Redbooks

ibm.com/redbooks

Implementing IBM Spectrum Virtualize software only

Peter Mescher

John Sexton

Megan Gilge

Cloud

Storage

IBM

Redpaper

International Technical Support Organization

Implementing IBM Spectrum Virtualize software only

May 2017

REDP-5392-01

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

Second Edition (May 2017)

This edition applies to Version 7, Release 8, Modification 1 of IBM Spectrum Virtualize.
This document was created or updated on May 19, 2017.

Contents

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

INTERNATIONAL BUSINESS MACHINES CORPORATION 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. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

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

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

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

DB2®
Easy Tier®
FlashCopy®
HyperSwap®
IBM®
IBM FlashSystem®
IBM Resiliency Services™
IBM Spectrum™
IBM Spectrum Control™
IBM Spectrum Protect™
IBM Spectrum Storage™
IBM Spectrum Virtualize™
Insight™
MicroLatency®
Passport Advantage®
Power Systems™
Real-time Compression™
Redbooks®
Redpaper™
Redbooks (logo)  ®
Storwize®
System i®
System Storage®
XIV®

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, 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

This IBM® Redpaper™ publication provides a broad understanding of IBM Spectrum Virtualize™ software only and how it fits into to the IBM SAN Volume Control and IBM Storwize® families. It also provides use cases for cloud, Cloud Service Provider (CSP), and Managed Service Provider (MSP) implementations.

This publication helps storage and networking administrators install, tailor, and configure IBM Spectrum Virtualize software only. It also provides a detailed description of supported hardware and troubleshooting tips.

In April 2017, this paper was updated to include information about Version 7.8.1 and Supermicro SYS-2028U-TRTP+ server implementation.

Authors

This paper was produced by a team of specialists from around the world working at IBM Hursley, United Kingdom.

Peter Mescher is a Senior IT Architect for IBM Resiliency Services™, where he designs dedicated compute and storage resiliency solutions. He has 16 years of experience with IBM Storage Area Networking (SAN) products, and has worked with the IBM Spectrum Virtualize family since its inception in 2003. He has special areas of expertise with the Fibre Channel (FC) protocols and with IBM remote replication technologies, especially Global Mirror. He has a bachelors degree in Computer Engineering from Lehigh University in Bethlehem, PA.

John Sexton is a Certified Consulting IT Specialist, based in Wellington, New Zealand, and has over 25 years of experience working in IT. He has worked at IBM for the last 20 years. His areas of expertise include IBM Power Systems™, virtualization, SAN, cloud, IBM Spectrum Storage™, and business continuity. He provides pre-sales support and technical services for clients throughout New Zealand, including consulting, architecture solution design, performance monitoring, system migration, and training. Before joining IBM in New Zealand, John worked in the United Kingdom building and maintaining systems in the UK financial and advertising industries.

Megan Gilge is a Project Leader in IBM Technical Content Services. Before joining the Redbooks organization four years ago, she was an Information Developer in the IBM Semiconductor Solutions and IBM System i® areas.

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

Jon Tate
International Technical Support Organization

Robin Findlay
Joshua Vines
Stephen Wright
IBM Hursley

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

- Use the online **Contact us** review Redbooks form found at:

  ibm.com/redbooks
- Send your comments in an email to:

  redbooks@us.ibm.com
- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks
- Follow us on Twitter:

  http://twitter.com/ibmredbooks
- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806
- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm
- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

1

Introduction

This chapter describes IBM Spectrum Virtualize and compares IBM Spectrum Virtualize software only to other products in the IBM Spectrum Virtualize family. It also describes how IBM Spectrum Virtualize software only can be used with both existing IBM Storwize environments and additional cloud infrastructures, including integration with the OpenStack platform.

IBM Spectrum Virtualize is available as software only starting with IBM Spectrum Virtualize V7.7.1. This publication describes version 7.8.1.

This chapter provides the following information:

1.1 Introduction to IBM Spectrum Virtualize

Designed for software-defined storage (SDS) environments, IBM Spectrum Virtualize software only is part of the IBM Storwize family and includes technologies that both complement and enhance virtual environments. It has built-in functions such as IBM Real-time Compression™ and IBM Easy Tier® technology that deliver extraordinary levels of efficiency. Available in a wide range of storage systems, the Storwize family delivers sophisticated capabilities that are easy to deploy, yet also help control costs for growing businesses.

The term *virtualization* is used widely in IT and applied to many of the associated technologies. Its usage in storage products and solutions is no exception. IBM defines storage virtualization in the following manner:

- Storage virtualization is a technology that makes one set of resources resemble another set of resources, preferably with more desirable characteristics.
- It is a logical representation of resources that is not constrained by physical limitations and hides part of the complexity. It also adds or integrates new functions with existing services and can be nested or applied to multiple layers of a system.

IBM Spectrum Virtualize is at the heart of IBM SAN Volume Controller, the IBM Storwize family, IBM FlashSystem® V9000, and VersaStack. It provides virtualization at the disk layer (block-based) of the I/O stack only. It is different than that of virtualization at the file system layer (file-based). IBM Spectrum Virtualize is referred to as block-level virtualization, or the block aggregation layer.

IBM Spectrum Virtualize is designed to address the following requirements:

- Virtualize block storage arrays with Fibre Channel (FC) and ISCSI front-end ports to centralize management of storage volumes for management from a single point.
- Improve utilization of storage capacity so that businesses can tap into previously unused disk capacity.
- Use breakthrough ease of use and management.
- Avoid downtime for backups, maintenance, and upgrades.
- Support data migration without disruption to applications.
- Enable virtualized storage volumes to be managed as a single pool of storage.
- Enable advanced functions to be used across multiple virtualized storage arrays, even if the virtualized storage array does not have these advanced functionalities natively. These functions include Real-time Compression (in-line, real-time compression), stretched cluster and IBM HyperSwap® (high-availability solution), IBM Easy Tier (automatic and dynamic data tiering), encryption of external virtualized capacities, IBM FlashCopy® (snapshot), and remote data replication.
- Enable a high-availability storage infrastructure through a stretched cluster or HyperSwap.
- Improve the efficiency of the virtualized block storage infrastructure with advanced technologies such as Real-time Compression for up to 80% data compression rates without performance impact, and high-performance thin provisioning with zero detection.
- Move the most accessed data to higher performance tiers automatically and dynamically through Easy Tier.

1.1.1 Features of IBM Spectrum Virtualize

Virtualizing storage with IBM Spectrum Virtualize helps make new and existing heterogeneous storage arrays more effective by including many functions that are traditionally deployed within disk array systems. By including these functions in a virtualization system, IBM Spectrum Virtualize software only standardizes functions across virtualized storage for greater flexibility and potentially lower costs. Table 1-1 describes IBM Spectrum Virtualize features and benefits.

Table 1-1 IBM Spectrum Virtualize features and benefits

Feature	Benefits
Single point of control for storage resources	► Designed to increase management efficiency ► Designed to help support business application availability
Pools the storage capacity of multiple storage systems on a SAN	► Helps you manage storage as a resource to meet business requirements, and not just as a set of boxes ► Helps administrators better deploy storage as required beyond traditional “SAN islands” ► Can help increase usage of storage assets ► Insulates applications from physical changes to the storage infrastructure
Clustered pairs of Intel servers that are configured as IBM Spectrum Virtualize data engines	► Use of generic Intel hardware foundation ► Designed to avoid single points of hardware failure
Real-time Compression	► Increases effective capacity of storage systems up to five times, helping to lower costs, and floor-space, power, and cooling requirements ► Can be used with a wide range of data, including active primary data, for dramatic savings ► Hardware compression acceleration helps transform the economics of data storage
Innovative and tightly integrated support for flash memory	► Designed to deliver ultra-high performance capability for critical application data ► Move data to and from flash memory without disruption; make copies of data onto hard disk drives (HDDs)
Support for IBM FlashSystem	Enables high performance for critical applications with IBM MicroLatency®, coupled with sophisticated functions
Easy-to-use IBM Storwize family management interface	► Single interface for storage configuration, management, and service tasks regardless of storage vendor ► Helps administrators use their existing storage assets more efficiently
IBM Storage Mobile Dashboard	Provides basic monitoring capabilities to securely check system health and performance
Dynamic data migration	► Migrate data among devices without taking applications that use that data offline ► Manage and scale storage capacity without disrupting applications
Manage tiered storage	Helps balance performance needs against infrastructure costs in a tiered storage environment

Feature	Benefits
Advanced network-based copy services	► Copy data across multiple storage systems with IBM FlashCopy ► Copy data across metropolitan and global distances as needed to create high-availability storage solutions
Integrated Bridgeworks SANrockIT technology for IP replication	► Optimize use of network bandwidth ► Reduce network costs or speed replication cycles, improving the accuracy of remote data
Enhanced stretch cluster configurations	► Provide highly available, concurrent access to a single copy of data from data centers up to 300 km apart ► Enable nondisruptive storage and virtual machine (VM) mobility between data centers
Thin provisioning and snapshot replication	► Dramatically reduce physical storage requirements by using physical storage only when data changes ► Improve storage administrator productivity through automated on-demand storage provisioning
IBM Spectrum Protect™ Snapshot application-aware snapshots	► Performs near-instant application-aware snapshot backups, with minimal performance impact for IBM DB2®, Oracle, SAP, VMware, Microsoft SQL Server, and Microsoft Exchange ► Provides advanced, granular restoration of Microsoft Exchange data

1.1.2 OpenStack and IBM Spectrum Virtualize

IBM Spectrum Virtualize integrates with widely supported cloud frameworks. OpenStack is the open source software for building and managing public, private, and hybrid clouds. The OpenStack environment is a cloud framework that controls large pools of compute, storage, and networking resources throughout a data center. The resources are managed through a dashboard that gives administrators control while empowering their users to provision resources through a web interface.

OpenStack technology is a key enabler of cloud infrastructure as a service (IaaS) capability. The OpenStack architecture provides an overall cloud preferred practices workflow solution that is readily installable, and supported by a large infrastructure of worldwide developers in the OpenStack open source community. These solutions have many benefits, including reduced operational cost and increased operational efficiency, the ability to deploy applications more quickly and with standardized APIs, and the ability to deploy these solutions by using the compute, storage, and networking infrastructure that is best suited for the workload.

The OpenStack platform controls large pools of compute, storage, and networking resources throughout a data center. The resources are managed through a dashboard that gives administrators control while enabling their users to provision resources through a web interface.

For more information, see this website:

https://ibm.biz/Bd4Lnw

The OpenStack platform consists of a set of loosely coupled services, each delivering a specific set of functions through RESTful APIs. The core of OpenStack is made up of the following essential services:

- Nova: Management of Compute Resources
- Keystone: Identity Management
- Glance: Image Storage
- Cinder: Virtual Block Storage
- Swift: Object Storage
- Neutron: Management of Network Resources

There are many other optional services that can be configured in an OpenStack environment, including mature projects such as the Manila shared file system service, and newer projects such as Magnum for container management. For more information, see this website:

http://www.openstack.org/software/project-navigator

The OpenStack architecture is one implementation of a preferred practices cloud workflow. Regardless of the cloud operating system environment that is used, the following key summary points apply:

- Cloud platforms provide the necessary technology workflow to provide truly elastic, pay-per-use cloud services.
- OpenStack cloud software provides a vibrant open source cloud operating system that is growing quickly.

OpenStack storage integration

The OpenStack platform includes the following storage components, which support IBM storage products:

- The volume management driver for IBM Spectrum Virtualize and the IBM Storwize family provides OpenStack Compute instances with access to IBM Spectrum Virtualize and IBM Storwize family storage systems.

  These families support fully transparent live storage migration in OpenStack Havana and later releases:

  - No interaction with the host is required: All advanced Storwize features are supported and exposed to the Cinder system.
  - Real-time Compression with EasyTier supports iSCSI + FC attachment.
- IBM FlashSystem (Kilo release): The volume driver for IBM FlashSystem provides OpenStack Block Storage hosts with access to IBM FlashSystem storage.

Within the overall cloud workflow, the following OpenStack components support storage:

- Cinder (block storage)
- Swift (object storage)
- Manila (file storage)

IBM Spectrum Virtualize integration and OpenStack Cinder

Cinder is an OpenStack project to provide “block storage as a service” and provides an API to users to interact with different storage back-end solutions. The Cinder component provides support, provisioning, and control of block storage. The following framework is standard across all drivers for Cinder services to properly interact with a driver. IBM Spectrum Virtualize integrates with the OpenStack environment through Cinder.

Figure 1-1 shows where Cinder is placed in OpenStack and shows how IBM Spectrum Virtualize can integrate with the framework.

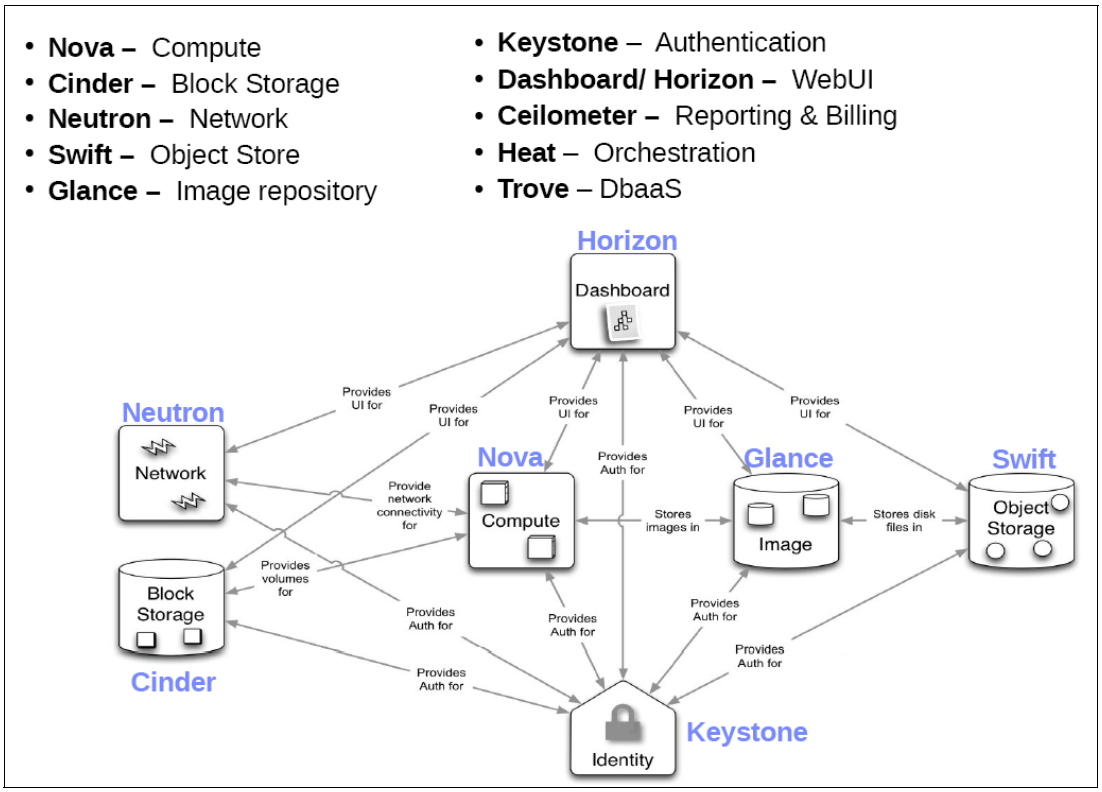

Figure 1-1 OpenStack storage components in the framework

IBM Spectrum Virtualize and the IBM Storwize family support these items:

- Folsom and later releases of OpenStack
- iSCSI and FC
- Advanced IBM Spectrum Virtualize features, such as Real-time Compression and Easy Tier
- Software-defined placement with the OpenStack filter scheduler
- Storage-assisted volume migration

OpenStack Havana and later releases include a new administrator feature for migrating volumes between Cinder instances. Volumes can be migrated with Host Assisted Data Migration or by Storage Assisted Data Migration with the IBM Storwize family.

Figure 1-2 on page 7 lists the common use cases for migrating volumes in the OpenStack environment. The function of IBM Spectrum Virtualize and the Storwize family is integrated with the OpenStack features.

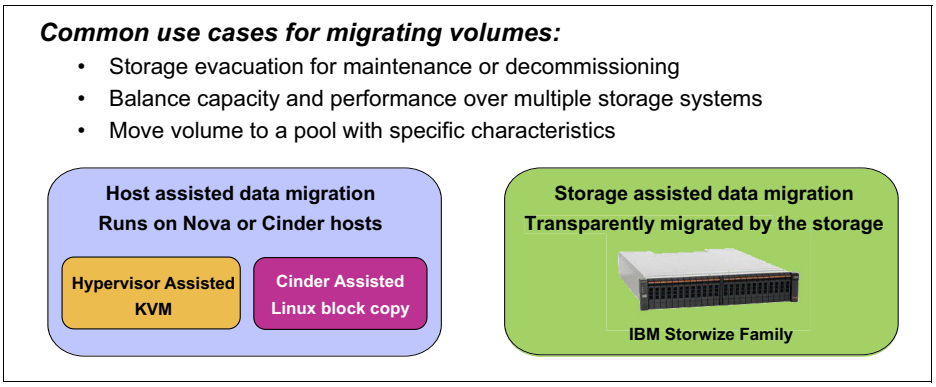

Figure 1-2 Common use cases for volume migration in OpenStack environment

IBM Spectrum Virtualize is the only storage in the Havana release to support storage assisted migration. Volumes move between two storage pools that are managed by a Storwize family system.

Here are the key benefits to using the IBM Spectrum Virtualize storage assisted migration:

- No interaction with the host
- No impact on VM and node
- Instantaneous
- No effect on VM operations or volume management

The volume management driver for IBM Spectrum Virtualize provides OpenStack Compute instances with access to IBM Spectrum Virtualize managed storage systems.

IBM Spectrum Virtualize supports fully transparent live storage migration in OpenStack (Havana and later). No interaction with the host is required: All advanced IBM Spectrum Virtualize features are supported and exposed to the Cinder system.

Real-time Compression with Easy Tier supports iSCSI and FC attachment.

1.2 IBM Spectrum Virtualize software only

IBM Spectrum Virtualize software only is an SDS implementation that provides all the capabilities and functions of the IBM SAN Volume Controller. It runs on supported Intel hardware that you supply.

This SDS layer is designed to virtualize and optimize storage within the data center or managed private cloud service. This publication addresses the features and functions of IBM Spectrum Virtualize software only.

IBM Spectrum Virtualize software only is installed on supported bare metal Intel servers. It is a self-contained binary file with the same features, options, GUIs, and the same CLI support. It is an attractive offering that is targeted for customers who are Managed Service Providers (MSPs), and private Cloud Service Providers (CSPs) that require more flexibility in hybrid cloud deployment. IBM Spectrum Virtualize software only is the same as IBM SAN Volume Controller, which is designed to manage storage and supports nearly 400 brands and models, including all major storage vendors.

Whether in an on-premises private or managed cloud service, this offering reduces the complexities and cost of managing SAN FC- or iSCSI-based storage while improving availability and enhancing performance.

Figure 1-3 shows the logical placement of IBM Spectrum Virtualize software only with a supported hardware configuration that is required to establish the functions that IBM Spectrum Virtualize provides.

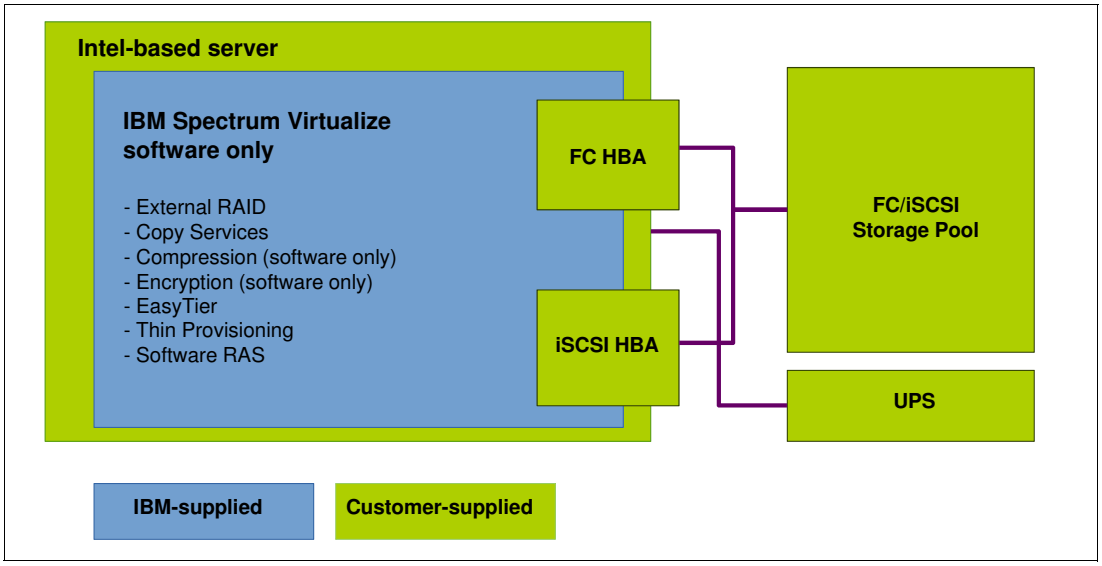

Figure 1-3 IBM Spectrum Virtualize software only with customer-supplied hardware

Note: As shown in Figure 1-3, the uninterruptible power supply (UPS) is customer-supplied. All planning for and provision of the power and environmental support that is required for IBM Spectrum Virtualize software only, including all storage engines, are also a customer responsibility.

IBM Spectrum Virtualize software only has many of the same capabilities as IBM SAN Volume Controller (see Figure 1-3). The main difference is that you provide the hardware.

1.2.1 Benefits and features of Spectrum Virtualize software only

IBM Spectrum Virtualize software only offers a powerful value proposition for enterprise IT and service providers who are searching for more flexible and agile ways to deploy block storage on commodity Intel hardware. Using these standard Intel servers, IBM Spectrum Virtualize software only can be easily added to existing infrastructures to deliver greater storage efficiency through virtualization.

