

Implementation guide for IBM Spectrum Virtualize for Public Cloud Version 8.3

Angelo Bernasconi

Eric Goodall

Jackson Shea

Jimmy John

Jordan Fincher

Nicolo Lorenzoni

Pierluigi Buratti

Vasfi Gucer



 **Cloud**

Storage



International Technical Support Organization

**Implementation guide for IBM Spectrum Virtualize for
Public Cloud Version 8.3**

June 2019

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

Second Edition (June 2019)

This edition applies to IBM Spectrum Virtualize for Public Cloud Version 8.3.

This document was created or updated on May 11, 2020.

© Copyright International Business Machines Corporation 2017, 2019. All rights reserved.

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

Contents

Notices	vii
Trademarks	viii
Preface	ix
Authors	ix
Now you can become a published author, too	xi
Comments welcome	xii
Stay connected to IBM Redbooks	xii
Chapter 1. Introduction	1
1.1 Introduction to IBM Spectrum Virtualize for Public Cloud	2
1.1.1 The evolution of IBM SAN Volume Controller	2
1.2 IBM Spectrum Virtualize for Public Cloud	4
1.2.1 Primers of storage virtualization and software-defined-storage	4
1.2.2 Benefits of IBM Spectrum Virtualize for Public Cloud	5
1.2.3 Features of IBM Spectrum Virtualize for Public Cloud	7
1.2.4 IBM Spectrum Virtualize on IBM Cloud	9
1.3 Use cases for IBM Spectrum Virtualize for Public Cloud	11
1.3.1 Hybrid scenario: on-premises to IBM Cloud	12
1.3.2 Cloud-native scenario: Cloud to cloud	14
1.4 Licensing	14
Chapter 2. Solution architecture	17
2.1 IBM Cloud	18
2.2 Storage virtualization	19
2.3 IBM Spectrum Virtualize	19
2.3.1 Nodes	20
2.3.2 I/O groups	20
2.3.3 System	21
2.3.4 MDisks	21
2.3.5 Storage pool	22
2.3.6 Volumes	22
2.3.7 Cache	23
2.3.8 IBM Easy Tier	25
2.3.9 Hosts	25
2.3.10 Host cluster	26
2.3.11 iSCSI	26
2.3.12 IP replication	26
2.3.13 Synchronous or asynchronous remote copy	27
2.4 Environment used for this book	28
Chapter 3. Planning and preparation for the IBM Spectrum Virtualize for Public Cloud deployment	31
3.1 Provisioning cloud resources	32
3.1.1 Ordering servers	32
3.2 Provisioning IBM Cloud Block Storage	48
3.2.1 Cloud Block Storage overview	48
3.2.2 Provisioning Block Volumes	51
3.3 IBM Spectrum Virtualize networking considerations	57

3.3.1 Provisioning a Network Gateway Appliance	64
Chapter 4. Implementation	71
4.1 IBM Spectrum Virtualize for Public Cloud installation	72
4.1.1 Downloading the One-click installer	73
4.1.2 Fully Automated installation	74
4.1.3 Semi Automated installation	78
4.2 Configuring Spectrum Virtualize	82
4.2.1 Log in to cluster and complete the installation	82
4.2.2 Configure Cloud quorum.	92
4.2.3 Installing the IP quorum application	94
4.2.4 Configure the back-end storage	94
4.2.5 Configuring Call Home with CLI	96
4.2.6 Upgrading to second I/O group.	97
4.3 Configuring replication from on-prem IBM Spectrum Virtualize to IBM Spectrum Virtualize for IBM Cloud	105
4.4 Configuring Remote Support Proxy	115
Chapter 5. Typical use cases for IBM Spectrum Virtualize for Public Cloud	119
5.1 Whole IT services deployed in the Public Cloud.	120
5.1.1 Business justification.	120
5.1.2 Highly available deployment models.	121
5.2 Disaster recovery	124
5.2.1 Business justification.	125
5.2.2 Common DR scenarios with IBM Spectrum Virtualize for Public Cloud	126
5.3 IBM FlashCopy in the Public Cloud.	128
5.3.1 Business justification.	128
5.3.2 FlashCopy mapping	129
5.3.3 Consistency groups	130
5.3.4 Crash consistent copy and hosts considerations	130
5.4 Workload relocation into the Public Cloud	131
5.4.1 Business justification.	131
5.4.2 Data migration.	132
5.4.3 Host provisioning	134
5.4.4 Implementation considerations	134
Chapter 6. Supporting the solution	135
6.1 Who to contact for support	136
6.2 Working with IBM Cloud Support	137
6.3 Working with IBM Spectrum Virtualize Support.	138
6.3.1 Email notifications and the Call Home function.	138
6.3.2 Disabling and enabling notifications	142
6.3.3 Collecting Diagnostic Data for IBM Spectrum Virtualize	143
6.3.4 Uploading files to the Support Center.	147
6.3.5 Service Assistant Tool	148
6.3.6 Remote Support Assistance	150
Appendix A. Guidelines for disaster recovery solution in the Public Cloud.	157
Plan and design for the worst case scenario	158
Recovery tiers	158
Design for production	159
On-premise to DR cloud considerations	162
Cloud to DR cloud considerations.	163
Cloud to DR on-premises considerations	164

Common pitfalls	165
Plan and design for the best scenarios	165
Single-points-of-failure	165
Only have a Plan A	165
Poor DR testing methodology	166
Networking aspects	166
Five networks	166
DR test	167
Failover or emergency	168
Fallback	168
Full versus partial failover	168
User-to-DR server	169
Server-to-DR servers	169
Network virtualization	170
Network function virtualization	170
Bring Your Own IP	170
Related publications	173
IBM Redbooks	173
Online resources	173
Help from IBM	174

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.

AIX®	IBM Cloud™	Passport Advantage®
Aspera®	IBM FlashSystem®	PowerVM®
Bluemix®	IBM Resiliency Services®	Real-time Compression™
DB2®	IBM Spectrum™	Redbooks®
Easy Tier®	IBM Spectrum Control™	Redpaper™
FlashCopy®	IBM Spectrum Protect™	Redbooks (logo)  ®
HyperSwap®	IBM Spectrum Storage™	Storwize®
IBM®	IBM Spectrum Virtualize™	System Storage®

The following terms are trademarks of other companies:

SoftLayer, are trademarks or registered trademarks of SoftLayer, Inc., an IBM Company.

Intel, 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.

ITIL is a Registered Trade Mark of AXELOS Limited.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

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

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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

Preface

IBM® Spectrum Virtualize is a key member of the IBM Spectrum™ Storage portfolio. It is a highly flexible storage solution that enables rapid deployment of block storage services for new and traditional workloads, on-premises, off-premises and in a combination of both.

IBM Spectrum Virtualize™ for Public Cloud provides the IBM Spectrum Virtualize functionality in IBM Cloud™. This new capability provides a monthly license to deploy and use Spectrum Virtualize in IBM Cloud to enable hybrid cloud solutions, offering the ability to transfer data between on-premises private clouds or data centers and the public cloud.

This IBM Redpaper™ publication gives a broad understanding of IBM Spectrum Virtualize for Public Cloud architecture and provides planning and implementation details of the common use cases for this product.

This publication helps storage and networking administrators plan and implement install, tailor, and configure IBM Spectrum Virtualize for Public Cloud offering. It also provides a detailed description of troubleshooting tips.

IBM Spectrum Virtualize is also available on AWS. For more information, see *Implementation guide for IBM Spectrum Virtualize for Public Cloud on AWS*, REDP-5534.

Authors

This paper was produced by a team of specialists from around the world working at the International Technical Support Organization, Austin Center.




Angelo Bernasconi is an Executive Certified Storage and SDS IT Specialist for IBM Systems Italy. He has 30 years of experience in the delivery of professional services and solutions for IBM Enterprise customers in open systems. He holds a degree in Electronics, and his areas of expertise are the IBM Storage Portfolio and SDS, and design and implementation of storage solutions. He has written extensively about SAN and storage products in various IBM publications. He is also a member of the Italy SNIA committee and a member of the Italy TEC.




Eric Goodall is an IBM Cloud Solution Architect, working in the IBM Global Cloud Technical Sales and Solutioning with a focus on storage solutions on the IBM Bluemix® IaaS platform. In 35 years at IBM, Eric has worked in storage system and storage intensive application architecture design for federal, healthcare, and commercial customers. He currently provides cross business unit architecture support for IBM offering and cloud sales teams in worldwide geographies.



Jackson Shea is a level 2 certified IBM Information Technology Specialist/Architect, performing design and implementation engagements through Lab Services. He's been with IBM since April 2010, before which he was a Lead Storage Administrator with a large health insurance consortium in the Pacific Northwest, working with IBM equipment since 2002. He has had over 12 years of experience with Spectrum Virtualize, formerly known as the IBM SAN Volume Controller and related technologies. Jackson is based out of Portland, Oregon, and got his Bachelor of Science degree in Philosophy with minors in Communications and Chemistry from Lewis & Clark College. Jackson's professional focus has continued to be Spectrum Virtualize, but is conversant in storage area network design, implementation, extension, and storage encryption. In addition, he is trying to stay alert to new opportunities for storage to play a role in emerging IBM technologies such as Hyperledger, AI, and Quantum Computing. In his spare time, he is an avid video game enthusiast.



Jimmy John is IBM System Storage® Technical Advisor and SME for Spectrum Virtualize, SVC, IBM Storwize® products. He worked as Advisory Software Engineer /PFE for SVC and Product Focal for Storwize platform for 6+ years. His past assignments include as Operational Team Lead and PFE for IBM Modular Storage Products. He has been with IBM for 18 years at various positions. His current interests are Cloud Implementation and Analytics.



Jordan Fincher is a Product Field Engineer working in Storage Support at IBM. He received his Bachelor of Science in Information Security from Western Governor's University. Jordan first started his IBM career in 2012 as a Systems Engineer for the IBM Business Partner e-TechServices doing pre-sales consulting and implementation work for many IBM accounts in Florida. In 2015, Jordan started working in his current role as a Product Field Engineer for IBM Spectrum Virtualize storage products.



Nicolo Lorenzoni is an IBM Cloud Architect within the banking industry designing architectures for regulatory computing platforms to securely process regulated data. Nicolò focuses on the security and compliance domain, and has collaborated with IBM Systems to create the patterns for the secure deployment of software-defined-storage solutions capabilities into IBM Cloud. He received his M.S. in Engineering from Politecnico of Milano, Italy and Keio University, Japan. In 2017, Nicolò moved to IBM Austin, TX to drive cloud initiatives and activities with major global systemically important banks. He is also a member of IBM's Cloud Compliance Advisory Body, and a leader for Young Technical Expert Council Italy.



Pierluigi Buratti is an Executive IT Architect at IBM Italy – Resiliency Services, with more than 20 years spent in designing and implementing Disaster Recovery and Business Continuity solutions for Italian and European clients, using the latest technologies for data mirroring and data center interconnectivity. Since 2011, he is working in the Global Development of IBM Resiliency Services®, leading the design of DRaaS technologies and solutions.



Vasfi Gucer is Pierluigi Buratti an IBM Technical Content Services Project Leader with the Digital Services Group. He has more than 20 years of experience in the areas of Systems Management, networking hardware, and software. He writes extensively and teaches IBM classes worldwide about IBM products. His focus has been primarily on cloud computing, including cloud storage technologies for the last 6 years. Vasfi is also an IBM Certified Senior IT Specialist, Project Management Professional (PMP), IT Infrastructure Library (ITIL) V2 Manager, and ITIL V3 Expert.

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

Jon Tate, Debbie Willmschen, Erica Wazewski, Ann Lund
Digital Services Group, Technical Content Services

Michelle Tidwell, Mary J Connell
IBM USA

Robin Findlay
IBM UK

Henry E Butterworth, Long Wen Lan, Yu Yan Chen, Xiaoyu Zhang, Qingyuan Hou
IBM China

The team would like to express thanks to IBM Gold Partner **e-TechServices** for providing infrastructure as a service utilizing e-TechServices' cloud systems as a contribution to the development and test environment for the use cases covered in this book. Special thanks to Javier Suarez, Senior Systems Engineer, Marc Spindler, CEO, and Mario Ariet, President.

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:

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:

ibm.com/redbooks

- ▶ Send your comments in an email:

redbooks@us.ibm.com

- ▶ Mail your comments:

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>



Introduction

This chapter describes IBM Spectrum Virtualize implemented in a cloud environment and referred to as IBM Spectrum Virtualize for Public Cloud. A brief overview of the technology behind the product introduces the drivers and business values of using IBM Spectrum Virtualize in the context of public cloud services. It also describes how the solution works from a high-level perspective.

IBM Spectrum Virtualize Software only is available starting with IBM Spectrum Virtualize V7.7.1. This publication describes IBM Spectrum Virtualize for Public Cloud V8.1.1.

This chapter includes the following topics:

- ▶ 1.1, “Introduction to IBM Spectrum Virtualize for Public Cloud” on page 2
- ▶ 1.2, “IBM Spectrum Virtualize for Public Cloud” on page 4
- ▶ 1.3, “Use cases for IBM Spectrum Virtualize for Public Cloud” on page 11
- ▶ 1.4, “Licensing” on page 14

1.1 Introduction to IBM Spectrum Virtualize for Public Cloud

Companies are currently undergoing a digital transformation and taking architecture decisions that determine how their business is going to operate in the next couple of years. They recognize the value to deliver services via cloud, and the majority of them are already using public cloud to some degree. The role of the cloud is maturing, and more considered as a platform for innovation and business value. The cloud is a key enabler to drive transformation and innovation for IT agility and new capabilities.

Nevertheless, one of the challenges for these organizations is how to integrate those public cloud capabilities with the existing back-end. Organizations want to retain flexibility without introducing new complexity or requiring significant new capital investment.

Cloud integration can occur between different endpoints (cloud-to-cloud, on-premises to off-premises or cloud to non-cloud) and at different levels within the cloud stack: infrastructure layer, service layer and, for example, at the application layer or at the management one. Within the infrastructure as a service (IaaS) domain, storage layer integration is often the most attractive approach for ease of migration and replication of heterogeneous resources and data consistency.

In this sense, coming from the IBM Spectrum Storage™ family, IBM Spectrum Virtualize for Public Cloud supports clients in their IT architectural transformation and transition towards the cloud service model, enabling hybrid cloud strategies or, for cloud-native workload, providing the benefits of familiar and sophisticated storage functionality on the public cloud data centers, enhancing the existing cloud offering.

Running on-premises (on-prem), IBM Spectrum Virtualize software supports capacity built into storage systems, and capacity in over 400 different storage systems from IBM and other vendors. This wide range of storage support means that the solution can be used with practically any storage in a data center today and integrated with its counterpart IBM Spectrum Virtualize for Public Cloud, which supports IBM Cloud block storage offering in its two variants: Performance and Endurance storage options.

1.1.1 The evolution of IBM SAN Volume Controller

IBM Spectrum Virtualize represents the software engine of IBM SAN Volume Controller. The core software, when extracted from the IBM SAN Volume Controller storage appliance, is deployable as a software solution on general-purpose hardware. The software version further extended to include and support standardized public cloud infrastructure is represented by IBM Spectrum Virtualize for Public Cloud, and described in this publication.

IBM SAN Volume controller is based on an IBM project started in the second half of 1999 at the IBM Almaden Research Center. The project was called COMmodity PArts Storage System or COMPASS. However, most of the software has been developed at the IBM Hursley Labs in the UK. One goal of this project was to create a system that was almost exclusively composed of commercial off the shelf (COTS) standard parts. Yet, it had to deliver a level of performance and availability that was comparable to the highly optimized storage controllers of previous generations.

COMPASS also had to address a major challenge for the heterogeneous open systems environment, namely to reduce the complexity of managing storage on block devices. The first documentation that covered this project was released to the public in 2003 in the form of the IBM Systems Journal, Vol. 42, No. 2, 2003, "*The software architecture of a SAN storage control system*" by J. S. Glider, C. F. Fuente, and W. J. Scales.

The article is available at [this website](#).

The first release of IBM System Storage SAN Volume Controller was announced in July 2003.

IBM Spectrum Virtualize Software only is a *software-defined storage* (SDS) implementation that provides all the capabilities and functions of the IBM SAN Volume Controller and was announced in 2016. It runs on supported Intel hardware that the customer supplies.

This SDS layer is designed to virtualize and optimize storage within the data center or managed private cloud service. 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. For more information, see *Implementing IBM Spectrum Virtualize software only*, REDP-5392.

Part of the IBM Spectrum family, IBM Spectrum Virtualize for Public Cloud (released in 2017) is the solution adapted for public cloud implementations of IBM Spectrum Virtualize Software only. At the time of the writing of this book, the software supports the deployment on any Intel-based cloud bare-metal servers (not virtualized environment) and is backed by the storage available on the public cloud catalog.

The license pricing aligns with the monthly consumption model of both servers and back-end storage within IBM Cloud. IBM Spectrum Virtualize for Public Cloud provides a new solution to combine on-premises and cloud storage for higher flexibility at lower cost for a comprehensive selection of use cases complementing the existing implementation and options for IBM Spectrum Virtualize and IBM SAN Volume Controller.

Table 1-1 shows the features of IBM Spectrum Virtualize for both on-premises and public cloud products at a glance.

Table 1-1 IBM Spectrum Virtualize features at-a-glance

	On-premises	Public Cloud
Storage supported	Built- in and more than 400 different systems from IBM and others	IBM Cloud Performance and Endurance storage
Licensing approach	Tiered cost per TB (IBM SAN Volume Controller) or per enclosure (Storwize family)	<ul style="list-style-type: none"> ▶ Simple, flat cost per capacity ▶ Monthly licensing
Platform	<ul style="list-style-type: none"> ▶ IBM SAN Volume Controller, Storwize family, IBM FlashSystem® V9000, FlashSystem 9100 ▶ VersaStack, software only 	IBM Cloud bare-metal server infrastructure
Reliability, availability, and serviceability (RAS)	Integrated RAS capabilities	Flexible RAS: IBM Cloud and software RAS capabilities
Service	IBM support for hardware and software	IBM support for software in the IBM Cloud environment

1.2 IBM Spectrum Virtualize for Public Cloud

Designed for software-defined storage environments, IBM Spectrum Virtualize for Public Cloud represents the solution for public cloud implementations and includes technologies that both complement and enhance public cloud offering capabilities.

For example, traditional practices that provide data replication simply by copying storage at one facility to largely identical storage at another facility aren't an option where public cloud is concerned. Also, using conventional software to replicate data imposes unnecessary loads on application servers. More detailed use cases will be analyzed further in Chapter 5, "Typical use cases for IBM Spectrum Virtualize for Public Cloud" on page 119.

IBM Spectrum Virtualize for Public Cloud delivers a powerful solution for the deployment of IBM Spectrum Virtualize software in public clouds, starting with IBM Cloud. This new capability provides a monthly license to deploy and use IBM Spectrum Virtualize in IBM Cloud to enable hybrid cloud solutions, offering the ability to transfer data between on-premises data centers using any IBM Spectrum Virtualize-based appliance and IBM Cloud.

With a deployment designed for the cloud, the IBM Spectrum Virtualize for Public Cloud can be deployed in any of over 25 IBM Cloud data centers around the world where, after provisioning the infrastructure, an install script automatically installs the software.

1.2.1 Primers of storage virtualization and software-defined-storage

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.

The aggregation of volumes into storage pools enables us to better manage capacity, performance, and multiple tierings for the workloads. IBM Spectrum Virtualize for Public Cloud provides virtualization only at the disk layer (block-based) of the I/O stack, and for this reason is referred to as *block-level virtualization*, or the block aggregation layer. For the sake of clarity, the block-level volumes provided by the IBM Cloud are exposed as iSCSI target volumes, and are seen by IBM Spectrum Virtualize managed disk (MDisk).

These MDisks are then aggregated into a storage pool, sometimes referred to as a managed disk group (mdiskgrp). IBM Spectrum Virtualize then creates logical volumes (referred to as volumes or VDIs) which are striped across all of the MDisks inside of their assigned pool.

The virtualization terminology is included into the wider concept of software-defined storage (SDS), an approach to data storage in which the programming that controls storage-related tasks is decoupled from the physical storage hardware. This separation allows SDS solutions to be placed over any existing storage systems or, more generally, installed on any commodity x86 hardware and hypervisor.

Shifting to a higher level in the IT stack allows for a deeper integration and response to application requirements for storage performance and capabilities. SDS solutions offer a full suite of storage services (equivalent to traditional hardware systems) and federation of multiple persistent storage resources: internal disk, cloud, other external storage systems, or cloud/object platforms.

In general, SDS technology leverages the following concepts:

- ▶ Shared nothing architecture (or in some cases a partial or fully shared architecture): with no single point of failure and nondisruptive upgrade.
- ▶ Scale-up or scale-out mode: adding building blocks for predictable increase in capacity, performance and resiliency.
- ▶ Multiple classes of service: file-based, object-based, block-based, auxiliary/storage support service. SDS solutions maybe also be integrated together into a hybrid or composite SDS solution.
- ▶ High availability (HA) and disaster recovery (DR): able to tolerate level of availability and durability as self healing and adjusting.
- ▶ Lower TCO: lowering the TCO for those workloads capable of using SDS.

1.2.2 Benefits of IBM Spectrum Virtualize for Public Cloud

IBM Spectrum Virtualize for Public Cloud offers a powerful value proposition for enterprise and cloud users who are searching for more flexible and agile ways to deploy block storage on cloud. Using standard Intel servers, IBM Spectrum Virtualize for Public Cloud can be easily added to existing cloud infrastructures to deliver additional features and functionalities, enhancing the storage offering available on public cloud catalog. The benefits of deploying IBM Spectrum Virtualize on a public cloud platform are two-fold:

- ▶ Public cloud storage offering enhancement

IBM Spectrum Virtualize for Public Cloud enhances the public cloud catalog by increasing standard storage offering capabilities and features improving specific limitations:

- Snapshots: a volume's snapshots occur at high-tier storage with no options for lower-end storage tier. Using IBM Spectrum Virtualize, the administrator has more granular control, enabling a production volume to have a snapshot stored lower-end storage.
- Volume size: most cloud storage providers have a maximum volume size (typically a few TB) that can be provided which can be mounted by a few nodes. At the time of writing, IBM Spectrum Virtualize allows for up to 256 TB and up to 20,000 host connections.
- Native storage-based replication: replication features are natively supported but are typically limited to specific data center pairs, to a predefined minimum recovery point objective (RPO). They are accessible only when the primary volume is down. IBM Spectrum Virtualize provides greater flexibility in storage replication allowing for user-defined RPO and replication between any other system running IBM Spectrum Virtualize.

- ▶ New features for public cloud storage offering

IBM Spectrum Virtualize for Public Cloud introduces to the public cloud catalog new storage capabilities as those features available on IBM SAN Volume Controller and IBM Spectrum Virtualize, not available by default. These additional features provided on public cloud are mainly related to hybrid cloud scenarios and its support to foster all those solutions for improved hybrid architectures:

- Replication or migration of data between on-premises storage and public cloud storage
In a heterogeneous environment (VMware, bare metal, Hyper-V, and so on), replication consistency is usually achieved through storage-based replica peering cloud storage with primary storage on premises. Due to standardization of storage service model and inability to move its own storage to a cloud data center, the storage-based replica is usually achievable only by involving an SDS solution on premises.

In this sense, IBM Spectrum Virtualize for Public Cloud not only offers data replication between Storwize family, FlashSystem V9000, FlashSystem 9100, IBM SAN Volume Controller, or VersaStack and Public Cloud, but extends replication to all types of supported virtualized storage on-premises. Working together, IBM Spectrum Virtualize and IBM Spectrum Virtualize for Public Cloud support synchronous and asynchronous mirroring between the cloud and on-premises for more than 400 different storage systems from a wide variety of vendors. In addition, they support other services, such as IBM FlashCopy® and IBM Easy Tier®.

- Disaster recovery strategies between on-premises and public cloud data centers as alternative DR solutions

One of the reasons to replicate is to have a copy of the data from which to restart operations in case of emergency. IBM Spectrum Virtualize for public cloud enables this for virtual and physical environments, thus adding new possibilities compared to software replicators in use today that handle virtual infrastructure only.

- Benefit from familiar, sophisticated storage functionality in the cloud to implement reverse mirroring

IBM Spectrum Virtualize enables the possibility to reverse data replication to offload from Cloud Provider back to on-premises or to another Cloud provider.

IBM Spectrum Virtualize, both on-premises and on cloud, provides a data strategy that is independent of the choice of infrastructure, delivering tightly integrated functionality and consistent management across heterogeneous storage and cloud storage. The software layer provided by IBM Spectrum Virtualize on premises or in the cloud can provide a significant business advantage by delivering more services faster and more efficiently, enabling real-time business insights and supporting more customer interaction.

Capabilities such as rapid, flexible provisioning; simplified configuration changes; nondisruptive movement of data among tiers of storage; and a single user interface help make the storage infrastructure (and the hybrid cloud) simpler, more cost-effective, and easier to manage.

IBM Spectrum Virtualize Software only

IBM Spectrum Virtualize for Public Cloud can work together with IBM SAN Volume Controller and IBM Spectrum Virtualize (on-premises only). IBM Spectrum Virtualize Software only is an SDS implementation that provides all the capabilities and functions of the IBM SAN Volume Controller. IBM Spectrum Virtualize software is installed on supported bare-metal Intel servers.

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 following benefits of storage virtualization and advanced storage capabilities:

- ▶ Support for more than 400 different storage systems from a wide variety of vendors
- ▶ Storage pooling and automated allocation with thin provisioning
- ▶ Easy Tier automated tiering
- ▶ IBM Real-time Compression™, enabling storing up to five times as much data for even the most demanding applications
- ▶ Software encryption to improve data security on existing storage (IBM Spectrum Virtualize for Public Cloud uses cloud infrastructure encryption services)
- ▶ FlashCopy and remote mirror for local and remote replication

- ▶ Support for virtualized and containerized server environments including VMware (VVOL), Microsoft Hyper-V, IBM PowerVM®, Docker, and Kubernetes

Figure 1-1 shows an overview of the IBM Spectrum Virtualize software-only solution.

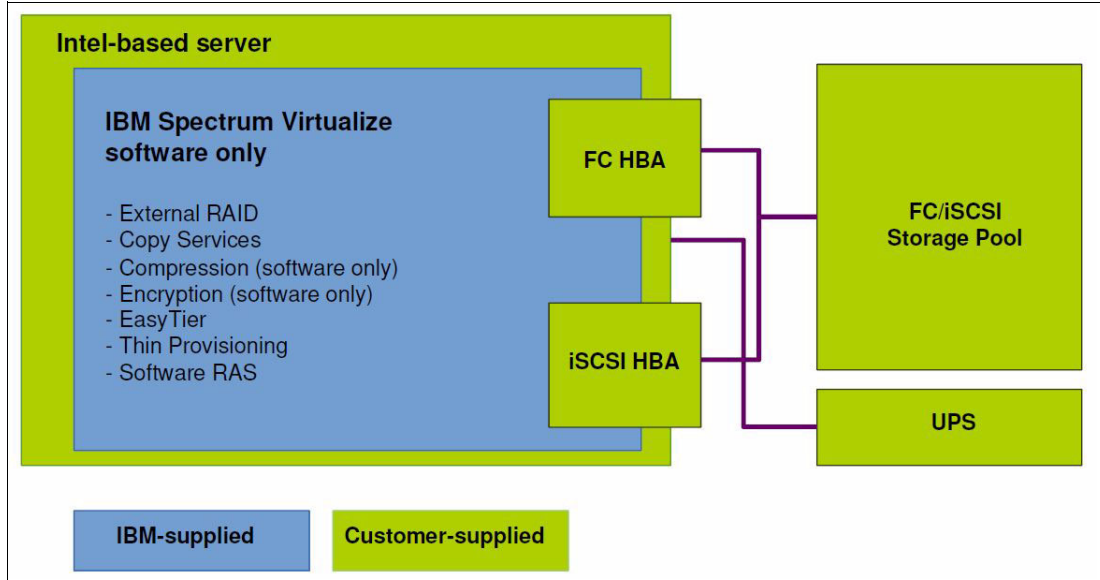


Figure 1-1 IBM Spectrum Virtualize Software only with customer provided hardware

1.2.3 Features of IBM Spectrum Virtualize for Public Cloud

IBM Spectrum Virtualize for a Public Cloud helps make cloud storage volumes (block-level) more effective by including functions that are not natively available on the public cloud catalogs and that are traditionally deployed within disk array systems in the on-premises environment. For this reason, IBM Spectrum Virtualize for Public Cloud improves and expands the existing capabilities of the cloud offering.

Table 1-1 summarizes IBM Spectrum Virtualize for Public Cloud features and benefits.

Table 1-2 IBM Spectrum Virtualize for Public Cloud features and benefits

Feature	Benefits
Single point of control for cloud storage resources	<ul style="list-style-type: none"> ▶ Designed to increased management efficiency ▶ Designed to help to support application availability
Pools the capacity of multiple storage volumes	<ul style="list-style-type: none"> ▶ Helps to overcome the volume size limitations ▶ Helps to manage storage as a resource to meet business requirements, and not just as a set of independent volumes ▶ Helps administrator to better deploy storage as required beyond traditional “islands” ▶ Can help to increase the use of storage assets ▶ Insulate applications from maintenance or changes to storage volume offering
Clustered pairs of Intel servers that are configured as IBM Spectrum Virtualize for Public Cloud data engines	<ul style="list-style-type: none"> ▶ Use of cloud-catalog Intel servers foundation ▶ Designed to avoid single point of hardware failures

Feature	Benefits
Manage tiered storage	<ul style="list-style-type: none"> ▶ Helps to balance performance needs against infrastructures costs in a tiered storage environment ▶ Automated policy-driven control to put data in the right place at the right time automatically among different storage tiers/classes
Easy-to-use IBM Storwize family management interface	<ul style="list-style-type: none"> ▶ Single interface for storage configuration, management, and service tasks regardless the configuration available from public cloud portal ▶ Helps administrators use storage assets/volumes more efficiently ▶ IBM Spectrum Control™ Insights and IBM Spectrum Protect™ for additional capabilities to manage capacity and performance
Dynamic data migration	<ul style="list-style-type: none"> ▶ Migrate data among volumes/LUNs without taking applications that use that data offline ▶ Manage and scale storage capacity without disrupting applications
Advanced network-based copy services	<ul style="list-style-type: none"> ▶ Copy data across multiple storage systems with IBM FlashCopy ▶ Copy data across metropolitan and global distances as needed to create high-availability storage solutions between multiple data centers
Thin provisioning and snapshot replication	<ul style="list-style-type: none"> ▶ Reduce volume requirements by using storage only when data changes ▶ Improve storage administrator productivity through automated on-demand storage provisioning ▶ Snapshots available on lower tier storage volumes
IBM Spectrum Protect Snapshot application-aware snapshots	<ul style="list-style-type: none"> ▶ 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
Integrated Bridgeworks SANrockIT technology for IP replication	<ul style="list-style-type: none"> ▶ Optimize use of network bandwidth ▶ Reduce network costs or speed replication cycles, improving the accuracy of remote data
Third parties native integration	Integration with VMware vRealize and Site Recovery Manager

Note: The following features are not supported in the first IBM Spectrum Virtualize for Public Cloud release:

- ▶ Stretched cluster
- ▶ IBM HyperSwap®
- ▶ Real-time Compression
- ▶ Data deduplication
- ▶ Encryption
- ▶ Data reduction
- ▶ Unmap
- ▶ Cloud backup
- ▶ Transparent cloud tiering
- ▶ Hot spare node
- ▶ Distributed RAID
- ▶ N-Port ID Virtualization

Some of these features are already in the plan for future releases and will be prioritized for implementation based on customer feedback.

1.2.4 IBM Spectrum Virtualize on IBM Cloud

The initial release of IBM Spectrum Virtualize for Public Cloud is available on IBM Cloud and is designed for deployment on other cloud service providers at later code versions. Block virtualization further leverages public cloud infrastructure for various types of workload deployments whether it is new or traditional. The following features are supported on IBM Cloud infrastructure:

- ▶ Offers data replication with Storwize family, V9000, IBM SAN Volume Controller, or Versastack, and between public clouds
- ▶ Supports two, four, six, or eight Node Clusters in IBM Cloud
- ▶ Data Services for IBM Cloud Block Storage
- ▶ Common Management: IBM Spectrum Virtualize GUI
- ▶ Deployment in more than 25 globally distributed data centers

IBM Cloud infrastructure (shown in Figure 1-2 on page 10) is a proven, established platform for today's computing needs, by deploying IBM Spectrum Virtualize on a cloud platform, features of IBM SAN Volume Controller and IBM Spectrum Virtualize software only are further enhanced for the changing environments. Customers can decide which configuration to start with, just like a regular IBM SAN Volume Controller or Storwize product a two node cluster can be upgraded all the way up to an eight node cluster dynamically without impacting the production environment.

