.
- Advanced Settings:**
 - I/O Group:** Four checkboxes labeled 'io_grp0', 'io_grp1', 'io_grp2', and 'io_grp3', all of which are checked.
 - Host Type:** Four radio buttons labeled 'Generic (default)', 'HP/UX', 'OpenVMS', and 'TPGS'. 'Generic (default)' is selected.
- Buttons:** 'Advanced' (checked), 'Create Host', and 'Cancel'.

Figure 3-37 Entering required information

After the host is added to the storage system, complete the following steps to allocate LUNs to it:

1. Hover over the Hosts menu icon on the left side of the window and click **Host Mappings**, as shown in Figure 3-38.



The screenshot shows the 'Hosts' menu in the ITSO_V7000 interface. The menu is open, showing options: 'Hosts', 'Ports by Host', 'Host Mappings', and 'Volumes by Host'. The 'Host Mappings' option is highlighted by the mouse cursor. In the background, a table lists the hosts and their status.

Name	Status	Host Type	# of Ports	Host Mappings
HS22-ESXi-1	Online	Generic	1	Yes
HS22-ESXi-2	Online	Generic	1	Yes
HS22-Hyper-V-1	Degraded	Generic	1	No
HS22-Hyper-V-2	Degraded	Generic	1	No
hs22-mgmt-sc	Online	Generic	1	Yes
HS23-ESXi-1	Online	Generic	1	Yes
HS23-ESXi-2	Online	Generic	1	Yes
HS23-Hyper-V-1	Online	Generic	1	Yes
HS23-Hyper-V-2	Online	Generic	1	Yes
x240-ESXi-1	Online	Generic	1	No

Figure 3-38 Host Mappings option

2. In the Host Mappings window, right-click the Host that you added, and click **Modify Mappings**, as shown in Figure 3-39.

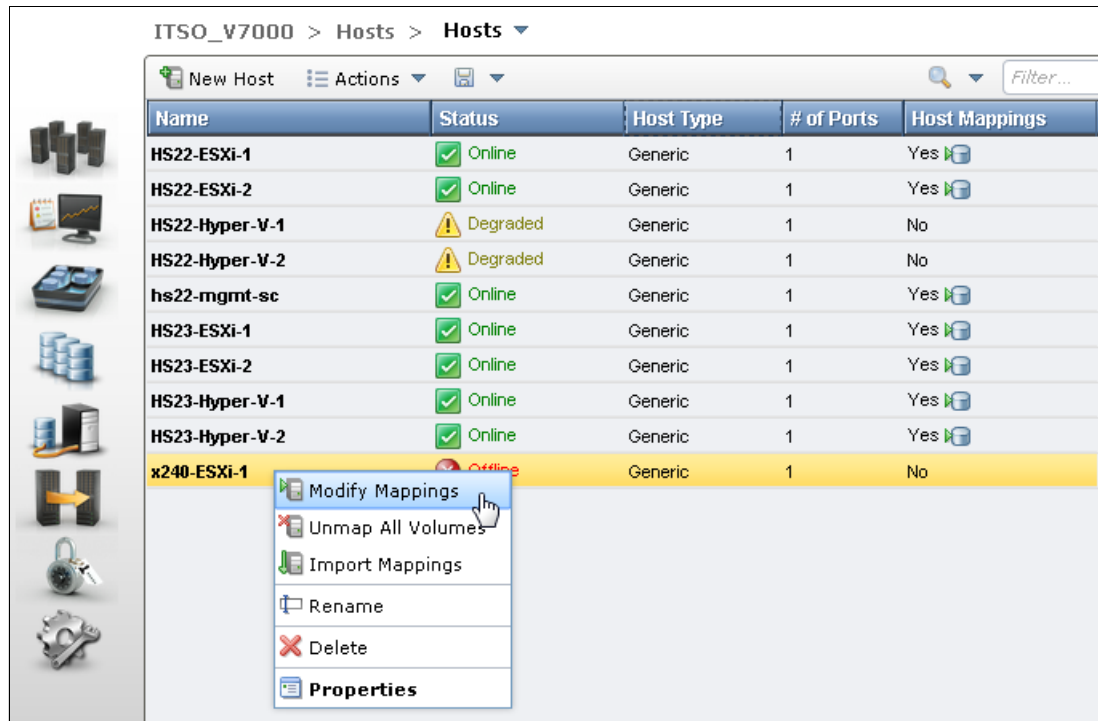


Figure 3-39 Modify Mappings menu

3. The Modify Host Mappings window opens. The available LUNs are listed in the left pane of the window. Select the LUNs that you can to allocate to the host and click the arrow icon in the center of the window, as shown in Figure 3-40. Click **Apply**.

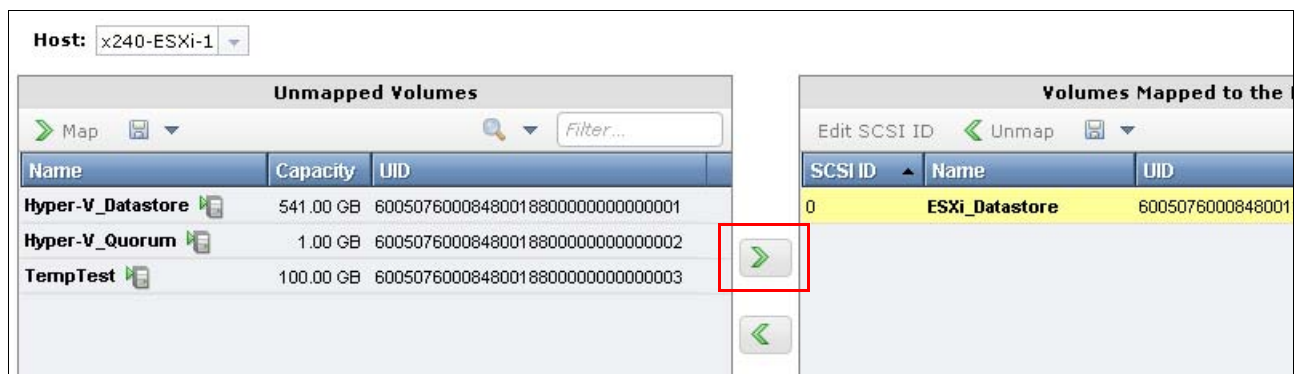


Figure 3-40 Allocate LUNs

4. If the LUN you assign to a host is assigned to any other host, you see a warning message in which you are prompted to confirm the action. Click **Map All Volumes**, as shown in Figure 3-41.

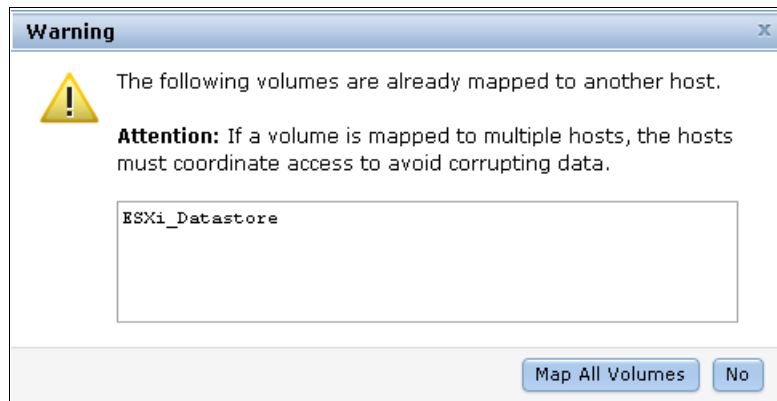


Figure 3-41 Warning message if the LUN is allocated to several hosts

The LUN is now allocated to your host. If your operating system is configured to discover LUNs on this storage device, perform a Rescan and the LUN becomes usable.

3.3 Integrating VMware vSphere

In this section, we describe the integration of VMware ESXi that is installed on Flex System Compute Node node to an existing VMware vSphere environment that was built on BladeCenter chassis. For more information about the full installation process of VMware ESXi, see the following resources:

- ▶ ESXi 5.0:
<http://pubs.vmware.com/vsphere-50/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcen-ter-server-50-installation-setup-guide.pdf>
- ▶ ESXi 5.1:
<http://pubs.vmware.com/vsphere-51/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcen-ter-server-51-installation-setup-guide.pdf>
- ▶ ESXi 5.5:
<http://pubs.vmware.com/vsphere-55/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcen-ter-server-55-installation-setup-guide.pdf>

3.3.1 USB memory key with ESXi

All Flex System Compute Nodes can have an integrated USB memory key. The USB memory key can be used for VMware ESXi installation. For more information about installing a USB key on a particular model of Flex System Compute Node, see the Flex System Information Center, which is available at this website:

<http://pic.dhe.ibm.com/infocenter/flexsys/information/index.jsp>

If a USB memory key is not used, the following installation methods can be used:

- ▶ Installation on local disk
- ▶ Boot from SAN

Use Lenovo customized VMware ESXi image to ensure that it contains all of needed hardware drivers and modules. Customized ESXi images can be downloaded from this website:

<http://www-03.ibm.com/systems/x/os/vmware/>

3.3.2 Configure management network for VMware ESXi

You should configure the Management Network to administer VMware ESXi that uses dedicated tools, such as VMware vSphere Client or VMware vCenter Server. After the installation is complete and the ESXi hosts are rebooted, you see the Direct Console User Interface (DCUI) of VMware ESXi.

Complete the following steps to configure the Management Network:

1. Press F2 and enter `root` as the user credentials, as shown in Figure 3-42.




Figure 3-42 Enter root credentials

2. In the DCUI, select the **Configure Management Network**, as shown in Figure 3-43.




Figure 3-43 Select Configure Management Network

3. In the Configure Management Network menu, complete the following configuration changes:
 - a. Select **Network Adapters** and choose the network card that is used for Management Network traffic. In our example, we use vmnic0, as shown in Figure 3-44. Press **Enter**.

```

Network Adapters

Select the adapters for this host's default management network
connection. Use two or more adapters for fault-tolerance and
load-balancing.
  
```

Device Name	Hardware Label (MAC Address)	Status
<input checked="" type="checkbox"/> vmnic0	ServerEngin... (...be:7d:00)	Connected (...)
<input type="checkbox"/> vmnic1	N/A (34:40:b5:be:7d:04)	Connected
<input type="checkbox"/> vusb0	N/A (36:40:b5:bf:4d:73)	Connected (...)

Figure 3-44 Selecting NIC that is used for Management traffic

- b. In the VLAN menu (if applicable), set the VLAN number that is used for Management Network traffic. In our example, we use VLAN 42, as shown in Figure 3-45. Press **Enter**.

```

VLAN (optional)

If you are unsure how to configure or use a VLAN, it is safe to
leave this option unset.
  
```

VLAN ID (1-4094, or 4095 to access all VLANs): [42_]

Figure 3-45 Setting VLAN for Management Network interface

- c. In the IP Configuration menu, do not change the default settings if DHCP is used. If you static IP configuration is used, select the Set static IP address and network configuration options, enter the IP address, Subnet Mask, and Default Gateway that is to be used for this ESXi host, as shown in Figure 3-46. Press **Enter**.

```

IP Configuration

This host can obtain network settings automatically if your network
includes a DHCP server. If it does not, the following settings must be
specified:
  
```

() Use dynamic IP address and network configuration
 (o) Set static IP address and network configuration:

IP Address	[9.42.171.26]
Subnet Mask	[255.255.254.0]
Default Gateway	[9.42.170.1]

Figure 3-46 Configure IP settings for Management Network

- d. In the DNS Configuration window, enter the IP addresses of your primary and secondary DNS servers and set the host name of VMware ESXi server, as shown in Figure 3-47. Press **Enter**.

```
DNS Configuration

This host can only obtain DNS settings automatically if it also obtains
its IP configuration automatically.

( ) Obtain DNS server addresses and a hostname automatically
(o) Use the following DNS server addresses and hostname:

Primary DNS Server [ 19.42.171.27 ]
Alternate DNS Server [                ]
Hostname           [ x240-ESXi-1.lab.local ]
```

Figure 3-47 Configure host names and DNS resolution on VMware ESXi

- e. In the Configure Management Network menu, press Esc and select Yes to save your changes and return to DCUI main menu, as shown in Figure 3-48.

```
Configure Management Network: Confirm

You have made changes to the host's management network.
Applying these changes may result in a brief network outage,
disconnect remote management software and affect running virtual
machines. In case IPv6 has been enabled or disabled this will
restart your host.

Apply changes and restart management network?

<Y> Yes <N> No                                <Esc> Cancel
```

Figure 3-48 Save settings and return to DCUI main menu

4. In the DCUI main menu, select **Test Management Network** and press Enter. Press Enter again to perform the tests. If some tests fail, see step 3 that is described in 3.3.2, “Configure management network for VMware ESXi”.

3.3.3 Joining hypervisors to the existing environment

Several choices are available to integrate a VMware vSphere environment that is deployed on BladeCenter with Flex System. You can join your new Flex System nodes to the existing BladeCenter vSphere cluster, or create a separate cluster that contains Flex System nodes only. These scenarios are described in this section.

Integrating new Flex System nodes with the existing vSphere cluster

Consider the following integration aspects:

- ▶ VMware ESXi hosts and new Flex System nodes with ESXi installed on them should access same shared storage.
- ▶ BladeCenter VMware ESXi hosts and new Flex System nodes with ESXi installed on them should have access to same networks (VLANs).
- ▶ Enhanced vMotion Compatibility (EVC) should be enabled on the cluster.

These aspects are required to have VMware vSphere clustered features working, including High Availability (HA) and Distributed Resource Scheduler (DRS).

Adding ESXi that is installed on the compute node to vCenter inventory

To join an ESXi host that is installed on the Flex System compute node to your VMware vCenter inventory, complete the following steps:

1. Log in to vCenter via Web Client and click **Hosts and Clusters**, as shown in Figure 3-49.

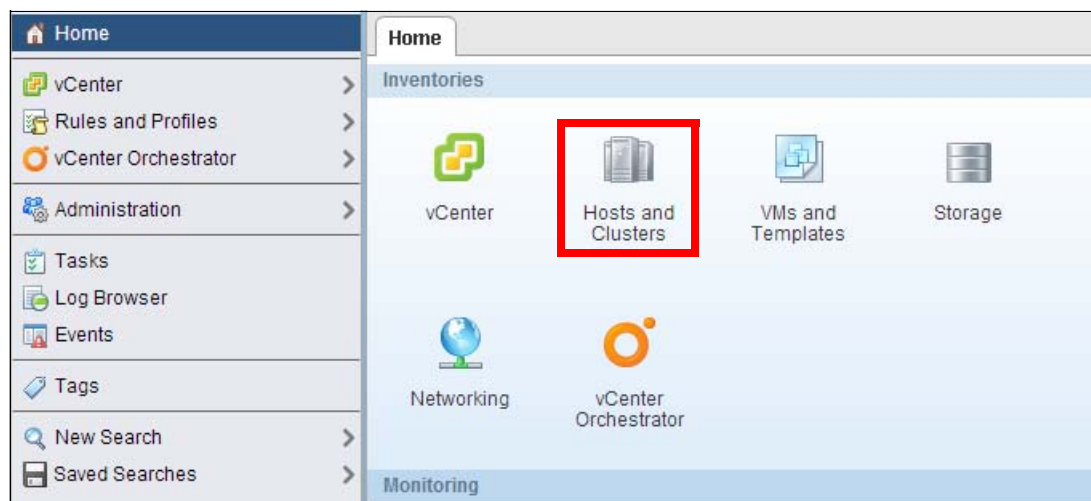


Figure 3-49 Hosts and Clusters view

2. Expand your vCenter object to see the data centers. Right-click the needed data center and click **Add Host**, as shown in Figure 3-50.

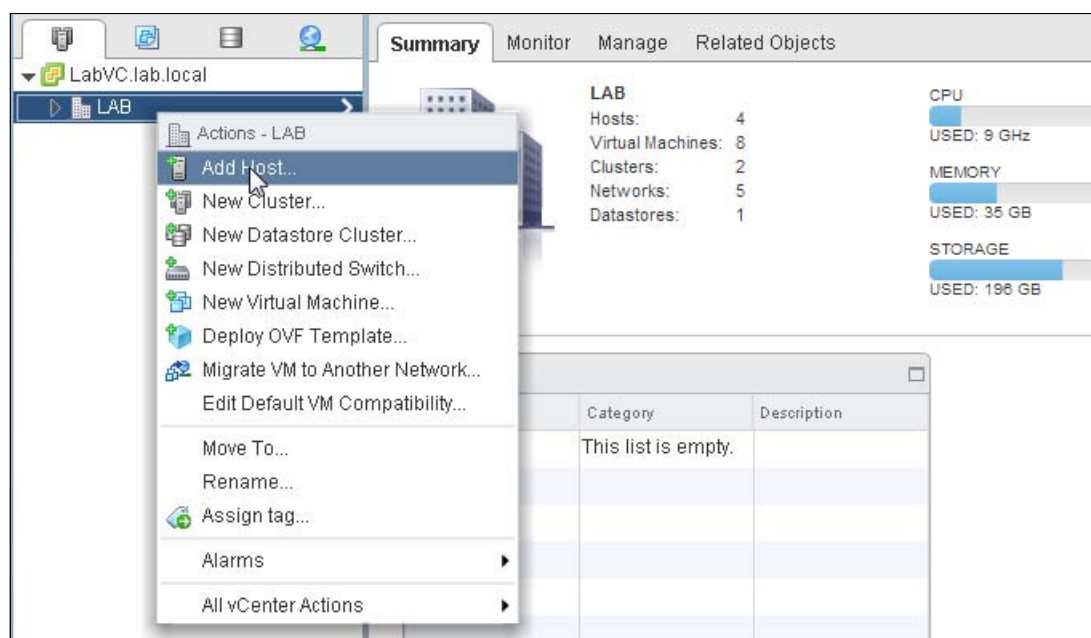


Figure 3-50 Add Host

3. The Add Host wizard opens. Complete the following steps:
 - a. Enter the host name of your VMware ESXi host that you want to add. Create a cluster or join existing clusters. Click **Next**.

- b. Enter root as the user credentials. Click **Next**.
- c. The Security Alert window opens, in which it is stated that it was not possible to verify the authority of SSL certificate. This issue is normal behavior when self-signed certificates are used. Click **Yes** to allow connection to the host.
- d. On the Host Summary page, click **Next**.
- e. On the Assign License page, select the license that you want to assign to this ESXi host, or select **(No License Key)**, if you want to activate Evaluation mode, as shown in Figure 3-51. Click **Next**.



Figure 3-51 Select License key to assign to ESXi host

- f. On the next page, select the Enable lockdown mode option, if required. Click **Next**.

Note: Enabling lockdown mode limits access to your ESXi host. For more information about enabling and disabling lockdown mode, see the VMware Knowledge Base article that is available at this website:

<http://kb.vmware.com/kb/1008077>

- g. If there are virtual machines (VMs) that are running on the ESXi host that you are adding, select the folders or resource pools in which those components should be placed. If no VMs are present, no action is required. Click **Next**.
- h. On the Ready to complete page, review the details and click **Finish**.

This operation takes 1 - 2 minutes to complete. After it is complete, you can see the added host in the vCenter inventory.

Joining ESXi installed on the compute node to the existing cluster

After the ESXi host is added to vCenter inventory, you can proceed with the configuration. The following basic configuration procedures for storage and network are described:

- ▶ Storage configuration, in: 3.2.3, "Configuring storage" on page 42
- ▶ Network configuration, in: 3.2.1, "Configuring the network" on page 34

If network and storage are not yet configured, it is recommended that you enter the ESXi host into Maintenance Mode before joining to the cluster. To enter Maintenance Mode, right-click the host in vSphere web client and click **Enter Maintenance Mode**.

To join ESXi host to the existing cluster, complete the following steps:

1. Right-click the host that you want to join to the existing cluster and click **Move To**.

2. A window opens. Expand the vCenter object, expand the data center object to which your cluster belongs, and select the cluster that you want to join.
3. Click **OK**.

Your new ESXi host is now member of your existing cluster.

Configuring virtual switch and creating needed port groups

In our scenario, we use VMware vNetwork Distributed Switch. If your environment is built by using VMware vNetwork Standard Switch or you need more in-depth information about VMware vDS, see the VMware official documentation.

The documentation for each version of VMware is available at the following websites:

- ▶ ESXi 5.0:
<http://pubs.vmware.com/vsphere-50/index.jsp>
- ▶ ESXi 5.1:
<http://pubs.vmware.com/vsphere-51/index.jsp>
- ▶ ESXi 5.5:
<http://pubs.vmware.com/vsphere-55/index.jsp>

For more information about vDS concepts, see this website:

<http://kb.vmware.com/kb/1010555>

To join ESXi host to the existing vDS, complete the following steps:

1. Log in to vSphere web client. From the home menu, click **Networking**.
2. Right-click the vDS to which you want to add the new ESXi host and click **Add and Manage Hosts**.
3. The Add and Manage Hosts wizard starts. Complete the following steps:
 - a. On the first page, select the **Add new and migrate host networking (advanced)** option and click **Next**.
 - b. On the second page, click **New Host** near the green + symbol and select the host that you want to add, as shown in Figure 3-52. Click **Next**.

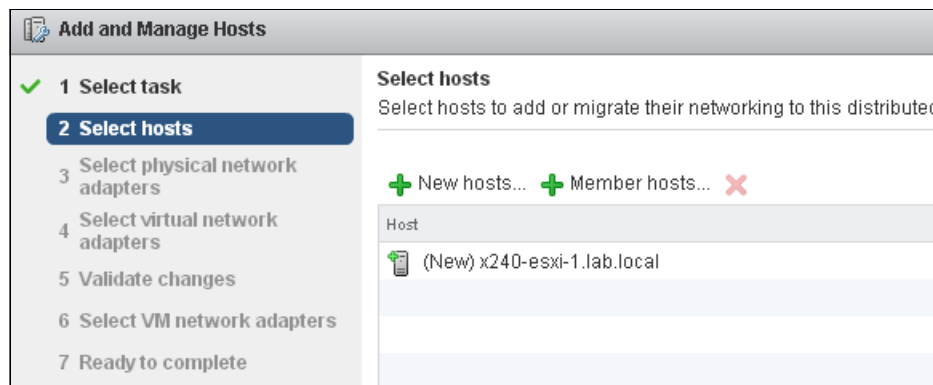


Figure 3-52 Add host to vDS

- c. On the next page, select the physical Network Interface Card that you want to use in your vDS as uplink. In our example, we are adding **vmnic0**, which is used as an uplink in vSS vSwitch0, as shown in Figure 3-53. Click **Next**.

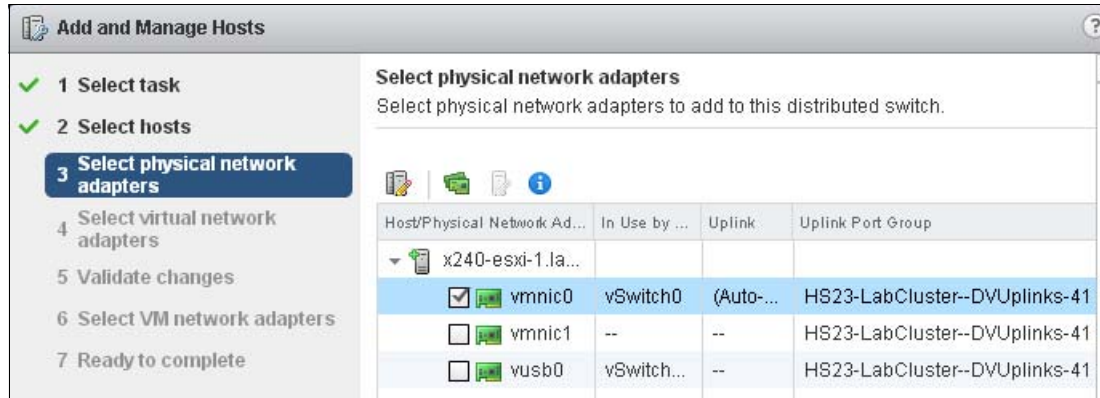


Figure 3-53 Adding NIC to vDS

- d. Select the Virtual Network Adapters to which you want to migrate vDS, if any. In our example, we migrate vmk0, which is used for Management Network on the host. To migrate the adapter, click the Virtual Network Adapter that you want to migrate, click the **Assign Port Group** icon, and select the Port group to which Virtual Network Adapter should be migrated, as shown in Figure 3-54. Click **Next**.

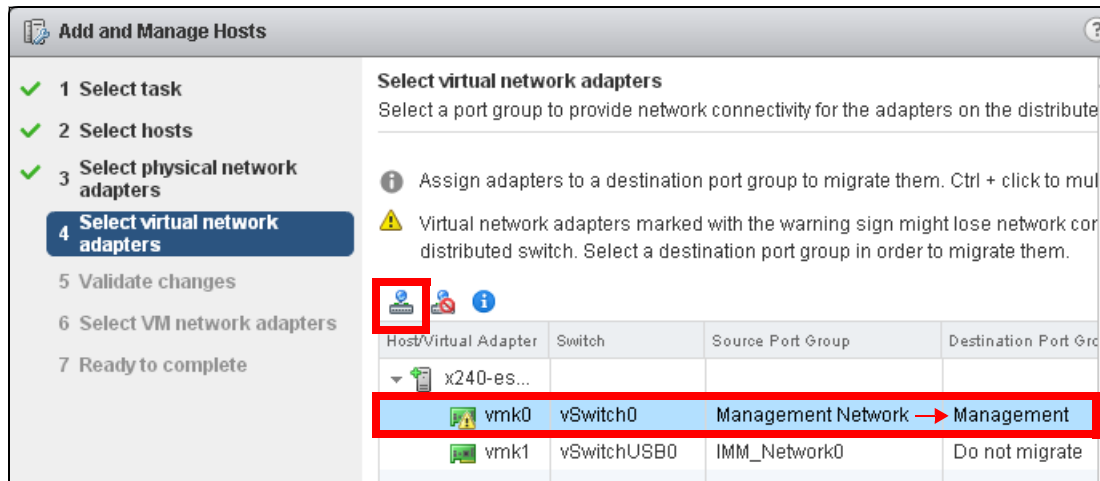


Figure 3-54 Migrate Virtual Network adapters vDS

- e. Verify your changes and that the Overall validation status is shown as Passed.
- f. If you have VMs running on the host to which you are adding vDS, you can perform VM network migration. (This step can be skipped for now.) Click **Next**.
- g. Click **Finish** on the Ready to complete page to apply your changes. The process takes 1 - 2 minutes to complete.

Creating a Virtual Network Adapter for vMotion

vMotion Port Group is required for creating Virtual Network Adapter with vMotion capabilities. Consider having vMotion network on an isolated VLAN in separate subnet. If such a solution is not possible, vMotion can work by using a management network.

In our example, we consider that vMotion Port group exists on the vDS.

To create a Virtual Network Adapter for vMotion, complete the following steps:

1. Log in to vSphere Web Client and enter the Hosts and Clusters view. Click the ESXi host on which you must create the Virtual Network Adapter.
2. In the central pane, click the **Manage** tab, click **Networking**, then select the **Virtual adapters** menu item and click the **Add host networking** icon, as shown in Figure 3-55.

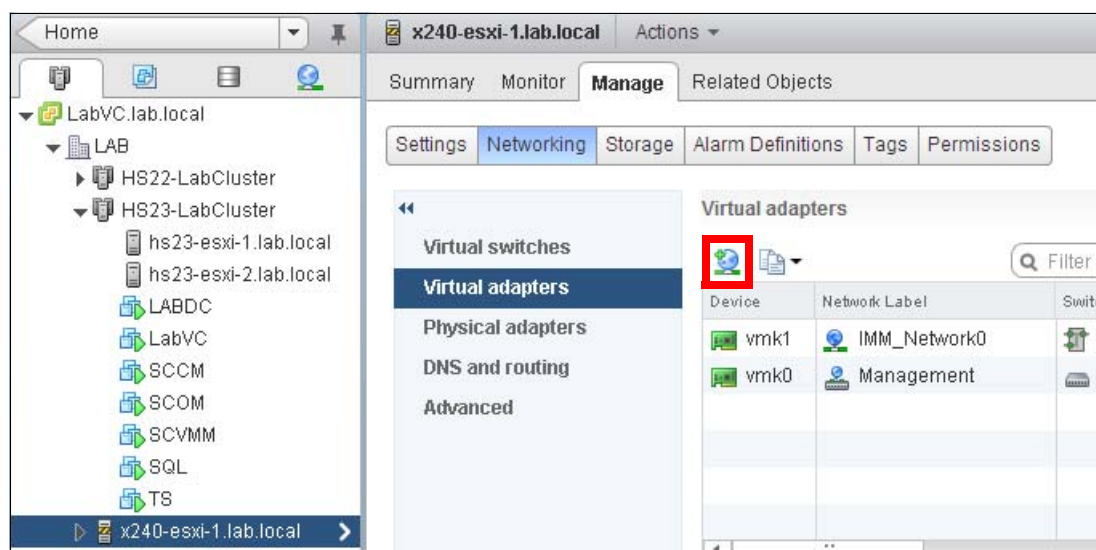


Figure 3-55 Click Add host networking icon in Virtual adapters menu

3. Configuration wizard opens. Complete the following steps:
 - a. On the Select connection type page of the wizard, select **VMkernel Network Adapter** and click **Next**.
 - b. On the Select target device page, select **Select and existing distributed port group** and click **Browse**. In the opened window, select the port group that you are planning to use for vMotion, as shown in Figure 3-56. Click **Next**.

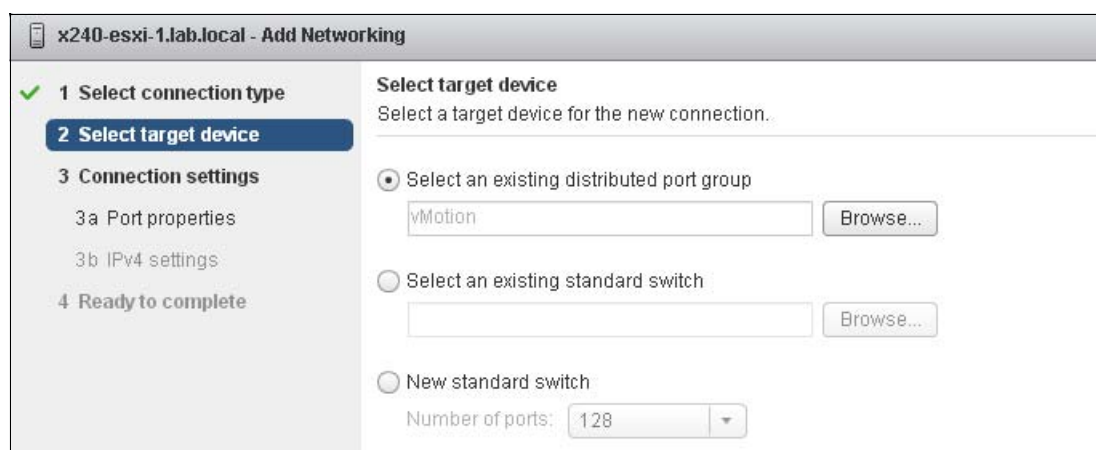


Figure 3-56 Select port group for vMotion

- c. On the Port properties page (as shown in Figure 3-57), enable this Virtual Network Adapter for vMotion by selecting the **vMotion traffic** option. In IP settings, select IPv4 or IPv6 (based on your environment). In our example, we use IPv4. Click **Next**.