This software only version of the established IBM Storwize family provides a compelling solution to how SDS can be implemented in numerous types of solutions for storage environments. IBM Spectrum Virtualize provides the benefits of storage virtualization and advanced storage capabilities, including:

- EasyTier
- Real-time Compression
- The Compretimator tool
- FlashCopy snapshots
- Remote data replication

- HyperSwap
- Encryption
- High availability
- iSCSI Virtualization
- VMware Virtual Volumes (VVOL) support

1.2.2 Benefits of IBM Spectrum Virtualize in a software-defined infrastructure

IBM Spectrum Virtualize provides the following benefits:

- Facilitates IT automation to improve business and IT agility at lower cost
- Optimized systems administration and control to allow effective and efficient resource utilization that lowers cost and supports business requirements
- Speed of deployment
- Performance tuning and storage tiering
- Capacity planning
- Enables advanced application deployment by employing Systems of Engagement and Systems of Insight™ that use Systems of Record:
  - Cloud
  - Analytics
  - Mobile
  - Social
  - Security
- Simplified architecture to reduce specialized components and skills requirements
- Users who are interested in a hybrid cloud environment targeting a lower cost, yet highly available and robust environment for their secondary site
- Service Providers (MSPs or CSPs) who might offer hybrid cloud services to their users

From an architectural perspective where SDS is an integral component of an overall software-defined infrastructure (SDI) framework, other components are software-defined compute and software-defined network. Figure 1-4 shows where SDS aligns with the overall software-defined framework environment contributing to the benefits that it provides.

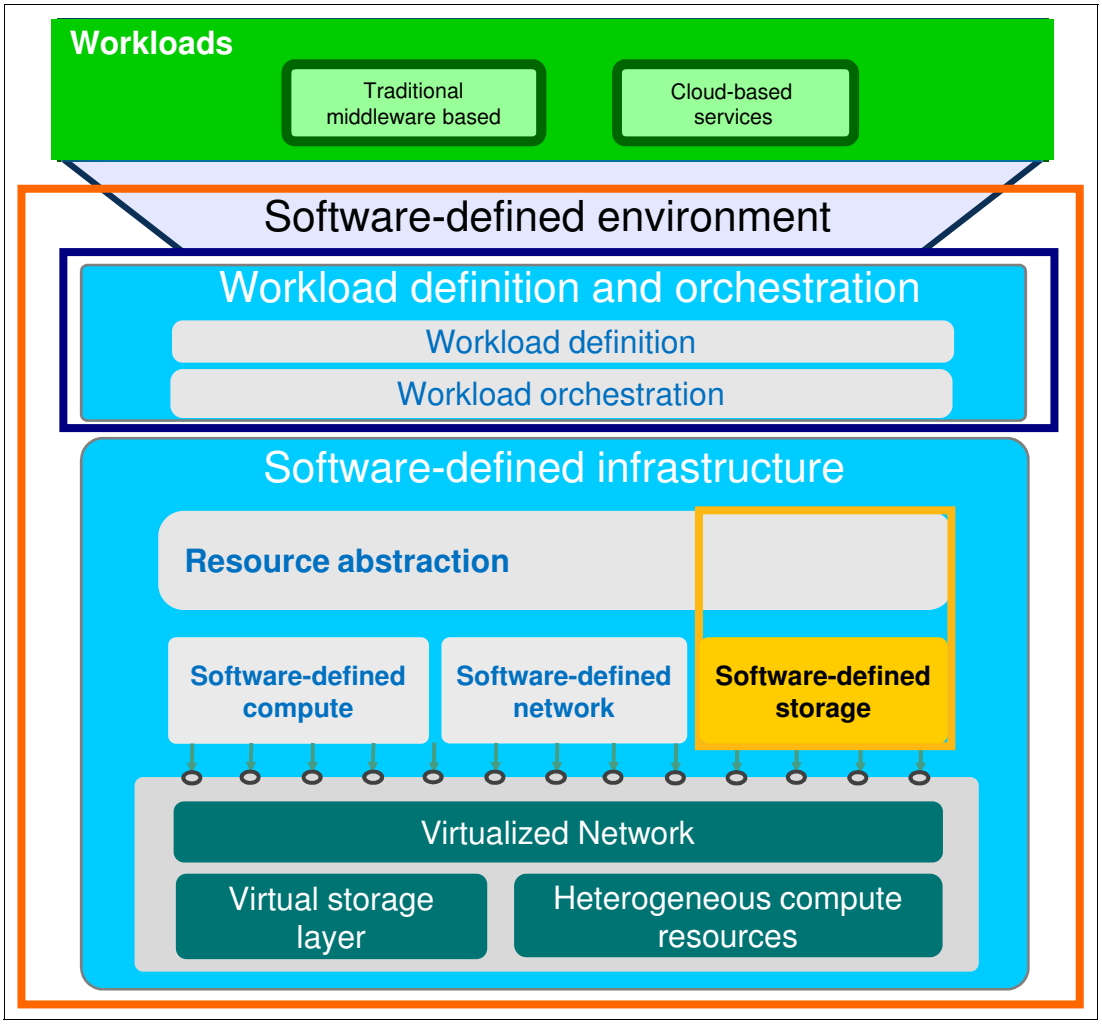

Figure 1-4 Software-defined infrastructure framework

The SDI framework drives efficiency by optimizing the connections between workloads and resources based on the needs of the business. As shown in Figure 1-4, SDS contributes in providing some of the major attributes of SDI:

- Agility

  IT resource customers expect to use infrastructure resources on demand based on immediate business requirements.

- Standardization

  Ability to create uniformity by automating, standardizing, and integrating IT infrastructures.

- Provisioning and Orchestration

  Ability to configure pools of resources and put them together in a way that can be dynamically delivered (by using software) with service-level-oriented interfaces appropriate to IT consumers.

In this software-defined framework, workloads are defined and orchestrated based on usage patterns, and resources are managed and deployed according to business rules and policies. IT enterprises employing an SDI approach have several core advantages over processes that have traditionally been handled manually.

1.3 Use cases for IBM Spectrum Virtualize software only

This section describes common use cases where IBM Spectrum Virtualize software only can be implemented to improve existing services, functions, and features for an organization. That organization might be a CSP or MSP where their requirements might differ. The following use cases that address the differing requirements are all based on the added functions of IBM Spectrum Virtualize:

- Optimizing CSP or MSP data centers with IBM Spectrum Virtualize by adding optimization and functions for heterogeneous storage
- Offering collocation solutions for IBM SAN Volume Controller customers where bare metal Intel servers are used instead of IBM SAN Volume Controller appliances at the second site
- Extending Storage as a Service capability for the service provider

1.3.1 Optimizing cloud data centers

IBM Spectrum Virtualize software can be used only with existing CSPs or MSPs typically managing iSCSI storage environments. By using installed and supported Intel hardware, IBM Spectrum Virtualize software only can be implemented at a lower cost and reduced effort. This enhances the CSP or MSP infrastructure with added functions that IBM Spectrum Virtualize offers plus fulfilling an SDS strategy.

Figure 1-5 shows the cloud data center storage that is managed by IBM Spectrum Virtualize software functions where advantages for cloud offerings include the following items:

- Use the existing infrastructure and improve the environment with software
- Improve security (encryption of data at rest)
- Pay as you grow with monthly licensing
- Optimize use of heterogeneous storage

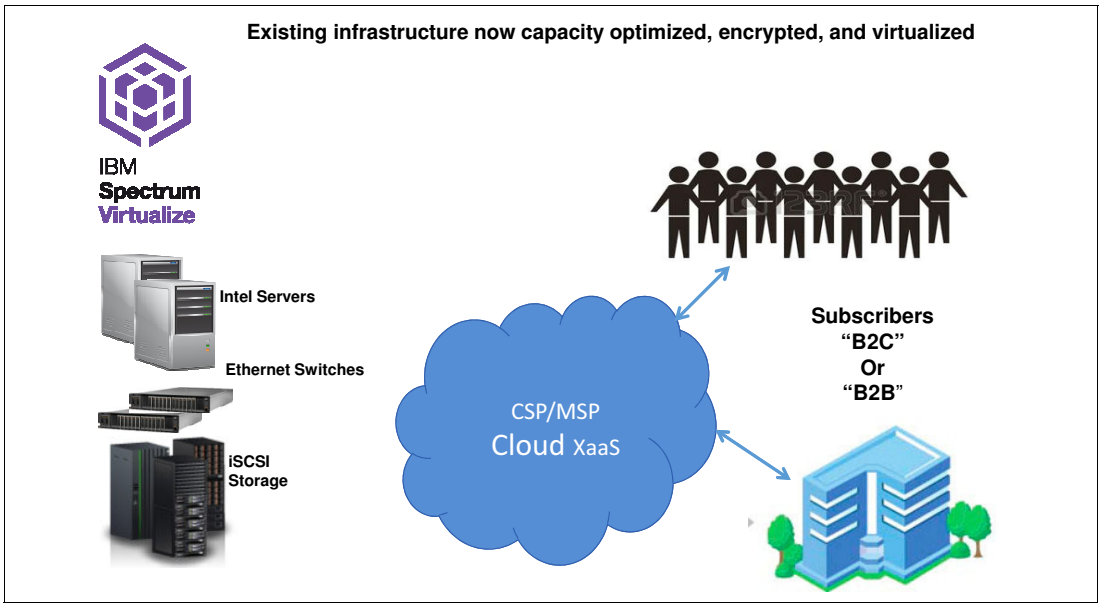

Figure 1-5 Optimize existing cloud data centers

Note: In Figure 1-5, XaaS stands for Everything as a Service.

The CSP or MSP can offer additional functions with IBM Spectrum Virtualize for users or subscribers. For example, an offering for subscribers with security concerns is to encrypt data on the various heterogeneous back-end storage.

1.3.2 Optimizing collocation cloud data centers

Some MSPs and even public clouds offer collocation options for their clients. In the case of IBM Spectrum Virtualize, existing or new IBM SAN Volume Controller or Storwize clients can benefit from a single management layer that can use functions such as replication, backup, and storage tiers.

Figure 1-6 on page 13 shows a client with their own IBM SAN Volume Controller managed storage infrastructure with a collocation of clients equipment that is housed at an MSP site.

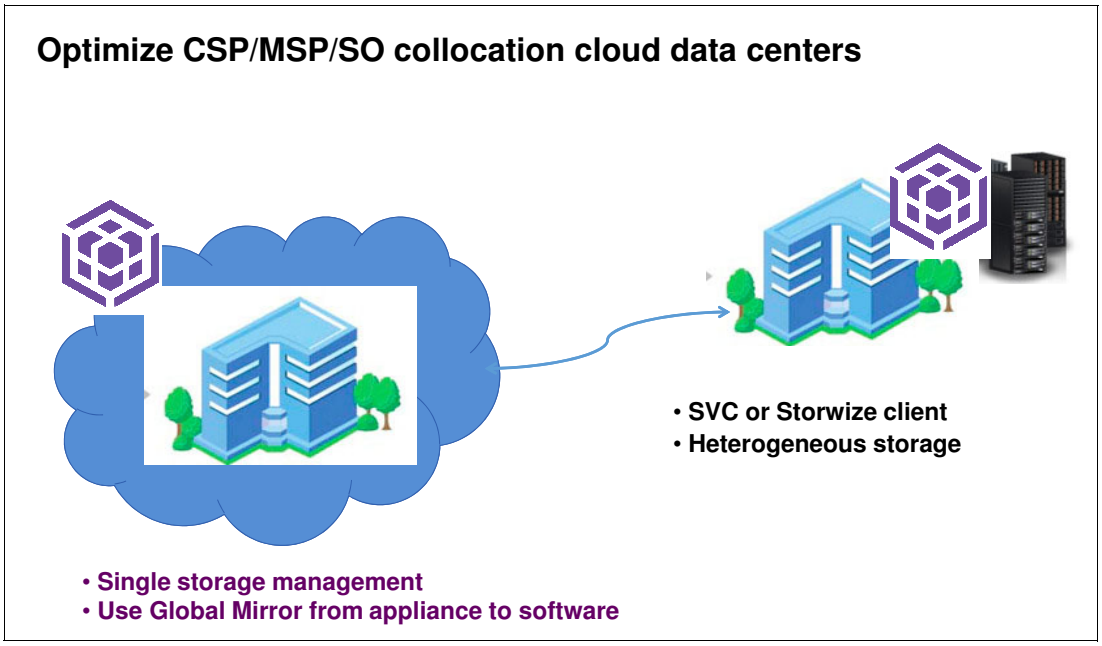

Figure 1-6 Optimize collocation data centers

With IBM Spectrum Virtualize software only managing storage at the MSP on supported Intel servers, a single management layer for storage is possible across both data centers that are managed by IBM Storwize or IBM SAN Volume Controller appliances, which provides full IBM Spectrum Virtualize functions between the locations. For example, by using storage replication that uses Global Mirror, you can set up and control storage replication within the single management layer.

1.3.3 Implementing Storage as a Service

For MSPs and CSPs that offer Storage as a Service or disaster recovery as a service (DRaaS), adding IBM Spectrum Virtualize to their environment provides more benefits. Similar to 1.3.1, “Optimizing cloud data centers” on page 11, it can greatly improve the efficiency of their existing storage services. With iSCSI support, the solution becomes more viable for service providers.

Figure 1-7 shows an MSP or CSP providing additional storage-related *as-a-service* offerings at their site, which allows them to fully manage the clients storage services by using their own supported Intel servers.

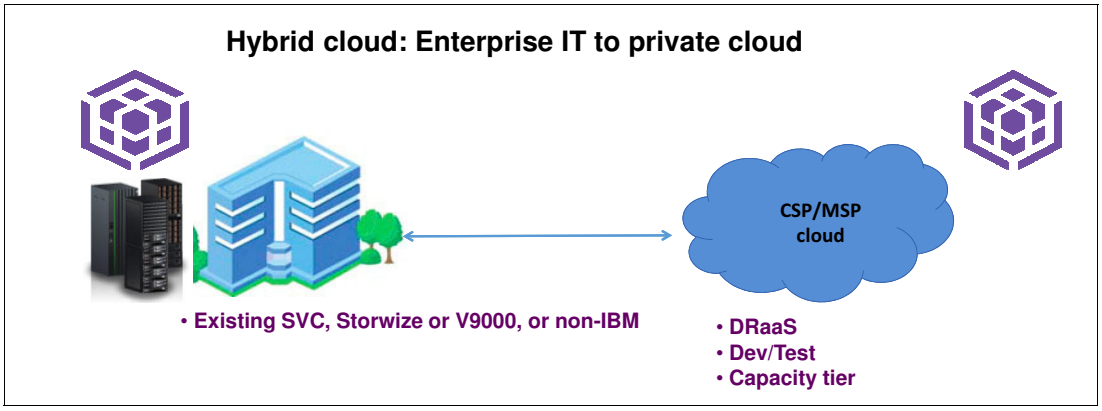

Figure 1-7 Storage as a Service

For example, in Figure 1-7 the client has an existing IBM SAN Volume Controller or Storwize infrastructure and can use MSP for an alternative data center without having to purchase hardware. The MSP uses bare metal Intel servers at their data center with IBM Spectrum Virtualize software only. This fully managed service allows the MSP to provide DRaaS, capacity and tier management, plus a test and development environment across both data centers.

1.4 Licensing

IBM Spectrum Virtualize software only is offered as a fully inclusive, perpetual, or monthly license. You do not need to purchase an optional license for Copy Services, Encryption, Real-time Compression, FlashCopy, and Remote Mirror. The option for monthly licensing delivers increased flexibility over capital versus operational expense budgets, varying storage needs, and software maintenance requirements.

The license can be used across multiple data centers and IBM Spectrum Protect clusters, but is limited per organization and within a country. If an organization operates in multiple countries, a license is required for each country. Orders are placed through the IBM Passport Advantage® system in each country where it is required.

Licensing for the software itself is per storage capacity unit (SCU). You order the number of SCUs that are needed to cover your storage pool. Space-Efficient and Compressed volumes are based on the amount back-end storage pool they use, *not* the total capacity they present as volumes; this allows for more predictable planning and purchasing.

There is no entitlement or conversion for IBM Spectrum Virtualize licenses on the IBM SAN Volume Controller. A new license for IBM Spectrum Virtualize software only must be purchased.

1.4.1 Storage Capacity Units

SCUs are defined in terms of the category of the storage capacity:

- Category 1: Flash and solid-state drives (SSDs)
- Category 2: Serial-attached SCSI (SAS) drives, FC drives, and systems that use Category 3 drives with advanced architectures to deliver high-end storage performance
- Category 3: Near Line SAS (NL-SAS) and Serial ATA (SATA) drives

Any storage use case that is not listed here is classified as Category 1.

For each SCU, the following number of tebibytes (TiB[1]) by storage classification is applicable:

- 1 SCU equates to 1.00 TiB usable of Category 1
- 1 SCU equates to 1.18 TiB usable of Category 2
- 1 SCU equates to 4.00 TiB usable of Category 3

Licensing is purchased based on the total capacity of your storage pools. When calculating the count of SCUs per category, fractions must be rounded up to the next higher integer number.

For the IBM Spectrum Virtualize license, a sufficient number of SCUs is required to cover actual managed disk capacity that is consumed by the compressed volumes.

1.4.2 License examples

The following examples show how the SCU value is calculated. This calculation is the same as for the IBM SAN Volume Controller.

Example 1

A user with 42 TiB usable storage in a FlashSystem 9000, 400 TiB usable storage in SAS drives, and 800 TiB usable storage in NL-SAS drives. Table 1-2 shows how to calculate for IBM Spectrum Virtualize license SCUs for a hybrid storage array.

Table 1-2 Example 1 SCU requirement calculation1

SCU category	TiB usable	TiB usable per SCU	SCU
Category 1	42	1	42
Category 2	400	1.18	339
Category 3	800	4	200
Total	**1242**		**581**

[1] A tebibyte (TiB) is 2 to the 40th power (1,099,511,627,776) bytes. It is not the same as a terabyte definition, which is 10 to the 12th power (1,000,000,000,000) bytes.

Example 2

A user with 10 TiB usable storage in a FlashSystem 9000, 20 TiB usable storage in SSDs, 2400 TiB usable storage in NL-SAS drives in a Hybrid Storage array, and 200 TiB usable storage in SAS drives in a second disk storage array. Table 1-3 shows how you would calculate for IBM Spectrum Virtualize SCUs.

Table 1-3 Example 2 SCU requirement calculation

SCU category	TiB usable	TiB usable per SCU	SCU
Category 1	30	1	42
Category 2	200	1.18	170
Category 3	2400	4	600
Total	**2630**		**800**

Example 3

A user has 906 TiB usable storage in an IBM XIV®, and 300 TiB usable storage in NL-SAS drives in a second disk storage array. As in the Differential Licensing Category definition, the XIV use case is Category 2. Calculations for IBM Spectrum Virtualize SCUs are shown in Table 1-4.

Table 1-4 Example 3 SCU requirement calculation

SCU category	TiB usable	TiB Usable per SCU	SCU
Category 1	0	1	0
Category 2	906	1.18	768
Category 3	300	4	75
Total	**1206**		**843**

IBM Spectrum Virtualize software only (5725-Z61) has the same software prerequisites as IBM SAN Volume Controller. IBM includes one year of software maintenance with each program license that is acquired. The initial period of software maintenance can be extended by the purchase of a renewal option when ordered through Passport Advantage or Passport Advantage Express.

Important: IBM Spectrum Virtualize software can be ordered only through Passport Advantage. All software maintenance renewals are managed and invoiced through Passport Advantage.

1.4.3 Ordering information

IBM Spectrum Virtualize Software Version 7.7.1 is required for Lenovo System x3650 M5 server. Version 7.8.0.2 or later is required for the Supermicro SYS-2028U-TRTP+ server.

Available from V7.7, IBM Spectrum Virtualize Software only has the same software prerequisites as IBM SAN Volume Controller, and is available through both monthly license, with per-month service and support, or through perpetual license, including one year of software maintenance with each license that is acquired. The initial period of either option can be extended by the purchase of a renewal option when ordered through Passport Advantage or Passport Advantage Express.

Table 1-5 lists the ordering codes for IBM Spectrum Virtualize software only.

Table 1-5 Ordering information

Description	Part number
IBM Spectrum Virtualize Software Stg Cap Unit Monthly	D1PHVLL
IBM Spectrum Virtualize Software Stg Cap Unit Lic + SW S&S 12 Mo	D1PHTLL
IBM Spectrum Virtualize Software Stg Cap Unit Annual SW S&S RNWL	E0MZNLL
IBM Spectrum Virtualize Software Stg Cap Unit SW S&S Reinstate 12 Mo	D1PHULL

2

Planning

This chapter describes the planning process for a successful IBM Spectrum Virtualize software only installation. It also points out key differences between IBM Spectrum Virtualize software only and IBM SAN Volume Controller implementations.

This chapter provides the following information:

2.1 Planning overview

The planning process for IBM Spectrum Virtualize software only is similar to the process for IBM SAN Volume Controller nodes (for example, the 2145-SV1.) This chapter describes the planning steps that are unique to deploying IBM Spectrum Virtualize on hardware that you supply, as compared to configuring IBM SAN Volume Controller nodes, to create the IBM Spectrum Virtualize cluster.

This publication is designed to be used with *Implementing the IBM System Storage SAN Volume Controller with IBM Spectrum Virtualize V7.6*, SG24-7933. The “Planning and Configuration” chapter is a required reference to complete the planning process. That publication includes planning for the post-installation steps of configuring storage, hosts, and replication.

The most notable difference between using an IBM SAN Volume Controller node and IBM Spectrum Virtualize software only is the physical planning process. The process is specific for IBM SAN Volume Controller because all hardware is supplied by IBM and is described in *Implementing the IBM System Storage SAN Volume Controller with IBM Spectrum Virtualize V7.6*, SG24-7933. The process is similar for IBM Spectrum Virtualize software only. However, in this case, you choose the hardware and must provide the details yourself.