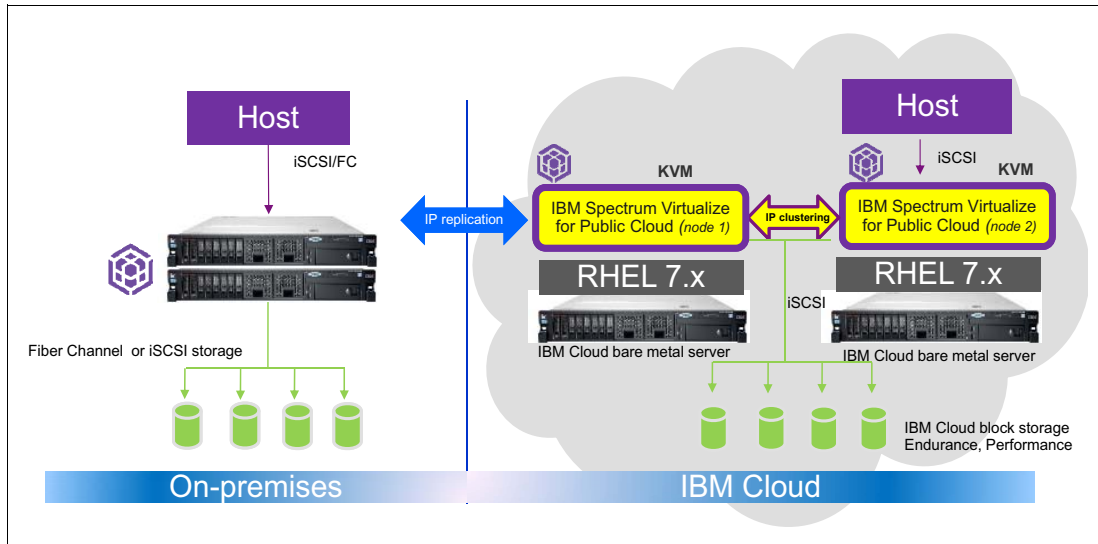


Figure 1-2 High-level architecture of IBM Spectrum Virtualize for Public Cloud

Among multiple infrastructure deployment models on IBM Cloud, IBM Spectrum Virtualize is supported on *bare-metal servers*. IBM Cloud bare-metal servers provide the raw horsepower for processor-intensive and disk I/O-intensive workloads. They're also completely customizable, down to the exact specifications, which enables unmatched control of the cloud infrastructure.

IBM Cloud bare metal provides 10 Gbps network interfaces that sit on a *Triple Network Architecture* with dedicated backbone network. IBM Cloud offers a wide range of data centers and *Network Points of Presence (PoPs)* throughout the world. Connecting these data centers and PoPs is a redundant, high-speed network. Therefore, no traffic between data centers or PoPs is ever routed over the Internet but rather stays in IBM Cloud's private network. Better yet, all network traffic on the internal network is unmetered and therefore without incremental cost.

This creates compelling deployment architecture opportunities, especially for failover and disaster recovery where it is now possible to mirror data between data centers without having to pay for the (sometimes significant) traffic between data centers. The Triple Network Architecture provides three network interfaces to every server regardless of if it is a bare-metal server or a virtual compute instance. Each server is complimented with a five physical network interface card (NIC) configuration:

- ▶ Public internet access
- ▶ Private network
- ▶ Management access

The available private interfaces are 2 x 10 Gbps and currently it is not possible to move public interfaces to private network to increase the number of NIC private ports. Because in IBM Cloud native features there is a physical separation between public and private network interfaces, and in the IBM Spectrum Virtualize deployment, all the intra-nodes traffic is routed within the private network.

Networking for IBM Spectrum Virtualize for Public Cloud is all IP based with no Fibre Channel (FC), which is not supported by IBM Cloud. This includes inter-node communication and inter-cluster replication (remote replication, on-premises to cloud or cloud to cloud).

IBM Cloud provides IBM Spectrum Virtualize with flash-backed block storage on high-performance iSCSI targets. The storage is presented as a block-level device that customers can format to best fit their needs. The iSCSI storage resides on the private network and does not count toward public and private bandwidth allotments. Options available are either Performance in granular IOPS (Input/output operations per second) increments 1,000 - 48,000 or predefined per GB Endurance tier.

Both are available as volumes sized 20 GB - 12 TB. All volumes are encrypted by default with IBM Cloud-managed encryption.

Note: IBM Cloud Endurance and Performance volumes when used as back-end storage for IBM Spectrum Virtualize on IBM Cloud are no different from a technical stand point. Once the IOPS profile fits the application requirements from an IBM Spectrum Virtualize perspective the two solutions are identical. The only notable advantage is the IBM Cloud Performance storage granularity during the definition of the IOPS profile. This allows for much more accurate capability estimation, which minimizes waste.

An install script automatically installs IBM Spectrum Virtualize for Public Cloud software, and the software installation must be performed after first purchasing the bare-metal servers using an IBM Cloud IaaS account (formerly known as SoftLayer® and Bluemix).

For more information about architecture and networking, see Chapter 2, “Solution architecture” on page 17. For more information about installation steps, see Chapter 4, “Implementation” on page 71.

1.3 Use cases for IBM Spectrum Virtualize for Public Cloud

From a high-level perspective, IBM Spectrum Virtualize (on premises or in the cloud) delivers leading benefits that improve how to use storage in three key ways:

- ▶ Improving data value
IBM Spectrum Virtualize software helps reduce the cost of storing data by increasing utilization and accelerating applications to speed business insights.
- ▶ Increasing data security
IBM Spectrum Virtualize helps enabling a high-availability strategy that includes protection for data and application mobility and disaster recovery.
- ▶ Enhancing data simplicity
IBM Spectrum Virtualize provides a data strategy that is independent of infrastructure, delivering tightly integrated functionality and consistent management across heterogeneous storage.

Figure 1-3 shows the deployment models for IBM Spectrum Virtualize.

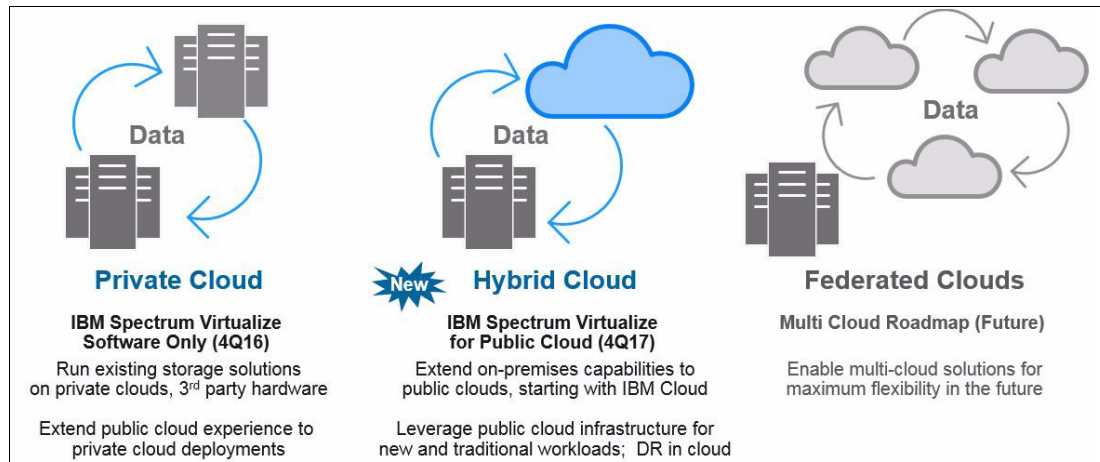


Figure 1-3 IBM Spectrum Virtualize in the Public Cloud deployment models

These three key benefits span over multiple use cases where IBM Spectrum Virtualize applies. In fact, IBM Spectrum Virtualize for Public Cloud provides a new solution to combine on-premises and cloud storage for higher flexibility at lower cost for a comprehensive selection of use cases both for hybrid cloud solutions and cloud native architectures. This includes and is not limited to:

- ▶ Data migration and disaster recovery (DR) to Public Cloud
- ▶ Data Center extension or consolidation to Public Cloud
- ▶ Data migration and disaster recovery (DR) between Public Cloud data centers
- ▶ Federation to/between multiple cloud providers (statement of direction)

1.3.1 Hybrid scenario: on-premises to IBM Cloud

When discussing integration to IBM Cloud Infrastructure, there are multiple angles to approach it. Building a hybrid cloud solution, traditional practices that provide data replication simply by copying storage at one facility to largely identical storage at another facility aren't an option where public cloud is concerned. And using conventional software to replicate data imposes unnecessary loads on application servers.

As shown in Figure 1-4 on page 13, many of the existing on-premises environments are heterogeneous and composed of several different technologies, such as VMware, Hyper-V, KVM, Oracle, IBM. In order to achieve data consistency when migrating or replicating, the usage of a storage-based replica is the preferred solution rather than multiple specific tools working all together that introduce complexity in their steady state management.

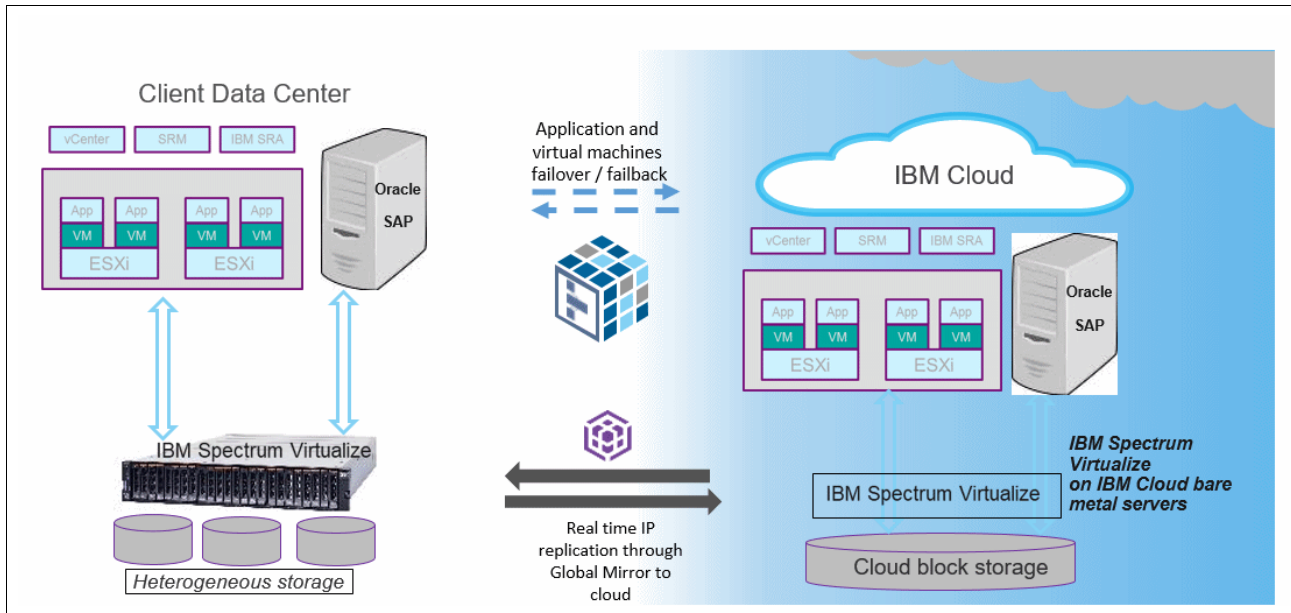


Figure 1-4 Heterogeneous environment are migrated or replicated to IBM Cloud

IBM Spectrum Virtualize on IBM Cloud enables hybrid cloud solutions, offering the ability to transfer data between on-premises data centers using any IBM Spectrum Virtualize based appliance, such as IBM SAN Volume Controller, Storwize family products, FlashSystem V9000, FlashSystem 9100 and VersaStack with Storwize family or IBM SAN Volume Controller appliances, or IBM Spectrum Virtualize Software Only and IBM Cloud. Other non-IBM vendors are also supported.

Figure 1-5 shows a typical scenario. Through IP-based replication with Global or Metro Mirror, or Global Mirror with Change Volumes, users can create secondary copies of their on-premises data in the public cloud for disaster recovery, workload redistribution, or migration of data from on premises data centers to the public cloud.

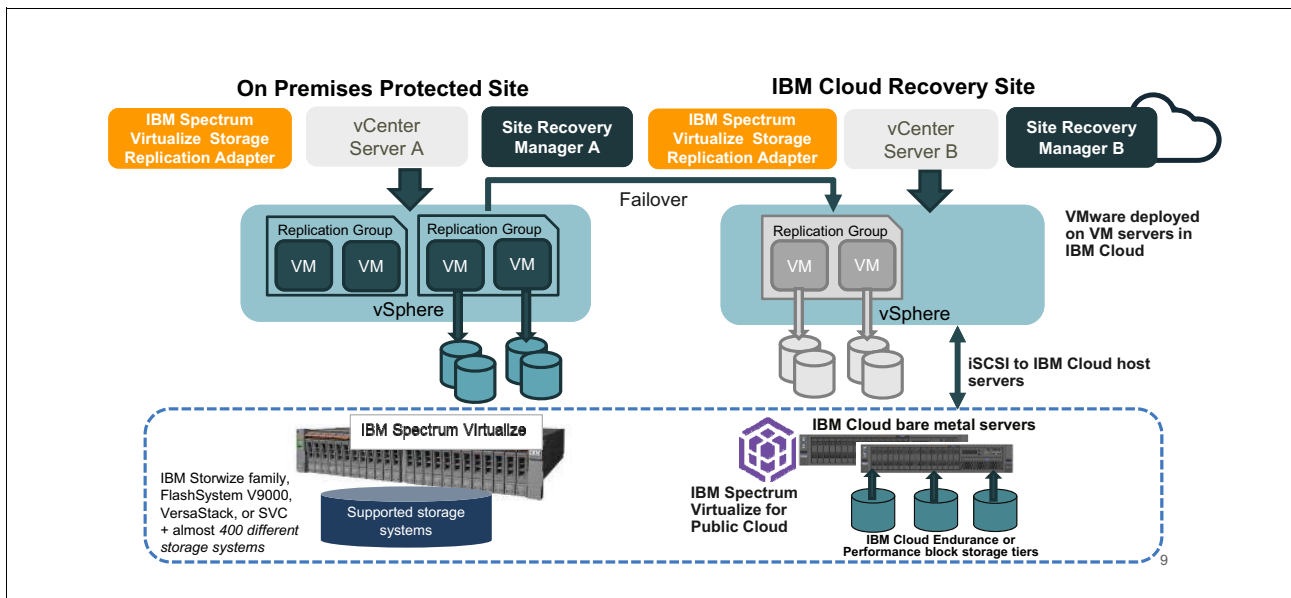


Figure 1-5 Hybrid scenario: VMware on IBM Cloud with Site Recovery Manager solution

In this sense, IBM Spectrum Virtualize for Public Cloud represents the ideal target, by abstracting the storage layer, to avoid any dependency with specific vendor (both at storage and application layer), even if multiple storage technologies are involved.

You can create two-node to eight-node high-availability clusters, similar to on premises IBM SAN Volume Controller appliances on IBM Cloud. IBM Cloud block storage can be easily managed through IBM Spectrum Virtualize for Public Cloud for persistent data storage, and as the target of remote copy services.

1.3.2 Cloud-native scenario: Cloud to cloud

IBM Spectrum Virtualize for Public Cloud also supports those use cases where cloud users are moving the entire workload/application to cloud. In this sense, there is no integration between the on premises and the cloud but the application is *cloud-native*.

In this case, the replication scenario from multiple storage resources is not applicable but the complex and heterogeneous environment still applies. IBM Spectrum Virtualize for Public Cloud extends and enhances the capabilities of IBM Cloud block storage by using enterprise-class features, such as local and remote copy services. Replication features are often available among cloud providers but limited to specific RPOs and feature that are not negotiable or adjustable because of the standardized nature of public clouds.

For this reason IBM Spectrum Virtualize for Public Cloud overcomes some of the limitations of the public cloud catalog, is a good fit also for cloud to cloud scenario. As shown in Figure 1-5, the VMware distributed over multiple IBM Cloud data centers is also a common use case.

Last but not least, is the necessity of the users to maintain the existing skills and tools and access them on cloud as it was their own data center, for a smooth transition to off-premises environments.

1.4 Licensing

Within the public cloud model, servers and storage resources are provisioned and prices based either on a monthly or hourly usage. In order to adapt to the public cloud flexibility, IBM Spectrum Virtualize for Public Cloud has monthly licensing based on the number of terabytes of IBM Cloud block storage that is managed by IBM Spectrum Virtualize for Public Cloud. Additional options for metered usage are also available so you can purchase more capacity as needed.

The IBM Cloud compute, bare-metal nodes and backend storage capacity are part of a separate purchase and, with network equipment, are not included in the licensing. For IBM Spectrum Virtualize for Public Cloud, licenses are available on IBM Marketplace and Passport Advantage® under IBM Spectrum Virtualize for Public Cloud (5737-F08) and two licensing models (it is not available as shrink wrap). These are all-inclusive and flat \$/TB:

- ▶ IBM Spectrum Virtualize for Public Cloud 10 Terabytes Monthly Base License
- ▶ IBM Spectrum Virtualize for Public Cloud 1 Terabyte Monthly License Incremental capacity

After the licenses are purchased, the deployment model follows a Bring Your Own License to IBM Cloud. A semi-automated install script for IBM Spectrum Virtualize will install the software according to the procedure described in 4.1.3, “Semi Automated installation” on page 78. The automated capacity metering each month is available through Call Home for additional capacity purchase.

Note: Additional capacity purchases are available through Passport Advantage according to the license terms for this offering. Automated billing is not enabled at this time.



Solution architecture

In this chapter, we provide a technical overview of the IBM Cloud environment and the solution that is implemented in this book. This chapter is meant to provide technical information about the solution components and propose a reference architecture that can be used.

However, each environment is unique and as such, it is important to review the planning considerations in Chapter 3, “Planning and preparation for the IBM Spectrum Virtualize for Public Cloud deployment” on page 31 before designing your solution.

This chapter includes the following topics:

- ▶ 2.1, “IBM Cloud” on page 18
- ▶ 2.2, “Storage virtualization” on page 19
- ▶ 2.3, “IBM Spectrum Virtualize” on page 19
- ▶ 2.4, “Environment used for this book” on page 28

2.1 IBM Cloud

The IBM Cloud IaaS offering provides a robust environment, as shown in Figure 2-1.

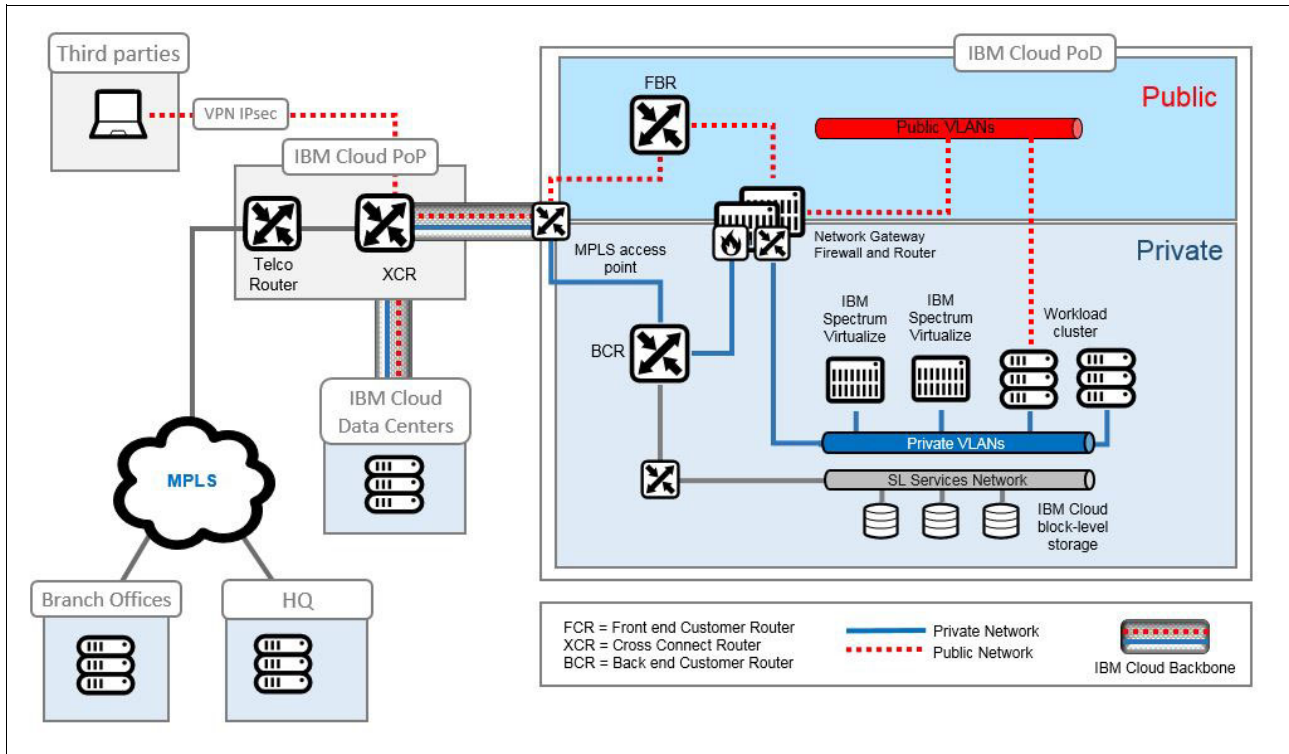


Figure 2-1 IBM Cloud networking

Each system provisioned in the IBM Cloud is connected to a public and private network in the IBM Cloud PoD (point of delivery). The public network is internet routable and accessible by default. This network is used to serve internet and web applications in the cloud and to allow users to access the internet if needed.

The private network is internal to IBM Cloud and is used typically for services provided by IBM in the cloud. This includes access to block, file, and Object Storage. This also includes communication between servers that are provisioned in the cloud. Additionally, clients can end a Multiprotocol Label Switching (MPLS) connection into the private network and allow access between their on-premises data center resources and IBM Cloud.

When deployed in the cloud environment, a network gateway appliance as seen in Figure 2-1 controls access going to and from both the public and private networks provisioned to a particular environment. This appliance can be used to secure cloud servers and applications. This appliance can also be used to terminate IPsec VPN connections between sites over the internet.

Both the public and the private networks converge at the IBM Cloud backbone network environment and the POP (Point of Presence). These network connections serve as the gateway for external connections coming into the IBM Cloud through MPLS, the termination point for internet access, and as a network to link together multiple IBM Cloud data centers.

2.2 Storage virtualization

Storage virtualization is a term that is used extensively throughout the storage industry. It can be applied to various technologies and underlying capabilities. In reality, most storage devices technically can claim to be virtualized in one form or another. IBM describes storage virtualization as 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 function with existing services and can be nested or applied to multiple layers of a system.

When the term storage virtualization is mentioned, it is important to understand that virtualization can be implemented at various layers within the I/O stack. There must be a clear distinction between virtualization at the disk layer (block-based) and virtualization at the file system layer (file-based).

The focus of this publication is virtualization at the disk layer, which is referred to as *block-level virtualization* or the *block aggregation layer*. A description of file system virtualization is beyond the intended scope of this book.

Figure 2-2 shows an overview of block-level virtualization.

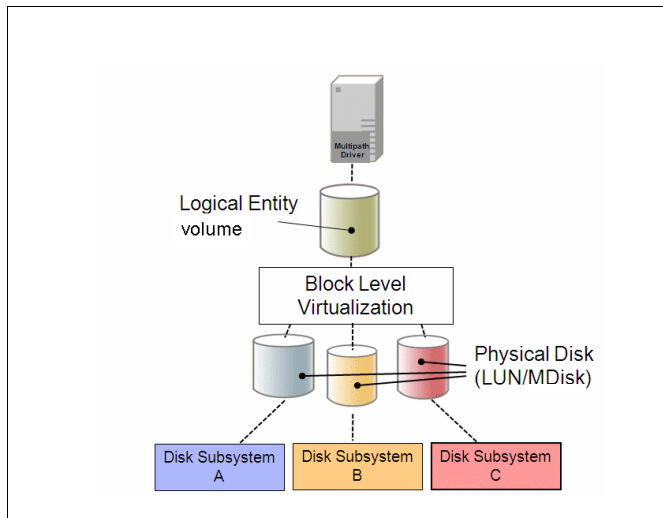


Figure 2-2 Block-level virtualization overview

2.3 IBM Spectrum Virtualize

IBM Spectrum Virtualize is a software-enabled storage virtualization engine that provides a single point of control for storage resources within the data centers. IBM Spectrum Virtualize is a core software engine of well-established and industry-proven IBM storage virtualization solutions, such as IBM SAN Volume Controller and all versions of IBM Storwize family of products (IBM Storwize V3700, IBM Storwize V5000, IBM Storwize V7000, IBM FlashSystem V9000 and 9100). This technology is now available in the IBM Cloud, providing increased flexibility in data center infrastructure and cloud systems. This section describes the components of Spectrum Virtualize as they are deployed in the cloud.

2.3.1 Nodes

IBM Spectrum Virtualize software is installed on bare-metal servers provisioned in the IBM Public Cloud. Each bare-metal server unit is called a *node*. The node provides the virtualization for a set of volumes, cache, and copy services functions. The nodes are deployed in pairs (*I/O groups*) and 1 - 4 pairs make up a *clustered system*.

One of the nodes within the system is assigned the role of the *configuration node*. The configuration node manages the configuration activity for the system and owns the cluster IP address that is used to access the management Graphical User Interface (GUI) and Command Line Interface (CLI) connections. If this node fails, the system chooses a new node to become the configuration node.

Because the active nodes are installed in pairs, each node maintains cache coherence with its partner to provide seamless failover functionality and fault tolerance, which is described next.

2.3.2 I/O groups

Each pair of Spectrum Virtualize nodes is referred to as an *I/O group*. As with all Spectrum Virtualize products, the version for IBM Public Cloud can support clusters of up to four I/O groups.

A specific *volume* is always presented to a host server by a single I/O group in the system. When a host server performs I/O to one of its volumes, all the I/Os for a specific volume are directed to one specific I/O group in the system. Under normal conditions, the I/Os for that specific volume are always processed by the same node within the I/O group. This node is referred to as the *preferred node* for this specific volume. As soon as the preferred node receives a write into its cache, that write is mirrored to the partner node before the write is acknowledged back to the host. Reads are serviced by the preferred node. More on this in section 2.3.7, “Cache” on page 23.

Both nodes of an I/O group act as the preferred node for their own specific subset of the total number of volumes that the I/O group presents to the host servers. However, both nodes also act as failover nodes for their respective partner node within the I/O group. Therefore, a node takes over the I/O workload from its partner node, if required. For this reason, it is mandatory for servers that are connected to use multipath drivers to handle these failover situations.

If required, host servers can be mapped to more than one I/O group within the Spectrum Virtualize system. Therefore, they can access volumes from separate I/O groups. You can move volumes between I/O groups to redistribute the load between the I/O groups. Modifying the I/O group that services the volume can be done concurrently with I/O operations if the host supports nondisruptive volume moves and is zoned to support access to the target I/O group.

It also requires a rescan at the host level to ensure that the multipathing driver is notified that the allocation of the preferred node changed, and the ports by which the volume is accessed changed. This modification can be done in the situation where one pair of nodes becomes overused.

2.3.3 System

The current IBM Cloud Spectrum Virtualize system or clustered system consists of 1 - 4 I/O groups. Certain configuration limitations are then set for the individual system. For example, at the time of writing, the maximum number of volumes that is supported per system is 10000, or the maximum managed disk that is supported is ~28 PIB (pebibytes) or 32 PB (petabytes) per system.

All configuration, monitoring, and service tasks are performed at the system level. Configuration settings are replicated to all nodes in the system. To facilitate these tasks, a management IP address is set for the system.

Note: The management IP is also referred to as the system or cluster IP and is active on the configuration node. Each node in the system is also assigned a service IP to allow for individually interacting with the node directly.

A process is provided to back up the system configuration data onto disk so that it can be restored if there is a disaster. This method does not back up application data. Only the Spectrum Virtualize system configuration information is backed up.

For the purposes of remote data mirroring, two or more systems must form a *partnership* before relationships between mirrored volumes are created.

For more information about the maximum configurations that apply to the system, I/O group, and nodes, see the IBM Spectrum Virtualize 8.2.1 [configuration limits web page](#).

2.3.4 MDisks

The IBM Spectrum Virtualize system and its I/O groups view the storage that is presented to the LAN by the back-end controllers as several disks or LUNs, which are known as *managed disks* or *MDisks*. Because Spectrum Virtualize does not attempt to provide recovery from physical failures within the back-end controllers, an MDisk often is typically provisioned from a RAID array.

However, the application servers do not see the MDisks at all. Rather, they see several logical disks, which are known as *virtual disks* or *volumes*, which are presented by the I/O groups through the LAN (iSCSI) to the servers. The MDisks are placed into storage pools where they are divided into several extents used to create the *virtual disks* or *volumes*.

For more information about the total storage capacity that is manageable per system regarding the selection of extents, see the IBM Spectrum Virtualize 8.2.1 [configuration limits web page](#).

MDisks presented to Spectrum Virtualize can have the following modes of operation:

- ▶ Unmanaged MDisk

An MDisk is reported as unmanaged when it is not a member of any storage pool. An unmanaged MDisk is not associated with any volumes and has no metadata that is stored on it. Spectrum Virtualize does not write to an MDisk that is in unmanaged mode, except when it attempts to change the mode of the MDisk to one of the other modes.

- ▶ Managed MDisk

Managed MDisks are members of a storage pool and they contribute extents to the storage pool. This mode is the most common and normal mode for an MDisk.

- ▶ Image mode MDisk

Image mode provides a direct block-for-block translation from the MDisk to the volume by using virtualization. This mode is provided to satisfy the following major usage scenarios:

- Image mode enables the virtualization of MDisks that already contain data that was written directly and not through an IBM Spectrum Virtualize. Rather, it was created by a direct-connected host. This mode enables a client to insert IBM Spectrum Virtualize into the data path of an existing storage volume or LUN with minimal downtime.

Image mode enables a volume that is managed by IBM Spectrum Virtualize to be used with the native copy services function that is provided by the underlying RAID controller. To avoid the loss of data integrity when IBM Spectrum Virtualize is used in this way, it is important that you disable IBM Spectrum Virtualize cache for the volume.

- IBM Spectrum Virtualize provides the ability to migrate to image mode, which enables IBM Spectrum Virtualize to export volumes and access them directly from a host without IBM Spectrum Virtualize in the path.
- Most typically, image mode is used to import server data into the SVC for migrating that data to a fully managed pool, and then releasing the image mode copy or copies of that data to complete the migration.
- In rare cases, IBM Spectrum Virtualize is used simply as a migration conduit where *both source and destination* are image mode MDISKS, and IBM Spectrum Virtualize is removed from the environment after migration. Given the many benefits of IBM Spectrum Virtualize, most migrations are of the previous design where the data is expected to remain within IBM Spectrum Virtualize.

2.3.5 Storage pool

A *storage pool* or *mdiskgroup* is a collection of MDisks that provides the pool of storage from which volumes are provisioned. The size of these pools can be changed (expanded or shrunk) nondisruptively by adding or removing MDisks, without taking the storage pool or the volumes offline. At any point, an MDisk can be a member in one storage pool only.

Each MDisk in the storage pool is divided into extents. The size of the extent is selected by the administrator when the storage pool is created, and cannot be changed later. The size of the extent can be 16 MiB (mebibyte) - 8192 MiB, with the default being 1024 MiB.

It is a preferred practice to use the same extent size for all storage pools in a system. This approach is a prerequisite for supporting volume migration between two storage pools. If the storage pool extent sizes are not the same, you must use volume mirroring to copy volumes between pools.

2.3.6 Volumes

Volumes are logical disks that are presented to the host or application servers by the Spectrum Virtualize. The hosts cannot see the MDisks; they can see only the logical volumes that are created from combining extents from a storage pool or passed through Spectrum Virtualize in the case of *Image Mode* objects.