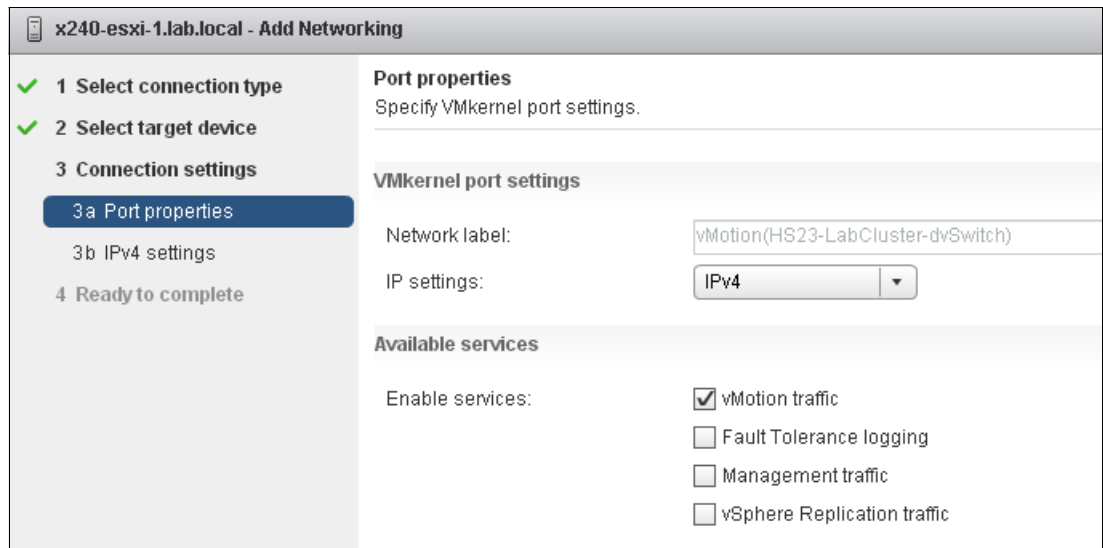


Figure 3-57 Enable vMotion traffic, and select IP version

- d. On the IPv4 settings page, configure your IP settings, depending on the IP protocol version that you selected in the previous step. Figure 3-58 shows an example of configuration for IPv4 with static IP assignment. Click **Next**.

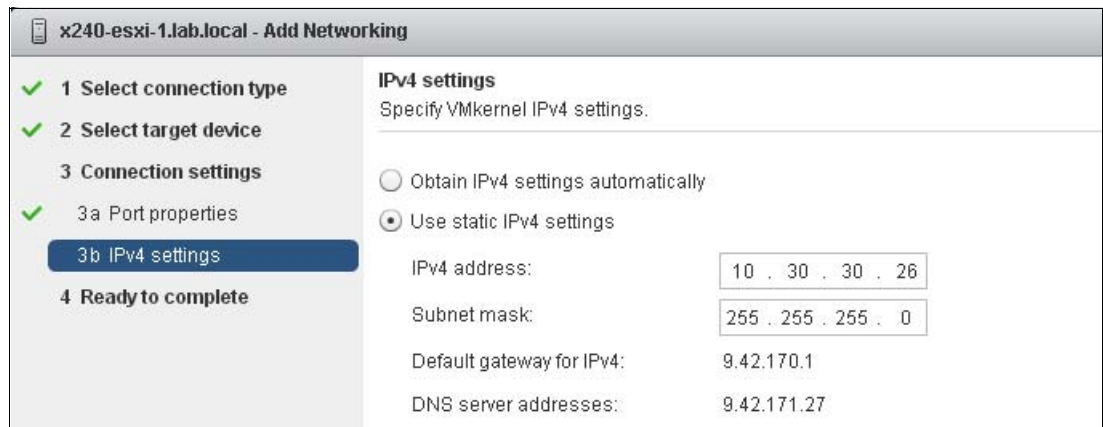


Figure 3-58 Set IP for vMotion

- e. On the Ready to complete page, review the configuration and click **Finish**.

The creation process takes several minutes. When finished, you can migrate VMs from your existing ESXi hosts to your newly added ESXi host that is installed on the Flex System Compute Node.

3.4 Integrating Microsoft Hyper-V

In this section, we describe the integration of the Microsoft Hyper-V server that is running on Flex System compute node into an existing Hyper-V cluster that is based on BladeCenter HS23. This approach can also be used to integrate Flex System into an existing physical (non-virtualized) Microsoft Windows Server based environment that is running on BladeCenter.

This section includes the following topics:

- ▶ 3.4.1, “BladeCenter environment”
- ▶ 3.4.2, “Deployment considerations” on page 59
- ▶ 3.4.3, “Connecting to a Flex System compute node using Remote Console” on page 60
- ▶ 3.4.4, “Installing the operating system” on page 61
- ▶ 3.4.5, “Management network configuration” on page 65
- ▶ 3.4.6, “Joining the node to System Center Virtual Machine Manager” on page 68
- ▶ 3.4.7, “VM live migration compatibility” on page 75

For more information about Microsoft solution considerations, see the following resources:

- ▶ System Center Documentation:
<http://technet.microsoft.com/en-us/library/cc507089.aspx>
- ▶ *Deploy a Hyper-V Cluster*:
<http://technet.microsoft.com/en-us/library/jj863389.aspx>
- ▶ *Network Recommendations for a Hyper-V Cluster in Windows Server 2012*:
<http://technet.microsoft.com/en-us/library/dn550728.aspx>

3.4.1 BladeCenter environment

Infrastructure is managed with Microsoft System Center 2012 R2 with Upward Integration For Microsoft System Center Bundle.

The following management components are installed:

- ▶ Microsoft System Center Operations Manager (SCOM)
- ▶ Microsoft System Center Virtual Machine Manager (SCVMM)
- ▶ Microsoft System Center Configuration Manager
- ▶ Lenovo Hardware Management Pack for Microsoft System Center Operations Manager
- ▶ Lenovo Hardware Performance and Resource Optimization Pack for Microsoft System Center Virtual Machine Manager
- ▶ Upward Integration Modules Add-in for Microsoft System Center Virtual Machine Manager
- ▶ Lenovo System Updates for Microsoft System Center Configuration Manager
- ▶ Lenovo Inventory Tool for Microsoft System Center Configuration Manager

Our Hyper-V cluster consists of two HS23 servers that are connected to the SAN storage by using iSCSI.

There are four integrated NICs in the HS23 default configuration (2x 1 GbE and 2x 10 GbE) that are connected in the following way (see Figure 3-59):

- ▶ 1 GbE NICs are teamed and used for the servers in-band management (AD Domain Network).
- ▶ 10 GbE NICs are teamed by using SCVMM Logical Switch with virtual adapters that are created on top of the Hyper-V extensible switch that is used for Management OS and VMs traffic. VLAN and network subnet separation is used to isolate traffic.

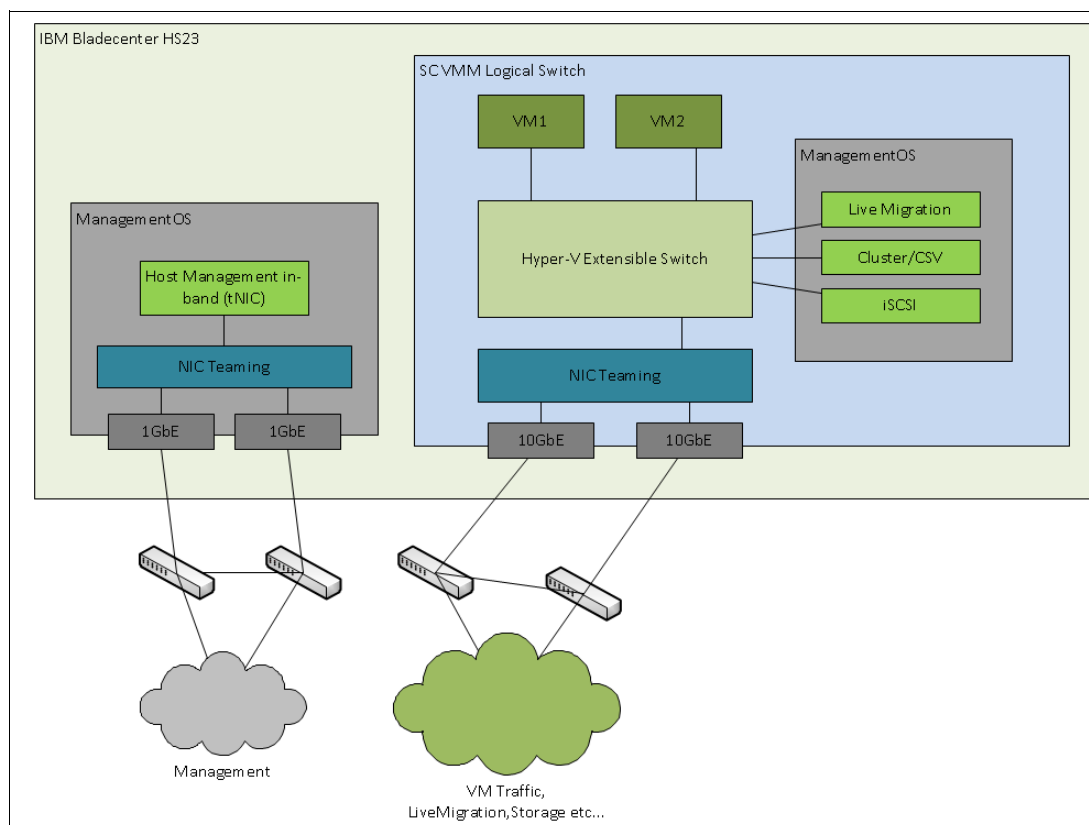


Figure 3-59 Hyper-V network topology

The operating system is installed on the local hard disk drives.

3.4.2 Deployment considerations

Although Flex System compute nodes can be considered as similar to BladeCenter nodes, the following differences must be considered for merging systems into one Hyper-V cluster:

- ▶ Processor compatibility for live migration: As live migration is performed on the online VM, migration between the processors with different instruction sets can cause a VM to crash. Migration between different vendor processor family is not supported. Consider the following points:
 - Mixed clusters with different CPU versions are supported by using CPU Compatibility mode; however, such clusters might not benefit from new generation processor features.
 - To achieve best performance, consider creating separate clusters for each processor that is generated.

- Network configuration: Every Hyper-V Switch must exist across all members of the cluster; therefore, if your current environment has three Hyper-V Switches, you also must deploy them on Flex System compute nodes. Each Hyper-V switch in the external mode requires a dedicated uplink interface (teamed or stand-alone).

Providing this dedicated uplink interface might be a challenge; for example, some Flex System x240 nodes include dual-port embedded 10 GbE LOM adapters in the default configuration instead of four adapter ports in the HS23. For such cases, you can use the Unified Fabric Port (UFP) feature of the Flex System to create up to four virtual NICs per physical NIC.

3.4.3 Connecting to a Flex System compute node using Remote Console

Use the CMM web interface for connecting to the Flex System node. Right-click the wanted node and click **Launch Compute Node Console**, as shown in Figure 3-60. Then, click **Launch** and you are redirected to the IMM interface of the node.

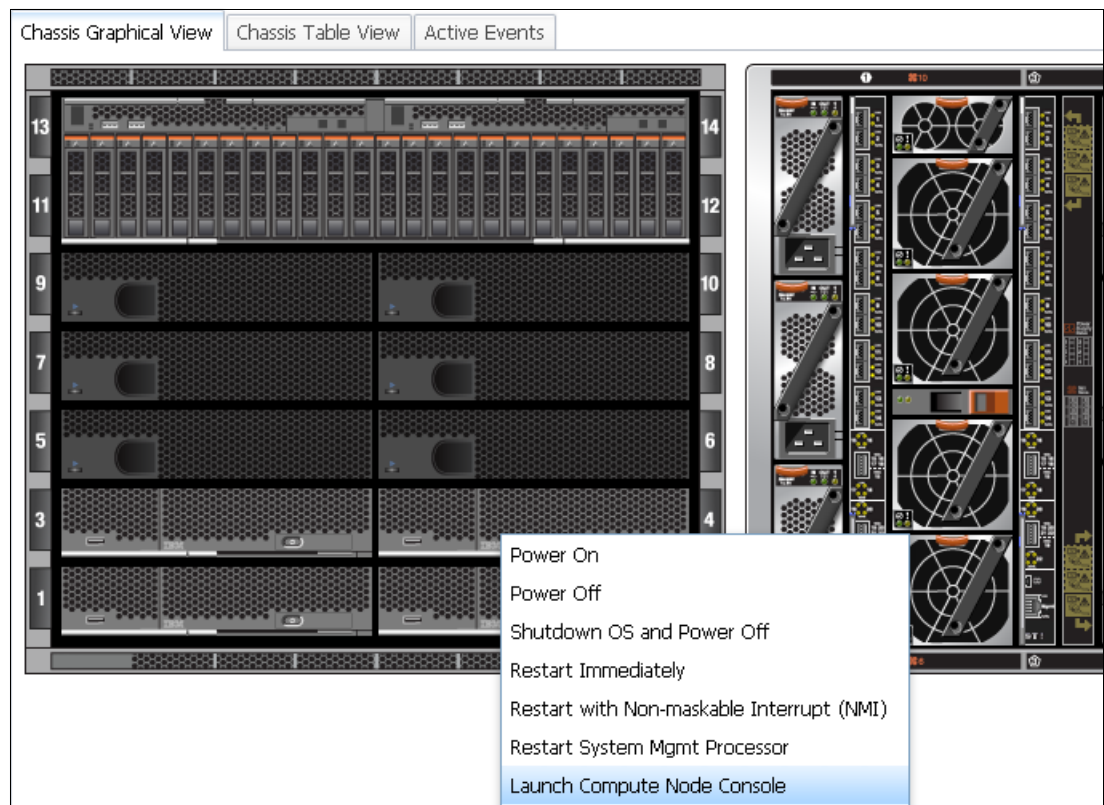


Figure 3-60 Selecting the Launch Compute Node Console option

From the IMMv2, you can attach virtual media for the operating system installation. You also can control the power state of the system. You can start the console in single- or multi-user mode, as shown in Figure 3-61.

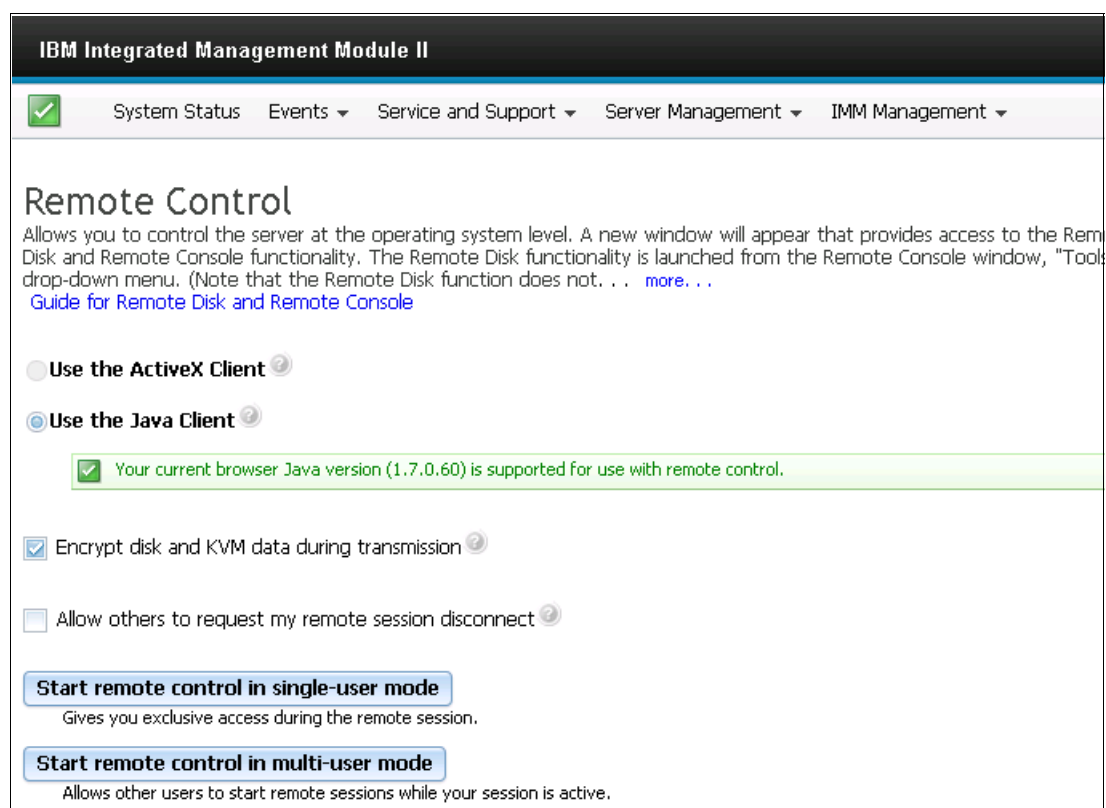


Figure 3-61 Start remote control

For more information about the use of the IMMv2 web interface, see this website:

<https://www-947.ibm.com/support/entry/myportal/docdisplay?lnidocid=MIGR-5086346>

3.4.4 Installing the operating system

To include all required drivers and configure system components with minimal user intervention, consider the use of ServerGuide when Microsoft Windows Server is deployed. For more information, see this website:

<https://www-947.ibm.com/support/entry/myportal/docdisplay?lnidocid=SERV-GUIDE>

Alternatively, if your environment is configured to use SCCM to deploy operating systems on blades, you can use Lenovo Deployment Pack For MS SCCM. For more information, see this website:

<https://www-947.ibm.com/support/entry/myportal/docdisplay?lnidocid=MIGR-5082208>

Deploying by using ServerGuide boot image

To deploy a compute node by using ServerGuide, complete the following steps:

1. Connect to the IMM2 remote console of the node.
2. Power on the server by clicking **Video Viewer Tools** → **Power** → **Power On**.

3. Press F1 when prompted to enter the UEFI setup and configure virtual NICs in IBM UFP mode (if needed). You can also configure hardware iSCSI and FCoE. For more information about storage configuration, see 3.2.3, “Configuring storage” on page 42.
4. Save the settings and reboot the server.
5. Start the Virtual Media and add a ServerGuide image in Video Viewer by clicking **Tools** → **Virtual Media** → **Add Image** → **Mount Selected**, as shown in Figure 3-62.




Figure 3-62 Mount ISO

6. If you did not set the boot order before, press F12 when prompted to select the boot device from which to boot (CD/DVD), as shown in Figure 3-63.




Figure 3-63 Boot settings

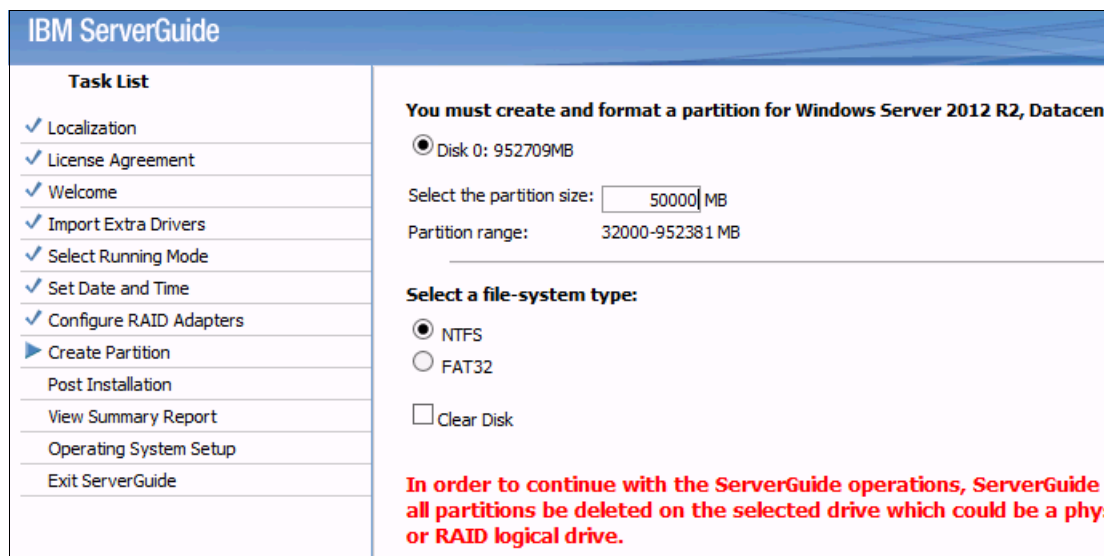
- After ServerGuide is booted, select language, localization, and accept the license to continue. Then, select the running mode, as shown in Figure 3-64.

Figure 3-64 Select running mode

- Select **RAID configuration and Windows OS deployment**, select the wanted operating system and version, and click **Next**.
- Configure the Date and Time settings.
- Select the wanted RAID configuration or keep the existing configuration (see Figure 3-65).

Figure 3-65 Configure RAID

11. Select the wanted partition size for the operating system and click **Next**, as shown in Figure 3-66.

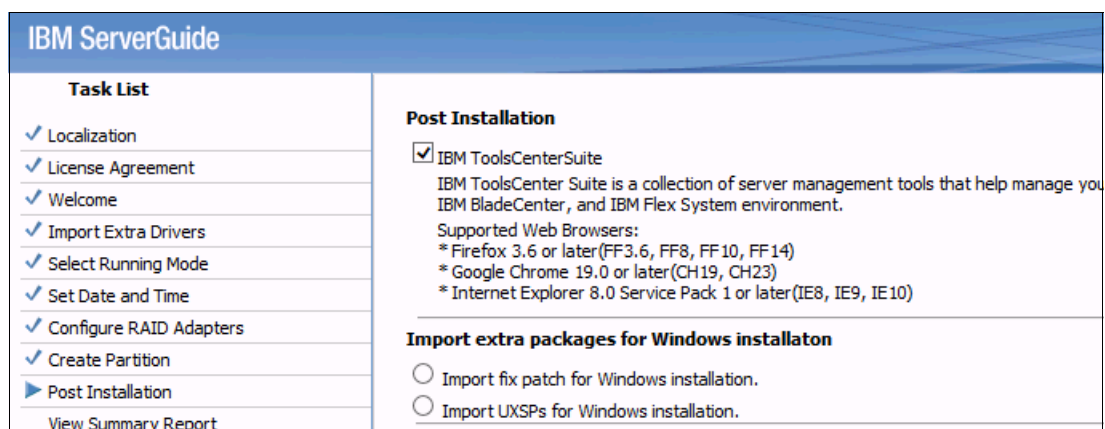


IBM ServerGuide

Task List	Configuration
✓ Localization	<p>You must create and format a partition for Windows Server 2012 R2, Datacenter</p> <p><input checked="" type="radio"/> Disk 0: 952709MB</p> <p>Select the partition size: <input type="text" value="50000"/> MB</p> <p>Partition range: 32000-952381 MB</p> <p>Select a file-system type:</p> <p><input checked="" type="radio"/> NTFS</p> <p><input type="radio"/> FAT32</p> <p><input type="checkbox"/> Clear Disk</p> <p>In order to continue with the ServerGuide operations, ServerGuide all partitions be deleted on the selected drive which could be a physical or RAID logical drive.</p>
✓ License Agreement	
✓ Welcome	
✓ Import Extra Drivers	
✓ Select Running Mode	
✓ Set Date and Time	
✓ Configure RAID Adapters	
▶ Create Partition	
Post Installation	
View Summary Report	
Operating System Setup	
Exit ServerGuide	

Figure 3-66 Partition size

12. You can include ToolsCenter Suite, which can help you with managing a System x server. You also can include more Windows hotfixes and UpdateXpress updates, as shown in Figure 3-67. Click **Next**.

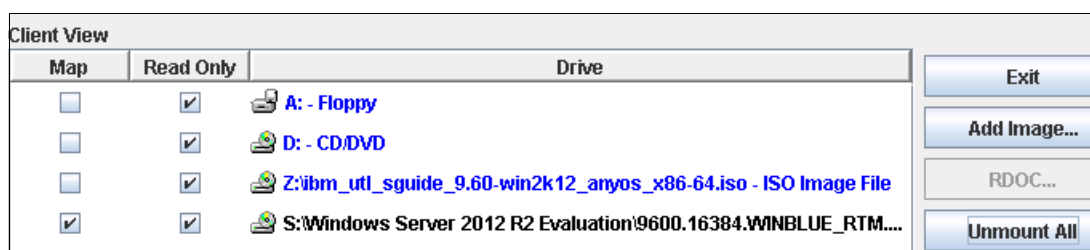


IBM ServerGuide

Task List	Post Installation
✓ Localization	<p>Post Installation</p> <p><input checked="" type="checkbox"/> IBM ToolsCenter Suite</p> <p>IBM ToolsCenter Suite is a collection of server management tools that help manage your IBM BladeCenter, and IBM Flex System environment.</p> <p>Supported Web Browsers:</p> <ul style="list-style-type: none"> * Firefox 3.6 or later (FF3.6, FF8, FF10, FF14) * Google Chrome 19.0 or later (CH19, CH23) * Internet Explorer 8.0 Service Pack 1 or later (IE8, IE9, IE10) <p>Import extra packages for Windows installation</p> <p><input type="radio"/> Import fix patch for Windows installation.</p> <p><input type="radio"/> Import UXSPs for Windows installation.</p>
✓ License Agreement	
✓ Welcome	
✓ Import Extra Drivers	
✓ Select Running Mode	
✓ Set Date and Time	
✓ Configure RAID Adapters	
✓ Create Partition	
▶ Post Installation	
View Summary Report	

Figure 3-67 Postinstallation tasks: IBM ToolsCenter Suite

13. Review summary report and mount the operating system image when prompted. You must unmount the ServerGuide image first by clicking **Unmount All**, as shown in Figure 3-68.



Client View

Map	Read Only	Drive	
<input type="checkbox"/>	<input checked="" type="checkbox"/>	A: - Floppy	<p>Exit</p> <p>Add Image...</p> <p>RDOC...</p> <p>Unmount All</p>
<input type="checkbox"/>	<input checked="" type="checkbox"/>	D: - CD/DVD	
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Z:\ibm_util_sgguide_9.60-win2k12_anyos_x86-64.iso - ISO Image File	
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	S:\Windows Server 2012 R2 Evaluation\9600.16384.WINBLUE_RTM....	

Figure 3-68 Map operating system installation image

- 14. Follow instructions and enter the operating system details, accept the license, and click **Next** to install the operating system.
- 15. After the installation process completes, log in to the operating system and wait until any postinstallation tasks complete. Then, reboot the operating system.

3.4.5 Management network configuration

Note: Depending on your virtual NIC configuration, you might see multiple adapters in your operating system after installation. In our example, we use only one port on Emulex 10GbE VFA (the second is disabled in UEFI) by using virtual NICs with iSCSI personality. If we do not use the iSCSI personality of the adapter, there are four NICs in total; however, iSCSI or FCoE personality is always assigned to the second vNIC function (vPort2) of particular physical NIC and that port cannot be used for anything else. “IBM USB Remote NDIS Network Device” is the server internal network that is used for server management tasks, such as online firmware updates.

If you open Network Connections, you see a list of available network interfaces, as shown in Figure 3-69.

Name	Status	Device Name	Connectivity
Ethernet	Unidentified network	Emulex OCI11102-F-X Virtual Fa...	No Internet a
Ethernet 2	Identifying...	Emulex OCI11102-F-X Virtual Fa...	No network .
Ethernet 3	Unidentified network	Emulex OCI11102-F-X Virtual Fa...	No network .
Local Area Connection	Disabled	IBM USB Remote NDIS Network ...	

Figure 3-69 Available network connections after the operating system is installed

You can use Windows Device Manager to identify adapters in the Windows operating system, as shown in Figure 3-70. Device 0 means pNIC 1 and Function 6 means vPort 4.

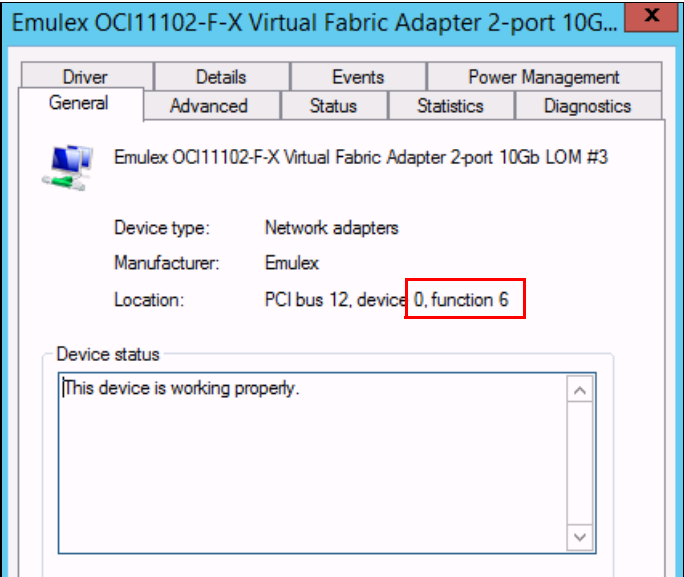


Figure 3-70 Network Function Identification in Device Manager

Complete the following steps to configure your management network:

1. Rename the adapters as wanted, as shown in Figure 3-71.

Name	Status	Device Name
Conv1	Unidentified network	Emulex OCI11102-F-X Virtual Fabric Adapter 2-port 10Gb LOM
Ethernet 3	Network cable unplugg...	Emulex OCI11102-F-X Virtual Fabric Adapter 2-port 10Gb LOM #3
Local Area Connection	Disabled	IBM USB Remote NDIS Network Device
Mngmt1	Network	Emulex OCI11102-F-X Virtual Fabric Adapter 2-port 10Gb LOM #2

Figure 3-71 Renaming network connections

2. Create NIC Team for the Management (Domain) network, as shown in Figure 3-72.
Click **OK**.

New team

Team name:

Member adapters:

In Team	Adapter	Speed	State	Reason
<input type="checkbox"/>	Conv1	10 Gbps		
<input type="checkbox"/>	Ethernet 3	Disconnected		
<input type="checkbox"/>	Local Area Connection	Disabled		
<input checked="" type="checkbox"/>	Mngmt1	2 Gbps		

Additional properties

Teaming mode: Switch Independent

Load balancing mode: Dynamic

Standby adapter: None (all adapters Active)

Primary team interface: [MngmtTeam: Default VLAN](#)

Figure 3-72 Create team for management network

The team is created, as shown in Figure 3-73.