2.1.1 Node specification

Before your purchase nodes, you must consult the IBM System Storage® Interoperation Center (SSIC) for a list of current compatible hardware:

https://www.ibm.com/systems/support/storage/ssic/interoperability.wss

Also, consult IBM Knowledge Center:

https://www.ibm.com/support/knowledgecenter/STVLF4_7.8.1/spectrum.virtualize.781.doc/svirt_ichome_781.html

2.1.2 Node model and capabilities

At the time of writing, the Lenovo x3650 M5 and Supermicro SYS-2028U-TRTP+ servers are supported. For the current list, see the following website:

http://www.ibm.com/support/docview.wss?uid=ssg1S1009260

The server capabilities that are supported by the software are:

- One or two 8-core Intel Xeon E5-2667 v3 CPUs (two CPUs are required for compression support).
- Thirty-two or 64 GB of RAM (64 GB is required for two CPUs).
- One to four Dual-Port 16 Gb PCIe Fibre Channel (FC) HBAs (Emulex LPe16002b) (FCoE is not supported). Use two or more HBAs for resilience.
- One or two Dual-Port 10 Gb PCIe Ethernet NIC (Intel X520 SFP+ or Intel X710-DA2 SFP+).

Note: These specifications are expected to change as the capabilities of the product are enhanced. See the SSIC for a list of current compatible hardware:

https://www.ibm.com/systems/support/storage/ssic/interoperability.wss

Also, see IBM Knowledge Center:

https://www.ibm.com/support/knowledgecenter/en/STVLF4_7.8.1/spectrum.virtualize.781.doc/svc_hardware_planning.html

2.2 Physical planning

You must consider several key factors when you are planning the physical site of an IBM Spectrum Virtualize software only installation. The physical site must have the following characteristics:

- Power, cooling, and location requirements are met for the nodes and the uninterruptible power supply (UPS) units. (For more information about requirements for the UPS, see 2.2.1, “Uninterruptible power supply requirements” on page 21.)
- The nodes and their UPS units must be in the same rack.
- You must consider the maximum power rating of the rack; do not exceed it. For more information about the power rating and requirements of your nodes, and the capabilities of your rack cabinet, consult your hardware vendor.

2.2.1 Uninterruptible power supply requirements

As a software-only product, the user is responsible for ensuring that continuous power is supplied to the nodes.

Attention: An orderly shutdown must be performed before the nodes lose power. A sudden power loss to both nodes in an I/O group might lead to data loss, data corruption, or both.

UPS sizing

When sizing the UPS, consider the amount of time it takes the nodes to complete an orderly power-off process. As of this writing, this period is 2.5 minutes for solid-state drive (SSD) boot drives, and 3 minutes for hard disk drives (HDDs). For the latest information, see IBM Knowledge Center:

https://www.ibm.com/support/knowledgecenter/en/STVLF4_7.8.1/spectrum.virtualize.781.doc/svc_hardware_planning.html

Provide the node power-off time and the system power consumption to your UPS vendor to select an appropriate model.

Cabling IBM Spectrum Virtualize nodes to UPSs

To optimize availability for UPS attachment, avoid single points of failure whenever possible. The IBM Spectrum Virtualize nodes must be installed in pairs to provide high availability, and each node in the clustered system must be connected to a separate UPS unit.

Be aware of the following considerations:

- Each IBM Spectrum Virtualize node in an I/O Group must be connected to a separate UPS unit.
- Each UPS unit pair that supports a pair of nodes must be connected to a separate power domain (if possible) to reduce the chances of input power loss.

Shutdown process

At the time of writing, IBM Spectrum Virtualize software only does not have any special interface to communicate with user-supplied UPS units. Your UPS installation should include the ability to run IBM SAN Volume Controller commands; this is a common feature of UPS software for management and monitoring.

To shut down the system, your UPS should run a script to run the **`svctask stopsystem -force`** command on the cluster, followed by `yes`.

Although the exact command that must be run varies based on your SSH client, it looks something like the following example:

```
ssh -p 22 userid@1.2.3.4 "echo yes | svctask stopsystem -force
```

The **`echo yes |`** before the **`stopsystem`** is to bypass the interactive portion of the command. This might change to a dedicated AC power-loss shutdown command in a future release; consult IBM Knowledge Center for the latest details.

2.2.2 IP network planning

Although the wide variation in possible slot port configurations makes definitive rules for port assignment impossible, there are some guidelines to use in your IP network planning.

Ethernet ports

The Ethernet port setup is different than what you might be used to if you are familiar with IBM SAN Volume Controller nodes, such as the 2145-DH8.

Technician port

The lowest-numbered user port on your IBM Spectrum Virtualize software only node is reserved for use as the technician port and must be connected directly to a workstation with an Ethernet cable when used. It is used for initial node setup and for the Service GUI. (You may also access the Service GUI through a Service IP that is assigned to the management port.) Connecting the technician port to a router or switch causes it to be disabled and raises an error.

The rear panel diagrams in Appendix A, “Lenovo System x3650 M5 server implementation” on page 63 and Appendix B, “Supermicro SuperServer 2028U-TRTP+ server implementation” on page 69 show the location of the technician port for the Lenovo x3650 M5 and Supermicro SYS-2028U-TRTP+ servers.

Notes:

- Your node might have another port that is dedicated to accessing the low-level management interface for the server hardware to provide functions such as hardware monitoring and remote presence. This port is not the same thing as the technician port, and is not used by IBM Spectrum Virtualize.
- As an example, for Lenovo x3650 M5 implementations, the technician port is the port that is marked “1” (the lowest-numbered port), but this might not be the case with all implementations.

Attention: The technician port can be used to reset the Superuser passwords for the IBM Spectrum Virtualize cluster to the default value. Therefore, it is *extremely* important that you secure physical access to the cabinet that contains the nodes.

Management ports

On a 2145-DH8 IBM SAN Volume Controller Node, the management ports are the ports that are labeled “1” and “2”. On an IBM Spectrum Virtualize software only node, the available management ports are the two “next” ports after the technician port.

Note: On the initial Lenovo x3650 M5 implementation, the ports for management are the ports that are physically labeled “2” and “3”, but this might not be true for all implementations.

The node in your cluster that is the current config node responds to ARP requests for your management IPs. Although the role of config node moves to another node if the original config node fails (meaning a new node takes ownership of the management IP), a loss of network connectivity (with the config node still online) will *not* cause movement of the management IP. For this reason, assign management IPs to both sets of management ports.

Each set of management ports (for example, all “Port 2”s) should be connected to the same subnet in your network to facilitate the movement of the config node role between nodes.

iSCSI and IP replication ports

Except for the technician port, all remaining Ethernet ports can be used for iSCSI and IP replication connectivity. The management ports can be used for both management and iSCSI/IP replication traffic, but this configuration is not recommended.

Managing network ports in the GUI and CLI

The IBM Spectrum Virtualize software (both the CLI and the GUI) starts numbering the Ethernet ports on the node at “1”. The “1” refers to the first *usable* Ethernet port on the node. For the purposes of this numbering scheme, the technician port is not considered a usable port. That means that although the ports on the back panel of your node might be labeled, for example, 1 - 4, in the management interfaces they are referred to as ports 1 - 3. In such a node, the port on the rear of the system that is labeled as port “3” is managed in the GUI and CLI as port “2”.

IP addresses

You must allocate the following IP addresses:

- Management Ports. Allocate at least one address, preferably two, so that you have a backup management port on each node. (The management IP is often referred to as the “Cluster IP”.)
- Service GUI. Allocate one address for each node to access the Service GUI without having to physically attach a computer to the technician port.
- iSCSI addresses (if applicable).

TCP/IP planning

If there is a firewall in your administrative or access network, you must plan for specific TCP/IP port access to allow connectivity to the IBM Spectrum Virtualize cluster nodes from your browser.

Mandatory TCP ports include ports 22 (SSH) and 443 (HTTPS). Most users also configure NTP (123) and SMTP (25). For the full list of network services that are used by the software, see “Network requirements for IBM SAN Volume Controller” in *Implementing the IBM System Storage SAN Volume Controller with IBM Spectrum Virtualize V7.6*, SG24-7933.

The various services and options for your IBM Spectrum Virtualize software only internet protocol network are shown in Figure 2-1.

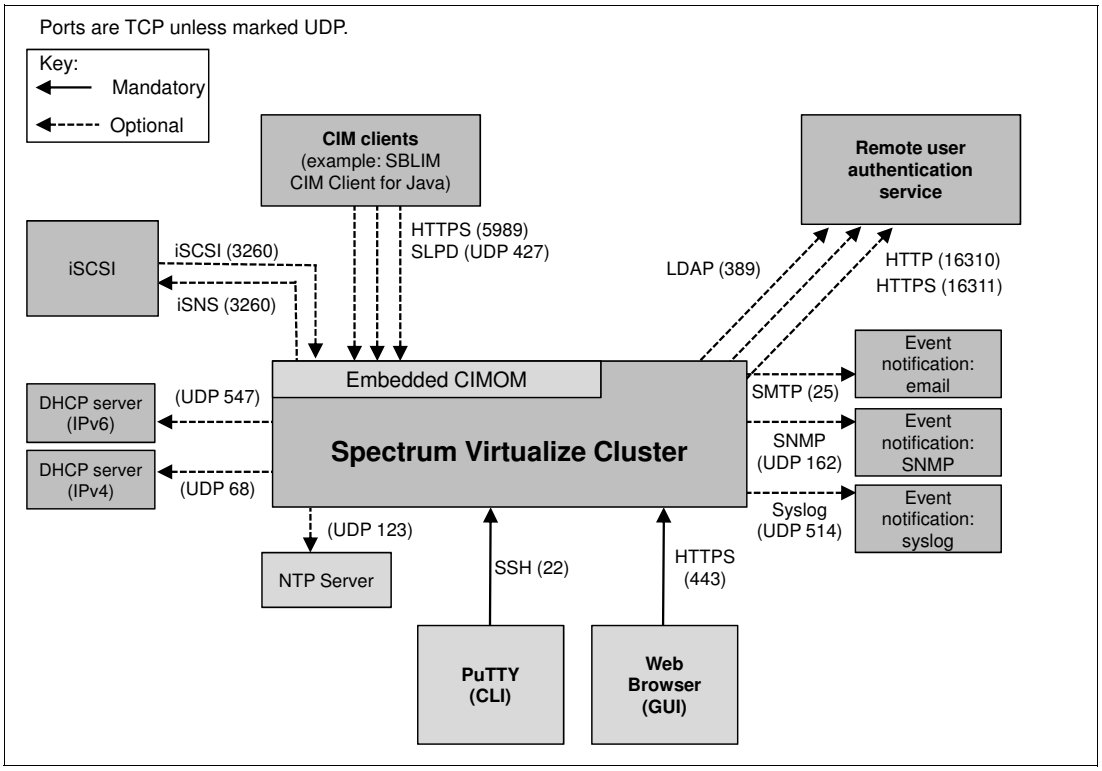

Figure 2-1 Mandatory and optional TCP/IP ports to the IBM Spectrum Virtualize cluster

2.2.3 SAN planning

The SAN planning process for IBM Spectrum Virtualize software only is similar to that of an installation with IBM SAN Volume Controller nodes. Although the configurations are not identical, you can use Chapter 2, “SAN topology”, in *IBM System Storage SAN Volume Controller and Storwize V7000 Best Practices and Performance Guidelines*, SG24-7521 as a guide. Be sure to read 2.1.2, “Topology basics”.

As with the 2145-DH8 and SV1 nodes, you cannot determine the FC WWPNs before installation. For cabling your nodes to the SAN, it is vital to plan and document first, then track which SAN ports are connected to which node ports. With this information, you can accurately assign your SAN ports to zones. For more information about IBM Spectrum Virtualize zoning, see “SAN zoning and SAN connections” in *Implementing the IBM Storwize V7000 and IBM Spectrum Virtualize V7.6*, SG24-7938.

IBM Spectrum Virtualize software only supports 2 - 8 FC ports on 1 - 4 FC HBA cards. The following dual-port cards are supported:

- Lenovo Emulex 16GB FC Dual-port HBA (81Y1662)
- IBM Spectrum Virtualize software supports 2 - 16 FC ports on 1 - 4 FC HBA cards, depending in the hardware platform. For information about which cards are supported, and the platforms that support them, see the Hardware Requirements page on the IBM Support website:

  `http://www.ibm.com/support/docview.wss?uid=ssg1S1009260`

Resilience: For a resilient solution, use a minimum of two FC HBA cards.

The WWPNs of the FC ports follow similar numbering rules as IBM SAN Volume Controller nodes. The format that is used for IBM Spectrum Virtualize software only nodes is `5005076071SP`*`XXXX`*, where “S” is the slot number, “P”, the port number, and “XXXX” is the sequence number (this name is generated by the software and is based on the installation signature file for each node).

Before you start the installation, prepare your SAN switches with (empty) zoning aliases for the IBM Spectrum Virtualize ports.

During the installation process, after each node reaches the installation stage where the technician port is live, you must check your SAN fabric to determine the WWPNs that are assigned to the ports and map them to the appropriate aliases.

2.3 Initial cluster installation planning checklist

Use the following checklist to help prepare for a successful installation of IBM Spectrum Virtualize software only:

1. Produce and procure a configuration for your IBM Spectrum Virtualize nodes by using the guidelines in the SSIC and IBM Knowledge Center. The most prominent difference between using an IBM SAN Volume Controller node and IBM Spectrum Virtualize software only is the physical planning process.

   The process is specific for IBM SAN Volume Controller because all hardware is supplied by IBM and is described in *Implementing the IBM System Storage SAN Volume Controller with IBM Spectrum Virtualize V7.6*, SG24-7933. The process is similar for IBM Spectrum

Virtualize software only. However, in this case, you choose the hardware and so you must provide the details yourself.

2. Consult the vendor for the node hardware to determine physical requirements, especially the power consumption.
3. Consult your UPS vendor to procure an appropriate UPS installation that can support the nodes for the appropriate period and can send the necessary node shutdown command. See 2.2.1, “Uninterruptible power supply requirements” on page 21.
4. Download the IBM Spectrum Virtualize `.iso` file and load it onto a USB flash drive. For more information, see 3.1.1, “Creating a bootable USB flash drive” on page 30.
5. Determine the IP addresses to be used for the management ports:

   First Management (Cluster) Ports IP: ___.___.___.___

   Second Management (Cluster) Ports IP: ___.___.___.___
6. Determine the IP addresses to be used for the service ports:

   Node 1 Service: ___.___.___.___

   Node 2 Service: ___.___.___.___

   Node 3 Service: ___.___.___.___

   Node 4 Service: ___.___.___.___

   Node 5 Service: ___.___.___.___

   Node 6 Service: ___.___.___.___

   Node 7 Service: ___.___.___.___

   Node 8 Service: ___.___.___.___
7. Prepare SAN zoning aliases for your FC ports by using a consistent naming scheme, such as SV_Cluster_Node_X_Port_Y. It is advantageous to use an alias for each single WWPN where subsequent zoning refers to the alias. This schema simplifies some troubleshooting resolution, as described in 4.3.2, “Replacing nodes disruptively” on page 60.
8. Prepare firewalls to pass the necessary TCP traffic. See “TCP/IP planning” on page 24.
9. Install the nodes and UPS in your rack cabinet.
10. Obtain a Keyboard, Video, and Mouse (KVM) to be used during the installation process. (It is not necessary after the installation is complete.)
11. Connect and label the necessary LAN and SAN cables. See “Ethernet ports” on page 22 and 2.2.3, “SAN planning” on page 25. Document all IBM Spectrum Virtualize ports that are attached to corresponding SAN switch ports.
12. Install the software on the first node and create the cluster. See 3.3.4, “Creating the first node of an IBM Spectrum Virtualize cluster” on page 35.
13. Install the software on the remaining nodes and leave them in candidate node status. See 3.5.4, “Adding a node to the IBM Spectrum Virtualize cluster” on page 46.
14. Using the SAN switch management tools, determine the WWPN belonging to each FC Port in your cluster.
15. Configure your SAN fabrics to insert the WWPNs for your nodes into their zone aliases.
16. Add the remaining nodes to the cluster.
17. Create a user account for the UPS system.
18. Perform a UPS shutdown test by removing the power feed to your UPS system, and verify that the UPS successfully initiates a graceful shutdown of the IBM Spectrum Virtualize

nodes. Ensure that the nodes power themselves down successfully before the battery in the UPS being depleted.

19. Create “as-built” documentation, including all physical connections for future reference.

2.4 Post-installation tasks

After you install your IBM Spectrum Virtualize cluster, you are ready to complete your configuration to turn it into a core storage management system. The details are beyond the scope of this publication. For more information, see *Implementing the IBM System Storage SAN Volume Controller with IBM Spectrum Virtualize V7.6*, SG24-7933.

Here are some of the tasks that you must complete:

- Add storage devices to storage pools.
- Create volumes from the storage pools.
- Assign hosts to the volumes.
- Configure system monitoring.

You must also configure optional features as required:

- Configure Easy Tier so that frequently accessed storage is assigned to faster pools.
- Configure Copy Services, if needed. Copy Services includes FlashCopy, Global Mirror, Metro Mirror, and Volume Mirroring.
- Implement the Stretch Cluster function, which uses the Volume Mirror function to provide a high-availability capability to IBM Spectrum Virtualize.
- Configure compression or space-efficient volumes to maximize the available capacity.

3

Initial installation

This chapter explains the initial configuration that is required for IBM Spectrum Virtualize software only on the customer-supplied Intel server.

This chapter provides the following information:

3.1 Setting up the IBM Spectrum Virtualize cluster

This section provides step-by-step instructions to create the IBM Spectrum Virtualize cluster with IBM Spectrum Virtualize software only. An IBM Spectrum Virtualize cluster must be created to provision and manage virtualized storage.

Prerequisite: You must follow the planning guide and complete the pre-installation check list as described in Chapter 2, “Planning” on page 19.

This process includes multiple phases:

1. The first phase of the installation is to create a cluster. This phase requires physical interaction with the Intel hardware. You must load IBM Spectrum Virtualize software only from a USB key onto the system’s internal disk storage.
2. The second phase consists of accessing the Intel system with the technician port from a device with browser capability, for example, your personal computer or notebook.
3. The last phase is the final configuration through the administration network to the IBM Spectrum Virtualize GUI access with the cluster IP address. The last phase uses the IBM Spectrum Virtualize web browser access method, as shown in Figure 3-28 on page 48.

3.1.1 Creating a bootable USB flash drive

As noted on the pre-installation checklist, the ISO image that is downloaded from IBM Passport Advantage is used to make a bootable USB drive. There are multiple methods of creating the bootable USB drive. For an example method, see IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/en/SS42VS_7.2.7/com.ibm.qradar.doc/t_siem_inst_create_bootable_usb_RHL.html

A reliable USB drive is necessary to successfully install the Intel machine with IBM Spectrum Virtualize software only to create the IBM Spectrum Virtualize node. The USB drive must have a capacity of at least 8 GB and is used in the first phase of installing each of the cluster nodes.

3.1.2 Hardware installation

The installation process in this chapter shows the installation of two servers (that are supported for IBM Spectrum Virtualize software only) to create a cluster that consists of a single I/O Group. The hardware is identical in terms of processors, memory, and I/O adapters. Each server is connected to a keyboard and monitor to complete the initial IBM Spectrum Virtualize software installation.

It is also assumed that the systems are physically connected to the LAN and SAN infrastructure, as described in Chapter 2, “Planning” on page 19.

3.2 Initial power on and system boot

This is the first phase of the installation. You should have a prepared USB drive to use for loading and configuring both servers. Power off each server before starting the installation process.

Follow these steps for the initial configuration phase:

1. Insert the USB key into an available USB port on the server.
2. Power on the server and watch the power-on routine on the monitor.
3. Interrupt the power-on sequence to manually select the boot device.
4. Select the USB port that has the USB key and continue with the boot process.

   The server continues its power-on cycle and boots from the USB key that is placed in the port, and shows the initial panel (see Figure 3-1) when the machine begins booting from the USB key.

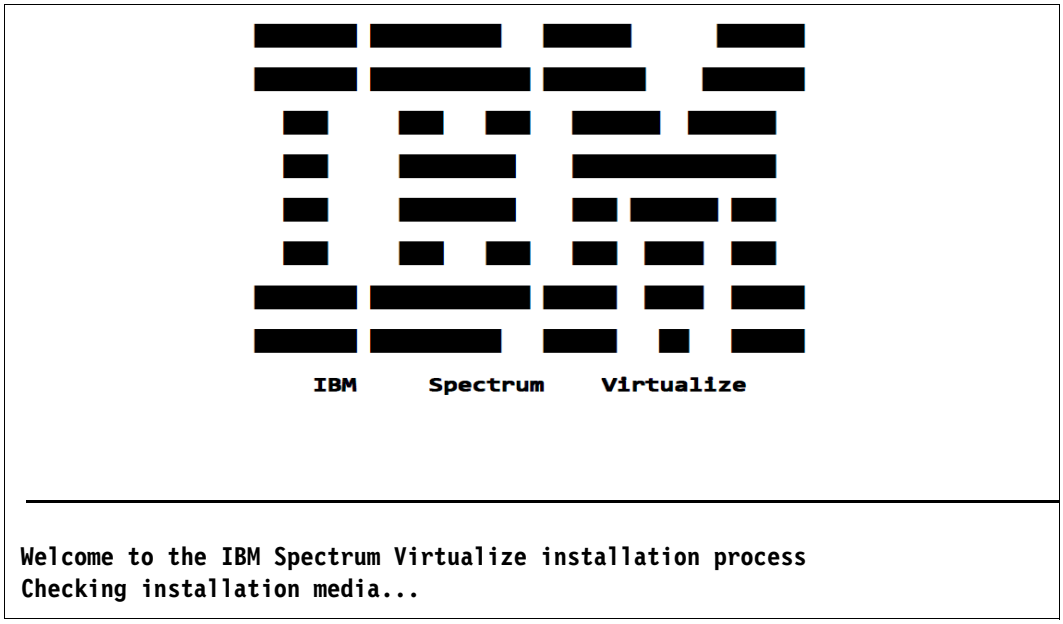

Figure 3-1 Initial boot panel with IBM Spectrum Virtualize

3.2.1 Gathering information to create the Spectrum Node Activation Key file

When the load from USB completes, a six-digit code is displayed. This code is unique for that specific machine. Write down this code because it is used in later steps to obtain a Unique Spectrum Virtualize Node ID (usvnid) or signature file to create the IBM Spectrum Virtualize node.