There are three types of volumes in terms of extent management:

- ▶ Striped

A striped volume is allocated one extent in turn from each MDisk in the storage pool. This process continues until the space required for the volume has been satisfied.

It is also possible to supply a list of MDisks to use. This is the default volume type.

- ▶ Sequential

A sequential volume is where the extents are allocated from one MDisk. This is usually only used when provisioning Storwize (V7000, V5000, and V3000) volumes as backend storage to Spectrum Virtualize systems and requires specification of the MDisks from which extents will be drawn. A second MDisk is specified if it is a mirrored sequential volume.

- ▶ Image mode

Image mode volumes are special volumes that have a direct relationship with one MDisk. The most common use case of image volumes is a data migration from your old (typically non-virtualized) storage to the Spectrum Virtualize-based virtualized infrastructure.

When the image mode volume is created, a direct mapping is made between extents that are on the MDisk and the extents that are on the volume. The logical block address (LBA) x on the MDisk is the same as the LBA x on the volume, which ensures that the data on the MDisk is preserved as it is brought into the clustered system.

Some virtualization functions are not available for image mode volumes, so it is often useful to migrate the volume into a new storage pool. After migration, the data then resides in a volume that is backed by a fully managed pool.

2.3.7 Cache

The primary benefit of storage cache is to improve I/O response time. Reads and writes to a magnetic disk drive experience seek and latency time at the drive level, which can result in 1 ms - 10 ms of response time (for an enterprise-class disk).

IBM Spectrum Virtualize provides a flexible cache model, and the node's memory can be used as read or write cache. The cache management algorithms allow for improved performance of many types of underlying disk technologies. IBM Spectrum Virtualize's capability to manage, in the background, the destaging operations that are incurred by writes (in addition to still supporting full data integrity) assists with IBM Spectrum Virtualize's capability in achieving good database performance.

The cache is separated into two layers: upper cache and lower cache.

Figure 2-3 on page 24 shows the separation of the upper and lower cache.

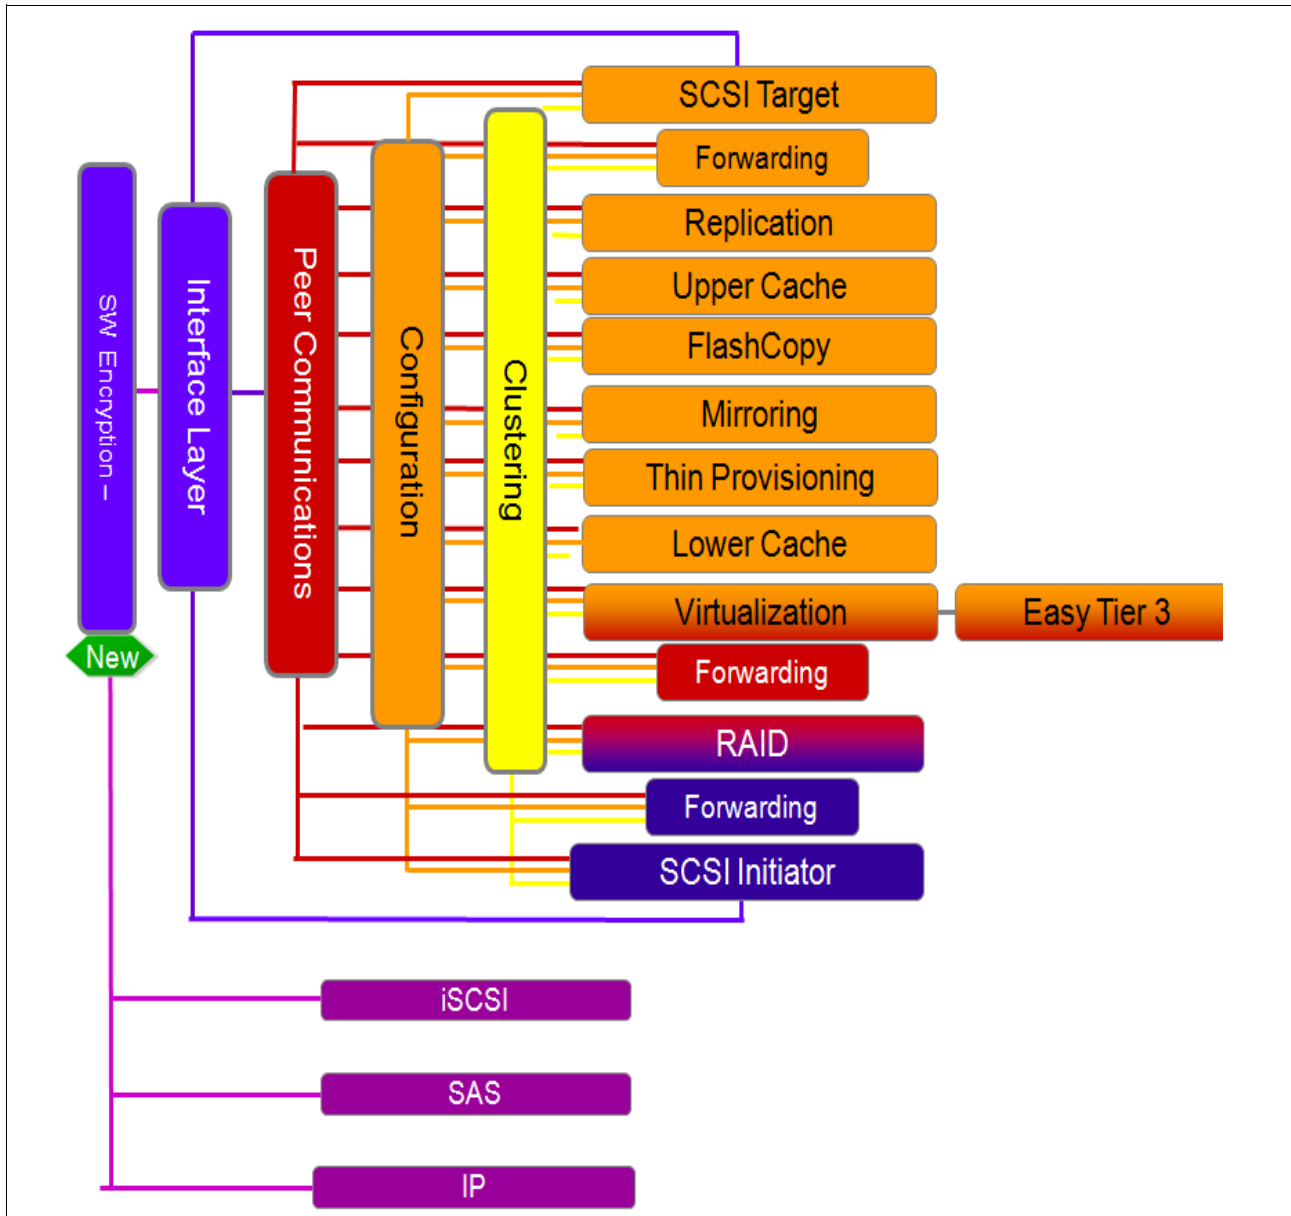


Figure 2-3 Separation of upper and lower cache

The upper cache delivers fast write response times to the host by being as high up in the I/O stack as possible. The lower cache works to help ensure that cache between nodes are in sync, pre-fetches data for an increased read cache hit ratio on sequential workloads, and optimizes the destaging of I/O to the backing storage controllers.

Combined, the two levels of cache also deliver the following functionality:

- ▶ Pins data when the LUN goes offline
- ▶ Provides enhanced statistics for IBM Spectrum Control or Storage Insights, and maintains compatibility with an earlier version
- ▶ Provides trace data for debugging
- ▶ Reports media errors
- ▶ Resynchronizes cache correctly and provides the atomic write functionality

- ▶ Ensures that other partitions continue operation when one partition becomes 100% full of pinned data
- ▶ Supports fast-write (two-way and one-way), flush-through, and write-through
- ▶ Integrates with T3 recovery procedures
- ▶ Supports two-way operation
- ▶ Supports none, read-only, and read/write as user-exposed caching policies
- ▶ Supports flush-when-idle
- ▶ Supports expanding cache as more memory becomes available to the platform
- ▶ Supports credit throttling to avoid I/O skew and offer fairness/balanced I/O between the two nodes of the I/O group
- ▶ Enables switching of the preferred node without needing to move volumes between I/O groups

2.3.8 IBM Easy Tier

IBM Easy Tier is a performance function that automatically migrates or moves extents off a volume to or from one MDisk storage tier to another MDisk storage tier. IBM Spectrum Virtualize code can support a three-tier implementation.

Easy Tier monitors the host I/O activity and latency on the extents of all volumes with the Easy Tier function that is turned on in a multitier storage pool over a 24-hour period.

Next, it creates an extent migration plan that is based on this activity, and then dynamically moves high-activity or hot extents to a higher disk tier within the storage pool. It also moves extents whose activity dropped off or cooled down from the high-tier MDisk back to a lower-tiered MDisk. The condition for hot extents is frequent small block (64 Kb or less) reads.

Easy Tier: The Easy Tier function can be turned on or off at the storage pool and volume level.

The automatic load-balancing (*auto rebalance*) function is enabled by default on each volume, and cannot be turned off using the GUI. This load-balancing feature is not considered the same as the Easy Tier function, however it uses the same principles. Auto rebalance evens the load for a pool across MDisk. Therefore, even the addition of new MDisk, or having MDisk of different sizes within a pool, does not adversely affect the performance.

The IBM Easy Tier function can make it more appropriate to use smaller storage pool extent sizes. The usage statistics file can be off-loaded from the Spectrum Virtualize nodes. Then, you can use the IBM Storage Advisor Tool (STAT) to create a summary report. STAT is available at no initial cost at [this website](#).

2.3.9 Hosts

Volumes can be mapped to a *host* to allow access for a specific server to a set of volumes. A host within the Spectrum Virtualize is a collection of iSCSI-qualified names (IQNs) that are defined on the specific server.

The iSCSI software in IBM Spectrum Virtualize supports IP Address failover when a node is shut down or rebooted. As a result a node failover (when a node is rebooted) can be handled without having a multipath driver that is installed on the iSCSI attached server.

An iSCSI attached server can reconnect after the node shutdown to the original target IP address, which is now presented by the partner node. However, to protect the server against link failures in the network, the use of a multipath driver is needed. As a result, it is suggested to implement multipathing on all hosts attaching to IBM Spectrum Virtualize systems.

2.3.10 Host cluster

Host cluster is a host object in IBM Spectrum Virtualize. A host cluster is a combination of two or more servers that is connected to IBM Spectrum Virtualize through an Internet SCSI (iSCSI) connection. A host cluster object can see the same set of volumes, therefore, volumes can be mapped to a hostcluster to allow all hosts to have a common mapping.

2.3.11 iSCSI

The *iSCSI function* is a software function that is provided by the IBM Spectrum Virtualize code, IBM introduced software capabilities to allow the underlying virtualized storage to attach to IBM Spectrum Virtualize using iSCSI protocol.

The major functions of iSCSI include encapsulation and the reliable delivery of Command Descriptor Block (CDB) transactions between initiators and targets through the Internet Protocol network, especially over a potentially unreliable IP network.

Every iSCSI node in the network must have an iSCSI name and address. An *iSCSI name* is a location-independent, permanent identifier for an iSCSI node. An iSCSI node has one iSCSI name, which stays constant for the life of the node. The terms *initiator name* and *target name* also refer to an iSCSI name.

An *iSCSI address* specifies not only the iSCSI name of an iSCSI node, but a location of that node. The address consists of a host name or IP address, a TCP port number (for the target), and the iSCSI name of the node. An iSCSI node can have any number of addresses, which can change at any time, particularly if they are assigned by way of Dynamic Host Configuration Protocol (DHCP). An IBM Spectrum Virtualize node represents an iSCSI node and provides statically allocated IP addresses.

2.3.12 IP replication

IP replication allows data replication between IBM Spectrum Virtualize family members. IP replication uses IP-based ports of the cluster nodes.

The configuration of the system is straightforward and IBM Storwize family systems normally find each other in the network and can be selected from the GUI.

IP replication includes *Bridgeworks SANSlide* network optimization technology, and is available at no additional charge. Remember, remote mirror is a chargeable option but the price does not change with IP replication. Existing remote mirror users have access to the function at no additional charge.

IP connections that are used for replication can have long latency (the time to transmit a signal from one end to the other), which can be caused by distance or by many “hops” between switches and other appliances in the network. Traditional replication solutions transmit data, wait for a response, and then transmit more data, which can result in network utilization as low as 20% (based on IBM measurements). In addition, this scenario gets worse the longer the latency.

Bridgeworks SANSlide technology, which is integrated with the IBM Storwize family, requires no separate appliances and so requires no additional cost or configuration steps. It uses artificial intelligence (AI) technology to transmit multiple data streams in parallel, adjusting automatically to changing network environments and workloads.

SANSlide improves network bandwidth utilization up to 3x. Therefore, customers can deploy a less costly network infrastructure, or take advantage of faster data transfer to speed replication cycles, improve remote data currency, and enjoy faster recovery.

2.3.13 Synchronous or asynchronous remote copy

The general application of remote copy seeks to maintain two copies of data. Often, the two copies are separated by distance, but not always. The remote copy can be maintained in either synchronous or asynchronous modes. With IBM Spectrum Virtualize, Metro Mirror and Global Mirror are the IBM branded terms for the functions that are synchronous remote copy and asynchronous remote copy.

Synchronous remote copy ensures that updates are committed at both the primary and the secondary volumes before the application considers the updates complete. Therefore, the secondary volume is fully up to date if it is needed in a failover. However, the application is fully exposed to the latency and bandwidth limitations of the communication link to the secondary volume. In a truly remote situation, this extra latency can have a significant adverse effect on application performance at the primary site.

Special configuration guidelines exist for SAN fabrics and IP networks that are used for data replication. There must be considerations about the distance and available bandwidth of the intersite links.

A function of Global Mirror designed for low bandwidth has been introduced in IBM Spectrum Virtualize. It uses change volumes that are associated with the primary and secondary volumes. These change volumes are used to record changes to the primary volume that are transmitted to the remote volume on an interval specified by the cycle period. When a successful transfer of changes from the master change volume to the auxiliary volume has been achieved within a cycle period, a snapshot is taken at the remote site from the auxiliary volume onto the auxiliary change volume to preserve a consistent state and a freeze time is recorded. This function is enabled by setting the *Global Mirror cycling mode*.

Figure 2-4 shows an example of this function where you can see the association between volumes and change volumes.

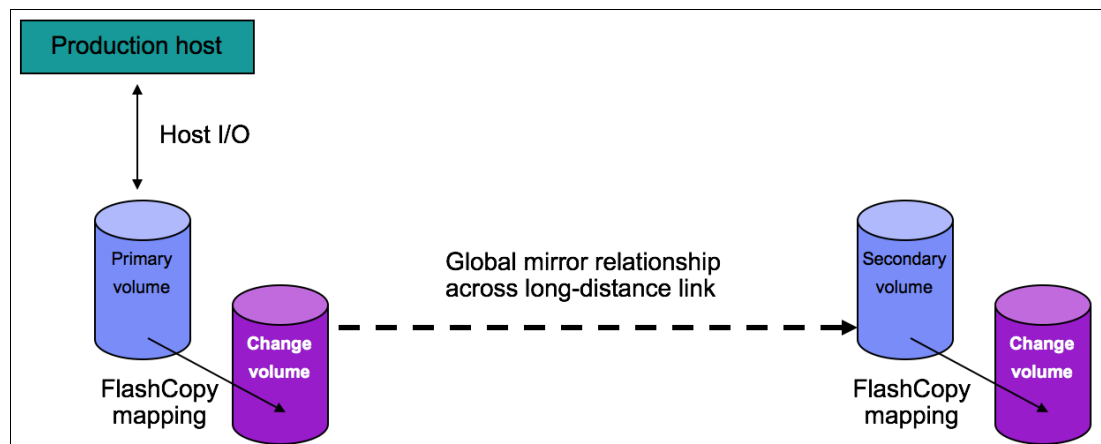


Figure 2-4 Global Mirror cycling mode

FlashCopy

FlashCopy is sometimes described as an instance of a time-zero (T0) copy or a point-in-time (PiT) copy technology.

FlashCopy can be performed on multiple source and target volumes. FlashCopy permits the management operations to be coordinated so that a common single point in time is chosen for copying target volumes from their respective source volumes.

With IBM Spectrum Virtualize, multiple target volumes can undergo FlashCopy from the same source volume. This capability can be used to create images from separate points in time for the source volume, and to create multiple images from a source volume at a common point in time. Source and target volumes can be thin-provisioned volumes.

Reverse FlashCopy enables target volumes to become restore points for the source volume without breaking the FlashCopy relationship, and without waiting for the original copy operation to complete. IBM Spectrum Virtualize supports multiple targets, and therefore multiple rollback points.

Most clients aim to integrate the FlashCopy feature for point in time copies and quick recovery of their applications and databases. An IBM solution to this is provided by IBM Spectrum Protect, which is described on [this website](#).

2.4 Environment used for this book

For the purpose of writing this book, we created a sample environment that involves an on-premises data center location replicating data to the IBM Cloud, as shown in Figure 2-5.




Figure 2-5 Solution overview

In this environment, the on-prem data center is connected to the IBM Cloud using an IPsec VPN that terminates to a network gateway appliance in the cloud. In addition to this, we are using native IP replication between a Storwize system and IBM Spectrum Virtualize. This storage-based replication allows for data consistency between sites.



Planning and preparation for the IBM Spectrum Virtualize for Public Cloud deployment

This chapter describes the preparation steps to provision network, server, and storage components on the IBM Cloud required for installation of IBM Spectrum Virtualize. Background information about the IBM Cloud networking architecture and storage offerings is also described to help the reader who is unfamiliar with the IBM Cloud plan for the IBM Spectrum Virtualize storage placement into the larger context of application environment on the IBM Cloud.

This chapter includes the following topics:

- ▶ 3.1, “Provisioning cloud resources” on page 32
- ▶ 3.2, “Provisioning IBM Cloud Block Storage” on page 48
- ▶ 3.3, “IBM Spectrum Virtualize networking considerations” on page 57

3.1 Provisioning cloud resources

All of the cloud resources required for an IBM Spectrum Virtualize (SV) implementation on the IBM Cloud must be provisioned before the SV software deployment can be performed. The IBM Cloud has APIs with PERL, Python, and Ruby bindings, which can be used to automate this provisioning. For this publication, the provisioning orders are shown as performed manually through the IBM Cloud Portal.

Note: At the time of this writing, the IBM Cloud had recently undergone multiple rebrandings from IBM SoftLayer to IBM Bluemix to the current IBM Cloud. Some of the screens illustrated continue to carry some SoftLayer and Bluemix branding. In the document, all references call the portal the *IBM Cloud Portal*.

All the following sections presume that the installer has a login ID to the IBM Cloud Portal, which can be accessed from <http://control.bluemix.net>.

Users ordering servers need privileges to provision servers, enable public network connections, configure VLANs and subnets, and order storage on the IBM Cloud Portal account.

3.1.1 Ordering servers

IBM Spectrum Virtualize for Public Cloud is enabled for two, four, six, or 8-node configurations. For this example, an order for the servers in a four node cluster is shown. Complete the following steps:

1. On the Portal Home window, select **Devices** under the **Order** pane, as shown in Figure 3-1.

The screenshot shows the IBM Cloud Portal interface. At the top, there is a navigation bar with links for HELP, CONTACT, NOTIFICATIONS, and OPEN TICKETS. The user is logged in as ANGELO BERNASCONI. The main content area is divided into several sections:

- Account Summary:** Shows Current Balance and Estimated Next Balance, both at \$0.00. It also displays the Current Payment Method as Manual.
- Order:** A section with five icons representing different services: Devices, Storage, Network, Security, and Add-On Services. The 'Devices' icon is highlighted.
- Tickets:** Shows two categories: Open (0) and Awaiting Response (0).
- Scheduled Maintenance & Updates:** Displays two calendar views for November and December 2017, showing no scheduled events.
- Bandwidth:** Shows 'No bandwidth data to display'.
- Account Snapshot:** A section at the bottom right.

Figure 3-1 Order page

Note: Because we used a special demo account for this configuration, the prices you see in this section might be different from the prices you see when provisioning your own resources on the IBM Cloud portal.

- At the top of the Server List page, an input field labeled **Select Data Center** with pull-down selection list enables selection of the cloud data center where the servers are to be provisioned. Select a data center before proceeding to the server selection. Care must be taken in selecting the Data Center to ensure that the chosen data center has the appropriate resources that are needed for Spectrum Virtualize, especially the network card configurations and backend storage options. Minimizing distance is an important consideration, but if the chosen datacenter does not contain the appropriate resources, proximity is irrelevant (see Figure 3-2).

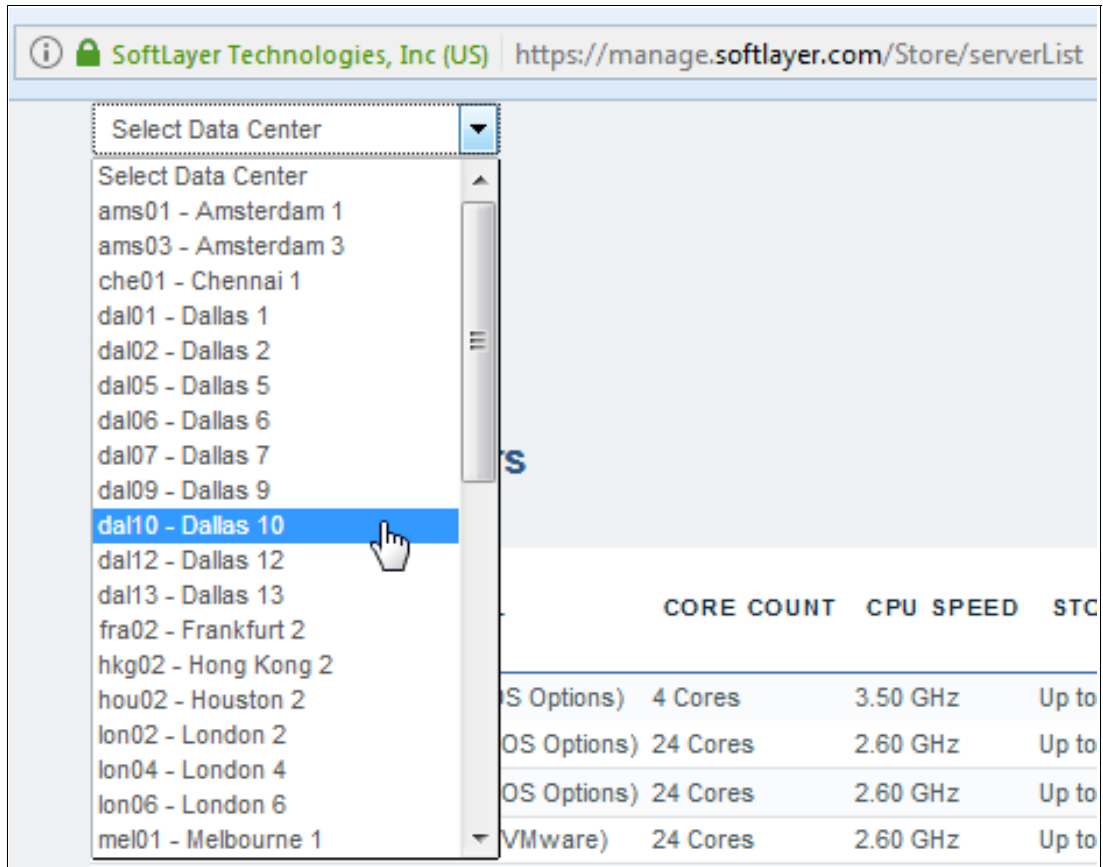


Figure 3-2 Select data center

After the data center is selected, a list of servers is updated to display only those models available in the selected data center.

Tip: If you choose a server from the server list before selecting the data center, your selection is lost when the data center selection updates the list.

3. Scroll down to the Dual Processor Multi-Core Servers. IBM Spectrum Virtualize requires a dual-processor server with a minimum of 12 cores. Servers in the E5-2600 model class are suggested. The selected server must be at least V3, support 64 GB RAM, and have slots for at least four disk drives. A server is selected by clicking the server monthly price, usually referred to as the monthly recurring cost (MRC) (see Figure 3-3).

Dual Processor Multi-Core Servers									
MANUFACTURER	CPU MODEL	CORE COUNT	CPU SPEED	STORAGE	RAM	PRIVATE NETWORK ONLY	GPU SUPPORT	STARTING PRICE PER MONTH	STARTING PRICE PER HOUR
Intel	Intel Xeon E5-2620 v3	12 Cores	2.40 GHz	Up to 4 drives	64 GB up to 768 GB	-	-	\$420.60	-
Intel	Intel Xeon E5-2650 v3	20 Cores	2.30 GHz	Up to 4 drives	64 GB up to 768 GB	-	-	\$478.60	-
Intel	Intel Xeon E5-2620 v4	16 Cores	2.10 GHz	Up to 4 drives	64 GB up to 1536 GB	-	-	\$492.73	-
Intel	Intel Xeon E5-2620 v4	16 Cores	2.10 GHz	Up to 1 drives	64 GB	-	-	\$503.88	\$0.760

Figure 3-3 Monthly recurring cost

In Figure 3-4 on page 35, the E5-2620-V4 server is selected for configuration and ordering. When the server configuration page is displayed, the selected server appears in a list of related models, and allows a change of model for the final selection.

Tip: Spectrum Virtualize does not support Intel TXT verification capabilities so the TXT option should be left cleared.

If the server allows multiple RAM configurations, it can be changed. Spectrum Virtualize currently cannot utilize more than 64 GB of RAM so there is no benefit in selecting more than the minimum 64 GB of RAM for the server. This might change in future releases, in which case additional RAM could be selected.

▼ System Configuration

Server

- Dual Intel Xeon E5-2620 v4 (16 Cores, 2.10 GHz) [\$192.13]
- Dual Intel Xeon E5-2650 v4 (24 Cores, 2.20 GHz) [\$250.13]
- Dual Intel Xeon E5-2690 v4 (28 Cores, 2.60 GHz) [\$315.38]

Server Security

- None
- Intel TXT (Trusted Execution Technology) [\$1.45]

RAM

- 64 GB RAM [\$287.10]
- 128 GB RAM [\$449.50]
- 256 GB RAM [\$635.10]
- 384 GB RAM [\$820.70]
- 512 GB RAM [\$983.10]
- 768 GB RAM [\$1,307.90]
- 1.5 TB RAM [\$2,143.10]

Figure 3-4 System configuration

4. In the **Operating System** select section, select the category Red Hat, and click the radio button for **Red Hat Enterprise Linux 7.x**, as shown in Figure 3-5 on page 36. The annotation indicating that a third-party agreement applies means that you need to accept the Terms and Conditions for Red Hat support on the order completion window.



Figure 3-5 Select Redhat

- In the **Hard Drive** section, the server comes preconfigured with a single 1 TB SATA (Serial Advanced Technology Attachment) drive, as shown in Figure 3-6. The standard configuration of Spectrum Virtualize requires two independent physical disk partitions: one for the system boot, operating system and application installation, and another for application data. *IBM Cloud preferred practices advise always protecting the boot partition with a RAID-1 mirrored configuration.* For this document, we are suggesting a similar RAID-1 configuration for the second application data disk. Technically, the installation works with only two JBOD (just a bunch of disks) disks in the system, but loss of either disk would be the equivalent of a cluster node loss. We recommend the use of RAID as defined in Figure 3-8 on page 37.

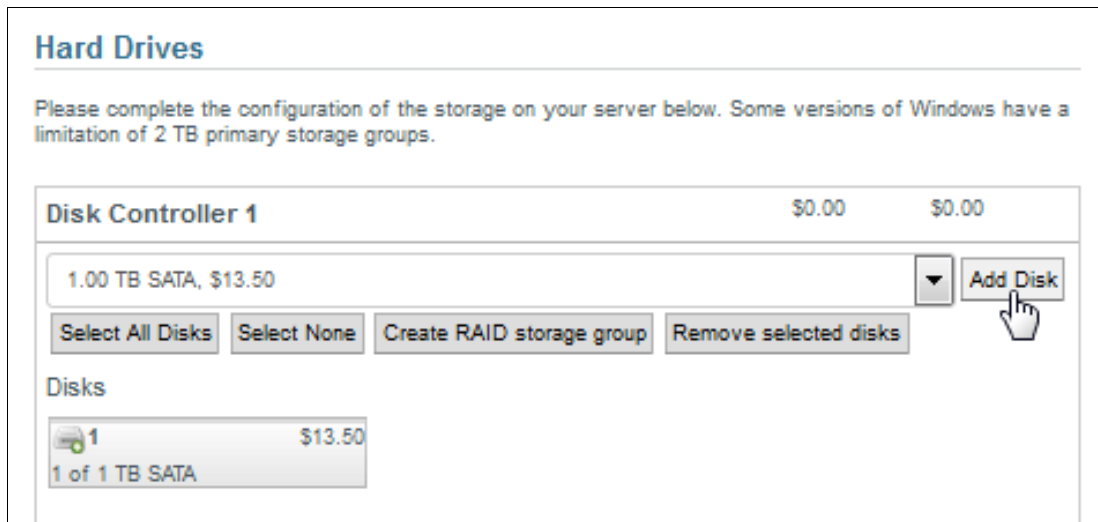


Figure 3-6 Hard Drives pane

- Click **Add Disk** three times to add three 1 TB drives. An icon for each added drive displays in the frame as you add them. Do not change the selection in the drive list for other drive types or sizes. If you have inadvertently selected and added another model drive, you can remove it by clicking the drive icon and clicking **Remove selected disks**.

After adding the four 1 TB drives, click two of the drives as shown in Figure 3-7, then select **Create RAID Storage Group**.

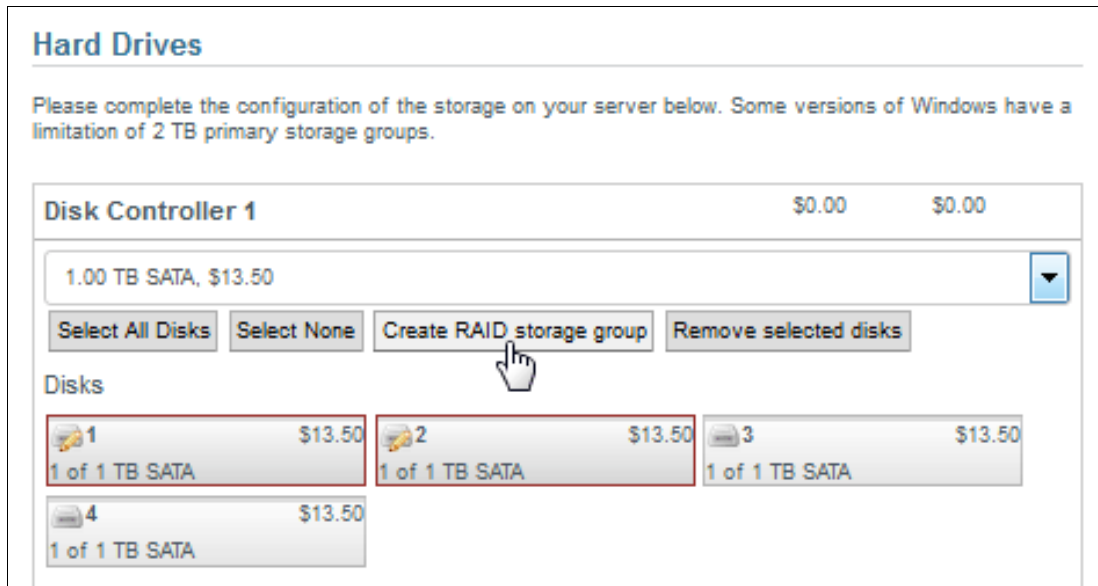


Figure 3-7 Adding the four 1 TB drives

7. Select **RAID-1** in the RAID Group configuration control selection as shown in Figure 3-8.

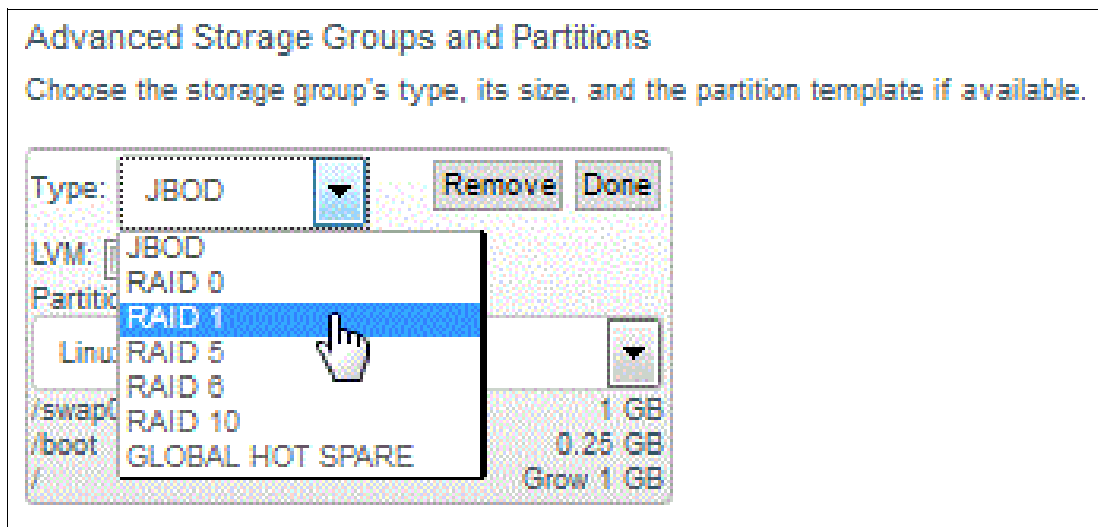


Figure 3-8 RAID configuration

8. Leave the selection for Linux Basic partition scheme set. *Do not* select the check box for the **Redhat Logical Volume Manager to be installed** on the disk partitioning, as indicated by the green box in Figure 3-9 on page 38. Select **Done** to complete the RAID partition configuration.