Name	Status	Device Name
Conv1	Unidentified network	Emulex OCI11102-F-X Virtual Fabric Adapter 2-port 10Gb LOM
Ethernet 3	Network cable unplugg...	Emulex OCI11102-F-X Virtual Fabric Adapter 2-port 10Gb LOM #3
Local Area Connection	Disabled	IBM USB Remote NDIS Network Device
Mngmt1	Enabled	Emulex OCI11102-F-X Virtual Fabric Adapter 2-port 10Gb LOM #2
MngmtTeam	Identifying...	Microsoft Network Adapter Multiplexor Driver

Figure 3-73 Team created

3. Configure the IP address for the teamed interface, as shown in Figure 3-74. Click **OK**.

General

You can get IP settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IP settings.

☐ Obtain an IP address automatically

☒ Use the following IP address:

IP address: 9 . 42 . 171 . 36

Subnet mask: 255 . 255 . 254 . 0

Default gateway: 9 . 42 . 170 . 1

☐ Obtain DNS server address automatically

☒ Use the following DNS server addresses:

Preferred DNS server: 9 . 42 . 171 . 27

Alternate DNS server: . . .

☐ Validate settings upon exit

Advanced...

Figure 3-74 Configure Management IP

Depending on your VLAN tagging configuration on the port, you might need to enable VLAN tagging in the device manager for the particular adapter. For example, if you must use VLAN 20 and 30 on the same port, you must configure it on each member of the team, as shown in Figure 3-75.

Driver Details Events Power Management

General Advanced Status Statistics Diagnostics

EMULEX Advanced Property Configuration

Configuration

- Class of Service (802.1p)
- Network Address
- Packet Size
- VLAN Identifier (802.1q)

Performance

- CPU Affinity
- Interrupt Moderation
- Performance Tuning
- Receive Buffers
- RSS Base Processor Group
- RSS Profile
- SR-IOV

Value:

☒ 20 30

☐ Not Present

Set to Default

Reset All to Default

Figure 3-75 Set VLAN tagging in Emulex driver

After you have your connection in the domain network working, you can join the node to the domain.

3.4.6 Joining the node to System Center Virtual Machine Manager

There are some prerequisites that must be met before you can join the computer to the Microsoft System Center environment, such as administrator privileges on the hosts and open ports for communication between management server. There are also multiple location and security scenarios that might require more configuration steps.

In this section, we describe a scenario in which the server is joined in a trusted Active Directory domain. For more information about the requirements and all scenarios, see this website:

<http://www.microsoft.com/en-us/download/details.aspx?id=6346>

Assuming that all prerequisites were met, you can add node into the SCVMM console. Complete the following steps:

1. In the Fabric Resources View, click **Add Resources** → **Hyper-V Hosts and Clusters**, as shown in Figure 3-76.




Figure 3-76 Add Hyper-V Hosts into SCVMM

2. In the Add Resource wizard, select **Windows Server computers in a trusted Active Directory domain**, as shown in Figure 3-77. Click **Next**.



Figure 3-77 Select computer location

3. Specify the credentials for discovery, as shown in Figure 3-78. Click **Next**.

Resource Location

Credentials

Discovery Scope

Target resources

Host Settings

Summary

Specify the credentials to use for discovery

The Run As account or credentials will be used to discover computers and to install the Hyper-V the Virtual Machine Manager agent if necessary.

☒ Use an existing Run As account

Run As account:

☐ Manually enter the credentials

User name:

Example: contoso\domainuser

Password:

Figure 3-78 Specify discovery credentials

4. Define the discovery scope that is based on your environment, as shown in Figure 3-79. Click **Next**.

Resource Location

Credentials

Discovery Scope

Target resources

Host Settings

Summary

Specify the search scope for virtual machine host candidates

Search for computers by whole or partial names, FQDNs, and IP addresses. Alternatively, you may generate an Active Directory query to discover the desired computers.

☒ Specify Windows Server computers by names

☐ Specify an Active Directory query to search for Windows Server computers

Enter the computer names of the hosts or host candidates that you want VMM to manage. Each computer name must be on a separate line.

Computer names:

☐ Skip AD verification

Examples: server1
server1.contoso.com
10.0.1.1
2a01:110:1e:3:f8ffce44:23

Figure 3-79 Define scope of discovery

5. Select the wanted computers, as shown in Figure 3-80.

Resource Location

Credentials

Discovery Scope

Target resources

Host Settings

Summary

Select the computers that you want to add as hosts

Discovered computers:

Computer Name	Operating System	Hypervisor
<input checked="" type="checkbox"/> x240-hyper-v-1.lab.local	Windows Server 2012 R2 Datacenter E...	Unknown

Figure 3-80 Select hosts

- Click **Next**. If there is no Hyper-V role that is installed in the server, SCVMM prompts you for confirmation to install and reboot the server.
- Specify the target Host Group (this group should be the group with the existing cluster); however, you can change this selection later. Optionally, you can define the default VM placement paths, as shown in Figure 3-81. Click **Next**.

Figure 3-81 Specify target host group

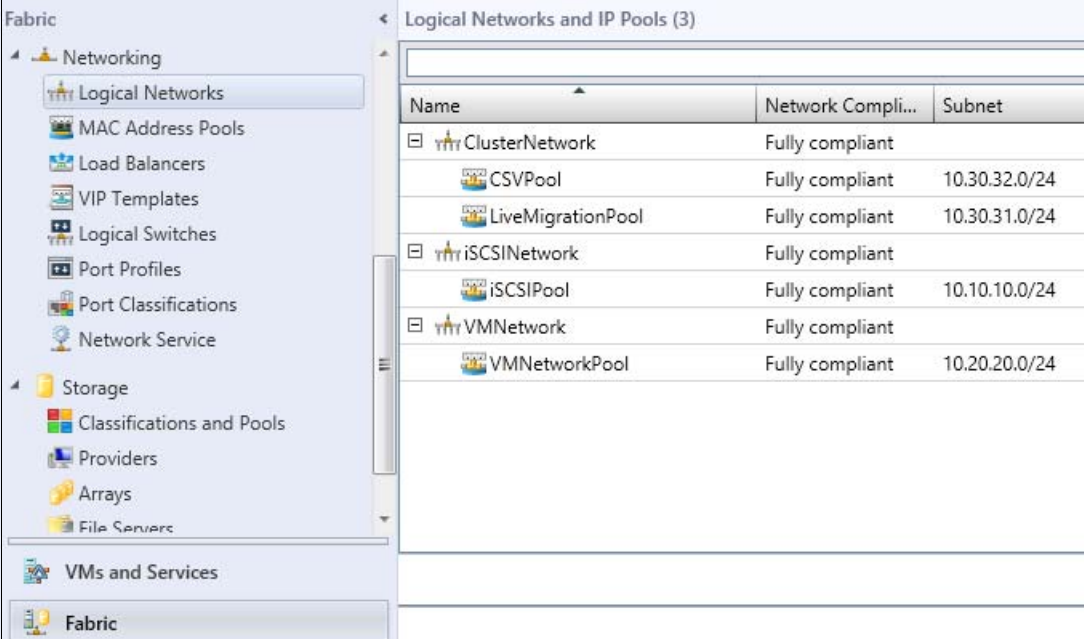
- Review the summary and confirm the settings, as shown in Figure 3-82. Wait for the joining process to finish; your host might be rebooted several times during the process.

Step	Name	Status
1	Add virtual machine host	Completed
1.1	Create undeployed host	Completed
1.2	Install Virtual Machine Manager agent	Completed
1.3	Refresh host	Completed
1.4	Enable Hyper-V	Completed
1.5	Install SCVMM Switch Port Settings	Completed
1.6	Install virtual switch extension	Completed
1.6.1	Deploy driver and install virtual switch...	Completed
1.6.1.1	Deploy file (using LAN)	Completed

Figure 3-82 Adding a host to SCVMM

Configuring cluster networks by using logical switch

We have the existing logical networks managed by SCVMM, as shown in Figure 3-83.

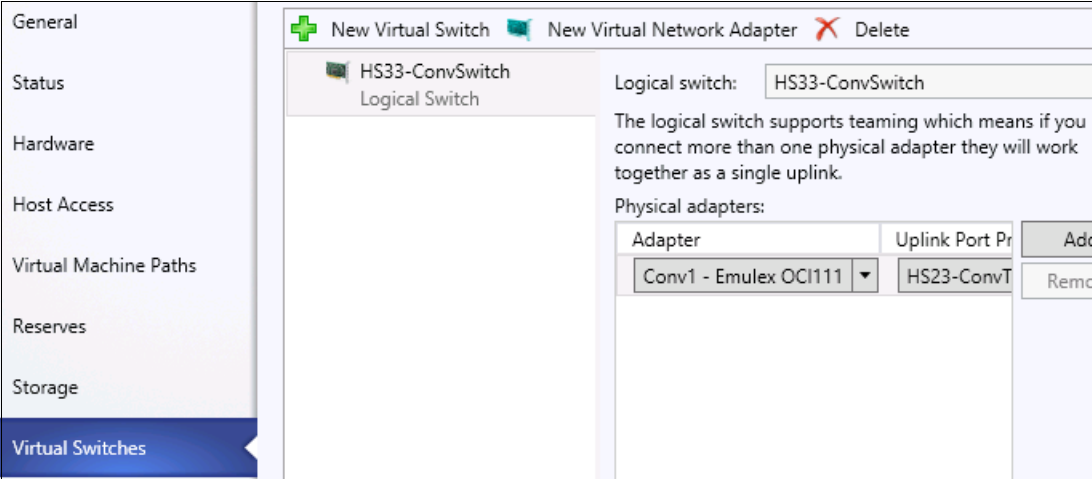


Name	Network Compli...	Subnet
ClusterNetwork	Fully compliant	
CSVPool	Fully compliant	10.30.32.0/24
LiveMigrationPool	Fully compliant	10.30.31.0/24
iSCSINetwork	Fully compliant	
iSCSIPool	Fully compliant	10.10.10.0/24
VMNetwork	Fully compliant	
VMNetworkPool	Fully compliant	10.20.20.0/24

Figure 3-83 Logical Networks settings

Complete the following steps to configure cluster networks by using logical switch:

1. To join the host into the existing logical switch, right-click the host (Fabric Resource Hosts view) that you want to configure and click **Properties** → **Virtual Switches** → **New Virtual Switch** → **Logical Switch**, as shown in Figure 3-84.



Adapter	Uplink Port Pr	Add	Remo
Conv1 - Emulex OCI111	HS23-ConvT		

Figure 3-84 Selecting the logical switch

2. Select the adapters that acts as uplinks for the logical switch. If you add multiple adapters, the adapters are teamed together.

Note: In our environment, we use virtual adapters for the cluster and live migration network (see Figure 3-59 on page 59).

3. To add the new virtual adapter, select logical switch to which this adapter is connected and click **New Virtual Network Adapter**. Specify the VM Network, Subnet VLAN, Port profile, and define the IP address (if needed), as shown in Figure 3-85.

The screenshot shows the VMware vSphere Client interface. On the left, a navigation pane lists various settings, with 'Virtual Switches' selected. The main area displays a tree view of network components: 'HS33-ConvSwitch' (Logical Switch) and 'CSV' (CSVNetwork). The 'CSV' component is selected, and the 'New Virtual Network Adapter' dialog is open. The dialog has tabs for 'New Virtual Switch', 'New Virtual Network Adapter', and 'Delete'. The 'New Virtual Network Adapter' tab is active, showing the following configuration:

- Name:** CSV
- ☐ This virtual network adapter inherits settings from the physical management adapter.
- Connectivity**
 - VM Network: CSVNetwork (with a 'Browse...' button)
 - VM Subnet: CSVNetwork_0
 - ☒ Enable VLAN
 - VLAN ID: 30
- Port profile**
 - Classification: Host Cluster Workload
- IP address configuration**
 - ☐ DHCP
 - ☒ Static
 - IPv4 pool: CSVPool (10.30.32.1 - 10.30.32.254)
 - IPv4 address: 10.30.32.36
 - IPv6 pool: Not Applicable

Figure 3-85 Configure Cluster/CSV network adapter

4. Depending on the number of required virtual adapters, repeat the previous step for each adapter, as shown in Figure 3-86.

The screenshot shows the VMware vSphere Client interface, similar to Figure 3-85, but for configuring a Live Migration network adapter. The 'LiveMigration' component is selected in the tree view, and the 'New Virtual Network Adapter' dialog is open. The configuration is as follows:

- Name:** LiveMigration
- ☐ This virtual network adapter inherits settings from the physical management adapter.
- Connectivity**
 - VM Network: LiveMigrationNetwork (with a 'Browse...' button)
 - VM Subnet: LiveMigrationNetwork_0
 - ☒ Enable VLAN
 - VLAN ID: 30
- Port profile**
 - Classification: Live migration workload
- IP address configuration**
 - ☐ DHCP
 - ☒ Static
 - IPv4 pool: LiveMigrationPool (10.30.31.1 - 10.30.31.254)
 - IPv4 address: 10.30.31.36
 - IPv6 pool: Not Applicable

Figure 3-86 Configure Live Migration network adapter

- After you are finished, click **OK**.
- You can verify that the new configuration was applied successfully directly on the server, as shown in Figure 3-87.

PROPERTIES For x240-hyper-v-1	
Computer name	x240-hyper-v-1
Domain	lab.local
Windows Firewall	Domain: Off, Public: Off
Remote management	Enabled
Remote Desktop	Enabled
NIC Teaming	Enabled
MngmtTeam	9.42.171.36, IPv6 enabled
vEthernet (CSV)	10.30.32.36, IPv6 enabled
vEthernet (LiveMigration)	10.30.31.36, IPv6 enabled

Figure 3-87 IP configuration of the virtual adapters

Joining servers to the existing Hyper-V Cluster

After all new Flex System servers have the same configuration and working connections with the current members of the cluster, you can add them as new members into the cluster.

In addition to the network connection, all nodes in the cluster must have access to the shared cluster storage. For more information about storage configuration, 3.2.3, “Configuring storage” on page 42.

Complete the following steps to join the compute nodes to the cluster:

- In the Fabric Resources Servers view, right-click the target cluster and click **Add Cluster Node**, as shown in Figure 3-88.

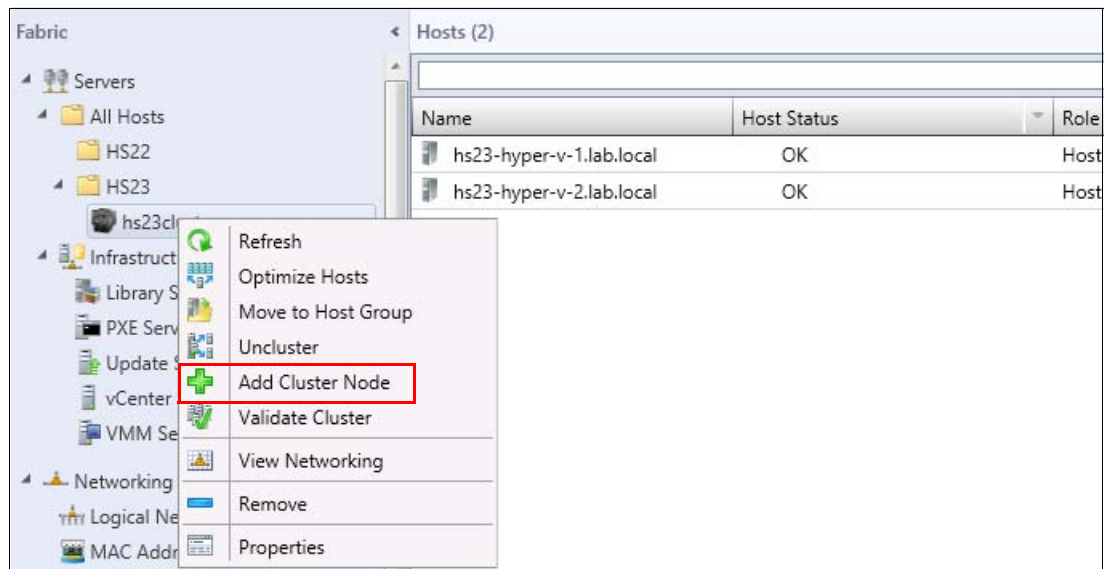


Figure 3-88 Add Cluster Node option

2. Select the wanted hosts (as shown in Figure 3-89) and click **Add**.

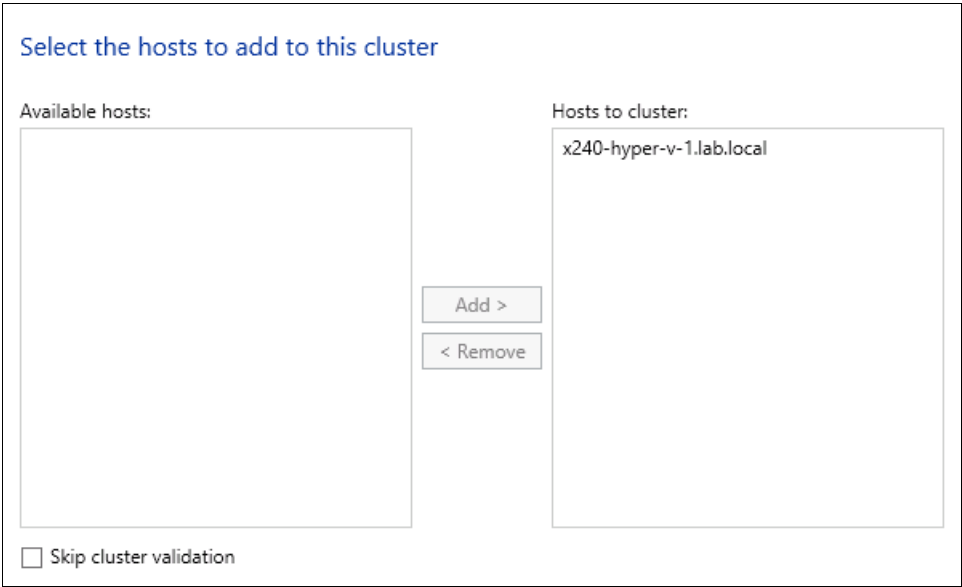


Figure 3-89 Select available hosts

3. Wait for SCVMM to finish (which includes the cluster validation process), as shown in Figure 3-90.


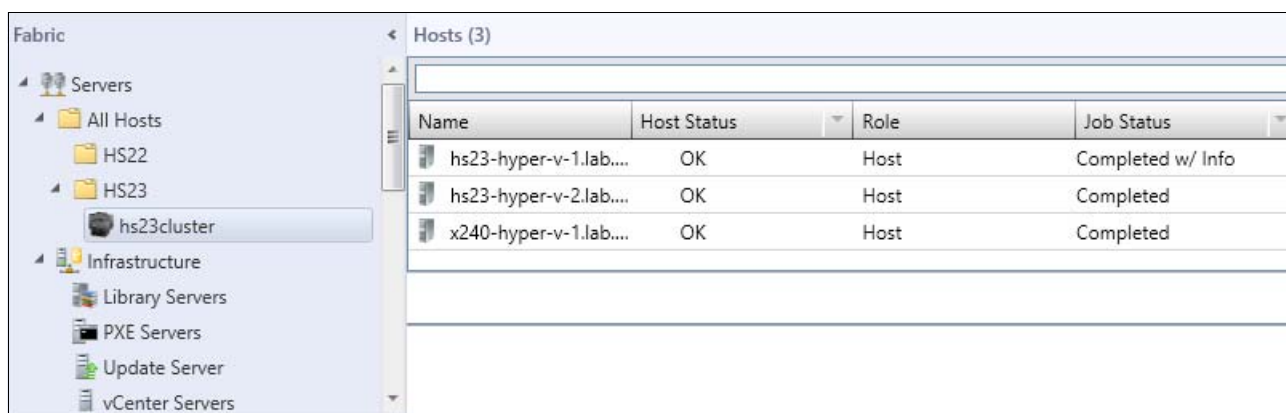


Figure 3-90 Add host to cluster job

The host is added to the cluster, as shown in Figure 3-91.



The screenshot shows the vSphere interface with the 'Fabric' tree on the left and the 'Hosts (3)' view on the right. The 'Hosts (3)' view displays a table with the following data:

Name	Host Status	Role	Job Status
hs23-hyper-v-1.lab....	OK	Host	Completed w/ Info
hs23-hyper-v-2.lab....	OK	Host	Completed
x240-hyper-v-1.lab....	OK	Host	Completed

Figure 3-91 Cluster Hosts view

3.4.7 VM live migration compatibility

If you created a cluster with the hosts that contain different processor versions, consider enabling the **Allow migration to a virtual machine host with a different processor version** option to ensure Live Migration compatibility between these hosts.

To enable this option, click **Virtual Machine Properties** → **Hardware Configuration** → **Processor**. Then, select the option, as shown in Figure 3-92.

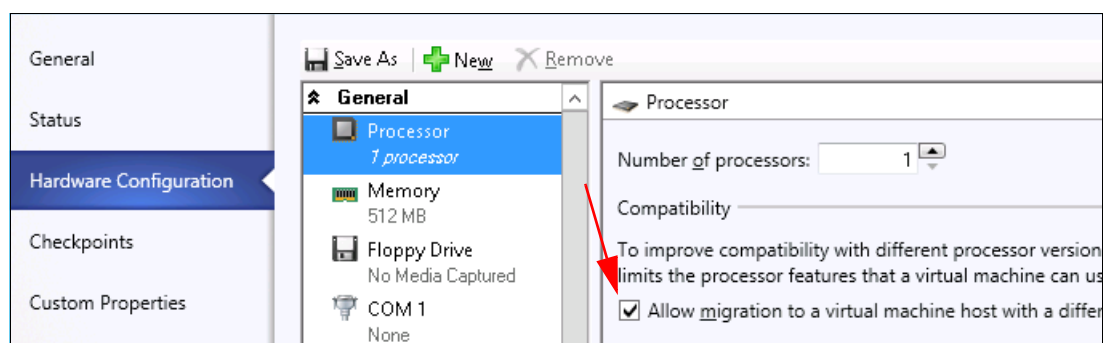


Figure 3-92 VM CPU compatibility

Note: This setting hides some of the CPU features from VMs to maintain compatibility, which can negatively affect the performance of your application. For more information about the potential affect of this setting on your application, consult with the application vendor.



Managing a combined Flex System and BladeCenter environment

In this chapter, we describe the use of Flex System and BladeCenter management extensions for the management tools for physical and virtual environments. Specifically, we describe the use of UIM for VMware vCenter to manage vSphere-based BladeCenter and Flex System environment, and UIM for Microsoft System Center to manage Microsoft Windows Server based physical and virtual environments.

This chapter includes the following topics:

- ▶ 4.1, “Managing a vSphere environment with UIM” on page 78
- ▶ 4.3, “Managing a Windows Server environment with UIM” on page 95

4.1 Managing a vSphere environment with UIM

When a BladeCenter environment is integrated with Flex System environment, consider the use of System x Upwards Integration Modules (UIMs) for VMware vSphere. By using UIMs for VMware vSphere, administrators can integrate the management features of the System x, BladeCenter, and Flex System with VMware vCenter. It also expands the virtualization management capabilities of VMware vCenter with System x hardware management functionality, which provides affordable, basic management of physical and virtual environments to reduce the time and effort that is required for routine system administration.

UIMs also provide the discovery, configuration, monitoring, event management, and power monitoring that is needed to reduce cost and complexity through server consolidation and simplified management.

For more information about UIMs for VMware vSphere, see this website:

<http://www-947.ibm.com/support/entry/portal/docdisplay?lnocid=migr-vmware>

In our scenario, it is assumed that you installed UIMs on your VMware vSphere environment.

Consideration: Consider the use of UIM for VMware vSphere for the unified hardware and software management of the combined Flex System and BladeCenter environment that is running VMware vSphere.

4.1.1 Enabling UIMs for a newly added ESXi host

Complete the following steps to enable UIM on your newly added ESXi host:

1. Log in to VMware vSphere web client.
2. Enter **Hosts and Clusters** view.
3. Click the cluster to which your ESXi host belongs.
4. Select the **Manage** tab and click **Upward Integration**.
5. In Overview tab of Upward Integration, select **Cluster Overview**.
6. From the list of ESXi hosts, select the check box next to the host that you want to enable.
7. From the drop-down list of ESXi hosts, select **Request Host Access**.
8. Enter the credentials when prompted.

These steps are shown in Figure 4-1.

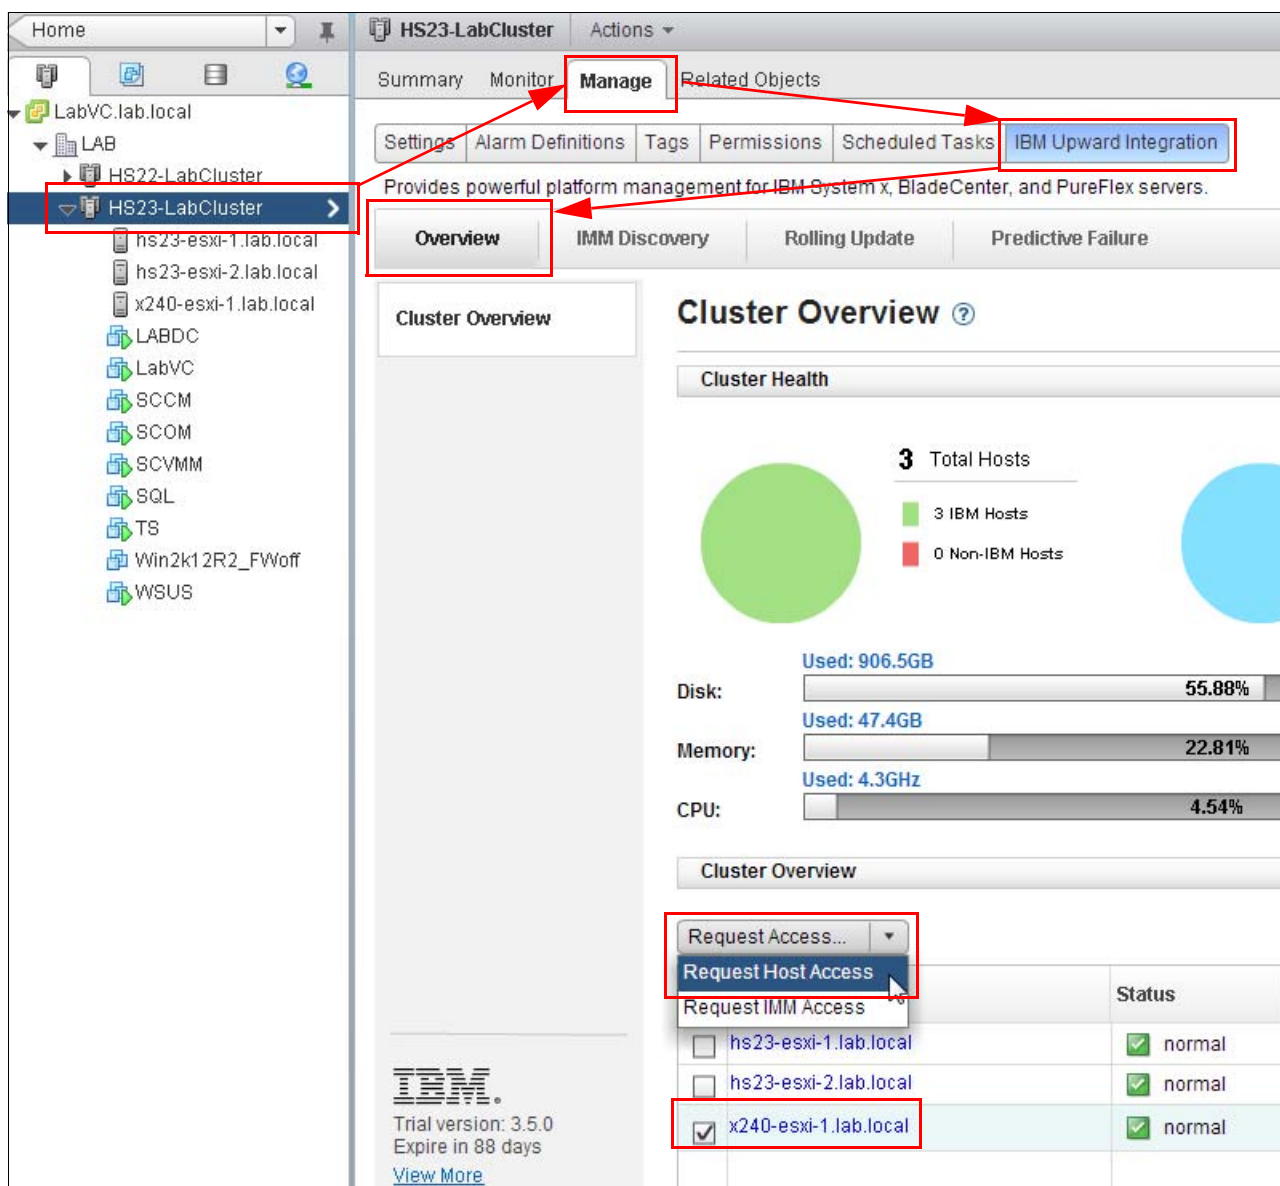


Figure 4-1 Request Host Access

4.1.2 Collecting system inventory with UIM

Complete the following steps to see the available information that is related to your VMware ESXi host:

1. Log in to VMware vSphere web client.
2. Enter the **Hosts and Clusters** view.
3. Click the ESXi host for which you want to gather the information.
4. Select the **Manage** tab and click **Upward Integration**.
5. In the System tab of Upward Integration, click **Collect** to collect hardware and software details. (The collection process can take several minutes.)

These steps are shown in Figure 4-2.