Figure 3-2 shows the code installation as complete. It also shows the six-digit code to be used to create the `usvnid` file that is referred to in 3.2.2, “Obtaining an IBM Spectrum Virtualize Node Activation Key file” on page 32.

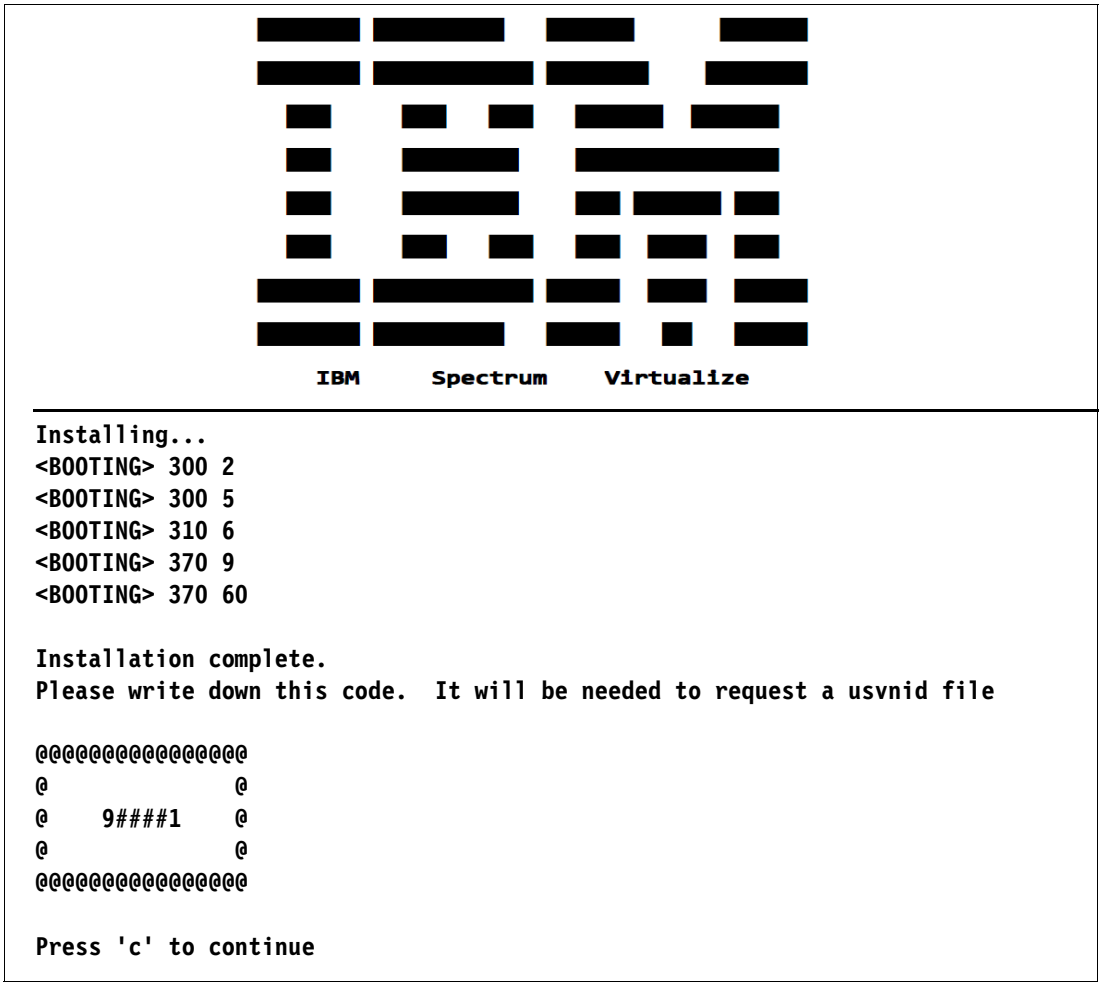

Figure 3-2 IBM Spectrum Virtualize installation complete panel

You are prompted to write down the six-digit code before continuing, and then to remove the USB key as instructed. When the installation resumes (the node boots more than once), the completion is indicated by a login prompt.

3.2.2 Obtaining an IBM Spectrum Virtualize Node Activation Key file

An activation key is required to continue. Go to the IBM website to obtain the usvnid activation key or Spectrum Node Activation Key file for that node. It can be generated at the following link:

`https://extbasicispr02.podc.sl.dst.ibm.com/support/home/spectrum-virtualize`

Follow the instructions that are provided.

When you enter the code that you obtained in 3.2.1, “Gathering information to create the Spectrum Node Activation Key file” on page 31, a file is generated and downloaded through your browser to your personal computer or notebook, which you can use to continue the installation process. Copy the file to the system that you will use for the next phase of the installation. This file is referred to as the Spectrum Node Activation Key.

3.3 Connecting to the technician port

IBM Spectrum Virtualize software only uses the feature that is called a *technician port,* which was introduced in the IBM SAN Volume Controller 2145-DH8 node. On that node, Ethernet port 4 is allocated as the technician port, and is marked with a ‘T’ on the node. However, for Spectrum Virtualize software only, the hardware ports cannot be determined. The first 1-Gbps Ethernet port that is defined by the hardware is allocated to be the technician port. Identify this port on the new nodes, and connect a cable from it to the Ethernet port of the notebook or computer that is used to continue the installation.

The next configuration steps for each node are done with the technician port. The port broadcasts a Dynamic Host Configuration Protocol (DHCP) service so that a notebook or computer is automatically assigned an IP address on connection to this port. The default IP address for a new node is `192.168.0.1`. However, you can also use a static IP address, which should be set to `192.168.0.2` on your notebook or computer Ethernet card.

Important: IBM Spectrum Virtualize software only does not provide IPv6 IP addresses for the technician port.

After the cluster configuration is complete, and the Ethernet cable is changed to the correct network, the technician port automatically routes the connected user directly to the correct GUI.

3.3.1 Completing the initial node configuration

Follow these steps for all nodes (including the first node in the IBM Spectrum Virtualize cluster and additional nodes):

1. Attach an Ethernet cable from the notebook or computer to the technician port. After you connect your PC or notebook to the technician port, validate that you have a correct IPv4 DCHP address, for example, `192.168.0.12` (the first IP address that the IBM Spectrum Virtualize node assigns).
2. Connect through your browser to the node (`http://192.168.0.1`).
3. Upload the file that is generated and obtained by following the directions in 3.2.2, “Obtaining an IBM Spectrum Virtualize Node Activation Key file” on page 32.

Note: During the initial configuration, you see certificate warnings because the Spectrum Virtual node certificates are self-issued. Accept these warnings to continue.

Figure 3-3 shows the initial window prompting for the Spectrum Node Activation Key.

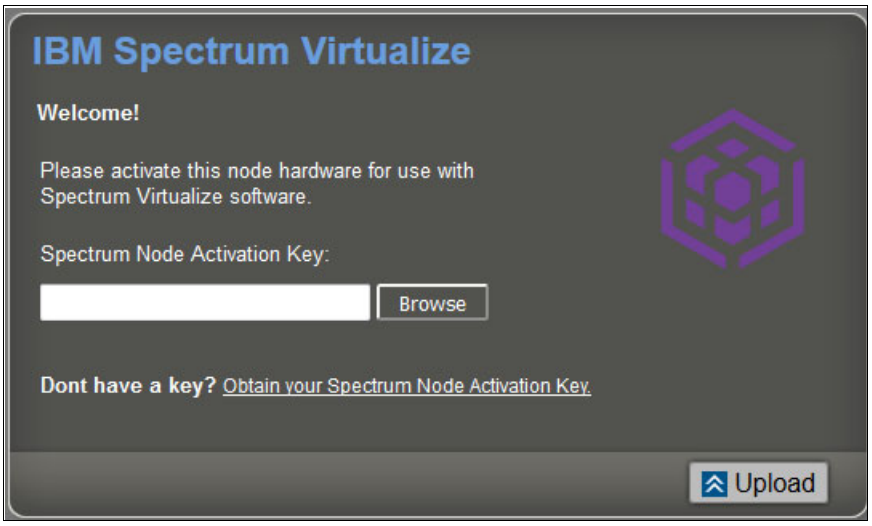

Figure 3-3 Prompt window to load Spectrum Node Activation Key

When the file is successfully uploaded and automatically verified, you are prompted to specify whether this node is the initial node of a cluster or an additional node for an existing cluster, as shown in Figure 3-4.

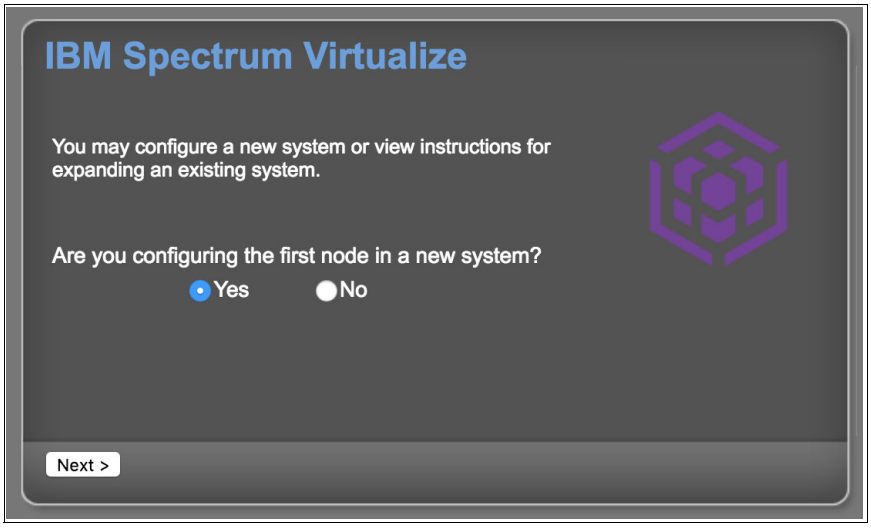

Figure 3-4 Prompt for new or existing cluster addition

3.3.2 SAN zoning requirements

At this point in the configuration process, the host bus adapters (HBAs) on the node (if cabled) can log in to the fabric switches. The worldwide port name (WWPN) can be obtained from the fabric switch. When the HBAs are logged in to the SAN switch administration interface, the WWPN is displayed on the physical SAN switch port that the HBA is cabled to.

The WWPN for each HBA is essential to create the fabric aliases and zones so that the IBM Spectrum Virtualize nodes can communicate with other devices on the SAN fabrics. It is assumed that ‘soft zoning’ or WWPN is used for all SAN fabric zoning. These devices include other IBM Spectrum Virtualize nodes, storage controllers, and host servers.

After the WWPNs are identified, they can be used to create zones to communicate between the IBM Spectrum Virtualize nodes, as described in *Implementing the IBM System Storage SAN Volume Controller with IBM Spectrum Virtualize V7.6*, SG24-7933. All zoning information should be included from the steps in Chapter 2, “Planning” on page 19 and any resulting worksheet. The WWPNs for the node can now be added to update the worksheet.

3.3.3 Completion of additional node for an existing cluster

Continuing from Figure 3-4 on page 34, if the new node is an addition to an existing IBM Spectrum Virtualize cluster and you click the **No** option, the message in Figure 3-5 is shown. This completes the installation for an additional node. After the SAN zoning is complete, it is ready to be added to the existing IBM Spectrum Virtualize cluster. This step is described in 3.5, “Adding a node to an existing IBM Spectrum Virtualize Cluster” on page 45.

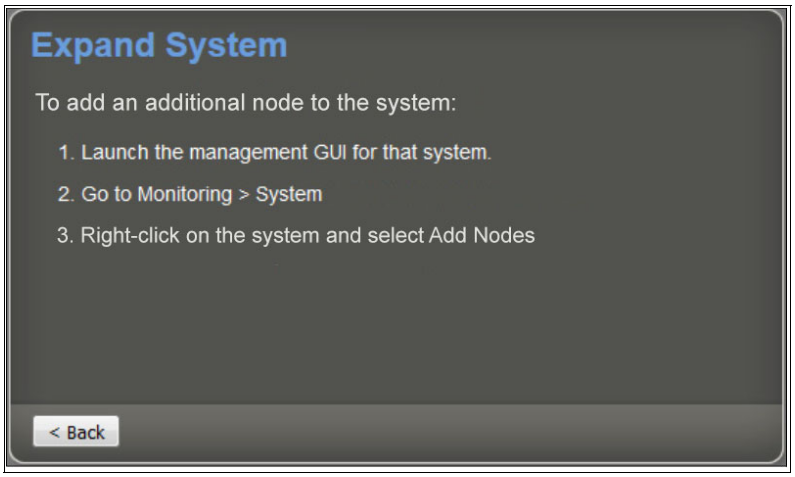

Figure 3-5 Configuration complete for additional node

3.3.4 Creating the first node of an IBM Spectrum Virtualize cluster

When you reach the window that is shown in Figure 3-4 on page 34, click the **Yes** option to continue the configuration and create the first node of an IBM Spectrum Virtualize cluster. You are prompted for more details to continue the configuration.

Figure 3-6 shows where you add the cluster IP address details from your planning worksheet.

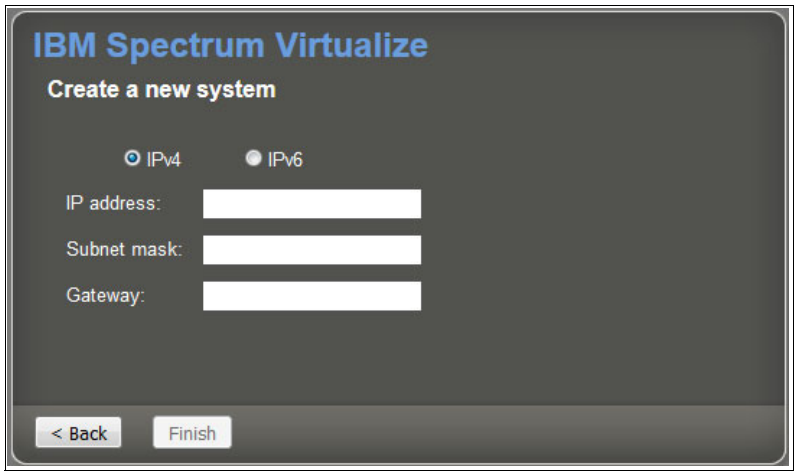

Figure 3-6 Request cluster IP information

When the cluster IP details are entered, click **Finish**, and the node initializes the configuration, as shown in Figure 3-7. This completes the configuration work from the technician port.

Figure 3-7 Initialization window

Figure 3-8 on page 37 shows the initialization completion window, which is shown when the process completes. By following the instructions and clicking **OK**, you can continue the configuration by using the management GUI or the IBM Spectrum Virtualize web browser interface.

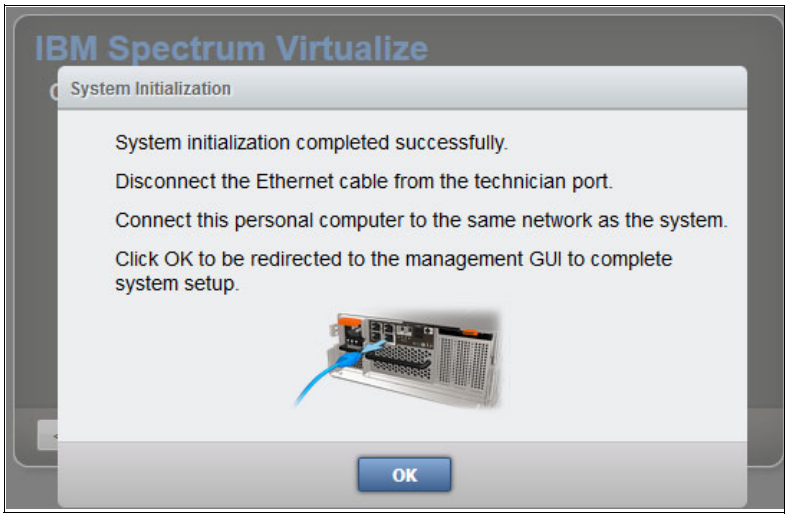

Figure 3-8 Initialization is complete

3.4 Configuring the first IBM Spectrum Virtualize node

In this phase of the IBM Spectrum Virtualize installation and configuration, you complete the configuration of the first node and add a second node to create the first cluster I/O Group. When this configuration is complete, the IBM Spectrum Virtualize Cluster is available to attach and manage storage resources through a SAN. It also is available to attach host server devices through SAN and provision and manage storage.

3.4.1 Connecting to the IBM Spectrum Virtualize cluster

The cluster IP address that is configured in 3.3.4, “Creating the first node of an IBM Spectrum Virtualize cluster” on page 35 is added to the first *available* 1-Gb Ethernet port, which does not include the technician port. Depending on the Intel hardware type, the first available port is physically next to the technician port or, if the ports are labeled, the next incremental port number.

You can access the management GUI by opening any supported web browser. Complete the following steps:

1. Open the web GUI from a supported web browser on any workstation that can communicate with the cluster on the internet protocol network.
2. Point to the IP address that you entered in Figure 3-6 on page 36:

   `http://<svclusteripaddress>/`

   (You are redirected to `https://<svclusteripaddress>/`, which is the default address for access to the IBM Spectrum Virtualize cluster.)

When connected, you are prompted to accept the license agreement, as shown in Figure 3-9. Read the license agreement and click the **Accept** arrow.

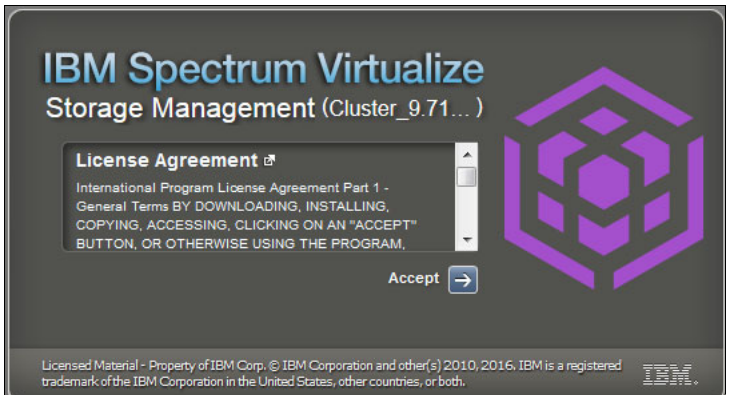

Figure 3-9 License agreement

You are prompted to create and confirm a password for the superuser account, as shown in Figure 3-10.

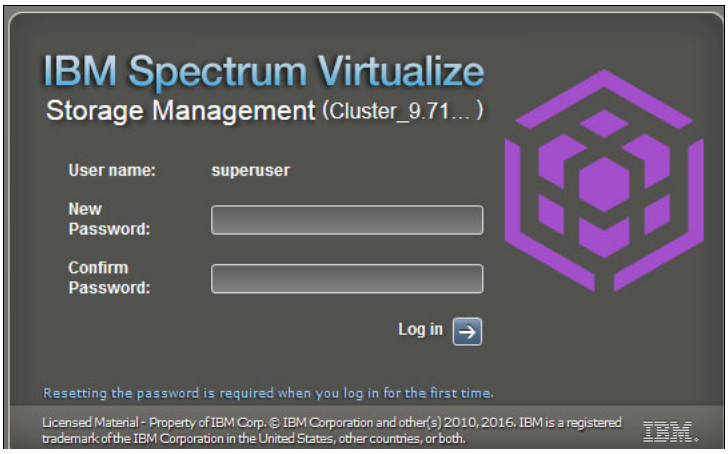

Figure 3-10 Create a superuser password

3. Type a new password and type it again to confirm it as instructed. The password length is 6 - 63 characters. *The password cannot begin or end with a space*. After you type the password twice, click the **Log in** arrow.

3.4.2 System setup tasks for IBM Spectrum Virtualize

After the superuser credentials are created, the Welcome to System Setup window opens (Figure 3-11 on page 39).

Complete the following steps:

1. Click **Next** to continue the installation process (Figure 3-11).

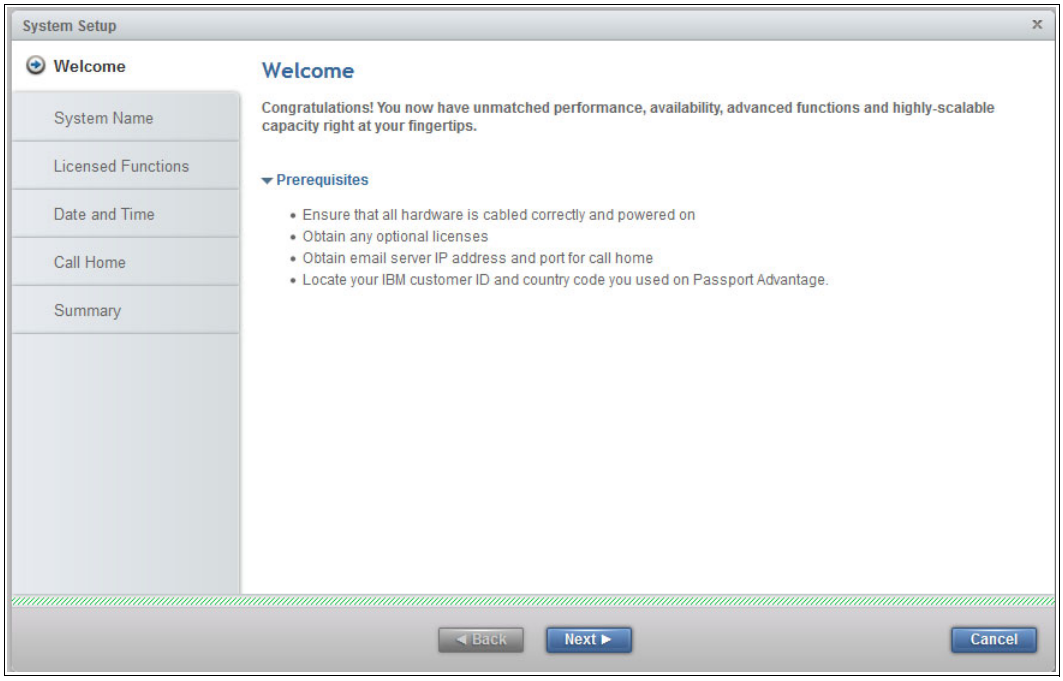

Figure 3-11 Welcome to System Setup window

You can choose to give the cluster a new name. As shown in Figure 3-12, the name *Cluster_Redbook_1* is given to the cluster.