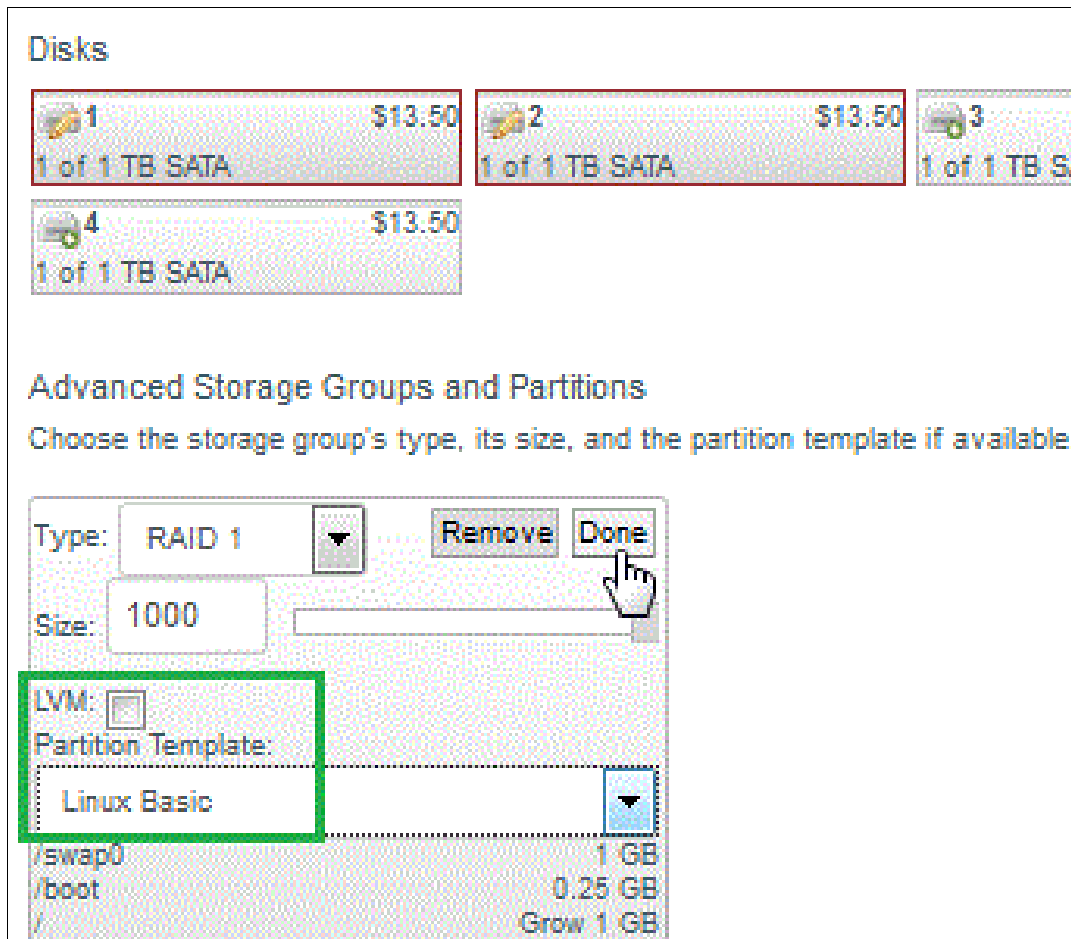


Figure 3-9 Partition Template: Linux Basic

- Repeat the process for the 3rd and 4th drives, clicking the drive icons to select them, then click **Create RAID Partition**. When the 2nd partition appears in the **Advanced Storage Groups and Partitions** control, select **RAID-1** from the pull-down menu, leave the LVM box cleared, and click **Done** to complete the process. For the second RAID partition, you are not offered a partition template but otherwise all the other choices remain the same for the second pair of drives. The resulting configuration should show two 1000 GB RAID-1 partitions configured as shown in Figure 3-10 on page 39.

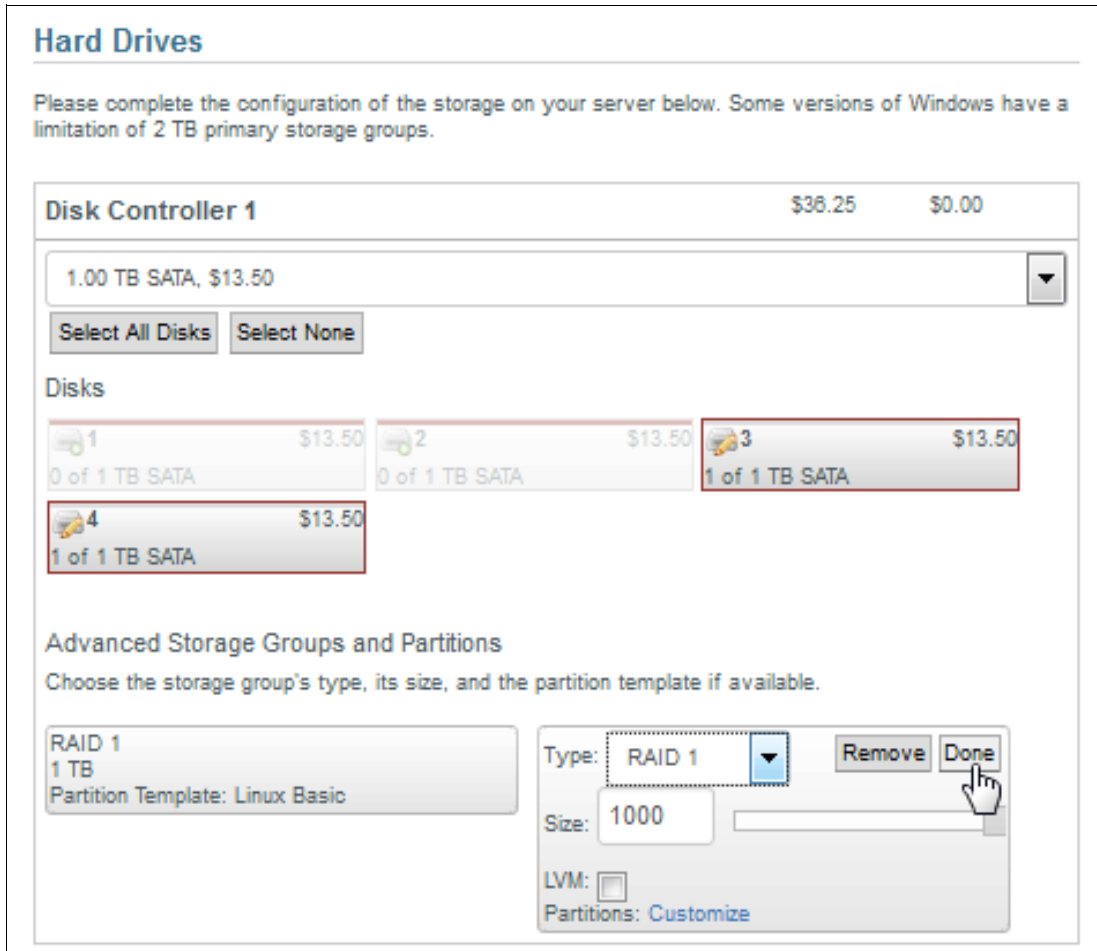


Figure 3-10 Configuring the 2nd RAID partition

10. For the networking options, several different choices are being selected. The first choice, Public Bandwidth, is presented under each of three different categories: **Limited**, **Unlimited**, and **Private Network Only**. In most circumstances, the nodes of IBM Spectrum Virtualize cluster are located in the Data Tier of a multitier virtual data center architecture; and, as a result would not be connected to the public network at all. See Figure 3-11 on page 40.

Public Bandwidth

Limited
Unlimited
Private Network Only

- 500 GB Bandwidth
- 1000 GB Bandwidth [\$32.63]
- 5000 GB Bandwidth [\$268.25]
- 10000 GB Bandwidth [\$494.45]
- 20000 GB Bandwidth [\$724.28]

Uplink Port Speeds

Public & Private

- 10 Gbps Public & Private Network Uplinks (Non Datacenter Restricted) [\$72.50]
- 100 Mbps Public & Private Network Uplinks
- 1 Gbps Public & Private Network Uplinks [\$10.00]
- 10 Gbps Public & Private Network Uplinks [\$72.50]
- 100 Mbps Redundant Public & Private Network Uplinks [\$10.00]
- 1 Gbps Redundant Public & Private Network Uplinks [\$35.00]
- 10 Gbps Redundant Public & Private Network Uplinks [\$145.00]
- 1 Gbps Public & Private Network Uplinks (Unbonded) [\$14.50]
- 100 Mbps Dual Public & Private Network Uplinks (Unbonded)
- 10 Gbps Dual Public & Private Network Uplinks (Unbonded) [\$145.00]

Figure 3-11 Networking options

When a server is connected to the public network, outgoing network traffic is metered and is subject to extra charges if the monthly data transfer exceeds the no-charge 500 GB monthly transfer volume. More data volumes can also be purchased for additional charges. The unlimited section has an option for paying a fixed fee for unmetered bandwidth. There are also more options in the IBM Cloud networking for Bandwidth Pooling; the bandwidth allocations for multiple servers can be pooled and used by a single server.

The full discussion of this feature is beyond scope of this section. However, if an internet VPN (persistent bidirectional IPsec tunnel, not simply a client VPN to access the environment for management) is to be used for replication with a remote site, the implementer should consider pooling the bandwidth of the Spectrum Virtualize servers with the network gateway (that is, a Virtual Router Appliance) required for an internet VPN.

With Bandwidth pooling, the 500 GB per month of each of the cluster servers and the network gateway server can be used by the network gateway before any excess bandwidth charges would accrue.

For this section, we selected the no charge **500 GB bandwidth** option. As of this writing, you select the **Private Network Only** option, and only have options of the number and speed for connections on the private network. Server network interfaces can be selected in either single or dual NICs.

If the server was enabled for both public and private, the default selection is for a single 1 Gbps network port on both the Public and Private networks, or just a single 1 Gbps on the Private network if Private Networks Only is selected. Dual network connections are available for both 1 Gbps and 10 Gbps Ethernet. The dual connections can be selected in HA/Bonded mode or Dual Unbonded. IBM Spectrum Virtualize requires Dual Unbonded 10 Gbps connections.

At the time of initial publication of this document, the IBM Spectrum Virtualize for Public Cloud installation procedure required access to the internet to install the software from an internet-based install server. These interfaces then needed to be protected with a firewall or disconnected from the public network.

11. Currently, it is recommend that the implementer select the **Private Network Only** option, which requires a selection of **0 GB** public bandwidth and provide a reduced set of selection options for private network connections only (see Figure 3-12 on page 42).

This means that the installation command must be run from a machine with access to the IBM Cloud private network. This can easily be accomplished by configuring the implementer's account with client VPN access and installing it on the implementer's notebook.

Alternatively, a bare-metal or virtual server can be provisioned in addition to the Spectrum Virtualize nodes and the installation can be started from that host. Getting to that host can be done again by using the implementer's machine with the VPN client or the installation host can have a public interface that is accessible from the internet.

▼ Network Options

Public Bandwidth

Private Network Only

0 GB Bandwidth

Uplink Port Speeds

Private Only

10 Gbps Private Network Uplink (Non Datacenter Restricted) [\$72.50]

100 Mbps Private Network Uplink

1 Gbps Private Network Uplink [\$5.00]

10 Gbps Private Network Uplink [\$72.50]

100 Mbps Redundant Private Network Uplinks [\$10.00]

1 Gbps Redundant Private Network Uplinks [\$29.00]

10 Gbps Redundant Private Network Uplinks [\$145.00]

10 Gbps Dual Private Network Uplinks (Unbonded) [\$145.00]

Figure 3-12 Network Options

12. After the networking interface selection, in most cases no further selections are required. In most data centers all servers are deployed with redundant power supplies and there is no option to select or deselection. However, if redundant power supply appears as an option for your server, it should be selected, as shown in Figure 3-13.

Power Supply

None

Redundant Power Supply [\$7.25]

Figure 3-13 Power supply

13. After completing the disk configuration, none of the subsequent additional service options are relevant for the IBM Spectrum Virtualize servers. However, in the service add-ons it is important to modify the response to the automatically included Cloud monitoring. **Automated Notification** should be selected rather than **Automated Reboot** as the response from a failed monitor detection.

If public network interfaces are provisioned on your server, the monitoring is performed on the public interface. The current recommendation is to configure the Spectrum Virtualize nodes with private only interfaces. However, this information might be useful for other hosts that are configured in the IBM Cloud account.

Note: Disabling or firewalling the public interface could result in your server being rebooted when the server stops responding to ping.

By default, the Cloud Account Master Account email is notified when the host fails to respond to a ping. Cloud account IDs to be notified can be added or deleted under the **Portal Accounts** → **Manage** → **Subscriptions** option after the server is provisioned (see Figure 3-14).

Service Addons

Monitoring

- Host Ping and TCP Service Monitoring
- Host Ping

Response

- Automated Reboot from Monitoring
- Automated Notification

Insurance

- None
- Business Continuity Insurance

ADD TO ORDER

Figure 3-14 Service Actions

14. When all options are selected, click **ADD TO ORDER**. An order validation window is briefly displayed. If the order validates, a new page for order completion is opened. If the validation finds missing or inconsistent server specifications, the order page is displayed with a message identifying the problematic parameters. When the order verifies, the Check Out window is displayed.

Some additional specifications are required before the order can finally be submitted. See Figure 3-15.

CHECKOUT Continue Shopping

Expand All | Collapse All

▼ Dual E5-2600 v4 Series (4 Drives) One-Time: \$0.00 Monthly: \$3,118.92
 Domain: Edit domains information for these servers Quantity: 4

Got Questions?
Speak with a Salesperson →

System Configuration Reconfigure System Configuration

Location: Dallas 10

Server: Dual Intel Xeon E5-2620 v4 (16 Cores, 2.10 GHz) \$192.13

Second Processor: Intel Xeon E5-2620 v4 (16 Cores, 2.10 GHz)

RAM: 64 GB RAM \$287.10

Operating System: Red Hat Enterprise Linux 7.x (64 bit) (per-processor licensing) \$65.25

Disk Controller: RAID \$36.25

First Hard Drive: 1.00 TB SATA \$13.50

Second Hard Drive : 1.00 TB SATA \$13.50

Third Hard Drive : 1.00 TB SATA \$13.50

Fourth Hard Drive : 1.00 TB SATA \$13.50

Network Options Reconfigure Network Options

Public Bandwidth: 500 GB Bandwidth

Uplink Port Speeds: 10 Gbps Dual Public & Private Network Uplinks (Unbonded) \$145.00

Private Network Port : 10 Gbps Dual Private Uplinks (Unbonded)

Public Network Port : 10 Gbps Dual Public Uplinks (Unbonded)

Service Addons Reconfigure Service Addons

Monitoring: Host Ping and TCP Service Monitoring

Response: Automated Reboot from Monitoring

Promotional Code:

Enter Code Apply

Totals	USD
Items	4
Total (Pre-tax)	\$3,118.92
Setup Total	\$0.00
Taxes	\$0.00
Monthly Total (Pre-tax)	\$3,118.92
Final Monthly Total	\$3,118.92
Prorated Total	\$2,599.12
Prorated Initial Charge	\$2,599.12

By clicking the Check Box, I acknowledge that the Cloud Service terms apply to this order.

I have read and agree to the Third-Party Service Agreement listed below:

- 3rd Party Software Terms Red Hat

Submit Order

[Save as Quote](#)

Totals will vary depending on current currency exchange rates.

Figure 3-15 Checkout window

15. The Checkout page provides cost elements broken out for the servers being ordered. On the right side of the page, Check boxes must be selected accepting the terms and conditions for the IBM Cloud Master Services agreement and for any third party licensed components included in the Cloud billing.

In the case of the IBM Spectrum Virtualize servers, the licensing terms for the Red Hat operating system must be accepted. This is shown in Figure 3-16. Before the order can be completed and submitted for provisioning, some further specification is required.

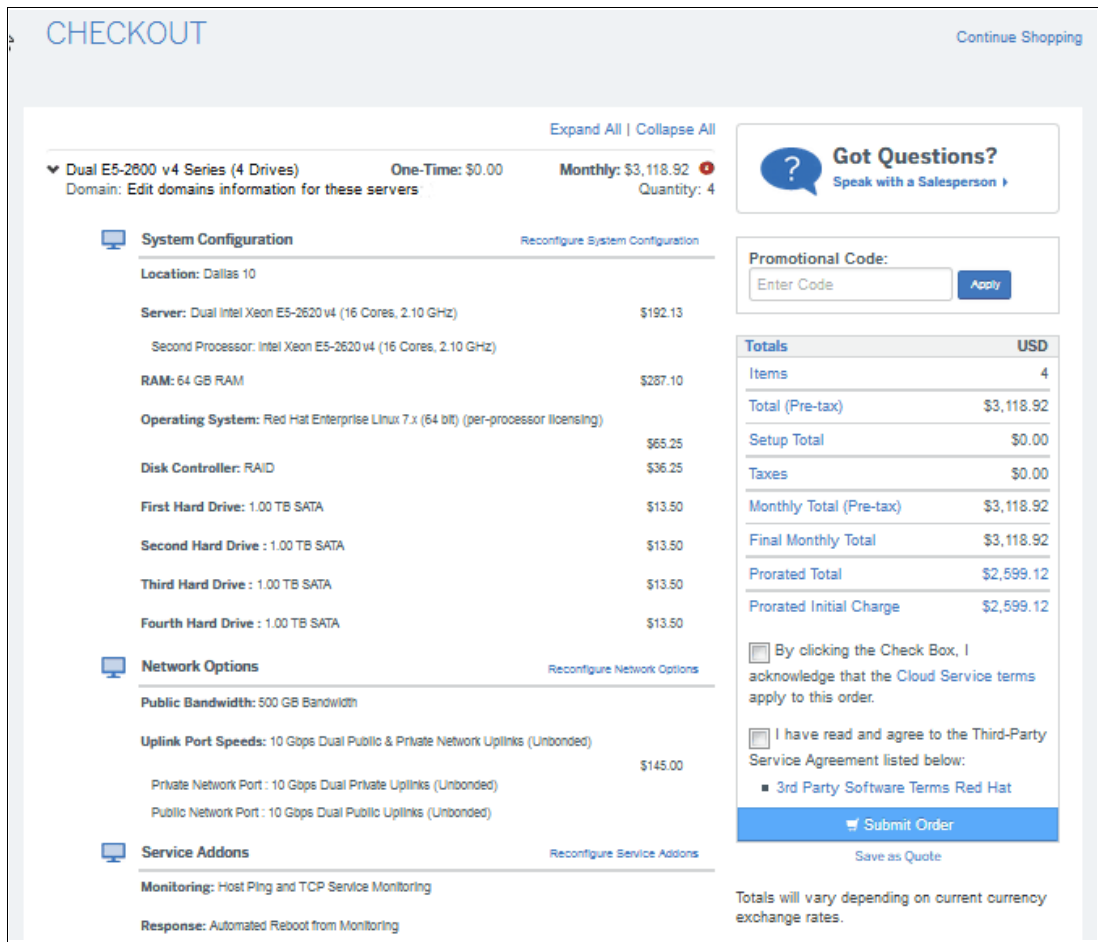


Figure 3-16 Checkout window

16. In this example, four servers are being provisioned in a single order. The check out page requires specification of host and domain names, VLANs, optional SSH keys for server logon. In this case, the servers were being ordered on IBM Cloud account in a data center where the account has not previously provisioned servers or requested pre-provisioning of VLANs. Thus fields for selecting the Front-End VLAN (Public Network) and Back-End VLAN (Private Network) for each server were not offered.

If IBM Spectrum Virtualize was being ordered for an account where servers had already been provisioned and placed on existing VLANs, input fields enabling specification for which of the existing account VLANs the servers should be placed on would have been included on the submission form.

For each server, a host name and domain name for the servers must be specified. The domain name does not have to exist or have been previously registered. The host/domain name will be associated with the private network primary subnet IP addresses assigned to the hosts on the IBM Cloud internal network but they will not be forwarded or part of external name resolution unless the user registers them (see Figure 3-17).

	Hostname	Domain	
Server 1:	ITSO_DAL10_SV_N1	ibmsvdemo.com	ex: server1.domain.com
Server 2:	ITSO_DAL10_SV_N2	ibmsvdemo.com	ex: server2.domain.com
Server 3:	ITSO_DAL10_SV_N3	ibmsvdemo.com	ex: server3.domain.com
Server 4:	ITSO_DAL10_SV_N4	ibmsvdemo.com	ex: server4.domain.com

Figure 3-17 Host and Domain names

17. When all information has been completed. The order is submitted with the **Submit** button in the right costs panel, as shown in Figure 3-18.

By clicking the Check Box, I acknowledge that the Cloud Service terms apply to this order.

I have read and agree to the Third-Party Service Agreement listed below:

- 3rd Party Software Terms Red Hat

Totals will vary depending on current currency exchange rates.

Figure 3-18 Submitting the order

18. Following the order submission, a ticket is opened for the provisioning request. The status of the server provisioning can be monitored through the account Devices menu as shown in Figure 3-19, Figure 3-20, Figure 3-21, and Figure 3-22 on page 48.

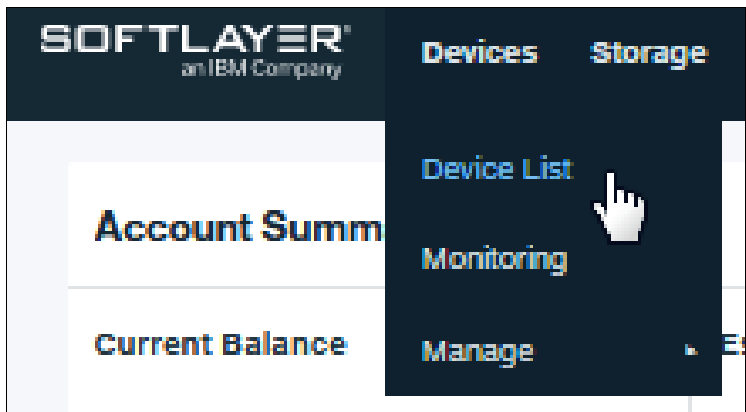


Figure 3-19 Devices menu

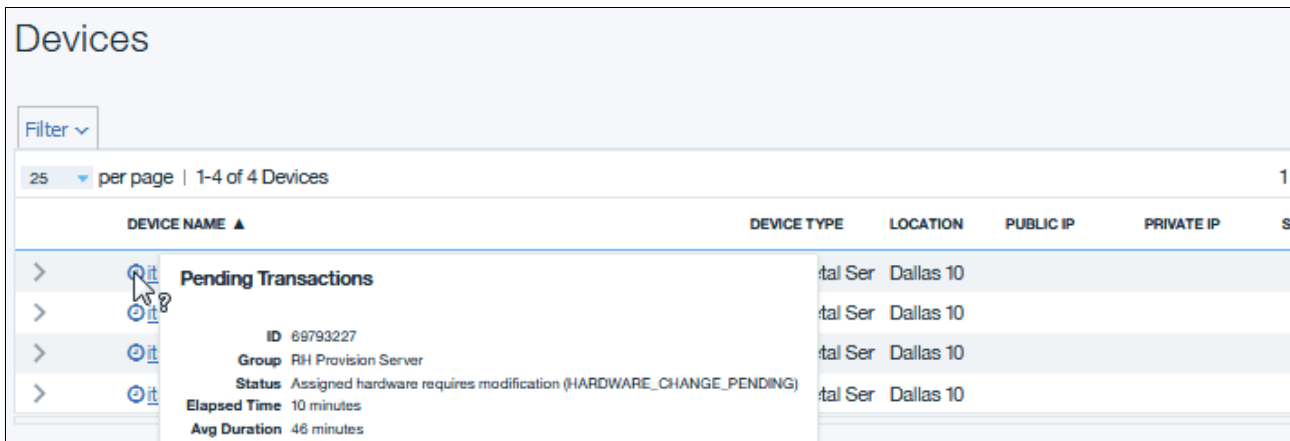


Figure 3-20 Pending transactions

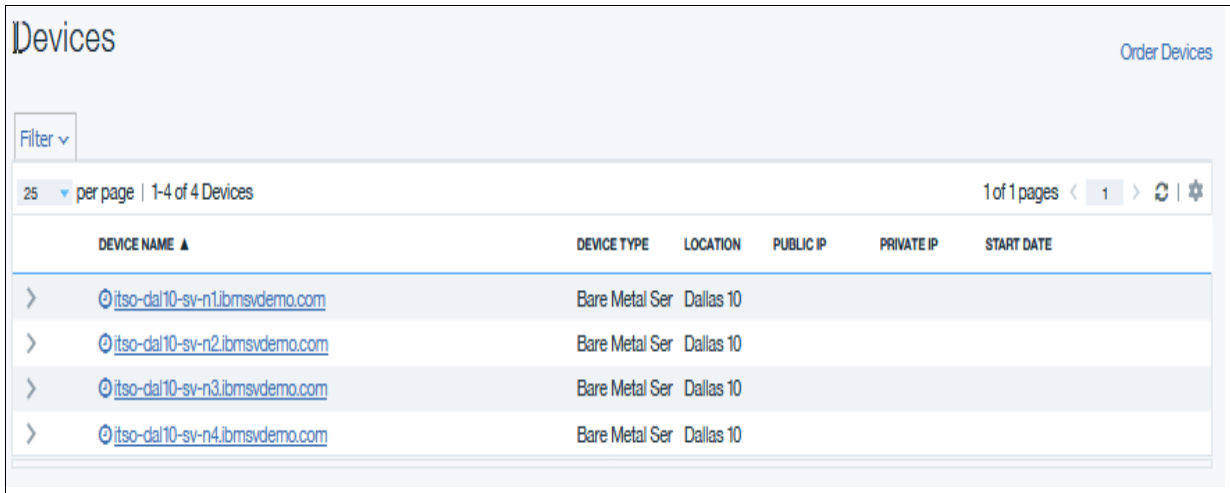


Figure 3-21 Devices list

3.2 Provisioning IBM Cloud Block Storage

This section covers provisioning of IBM Cloud Block Storage for our scenario.

3.2.1 Cloud Block Storage overview

IBM Cloud Block storage is available in two different offering products: Endurance storage and Performance storage. Both offerings provide iSCSI block storage LUNs in sizes ranging 20 GB - 12 TB in a range of input/output operations per second (IOPS) levels. The only difference between the two offerings is in how the IOPS is delivered for a given size storage volume.

With the Performance Storage offering, you can select the wanted storage volume size and then separately select the number of IOPS entitled on the volume. IOPS are provisioned in increments of 100 and can range from as low as 100 to as high as 48,000 per single volume. However, the complete range of IOPS values is not available for all volume sizes.

In Figure 3-22, The ranges of available IOPS for each of the defined LUN volume sizes is shown. Smaller volumes have a lower maximum IOPS that can be provisioned. Conversely, larger volumes have higher minimum IOPS that can be provisioned. The gray areas in Figure 3-22 indicate unsupported volume size and IOPS combinations.

		Available IOPS Range Per LUN											
		100	200	300	500	1000	2000	4000	6000	10000	20000	40000	48000
Storage LUN Size in GB	20	5.00	10.00	15.00	25.00	50.00							
	40	2.50	5.00	7.50	12.50	25.00	50.00						
	80	1.25	2.50	3.75	6.25	12.50	25.00						
	100	1.00	2.00	3.00	5.00	10.00	20.00	40.00					
	250	0.40	0.80	1.20	2.00	4.00	8.00	16.00					
	500	0.20	0.40	0.60	1.00	2.00	4.00	8.00	12.00	20.00			
	1000	0.10	0.20	0.30	0.50	1.00	2.00	4.00	6.00	10.00	20.00		
	2000		0.10	0.15	0.25	0.50	1.00	2.00	3.00	5.00	10.00	20.00	
	3000		0.07	0.10	0.17	0.33	0.67	1.33	2.00	3.33	6.67	13.33	16.00
	4000			0.08	0.13	0.25	0.50	1.00	1.50	2.50	5.00	10.00	12.00
	5000			0.06	0.10	0.20	0.40	0.80	1.20	2.00	4.00	8.00	9.60
	6000			0.05	0.08	0.17	0.33	0.67	1.00	1.67	3.33	6.67	8.00
	7000			0.04	0.07	0.14	0.29	0.57	0.86	1.43	2.86	5.71	6.86
	8000				0.06	0.13	0.25	0.50	0.75	1.25	2.50	5.00	6.00
	9000				0.06	0.11	0.22	0.44	0.67	1.11	2.22	4.44	5.33
	10000					0.10	0.20	0.40	0.60	1.00	2.00	4.00	4.80
11000					0.09	0.18	0.36	0.55	0.91	1.82	3.64	4.36	
12000					0.08	0.17	0.33	0.50	0.83	1.67	3.33	4.00	

Figure 3-22 Performance storage

The green areas indicate volume size and IOPS ranges that are only available in the newer, higher performance storage data centers. The numbers in the cells indicate the equivalent IO density of IOPS/volume size combinations. As a rule, any I/O Density greater or equal to 2 is implemented on SSD and has a lower latency than the storage with I/O density lower than 2.

Endurance storage offers the same set of predefined volume sizes as Performance Storage. With Endurance storage, volumes are provisioned in one of four Storage Tiers, which are defined by their IO Density: 0.25, 2.0, 4.0, and 10.0 IOPS per GB, with the first tier only on spinning disk and remaining tiers on SSD.

With the storage tiers defined in IOPS per GB of capacity, the IOPS delivered on a given LUN depends on the size of the LUN and storage tier in which it is provisioned.

In Figure 3-23 the IOPS delivered for each of the standard LUN sizes and storage tiers is shown. As described previously, the green cells indicate combinations of LUN size and storage tier that are only available in the high-performance storage IBM Cloud data centers.

		IOPS Per LUN by Tier			
		0.25	2	4	10
Storage LUN Size in GB	Storage Tiers-->				
	20	5	40	80	200
	40	10	80	160	400
	80	20	160	320	800
	100	25	200	400	1000
	250	62.5	500	1000	2500
	500	125	1000	2000	5000
	1000	250	2000	4000	10000
	2000	500	4000	8000	20000
	3000	750	6000	12000	30000
	4000	1000	8000	16000	40000
	5000	1250	10000	20000	
	6000	1500	12000	24000	
	7000	1750	14000	28000	
	8000	2000	16000	32000	
	9000	2250	18000	36000	
	10000	2500	20000	40000	
11000	2750	22000	44000		
12000	3000	24000	48000		

Figure 3-23 Endurance Storage

In planning for your IBM Spectrum Virtualize deployment, the implementer should have a target capacity and anticipated IO storage density in mind from the customer requirements. In a typical customer environment, the data center has multiple tiers of storage and usage profile defining the percentage of total capacity in each tier.

An example profile of storage classes and customer allocation is shown in Table 3-1.