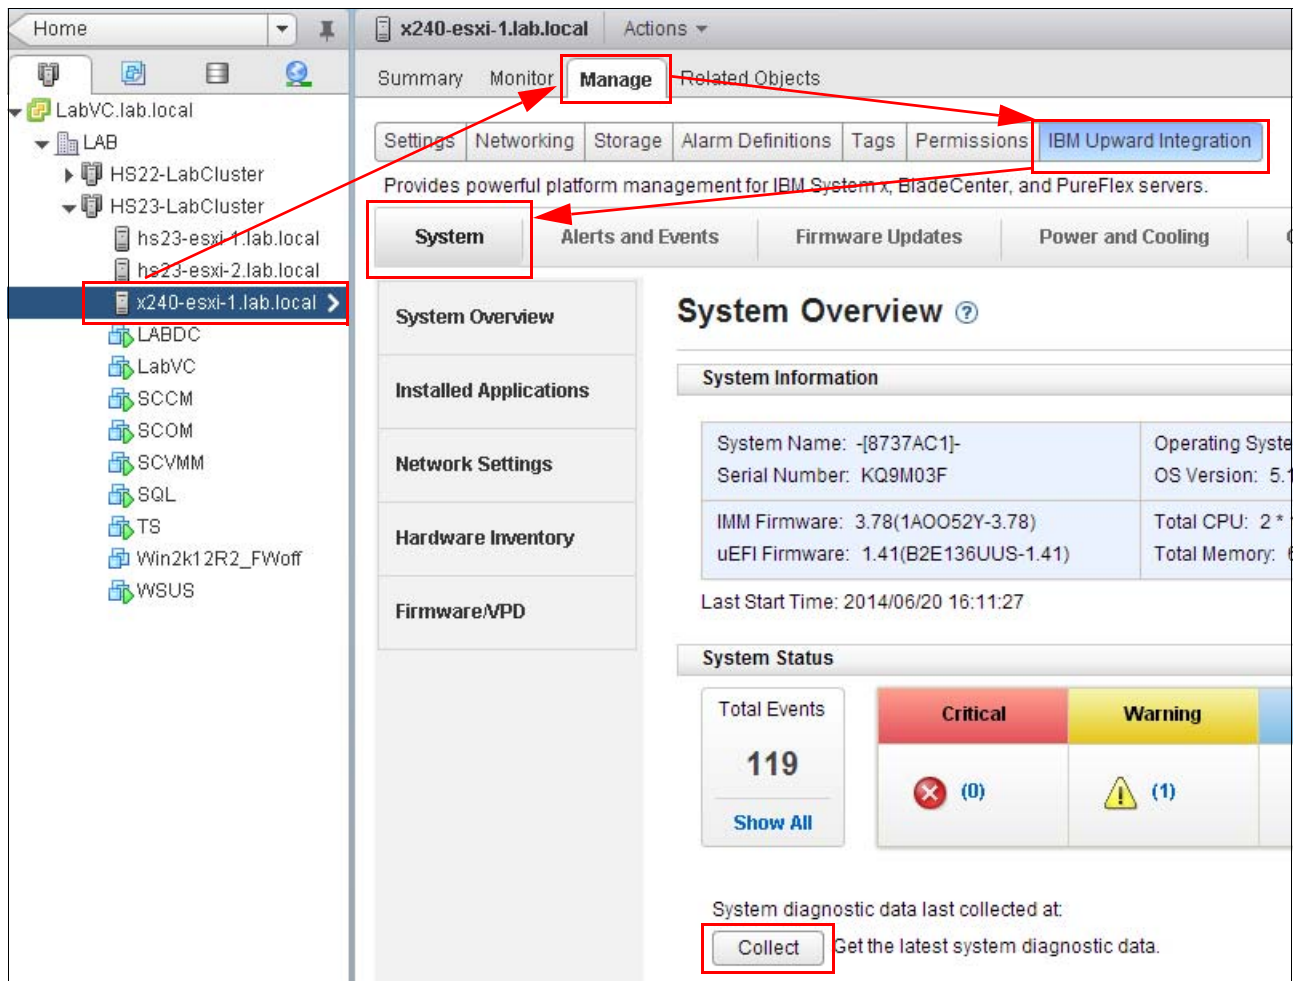


Figure 4-2 Collecting details about hardware and software of ESXi host

After the collection process is finished, you can access hardware and software details by clicking menu items on the left side of the Upward Integration page.

The following views are available:

- The Installed Applications view is shown in Figure 4-3.

x240-esxi-1.lab.local
Actions

Summary
Monitor
Manage
Related Objects

Settings
Networking
Storage
Alarm Definitions
Tags
Permissions
IBM Upward Integration

Provides powerful platform management for IBM System x, BladeCenter, and PureFlex servers.

System
Alerts and Events
Firmware Updates
Power and Cooling
Configuration
Help

System Overview
Installed Applications
Network Settings
Hardware Inventory
Firmware/VPD

Installed Applications ?

Name	Version	Caption	Install Date
brcm	500.2.0.3-000000	brcm	2014061314332
misc-cnic-register	1.72.1.v50.2-10EM.500.0.0.472	misc-cnic-register	2014061314332
net-bnx2	2.2.3e.v50.1-10EM.500.0.0.472	net-bnx2	2014061314332
net-bnx2x	1.74.22.v50.1-10EM.500.0.0.472	net-bnx2x	2014061314332
net-cnic	1.74.04.v50.3-10EM.500.0.0.472	net-cnic	2014061314332
net-tg3	3.135b.v50.1-10EM.500.0.0.472	net-tg3	2014061314332
scsi-bnx2fc	1.74.02.v50.2-10EM.500.0.0.472	scsi-bnx2fc	2014061314332
scsi-bnx2i	2.74.07.v50.1-10EM.500.0.0.472	scsi-bnx2i	2014061314332
brcdprovider	3.2.0.0-0	brcdprovider	2014061314332
net-bna	3.2.0.0-10EM.500.0.0.472560	net-bna	2014061314332
scsi-bfa	3.2.0.0-10EM.500.0.0.472560	scsi-bfa	2014061314332
emulex-cim-provider	3.8.21.1-01	emulex-cim-provider	2014062016054
ima-be2iscsi	4.6.142.2-10EM.500.0.0.47262	ima-be2iscsi	2014061314332
net-be2net	4.6.142.10-10EM.510.0.0.8022	net-be2net	2014061314332
scsi-be2iscsi	4.6.142.2-10EM.500.0.0.47262	scsi-be2iscsi	2014061314332
scsi-lpfc820	8.2.4.151.65-10EM.500.0.0.472	scsi-lpfc820	2014061314332
concreteah	500.24CE22CUE	concreteah	2014062016054

Trial version: 3.5.0
Expire in 88 days
[View More](#)
©2013, All Rights Reserved

Figure 4-3 Installed Applications view

- The Network Settings view is shown in Figure 4-4.

x240-esxi-1.lab.local

Actions ▾

Summary

Monitor

Manage

Related Objects

Settings

Networking

Storage

Alarm Definitions

Tags

Permissions

IBM Upward Integration

Provides powerful platform management for IBM System x, BladeCenter, and PureFlex servers.

System

Alerts and Events

Firmware Updates

Power and Cooling

Configuration

Help

System Overview

Installed Applications

Network Settings

Hardware Inventory

Firmware/VPD

Network Settings ?

Physical Network Ports

Name	vmnic0	vmnic1
DeviceID	vmnic0	vmnic1
OtherIdentifyingInfo	vmklinux,0x12,0x0,0x0,0x19a2,0x 710	vmklinux,0x12,0x0,0x1,0
LinkTechnology	Ethernet	Ethernet
PermanentAddress	3440B5BE7D00	3440B5BE7D04
NetworkAddresses	3440B5BE7D00	3440B5BE7D04
ActiveMaximumTransmissionUnit	1.5 Kilobytes	1.5 Kilobytes
EnabledState	Enabled	Enabled
FullDuplex	true	true

IPv4 Endpoint

Name	vmk0	vmk1	vmk2
TransitioningToState	Not Applicable	Not Applicable	Not Applicable
SubnetMask	255.255.254.0	255.255.255.0	255.255.255.0
RequestedState	No Change	No Change	No Change
ProtocolIFType	IPv4	IPv4	IPv4
IPv4Address	9.42.171.26	169.254.95.120	10.30.30.26

IBM

Trial version: 3.5.0

Expire in 88 days

[View More](#)

©2013, All Rights Reserved

Figure 4-4 Network Settings view

- ▶ The Hardware Inventory view is shown in Figure 4-5.

x240-esxi-1.lab.localActions

SummaryMonitorManageRelated Objects

SettingsNetworkingStorageAlarm DefinitionsTagsPermissionsIBM Upward Integration

Provides powerful platform management for IBM System x, BladeCenter, and PureFlex servers.

SystemAlerts and EventsFirmware UpdatesPower and CoolingConfiguration? Help

System Overview

Installed Applications

Network Settings

Hardware Inventory

Firmware/VPD

Hardware Inventory ?

Memory

Manufacturer	Samsung	Samsung	Samsung	Samsu
Capacity	8589934592	8589934592	8589934592	858993
BankLabel	Bank 1	Bank 4	Bank 9	Bank 12
SerialNumber	33F8CC8A	33F8CCD5	33F8CC89	33F8CE
Model	DDR3	DDR3	DDR3	DDR3
Speed	1600	1600	1600	1600
PartNumber	M393B1K70DH0-CK0	M393B1K70DH0-CK0	M393B1K70DH0-CK0	M393B1
Description	DIMM 1	DIMM 4	DIMM 9	DIMM 12

Processor

Name	Processor 1	Processor 2
Family	179	179
CPUStatus	1	1
NumberOfEnabledCores	8	8
CurrentClockSpeed	2000	2000
OtherFamilyDescription	Intel(R) Xeon(R) CPU E5-2650 0 @ 2.00GHz	Intel(R) Xeon(R) CPU E

Trial version: 3.5.0
Expire in 88 days
[View More](#)

©2013, All Rights Reserved

Figure 4-5 Hardware Inventory view

- The Firmware/VPD view is shown in Figure 4-6.

Provides powerful platform management for IBM System x, BladeCenter, and PureFlex servers.

System Alerts and Events Firmware Updates Power and Cooling Configuration ? Help

Firmware/VPD ?

Software Identity

Description	ElementName	IdentityInfoType	IdentityInfoValue
IMM2 Firmware	IMM2	SoftwareID;SoftwareStatus	1A00;2,6
IMM2 Backup Firmware	IMM2-Backup	SoftwareID;SoftwareStatus	1A00;6
UEFI Firmware/BIOS	UEFI	SoftwareID;SoftwareStatus	B2E1;2,6
UEFI Backup Firmware/B	UEFI-Backup	SoftwareID;SoftwareStatus	B2E1;6
DSA Diagnostic Software	DSA	SoftwareID;SoftwareStatus	DSYT;2,3,6

Figure 4-6 Firmware/VPD view

4.1.3 Monitoring hardware status

By using UIMs, the vSphere administrator can get a detailed view of the hardware's health. You can view your Hardware event logs directly from your vSphere web client, and there is no need to log in to IMM.

The System Health view is in the Alerts and Events tab of the UIM, as shown in Figure 4-7.

Provides powerful platform management for IBM System x, BladeCenter, and PureFlex servers.

System Alerts and Events Firmware Updates Power and Cooling Configuration ? Help

System Health ?

Critical(0)
 Warning(1)
 Information(499)

Filter by: All

Message ID	Severity	Time Stamp	Message Detail
PLAT0188	Information	2014-06-13 09:10:55	The System IBM Flex System x240 with 10Gb...
IMM0001	Information	2014-06-13 09:13:24	Management Controller SN# Network Initial...
IMM0025	Information	2014-06-13 09:13:30	LAN: Ethernet[IMM:ep1] interface is now activ...

Figure 4-7 System Health view

Information and statistics about ESXi host power usage can be found in the Power and Cooling tab of the UIM, as shown in Figure 4-8.

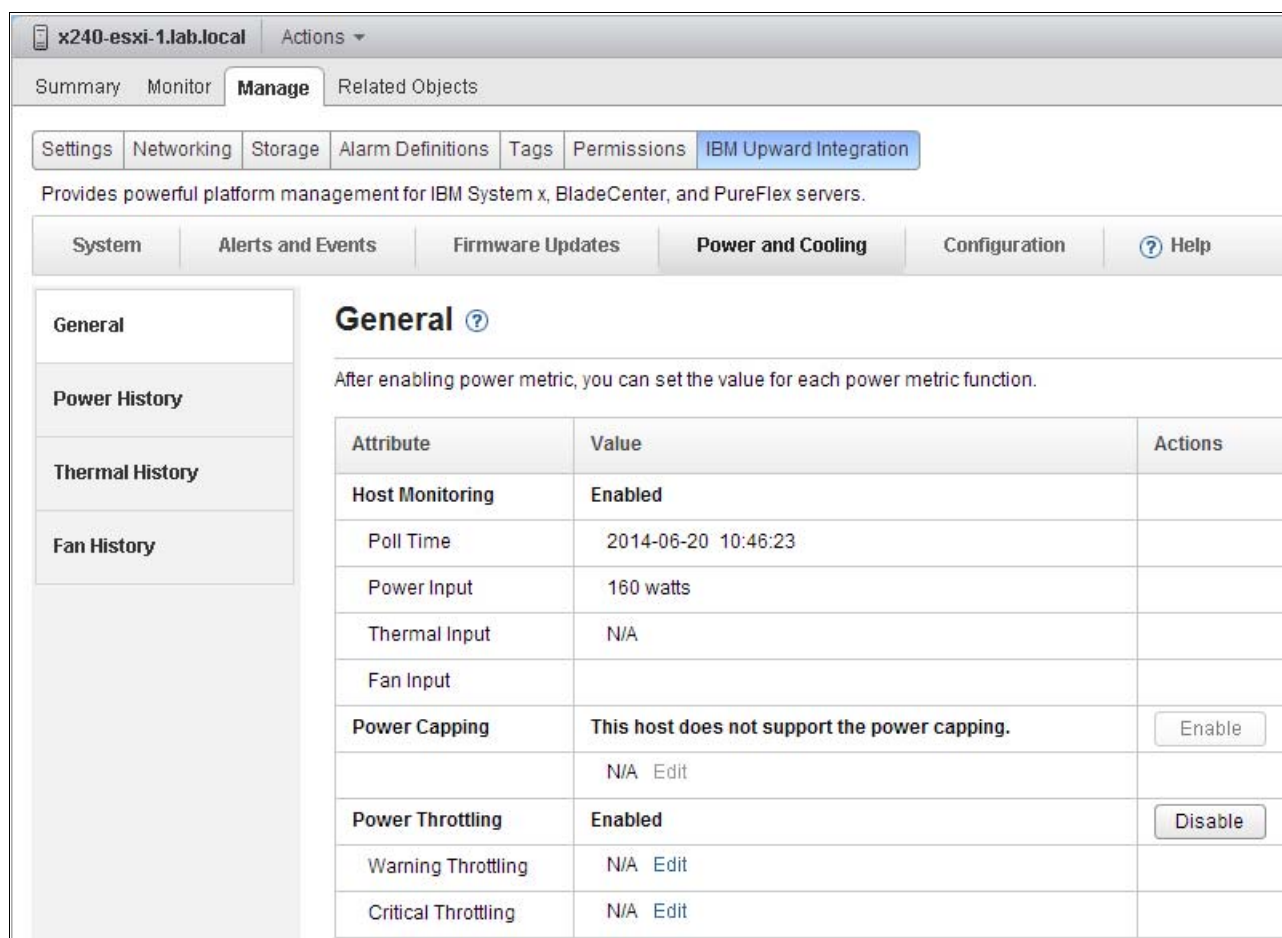


Figure 4-8 General View of Power and Cooling tab

4.2 Using PFA alert to move VMs to another ESXi host

In this section, we describe how to use predictive failure management on the vSphere web client to protect your running workload. By using the Policy and Rules page, you can set management policies for a server that is based on a hardware Predictive Failure Alert (PFA).

Based on a defined policy, the Upward Integration for VMware vSphere evacuates VMs from the server to other hosts in the cluster in response to a PFA. You can view PFAs from the server and the triggered policy history on the Predictive Failures page.

Before you begin

Before predictive failure management is used, verify that the following prerequisites are met:

- ▶ The predictive failure management policy can be set until you discover the IMMs and request the IMMs access.
- ▶ Predictive failure management relies on the hardware PFA capability. The IMM of the server must send out Predictive Failure Alerts when a failure is detected.

- ▶ Proper configuration of the network management policy on the vCenter server is required to enable TCP on the https port that you selected when IVP was installed (the default port is 9500). Upward Integration for VMware vSphere listens on this port for incoming indications.
- ▶ The host must be put in a properly configured cluster. There must be a host available with vMotion enabled in this cluster. Upward Integration evacuates VMs to other hosts in the cluster, and then puts the host in maintenance mode.

Setting a new policy

You can set an RAS policy on each supported server in the cluster. A policy defines the hardware event categories that you want to monitor and the corresponding action when the event occurs.

To implement this task, click your cluster object in the Hosts and Clusters view, select the **Manage** tab, and click **Upward Integration**. Then, select the **Predictive Failure** tab and you the Policy and rules page opens.

Complete the following steps to set up a policy:

1. Select one or more nodes.
2. Click **Set policy**. The Manage RAS Policy page is displayed, as shown in Figure 4-9.

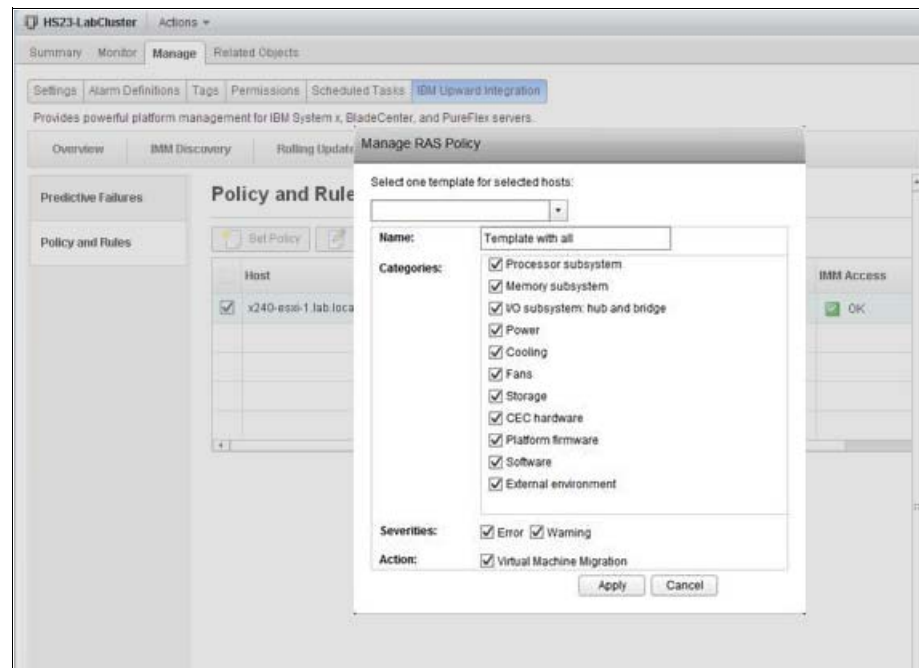


Figure 4-9 Manage RAS policy

3. Select the following event categories, severities, and action:
 - Event categories

The Table 4-1 lists the Predictive Failure Alert Event categories that are used on the Manage RAS Policy page.

Table 4-1 Predictive Failure Alert Event categories

PFA Event	Description
Processor subsystem	Processor subsystem includes the CPU and its internal circuits, such as cache, the bus controller, and external interface.
Memory subsystem	Memory subsystem includes the memory controller, memory buffer, memory bus interface, memory card, and DIMM.
I/O subsystem	I/O subsystem includes: IO Hub, IO bridge, IO bus, IO processor, IO adapters for various IO protocols, such as PCI and InfiniBand.
Power	Power includes the power supply and power control hardware.
Cooling	All thermal-related events.
Fans	Includes the fan and blower.
Storage	Includes the storage enclosure, storage controller, raid controller, and media (disk, flash).
Platform firmware	Platform firmware includes IMM and uEFI.
Software	Operating system software and application software.
External environment	All events of an external-related environment including: AC power source, Room ambient temperature, and user error.

– Event severity

Table 4-2 lists the PFA Event severity levels.

Table 4-2 Predictive Failure Alert severity levels.

Severity	Description
Warning	An indication of a failure, which can have no effect on performance. Service action is necessary.
Error	A failure that causes a loss of performance and can cause machines to be inoperable. Immediate service action is necessary.

– Action

The Virtual Machine Migration action evacuates all of the VMs from the server and puts the server in maintenance mode.

After setting the event categories and corresponding action, click **Apply** to apply the policy to the host.

Note: The newly created policy is saved as a template automatically so that for any other hosts, you can choose a template from the top template drop-down list to apply the same policy.

Editing a policy

You can modify a policy that is defined on a host by using the Edit policy function. Complete the following steps:

1. Select a host.
2. Click **Edit policy**.

Note: When the policy is modified and the policy also used by other hosts, a warning message is displayed with which you can apply the changes to other hosts or save the changed policy with a different policy name.

Disabling a policy

You can remove a policy from one or more hosts by using the Disable policy function. Complete the following steps:

1. Select one or more hosts.
2. Click **Disable policy**.
3. Click **Disable** to confirm the deletion of the policy from the hosts.

Viewing predictive failure alert events and the Action History table

Upward Integration for VMware vSphere with vSphere Client monitors Predictive Failure Alerts (PFAs) from the IMM. All predictive failure events are listed in the Event Log table. When the conditions of a rule are met, the defined action of the rule is started on the managed endpoint. All of the triggered rules and action results are listed in the Action History table, as shown in Figure 4-10.

Getting Started Summary Monitor **Manage** Related Objects

Settings Alarm Definitions Tags Permissions Scheduled Tasks **IBM Upward Integration**

Provides powerful platform management for IBM System x, BladeCenter, and PureFlex servers.

Overview IMM Discovery Rolling Update **Predictive Failure**

Predictive Failures ?

View Predictive Failure event log and action history.

Event Log

Host	Message ID	Severity	Time Stamp
2002:97b:c2bb:830:20a:f7ff:fe26:9a32	PLAT0138	Error	04:46:43 05/0

Action History

Host	Message ID	Status	Start Time
2002:97b:c2bb:830:20a:f7ff:fe26:9a32	PLAT0138	Success Detail...	12:56:27 05/0

IBM
Version information: 3.5.0
[View More](#)

Figure 4-10 Viewing Predictive Failures

4.2.1 Rolling firmware upgrades

You can upgrade your firmware by using Update manager. There are two possibilities: you can manually upgrade each ESXi host individually, or you can schedule a rolling update so that update is pushed to the servers at a scheduled time. UIM manages evacuating the ESXi host before the firmware is updated.

Complete the following steps to create a rolling update:

1. In your vSphere web client, browse to the Hosts and Clusters view, click your cluster, select the **Manage** tab, click **Upward Integration**, and then select the **Rolling Update** tab, as shown in Figure 4-11 on page 89. Click **Create**.

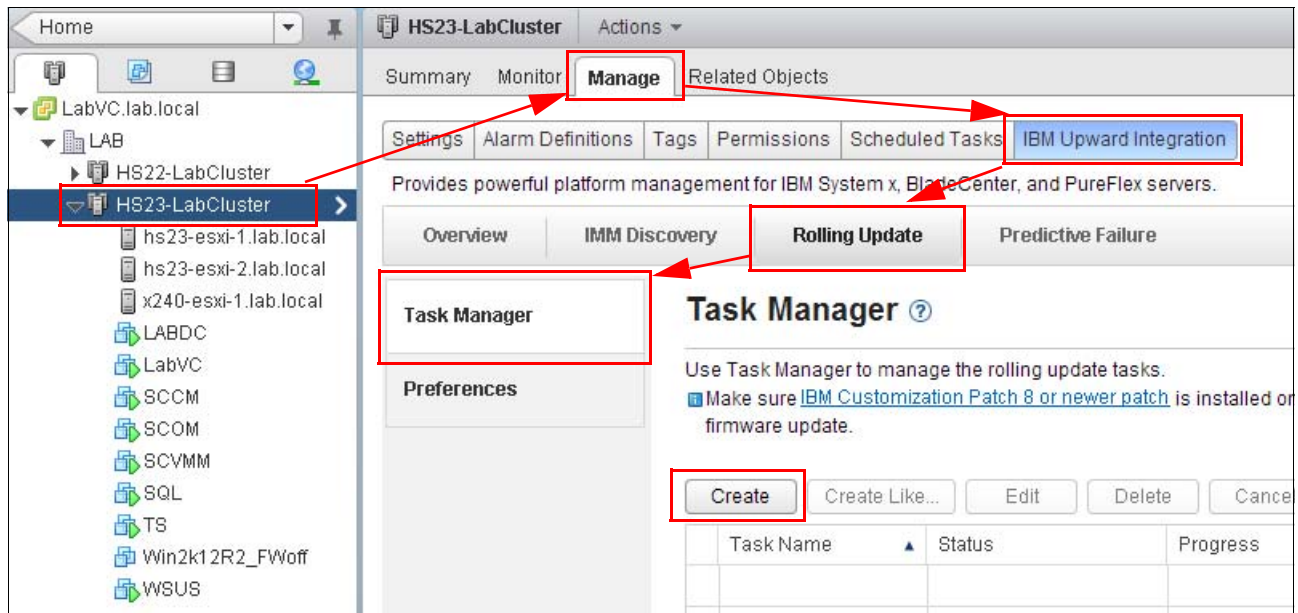


Figure 4-11 Rolling Update in UIM

2. A wizard opens. Complete the following steps:
 - a. Enter a Task Name for the rolling update job. Select the Task Type and click **Next**, as shown in Figure 4-12.

Rolling System Update		
1. Name and Type	2. Select hosts and firmware	3. Update options and schedules
<div>Task Name: <input type="text" value="Rolling_Update1"/></div> <div>Task Type: <input checked="" type="radio"/> Update and Reboot <input type="radio"/> Update Only <input type="radio"/> Reboot Only</div>		

Figure 4-12 Select Name and Type of Rolling update.

- b. Select the ESXi host the updates that you want to apply, as shown in Figure 4-13. Click **Next**.

Rolling System Update

1. Name and Type **2. Select hosts and firmware** 3. Update options and schedules

▼ ☐ -[7875AC1]-

☐ hs23-esxi-2.lab.local

☐ hs23-esxi-1.lab.local

▼ ☐ -[8737AC1]-

☒ x240-esxi-1.lab.local (3 selected items)

Available firmware for x240-esxi-1.lab.local

	Firmware Name	New Versions	1 ▲ Install
<input checked="" type="checkbox"/>	▼ UXSP		
<input checked="" type="checkbox"/>	IBM Dynamic System Analysis (DE	DSYTE0R-9.60	DSYTE
<input checked="" type="checkbox"/>	IBM Flex System x240 UEFI Flash	B2E142A-1.50	B2E14
<input checked="" type="checkbox"/>	Integrated Management Module 2	1A0058R-4.20	1A005
<input type="checkbox"/>	▼ Individual		
<input type="checkbox"/>	IBM Dynamic System Analysis (DE	DSYTC4P-...	DSYTE
<input type="checkbox"/>	IBM Flex System x240 UEFI Flash	B2E136U-1...	B2E14
<input type="checkbox"/>	Integrated Management Module 2	1A0056G-...	1A005

Figure 4-13 Select host and firmware

- c. You can update several hosts at the same time if your cluster resources can manage the workload. To do so, select the **Update Parallelization** option and enter the number of hosts that you want to update at the same time. You can force the downgrade of the firmware by selecting the **Force Downgrade** option. If you want to schedule this update instead of running it immediately, select the **Schedule** option and enter the date and time that you want to run the update, as shown in Figure 4-14. Click **Next**.

Rolling System Update

1. Name and Type 2. Select hosts and firmware **3. Update options and schedules**

☐ Update Parallelization

Scale: Make sure the value is set according to the current available system resources of the cluster.

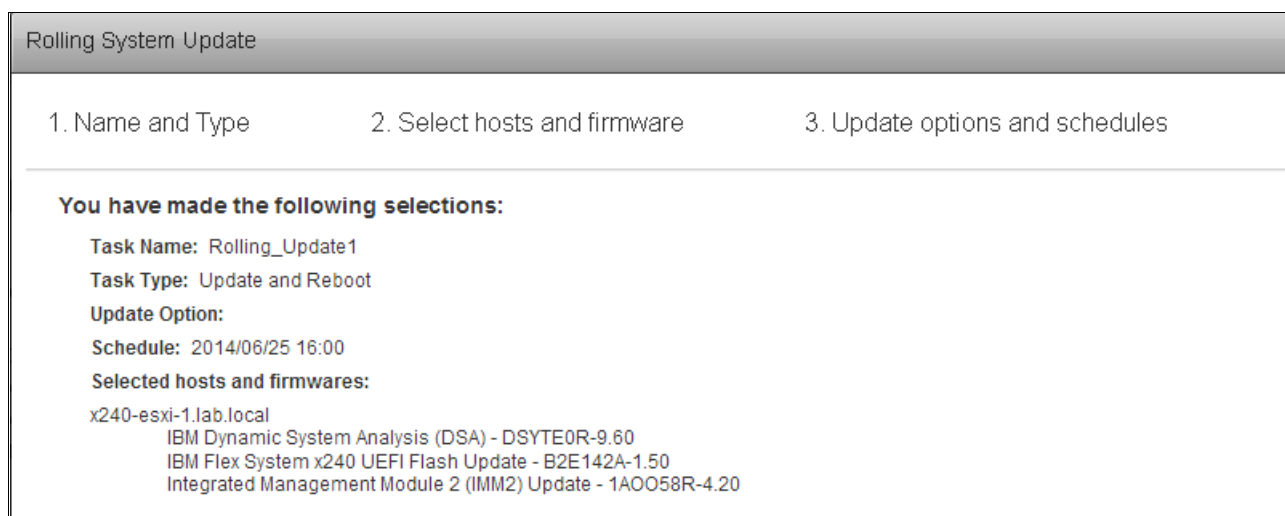
☐ Force downgrade

☒ Schedule

☐ Now
 ☒ Schedule

Figure 4-14 Update options and schedules

- d. On the last page, review the summary of the created job and click **Finish**, as shown in Figure 4-15.