2. Click **Apply and Next** after you enter your cluster name.

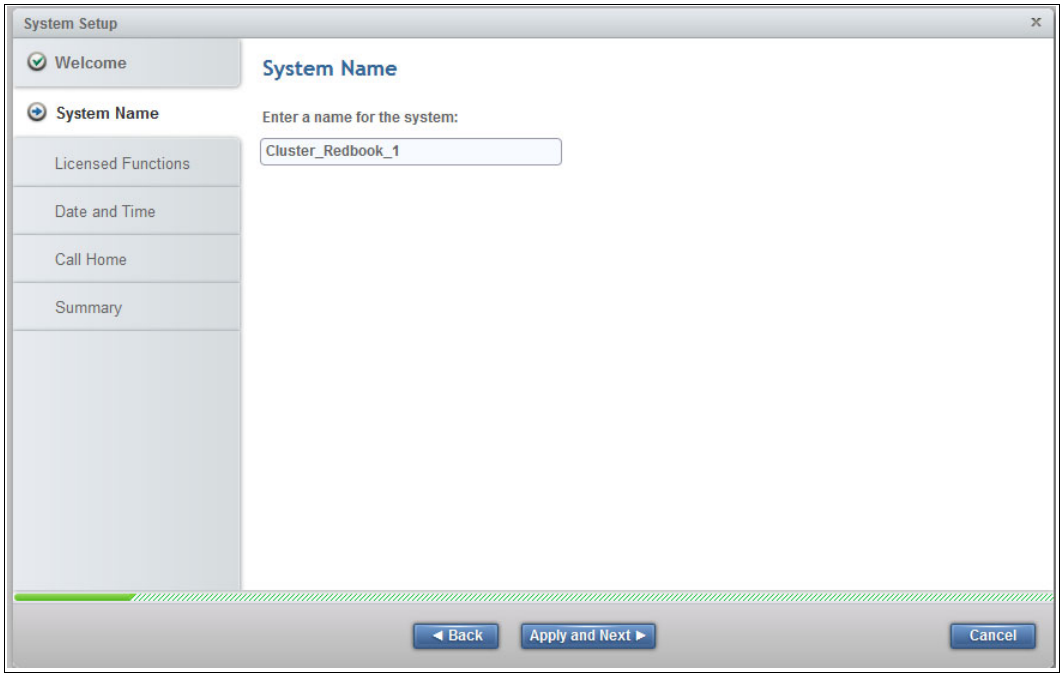

Figure 3-12 Cluster System Name

The System Setup action is confirmed by a message, as shown in Figure 3-13.

Click **Close** to continue.

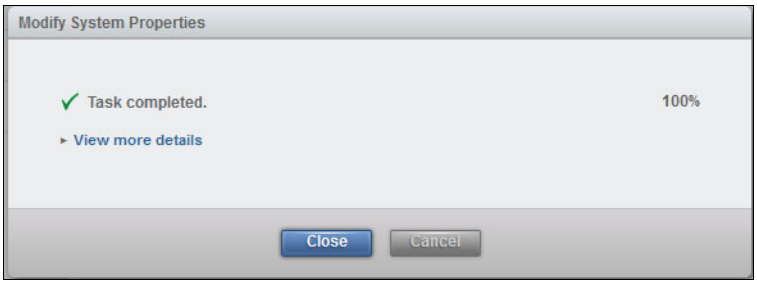

Figure 3-13 Task completed

A similar message is repeated after all actions from the System Setup Menu, where you click **Close** to continue the configuration process.

3. Enter the total purchased capacity for your system as authorized by your license agreement in the Licensed Functions window (see Figure 3-14). Click **Apply and Next** when done.

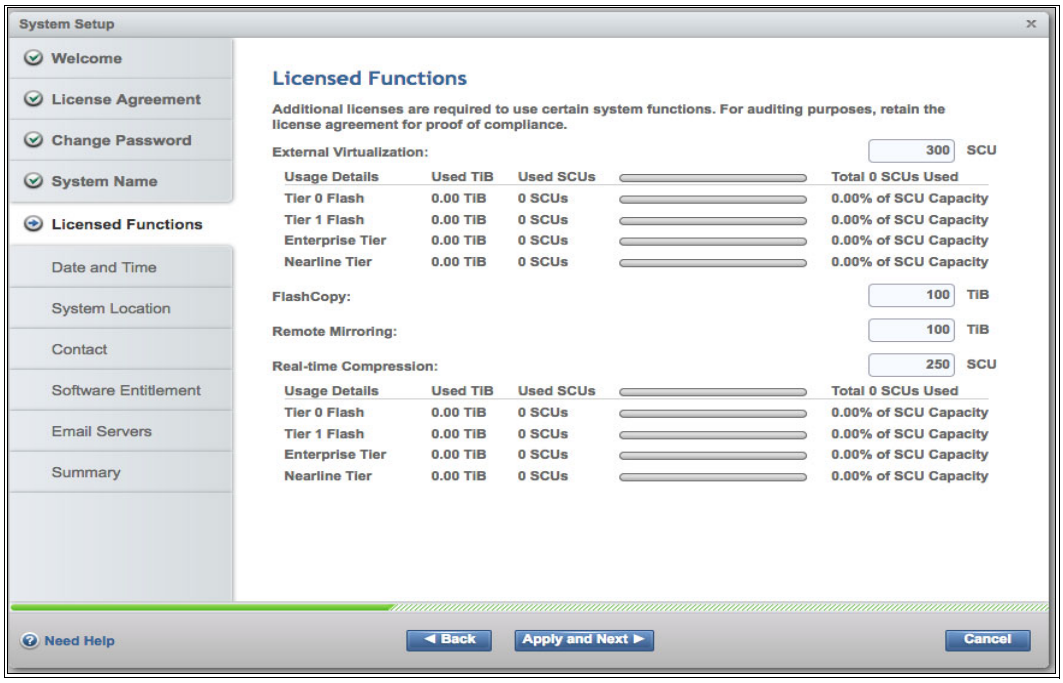

Figure 3-14 Licensed Functions

4. Set the time and date, as shown in Figure 3-15 on page 41. In this example, the date and time were set manually by selecting the **Manually** option and **Use Browser Settings**. At this time, you cannot choose to use the 24-hour clock. You can change to the 24-hour clock after you complete the initial configuration.

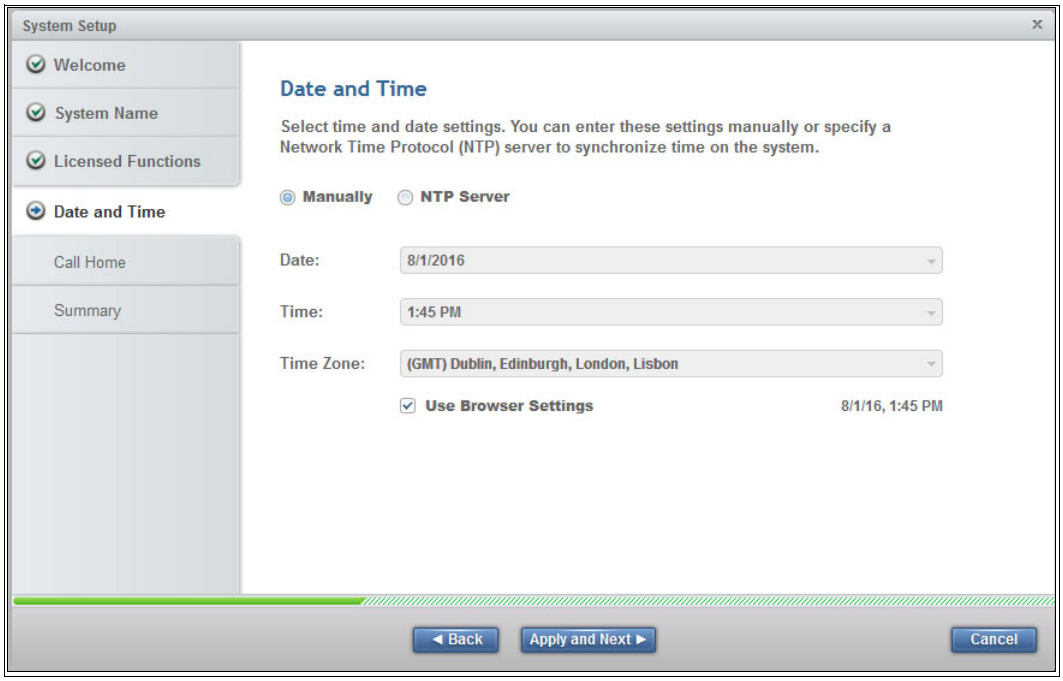

Figure 3-15 Manual browser-based option for Date and Time

5. To ensure that all of your SAN and storage devices have a common time stamp for troubleshooting, use a Network Time Protocol (NTP) server option. If you choose the **NTP Server** option, you must add only the IP address, as shown in Figure 3-16. Click the **Apply and Next** arrow when done.

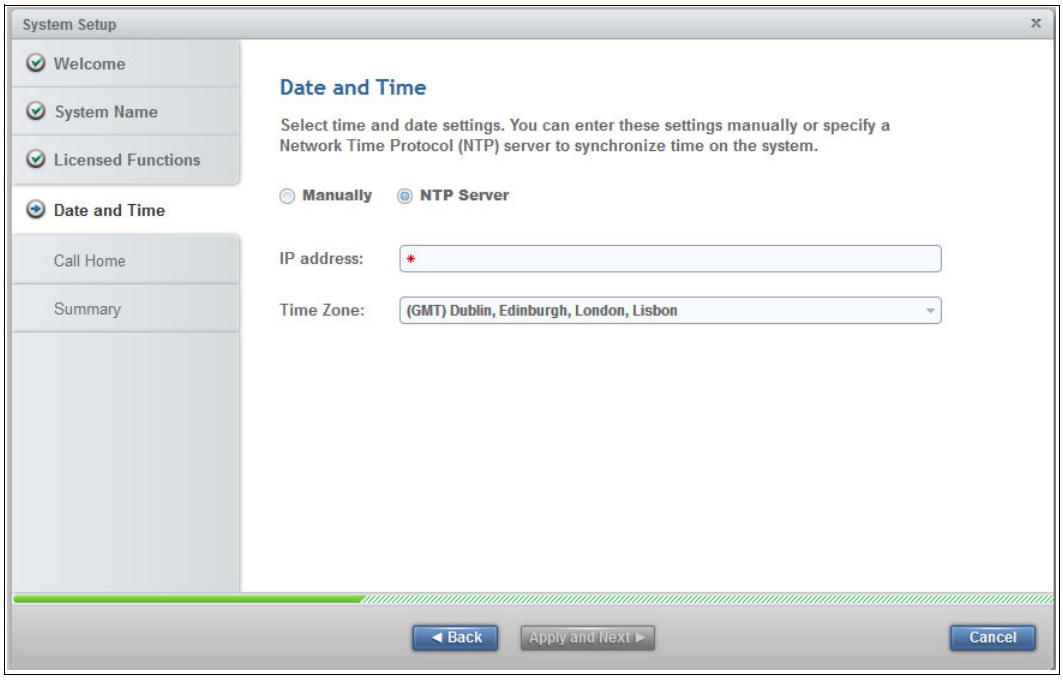

Figure 3-16 NTP option for Date and Time

6. Configure the Call Home settings.
   a. In the first window, set the System Location information, as shown in Figure 3-17. All fields are required in this step before you click the **Next** arrow.

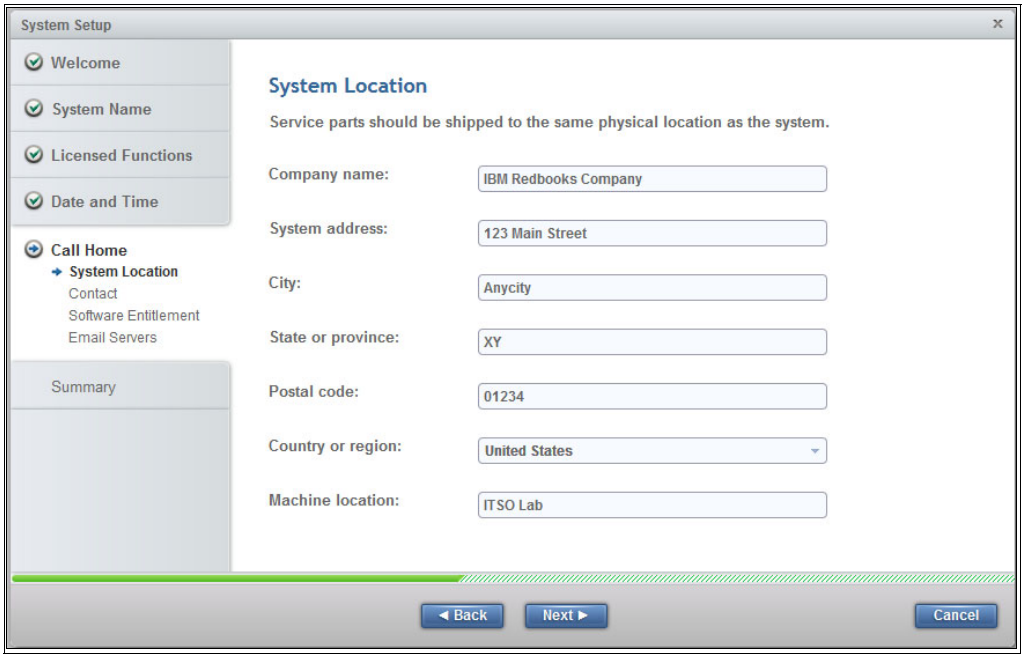

Figure 3-17 System Location

   b. Enter the Contact information, as shown in Figure 3-18. All fields are required, except for Phone (alternate). Click the **Apply and Next** arrow when done.

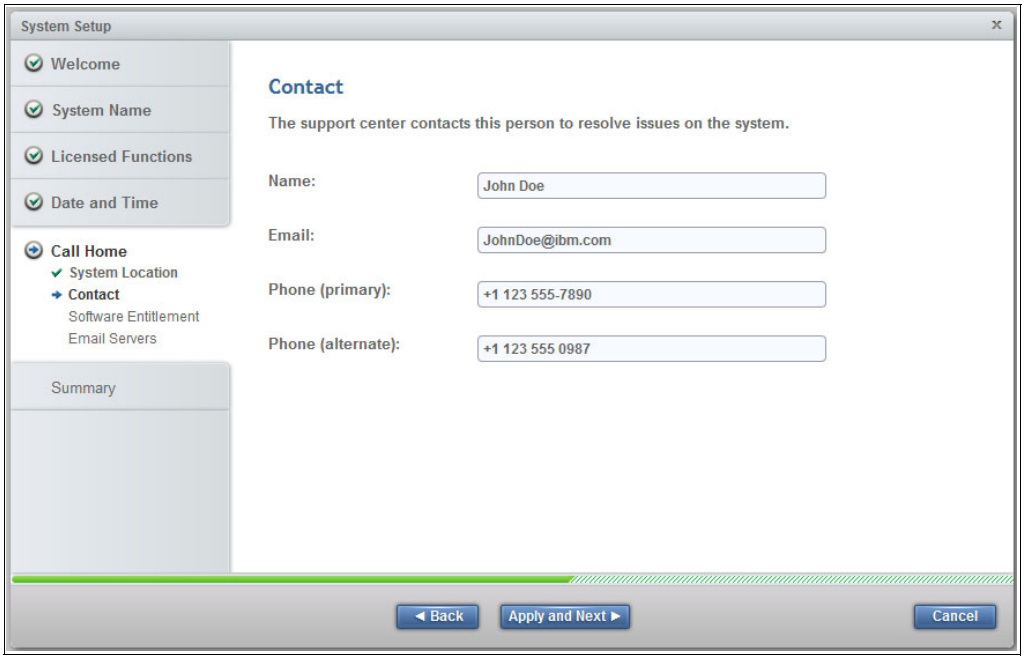

Figure 3-18 Contact details

c. Enter the Software Entitlement details, as shown in Figure 3-19. Click **Apply and Next** when done.

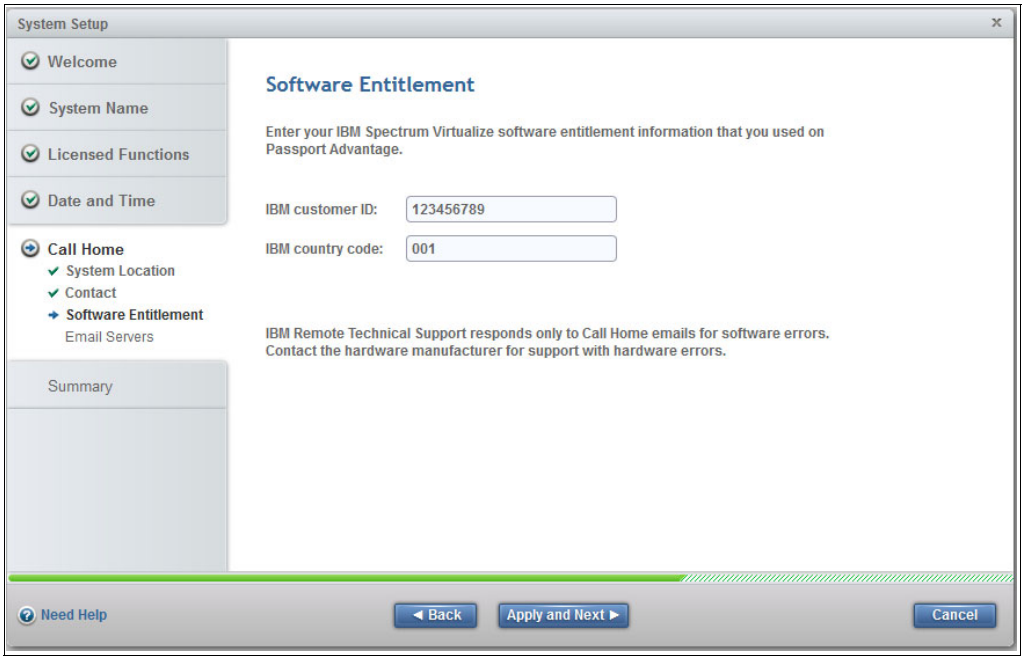

Figure 3-19 Software entitlement

d. Provide the IP address details of the email server, as shown in Figure 3-20. Click **Apply and Next** when done.

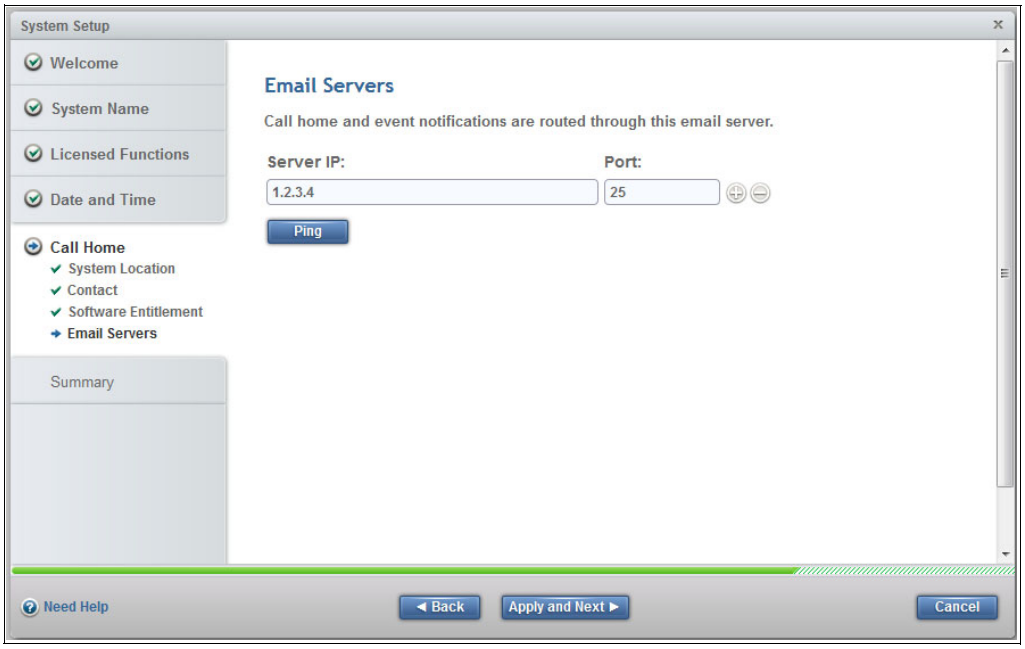

Figure 3-20 Email Servers

7. A summary of all the System Setup information is shown (Figure 3-21). Click **Finish** to complete the setup.

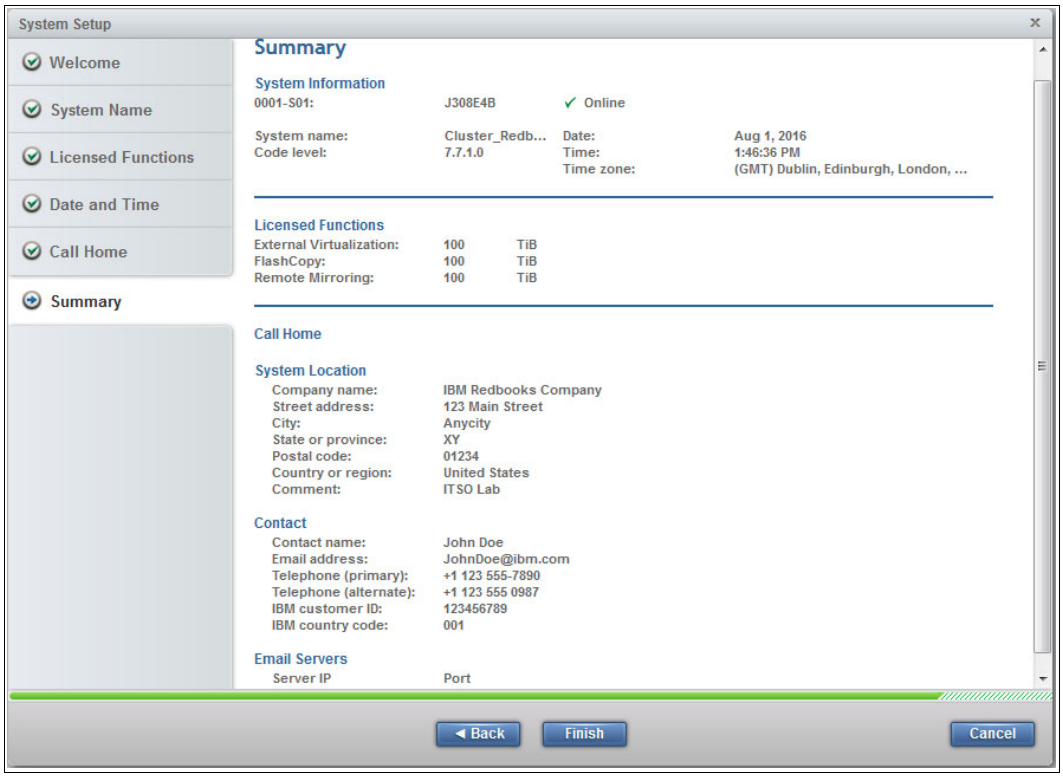

Figure 3-21 System Setup Summary

Figure 3-22 confirms that the setup is complete and you are redirected to the management GUI.