Table 3-1 An example profile of storage classes and customer allocation

Class	Capacity	Density	Latency
A	2 TB	8 IOPS/GB	< 1 ms
B	20 TB	4 IOPS/GB	< 2 ms
C	50 TB	2 IOPS/GB	< 5 ms
D	0	0.25 IOPS/GB	> 5 ms

When planning implementation of a IBM Spectrum Virtualize infrastructure the implementer should provision multiple storage volumes whose combined capacity provides the capacity and IOPS totals required for the storage tier. IBM Spectrum Virtualize distributes the I/O workload over multiple volumes allocated in disk group, with a preference for multiples of four.

So, for the example above four 500 GB LUNs, where each LUN delivers 4000 IOPS combined in a disk group provides 2 TBs and 16,000 IOPS, that is the 8 IOPS per GB storage density required for class A. Alternatively, 8 x 250 GB volumes where each LUN delivers 2000 IOPS would similarly meet the customer Class A space and performance requirement. Where exact multiples of four disks do not meet requirements, multiples of 2 disks can also be used. Disk groups should not be provisioned with odd numbers of disks in the group.

Before provisioning the storage for your implementation, you should determine appropriate number, size, and IOPS for the volumes in your disk groups. From the tables of available LUN sizes and IOPS provisioning listed above, it should be clear Performance Storage provides much greater flexibility for selecting both a capacity and IOPS per LUN.

However, if the customer only needs storage at a specific storage tier, Endurance storage can provide a simpler solution. Both offerings can be used by IBM Spectrum Virtualize interchangeably. Recognize, too, that IBM Spectrum Virtualize aggregates both the capacity and the IOPS of the LUNs that are configured together in a disk group.

3.2.2 Provisioning Block Volumes

This section walks through the IBM Cloud Portal options and screens used to provision block storage volumes.

1. From the IBM Cloud Portal home window, select and click **Storage**, as shown in Figure 3-24.

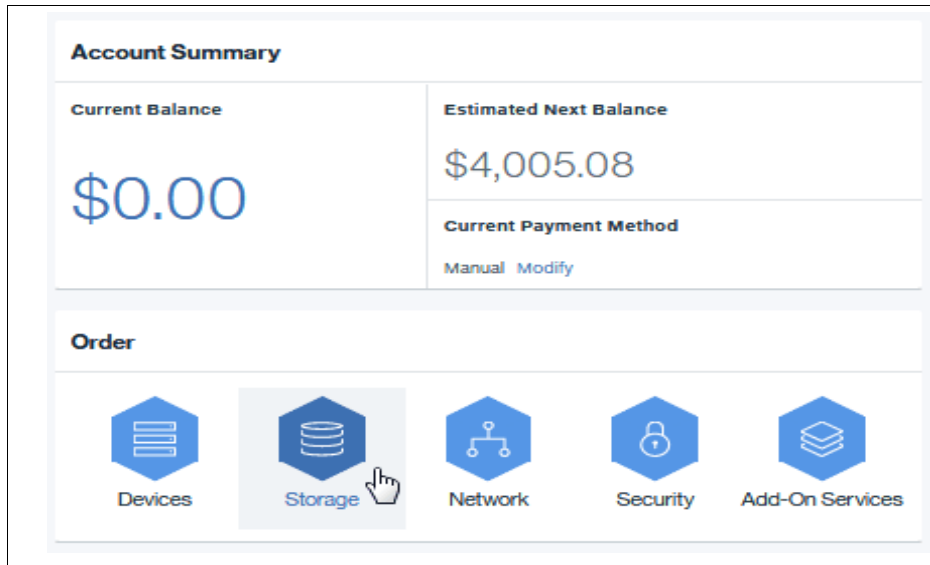


Figure 3-24 Select Storage

2. On the **Block Storage** window, a list of all the storage volumes already allocated on the account are shown. Click **Order Block Storage**, as shown in Figure 3-25.

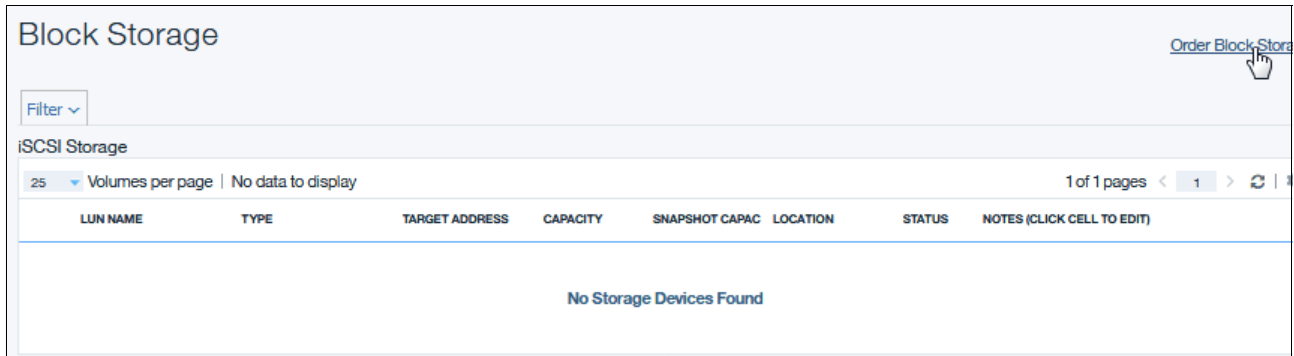


Figure 3-25 Order Block Storage

3. In the Order Block Storage control, click the selection list in the **Select Storage Type** field. For this example, we are using **Performance** storage. **Endurance** storage is also an option. **Portable** Storage is not recommended and is not available in newer data centers. Select **Monthly billing** unless you are planning to use the storage for a short time only.

Hourly billing allows storage to be deprovisioned without charging through the end of the month, but is more expensive than Monthly Billing if it used or provisioned for a full month period. There might be scenarios where this is desirable, but you should be aware of the higher costs if you use it (see Figure 3-26).

Order Block Storage

Select Storage Type **All fields are required.**

Performance

Endurance

Performance

Portable Storage

DAL10 - Dallas

Data at rest is encrypted by default for this datacenter.

Billing Method

*Hourly billing is only available in updated datacenters with improved capabilities.

Monthly Billing

Hourly Billing

Figure 3-26 Monthly billing

- When prompted, select the data center where the storage should be provisioned. This must be the same data center where your IBM Spectrum Virtualize servers are provisioned. Note the asterisks next to some data center names. These are the data centers where the high-performance storage options are deployed: storage with the capacity and performance options indicated in the green cells of the tables in the Cloud Storage overview (see Figure 3-27).

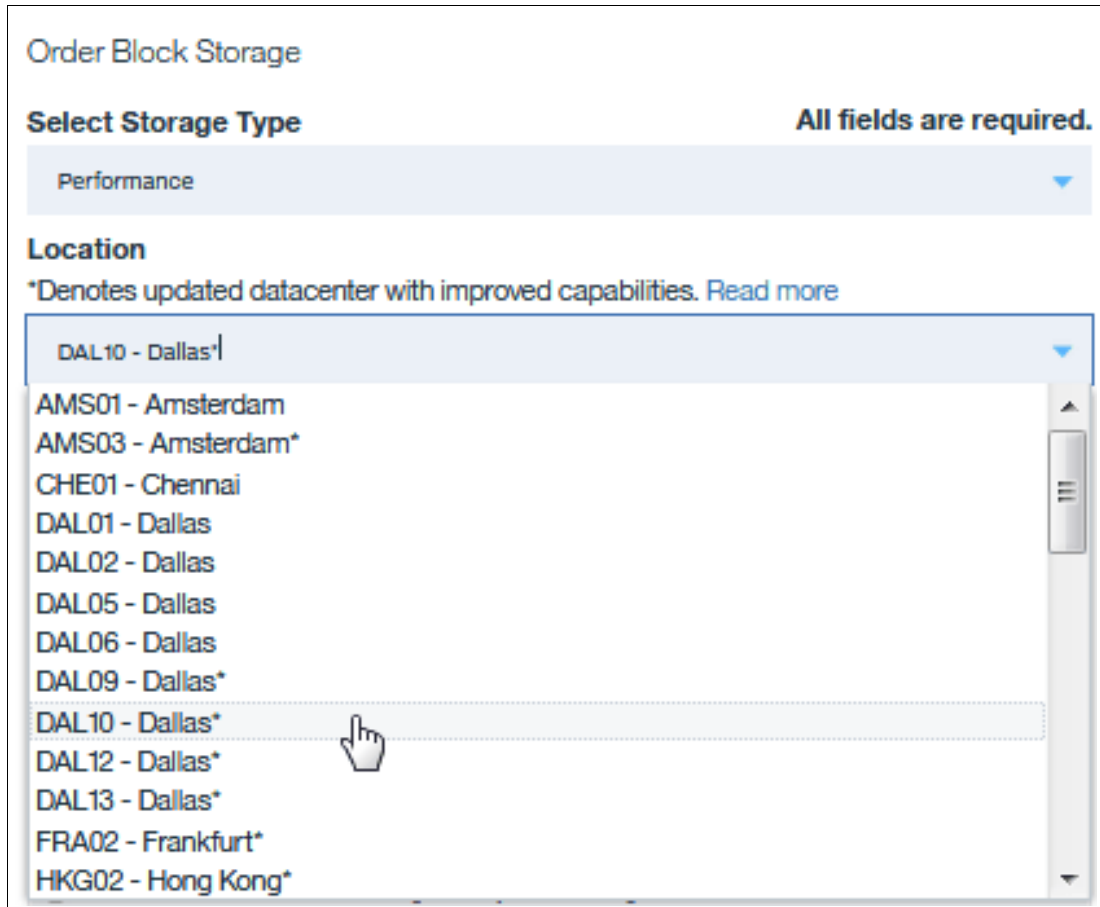


Figure 3-27 Select the data center where the storage should be provisioned

- Select the size of the volume to provision. Note the range of minimum and maximum IOPS that are available for the size of the volume you have selected. Ensure the IOPS level required is available for the size being selected. If the IOPS required is not available, reconsider the number, size, and IOPS for the volumes intended for your disk group (see Figure 3-28 on page 54).

Select Storage Size

DESCRIPTION
<input type="radio"/> 20 GB - 100 to 1000 IOPS [\$1.45 per month]
<input type="radio"/> 40 GB - 100 to 2000 IOPS [\$2.90 per month]
<input type="radio"/> 80 GB - 100 to 4000 IOPS [\$5.80 per month]
<input type="radio"/> 100 GB - 100 to 6000 IOPS [\$7.25 per month]
<input type="radio"/> 250 GB - 100 to 6000 IOPS [\$18.13 per month]
<input checked="" type="radio"/> 500 GB - 100 to 10000 IOPS [\$36.25 per month]
<input type="radio"/> 1000 GB - 100 to 20000 IOPS [\$72.50 per month]
<input type="radio"/> 2000 GB - 200 to 40000 IOPS [\$145.00 per month]
<input type="radio"/> 3000 GB - 200 to 48000 IOPS [\$217.50 per month]
<input type="radio"/> 4000 GB - 300 to 48000 IOPS [\$290.00 per month]
<input type="radio"/> 5000 GB - 300 to 48000 IOPS [\$362.50 per month]
<input type="radio"/> 6000 GB - 300 to 48000 IOPS [\$435.00 per month]
<input type="radio"/> 7000 GB - 300 to 48000 IOPS [\$507.50 per month]
<input type="radio"/> 8000 GB - 500 to 48000 IOPS [\$580.00 per month]
<input type="radio"/> 9000 GB - 500 to 48000 IOPS [\$652.50 per month]
<input type="radio"/> 10000 GB - 1000 to 48000 IOPS [\$725.00 per month]
<input type="radio"/> 11000 GB - 1000 to 48000 IOPS [\$797.50 per month]
<input type="radio"/> 12000 GB - 1000 to 48000 IOPS [\$870.00 per month]

Figure 3-28 Select Storage Size

6. Enter the requested IOPS for the volume. Any value evenly divisible by 100 can be entered between the minimum and maximum IOPS allowed for the size volume selected. The quantity of *snapshot space* should always be left or set to zero. Snapshot space is for Cloud Storage functionality that is not used, and is incompatible with IBM Spectrum Virtualize. Select **Linux** formatting for the storage volume, as shown in Figure 3-29 on page 55.

Specify IOPS

Enter between 100 and 48000 by intervals of 100 as determined by your selection above

8000 [\$406.00 per month]

8000

Specify Snapshot Space Size

Defines how much storage you want to use for storing snapshots.

0 GB

Select OS Type

Defines the type of data the LUN will store.

Linux

Figure 3-29 Select IOPS

- After entering all the specifications for the storage, an order confirmation window prompts you to confirm your selections and acceptance of the Cloud Services terms and conditions, as shown in Figure 3-30.

Order Review

Item	Storage as a Service
Item	Block Storage
Item	500 GBs
Item	8000 IOPS
Location	Dallas 10
Quantity	1

I have read the [Master Service Agreement](#) and agree to the terms therein.

Previous
Cancel
Place Order

Figure 3-30 Order Review

- Upon placing the order, your provisioning begins. Usual provisioning time is approximately 5 minutes. When provisioning completes, the ordered volume displays in the device list under the **Storage** option on the portal home window (see Figure 3-31).

You must complete the above procedure for each storage LUN to be incorporated in your Disk Group. At this point, there is no way to request provisioning of multiple identical LUNs (see Figure 3-31).

LUN NAME	TYPE	TARGET ADDRESS	CAPACITY	SNAPSHOT CAPAC	LOCATION	STATUS	NOTES (CLICK CELL TO EDIT)
IBM02SL1205109-1	Performance	161.26.98.88	500.0 GB	0 GB	Dallas 10	Active	Actions
IBM02SL1205109-2	Performance	161.26.98.87	500.0 GB	0 GB	Dallas 10	Active	Actions
IBM02SL1205109-3	Performance	161.26.98.88	500.0 GB	0 GB	Dallas 10	Active	Actions
IBM02SL1205109-4	Performance	161.26.98.86	500.0 GB	0 GB	Dallas 10	Active	Actions
IBM02SL1205109-5	Performance	161.26.98.85	500.0 GB	0 GB	Dallas 10	Active	Actions
IBM02SL1205109-6	Performance	161.26.98.86	500.0 GB	0 GB	Dallas 10	Active	Actions
IBM02SL1205109-7	Performance	161.26.98.88	500.0 GB	0 GB	Dallas 10	Active	Actions
IBM02SL1205109-8	Performance	161.26.98.87	500.0 GB	0 GB	Dallas 10	Active	Actions

Figure 3-31 Block Storage

- In Figure 3-32 on page 57, the specification window for Endurance storage is shown. In the case of Endurance storage the location, size, storage tier, and format type are all specified in a single control. By selecting a Storage Tier based on storage density, the number of IOPS delivered on the LUN is a function of the size of the LUN, as shown in Figure 3-23 on page 49.

Order Block Storage

Select Storage Type **All fields are required.**

Endurance ▼

Designed for higher-intensity workloads.

Location

*Denotes updated datacenter with improved capabilities. [Read more](#)

Select Location ▼

Billing Method

*Hourly billing is only available in updated datacenters with improved capabilities.

Monthly Billing

Hourly Billing

Select Storage Package

DESCRIPTION
<input type="radio"/> 0.25 IOPS/GB ⓘ
<input type="radio"/> 2 IOPS/GB ⓘ
<input type="radio"/> 4 IOPS/GB ⓘ

Select Storage Size

▼

Specify Snapshot Space Size

Defines how much storage you want to use for storing snapshots.

▼

Select OS Type

Defines the type of data the LUN will store.

Linux ▼

Figure 3-32 Order Block Storage

3.3 IBM Spectrum Virtualize networking considerations

IBM Spectrum Virtualize for IBM Cloud supports two, four, six, or eight-node cluster configurations in IBM Cloud. All nodes in a cluster must be provisioned on the same public and private VLANs. This can be accomplished by either ordering all nodes in the cluster in a single order of multiple machines, or by specifying the front-end and back-end VLANs on verified order page of the Cloud portal device order.

The automated installation procedure automatically creates the appropriate portable subnet on the private VLAN that the servers for the IBM Spectrum Virtualize cluster have been provisioned on. For each server in the cluster, five unique IP addresses are allocated and configured onto the server on the portable private subnet. In addition, a sixth IP address is configured on each server, but it is a cluster address with the same address value shared by all the nodes in the cluster.

For reinstallations, the semi-auto or manual install must be used, unless the creation of a portable subnet is needed or wanted. If the IPsec tunnel between on-prem and IBM Cloud is configured for the original installation and the portable private subnet that is associated with that installation, a new portable private subnet requires reconfiguration of the Virtual Routing Appliance and the on-prem device that serves as the other end of the IPsec tunnel.

If the semi-automated or manual installation procedures are used instead of the automated install, the IT manager needs to request the portable subnet through the cloud portal (unless this is a reinstallation) and allocate the needed IP addresses. In this case, “allocating” the address is simply the choosing which IP addresses in the subnet range are assigned to each of the five addressees needed for each node and the single shared address for the cluster.

Address allocation and inventory keeping for portable subnets is not performed or maintained by the cloud portal. The cloud network routers accept and process which addresses are configured on the server NICs.

When VLANs are provisioned on the cloud account, they are initially setup with only a single subnet, called the primary subnet. Addresses on the on primary subnet can only be assigned by the cloud provisioning engine. When servers are provisioned with multiple, unbonded Network Interface Cards (NICs), only the first NIC on the network (Public or Private) is assigned an IP address. If additional IP addresses are needed for the server (for example, IP addresses for a second NIC, or IP addresses for an HA cluster or VMs running on the host), a portable or secondary subnet must be allocated on the VLAN. IBM Spectrum Virtualize automated install.

For IBM Spectrum Virtualize, five IP addresses are needed for each host node in the IBM Spectrum Virtualize cluster, plus a sixth address for the cluster that will be used by all nodes. The following section shows how to allocate a private subnet required for allocating an assigning these addresses. This procedure is only required when the manual or semi automated installation procedure is used for IBM Spectrum Virtualize. The fully automated procedure includes logic to allocate the subnet.

Complete the following steps:

1. From the Cloud Portal home window, select the **Network** option, as shown in Figure 3-33.

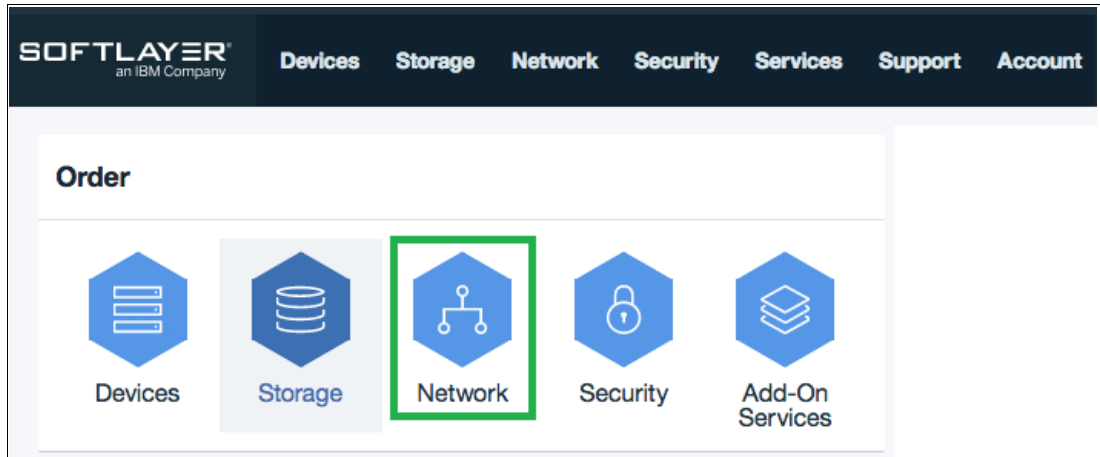


Figure 3-33 Select the Network option

2. At the bottom of the Network options page, select **Order** under the **Subnets/IPs** section (see Figure 3-34).

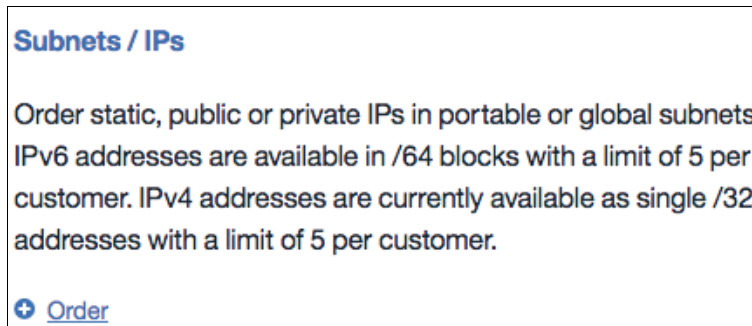


Figure 3-34 Subnets/IPs

3. The additional IP addresses for the IBM Spectrum Virtualize nodes is allocated on the private network. Therefore, a Portable Subnet is needed on the Private Network. Select **Portable Private**, as shown in Figure 3-35.

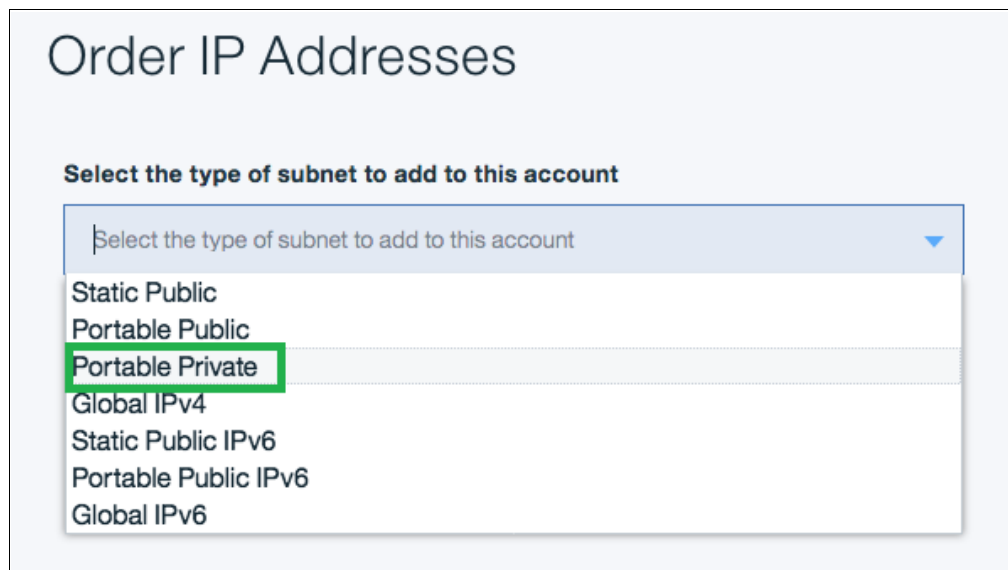


Figure 3-35 Order IP Addresses

4. For this example, we have allocated a 64-address subnet. The automated IBM Spectrum Virtualize installation script also always creates a 64-address subnet by default. This accommodates even an eight-node cluster, which requires 41 addresses. Additional addresses could potentially be allocated for storage clients to access IBM Spectrum Virtualize iSCSI targets on the same subnet, eliminating the need for routing between clients and IBM Spectrum Virtualize (see Figure 3-36).

The screenshot shows a web form titled "Order IP Addresses". Under the heading "Select the type of subnet to add to this account", there is a dropdown menu with "Portable Private" selected. Below this is a red asterisk and the text "Required field". Under the heading "Select Option", there is a list of five radio button options: "4 Portable Private IP Addresses", "8 Portable Private IP Addresses", "16 Portable Private IP Addresses", "32 Portable Private IP Addresses", and "64 Portable Private IP Addresses". The "64 Portable Private IP Addresses" option is selected, indicated by a blue dot in the radio button.

Figure 3-36 Select 64 Portable Private IP Addresses

After the number of IP addresses is selected, the user is presented with a list of Private Network VLANs already provisioned on the account. In this example, the demonstration cloud account was no longer available. Therefore, the data center and VLAN number is different than other examples. However, you are presented with a list of all VLANs available on the account, including those in other data centers.

5. Select the private VLAN where your IBM Spectrum Virtualize servers have been provisioned, as shown in Figure 3-37.

Order IP Addresses

Ordering **64 Portable Private IP Addresses**
[Go back](#) to change selection.
Select a VLAN to establish where the new IP addresses are routed. Expand each row

25 Items per page | 1-1 of 1 VLANs

VLAN NUMBER	PRIMARY ROUTER	LOCATION
> <input checked="" type="radio"/> 811	bcr06a.dal09	Dallas 9

Figure 3-37 Select the private VLAN

6. A justification questionnaire is required for the subnet. Complete the form with information appropriate to your intended use. Additional addresses are a semi-constrained resource on the IBM Cloud, and the information allows automated planning processes to determine when allocated address spaces can be deprovisioned (see Figure 3-38).

How many of these new IPs will be in use within the next 30 days:

* Required field

How many of these new IPs will be in use within the next 12 months:

* Required field

Description of your need for additional IPs:

* Required field

Contact Name:

* Required field

Contact Job Title:

* Required field

Contact Email:

* Required field

Contact Phone Number:

* Required field

I agree that the contact information is current and valid.

Figure 3-38 Justification questionnaire

7. Accept the Cloud Agreement terms and complete the order. The subnet is available immediately, as shown in Figure 3-39.

The screenshot shows an 'Order Review' section with a table of items. The table has two rows: 'Item' with the value '64 Portable Private IP Addresses' and 'Quantity' with the value '1'. Below the table is a checkbox that is currently unchecked, followed by the text 'I have read the Master Service Agreement and agree to the terms therein.' Below this text is a blue button labeled 'Previous'.

Order Review	
Item	64 Portable Private IP Addresses
Quantity	1

I have read the [Master Service Agreement](#) and agree to the terms therein.

Previous

Figure 3-39 Order Review

3.3.1 Provisioning a Network Gateway Appliance

Some scenarios for using IBM Spectrum Virtualize on the IBM Cloud involve replicating with an IBM Spectrum Virtualize or other IBM Virtualized storage appliance in an on-premises environment. Connectivity with external environments can be accomplished through either an internet VPN or a private Direct Link connection. If an internet VPN is used, it is necessary to deploy a network gateway appliance to terminate the VPN from the customer's environment. The Gateway Appliance also provides firewall and Network Address Translation capabilities, if required for networking with the customer environment.

A Network Gateway, sometimes called a Vyatta, is only required if the IBM Spectrum Virtualize is to be configured for replication through an internet VPN. If your IBM Spectrum Virtualize is for a single site only, is replicating with a IBM Spectrum Virtualize in another IBM Cloud data center, or a private Direct Link connection is being provisioned between the customer network and the IBM Cloud, then a Network Gateway is not required.

To provision a network gateway, complete the following steps:

1. Select the Network menu from the portal home window and select **Network Appliances**, as shown in Figure 3-40.

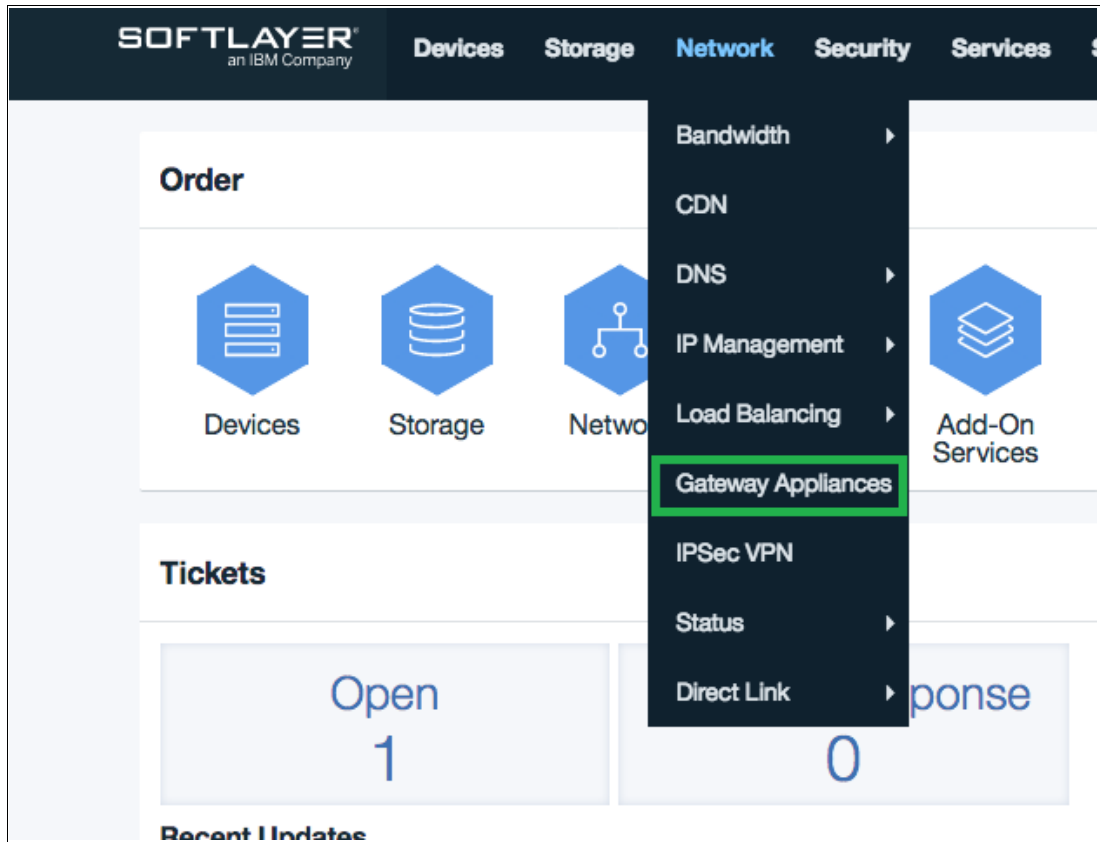


Figure 3-40 Select Network Appliances

2. A list of already provisioned appliances is displayed. Normally one would not expect any to be listed, but if the customer cloud account has been zoned into multiple firewalled separate VLANs, there might be multiple already provisioned. Select **Order Gateway** from the upper right corner of the page (see Figure 3-41).



Figure 3-41 Select Order Gateway

- The Network Gateway is just another instance of a cloud bare-metal server. The same list of servers available for provisioning, as was provided for the IBM Spectrum Virtualize node selection, is presented (see Figure 3-42).

Dual Processor Multi-Core Servers

MANUFACTURER	CPU MODEL	CORE COUNT	CPU SPEED	STORAGE	RAM	PRIVATE NETWORK ONLY	10 GBPS UPLINK SUPPORT	STARTING PRICE PER MONTH
Intel	Intel Xeon E5-2620 v416 Cores		2.10 GHz	Up to 12 drives	16 GB up to 256 GB-		Yes	\$452.41
Intel	Intel Xeon E5-2650 v424 Cores		2.20 GHz	Up to 12 drives	16 GB up to 256 GB-		Yes	\$510.00
Intel	Intel Xeon E5-2690 v428 Cores		2.60 GHz	Up to 12 drives	16 GB up to 256 GB-		Yes	\$575.66
Intel	Intel Xeon E5-2620 v312 Cores		2.40 GHz	Up to 12 drives	16 GB up to 256 GB-		Yes	\$665.56
Intel	Intel Xeon E5-2650 v320 Cores		2.30 GHz	Up to 12 drives	16 GB up to 256 GB-		Yes	\$723.56

Figure 3-42 Dual Processor servers

- For the gateway, a Dual Processor server is advised when an IPsec VPN will be terminated on the gateway. The RAM is suggested at 64 GB for IPsec VPN (see Figure 3-43).

Operating System

Brocade

Virtual Router Appliance 5.x (up to 1Gbps) Subscription Edition (64 Bit) [\$158.78]

The following Third-Party Service Agreement applies to the product selected above:
[3rd Party Software Terms Vyatta](#)

Figure 3-43 Operating System selection

- Only two drives are required, configured in RAID1 with Linux Basic partition map, as shown in Figure 3-44 on page 67.