Rolling System Update

1. Name and Type 2. Select hosts and firmware 3. Update options and schedules

You have made the following selections:

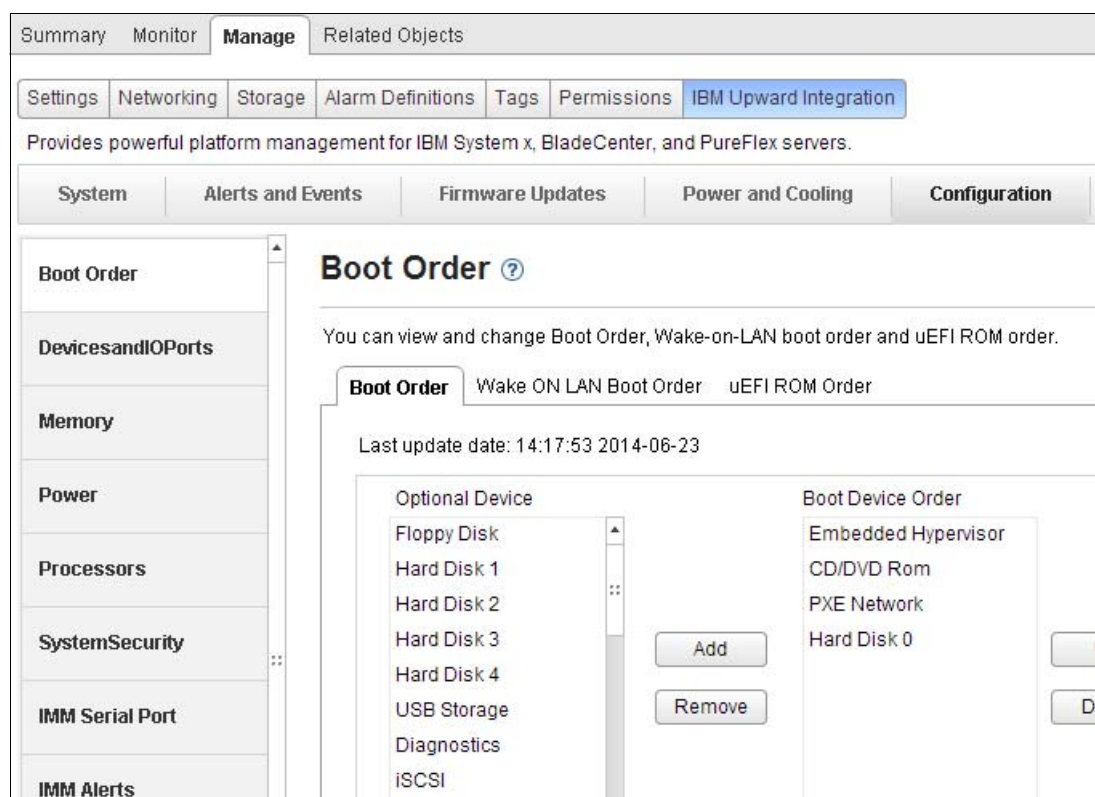
Task Name: Rolling_Update1
 Task Type: Update and Reboot
 Update Option:
 Schedule: 2014/06/25 16:00
 Selected hosts and firmwares:
 x240-esxi-1.lab.local
 IBM Dynamic System Analysis (DSA) - DSYTE0R-9.60
 IBM Flex System x240 UEFI Flash Update - B2E142A-1.50
 Integrated Management Module 2 (IMM2) Update - 1AO058R-4.20

Figure 4-15 Rolling Update job creation

4.2.2 Changing IMM and UEFI configuration

By using UIM, you can some of the IMM and UEFI parameters. To do so, browse to the Hosts and Clusters view, click your ESXi host, select the **Manage** tab, click **Upward Integration**, and then select **Configuration** tab, as shown in the following examples:

- Edit host boot order is shown in Figure 4-16.



Summary Monitor **Manage** Related Objects

Settings Networking Storage Alarm Definitions Tags Permissions **IBM Upward Integration**

Provides powerful platform management for IBM System x, BladeCenter, and PureFlex servers.

System Alerts and Events Firmware Updates Power and Cooling **Configuration**

Boot Order ?

You can view and change Boot Order, Wake-on-LAN boot order and uEFI ROM order.

Boot Order Wake ON LAN Boot Order uEFI ROM Order

Last update date: 14:17:53 2014-06-23

Optional Device	Boot Device Order
Floppy Disk	Embedded Hypervisor
Hard Disk 1	CD/DVD Rom
Hard Disk 2	PXE Network
Hard Disk 3	Hard Disk 0
Hard Disk 4	
USB Storage	
Diagnostics	
iSCSI	

Add Remove

Figure 4-16 Edit Boot Order window

- Manage Devices and IO ports is shown in Figure 4-17.

SummaryMonitorManageRelated Objects

SettingsNetworkingStorageAlarm DefinitionsTagsPermissionsIBM Upward Integration

Provides powerful platform management for IBM System x, BladeCenter, and PureFlex servers.

SystemAlerts and EventsFirmware UpdatesPower and CoolingConfiguration

Boot Order

DevicesandIOPorts

Memory

Power

Processors

SystemSecurity

IMM Serial Port

IMM Alerts

IMM Port Assignments

IMM SNMP

DevicesandIOPorts ?

Please save the changes when you finish the setting to make them effective.

Save

Refresh

Last update date: 14:20:17 2014-06-23

ActiveVideo	Add-in Device
COMPort1	Enable
COMPort2	Enable
Com1ActiveAfterBoot	Disable
Com1BaudRate	115200
Com1DataBits	8
Com1FlowControl	Disable

Figure 4-17 Devices and IO Ports window

- Manage your Memory modules configuration is shown in Figure 4-18.

SummaryMonitorManageRelated Objects

SettingsNetworkingStorageAlarm DefinitionsTagsPermissionsIBM Upward Integration

Provides powerful platform management for IBM System x, BladeCenter, and PureFlex servers.

SystemAlerts and EventsFirmware UpdatesPower and CoolingConfiguration

Boot Order

DevicesandIOPorts

Memory

Power

Processors

SystemSecurity

IMM Serial Port

IMM Alerts

IMM Port Assignments

IMM SNMP

IMM Security

Memory ?

Please save the changes when you finish the setting to make them effective.

Save

Refresh

Last update date: 14:20:53 2014-06-23

CKEThrottling	
CKSelfRefresh	
DIMM10onProcessor1	Enable
DIMM11onProcessor1	Enable
DIMM12onProcessor1	Enable
DIMM13onProcessor2	Enable
DIMM14onProcessor2	Enable
DIMM15onProcessor2	Enable

Figure 4-18 Memory settings

- Manage Power management settings is shown in Figure 4-19.

Summary	Monitor	Manage	Related Objects
---------	---------	---------------	-----------------

Settings	Networking	Storage	Alarm Definitions	Tags	Permissions	IBM Upward Integration
----------	------------	---------	-------------------	------	-------------	------------------------

Provides powerful platform management for IBM System x, BladeCenter, and PureFlex servers.

System	Alerts and Events	Firmware Updates	Power and Cooling	Configuration
--------	-------------------	------------------	-------------------	----------------------

<div>Boot Order</div> <div>DevicesandIOPorts</div> <div>Memory</div> <div>Power</div> <div>Processors</div> <div>SystemSecurity</div> <div>IMM Serial Port</div> <div>IMM Alerts</div> <div>IMM Port Assignments</div>	<h2>Power ?</h2> <p>Please save the changes when you finish the setting to make them effective.</p> <p> <input type="button" value="Save"/> <input type="button" value="Refresh"/> Last update date: 14:22:28 2014-06-23 </p> <table border="1"> <tr> <td>ActiveEnergyManager</td> <td>Capping Enabled ▼</td> </tr> <tr> <td>PlatformControlledType</td> <td>Efficiency - Favor Perfor... ▼</td> </tr> <tr> <td>PowerPerformanceBias</td> <td>Platform Controlled ▼</td> </tr> <tr> <td>S3Enable</td> <td>▼</td> </tr> <tr> <td>WorkloadConfiguration</td> <td>Balanced ▼</td> </tr> </table>	ActiveEnergyManager	Capping Enabled ▼	PlatformControlledType	Efficiency - Favor Perfor... ▼	PowerPerformanceBias	Platform Controlled ▼	S3Enable	▼	WorkloadConfiguration	Balanced ▼
ActiveEnergyManager	Capping Enabled ▼										
PlatformControlledType	Efficiency - Favor Perfor... ▼										
PowerPerformanceBias	Platform Controlled ▼										
S3Enable	▼										
WorkloadConfiguration	Balanced ▼										

Figure 4-19 Power management

- Manage SNMP configuration of the IMM is shown in Figure 4-20.

Summary Monitor **Manage** Related Objects

Settings Networking Storage Alarm Definitions Tags Permissions **IBM Upward Integration**

Provides powerful platform management for IBM System x, BladeCenter, and PureFlex servers.

System Alerts and Events Firmware Updates Power and Cooling **Configuration**

IMM SNMP ?

Please save the changes when you finish the setting to make them effective.

Save Refresh Last update date: 14:32:43 2014-06-23

SNMP Agent Port	161
SNMP Trap Port	162
SNMP Traps	Disabled
SNMPv3 Access Type	Set
SNMPv3 Authentication Protocol	HMAC-SHA
SNMPv3 Privacy Protocol	AES
SNMPv3 Trap Hostname	9.42.171.38

Figure 4-20 Configure SNMP on IMM

4.3 Managing a Windows Server environment with UIM

For managing Microsoft Windows server environment that is hosted on BladeCenter and System x, you can use the System x UIM for Microsoft System Center.

Important: Consider the use of UIM for Microsoft System Center for the unified hardware and software management of the combined Flex System and BladeCenter environment that is based on Windows Server infrastructure.

Lenovo expands Microsoft System Center server management capabilities by integrating System x hardware management functionality, which provides affordable, basic management of physical and virtual environments to reduce the time and effort that is required for routine system administration. It also provides the discovery, configuration, monitoring, event management, and power monitoring that is needed to reduce cost and complexity through server consolidation and simplified management.

For more information about UIM for Microsoft System Center, see this website:

<http://www-947.ibm.com/support/entry/portal/docdisplay?lnodocid=SYST-MANAGE>

4.3.1 Enabling Hardware Monitoring on the newly deployed Flex System

In this section, we describe how to discover an Flex System in Microsoft System Center Operations Manager 2012 (SCOM).

Setting up Flex System Chassis Management Module for discovery

Before you can monitor the hardware status of the Flex chassis components in SCOM, you must configure SNMP in CMM. Complete the following steps:

1. Log in to the CMM console as Administrator.
2. To change the SNMP settings, click **Mgt Module Management** → **Network** → **SNMP**. Select **Enabled for SNMPv3 Agent**. (You also can enter the Contact and Location information). Click **Apply**, as shown in Figure 4-21.

The screenshot shows the 'Network Protocol Properties' window with the 'SNMP' tab selected. Under 'Simple Network Management Protocol (SNMP)', the 'Enable SNMPv3 Agent' checkbox is checked. Below this, the 'Contact' tab is active, showing 'Contact and Location' information. The 'Contact person' field contains 'No Contact Configured' and the 'Chassis location' field contains 'No Location Configured'. There is an 'Apply' button at the top left of the window.

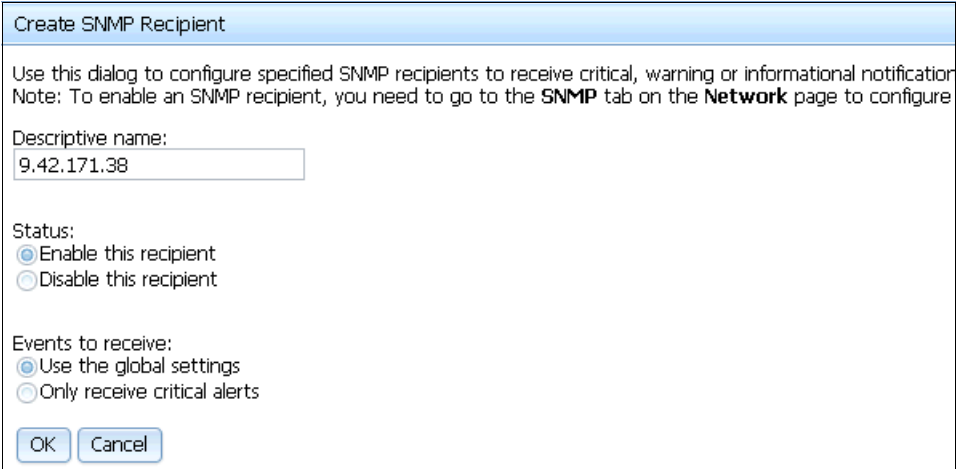
Figure 4-21 Enable SNMPv3 Agent

Note: There are two SNMP agent versions that can be selected for the SCOM to manage the Flex System chassis: SNMPv1 and SNMPv3. In our example, we show SNMPv3, which provides more security than SNMPv1.

To receive events from the management modules, a network connection must exist between the management module and the Microsoft System Center Operations Manager. You also must configure the management module to send events.

3. To define the SNMP recipient, click **Event** → **Event Recipients**.
4. Click **Create** → **Create SNMP Recipient**.

5. In the Create SNMP Recipient dialog box, enter the IP address of the SCOM server in Descriptive name field.
6. Select **Enable this recipient**.
7. Select **Use the global settings** or **Only receive critical alerts**, as shown in Figure 4-22. Click **OK** to return to the Event Recipients page.



Create SNMP Recipient

Use this dialog to configure specified SNMP recipients to receive critical, warning or informational notification
 Note: To enable an SNMP recipient, you need to go to the **SNMP** tab on the **Network** page to configure

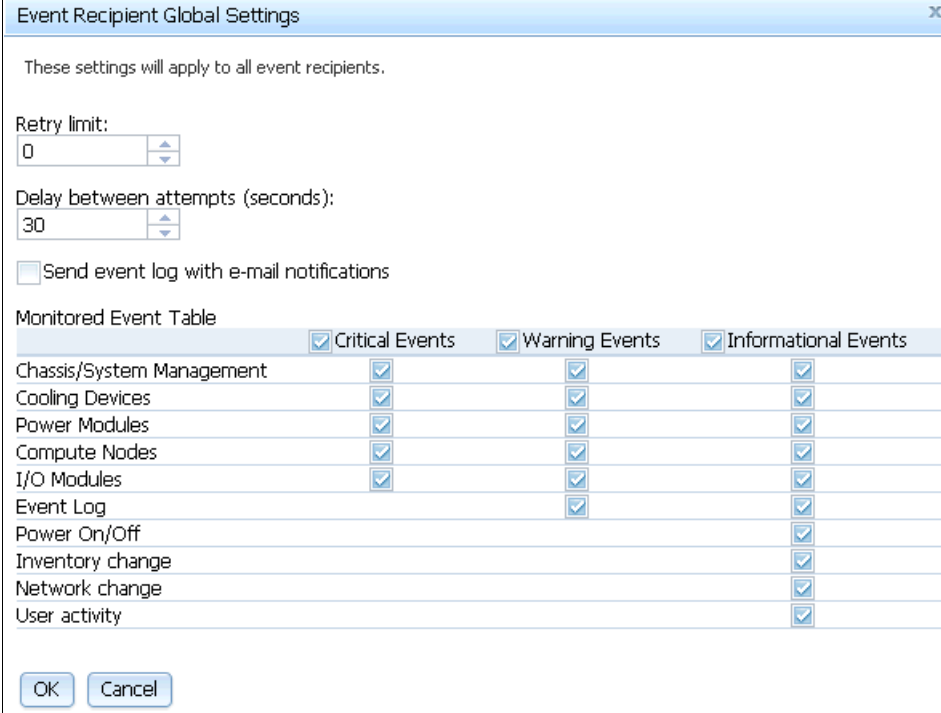
Descriptive name:

Status:
☒ Enable this recipient
☐ Disable this recipient

Events to receive:
☒ Use the global settings
☐ Only receive critical alerts

Figure 4-22 Create SNMP Recipient

If you selected **Use the global settings**, the Event Recipient Global Settings dialog box opens, as shown in Figure 4-23. Click **OK**.



Event Recipient Global Settings

These settings will apply to all event recipients.

Retry limit:

Delay between attempts (seconds):

☐ Send event log with e-mail notifications

Monitored Event Table

	<input checked="" type="checkbox"/> Critical Events	<input checked="" type="checkbox"/> Warning Events	<input checked="" type="checkbox"/> Informational Events
Chassis/System Management	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Cooling Devices	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Power Modules	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Compute Nodes	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
I/O Modules	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Event Log		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Power On/Off			<input checked="" type="checkbox"/>
Inventory change			<input checked="" type="checkbox"/>
Network change			<input checked="" type="checkbox"/>
User activity			<input checked="" type="checkbox"/>

Figure 4-23 Event: Recipient Global Settings window

8. To define the SNMPv3 user, click **Mgt Module Management** → **User Accounts**.
9. Click the existing user or **Create** to create a user.

10. In the General tab, enter the user name and password and click the **SNMPv3** tab.
11. Specify the security settings that are based on your company security policy. Set the Access type to Set and enter the IP address of the SCOM server for traps, as shown in Figure 4-24.

User Properties

General Permission Group **SNMPv3** SSH Client Public Key Node Account Mgmt

Context name:
context2

Authentication Protocol:
Hash-based Message Authentication Code (HMAC) - Secure Hash Algorithm (SHA)

☒ Use a privacy protocol
Encryption Method:
Advanced Encryption Standard (AES)

Privacy password:
.....

Confirm privacy password:
.....

Access type:
Set

IP address or host name for traps:
9.42.171.38

OK Cancel

Figure 4-24 SNMPv3 User Properties window

Setting up System Center Operations Manager 2012 for Discovery

There is only one discovery rule for network devices per SCOM management server allowed. Because we are integrating Flex System chassis into the existing environment, we describe modifying the existing rule in this section.

Note: If you are using dynamic discovery, your Flex System chassis might be discovered automatically if the CMM is in the previously defined discovery range with same the SNMP credentials that were assigned to the discovery rule.

Hardware monitoring by using SCOM requires the Lenovo Hardware Management Pack for Microsoft System Center Operations Manager to be imported in SCOM. Complete the following steps:

1. Log in to the Microsoft System Center Operations Manager operations console as Administrator.

Note: This feature supports a CMM IP address only. Do not use an IMM IP address.

2. Click **Administration** → **Network Management** → **Discovery Rules** to see the list of discovery rules.

3. Double-click a rule that you want to modify. In our example, we use the rule for BladeCenter AMM discovery, as shown in Figure 4-25.

Network Devices Discovery Wizard

General Properties

General Properties

Discovery Method

Default Accounts

Devices

Schedule Discovery

Summary

Completion

Specify general properties

Name:

Description (optional):

Select a management or gateway server

Select an Operations Manager management server or gateway server to run the discovery. A server can run only one network discovery. Servers that already run a network discovery do not appear in the list.

Available servers:

Select a resource pool

Select an Operations Manager resource pool for monitoring of discovered network devices.

Available pools:

Create Resource Pool

Add SNMP V3 Run As Account

Figure 4-25 Edit Discovery Rule

4. Edit the name (if wanted) and click **Next** twice to open the Devices page.
5. Click **Add**. The Add a Device window opens, as shown in Figure 4-26.

Add a Device

Specify the settings for the network device you want to discover.

Name or IP address:

Access mode:

SNMP version:

Port number:

SNMP V3 Run As account:

Add SNMP V3 Run As Account

Figure 4-26 Add a Device window

Specify the IP address of the CMM. Set the Access mode to ICMP and SNMP or SNMP and then select SNMP version **v3**. Select **Run As account** or **Add new** if you have different credentials for each device. Complete the following steps:

- a. To define a new Run As account, click **Add SNMP V3 Run As Account**. Then, click **Next** in the Introduction page and enter the name and description of the new account. Click **Next**, as shown in Figure 4-27.

Figure 4-27 Define display name

- b. Specify the credentials that were configured in CMM for SNMPv3 and click **Create**, as shown in Figure 4-28.

Figure 4-28 Credentials for SNMPv3 CMM account

6. You can add more devices or you can continue by clicking **Next**, as shown in Figure 4-29.

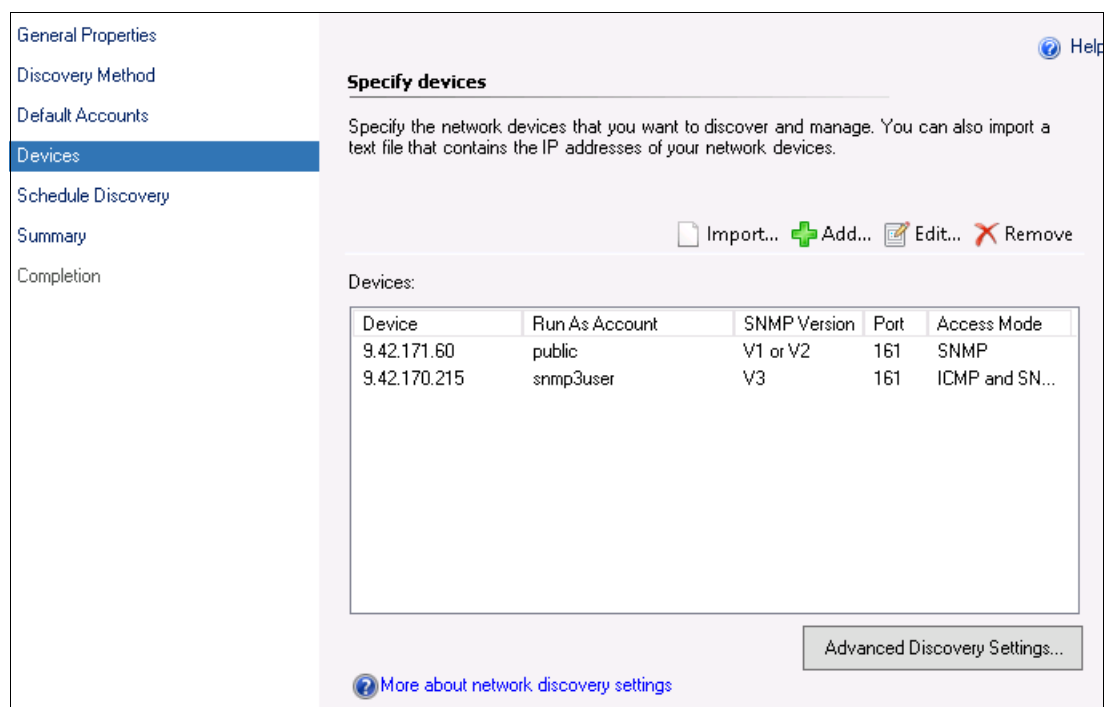


Figure 4-29 Specify devices window

7. Review the Schedule Discovery and Summary sections, or continue by clicking **Next**.
8. On the Completion page, select **Run the network discovery rule after the wizard is closed**. Click **Close**.

Note: It can take several hours for a new device to be discovered with all monitors enabled in SCOM. You can check whether the device discovery was successful in the Operations Manager logs that are in Windows Event Viewer.

After the discovery is completed, you will see your discovered BladeCenter and Flex System chassis in the Network Devices view of the Administration panel in the SCOM, as shown in Figure 4-30.

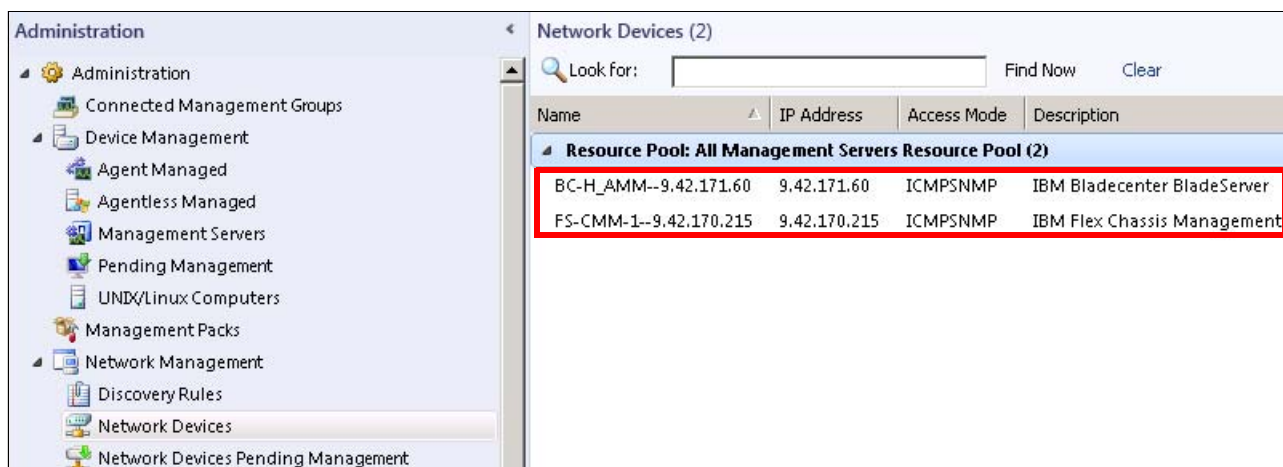


Figure 4-30 Network Devices view after the discovery completed

4.3.2 Deploying System Center agents for hardware monitoring and inventory

To enable operating system monitoring, enable the Lenovo Hardware Performance and Resource Optimization Pack for Microsoft System Center Virtual Machine Manager (SCVMM) or Lenovo Inventory Tool for Microsoft System Center Configuration Manager (SCCM). More management agents must be deployed to the Windows Operating system that is installed on a Flex System compute node.

Deploying Microsoft SCOM agent

SCOM agent is required to enable operating system monitoring with enabling Performance and Resource Optimization (PRO) tips in SCVMM.

Complete the following steps to install SCOM agent:

1. Log in to the Microsoft SCOM operations console as Administrator.
2. Click **Administration**. Right-click **Device Management** → **Discovery Wizard**, as shown in Figure 4-31.

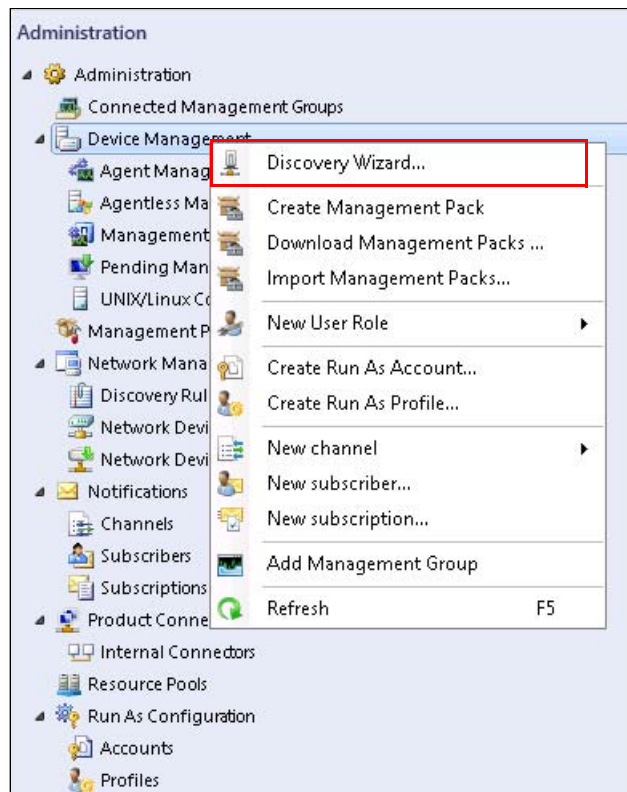


Figure 4-31 Selecting Discovery Wizard option

3. Select **Windows Computers** and click **Next**, as shown in Figure 4-32.

Discovery Type

Auto or Advanced?

Discovery Method

Administrator Account

Select Objects to Manage

Summary

Choose the type of computers or devices to discover and manage.

Windows computers
Discover Windows computers in your Active Directory environment and install agents on the ones you want to manage.

UNIX/Linux computers
This enables you to discover UNIX and Linux computers in your environment and install agents on the ones you want to manage.

Network devices
Discover and monitor network devices using Simple Network Management Protocol (SNMP).

Select a discovery type and click Next to continue.

Figure 4-32 Select Windows computers

4. Specify the discovery method. In larger environments, it might be faster to select **Advanced discovery**, as shown in Figure 4-33. Click **Next**.

Discovery Type

Auto or Advanced?

Discovery Method

Administrator Account

Select Objects to Manage

Summary

Choose automatic or advanced discovery

☐ Automatic computer discovery
Scans the "LAB" domain for all Windows-based computers.

☒ **Advanced discovery**
Allows you to specify advanced discovery options and settings.

Computer and Device Classes:
Servers Only

Note: This setting applies only when scanning Active Directory. You can configure how objects will be discovered, on the next screen(s).

Management Server
SCOM.lab.local

☒ Verify discovered computers can be contacted

Figure 4-33 Discovery Method

5. Select **Scan Active Directory** and click **Configure**. Enter the computer name or prefix, as shown in Figure 4-34. Click **OK**, then click **Next**.

The screenshot shows the 'Find Computers' dialog box with the 'Advanced' tab active. The 'Computer name' field contains the text 'x240'. The 'Owner' field is empty. The 'Role' dropdown menu is set to 'Any'. In the background, the 'Scan Active Directory' configuration window is visible, showing the 'Configure...' button.

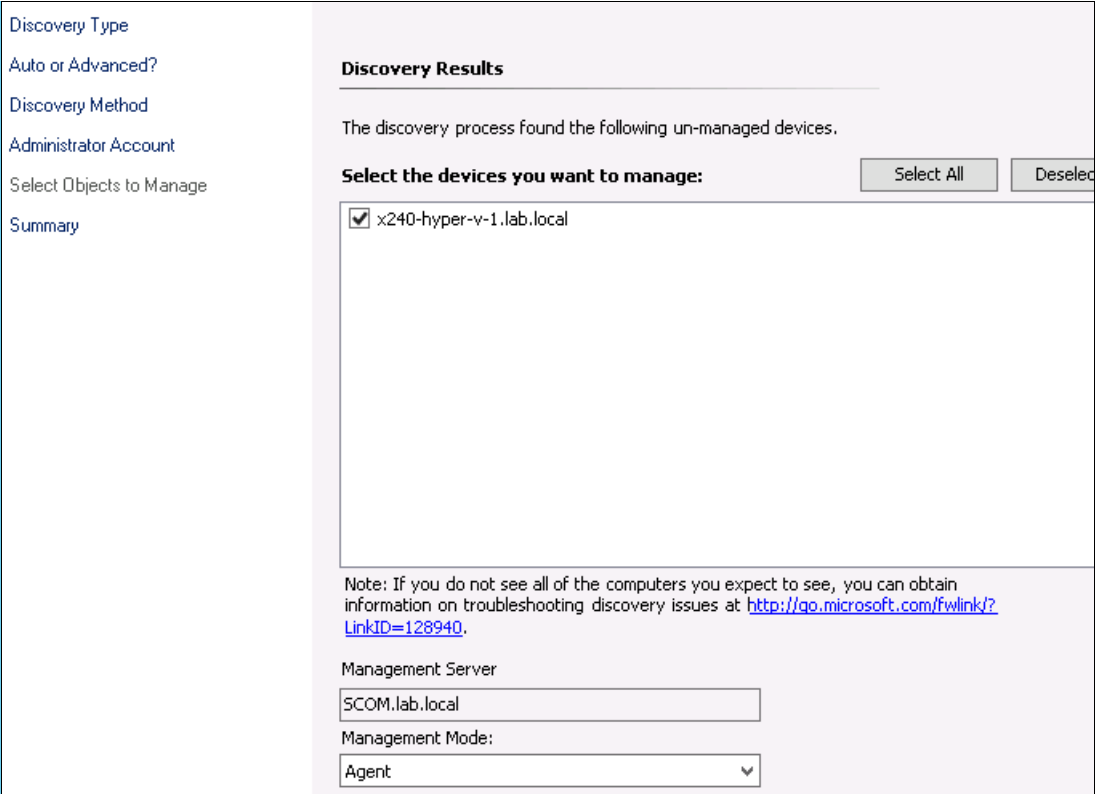
Figure 4-34 Specify computer name or prefix

6. Click **Discover** or specify another user account for discovery and agent installation, as shown in Figure 4-35. The user must have administrator privileges on the target server.