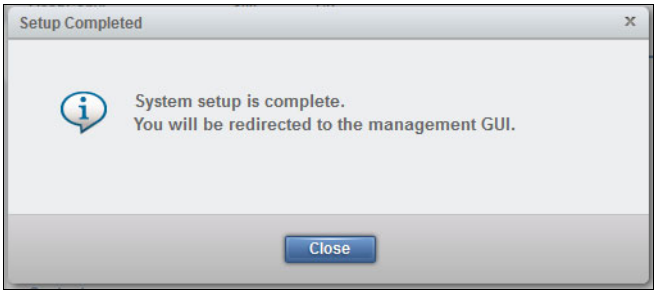

Figure 3-22 Setup Complete

Figure 3-23 shows the management GUI with a single IBM Spectrum Virtualize node.

Figure 3-23 Management GUI with a single node in the cluster

3.5 Adding a node to an existing IBM Spectrum Virtualize Cluster

This section continues the process to complete the first I/O Group configuration after the creation of the first node in the cluster. You can also repeat the following steps to add subsequent nodes to the cluster and create more I/O Groups. The following tasks are required:

- 3.5.1, “Installing IBM Spectrum Virtualize on a system to create an IBM Spectrum Virtualize node” on page 45
- 3.5.2, “Determining the WWPN from SAN switches” on page 45
- 3.5.3, “Verifying WWPNs and creating a connection to the node by SAN zoning” on page 46
- 3.5.4, “Adding a node to the IBM Spectrum Virtualize cluster” on page 46

3.5.1 Installing IBM Spectrum Virtualize on a system to create an IBM Spectrum Virtualize node

To create a second node to add to the initial cluster node, repeat the procedures in 3.2, “Initial power on and system boot” on page 31 continuing through 3.3.3, “Completion of additional node for an existing cluster” on page 35. These tasks must be completed on hardware that is identical to the first IBM Spectrum Virtualize node.

3.5.2 Determining the WWPN from SAN switches

When this task is completed, the WWPNs must be noted as described in 3.3.2, “SAN zoning requirements” on page 34. The node must then be zoned to communicate with the first IBM Spectrum Virtualize cluster node.

3.5.3 Verifying WWPNs and creating a connection to the node by SAN zoning

The WWPNs on the single IBM Spectrum Virtualize cluster node can be verified from the management GUI. To view it, select *Cluster_Name* and then click **Settings** → **Network** to view the Fibre Channel (FC) ports. An example is shown in Figure 3-24 in the *Local WWPN* column.

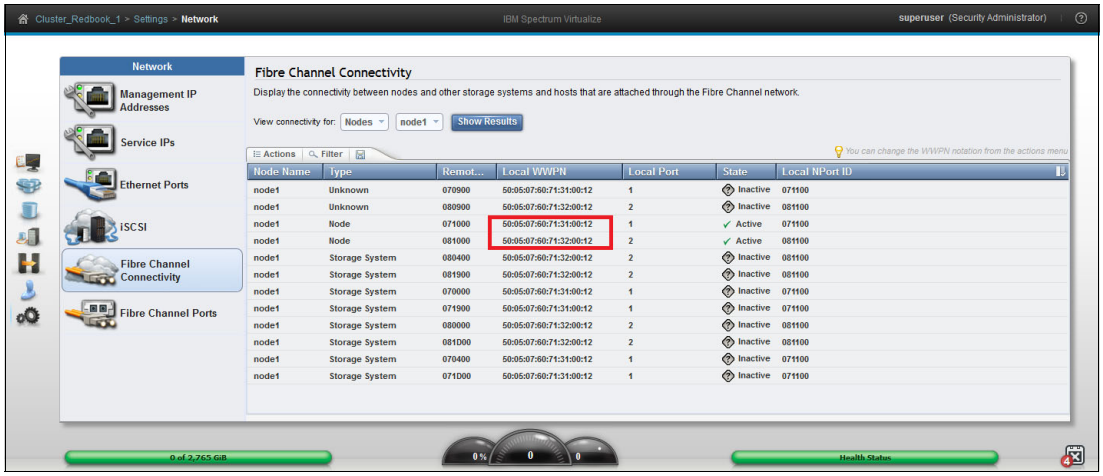

Figure 3-24 Fibre Channel Connectivity

The WWPNs in Figure 3-24 are highlighted for node1. They are listed as Active in the Status column. Use this information to verify the WWPNs with SAN port details when the node was installed (as described in 3.3.2, “SAN zoning requirements” on page 34).

When the SAN aliases zones are created and activated, the cluster node and new cluster candidate node can communicate by SAN. You can now add the second node to the cluster with the management GUI.

3.5.4 Adding a node to the IBM Spectrum Virtualize cluster

Complete the following steps to add a node to the IBM Spectrum Virtualize cluster:

1. Select *Cluster_Name*, click **Monitoring** → **System**, click **Actions** (in the upper left corner), and select **Add Node**. Figure 3-25 on page 47 shows that the new node is seen by the cluster. Select the node and click **Finish** to continue.

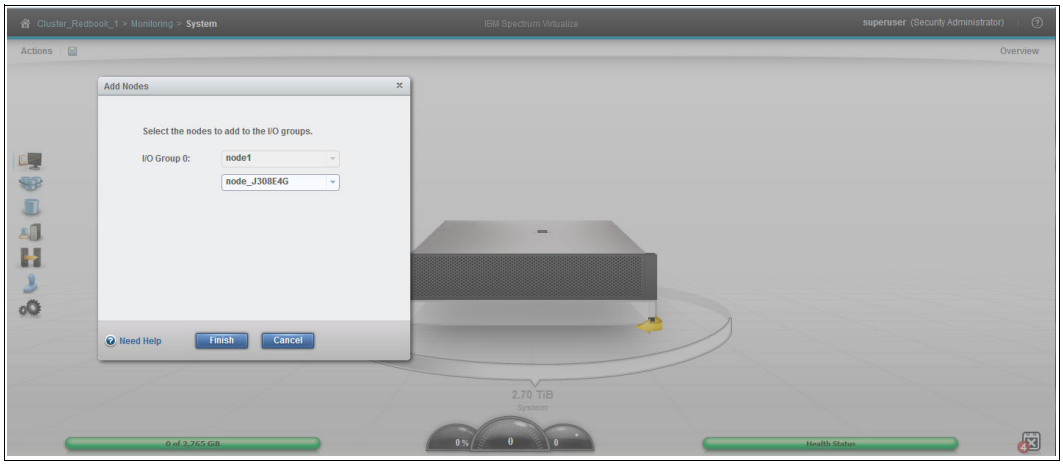

Figure 3-25 Add node to the cluster

2. Figure 3-26 shows the completed task. Click **Close** to continue.

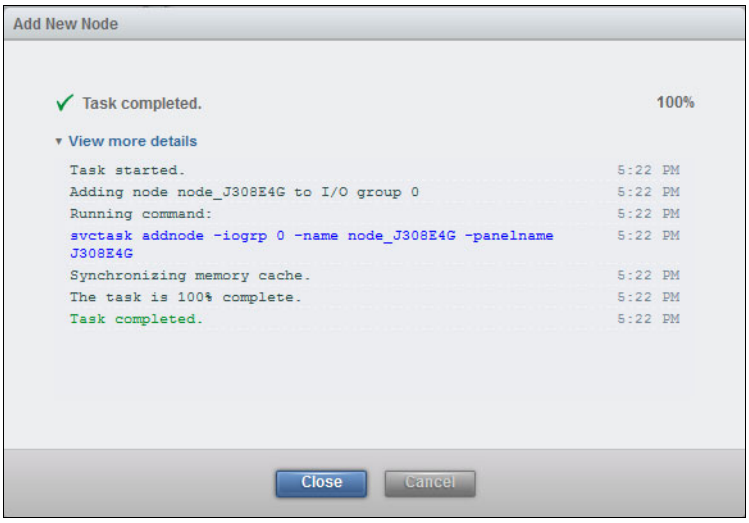

Figure 3-26 Add node task completed

The second node is then shown in the *Cluster_Name* → **Monitoring** → **System** display, as shown in Figure 3-27.

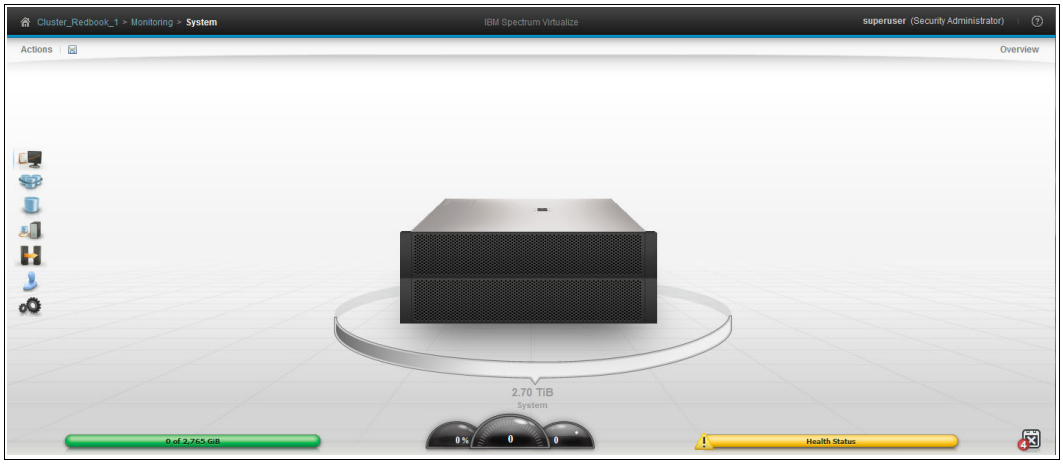

Figure 3-27 Second node added to cluster confirmation

3.6 Managing the cluster

You can manage the IBM Spectrum Virtualize cluster in many ways. The following methods are the most common ones:

- By using the IBM Spectrum Virtualize Web Management graphical user interface (GUI)
- By using a SSH-based IBM Spectrum Virtualize command-line interface (CLI)
- By using IBM Spectrum Control™

Figure 3-28 shows the three common options to manage IBM Spectrum Virtualize.

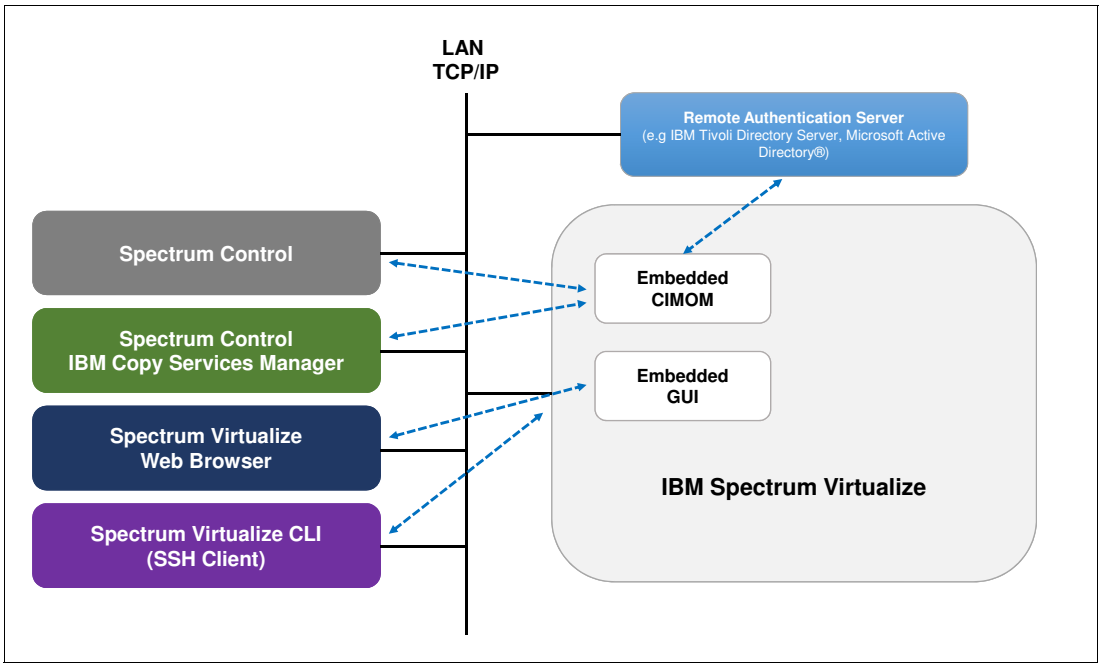

Figure 3-28 Common management options for IBM Spectrum Virtualize

You have full management control of the IBM Spectrum Virtualize cluster regardless of which method and *superuser* credentials that you use. IBM Spectrum Control is a robust software product with various functions (including performance and capacity features) that must be purchased separately.

If you decide to manage your IBM Spectrum Virtualize cluster with the IBM Storwize CLI, it does not matter whether you are using Secure Shell (SSH) access or IBM Spectrum Control server because the IBM Spectrum Virtualize CLI is on the cluster and accessed through SSH, which can be installed anywhere.

To access the IBM Spectrum Virtualize management GUI, direct a web browser to the system management IP address.

3.7 Post-installation tasks

Note: To ensure that your web browser is supported and has the appropriate settings enabled, see IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/en/STPVGU_7.7.0/com.ibm.storage.svc.console.770.doc/svc_configuringbrowser_1obg15.html

Although the following steps are optional to complete the IBM Spectrum Virtualize cluster configuration, it is a preferred practice to complete them at some point during the cluster implementation phase. For more information about these procedures, see *Implementing the IBM System Storage SAN Volume Controller with IBM Spectrum Virtualize V7.6*, SG24-7933, where these steps are described in detail along with the main functions of IBM Spectrum Virtualize:

1. Configure the SSH keys for the command-line user.
2. Rename nodes by using your site-naming convention.
3. Clearly label each individual physical node to avoid confusion should they require later identification.
4. Configure user authentication and authorization.
5. Set up event notifications and inventory reporting.
6. Back up configuration data.

When done, you can continue with your storage plan to provision and manage storage in your environment:

1. Create the storage pools.
2. Add managed disks (MDisks) to the storage pool.
3. Identify and create volumes.
4. Create host objects.
5. Map volumes to hosts.
6. Identify and configure the FlashCopy mappings and Metro Mirror relationships.
7. Set up a schedule to back up configuration data.

4

Event notification and troubleshooting

IBM Spectrum Virtualize software only is similar to an IBM Spectrum Virtualize installation with IBM SAN Volume Controller nodes (such as the 2145-DH8). There are some unique aspects of the day-to-day operation and infrequently run procedures of the product that differ from that of IBM SAN Volume Controller, which is detailed in this chapter.

This chapter provides the following information:

4.1 Event notification

An important difference between the IBM SAN Volume Controller nodes and nodes that are used for IBM Spectrum Virtualize software only is that IBM SAN Volume Controller has a greater level of server hardware monitoring and event notification.

IBM SAN Volume Controller uses the *call home* feature to automatically contact IBM for hardware and software alerts. However, IBM Spectrum Virtualize software only uses this feature to automatically contact IBM for software alerts only. The software alerts range from "out of resource" type errors, for example, disk full, log full, insufficient other internal resources, to WWNN-related errors. IBM Spectrum Virtualize software only uses the call home feature for approximately 15 alerts. For more information, see IBM Knowledge Center:

https://www.ibm.com/support/knowledgecenter/STVLF4_7.8.1/spectrum.virtualize.781.doc/svc_dataandeventnotifications_3qedwd.html

Call home alerting is set up when the IBM Spectrum Virtualize cluster is first configured (for more information, see 3.4.2, "System setup tasks for IBM Spectrum Virtualize" on page 38). These details are also used when you send a notification by email. An added feature of IBM Spectrum Virtualize V7.7.1 and later is that it provides a link to the latest IBM Spectrum Virtualize management information base (MIB) file when you configure SNMP server details. This feature is shown in Figure 4-1, where you can click the **Download MIB** option to receive the latest MIB file for IBM Spectrum Virtualize.

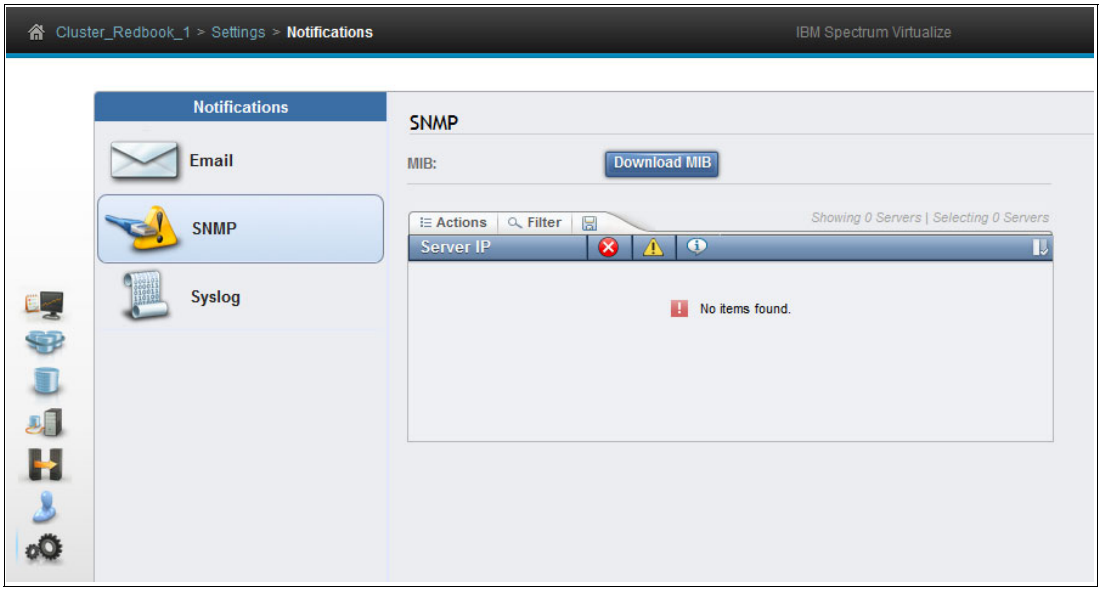

Figure 4-1 Download MIB option when you configure SNMP

The MIB file details, when used, allow the SNMP server to qualify messages that are sent from the IBM Spectrum Virtualize cluster. The predefined messages that are sent are classified into *error*, *warning,* and *informational* type traps. The SNMP server can alert support teams, depending on the trap type.

For more information about how to set up notifications, including call-home email notification that uses Simple Network Mail Protocol (SNMP), see *Implementing the IBM System Storage SAN Volume Controller with IBM Spectrum Virtualize V7.6*, SG24-7933.

4.2 The IBM Spectrum Virtualize Service Assistant Tool

IBM Spectrum Virtualize and IBM SAN Volume Controller have a tool for performing service tasks on individual cluster nodes. The function is called the Service Assistant Tool and requires you to enter the superuser password during login. You can service a node through an Ethernet connection by using a web browser to access a GUI interface. You can access the Service Assistant Tool with your browser by entering the service address IP of a specific node appended with `/service`. Alternatively, you can use the technician port, as described in “Replacing a node” on page 56.

Here are the browser-based methods to access the IBM Spectrum Virtualize Service Assistant Tool:

Technician port `https://install/service/node/home.action`
Service IP address `https://`*`node_service_ip_address`*`/service`

Figure 4-11 on page 59 shows an example of the Home window of the IBM Spectrum Virtualize Service Assistant Tool.

4.3 Replacing an IBM Spectrum Virtualize node

This section describes how to replace an IBM Spectrum Virtualize node that is already a member of an IBM Spectrum Virtualize cluster. The node might need to be replaced for either a hardware upgrade or a fault. This section describes replacing a node with identical hardware and components nondisruptively, and where hardware differences might cause a disruptive replacement.