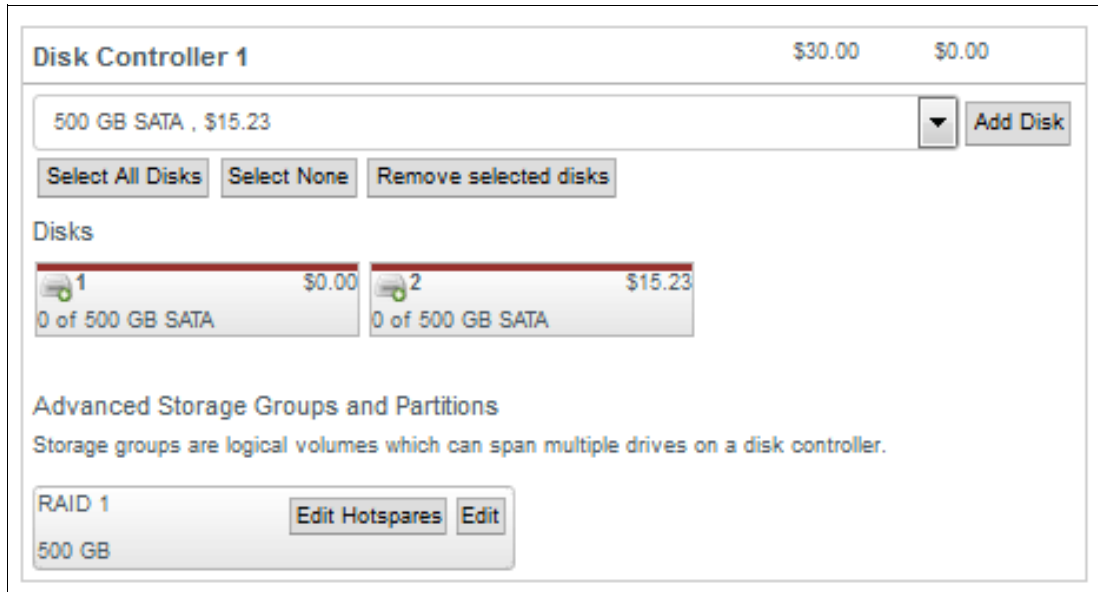


Figure 3-44 Disk Controller 1

6. A Network Gateway must be configured with both Public and Private networks. In Figure 3-45, a 1 Gbps Redundant (Bonded) network connection is selected. Depending on the intended replication data volume, a 10 Gbps connection might be wanted, but a redundant connection should be selected. This differs from the configuration used with the IBM Spectrum Virtualize cluster hosts.

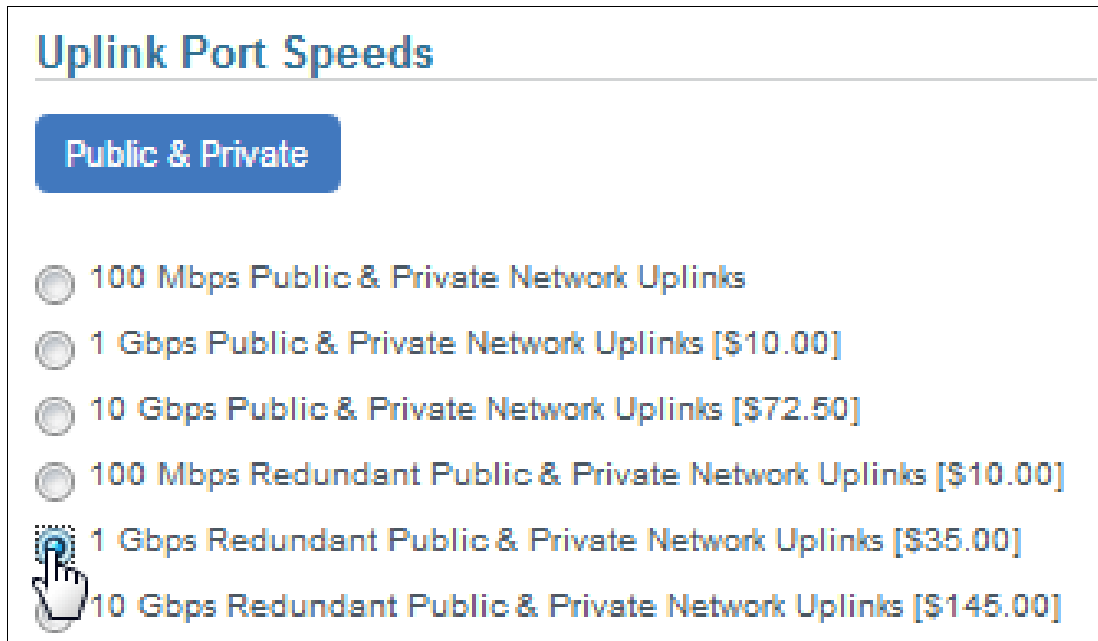


Figure 3-45 Uplink Port Speeds

7. Complete the server configuration and verify order (see Figure 3-46).

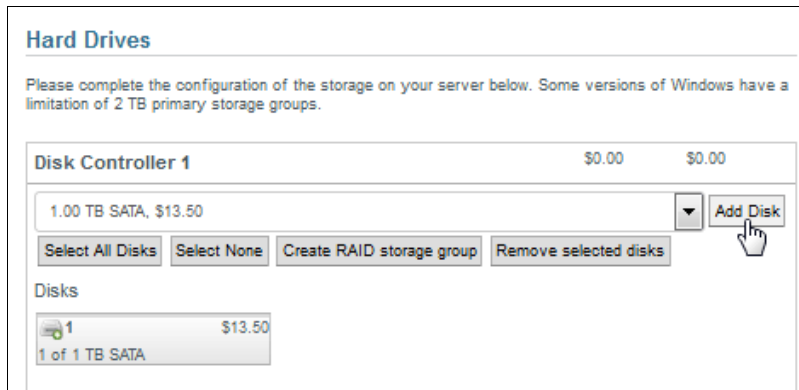


Figure 3-46 Complete the server configuration

8. On the order completion page, you are required to specify VLANs that the gateway will manage. The pull-down selection for the back-end and front-end VLANs allows either automatic assignment or selection from one of the existing VLANs on the account. The Network Gateway should be provisioned after the IBM Spectrum Virtualize cluster hosts have completed provisioning.

Select the back-end and front-end VLANs on which the IBM Spectrum Virtualize servers were placed when they provisioned. This action makes the Network Gateway server the default router for all subnets on those VLANs. This has several effects on the network environment that are explained in the Network Gateway configuration section of this document.

Assign a host and domain name to the Network Gateway. These names resolve in the IBM Cloud internal network, but are not published to any externally visible domain name servers. They are mainly used for naming within the IBM Cloud portal inventory and device listing screens (see Figure 3-47).


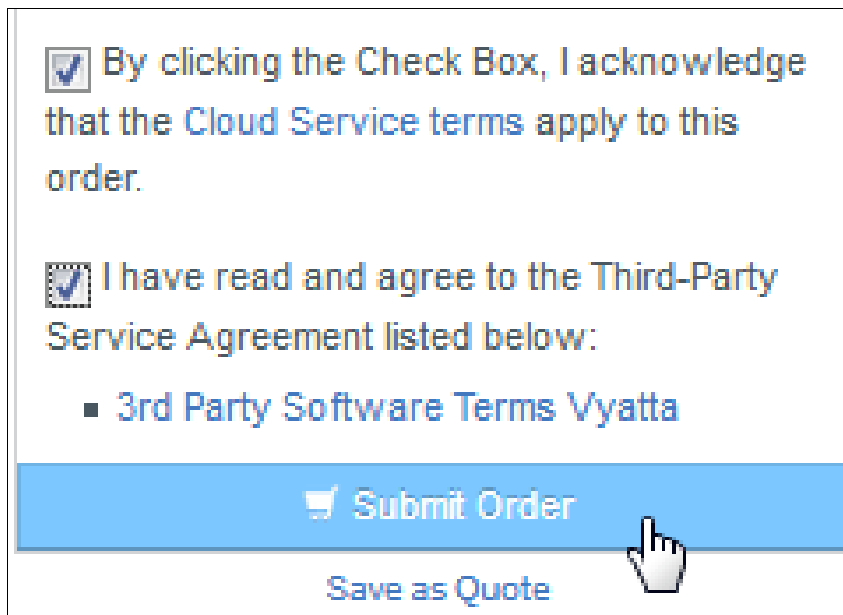


Figure 3-47 Host and Domain Names

9. Complete the order for the gateway, as shown in Figure 3-48.



By clicking the Check Box, I acknowledge that the Cloud Service terms apply to this order.

I have read and agree to the Third-Party Service Agreement listed below:

- 3rd Party Software Terms Vyatta

Figure 3-48 Submit Order

Servers might take as long as four hours to complete provisioning; however, simple, small servers, such as the Vyatta, often require less than one hour.



Implementation

This chapter describes how to implement IBM Spectrum Virtualize for Public Cloud environment and provides detailed instructions about the following topics:

- ▶ Downloading the One-click installer
- ▶ Fully Automated installation
- ▶ Semi Automated installation
- ▶ Configuring Spectrum Virtualize
- ▶ Configure Cloud quorum
- ▶ Configure the back-end storage
- ▶ Configuring Call Home with CLI

For more information about the Home Call configuration, see Chapter 6, “Supporting the solution” on page 135.

This chapter has the following topics:

- ▶ 4.1, “IBM Spectrum Virtualize for Public Cloud installation” on page 72
- ▶ 4.2, “Configuring Spectrum Virtualize” on page 82
- ▶ 4.3, “Configuring replication from on-prem IBM Spectrum Virtualize to IBM Spectrum Virtualize for IBM Cloud” on page 105
- ▶ 4.4, “Configuring Remote Support Proxy” on page 115

4.1 IBM Spectrum Virtualize for Public Cloud installation

This section includes instruction for implementing IBM Spectrum Virtualize in IBM Cloud. IBM Spectrum Virtualize for Public Cloud implementation starts from the following assumptions:

- ▶ The required bare-metal server resources that are described in Chapter 3, “Planning and preparation for the IBM Spectrum Virtualize for Public Cloud deployment” on page 31 are deployed.
- ▶ The required IBM Spectrum Virtualize back-end storage that is described in Chapter 3, “Planning and preparation for the IBM Spectrum Virtualize for Public Cloud deployment” on page 31 were made available in the same IBM Cloud Data Center where the bare-metal server were deployed.
- ▶ The required IBM Spectrum Virtualize Licenses was purchased and you can access IBM Passport Advantage.

When the bare-metal servers are ready, you can install IBM Spectrum Virtualize for Public Cloud by using the One-click installation methods that described in this section. *One-click cluster deployment* is a tool that helps the user to install IBM Spectrum Virtualize for Public Cloud automatically.

One-click cluster deployment has two modes:

- ▶ Fully Automated installation
- ▶ Semi Automated installation

Both modes install IBM Spectrum Virtualize for Public Cloud and create the cluster automatically.

The fully automated mode automatically determines the initial configuration parameters for configuration of IBM Spectrum Virtualize for Public Cloud, and automatically orders the required IP addresses from IBM Cloud. The semi-automated mode requires that the user orders the IP addresses (or uses an existing subnet) and generates and edits a configuration file to provide the installer script the needed parameters.

For each installation method, both GUI and command-line interface (CLI) are shown for comparison. Some common steps must be done, regardless of the installation method that is used. Figure 4-1 shows some common steps that must be completed.

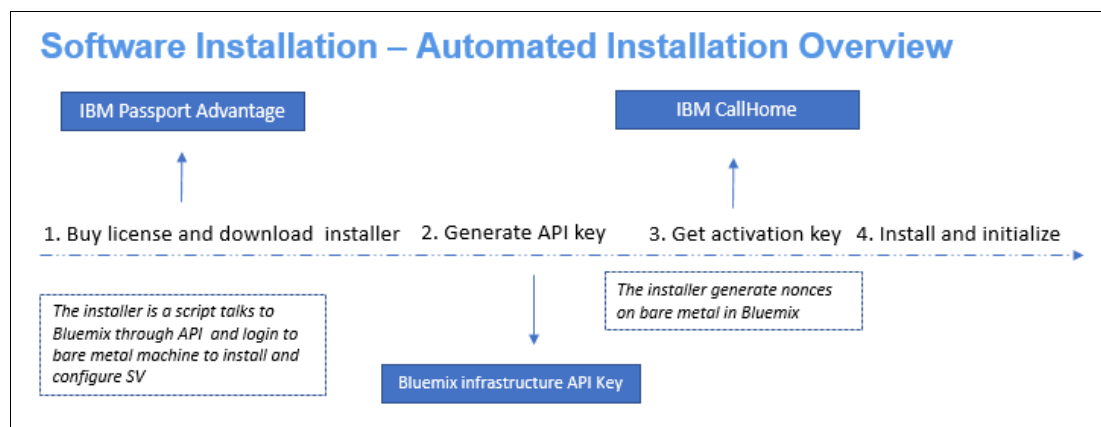


Figure 4-1 Software Installation overview

The first step for installing IBM Spectrum Virtualize for Public Cloud is to download the One-click installer, as described next.

4.1.1 Downloading the One-click installer

Before you can install the IBM Spectrum Virtualize for Public Cloud software, you must download the One-click installer for IBM Spectrum Virtualize for Public Cloud from IBM Marketplace and Passport Advantage.

In addition to the license, you must download the One-click installer, which is an application that runs on a local machine with internet access. This application facilitates the installation of the IBM Spectrum Virtualize for Public Cloud software on multiple bare-metal servers to create the system. You must download the One-click installer that is based on the operating system of the machine that you are using to run the installation. The One-click installer is available for Red Hat Enterprise Linux 7.x (RHEL 7.x), macOS, and Windows.

Downloading the installer for RHEL 7.x

To download the installer on a Linux system, complete the following steps:

1. Go to the IBM Passport Advantage [website](#) to obtain the One-click installer for IBM Spectrum Virtualize for Public Cloud.
2. Log in by using your Passport Advantage credentials.
3. Download the following installation package for RHEL 7.x to your Linux system:
SV_Cloud_Installer_RHEL_x.x.x.x.tar.gz (where x.x.x.x is the release identifier for the software).
4. On your Linux system, decompress the package by running the following command: `tar -zxvf SV_Cloud_Installer_RHEL_x.x.x.x.tar.gz` (where x.x.x.x is the release identifier for the software).

Downloading the installer for macOS

To download the installer for a MacOS system, complete the following steps:

1. Go to the IBM Passport Advantage [website](#) to obtain the One-click installer for IBM Spectrum Virtualize for Public Cloud.
2. Log in by using your Passport Advantage credentials.
3. Download the following installation package to your MacOS system:
SV_Cloud_Installer_MAC_x.x.x.x.tar.gz (where x.x.x.x is the release identifier for the software).
4. To decompress the package, enter the following command: `tar -zxvf SV_Cloud_Installer_MAC_x.x.x.x.tar.gz` (where x.x.x.x is the release identifier for the software).

Downloading the installer for Windows

To download the installer for a Windows system, complete the following steps:

1. Go to the IBM Passport Advantage [website](#) to obtain the one-click installer for IBM Spectrum Virtualize for Public Cloud.
2. Log in by using your Passport Advantage credentials.
3. Download the following installation package to your Windows system:
SV_Cloud_Installer_WIN_x.x.x.x.zip (where x.x.x.x is the release identifier for the software).

4. To decompress the package, use the Winzip application on your system.

Create a Classic Infrastructure API key

The automated installation script and the back-end storage configuration require an IBM Classic Infrastructure API key to use the IBM Cloud APIs to perform the automated software installation.

IBM Cloud includes different types of API keys. Here, we use an *infrastructure key*.

The API Key is used during the installation to access the API for the following purposes:

- ▶ Discover the passwords of the servers.
- ▶ Allocate a range of IP addresses for the cluster.
- ▶ Configure the storage at postinstallation.

Tip: It is best to generate the API key and perform the installation as a user without purchasing power to protect the client from the small risk of the installation script making purchases in error.

To create your API key, see this IBM Cloud [web page](#).

4.1.2 Fully Automated installation

Important: If in your environment you implemented a VPN firewall that is supplied by IBM Cloud IaaS (Vyatta) *do not use this procedure*. At the time of this writing, the Fully Automated procedure cannot interact with the Vyatta to make the subnet that was automatically allocated by the script by using API accessible. Instead, use the Semi Automated procedure that is described in 4.1.3, “Semi Automated installation” on page 78.

The Fully Automated installation requires the following steps to be run:

1. The user creates an infrastructure API key with minimum privilege, which allows the Install Script to query information of the nodes, as described in “Create a Classic Infrastructure API key” on page 74.
2. The user runs the installer on a notebook or server with an internet connection.
3. The script hides all the complexity of installation and first-time configuration by performing the following tasks:
 - a. Get the user name and password of the machine through the IBM Cloud API with the user-provided API key.
 - b. Allocate portable private IP addresses.
 - c. The script logs in to each node to generate a nonce and presents it to the customer to download an activation key.
 - d. The script logs in to each node and installs the IBM Spectrum Virtualize Cluster software with the activation keys that the client obtained from the IBM Call Home website.
 - e. The script creates the cluster and configures the iSCSI port IPs.
 - f. When finished, the script outputs a report about the installed cluster.

To run the Fully Automated installation, complete the following steps:

1. On your notebook or server:
 - a. For RHEL and MacOS hosts, change to the `one-click-install*` directory and run the command, as shown in Example 4-1.

Example 4-1 One-click-install example for RHEL and macOS

```
./SV_Cloud_Installer --user $username --apiKey $key --servers  
bm_server_name1 bm_server_name2 ...
```

- b. For Windows host, change to the `one-click-install*` directory, and run the command as shown in Example 4-2.

Example 4-2 One-click-install example for Windows

```
SV_Cloud_Installer.exe --user $username --apiKey $key --servers  
bm_server_name1 bm_server_name2 ...
```

Note: The `bm_server_name1/2` that is shown in Example 4-1 or Example 4-2 is the name you gave to your IBM Spectrum Virtualize Nodes when provided, as described in Chapter 3, “Planning and preparation for the IBM Spectrum Virtualize for Public Cloud deployment” on page 31. Use the host name only, not the Fully Qualified Domain Name (FQDN).

The parameter is case-sensitive, but is treated as *case-insensitive* unless ambiguity exists. For a two-node cluster, two bare-metal server names must be provided; four must be provided for a four node cluster.

Figure 4-2 shows an example of the `bm_server_name`.

Device Name	Device Type	Location	Public IP	Private IP	Start Date	Actions
fw1.agofelab.com	Gateway Member	Washington 4	169.47.129.107	10.183.3.16	2017-12-20	Actions
svcin1.agofelab.com	Bare Metal Serve	Washington 4	169.47.145.151	10.183.62.215	2017-12-20	Actions
svcin2.agofelab.com	Bare Metal Serve	Washington 4	169.47.145.154	10.183.62.194	2017-12-20	Actions

Figure 4-2 Hostname example

2. During the installation process, the output shows your nodes' *nonce* (number occurring once). Follow the steps that are described in the “ACTION REQUIRED” section to activate your nodes, as shown in Example 4-3.

Example 4-3 Fully Automated output example

```
C:\Users\IBM_ADMIN\Desktop\one-click-install-WIN(2)>SV_Cloud_Installer.exe  
--user 1205109_abernasconi@it.ibm.com --apiKey  
45e880b9ed885256aa3c9dc92c8d59eb4eef54010a75c6db434c65448e333249 --servers  
itso-dal10-sv-n1 itso-dal10-sv-n2 -f
```

Start deploying the SV_Cloud cluster, resource checking.

Allocating IP addresses for the SV_Cloud cluster.

Start deploying the SV_Cloud cluster, and the whole process will take about 20 minutes.

The SV_Cloud cluster will be deployed on: itso-dal10-sv-n1 itso-dal10-sv-n2
server name: nonce
itso-dal10-sv-n1: **D455F4**
itso-dal10-sv-n2: **D45D14**

ACTION REQUIRED

1. Please use server's nonce to get USVNID from <https://www.ibm.com/support/home/spectrum-virtualize>.
2. Put all USVNID files (such as D455F4.txt) into current working directory: C:\Users\IBM_ADMIN\Desktop\one-click-install-WIN(2)

Preparing for downloading.

```
.  
Downloading:  
Progress: |*****| 100% Complete  
Download completed, will start installing soon.  
Installing:  
Progress: |*****| 100% Complete
```

SV_Cloud nodes are installed.
Now starting nodes, should take a few minutes.
.
Activate nodes: ['itso-dal10-sv-n1', 'itso-dal10-sv-n2'].

SV_Cloud nodes started, will start making cluster soon.

Making cluster on itso-dal10-sv-n1, will take a few minutes.
.
Adding nodes to the cluster.
Adding node B01ELJ3 to IO_group 0.
.
The SV_Cloud cluster is ready, and the cluster IP is **10.93.135.196**.

3. To complete the required action that is described in Figure 4-1 on page 72, see [this page](#) of the IBM Support website (log in required). Complete the following steps:
 - a. Download the activation key for each node in the IBM Spectrum Virtualize cluster, as shown in Figure 4-3 on page 77.

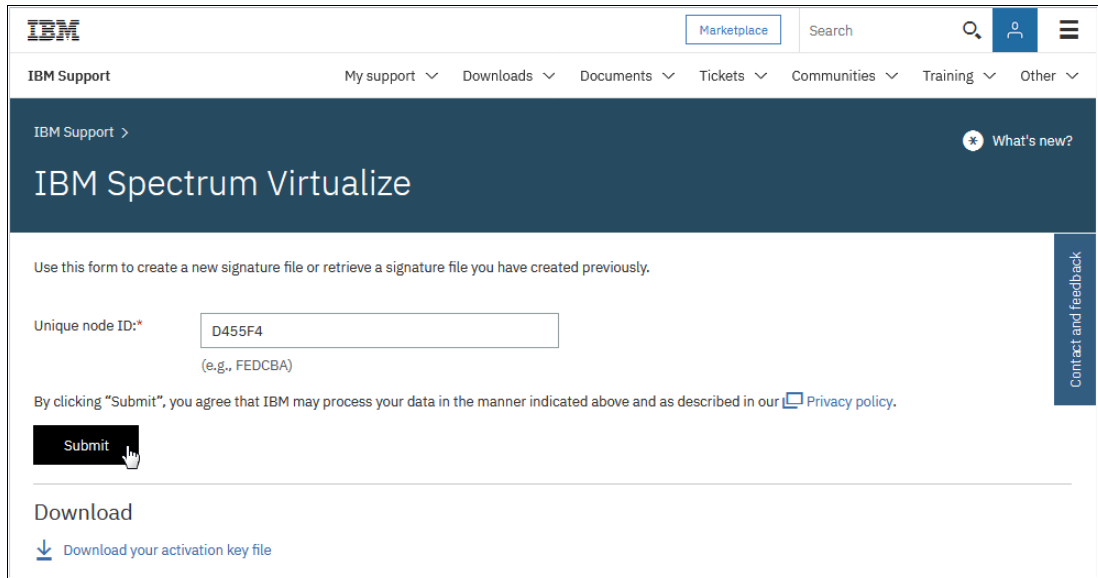


Figure 4-3 Downloading activation key example

- b. Save the downloaded keys in your one-click-install directory, as shown in Figure 4-4.

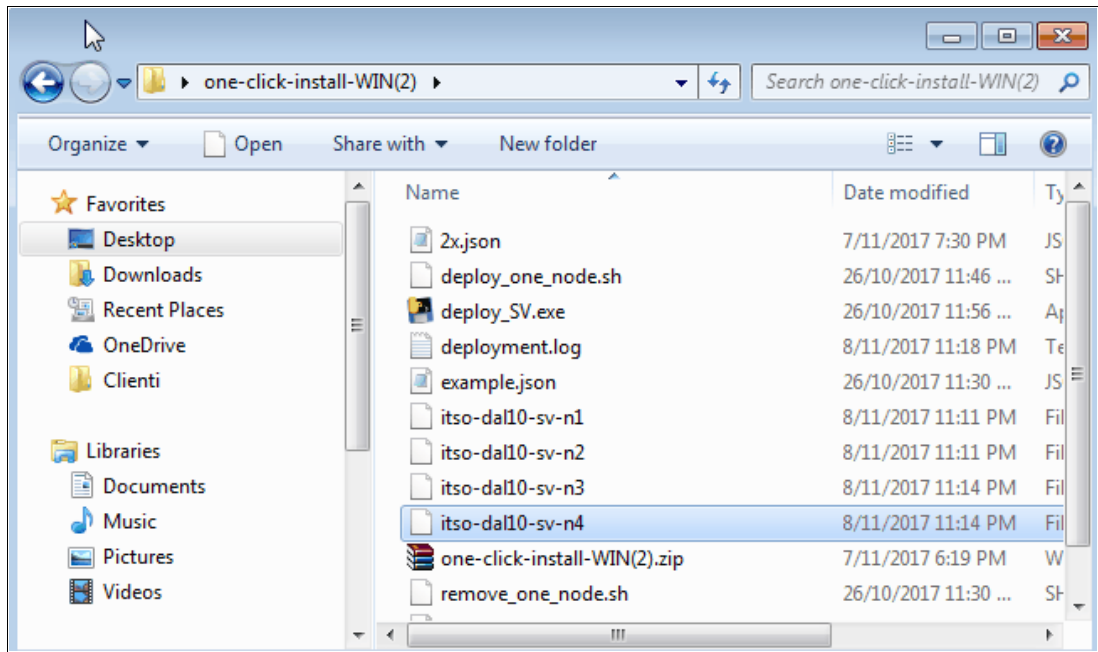


Figure 4-4 IBM Spectrum Virtualize activation key saving example

Note: If you do not save the activation keys in this manner, the Fully Automated script prompts you to confirm and does not continue until the keys are saved in the expected directory.

4. After approximately 20 minutes, an IBM Spectrum Virtualize for Public Cloud cluster is ready. The deployment script saves the cluster configuration report in the file `report.json` in the One-click installation working directory.

5. You can skip ahead to the sections **Log in to Cluster** to log in to the IBM Spectrum Virtualize for Public Cloud cluster and proceed with Configure Backend Storage.

4.1.3 Semi Automated installation

Semi Automated mode needs the user to provide more configuration information to install IBM Spectrum Virtualize for Public Cloud. A template is available in the One-click installer package: `example.yaml`. Carefully follow the next sequence to edit this file. Example 4-4 shows how to create an example configuration file to be edited for the installation.

Example 4-4 Generating a sample yaml configuration file

```
jfincher$ ./SV_Cloud_Installer -g sample.yaml
Generating config file: sample.yaml
```

Complete the following steps:

1. For each bare-metal server, collect the following information:
 - ID (assigned by yourself according to your naming convention, such as 1).
 - Server name (for example, `itso-dal10-sv-n1`).
 - `publicIpAddress`.
 - `privateIpAddress`.
 - Operating system user (for example, `root`).
 - Operating system password (you see `*****`). Double-click to make it visible.
 - Serial (for example, `SL01ELJ5`).

You can gather this information from the IBM Cloud portal by authenticating with your IBM Cloud user ID and password and selecting your bare-metal server, as shown in Figure 4-5.

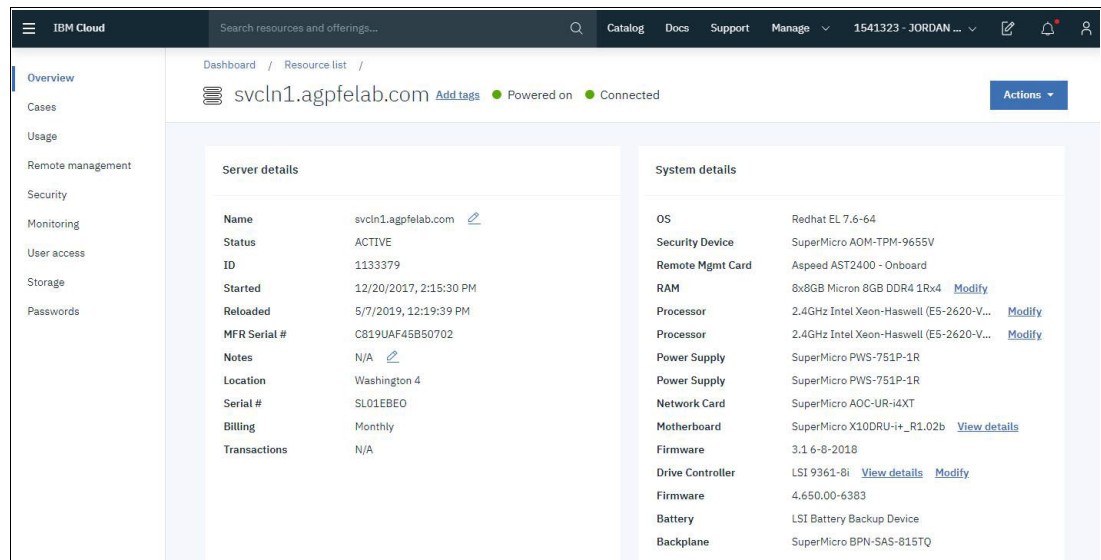


Figure 4-5 Device detail example

2. User and password can be collected from the window by expanding the device in the device list, as shown in Figure 4-5. You see the window that is shown in Figure 4-6 on page 79. Click show password to make it visible.

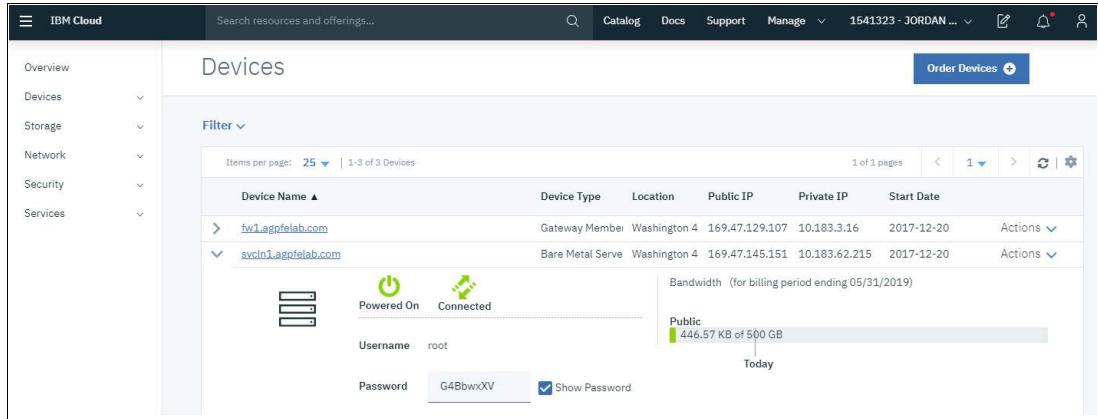


Figure 4-6 User and password example

- Allocate IP addresses for use by Spectrum Virtualize from the portable private subnet that is in the same VLAN the private network for the bare-metal servers. The required number of IP addresses for the installation is five IPs per node, plus one cluster IP address. For more information about ordering and provisioning a portable private subnet, see 3.3, “IBM Spectrum Virtualize networking considerations” on page 57.

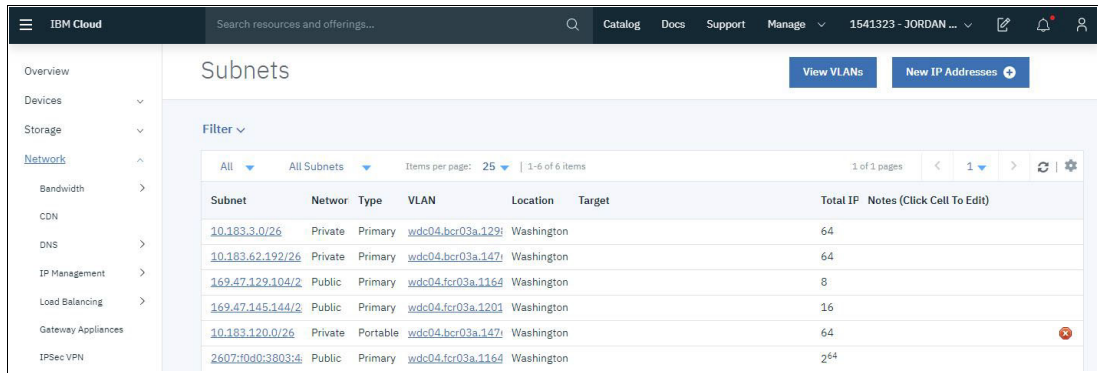


Figure 4-7 Subnet menu example

- Complete example.yaml by using your IP address configurations.

Keeping a spreadsheet (see Figure 4-8) in advance to assign your specific IP address to each specific role can be useful.