The screenshot shows the 'Administrator Account' configuration window. The 'Use selected Management Server Action Account' radio button is selected. The 'Other user account' radio button is unselected. The 'User name' field is empty. The 'Password' field is empty. The 'Domain' dropdown menu is set to 'LAB'.

Figure 4-35 Specify Administrator Account

7. Select discovered servers for agent installation and click **Next**, as shown in Figure 4-36.



Discovery Type

Auto or Advanced?

Discovery Method

Administrator Account

Select Objects to Manage

Summary

Discovery Results

The discovery process found the following un-managed devices.

Select the devices you want to manage: Select All Deselect All

- ☒ x240-hyper-v-1.lab.local

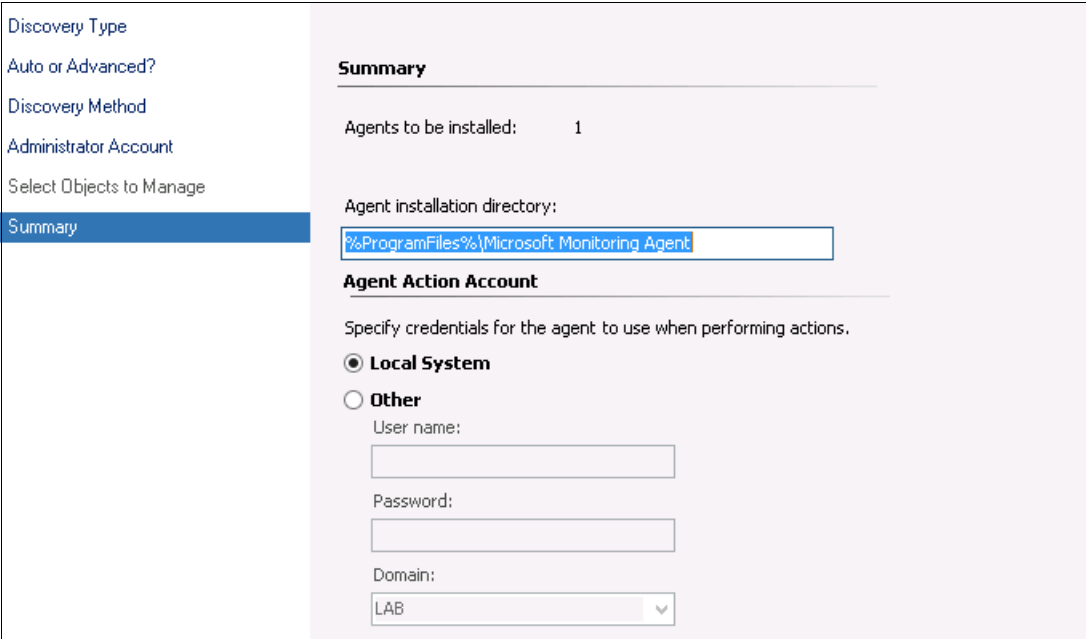
Note: If you do not see all of the computers you expect to see, you can obtain information on troubleshooting discovery issues at <http://go.microsoft.com/fwlink/?LinkID=128940>.

Management Server: SCOM.lab.local

Management Mode: Agent

Figure 4-36 Select discovered servers for agent deployment

8. Specify the agent installation folder and run as account for the agent that is based on your preferences and internal policies, as shown in Figure 4-37. Click **Finish**.



Discovery Type

Auto or Advanced?

Discovery Method

Administrator Account

Select Objects to Manage

Summary

Summary

Agents to be installed: 1

Agent installation directory: %ProgramFiles%\Microsoft Monitoring Agent

Agent Action Account

Specify credentials for the agent to use when performing actions.

☒ Local System

☐ Other

User name:

Password:

Domain: LAB

Figure 4-37 Installation Path and Run As policies

9. Monitor the deployment status as shown in Figure 4-38.

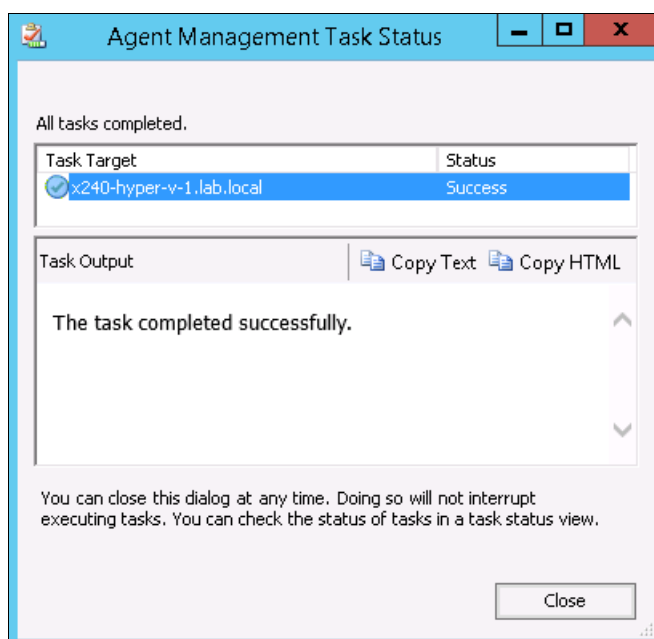


Figure 4-38 Agent deployed

If the target computer is a member of the Microsoft cluster, some management packs require management agent to be enabled in proxy mode. Complete the following steps:

1. Click **Administration** → **Device Management** → **Agent managed**. Right-click the wanted computer and click **Properties**.
2. Click **Security** and select **Allow this agent to act as a proxy and discover managed objects on other computers**, as shown in Figure 4-39.

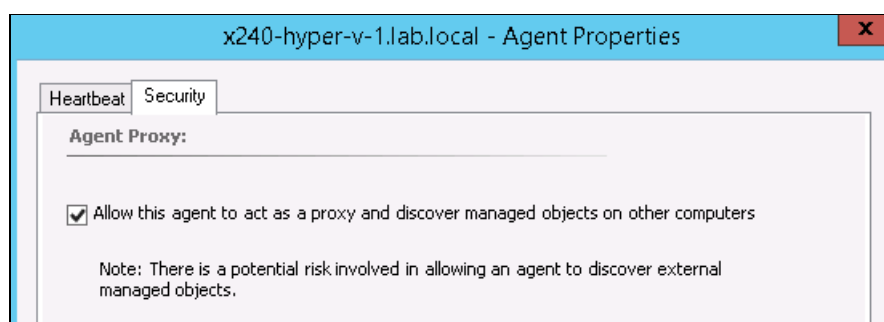


Figure 4-39 Allow agent to act as a proxy

Important: You must install Systems Director Platform agent for System x to enable monitoring of some Lenovo hardware components in SCOM.

For more information about Systems Director agents releases, see this website:

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

Deploying Microsoft System Center Configuration Manager agent

To enable Inventory Tool Client for SCCM, SCCM agent with Lenovo Inventory Tool Client must be deployed to client computer

Note: Lenovo Inventory Tool Client requires Microsoft .NET Framework Version 2.0 on client machines. It is included in .NET Framework 3.5 Features in Add Roles and Features Wizard in Windows Server operating system.

Based on your SCCM configuration, all agents can deploy automatically. In this section, we describe how to create a devices collection and how to deploy SCCM agent and Lenovo Inventory Tool Client.

Creating a Device Collection in SCCM

Complete the following steps to create a Device Collection:

1. Log in to the SCCM console as Administrator.
2. Click **Assets and Compliance** → **Device Collection** → **Create Device Collection**, as shown in Figure 4-40.

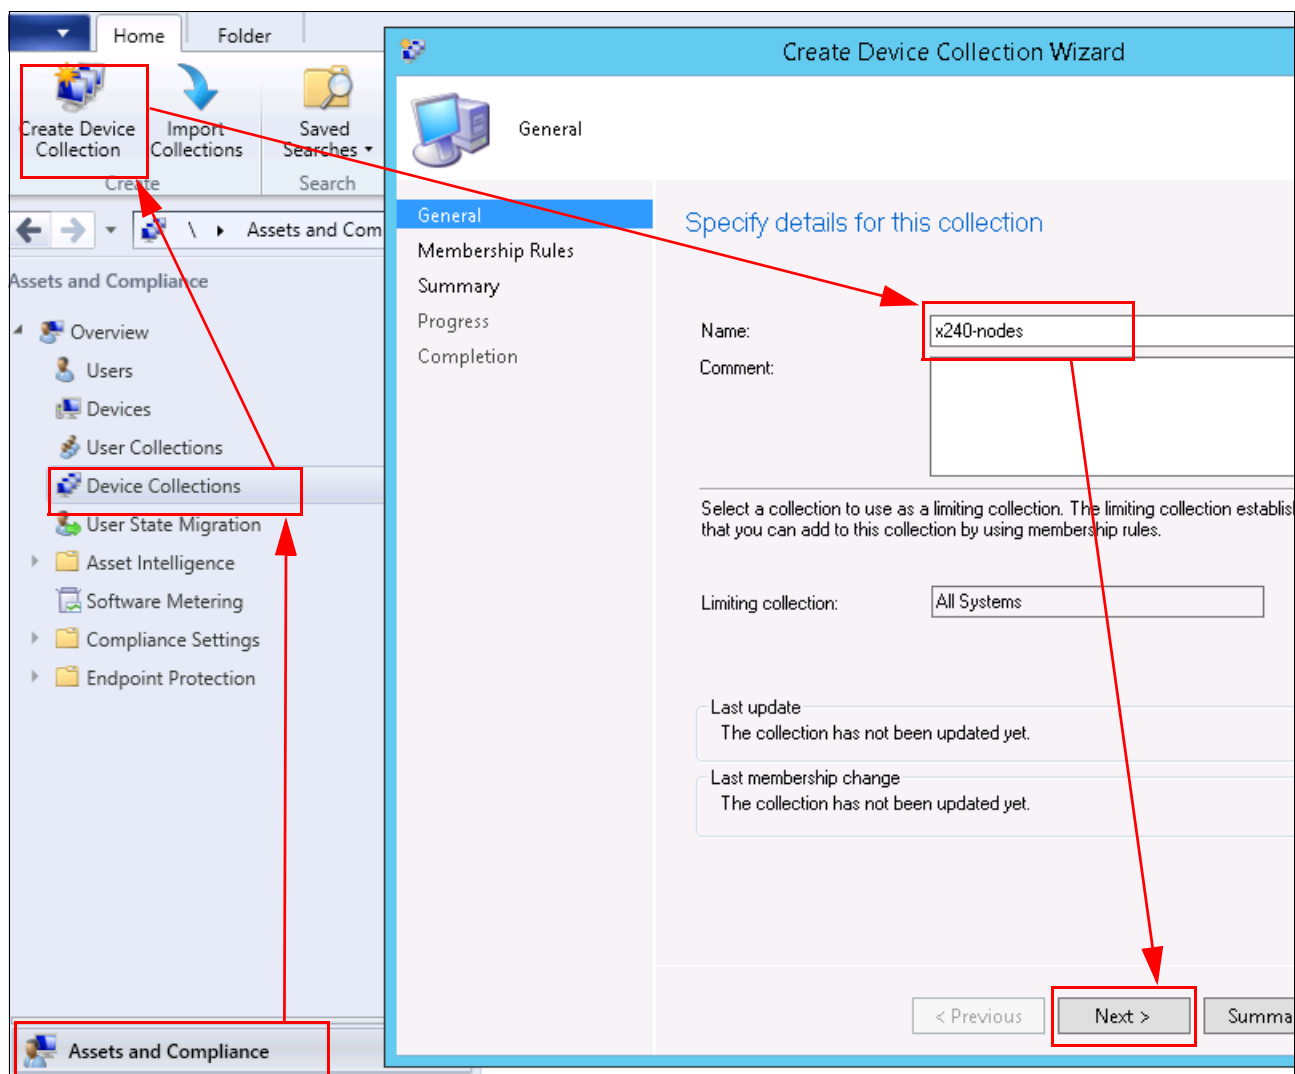


Figure 4-40 Create Device Collection wizard

3. Enter the Name of collection and click **Next** to open the Membership Rules page.
4. Click **Add Rule** → **Direct Rule**. The **Create Direct Membership Rule Wizard** window opens, as shown in Figure 4-41.

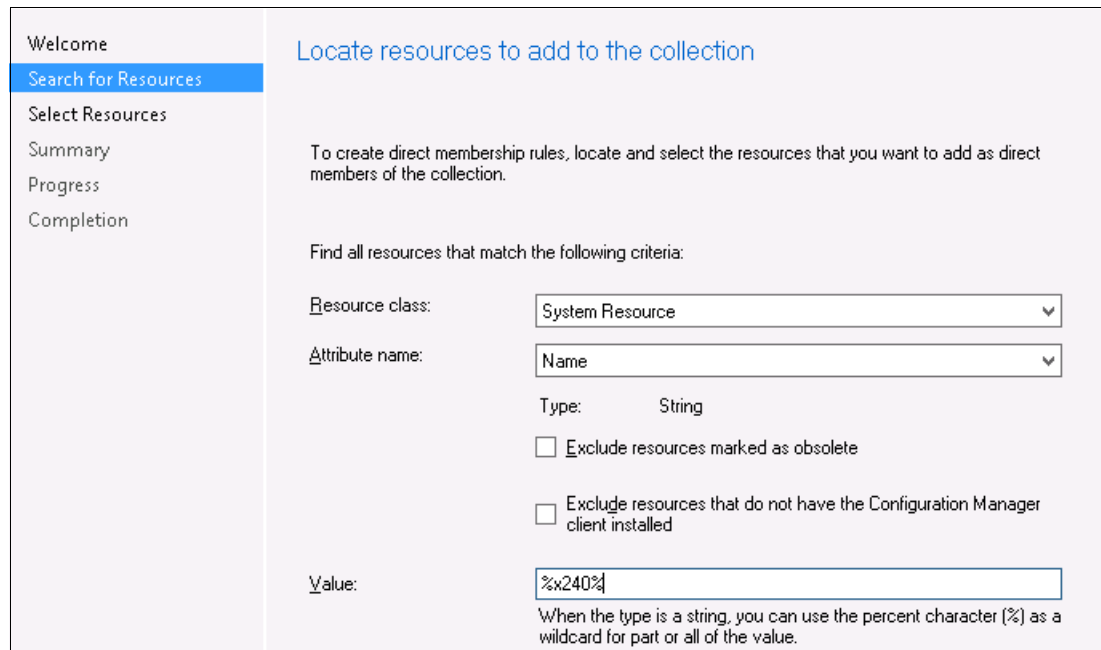


Figure 4-41 Create Direct Membership Rule wizard

5. To search for the computers of which the name includes the string x240, select **Resource Class**. Select **Name** as the Attribute name and enter %x240% into the Value field. Clear all check boxes. Click **Next**.
6. Select the wanted servers from searches, as shown in Figure 4-42. Click **Next** twice and wait for the rule to be created. Close the wizard.

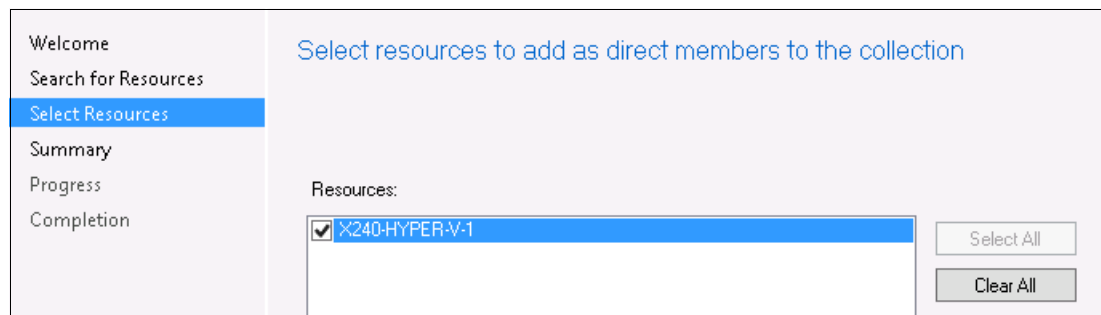


Figure 4-42 Select servers

7. In the Create Devices Collection wizard, select **Incremental updates for this collection**, as shown in as shown in Figure 4-43. You also can configure the schedule, review the membership rules, and click **Next**.

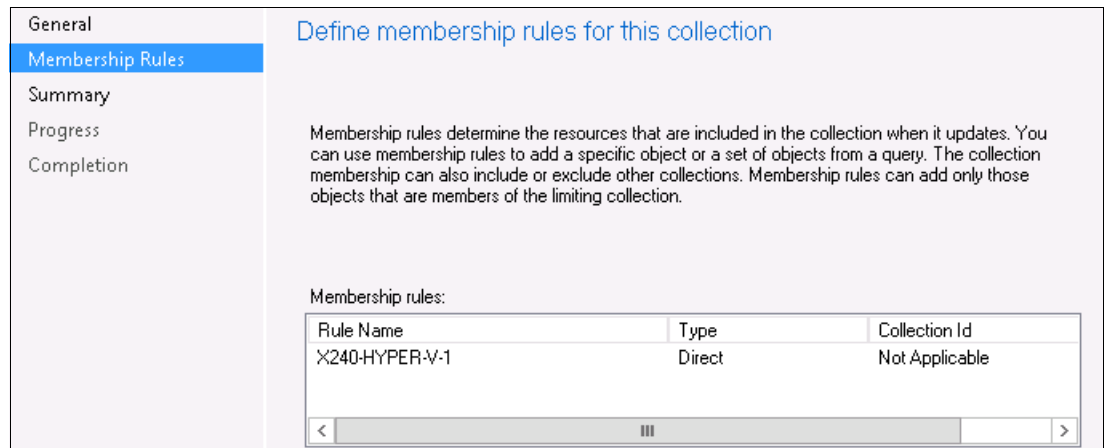


Figure 4-43 Review membership rules for collection

8. Complete the wizard and the collection is created.

Deploying SCCM agent

Complete the following steps to deploy SCCM agent:

1. In the Assets and Compliance view, select the newly created collection and click **Install Client**, as shown in Figure 4-44.

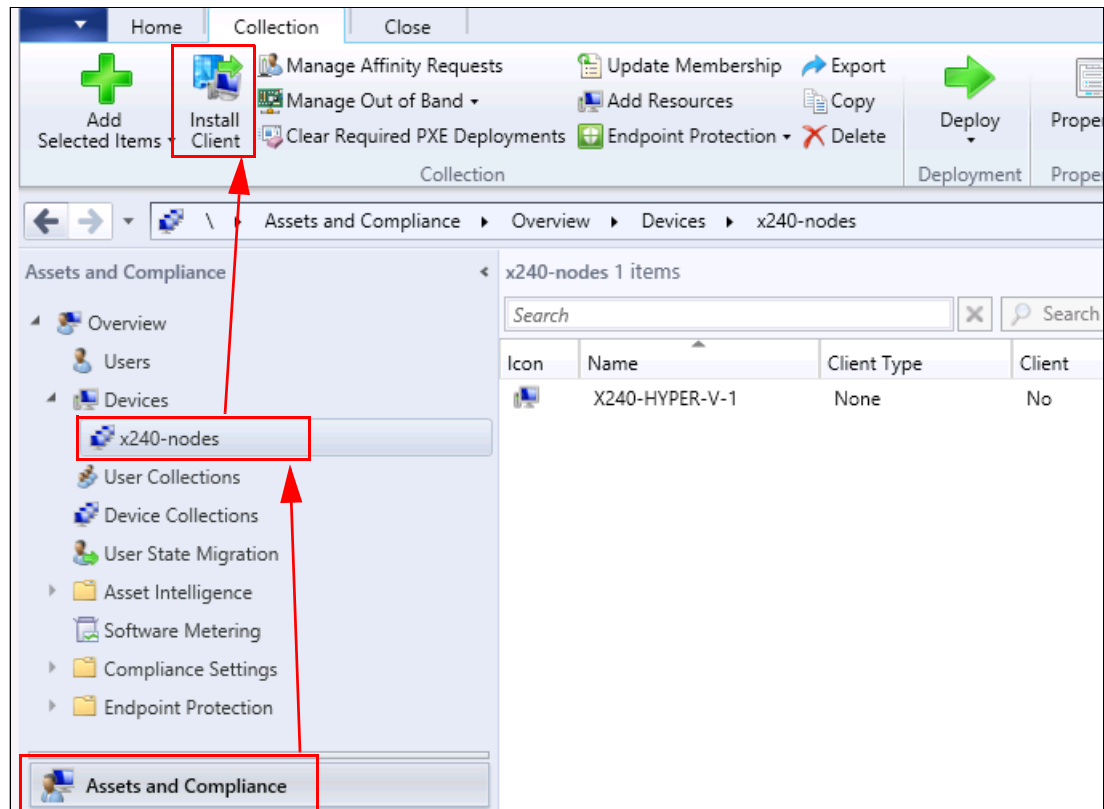


Figure 4-44 Devices collection view

2. Click **Next** in the Before You Begin page.
3. Keep the default selections in the Installation Options page or select some options if wanted (see Figure 4-45). Click **Next** twice.

Before You Begin

Installation Options

Summary

Progress

Completion

Specify Client Push Options

☐ Allow the client software to be installed on domain controllers
If you have configured client push installation to domain controllers in the Client Push Installation Properties dialog box, this option is unavailable.

☐ Always install the client software
When a computer already has the Configuration Manager client installed, you can repair, upgrade, or reinstall the client software.

☐ Uninstall existing Configuration Manager client before the client is installed

☐ Install the client software from a specified site

Site: LAB-Lab Headquarters

The site server in the specified site will install the client software. When you do not use this option, the site server in the assigned site for the resource will install the client software.

Figure 4-45 Install Configuration Manager Client wizard

4. Close the wizard after the process is completed.

Deploying Inventory Tool Client by using SCCM

Perform the following steps to deploy Inventory Tool Client:

1. In the Software Library window, click **Packages**.
2. Right-click **Inventory Tool Client** and click **Deploy**, as shown in Figure 4-46.

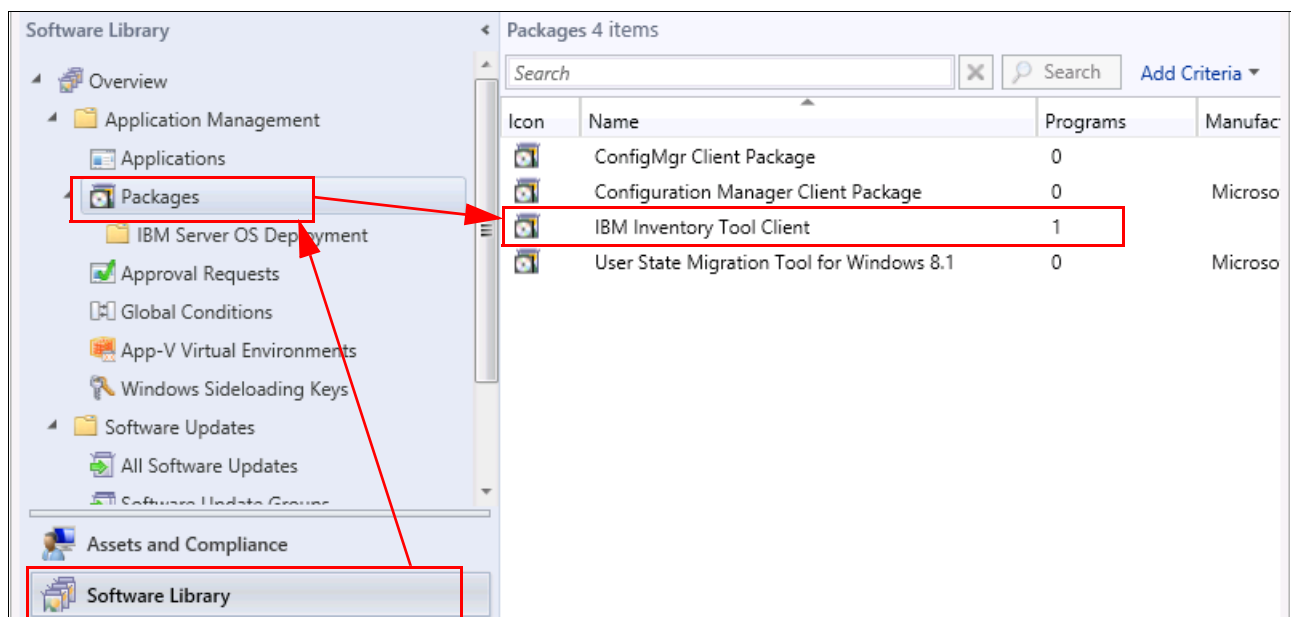


Figure 4-46 Locate Inventory Tool Client

3. Select the target collection and click **Next**, as shown in Figure 4-47.

General

Specify general information for this deployment

Software: IBM Inventory Tool Client (IBM.InventoryTool.Client.Setup.msi) Browse...

Collection: x240-nodes Browse...

☐ Use default distribution point groups associated to this collection

☐ Automatically distribute content for dependencies

Comments (optional):

Figure 4-47 Deploy Software wizard: General

4. In the Content page, click **Add** to add distribution points or distribution point groups, as shown in Figure 4-48. Click **Next**.

General

Content

Specify the content destination

Distribution points or distribution point groups that the content has been distributed to:

Name	Type
\\SCCM.lab.local	Distribution point

Additional distribution points, distribution point groups, and the distribution point groups that are currently associated with collections to distribute content to:

Filter...

Add

Name	Description	Associations
There are no items to show in this view.		

Remove

Figure 4-48 Deploy Software wizard: Content

5. In the Deployment Settings page, specify the settings that control how the software is deployed, as shown in Figure 4-49. Click **Next**.

General
Content
Deployment Settings
Scheduling
User Experience
Distribution Points
Summary
Progress
Completion

Specify settings to control how this software is deployed

Action:

Purpose:

☐ Pre-deploy software to the user's primary device

☐ Send wake-up packets

☐ Allow clients on a metered Internet connection to download content after the installation deadline, which might incur additional costs

Figure 4-49 Deploy Software wizard: Deployment settings

6. In the Scheduling page, configure the deployment and assignment schedule (if necessary). Select **As soon as possible** if you want to deploy the inventory tool client immediately, as shown in Figure 4-50. Click **Next**.

General
Content
Deployment Settings
Scheduling
User Experience
Distribution Points
Summary
Progress
Completion

Specify the schedule for this deployment

This program will be available as soon as it has been distributed to the content servers unless it is scheduled for a later time below. For required applications, specify the assignment schedule.

☐ Schedule when this deployment will become available:
 ☐ UTC

☐ Schedule when this deployment will expire:
 ☐ UTC

Assignment schedule:

Rerun behavior:

Figure 4-50 Deploy Software wizard: Scheduling

7. In the User Experience page, modify the settings as needed. Click **Next** twice. Close the wizard after the process is completed.

Note: It can take several hours for all of the agents and clients to be deployed and to reflect all monitors in the SCOM console. In addition, hardware inventory information in SCCM can appear 24 hours after the client is installed.

4.3.3 Monitoring hardware status in SCOM

After you deploy monitoring agents onto Flex System node's operating system, you can monitor the status of the systems hardware components for both BladeCenter and Flex System in SCOM monitoring.

Note: We are showing a few views only for demonstration purposes. For more information, see the Lenovo Hardware Management Pack for Microsoft System Center Operations Manager User's Guide, which is available at this website:

<http://www-947.ibm.com/support/entry/portal/docdisplay?lnodocid=MIGR-5082204>

In the Navigation pane, click **Monitoring** and expand Lenovo Hardware. Here, you can select from various monitors to see important information about your environment health.

Click **Lenovo Licensed System Group** under the Monitoring to see the list of your BladeCenter and Flex System managed servers, as shown in Figure 4-51.

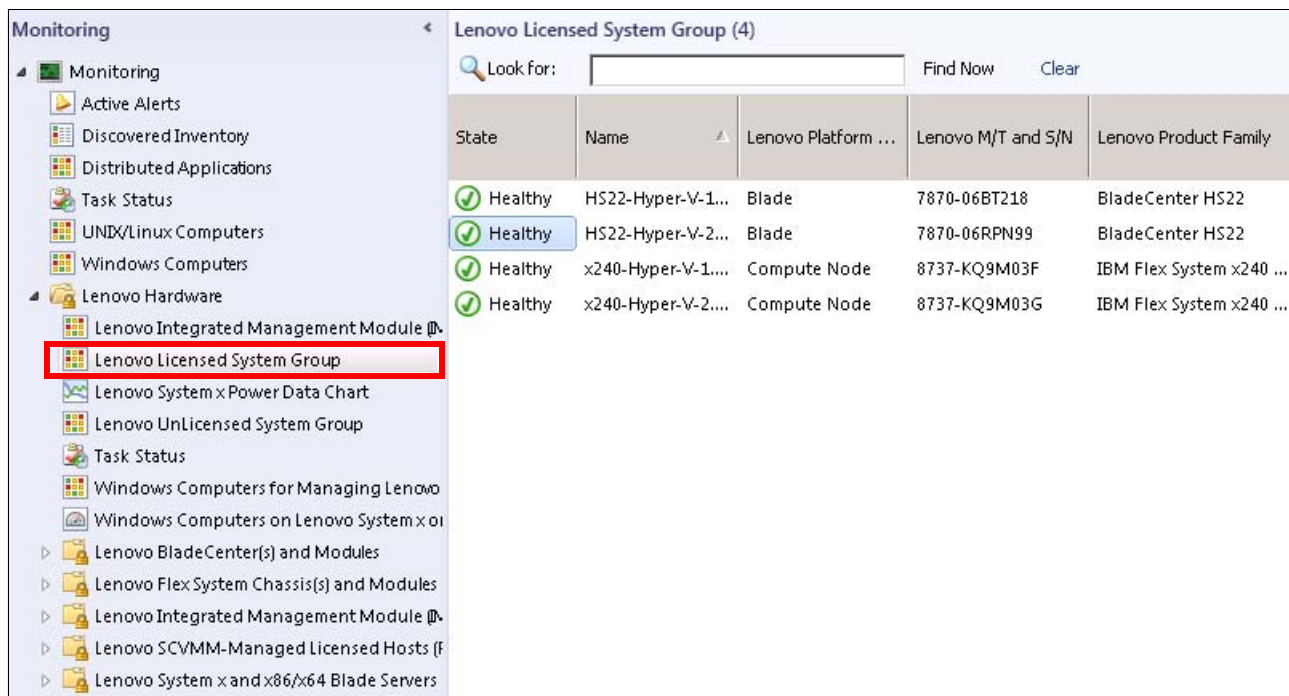


Figure 4-51 Lenovo Licensed System Group

Figure 4-51 also shows other available Lenovo groups to monitor BladeCenter and Flex System hardware components, such as:

- ▶ Lenovo BladeCenter(s) and Modules
- ▶ Lenovo Flex System Chassis(s) and Modules
- ▶ Lenovo Integrated Management Module (IMM)
- ▶ Lenovo SCVMM-Managed Licensed Hosts
- ▶ Lenovo System x and x86/x64 Blade Servers

For example, expand **Lenovo BladeCenter(s) and Modules**, expand **Lenovo BladeCenter Modules**, and click **Lenovo BladeCenter Chassis** to check the status of the BladeCenter chassis, as shown in Figure 4-52.