4.3.1 Replacing nodes nondisruptively

This procedure is nondisruptive because changes to your SAN environment are not required. The replacement new node uses the same worldwide node name (WWNN) as the node that you are replacing. The resulting worldwide port names (WWPNs) for the HBAs are identical to that of the IBM Spectrum Virtualize node that is being replaced. Hence, it is not necessary to undertake SAN work by modifying aliases and zoning across all attached fabrics. An alternative to this procedure is to replace nodes disruptively either by moving volumes to a new I/O group or by rezoning the SAN. However, the disruptive procedures require more work on the hosts, SAN fabric, or both.

Some loss of system performance might occur when the nodes are being replaced. This loss can occur because volumes that managed by the I/O group that contains the node to be replaced become degraded when one of the nodes is shut down at the start of this procedure. The volumes remain degraded until both new and existing nodes of the I/O group are fully functioning.

Prerequisites

This task assumes that the following conditions are met:

- The replacement node must be identical in hardware terms and components to the original, slot positions, and able to operate at the Fibre Channel (FC) or Ethernet connection speed of the node that it is replacing.
- An original bootable USB key must be used or new bootable USB key with same IBM Spectrum Virtualize version level as the existing cluster.

- The WWNN of the node to be replaced is recorded.
- All nodes that are configured in the system are present and online.
- All errors in the system event log are addressed and marked as fixed.
- There are no volumes, managed disks (MDisks), or external storage systems with a status of degraded or offline.
- You backed up the system configuration and saved the `svc.config.backup.xml` file.
- The FC device driver on each FC attached host should be set to time out a missing fibre path in 3 seconds or less. If it is not practical to check the parameters of the FC driver on each host, then you must reboot the new Spectrum Virtualize node shortly after it is added to the system so that the fibre paths to it stops long enough to ensure that they are recovered properly when the node is active.

Do not continue this task if any of the conditions that are listed are not met unless you are instructed to do so by IBM Remote Technical Support.

It is important to review all of the steps that follow before you proceed with this task:

1. Do not continue this task if you are not familiar with IBM Spectrum Virtualize or IBM SAN Volume Controller environments or the procedures that are described in this task.
2. If you plan to reuse the node that you are replacing, ensure that the WWNN of the node is set to a unique number on your SAN. If you do not ensure that the WWNN is unique, the WWNN and WWPN are duplicated in the SAN environment and can cause problems if the node is reintroduced to the SAN.
3. Change the node ID and possibly the node name during this task. After the system assigns the node ID, the ID cannot be changed. However, you can change the node name after this task is complete.

Removing a node from the cluster

When the node is free to be removed from the cluster, complete the following steps to remove it:

1. As the administrator, open the IBM Spectrum Virtualize GUI, select *Cluster_Name*, and click **Monitoring** → **System** to show the cluster nodes.
2. Hover the cursor over the node to be removed and right-click it to show the actions, as shown in Figure 4-2 on page 55. The action box can also be used to verify the procedures that are described in 4.3.1, “Replacing nodes nondisruptively” on page 53.

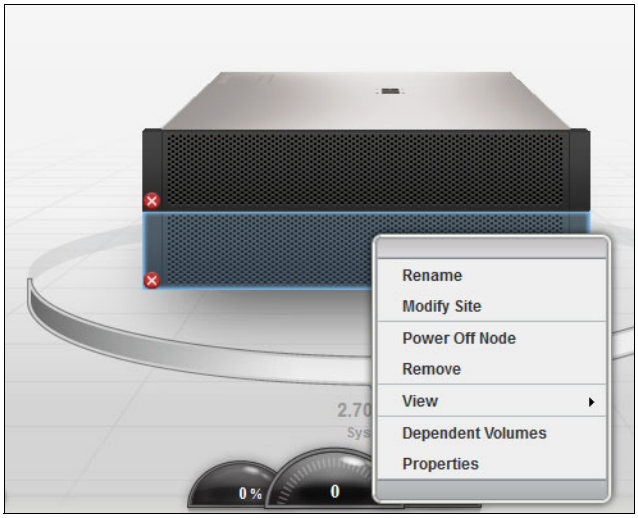

Figure 4-2 Remove a node from the cluster

3. When the **Remove** option is selected, a warning message is displayed, as shown in Figure 4-3.

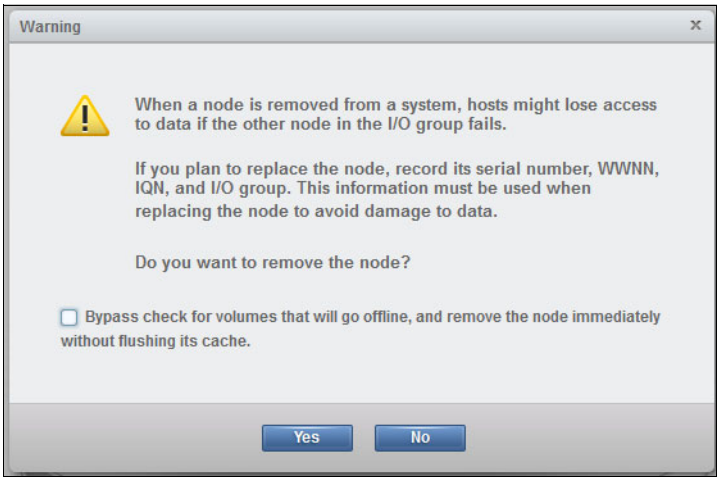

Figure 4-3 Warning regarding removing the node

Select **Yes** to remove the node. A completion message is displayed, as shown in Figure 4-4.

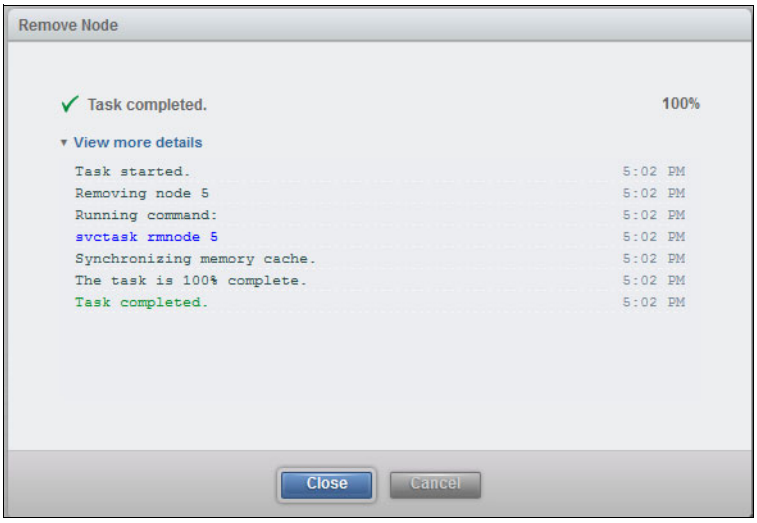

Figure 4-4 Node remove task completed

The node is removed successfully, as shown in Figure 4-5.

Figure 4-5 Node successfully removed from the cluster

Replacing a node

Although IBM Spectrum Virtualize software only has the same functions as IBM SAN Volume Controller, the node replacement procedure has an additional initial step, which is to load the IBM Spectrum Virtualize software onto the new hardware and obtain the IBM Spectrum Virtualize Node Activation Key file. This procedure is described in 3.3.1, “Completing the initial node configuration” on page 33.

When the node hardware has IBM Spectrum Virtualize software only installed with the Spectrum Node Activation Key successfully added, the node is ready to be added to the cluster. As it is replacing a node, there are many configuration changes to be made. This is achieved by connecting to the node through the technician port and accessing it through the IBM Spectrum Virtualize Service Assistant Tool.

Access and use of the technician port is described on 3.3, “Connecting to the technician port” on page 33.

Complete the following steps to replace a node:

1. Connect to the technician port and access the IBM Spectrum Virtualize Service Assistant Tool by entering the following address of the browser on your computer:

   `https://install/service/node/home.action`

   When you are connected, you are prompted for the superuser password, as shown in Figure 4-6. The default password is “passw0rd” (the sixth character is a *zero*). This password changes to the cluster superuser password when the node is added to the cluster.

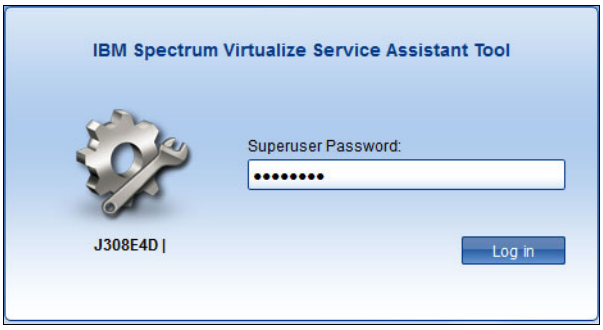

Figure 4-6 Password prompt

Figure 4-7 shows the Home view of the IBM Spectrum Virtualize Service Assistant Tool with the WWNN highlighted.

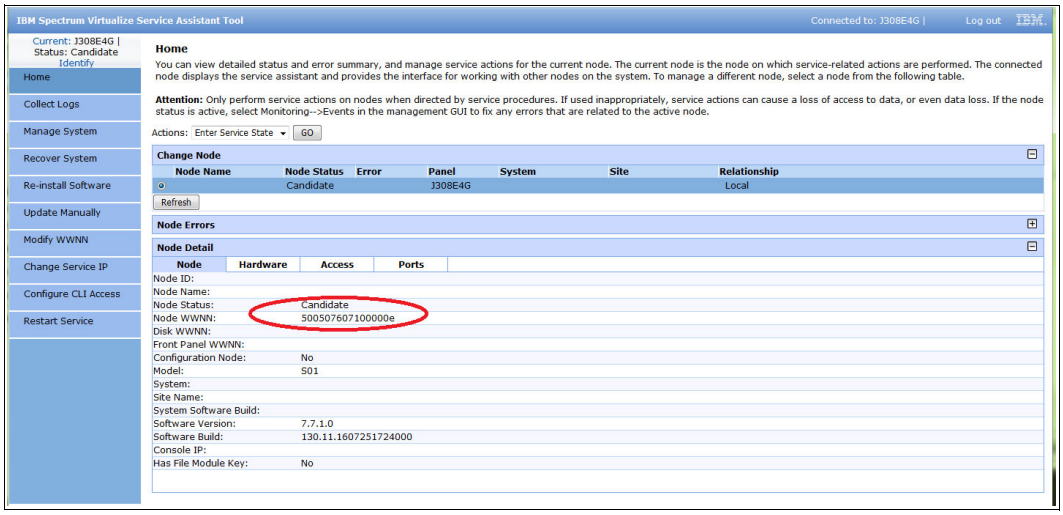

Figure 4-7 IBM Spectrum Virtualize Service Assistant Tool

The WWNN must be changed to update the correct WWPNs. The example WWNN that is marked in Figure 4-7 on page 57 is 500507607100000e, which must be changed. If you click the **Ports** tab just above the marked WWNN, you can view the WWPN, as highlighted in Figure 4-8.

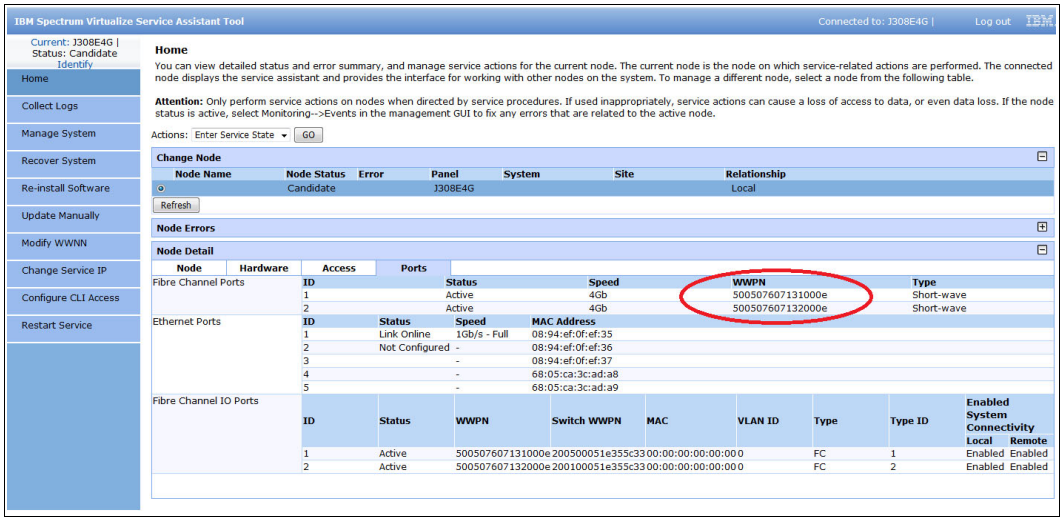

Figure 4-8 WWPNs shown

2. To change the WWNN, click **Modify WWNN** in the left column. You can edit the WWNN, as shown in Figure 4-9.

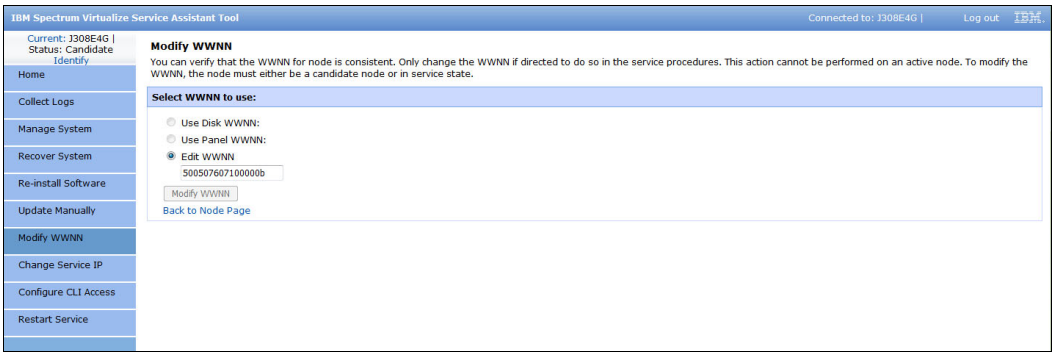

Figure 4-9 Modify WWNN

3. When the WWNN is changed, click **Modify WWNN** to continue. In this example, the WWNN is changed to 500507607100000b, which is the WWNN of the node that is to be replaced.

> **Note:** If the node that you are replacing can be reused, ensure that the WWNN of the node is changed and set to a unique number. If you do not do this task, the WWNN and WWPN are duplicated in the SAN environment if it is added later, which causes errors.

When the WWNN is updated, the node refreshes and you are returned to the superuser password prompt, as shown in Figure 4-6 on page 57.

4. Log in again with the superuser password where the Home view of the IBM Spectrum Virtualize Service Assistant Tool is shown again (Figure 4-7 on page 57). Check for the updated WWNN.

5. Change the service IP address by clicking the **Change Service IP** tab on the left where you can add the original node service IP address, as shown in Figure 4-10.

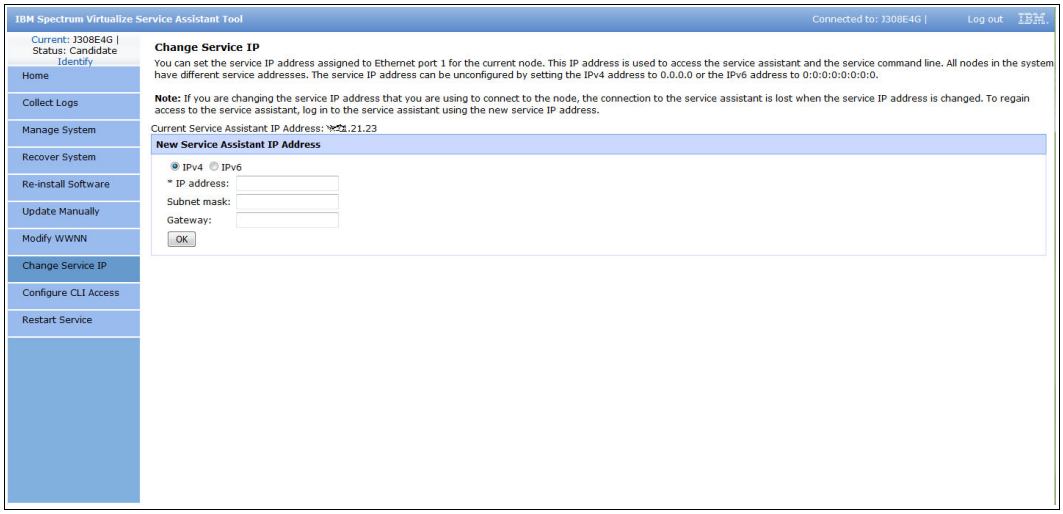

Figure 4-10 Change Service IP address

The service address IP allows you to view the node before and after the node is added to the cluster. The replacement node is now a candidate node and can be added to the cluster with the original node credentials of WWNN and WWPN. This can be checked from the Ports tab of the Home menu, as shown in Figure 4-11. In this example, the WWPNs are changed and they all have the suffix 000b.

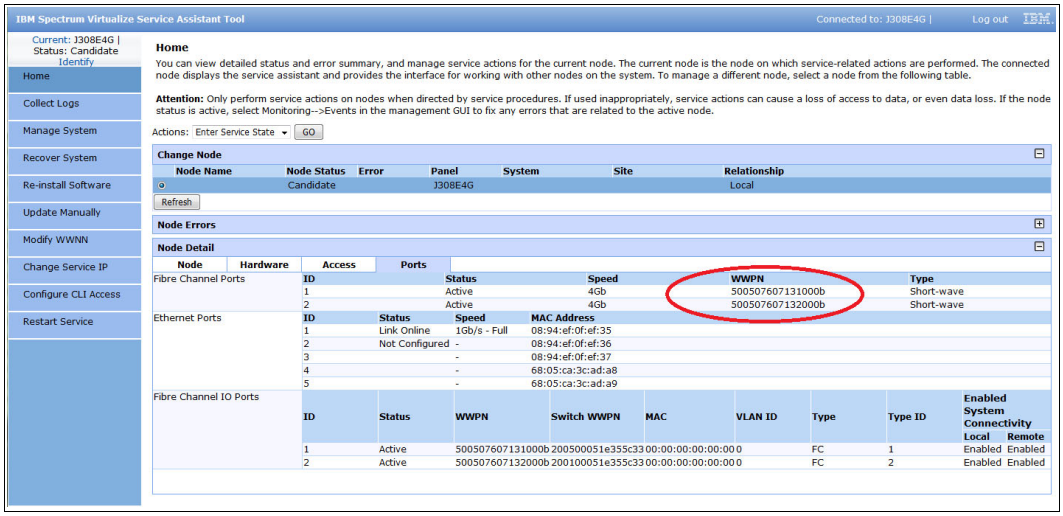

Figure 4-11 Changed WWPNs

6. To add the node to the cluster, for the instructions in 3.5, “Adding a node to an existing IBM Spectrum Virtualize Cluster” on page 45.

4.3.2 Replacing nodes disruptively

This process involves the WWPN needing to be changed along with the WWNN that is detailed in Figure 4-11 on page 59. The process that is described in 4.3.1, “Replacing nodes nondisruptively” on page 53 is repeated, where the resulting WWPNs are not the same as the original WWPNs that are used on the original node because the new node is not identical to the original node that is being replaced. This might include difference in host bus adapters (HBAs), slot placements, or a different type of node altogether (hardware upgrade).

The disruption is caused by the necessity of changing the zoning on fabrics connecting the IBM Spectrum Virtualize node and the hosts and storage controllers. This might mean that all hosts and storage controller zones accessing the node on the fabrics must be changed, or it might mean a smaller change, such as a few alias definitions. The SAN fabric reconfiguration that is involved with this change requires SAN fabric planning and an administrator, and, depending on the fabric complexity, might require a considerable amount of work.

If there are differences with the WWPN for the new nodes, you can use the IBM Spectrum Virtualize Service Assistant Tool (for example, the Ports tab in the Home menus that is shown in Figure 4-11 on page 59 can help you identify new WWPNs).

4.4 Identifying a node by activating the LED

IBM Spectrum Virtualize software only is compliant with Intelligent Platform Management Interface (IPMI). This is a set of hardware level interface specifications to monitor and manage host system capabilities independently of CPU, firmware, and operating system. The standardized interface and protocol allows systems management software a level of hardware control. It provides system administrators remote access to perform hardware-based functions. This access includes the ability to monitor platform status (such as system temperatures, voltages, fans, and power supplies) and to issue requests from a remote console for system power-down and rebooting. Most hardware vendors also comply with IPMI, which allows IBM Spectrum Virtualize software only to use some of the IPMI functions.

One such area that IBM Spectrum Virtualize uses is from the Service Assistant Tool home browser window to identify an IBM Spectrum Virtualize node. Using IPMI, it activates LED lights on the hardware to help you physically identify the node that is to be addressed.

Figure 4-12 on page 61 highlights the Identify option in the Service Assistant Tool browser display.