	Priv. IP	Pub. IP	Service IP	N2N IP I*	N2N IP II*	Port IP I*	Port IP II*	Cluster IP
svcln1	10.183.62.215	169.47.145.151	10.183.120.11	10.183.120.21	10.183.120.22	10.183.120.15	10.183.120.16	
svcln2	10.183.62.194	169.47.145.154	10.183.120.12	10.183.120.23	10.183.120.24	10.183.120.17	10.183.120.18	
Cluster								10.183.120.10
	Net	GW	Mask					
Public Network	169.47.145.144	169.47.145.145	225.255.255.192					
Private Network	10.183.62.192	10.183.62.193	225.255.255.192					
Portable Network	10.183.120.0	10.183.120.1	225.255.255.192					

Figure 4-8 IP address table document

Example 4-5 on page 80 shows a new yaml file that was completed with your IP address. This example is for an IBM Spectrum Virtualize cluster with only two nodes.

Note: The password and private key parameters are mutually exclusive.

Example 4-5 yml file example

```
# version=8.1.3.0-180512_1402
cluster:
  ipAddress: 10.183.120.10 # cluster ip
  gateway: 10.183.120.1
  netmask: 255.255.255.192
  site1:
    BareMetalServers:
      - servername: svcln1 # the name showed in cloud web portal
        publicIpAddress: 169.47.145.151
        privateIpAddress: 10.183.62.215
        user: root # username with root privilege
        password: G4BbwXV # login password for user.
        # privateKey: C:\Users\ADMIN\.ssh\bm01_private_key
        serial: SL01EBEO # Bare Metal server serial number
        id: 1 # select the SpecV node id for this node
        serviceIp:
          netmask: 255.255.255.192
          ipAddress: 10.183.120.11
          gateway: 10.183.120.1
        portIp:
          - netmask: 255.255.255.192
            ipAddress: 10.183.120.20
            gateway: 10.183.120.1
          - netmask: 255.255.255.192
            ipAddress: 10.183.120.21
            gateway: 10.183.120.1
        nodeIp:
          - netmask: 255.255.255.192
            ipAddress: 10.183.120.15
            gateway: 10.183.120.1
          - netmask: 255.255.255.192
            ipAddress: 10.183.120.16
            gateway: 10.183.120.1
      - servername: svcln2
        publicIpAddress: 169.47.145.154
        privateIpAddress: 10.183.62.194
        user: root # username with root privilege
        password: U8qp9t5w # login password for user
        # privateKey: C:\Users\ADMIN\.ssh\bm02_private_key
        serial: SL019TYC # Bare Metal server serial number
        id: 2
        serviceIp:
          netmask: 255.255.255.192
          ipAddress: 10.183.120.12
          gateway: 10.183.120.1
        portIp:
          - netmask: 255.255.255.192
            ipAddress: 10.183.120.22
            gateway: 10.183.120.1
          - netmask: 255.255.255.192
            ipAddress: 10.183.120.23
            gateway: 10.183.120.1
        nodeIp:
          - netmask: 255.255.255.192
```

```
    ipAddress: 10.183.120.17
    gateway: 10.183.120.1
  - netmask: 255.255.255.192
    ipAddress: 10.183.120.18
    gateway: 10.183.120.1
```

Note: If your configuration is with four Nodes, edit your yaml file and add two sections for the nodes.

5. Save your yaml file by using a useful name; for example, `sample.yaml`.
6. Run the configuration file validation command from the directory where you saved your installation files, as shown in Example 4-6.

Example 4-6 Deploy Spectrum Virtualize semi-automated example

```
jfincher$ ./SV_Cloud_Installer -c sample.yaml
Verifying config file: sample.yaml
Config file sample.yaml looks good, your are ready for installation.
```

7. After the validation process is complete, run the command to start the installation process that is shown in Example 4-7.

Note: As part of the installation procedure, you are presented with activation codes for each node. In this example, these codes are D2F9D8 and 3A0EC0. It is required as part of the installation process to download the activation keys per the instructions that are presented in the command output and store them in the same directory as the installation script, as explained in 4.1.2, “Fully Automated installation” on page 74.

At the time of this writing, the web page to download the key files works in Internet Explorer only. In Firefox and Chrome, pasting the link location in the address bar (and removing the `unsafe:` prefix from the URL) allows you to view the text file in the browser. Pasting it into a text file that is named `NONCE.txt` (replace `NONCE` with the string that is provided for the node) is sufficient to proceed.

Example 4-7 Semi automated node activation example

```
jfincher$ ./SV_Cloud_Installer -i -f --configFile sample.yaml
```

```
Start deploying the SV_Cloud cluster, resource checking.
Start deploying the SV_Cloud cluster, and the whole process will take about 20
minutes.
The SV_Cloud cluster will be deployed on: svc1n1 svc1n2
# svc1n1: D2F9D8
# svc1n2: 3A0EC0
```

ACTION REQUIRED

1. Please use server's nonce to get USVNID from <https://www.ibm.com/support/home/spectrum-virtualize>.
2. Put all USVNID files (such as `D2F9D8.txt`) into current working directory: `/Users/jfincher/Downloads/SV_Cloud_Installer`.

Preparing for downloading.

. . .

Downloading:

Progress: |*****| 100% Complete

```

Download completed, will start installing soon.
Installing:
Progress: |*****| 100% Complete

SV_Cloud nodes are installed.
Now starting nodes, should take a few minutes.
. . . . .
Please put USVNID file ['D2F9D8.txt', '3A0EC0.txt'] into
/Users/jfincher/Downloads/SV_Cloud_Installer.
Press anykey to continue or N to exit:

Activate nodes: ['svcln1', 'svcln2'].

SV_Cloud nodes started, will start making cluster soon.

Making cluster on svcln1, will take a few minutes.
. . . . .
Adding nodes to the cluster.
Adding node B019TYC to IO_group 0.
. . . . .
Setup portip for the cluster.
Setup DNS server for the cluster using IP 10.0.80.11.

The SV_Cloud cluster is ready, and the cluster IP is 10.183.120.10.

```

8. After the cluster is ready, see 4.2, “Configuring Spectrum Virtualize”.

4.2 Configuring Spectrum Virtualize

When the installation is complete, we can log in to IBM Spectrum Virtualize for Public Cloud for further configuration. The steps are described next.

4.2.1 Log in to cluster and complete the installation

Complete the following steps to log in to the cluster and complete the installation:

1. Log in to the cluster by using the GUI from your browser, as shown in Figure 4-9 on page 83.

With the GUI, you are guided to complete your cluster installation.

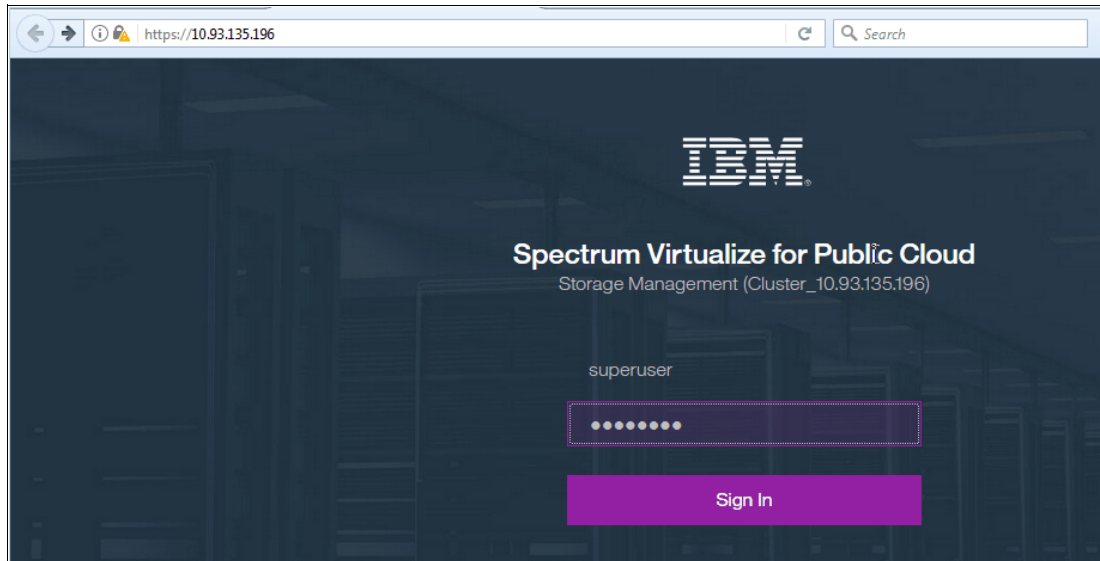


Figure 4-9 Logging in by using GUI

2. You are redirected to the Welcome window. Click **Next**, as shown in Figure 4-10.

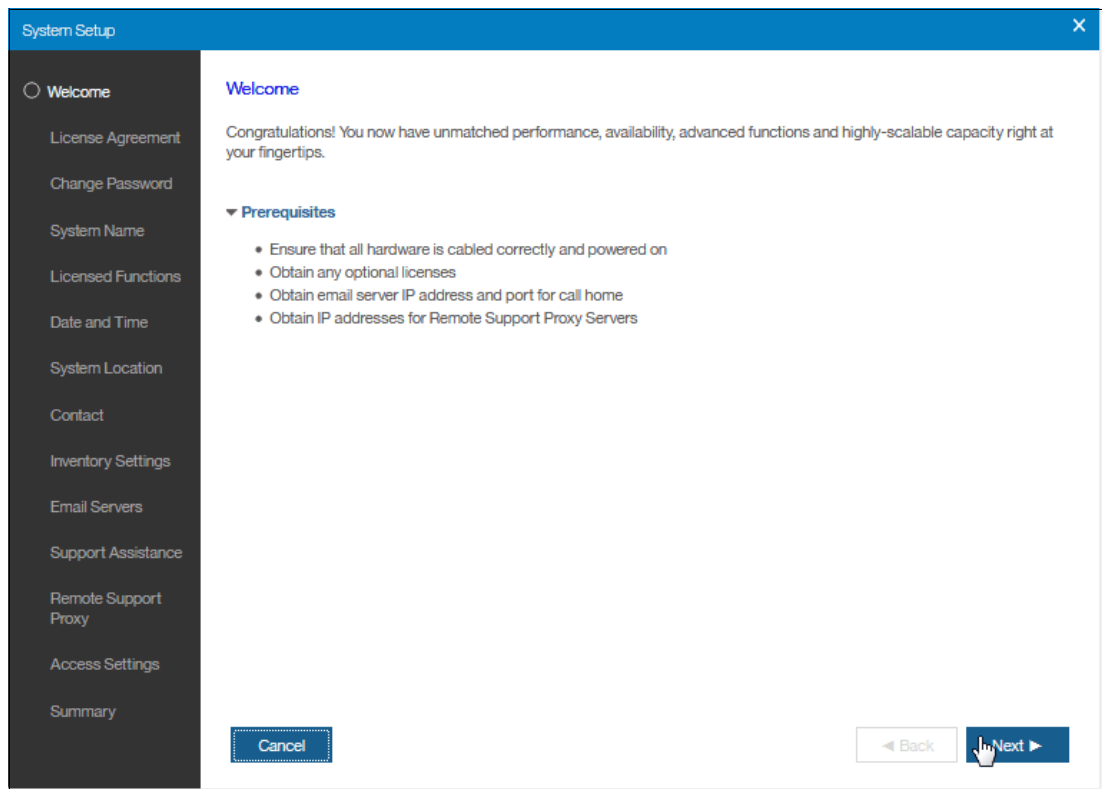


Figure 4-10 Welcome window

After the License Agreement window, you are redirected to the change password window, as shown in Figure 4-11.

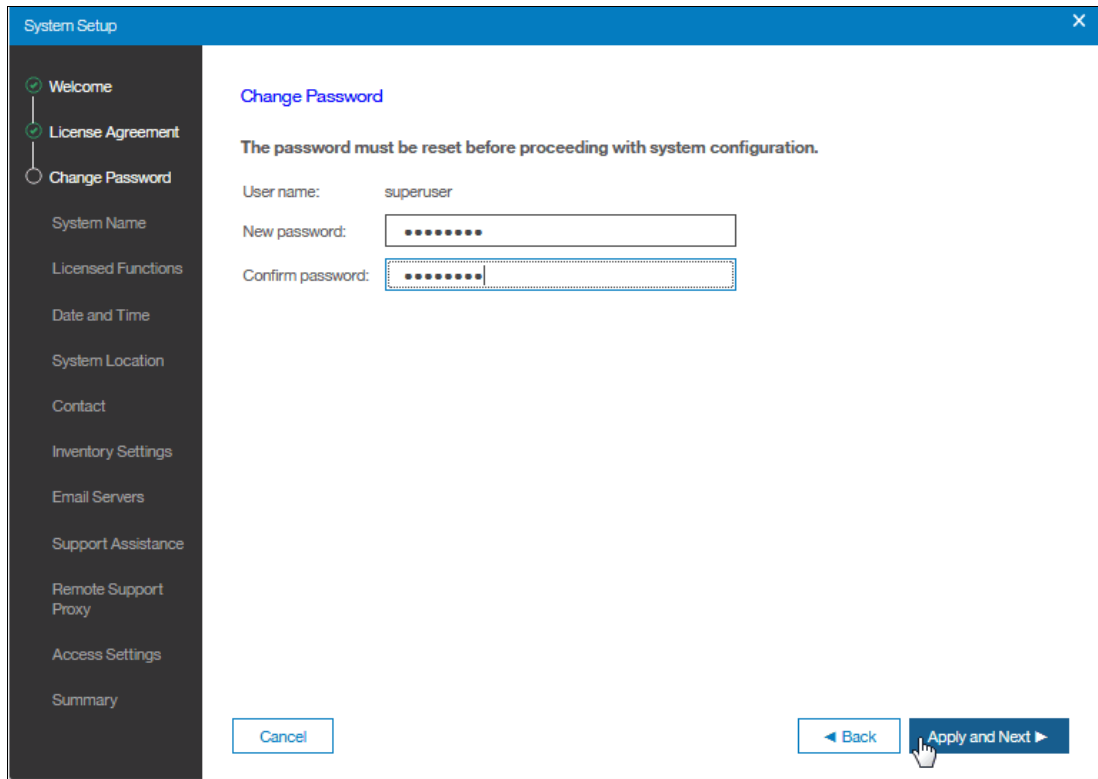


Figure 4-11 Change password

3. Change your password and then, click **Apply and Next**.

4. You can change your cluster default name, as shown in Figure 4-12. Then, click **Apply and Next**.

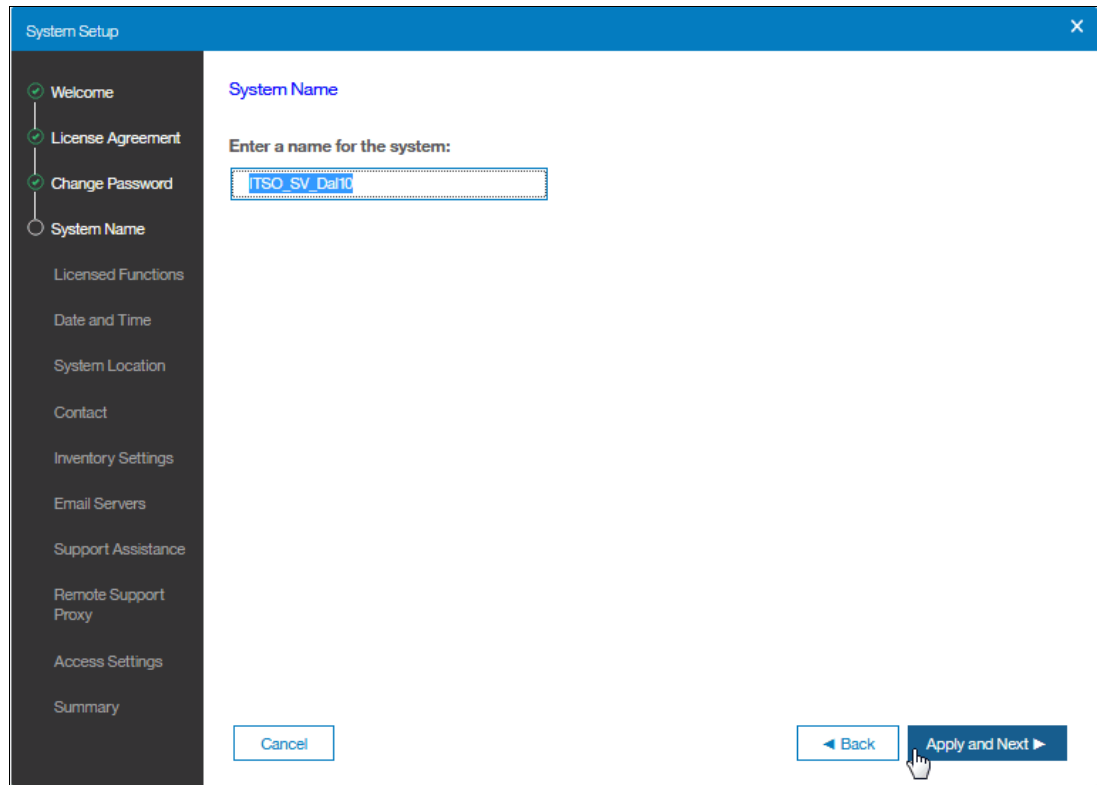


Figure 4-12 Cluster name change

5. Enter your capacity license in accordance with your IBM agreement, as shown in Figure 4-13. Then, click **Apply and Next**.

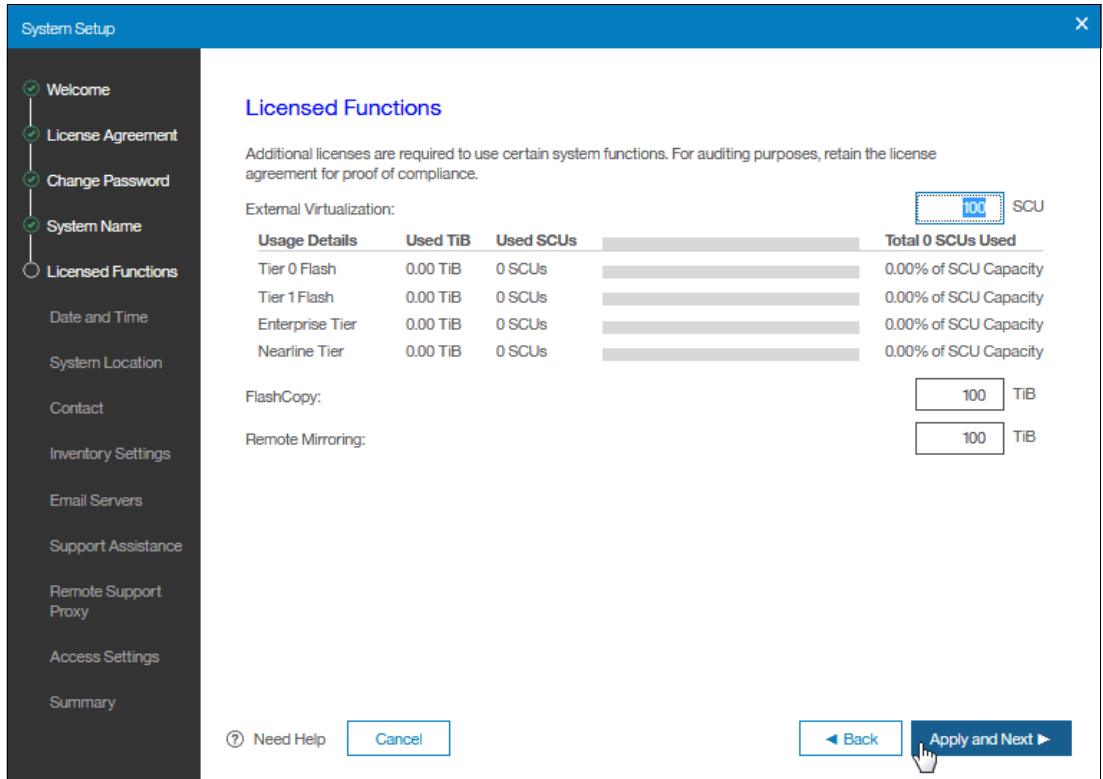


Figure 4-13 Capacity license window

6. Set your Date and Time in accordance to your specific policy. In our example, we configured Date and Time manually by using our environment Time Zone, as shown in Figure 4-14. Then, click **Apply and Next**.

We suggest the use of Network Time Protocol (NTP) and configuring it as shown in Figure 4-14.

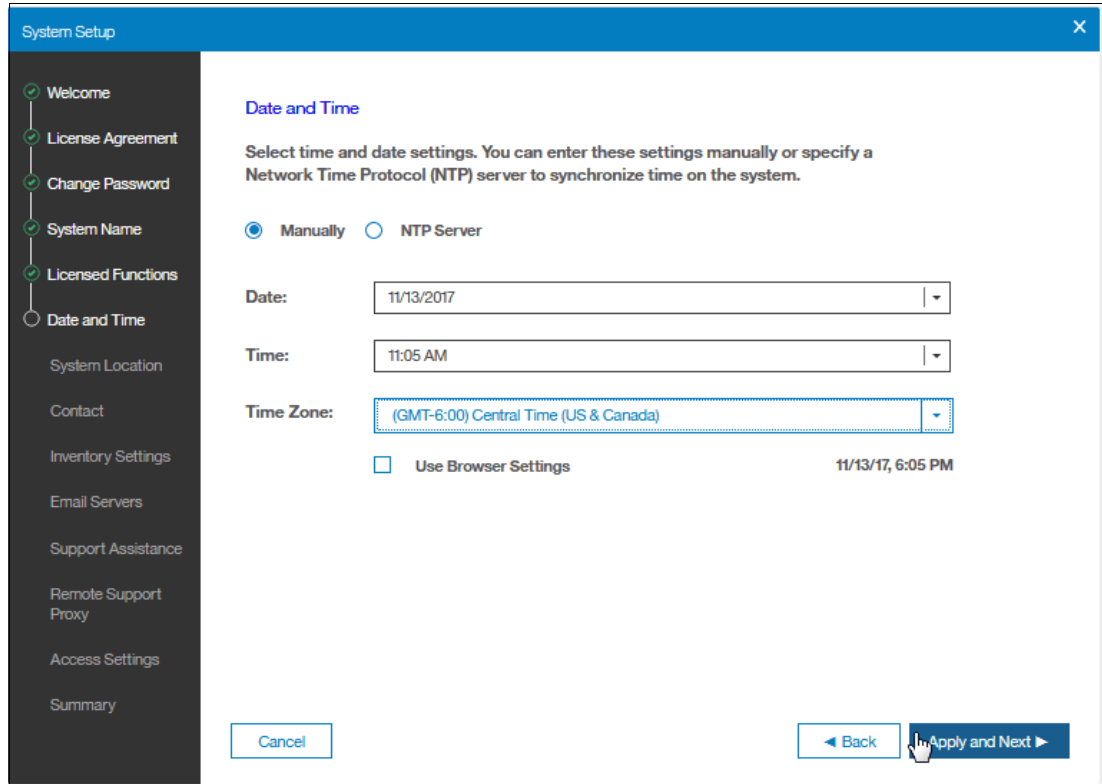


Figure 4-14 Date and Time window

- In the next windows, you are prompted to enter location information about your IBM Spectrum Virtualize cluster and some contact information, as shown in Figure 4-15 and Figure 4-16.

The screenshot shows the 'System Setup' window with the 'System Location' page selected in the left-hand navigation menu. The page title is 'System Location' and it includes a note: 'Service parts should be shipped to the same physical location as the system.' The form contains the following fields:

- Company name: IBM_ITO
- System address: Dal10 Data Center
- City: Dallas
- State or province: TX
- Postal code: 19283
- Country or region: United States (dropdown menu)
- Machine location: Dal10 DC roomX floorY BuildingW

At the bottom of the form, there are three buttons: 'Cancel', 'Back', and 'Next'. A mouse cursor is pointing at the 'Next' button.

Figure 4-15 System Location information

The screenshot shows the 'System Setup' window with the 'Contact' page selected in the left-hand navigation menu. The page title is 'Contact' and it includes a note: 'The support center contacts this person to resolve issues on the system.' The form contains the following fields:

- Name: MyName
- Email: MyNme@MyName.com
- Phone (primary): +10234987
- Phone (alternate): +19823546

At the bottom of the form, there are three buttons: 'Cancel', 'Back', and 'Apply and Next'. A mouse cursor is pointing at the 'Apply and Next' button.

Figure 4-16 Contact information

8. You are prompted to configure the Inventory Setting, as shown in Figure 4-17. We suggest setting Inventory Reporting and Configuration Reporting **ON** to help IBM Support Center better support you if potential issues and debugging occur.

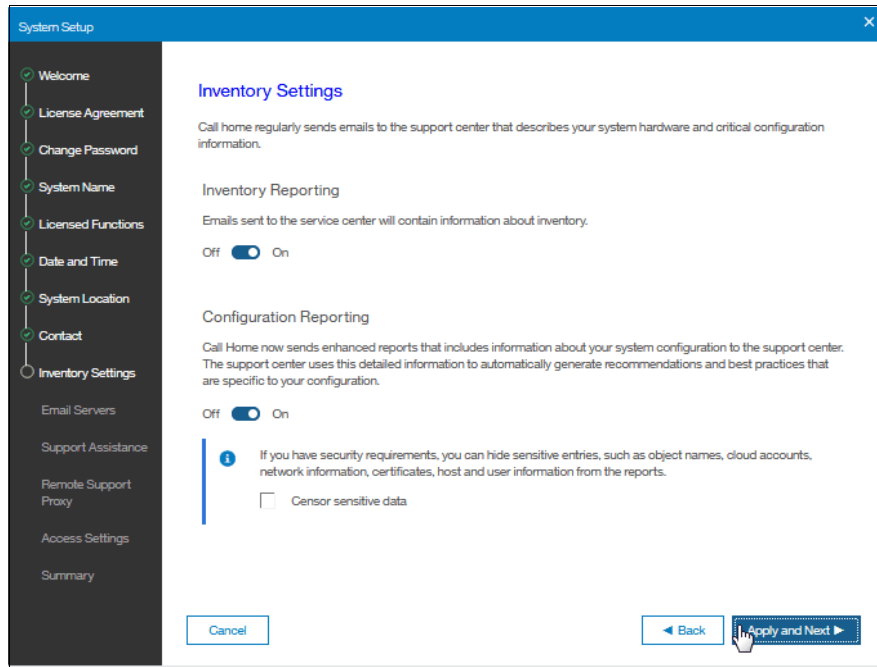


Figure 4-17 inventory setting example

9. You are prompted to configure Home Call (SMTP Server) alert and Support Assistance, as shown in Figure 4-18 and Figure 4-19 on page 90.

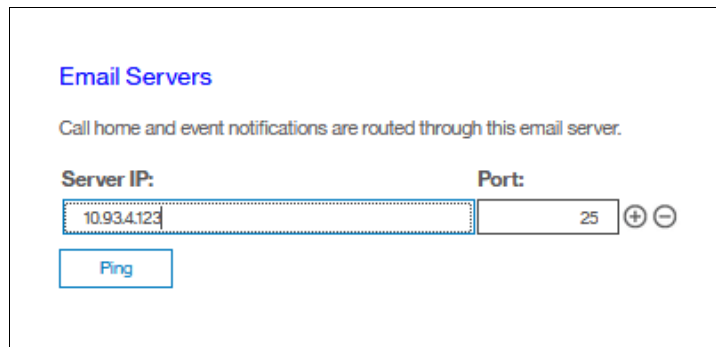


Figure 4-18 SMTP server window

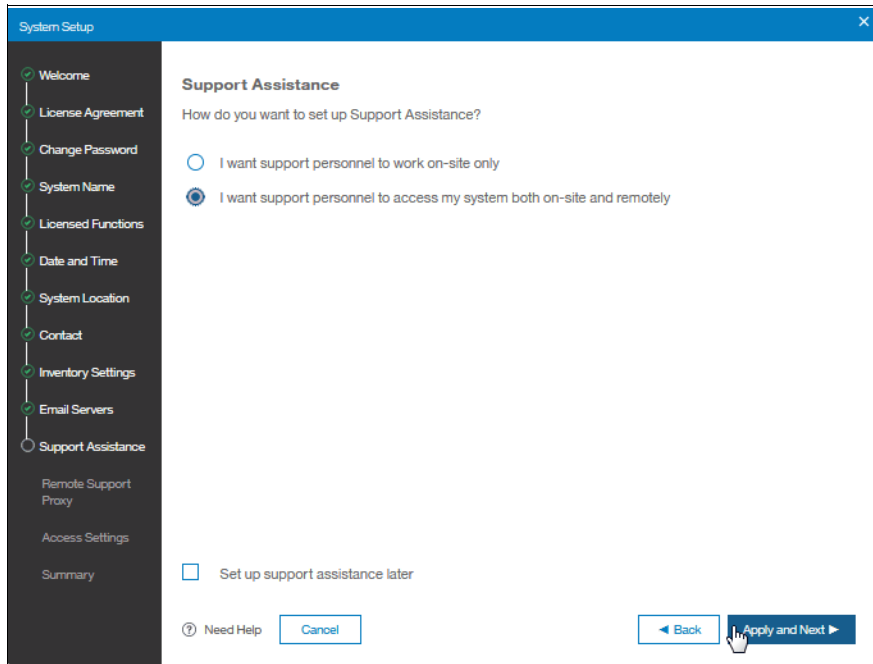


Figure 4-19 Support Assistance window

10. Configure your Remote Support Proxy, as shown in Figure 4-20.

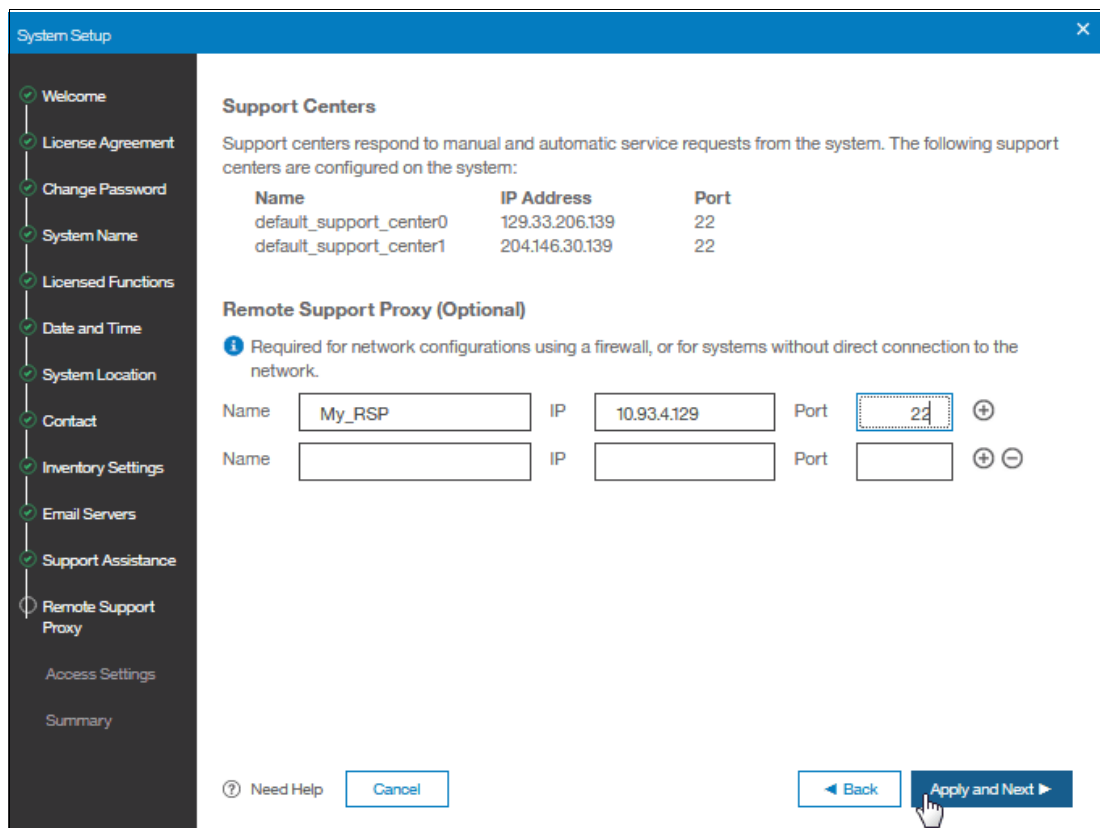


Figure 4-20 Remote Support Proxy window

Note: For more information about how to configure a Remote Support Proxy, see 4.4, “Configuring Remote Support Proxy” on page 115.

A summary of your configuration is shown. You cluster setup completes and you are redirected to your IBM Spectrum Virtualize GUI Dashboard, as shown in Figure 4-21, Figure 4-22, and Figure 4-23 on page 92.