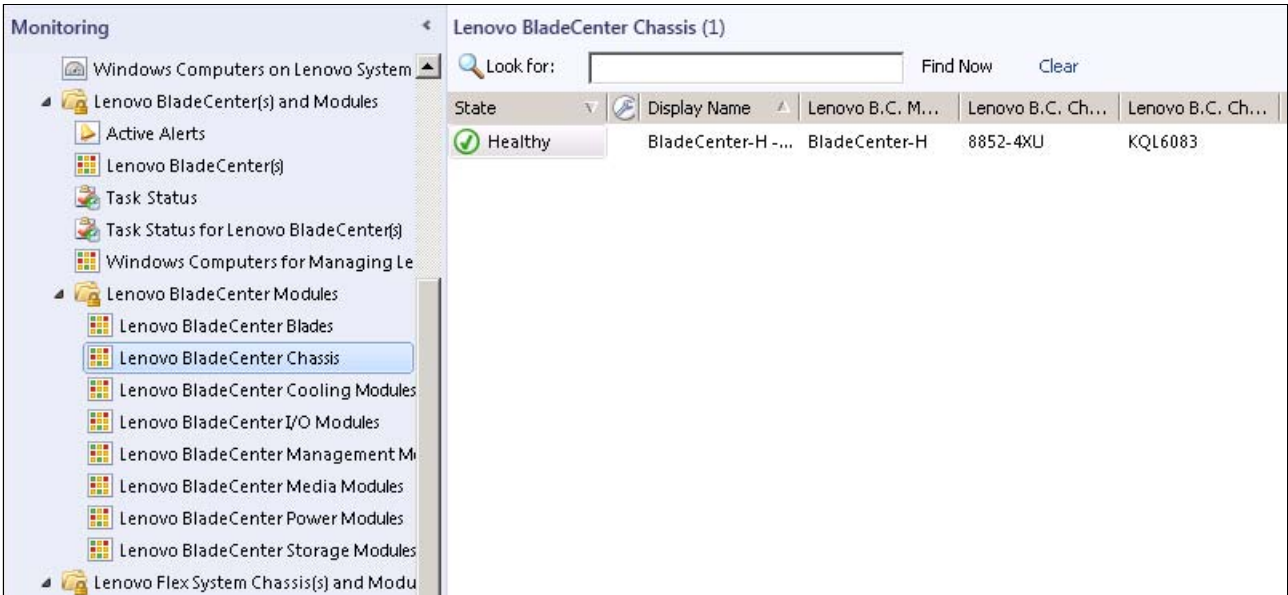


Figure 4-52 BladeCenter chassis status

You can check the status of other BladeCenter components by clicking the respective group. For example, the BladeCenter blade server status is shown under the **Lenovo BladeCenter Blades** group, as shown in Figure 4-53, and the I/O module status is shown under the **Lenovo BladeCenter I/O Modules** group, as shown in Figure 4-54 on page 115.

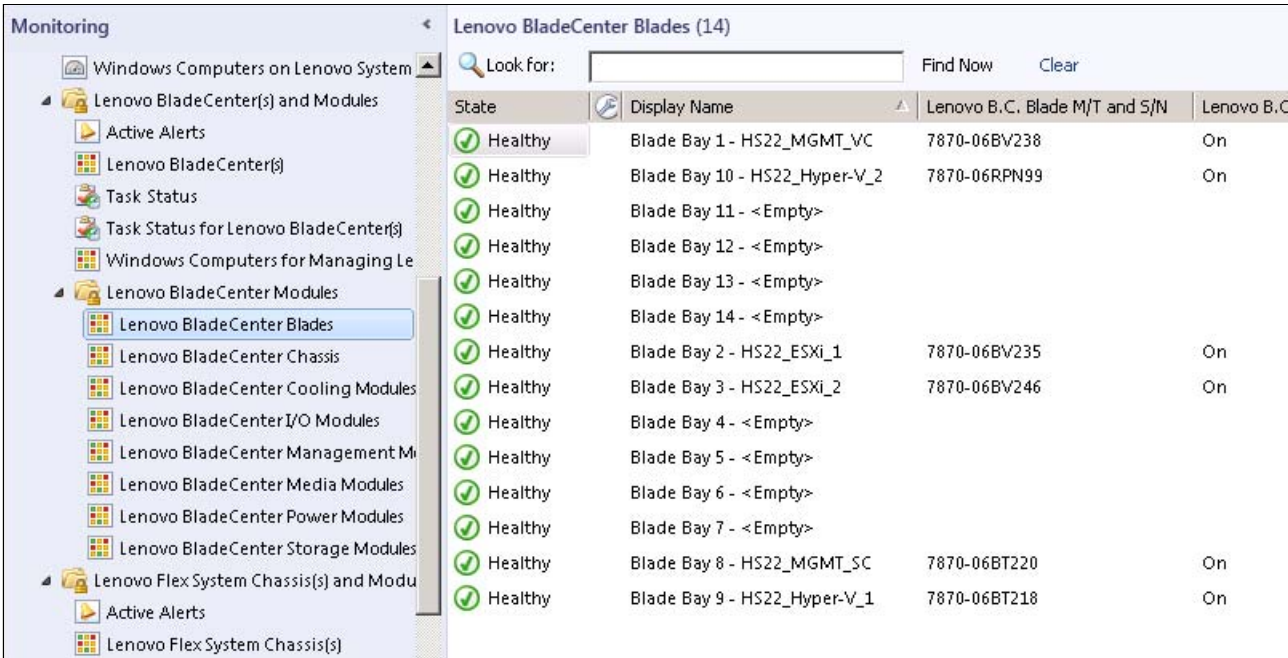


Figure 4-53 BladeCenter blades status

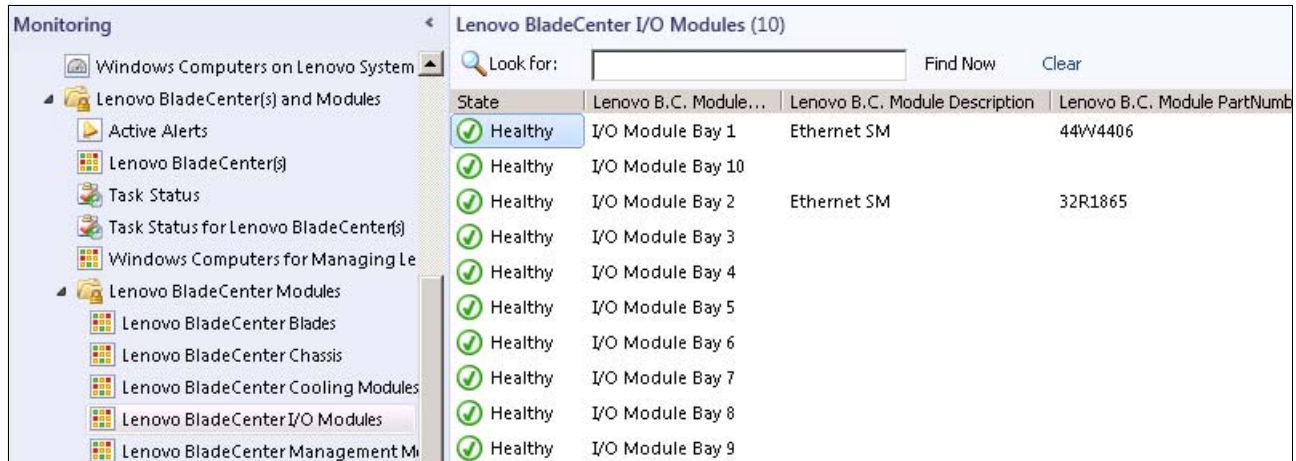


Figure 4-54 BladeCenter I/O modules status

Expand **Lenovo Flex System Chassis(s) and Modules** and click **Lenovo Flex System Chassis(s)** to check the status of the Flex System chassis, as shown in Figure 4-55.

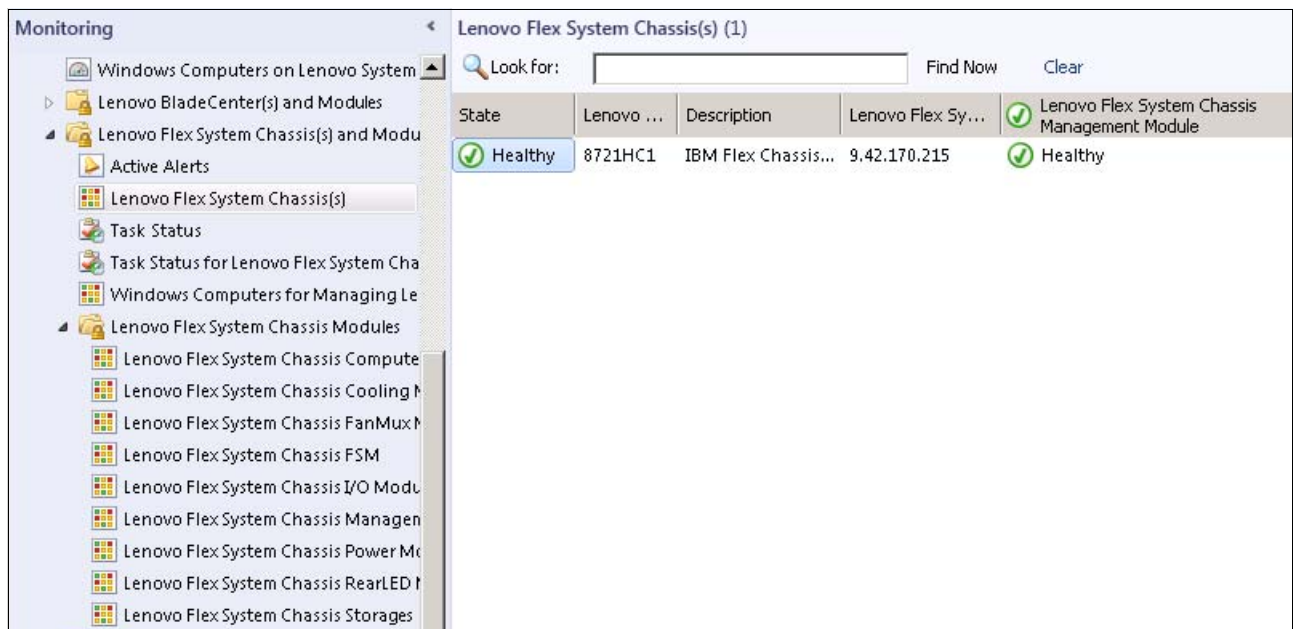


Figure 4-55 Flex System chassis status

You can check the status of other Flex System components by clicking the respective group. For example, the Flex System compute node status is shown under the Lenovo Flex System Compute Nodes group, as shown in Figure 4-56, and the I/O module status is shown under the Lenovo Flex System I/O Modules group, as shown in Figure 4-57.

Monitoring					
Lenovo Flex System Chassis Compute Nodes (14)					
Look for: <input type="text"/> Find Now Clear					
State	Lenovo Flex S...	Lenovo Flex Syst...	MachineTypeModel	Lenovo Flex Syste...	
Healthy	Node Bay 1	Flex System x240 ...	8737AC1	On	
Healthy	Node Bay 10				
Healthy	Node Bay 11				
Healthy	Node Bay 12				
Healthy	Node Bay 13				
Healthy	Node Bay 14				
Healthy	Node Bay 2	Flex System x240 ...	8737AC1	On	
Healthy	Node Bay 3	Flex System x240 ...	8737AC1	On	
Healthy	Node Bay 4	Flex System x240 ...	8737AC1	On	
Healthy	Node Bay 5				
Healthy	Node Bay 6				
Healthy	Node Bay 7				
Healthy	Node Bay 8				
Healthy	Node Bay 9				

Figure 4-56 Flex System compute nodes status

Monitoring					
Lenovo Flex System Chassis I/O Modules (4)					
Look for: <input type="text"/> Find Now Clear					
State	Lenovo Flex Sy...	Lenovo Flex System Module Description	Lenovo ...	Lenovo Flex	
Healthy	I/O Module Bay 2				
Healthy	I/O Module Bay 4				
Healthy	I/O Module Bay 1	EN4093 10Gb Ethernet Switch	49Y4272	On	
Healthy	I/O Module Bay 3	FC3171 8Gb SAN Switch	69Y1932	On	

Figure 4-57 Flex System I/O modules status

For a single view of all Lenovo x86 systems, including BladeCenter servers and Flex System compute nodes and the status of their hardware components, expand **Lenovo System x and x86/x64 Blade Servers** and click **All Lenovo System x and x86/x64 Blade Servers**, as shown in Figure 4-58.

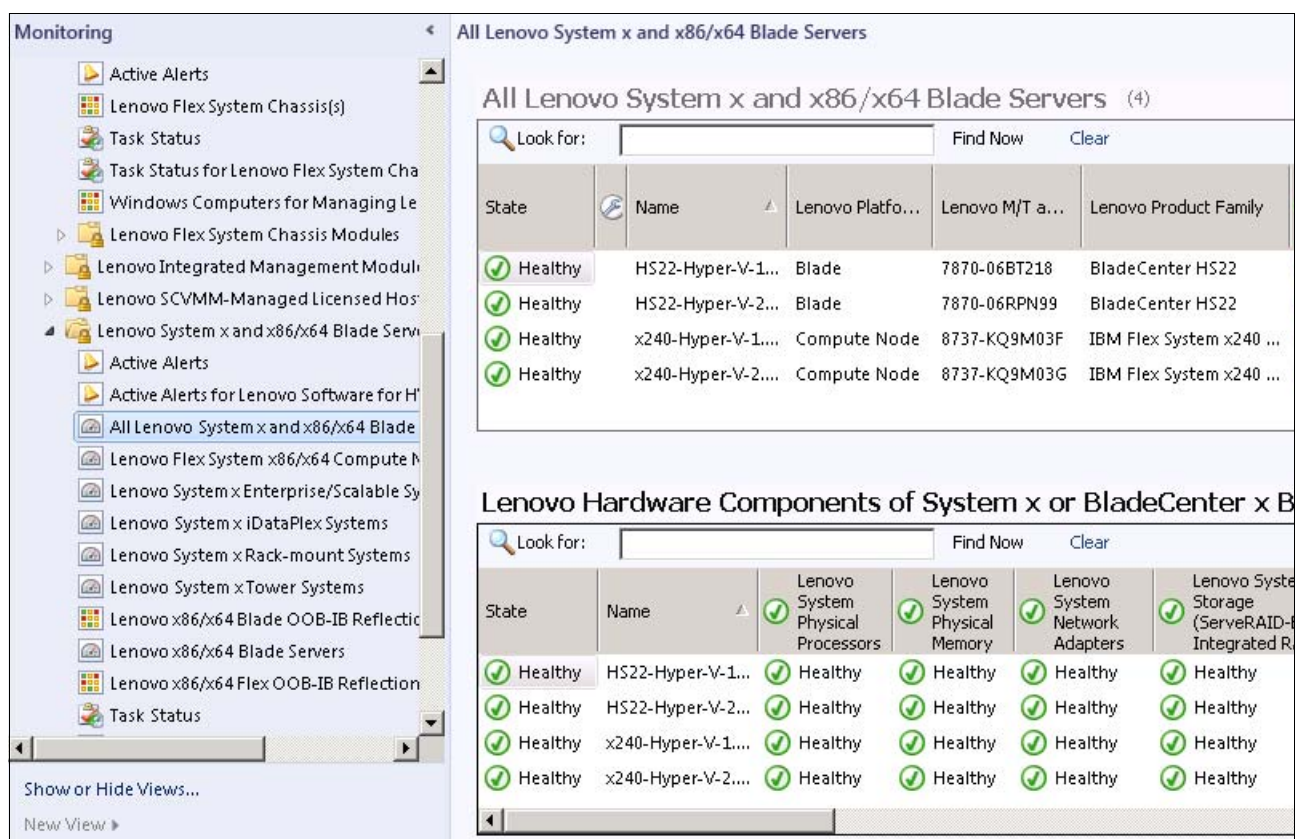


Figure 4-58 All System x and x86/x64 Blade Servers

4.3.4 Lenovo Hardware Performance and Resource Optimization Pack for VMM

By using the Lenovo Hardware Performance and Resource Optimization Pack (PRO) for Microsoft SCVMM, you can monitor and manage alerts for the physical host resources in a virtualized environment.

PRO includes the following key features:

- ▶ Automated VM Migration support. This support is based on hardware failure events or power consumption threshold exceptions for UEFI or IMM System x servers and blades that are running Windows 2012, Windows 2008 and 2008 R2, Hyper-V, or Virtual Server.
- ▶ Advisory PRO tips if existing or predictive hardware problems occur that warrant VMM administrative operations.

Note: For more information, see the following Lenovo Hardware Performance and Resource Optimization Pack for Microsoft System Center Virtual Machine Manager website:

<http://www-947.ibm.com/support/entry/portal/docdisplay?lnidocid=MIGR-5082203>

If PRO Monitors were enabled on the existing cluster or Host Group in SCVMM, they are enabled automatically after SCOM agent is deployed. You can verify the PRO status in the SCOM console by clicking **Monitoring** → **Lenovo Hardware** → **Lenovo SCVMM-Managed Licensed Hosts (PRO Views)** → **Licensed SCVMM-Managed Hosts PRO Status**, as shown in Figure 4-59.

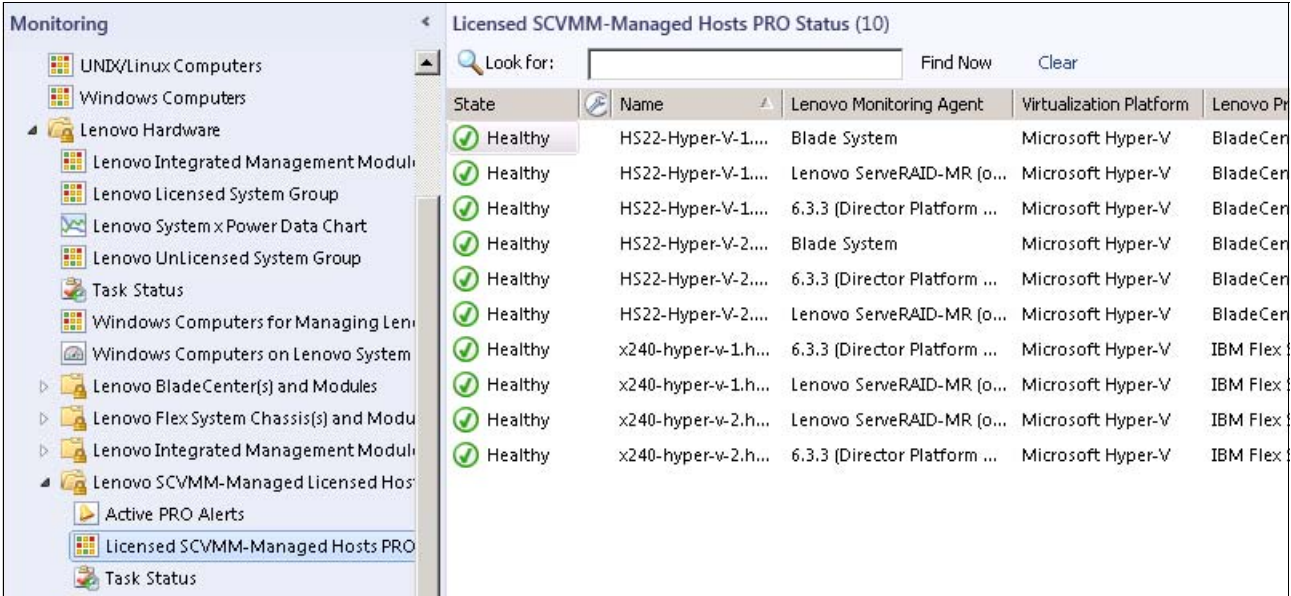


Figure 4-59 Licensed SCVMM-managed hosts status

Based on the SCVMM console settings, new PRO tips can appear as pop-up windows, as shown in Figure 4-60.

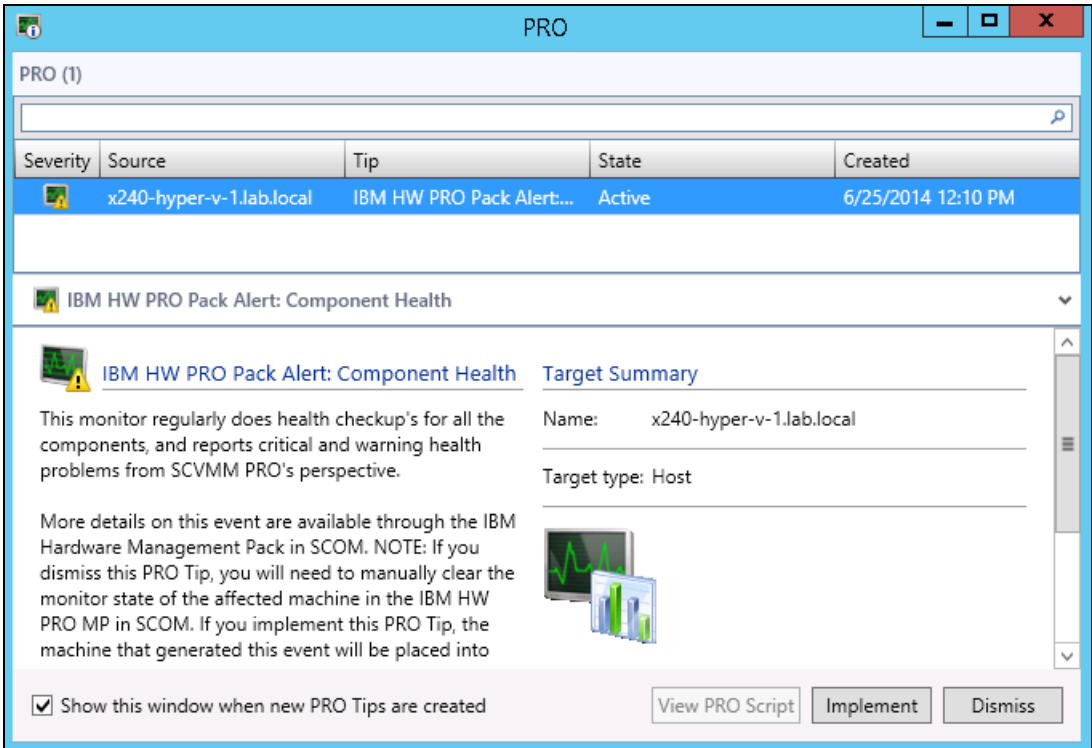


Figure 4-60 SCVMM PRO tip pop-up window

4.3.5 Rolling firmware upgrades by using UIM for System Center VMM

Upward Integration Modules Add-in for Microsoft System Center Virtual Machine Manager (VMM) provides non-disruptive system firmware updates in clustered environment.

Note: For more information, see the following Upward Integration Modules Add-in for Microsoft System Center Virtual Machine Manager website:

<http://www-947.ibm.com/support/entry/portal/docdisplay?lnodocid=MIGR-5095711>

After you join a new system to the cluster, you must set up the authentication information for the new hosts. Complete the following steps:

1. Start SCVMM console as OS administrator and login as SCVMM administrator.
2. In the Fabric view, select the wanted cluster and click the UIM icon.
3. Click the host that is marked red. Then, click **Set Auth Info**.

This process is shown in Figure 4-61.

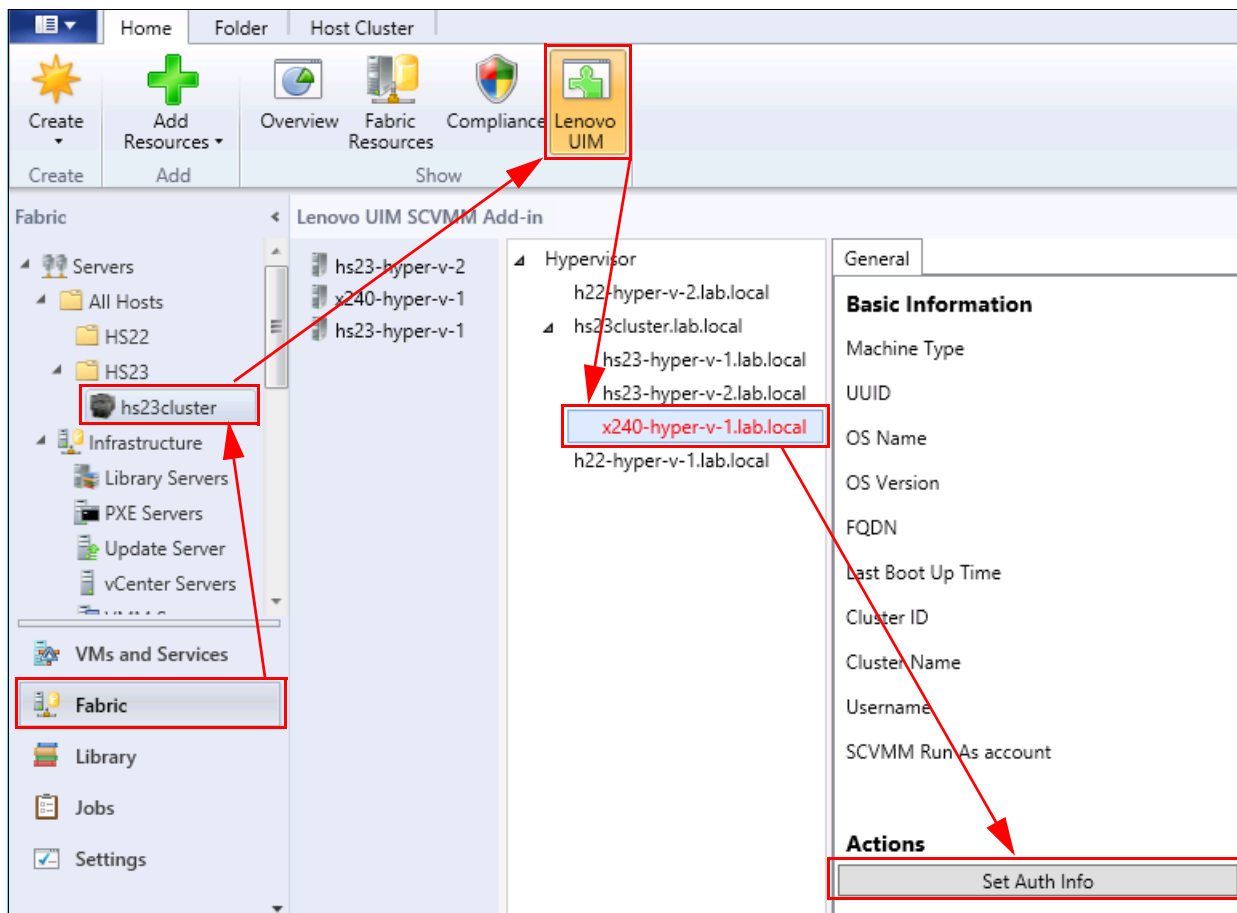
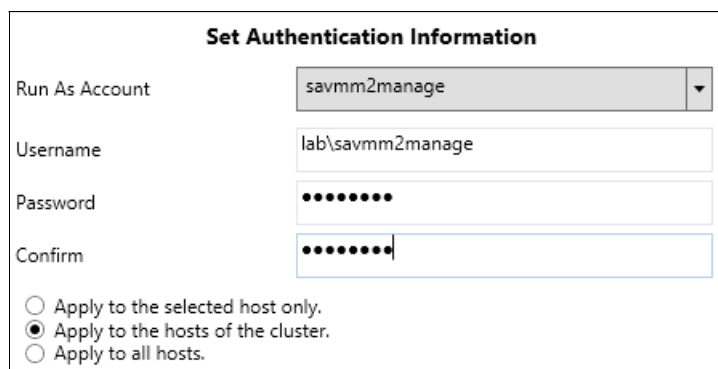


Figure 4-61 New system in UIM for SCVMM

- Specify Run As Account for SCVMM job and administrator account for new server. You can apply these credentials for All Hosts, Hosts in Cluster, or selected Host only, as shown in Figure 4-62. Click **OK**.