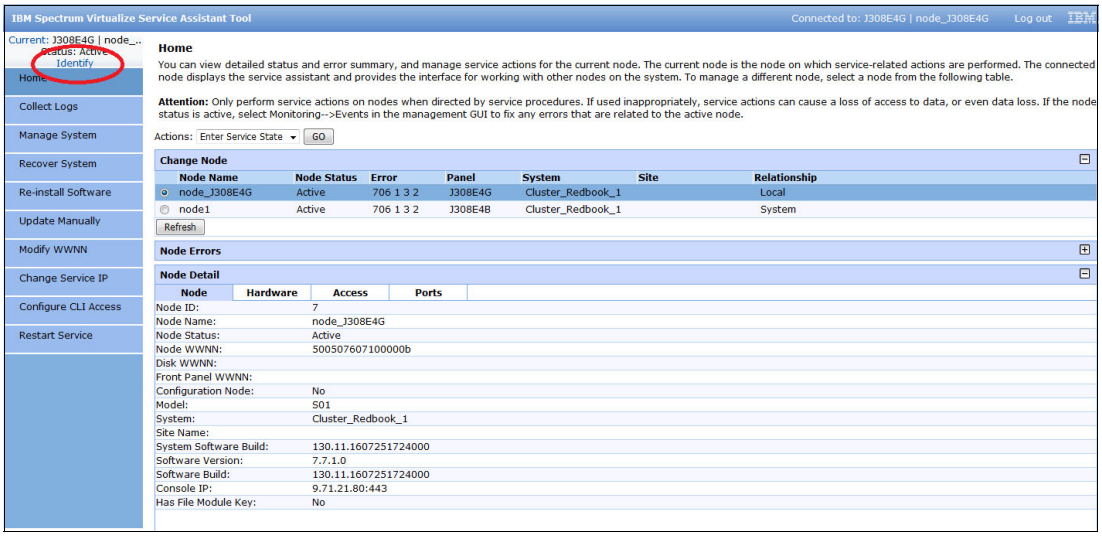

Figure 4-12 Identify and activate the LED on a node

Selecting the option lights the front LED of the physical IPMI-compliant node hardware so that you can identify it from other hardware and make physical changes. This feature is important because the IBM Spectrum Virtualize node hardware might look identical to all the surrounding hardware in the data center racks. There are no identifying features such as a machine type or serial number that could otherwise show where a specific node is. Without this feature, you might need to rely on the accuracy of data center documentation and diagrams to determine the correct hardware.

When you click the **Identify** option, the Service Assistant Tool lights the hardware LED, which you can use to identify the node in the data center. Figure 4-13 confirms that the LED is lit.

Figure 4-13 Node identity LED confirmation

To turn off the LED, click **Turn Off** in Figure 4-13. Completion is confirmed by the message that is shown Figure 4-14. Click **OK** to complete the task.

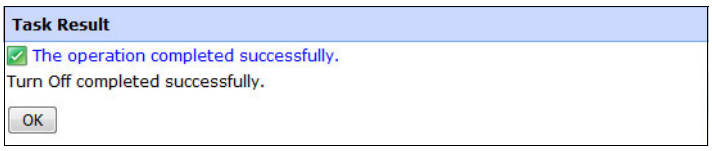

Figure 4-14 LED identify turned off successfully

A

Lenovo System x3650 M5 server implementation

This appendix describes the implementation of IBM Spectrum Virtualize software only on a Lenovo System x3650 M5 server (E5-2600 v3), Machine Type 5462. This server platform is used in this book.

This appendix provides the following information:

x3650 M5 server introduction

The Lenovo System x3650 M5 is the next generation of server beyond that used in the IBM SAN Volume Controller 2145-DH8 nodes. At the time of writing, there is one supported x3650 M5 Machine Type (M/T): the 5462 with the Intel Xeon E5-2600 v3 CPU family (known as “Haswell”).

The most notable difference between this server and the 2145-DH8 nodes is the absence of a built-in UPS in this product. In addition, the server has an LCD for status messages, and the DH8 uses LED status indicators.

The x3650 M5 is a 2U server that accepts up to two Intel XeonE5-2600 v3 processors (depending on the M/T), 24 DIMM slots, up to 8 PCIe 3.0 expansion slots, 8 - 16 drive bays (depending on the chassis that is chosen), and four built-in Gigabit Ethernet (GbE) ports. It features an Integrated Management Module (IMM) for remote monitoring and management.

Note: IBM Spectrum Virtualize software only requires specific features, functions, and options of the x3650 M5 server. See the IBM System Storage Interoperation Center (SSIC) for information about the current supported configuration:

https://www.ibm.com/systems/support/storage/ssic/interoperability.wss

For more information about the Lenovo x3650 M5 server with v3 CPUs, see the following website:

https://lenovopress.com/tips1193-lenovo-system-x3650-m5

x3650 M5 front panel

Figure A-1 shows the front panel of the x3650 M5 that is described in this appendix.

Figure A-1 x3650 M5 front panel

The front panel has the following key features:

- Drive slots. In Figure A-1 on page 64, the two SAS slots on the far left are used for the node boot drives. You must use the two far-left slots.

  Attention: Although your server might have hot-swap drive slots, IBM Spectrum Virtualize software only does not support this feature. Use IBM Spectrum Virtualize to power down a node before you replace a failed drive.

- Power button. As with most servers, press this button once to turn on the power and press it again to send a power-off request to the software. (The normal process to power down the node is to send a request through the software management interfaces, but the button works. However, an error event is logged when the server restarts.) You can hold the button to force an immediate power-off, but this method is not recommended.
- USB slots. These slots are used for loading the software image. The server has three USB ports; the one on the left is a USB 3.0 Super-Speed slot and the two slots on the right are USB 2.0 slots. The software installs from any of the three ports.
- LCD display. This display shows basic server status information. Unlike the display on some IBM SAN Volume Controller nodes, it does not display information specific to IBM Spectrum Virtualize software only, such as cluster and node errors or the cluster IP or the WWPN.
- VGA port. There is a VGA port on the front panel if you purchase M/T 5462 or select feature ATE9 with M/T 8871.
- Although there is a location to install an optical drive on some models, an optical drive is not used with IBM Spectrum Virtualize software only.

x3650 M5 rear panel

Figure A-2 shows the rear panel of the x3650 M5 server.

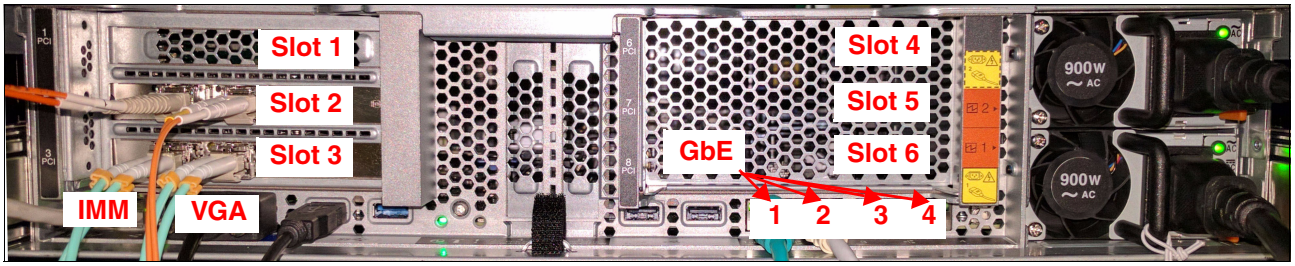

Figure A-2 x3650 M5 rear panel

The rear panel has the following key features:

- IMM port. The IMM port is used to manage the node hardware. You can use either a web interface or the management software of your choice for detailed unit monitoring. (The IBM Spectrum Virtualize software does some limited hardware monitoring; see Chapter 4, “Event notification and troubleshooting” on page 51.)
- VGA port. The VGA port is used for the initial node setup.

- Gigabit Ethernet (GbE) ports. The first GbE port (which is also labeled “IMM”) is used as the technician port, primarily during initial node setup (see 3.3, “Connecting to the technician port” on page 33 for more details.) This port is also used for certain node maintenance actions.
- PCIe Slots. Figure A-3 shows the placement of the PCI slots.

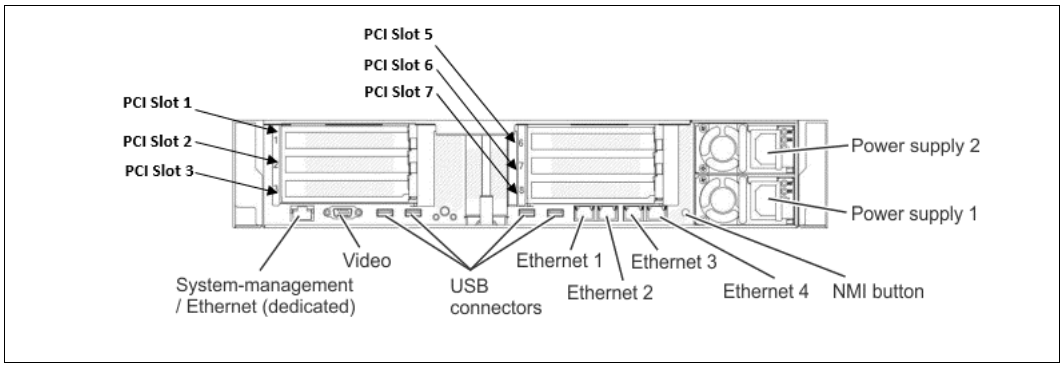

Figure A-3 PCI slot configuration

x3650 M5 sample configuration

The following x3650 M5 features are used for the examples in this publication:

- Base x3650 M5 (M/T 5462) System Unit
- One Intel Xeon E5-2667 v3 8C 3.2 GHz CPU
- Four 16 GB 2133 MHz RDIMMs
- ServeRAID M5210 SAS/SATA Controller (Feature Code A3YZ, which does not contain either the RAID Upgrade or Cache)
- Two 300 GB 10k RPM 2.5” SAS disks
- Two 900 W Power Supply Units
- One 8-slot 2.5” Hot-Swap HDD Assembly Kit (Feature Code A5G6, described as “Single RAID”)
- One 3-slot PCIe x8 Riser
- One Dual-Port Emulex 16Gb Fibre Channel (FC) HBA
- One Intel X710-DA2 Dual-Port 10GbE SFP+ Network Adapter

Power consumption

To calculate the power consumption of a node (which is necessary for uninterruptible power supply (UPS) sizing), Lenovo provides a power consumption calculator that is called the Power Configurator, found at the following website:

`https://support.lenovo.com/us/en/documents/LNVO-PWRCONF`

The calculator takes as an input your installed options, and then produces a power consumption figure.

Note: When you are using the Power Configurator, set the slider for CPU utilization to 100%. Although the cluster does not frequently run the CPUs at that level of utilization, you must plan for the worst case to ensure that the nodes can complete their shutdown procedure if power loss occurs.

More information

To access the complete vendor documentation for the x3650 M5, see the following website:

http://support.lenovo.com/us/en/products/servers/lenovo-x86-servers/lenovo-system-x3650-m5?tabName=Documentation

B

Supermicro SuperServer 2028U-TRTP+ server implementation

This appendix describes the implementation of IBM Spectrum Virtualize Software only on a Supermicro SuperServer 2028U-TRTP+ server. IBM Spectrum Virtualize Software version 7.8.0.2 or later is required.

This appendix provides the following information:

Supermicro SuperServer 2028U-TRTP+ server introduction

The Supermicro SuperServer 2028U-TRTP+ implementation is a traditional server-based implementation of IBM Spectrum Virtualize. At the time of writing, the only supported SuperServer 2028U-TRTP+ is M/T SYS-2028U-TRTP+, which supports both Intel Xeon E5-2600 v3 (Haswell) and v4 (Broadwell) CPU families.

This server differs from 2145-DH8 and 2145-SV1 implementations because it does not include a built-in UPS solution. In addition, this server has minimal LED notification capability, compared to the SVC appliance.

The SYS-2028U-TRTP+ is a 2U server that accepts up to 2 Intel Xeon E5-2600 v3 or v4 processors, 24 DIMM slots, up to 7 PCIe 3.0 expansion slots, 24 drive bays, and two built-in 10 Gbps SFP+ Ethernet ports. It also features a dedicated BMC management port for remote monitoring and management.

Note: IBM Spectrum Virtualize Software only requires specific features, functions, and options of the SuperServer 2028U-TRTP+ server. For information about the current supported configuration, see the IBM System Storage Interoperation Center (SSIC):

https://www.ibm.com/systems/support/storage/ssic/interoperability.wss

For more information about the SYS-2028U-TRTP+ server, see the following website:

https://www.supermicro.com/products/system/2U/2028/SYS-2028U-TRTP_.cfm

SYS-2028U-TRTP+ front panel

Figure B-1 shows the SYS-2028U-TRTP+ front panel.

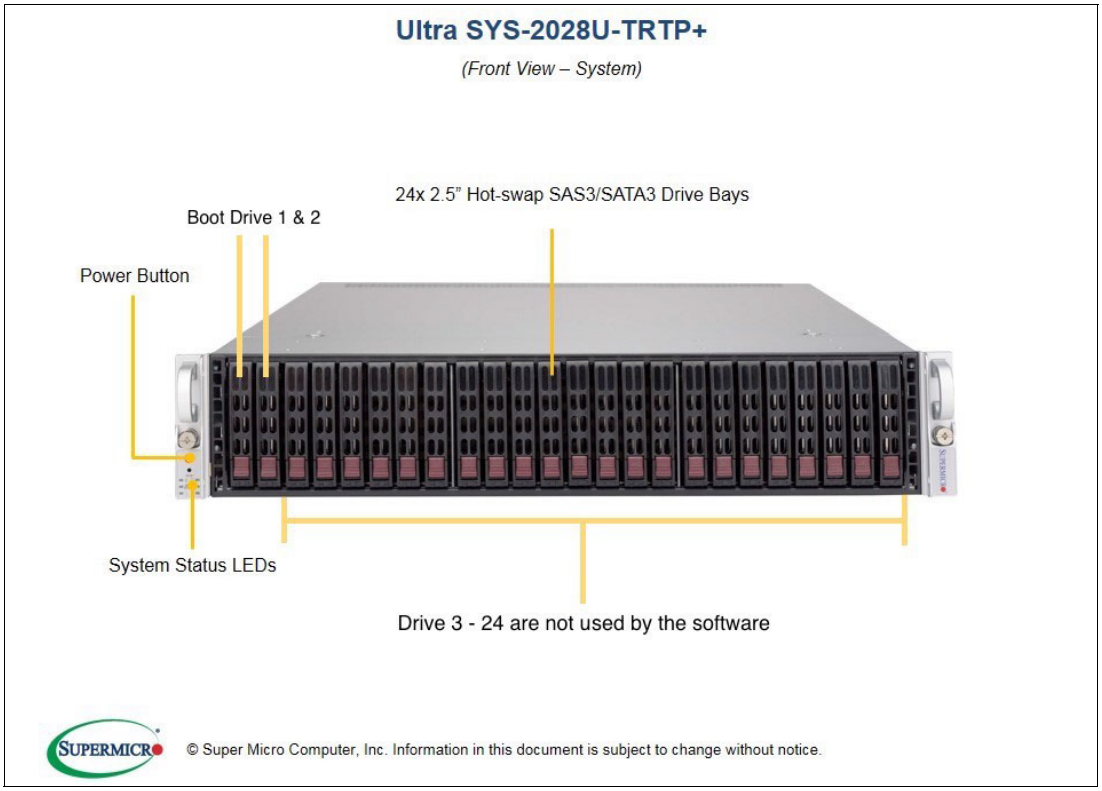

Figure B-1 SYS-2028U-TRTP+ front panel

The front panel has the following components:

- *Drive slots*. In Figure B-1, the two drive slots on the far left of the server are used for the node boot drives. You must use the two far-left slots (0 and 1). For more information about supported drives, see the following website:

  `http://www.ibm.com/support/docview.wss?uid=ssg1S1009260`

  Attention: Although your server might support hot-swappable drives, IBM Spectrum Virtualize Software only does not support this feature. Use the IBM Spectrum Virtualize software to power down a node before you replace a failed drive.

  In Figure B-1, the two leftmost slots are used. No other drives are used.

- *Power button*. As with most servers, press this button once to turn on the power and press it again to send a power-off request to the software. (The normal process to power down the node is to send a request through the software management interfaces, but the button works. However, an error event is logged when the server restarts.)

  Attention: You can hold the button to force an immediate power-off. However, this method is *not* recommended and might cause data loss.

- *System status LEDs*. The system status LEDs show basic server status information. These LEDs show the current power state of the server, identify the server through use of the

built-in command in IBM Spectrum Virtualize, and provide fault notification if a fault is found by means of the BMC interface. (The fault LED is not controlled by IBM Spectrum Virtualize software, and it does not illuminate if a software or fabric fault occurs.) For more information, see IBM Knowledge Center:

https://www.ibm.com/support/knowledgecenter/en/STVLF4_7.8.1/spectrum.virtualize.781.doc/svirt_supermicro.html

SYS-2028U-TRTP+ rear panel

Figure B-2 shows the rear panel of the SYS-2028U-TRTP+.

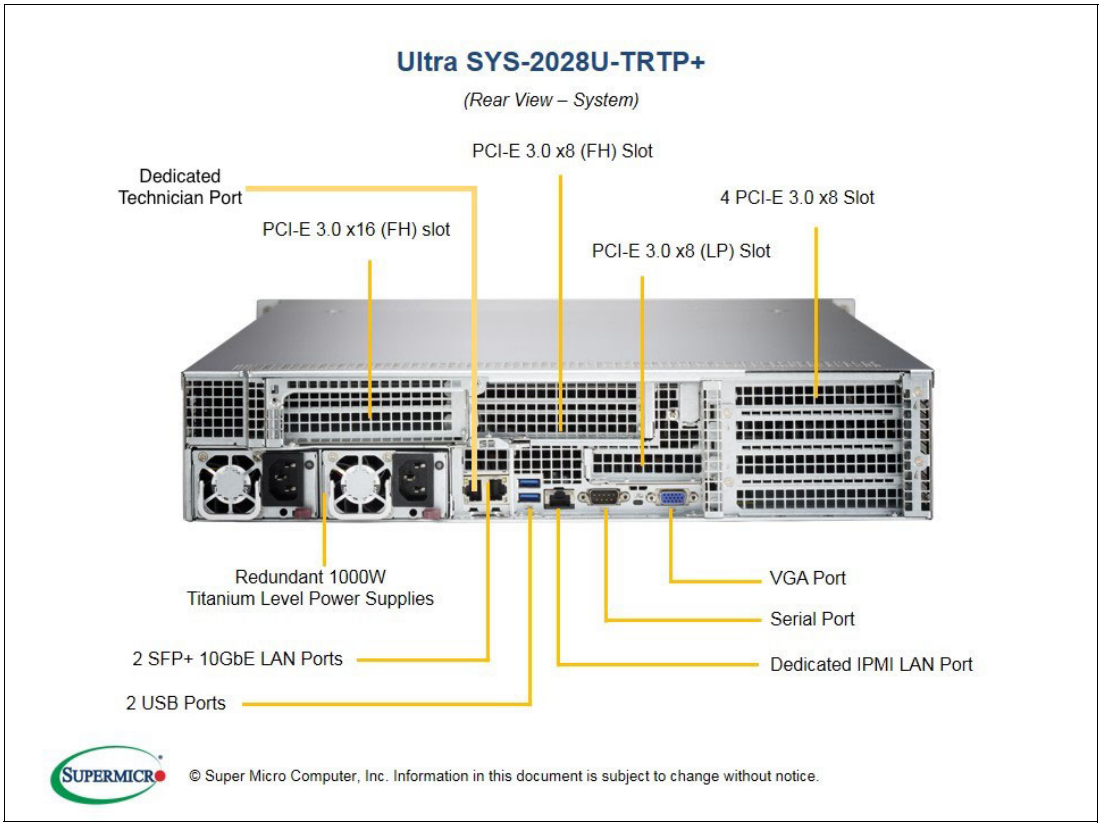

Figure B-2 SYS-2028U-TRTP+ Rear Panel

The rear panel has the following components:

- *IPMI LAN Port*. The IPMI LAN port is used by the low-level server firmware to enable management of the server hardware platform directly. This task can be done either directly with a web interface, or with appropriate server management software of your choice. (IBM Spectrum Virtualize software does some limited hardware monitoring. For more information, see Chapter 4, “Event notification and troubleshooting” on page 51.)
- *VGA Port*. For a direct console connection, the VGA port can be used to provide superuser access to the server directly with a keyboard and screen that are physically attached to the server.
- *SFP+ 10GbE Ports*. The server comes equipped with 2 SFP+ 10GbE Ports, which are used for the technician port and management port for IBM Spectrum Virtualize software. SFP modules are required to use these ports, and are sold separately from the server.

- *USB ports*. There are 2 USB 3.0 ports on the rear of the chassis. These can be used for peripheral or USB stick connection. If you use software encryption, you can use these ports to connect encryption keys to the system.
- *PCI-E slots*. There are 7 PCI-E 3.0 slots on the rear of the chassis. Their layout is shown in Figure B-2. For more information, see the following website:

  https://www.ibm.com/support/knowledgecenter/STVLF4_7.8.1/spectrum.virtualize.781.doc/svirt_supermicro.html

SYS-2028U-TRTP+ sample configuration

The following SYS-2028U-TRTP+ features are used for the examples in this publication:

- Base SYS-2028U-TRTP+ System Unit
- Two Intel Xeon E5-2667 v4 8C 3.2 GHz CPU
- Four 16 GB 2133 MHz RDIMMs
- Two 400 GB SATA SSD disks
- Two 1000 W Power Supply Units
- Two 10GbE SFP+ Modules
- One Intel X710-DA2 Dual-Port 10GbE SFP+ Network Adapter
- Two Emulex Lancer G6 4-port 16 GB FC Adapters

For more information, see the following website:

http://www.ibm.com/support/docview.wss?uid=ssg1S1009260

More information

To access the complete vendor documentation for the SYS-2028U-TRTP+, see the following website:

https://www.supermicro.com/products/system/2U/2028/SYS-2028U-TRTP_.cfm

Related publications

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

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Some publications that are referenced in this list might be available in softcopy only.

- *IBM System Storage SAN Volume Controller and Storwize V7000 Best Practices and Performance Guidelines*, SG24-7521
- *Implementing the IBM System Storage SAN Volume Controller with IBM Spectrum Virtualize V7.6*, SG24-7933
- *Implementing the IBM Storwize V7000 and IBM Spectrum Virtualize V7.6*, SG24-7938

You can search for, view, download, or order these documents and other Redbooks, Redpapers, web docs, draft and additional materials, at the following website:

ibm.com/redbooks

Other publications

This publication is also relevant as a further information source:

- *Lenovo System x3650 M5 (E5-2600 v3) Product Guide*, found at:

  https://lenovopress.com/tips1193-lenovo-system-x3650-m5

Online resources

This website is also relevant as a further information source:

- IBM System Storage Interoperation Center (SSIC):

  https://www.ibm.com/systems/support/storage/ssic

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

REDP-5392-01

ISBN 0738456101

Printed in U.S.A.