Set Authentication Information

Run As Account: savmm2manage

Username: lab\savmm2manage

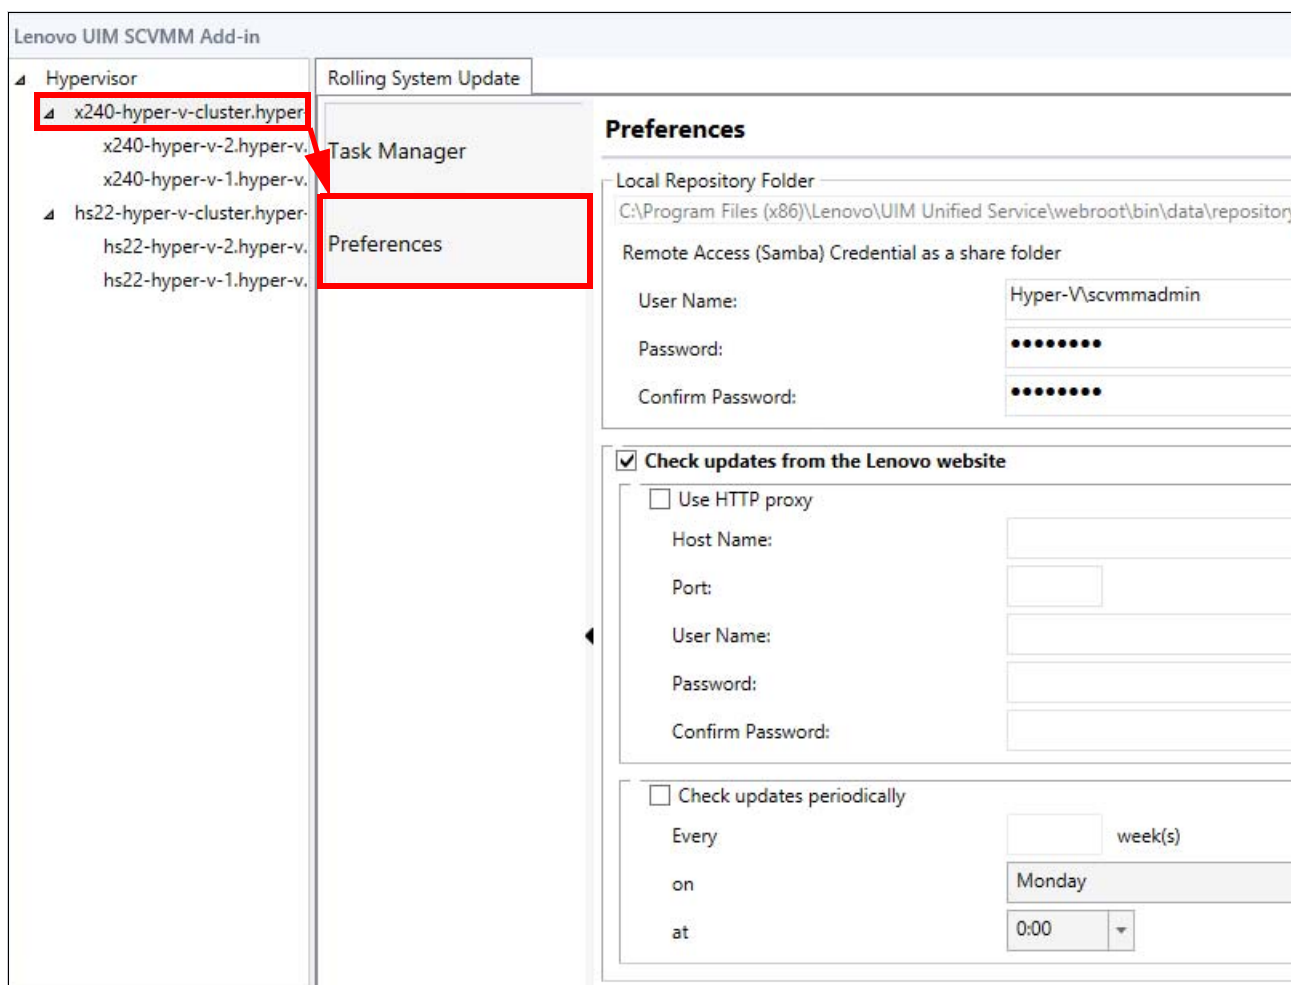
Password:

Confirm:

☐ Apply to the selected host only.
☒ Apply to the hosts of the cluster.
☐ Apply to all hosts.

Figure 4-62 Authentication Information window

Next, if not already done, you need to specify preferences for the local repository folder, including access credentials and firmware download schedule. Click the cluster name, then click **Preferences**, as shown in Figure 4-63. Click **Save**.



Lenovo UIM SCVMM Add-in

Hypervisor

- x240-hyper-v-cluster.hyper-v
 - x240-hyper-v-2.hyper-v
 - x240-hyper-v-1.hyper-v
- hs22-hyper-v-cluster.hyper-v
 - hs22-hyper-v-2.hyper-v
 - hs22-hyper-v-1.hyper-v

Rolling System Update

Task Manager

Preferences

Preferences

Local Repository Folder
C:\Program Files (x86)\Lenovo\UIM Unified Service\webroot\bin\data\repository

Remote Access (Samba) Credential as a share folder

User Name: Hyper-V\scvmmadmin

Password:

Confirm Password:

☒ Check updates from the Lenovo website

☐ Use HTTP proxy

Host Name:

Port:

User Name:

Password:

Confirm Password:

☐ Check updates periodically

Every week(s)

on Monday

at 0:00

Figure 4-63 Rolling System Update preferences

To update firmware on the Flex System compute nodes in the cluster, complete the following steps:

1. In UIM, click the cluster name. Then, click **Task Manager** and click **Create**, as shown in Figure 4-64.

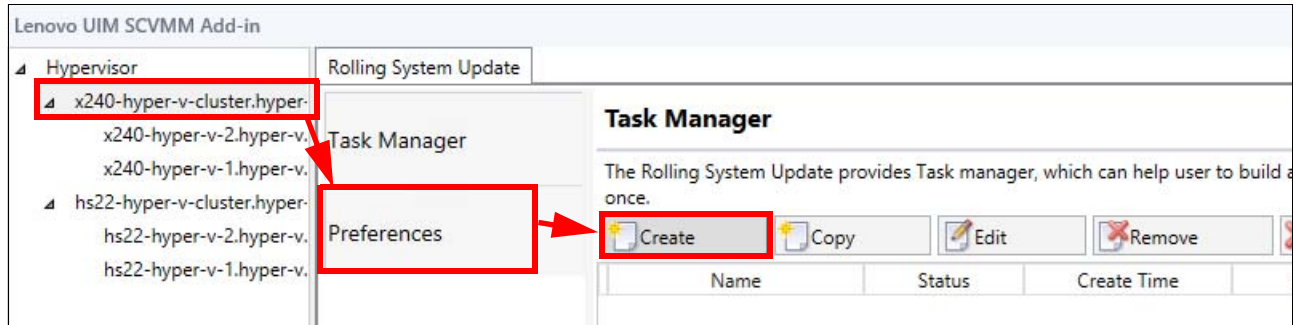


Figure 4-64 Create new task

2. Enter a task name. Select **Task type** as shown in Figure 4-65 and then click **Next**.

1. Name and Type

Task Name:

Task type: ☒ Update and Reboot ☐ Update Only ☐ Reboot Only

Figure 4-65 Task name and type

3. Select the wanted hosts to update and select the firmware and versions to update. You can specify per host or per host model, as shown in Figure 4-66.

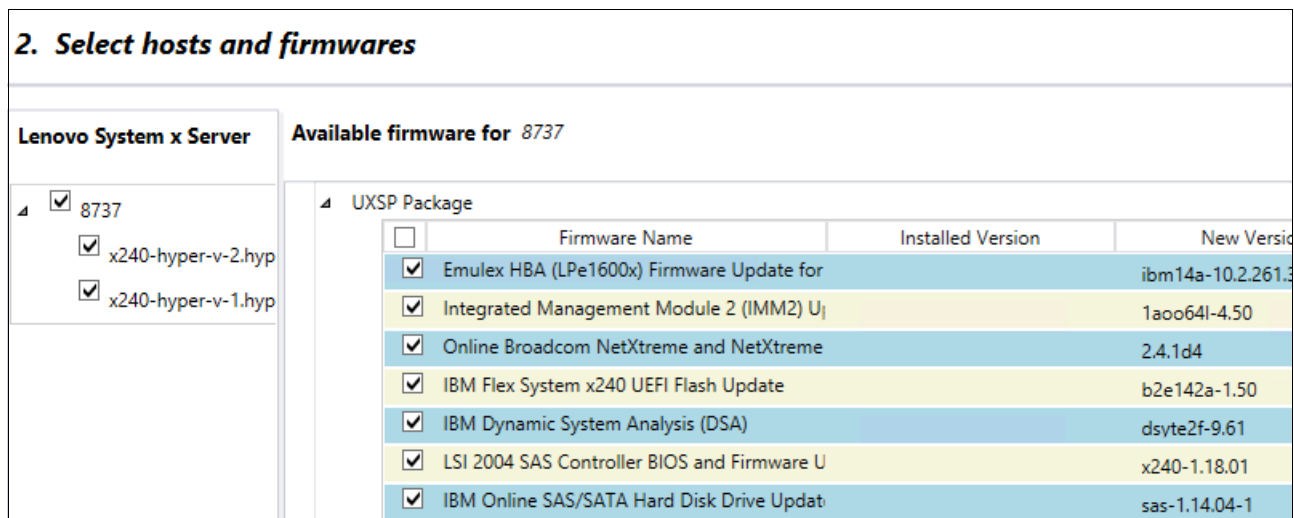


Figure 4-66 UIM for SCVMM: Select hosts and firmware

4. Define the wanted update options and schedule as shown in Figure 4-67 and click **Next**.

3. Update options and schedule

☐ Update Parallelization

Scale: Make sure the value is set according to the current available system resources of the cluster

☐ Force Downgrade

☒ Schedule

☒ Now

☐ Schedule Time

:

Figure 4-67 Update options and schedule

5. Review the Summary page as shown in Figure 4-68 and click **Save**.

4. Summary

You have made following selections:

Task name: Cluster Firmware Update

Task type: Update and reboot both

Update option:

Schedule: Now

Selected hosts and firmwares:

[x240-hyper-v-2.hyper-v.lenovopresslab.local:](#)

IBM Online SAS/SATA Hard Disk Drive Update Program

IBM Flex System x240 UEFI Flash Update

Online Broadcom NetXtreme and NetXtreme II Firmware Utility for Windows 2.4.1d4

Mellanox WinOF update for Windows 2012 R2 Server x86_64

Brocade BootCode Update for 16G FC HBA

Emulex OCE11xxx UCNA Firmware Update for Windows

Emulex HBA (LPe1600x) Firmware Update for Windows

IBM Flex System FC3172 2

IBM Flex System FC5172 2

Emulex HBA (LPe1205/LPe1200x) Firmware Update for Windows

IBM Dynamic System Analysis (DSA)

Figure 4-68 Review summary

You can monitor the progress and check the details of your task by clicking **Task** in Task Scheduler View window, as shown in Figure 4-69.

Task name: Cluster Firmware Update

Status: Running

Update Details:

Step 1: Download firmware

Status	Progress	Message	Start Time	End Time
Finished	100%	Download Completed	2/18/2015 11:03:02 AM	2/18/2015 11:04:20 AM

Step 2: Update progress

▷ x240-hyper-v-1.hyper-v.lenovopresslab.local Not Started
 ▲ x240-hyper-v-2.hyper-v.lenovopresslab.local Running Updating

Firmware Name	Installed Version	New Version	State	Message
IBM Online SAS/SATA Hard Disk Drive Updat	Undetected	sas-1.14.04-1	Not Start	The device is not
IBM Flex System x240 UEFI Flash Update	B2E142AUS-1.50	B2E142A-1.50	Not Start	The package ver
Emulex HBA (LPe1600x) Firmware Update for	Undetected	ibm14a-10.2.261.36-1	Not Start	The device is not
Firmware Update for ServeRAID M5115 PSoC	Undetected	m5115-68-1	Not Start	The device is not
Online Broadcom NetXtreme and NetXtreme	Undetected	2.4.1d4	Not Start	The device is not
IBM Flex System FC3172 2	Undetected	3.11af.d-8g-flex	Not Start	The device is not
Integrated Management Module 2 (IMM2) U	1A0058R-4.20	1A0064L-4.50	Running	Start Calling iFlas
Mellanox WinOF update for Windows 2012 R	Undetected	4.61.50000p4	Not Start	The device is not
IBM Flex System FC5172 2	Undetected	3.80.09-16g-flex	Not Start	The device is not
LSI 2004 SAS Controller BIOS and Firmware U		x240-1.18.01	Running	Package installat

Figure 4-69 UIM for SCVMM: Task Status details

When the cluster update task is completed successfully, it is reflected in the task status, as shown in Figure 4-70.

Task name: Cluster Firmware Update

Status: Finished

Update Details:

Step 1: Download firmware

Status	Progress	Message	Start Time	End Time
Finished	100%	Download Completed	2/18/2015 11:03:02 AM	2/18/2015 11:04:20 AM

Step 2: Update progress

▷ x240-hyper-v-2.hyper-v.lenovopresslab.local Finished Success
 ▷ x240-hyper-v-1.hyper-v.lenovopresslab.local Finished Success

Figure 4-70 Task completed successfully

You can perform the same firmware update actions for another clusters if required.

4.3.6 Publishing System Firmware to SCCM server

Lenovo System Updates for SCCM can be used if you want to update firmware on the Microsoft Windows servers that are not part of the cluster where usage of UIM Add-in for SCVMM is not possible.

Note: For more information, see the following Lenovo System Updates for Microsoft System Center Configuration Manager website:

<http://www-947.ibm.com/support/entry/portal/docdisplay?lnodocid=MIGR-5082209>

Complete the following steps:

1. Start Lenovo System Updates on SCCM server (you must run the program with administrator permissions).
2. In the My Machines view, click **Add**. To add a new machine, select the new Flex System type, as shown in Figure 4-71.

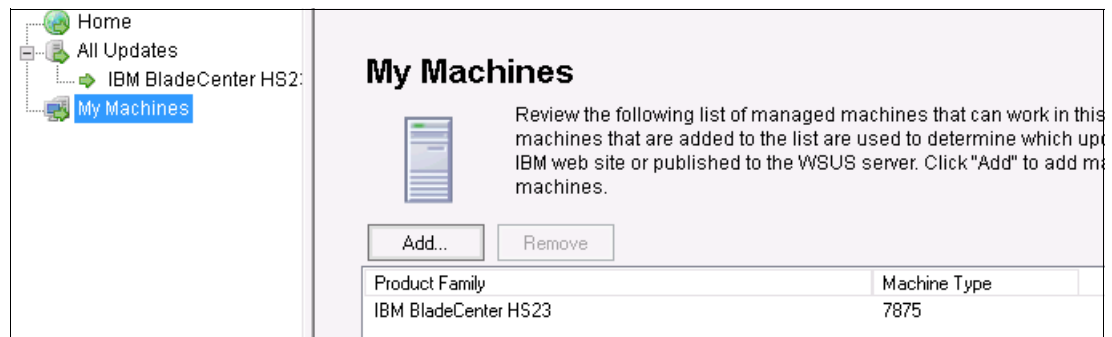


Figure 4-71 System Updates: My Machines view

3. In the Update retrieval action window, select **Check updates from IBM website now** and click **OK**, as shown in Figure 4-72.

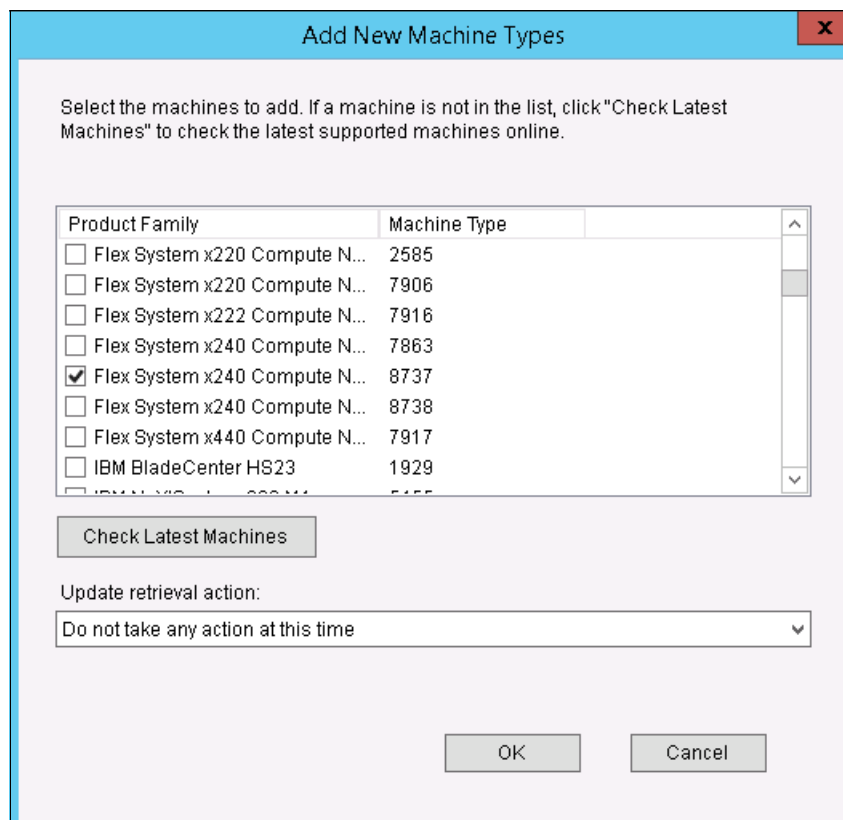


Figure 4-72 System Updates: Add new machine

- Expand All Updates in the Navigation pane and click the new machine type to show the available updates for this machine.
- Select the updates that you want to publish. Click **Actions** and then select **Download Selected updates from IBM website**, as shown in Figure 4-73. The Download and Publish wizard opens.

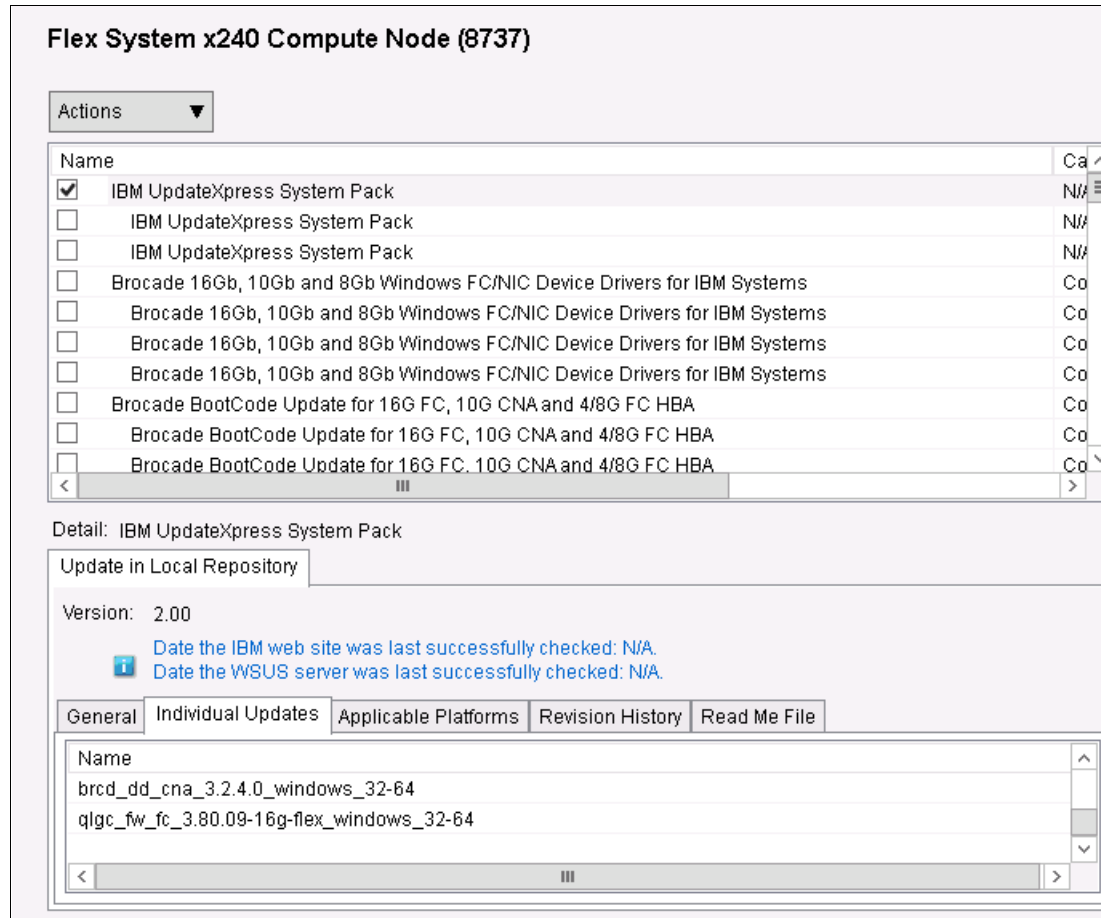


Figure 4-73 System Updates: Select updates

- Accept the license and click **Next**. The download starts. Click **Next** after the download completes to open the Publish wizard.

7. Confirm the updates that are to be deployed on the WSUS server. You also can select more update options that are based on preference, as shown in Figure 4-74. Click **Confirm**.

Confirm Updates Packages

Select and confirm that the following updates will be published to the WSUS server.

Confirm to publish the following updates

Name	Version	Size(KB)	Update ID
<input checked="" type="checkbox"/> IBM UpdateXpress System Pack	2.00	1182183	ibm_utl_uxsp_b2sp29p-2.0...

☒ Also select updates that are not installed or undetected
Select this option to detect all possible hardware in the system.
Select this option only if you have an unknown device in Windows that you want to resolve.

☒ Also select Host Bus Adapters(HBA)and Covered Network Adapter(CNA)
If this option is not selected, Brocade, Emulex, and QLogic HBAs and CNAs will not be updated.
Before you select this option, check the compatibility with your storage vendor.

☐ Allow updates to be installed as a downgrade

If you select this option, all published updates with the same update name will be expired.
Note: Only one version can be published at a time using the same update name.

< Back Confirm Finish Cancel

Figure 4-74 IBM System Updates: Confirm updates to publish

8. Publishing to the WSUS server begins. You can see the results after the process completes.

Note: You might need to wait some time for the updates to show in SCCM server. You also can run synchronization with WSUS server manually. For more information, see the SCCM documentation.

4.3.7 Inventory collection

After deploying Lenovo Inventory Tool Client onto a Flex System node's operating system, you can view inventory information in SCCM client.

Note: The inventory gathering cycle can vary based on your settings during the installation of the Lenovo Inventory Tool on SCCM server and SCCM client hardware inventory schedule. You can start Lenovo Inventory tool collection yourself by running the Inventory job that was created in the Windows Task Scheduler on client machine.

Complete the following steps to view inventory:

1. In the SCCM management console, open the Assets and Compliance view and find the computer on which you want to see its inventory.
2. Right-click **Computer name** → **Start** → **Resource Explorer**, as shown in Figure 4-75.

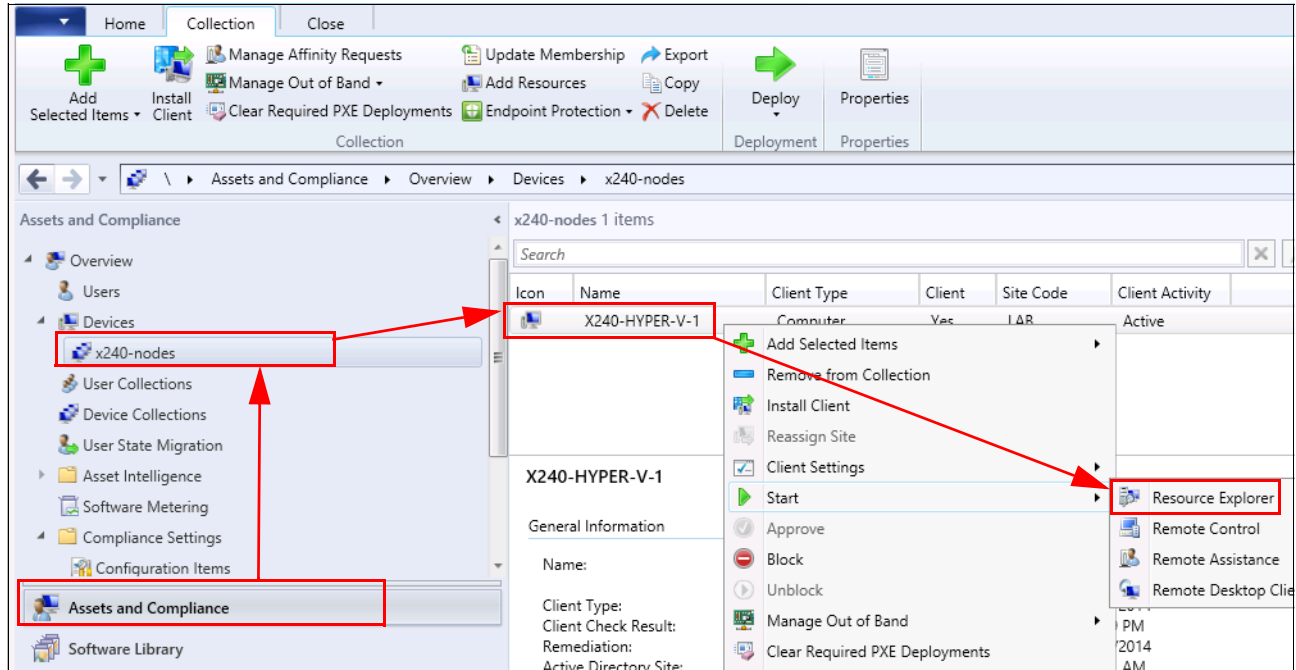


Figure 4-75 Inventory Collection: Resource explorer

You can see the inventory classes that were created in Resource Explorer, as shown in Figure 4-76.

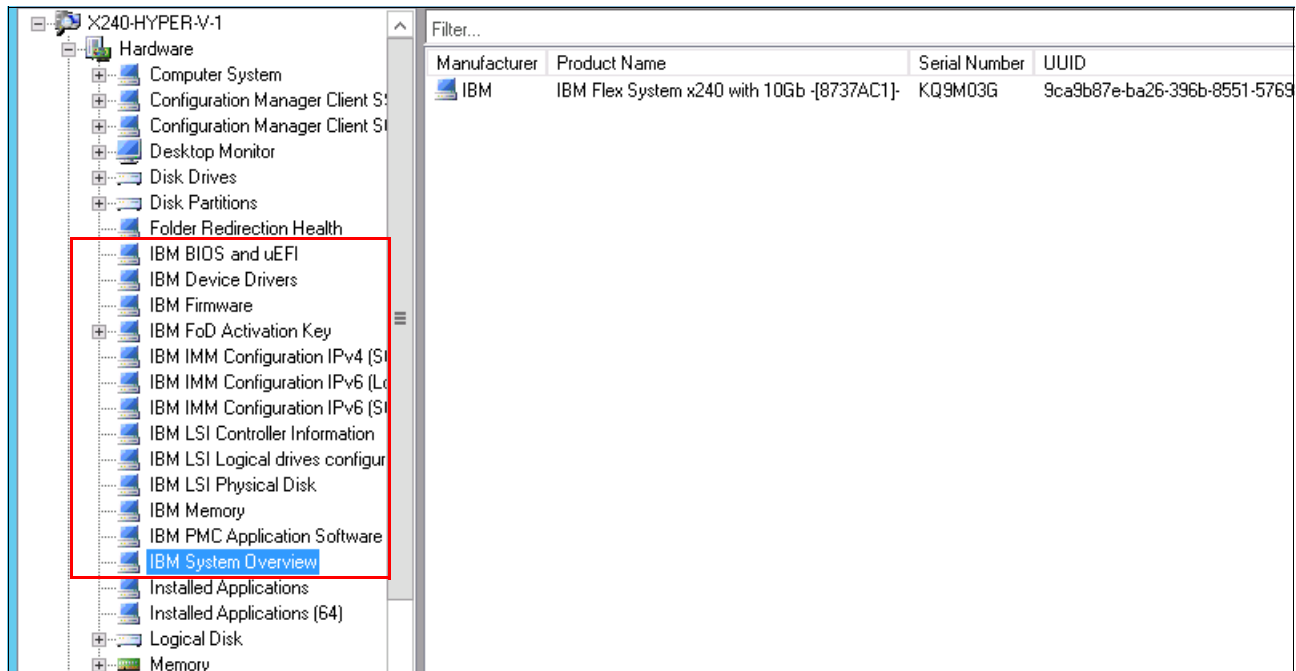


Figure 4-76 Resource Explorer: Inventory Classes

Abbreviations and acronyms

AMM	Advanced Management Module
CLI	command line interface
CMM	Chassis Management Module
CNA	Converged Network Adapter
DCUI	Direct Console User Interface
DRS	Distributed Resource Scheduler
EVC	Enhanced vMotion Compatibility
FC	Fibre Channel
FCF	Fibre Channel Forwarder
FCoE	Fibre Channel over Ethernet
FoD	Features on Demand
HA	High Availability
HBA	host bus adapter
IMM	Integrated Management Module
IQN	iSCSI qualified name
LOM	LAN on Motherboard
NPIV	N_Port ID Virtualization
PFA	Predictive Failure Alert
PRO	Performance and Resource Optimization
SCCM	System Center Configuration Manager
SCOM	System Center Operations Manager
SCVMM	System Center Virtual Machine Manager
SNMP	Simple network management protocol
SSIC	System Storage Interoperation Center
TCS	ToolsCenter Suite
UFP	Unified Fabric Port
UIM	Upward Integration Module
UXSP	UpdateXpress System Pack
VCI	VMware Certified Instructor
VLAN	virtual local area network
VM	virtual machine
VMM	Virtual Machine Manager
WAN	wide area network
isCLI	industry standard CLI

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 more information about the topic in this paper:

- ▶ *IBM Flex System Networking in an Enterprise Data Center*, REDP-4834
- ▶ *NIC Virtualization in Flex System Fabric Solutions*, SG24-8223
- ▶ *IBM PureFlex System and IBM Flex System Products and Technology*, SG24-7984
- ▶ *BladeCenter Products and Technology*, SG24-7523
- ▶ *Migrating from BladeCenter to Flex System*, REDP-4887

You can search for, view, or download these documents and other books, papers, and product guides at the following website:

<http://lenovopress.com>

Online resources

The following websites also are relevant as further information sources:

- ▶ Flex System Information Center:
<http://pic.dhe.ibm.com/infocenter/flexsys/information/index.jsp>
- ▶ BladeCenter Information Center:
<http://publib.boulder.ibm.com/infocenter/bladectr/documentation/index.jsp>
- ▶ System x Upwards Integration Modules (UIMs) for VMware vSphere:
<https://www-947.ibm.com/support/entry/myportal/docdisplay?ln docid=migr-vmware>
- ▶ System x Integration Offerings for Microsoft Systems Management Solutions:
<https://www-947.ibm.com/support/entry/myportal/docdisplay?ln docid=SYST-MANAGE>
- ▶ Switch Center 7.2.1 User Guide:
<http://www-01.ibm.com/support/docview.wss?uid=isg3T7000660>
- ▶ FastSetup:
<https://www-947.ibm.com/support/entry/myportal/docdisplay?ln docid=TOOL-FASTSET>