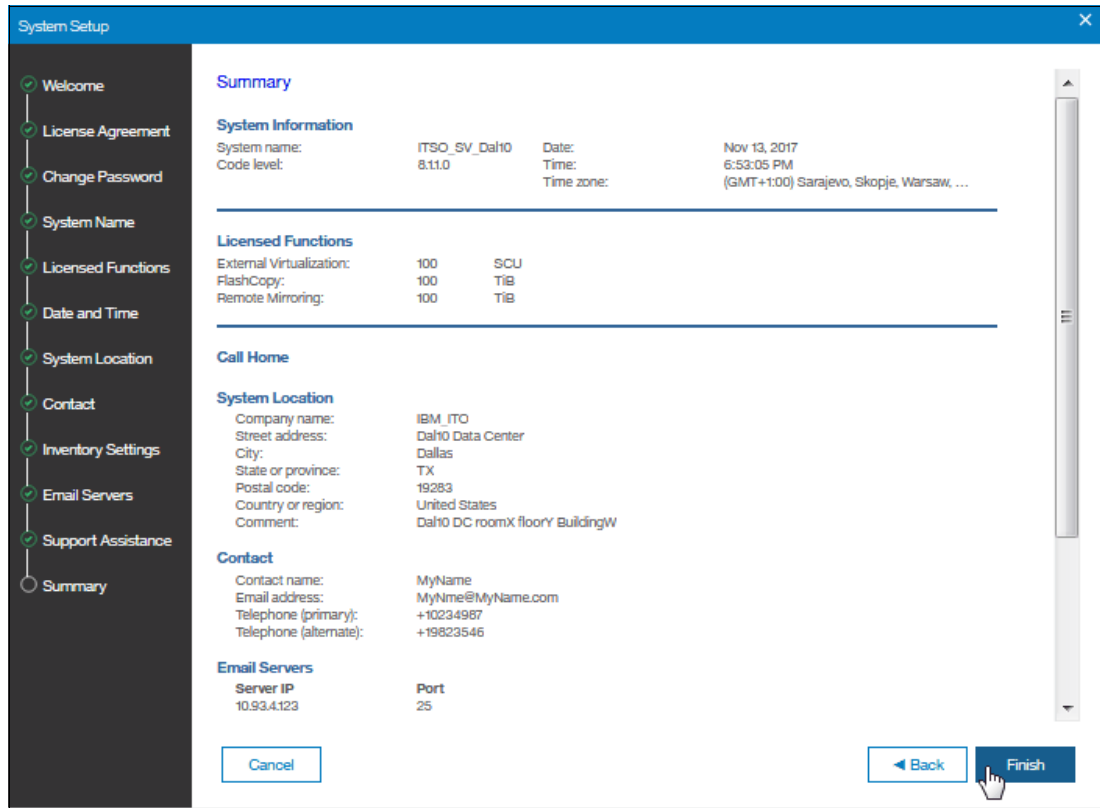


Figure 4-21 Summary example

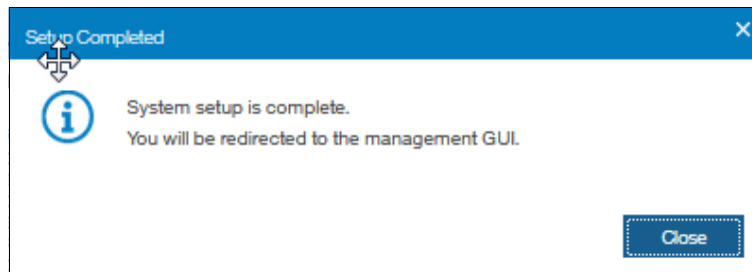


Figure 4-22 System setup complete example

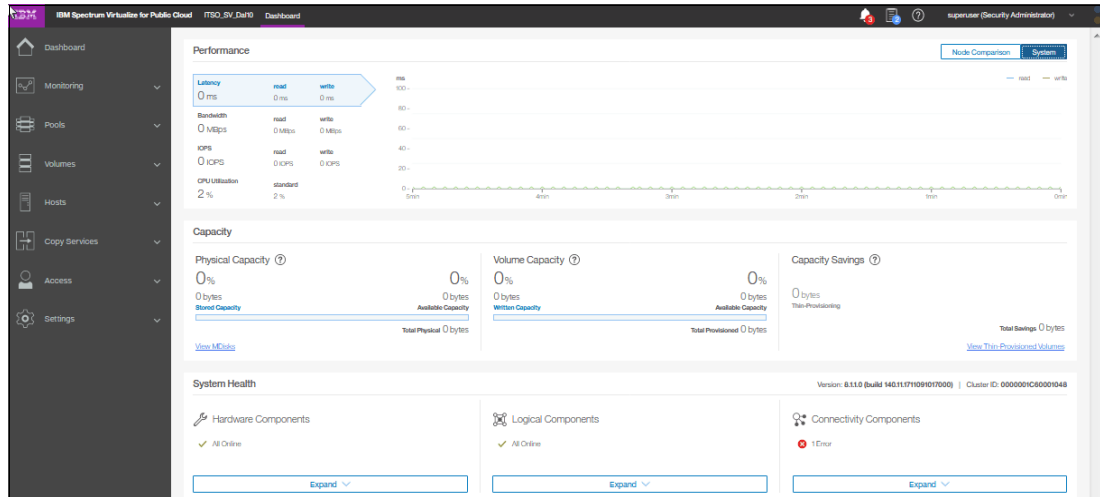


Figure 4-23 Spectrum Virtualize Dashboard example

Your IBM Spectrum Virtualize cluster now is complete. Configuring the Cloud quorum is described next.

4.2.2 Configure Cloud quorum

IP quorum applications are used in Ethernet networks to resolve failure scenarios in which half of the nodes on the system become unavailable. These applications determine which nodes can continue processing host operations and avoids split brain scenarios in which both halves attempt to independently service I/O, which causes corruption issues. IBM Spectrum Virtualize for Public Cloud requires at least one IP quorum application on a bare-metal or virtual server in IBM Cloud.

The IP quorum application is required for two- and four-node system in IBM Spectrum Virtualize for Public Cloud configurations. In two-node systems, the IP quorum application maintains availability after a node failure. In systems with four nodes, an IP quorum application is necessary handle with other failure scenarios. The IP quorum application is a Java application that runs on a separate bare-metal or virtual server in IBM Cloud.

There are strict requirements on the IP network for the use of IP quorum applications. All IP quorum applications must be reconfigured and redeployed to hosts when certain aspects of the system configuration change. These aspects include adding or removing a node from the system, when node service IP addresses are changed, changing the system certificate, or when an Ethernet connectivity issue occurs.

An Ethernet connectivity issue prevents an IP quorum application from accessing a node that is still online.

If an IP application is offline, it must be reconfigured because the system configuration changed.

To view the state of an IP quorum application in the management GUI, select **Settings** → **System** → **IP Quorum**, as shown in Figure 4-24 on page 93.

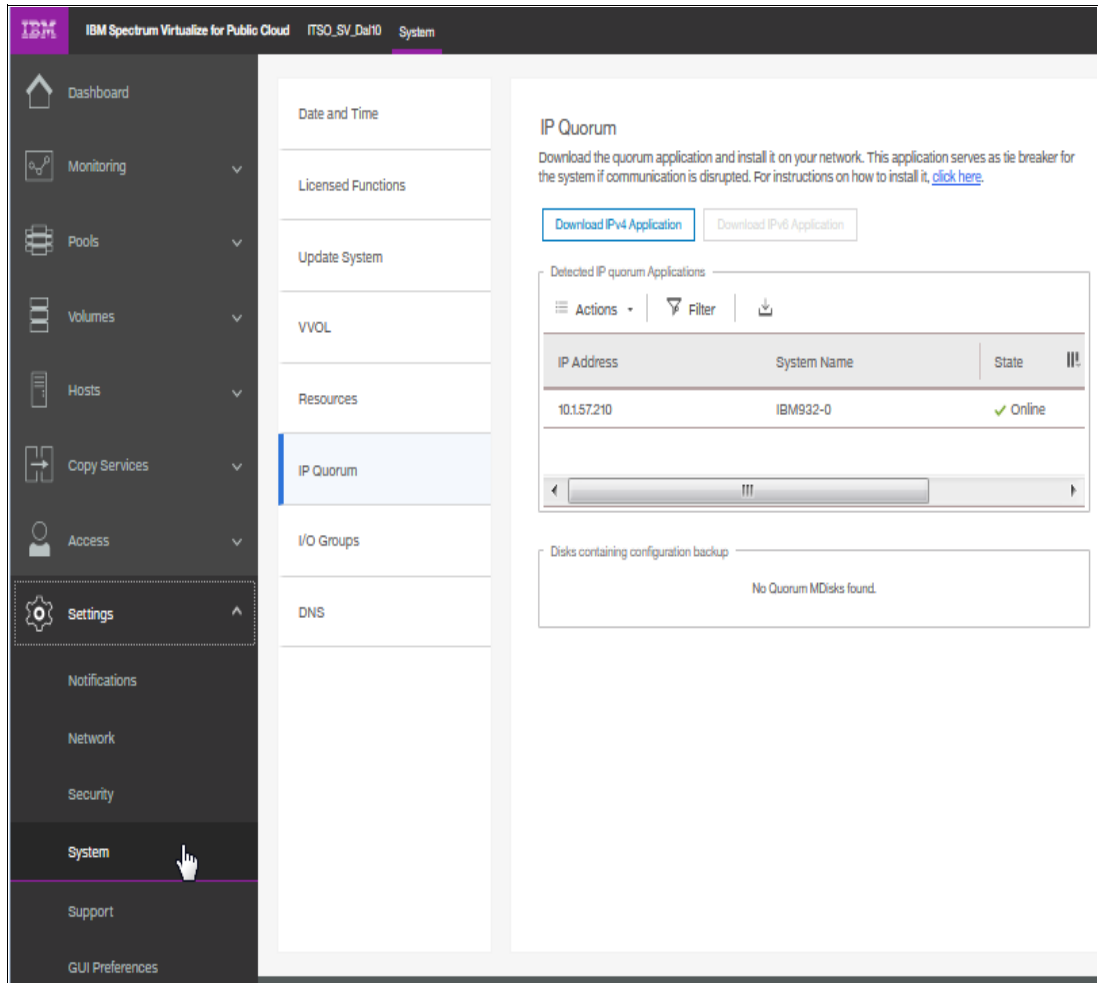


Figure 4-24 IP Quorum example from the GUI

Even with IP quorum applications on a bare-metal server, quorum disks are required on each node in the system. In a Cloud environment where IBM Spectrum Virtualize connectivity with its back-end storage is iSCSI, the Quorum disks cannot be on external storage or internal disk as in IBM SAN Volume Container or IBM Storwize. Therefore, they are automatically allocated on the bare-metal server internal disks.

The use of the IBM Spectrum Virtualize command `lsquorum` shows only the IP Quorum.

A maximum of five IP quorum applications can be deployed. Applications can be deployed on multiple hosts to provide redundancy.

For stable quorum resolutions, an IP network must meet the following requirements:

- ▶ Provide connectivity from the servers that are running an IP quorum application to the service IP addresses of all nodes.
- ▶ The network must also deal with possible security implications of exposing the service IP addresses because this connectivity can also be used to access the service assistant interface if the IP Network security is configured incorrectly.
- ▶ Port 1260 is used by IP quorum applications to communicate from the hosts to all nodes.

- ▶ The maximum round-trip delay must not exceed 80 milliseconds (ms), which means 40 ms each direction.
- ▶ A minimum bandwidth of 2 MB per second is guaranteed for node-to-quorum traffic.

For more information about IP quorum configuration, see [IBM Knowledge Center](#).

4.2.3 Installing the IP quorum application

If you are installing a new IBM Spectrum Virtualize system or changing the configuration by adding a node, changing a service IP address, or changing SSL certificates, you must download and install the IP quorum application again. To download and install the IP quorum application, complete the following steps:

1. Click **Download IPv4 Application** or **Download IPv6 Application** to create the IP quorum Java application. The application is stored in the dumps directory of the system with the file name `ip_quorum.jar`.
2. Transfer the IP quorum application from the system to a directory on the bare-metal server that hosts the IP quorum application.
3. Use the **ping** command on the host server to verify that it can establish a connection with the service IP address of each node in the system.
4. On the host, use the `java -jar ip_quorum.jar &` command to initialize the IP quorum application.
5. To verify that the IP quorum application is installed and active, select **Settings** → **System** → **IP Quorum**. The new IP quorum application is displayed in the table of detected applications.
6. To verify that the IP quorum application is installed and active by using IBM Spectrum Virtualize CLI, use `lsqlquorum` command.

The process to configure back-end storage is described next.

4.2.4 Configure the back-end storage

IBM Spectrum Virtualize for Public Cloud uses the back-end storage that is provided by IBM Cloud as external MDisk.

We assume that you ordered the back-end storage as described in Chapter 3, “Planning and preparation for the IBM Spectrum Virtualize for Public Cloud deployment” on page 31.

You can obtain at [this web page](#) the target IP address of the storage you just purchased, as shown in Figure 4-25.

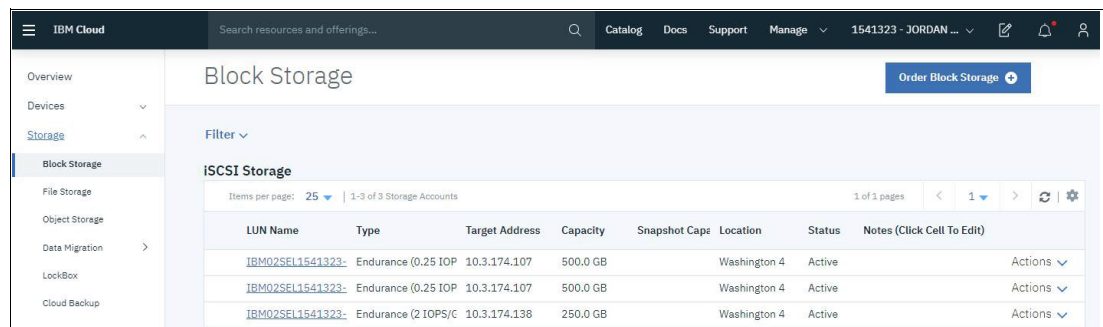


Figure 4-25 Block Storage menu

To configure your IBM Spectrum Virtualize back-end storage with GUI, complete the following steps:

1. In the management GUI, navigate to the **Pool** → **External Storage** menu in the GUI, as shown Figure 4-26.

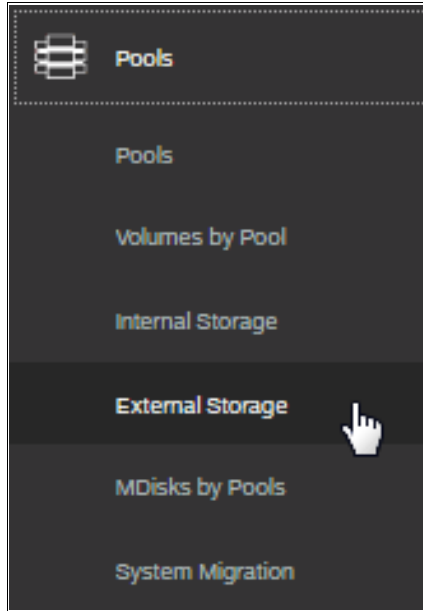


Figure 4-26 Pool menu example

2. Add external storage, as shown in Figure 4-27.

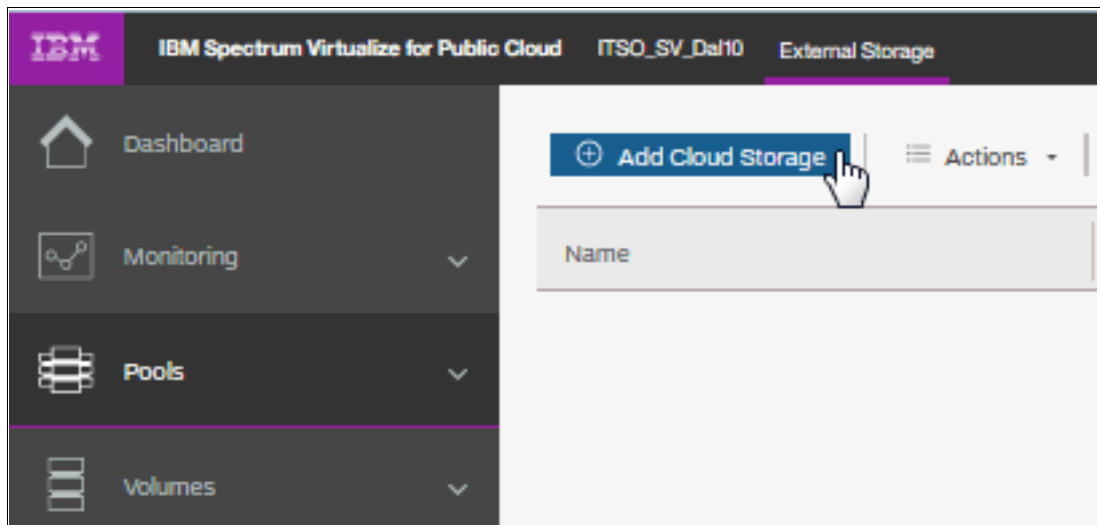


Figure 4-27 Add external storage example

3. Complete the following fields, as shown in Figure 4-28:
 - Bluemix API Username: Your IBM Cloud VPN user name.
 - Bluemix API key: Your IBM Cloud key that was created, as shown in Figure 4-28 on page 96.

- Node port ID: Your IBM Spectrum Virtualize node port ID. Ports ID 1 and ID2 are available. It is recommended that both be used in a round-robin fashion to get better workload balance and redundancy of each LUN (MDisk).

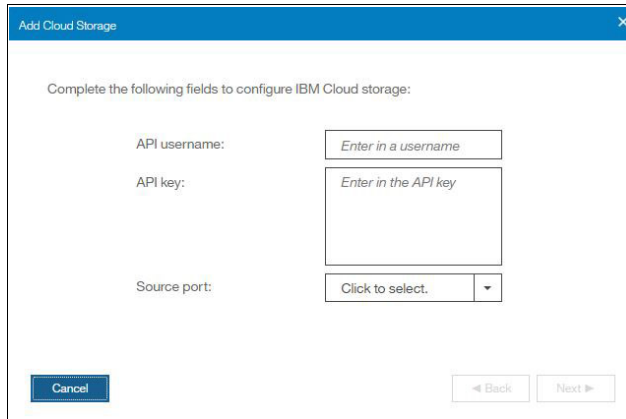


Figure 4-28 Storage configuration example

4. Select the storage you want to configure by right-clicking it and selecting **Include**, as shown in Figure 4-29.

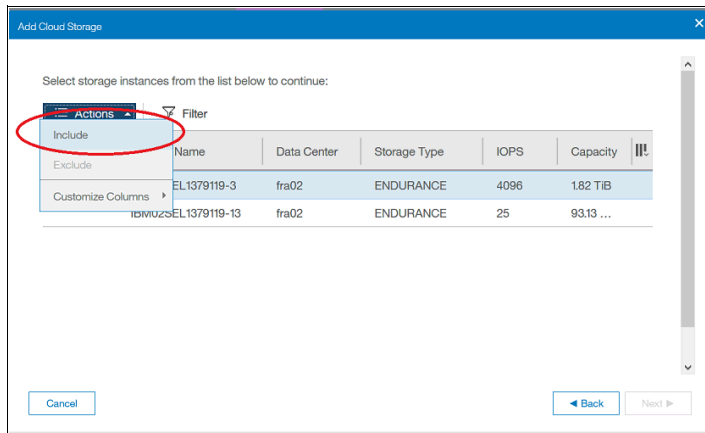


Figure 4-29 Selecting storage to configure

5. Click **Next** and complete the wizard to add the storage.

Now your back-end storage configuration is completed and you can create pools, volumes, and hosts as you do with any Spectrum Virtualize installation. For more information, see [IBM Knowledge Center](#).

4.2.5 Configuring Call Home with CLI

At the time of this writing, the Call Home configuration must to be run using CLI. In future releases, this is planned to be automated by using the one-click process. Currently, it requires a manual invocation of the command that is shown in Example 4-8.

Example 4-8 Call home config command example

```
cfgcloudcallhome -username <apiUser> -key <apiKey> -ip <ipOfIPQuorumServer>
-ibmcustomer <ibmCustomerNumber> -ibmcountry <ibmCompanyId>
```

The following elements are shown in Example 4-8:

- ▶ **username**: The name that you specified in the MDisk configuration (IBM Cloud user name)
- ▶ **key**: The same as you specified in MDisk configuration (IBM Cloud API Key)
- ▶ **IP**: IP address of IP quorum server
- ▶ **ibmcustomer**: Specifies the customer number that is assigned when a software license is automatically added to the entitlement database
- ▶ **ibmcountry**: Specifies the country ID used for entitlement and Call Home system

This command uses IBM Cloud APIs to get most of the information that are needed to enable the email functions, such as the contact information and detailed address of the machine (software datacenter).

The IP quorum server is chosen here because it is required to have public network access, and an SMTP server is suggested to be configured there. If necessary, the **chemailserver** or **mkemailserver** commands can be used after running this command to update it with another SMTP server, or add a new one.

For more information about the Home Call configuration, see Chapter 6, “Supporting the solution” on page 135 in this book.

4.2.6 Upgrading to second I/O group

In this section, we describe how to upgrade an IBM Spectrum Virtualize cluster from a two-nodes (single iogrp) configuration to a four-nodes (dual iogrp) configuration.

We are assuming that an IBM Spectrum Virtualize cluster with two nodes is up and running. To add two nodes to your cluster, complete the following steps:

1. Gather all of the information that is required for the sample yaml file as though you are performing the semi-automatic installation procedure.
2. Secure copy as part of the SSH/SFTP suite of tools (such as PuTTY) the `deploy_one_node.sh` script to the servers that will hold the Spectrum Virtualize instances to be added into the running cluster.
3. Open an SSH session to the bare-metal server by using an SSH client of your choosing.
4. Run the `deploy_one_node.sh` script, as shown in Example 4-9 on page 98.

Note: The following command arguments are available for the script to run in the order indicated:

- ▶ Service IP address
- ▶ Service IP default gateway
- ▶ Service IP subnet mask
- ▶ Node IP 1 address
- ▶ Node IP 1 gateway
- ▶ Node IP 1 mask
- ▶ Node IP 2 address
- ▶ Node IP 2 gateway
- ▶ Node IP 2 mask
- ▶ Port ID of node IP 2 (normally the value 2)
- ▶ Serial number, server name
- ▶ Node ID.

All of these parameters are required for the script to successfully run the **sntask initnode** command.

Example 4-9 Initializing a single node

```
[root@svcln3 ~]# /root/deploy_one_node.sh 10.183.120.3 10.183.120.1
255.255.255.192 10.183.120.4 10.183.120.1 255.255.255.192 1 10.183.120.5
10.183.120.1 255.255.255.192 2 SL019TYC svcln3 3
```

<omitting bulk of intermediate script output>

Downloading, it may take a few minutes.
Installing, it may take a few minutes.
Spectrum-virtualize node is successfully installed.
Please reboot to complete installation.

Tip: If the script fails to successfully run **sntask initnode** and you are prompted to add the force flag to the command, you can use a text editor to add **-f** to the last command in the script and re-run it so that it looks like the following example:

```
/usr/bin/sntask initnode -f -sip ${1} -gw ${2} -mask ${3} -nodeip1 ${4}
-nodegw1 ${5} -nodemask1 ${6} -nodeport1 ${7} -nodeip2 ${8} -nodegw2 ${9}
-nodemask2 ${10} -nodeport2 ${11} -serial ${12} -name ${13} -id ${14} ${15}
```

5. Restart the server on which you completed the installation.
6. Reconnect to the server and run the **sninfo lsnonce** command, as shown in Example 4-10.

Example 4-10 Activating the node

```
[root@svcln3 ~]# sninfo lsnonce
NONCEO
```

7. As described in Step 3 of 4.1.2, “Fully Automated installation” on page 74, use the **nonce** command to get the activation key.
8. Secure copy as part of the SSH/SFTP suite of tools (such as PuTTY) the activation key to the node’s service IP upgrade directory by using and activate the node software, as shown in Example 4-11.

Example 4-11 Activating the node

```
jfincher$ scp NONCE0.txt superuser@10.183.120.3:/upgrade
The authenticity of host '10.183.120.3 (10.183.120.3)' can't be established.
ECDSA key fingerprint is SHA256:4KDS/hL1/tDUdtK78SxdUMjdp2WWPKwaTEXcMPg41A.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '10.183.120.3' (ECDSA) to the list of known hosts.
Password:
NONCE0.txt
100% 446    2.1KB/s   00:00
jfincher$ ssh -l superuser 10.183.120.3
Password:
IBM_Spectrum_Virtualize::superuser>satask chvpd -idfile /upgrade/NONCE0.txt
IBM_Spectrum_Virtualize::superuser>Connection to 10.183.120.3 closed by remote
host.
Connection to 10.183.120.3 closed.
```

Tip: The superuser password for the node at this point is passw0rd.

- Repeat steps 1 - 8 for the second node to add into the cluster.
- Log in into your running IBM Spectrum Virtualize cluster by using the **lsnodecandidate** command. You see the two new nodes that were configured into candidate state and made visible to the cluster through their private IP links, as shown in Example 4-12.

Example 4-12 lsnodecandidate command example

```
IBM_Spectrum_Virtualize:Cluster_10.183.120.10:superuser>lsnodecandidate -delim :
id:panel_name:UPS_serial_number:UPS_unique_id:hardware:serial_number:product_mtm:
machine_signature
5005076071000850:NONCE0:::SW1:NONCE0:0002-SW1:5678-F1F1-3AB2-06B6
5005076071000851:NONCE0:::SW1:NONCE0:0002-SW1:5678-F1F1-3AB2-06B7
```

The same check can be done by using the GUI, as shown in Figure 4-30.

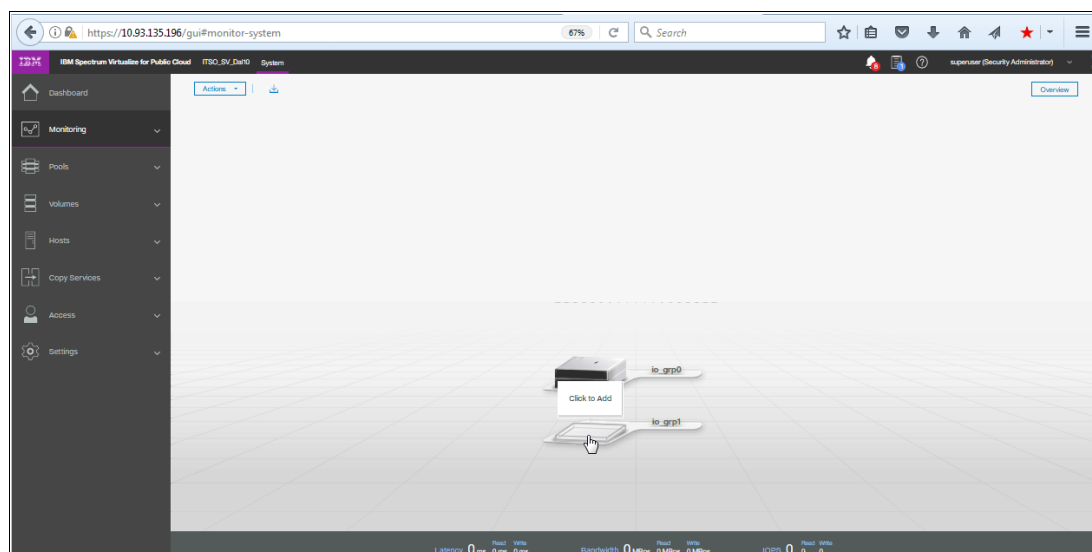


Figure 4-30 Candidate node example

- From the GUI (see Figure 4-30), click **Click to add** to add the two nodes. You are redirected to the next window, as shown in Figure 4-31.

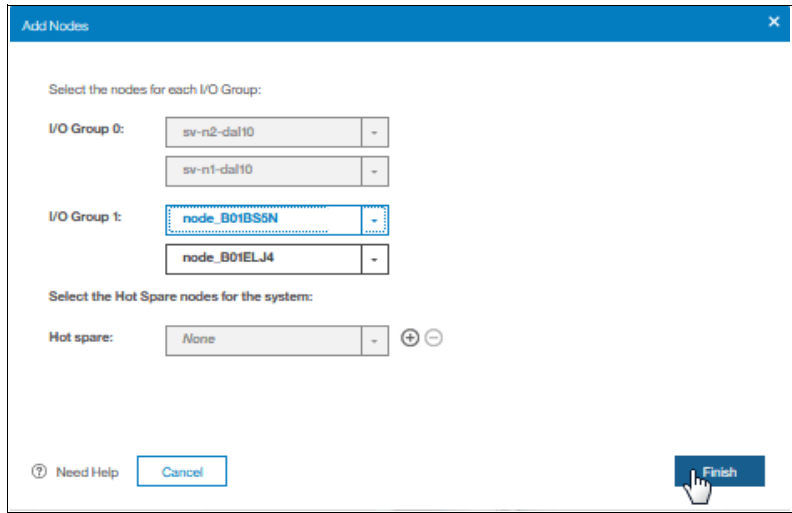


Figure 4-31 Adding nodes example

You are redirected to the completion window, as shown in Figure 4-32.

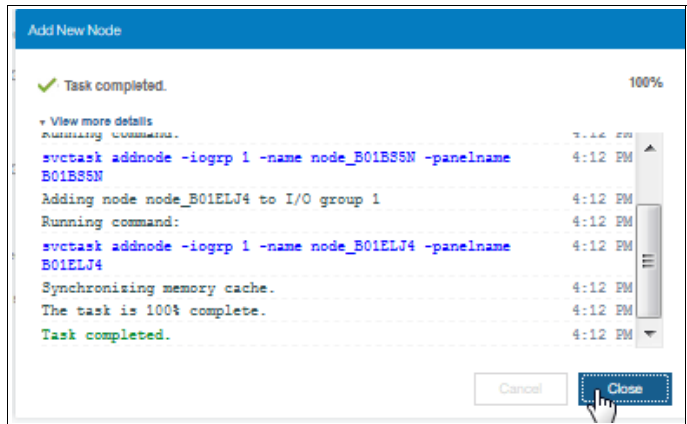


Figure 4-32 Completion window

12. You can now check that your new nodes are added to your cluster, as shown in Figure 4-33.

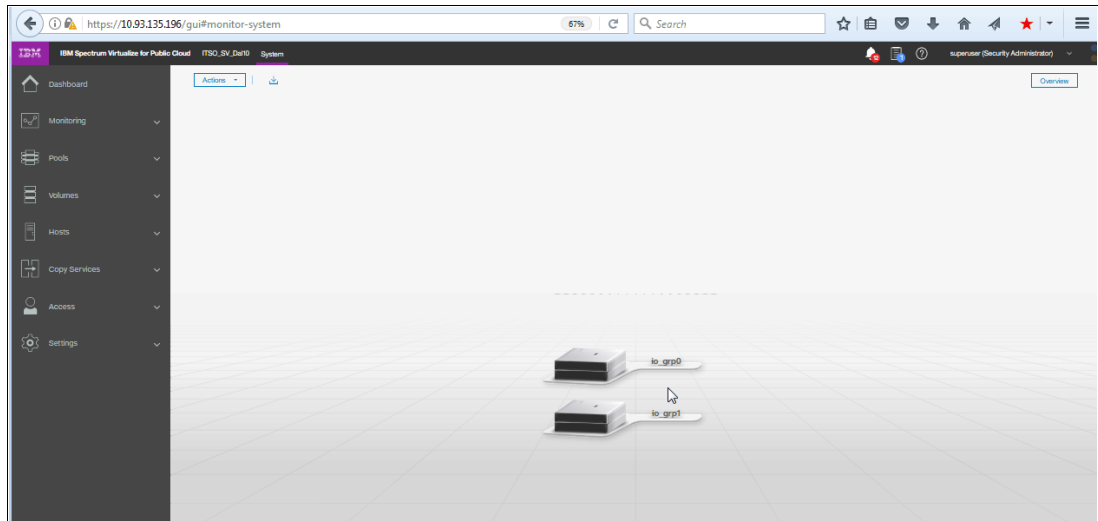


Figure 4-33 Two nodes added example

13. Change the node names according to your naming convention by using the GUI, as shown in Figure 4-34.

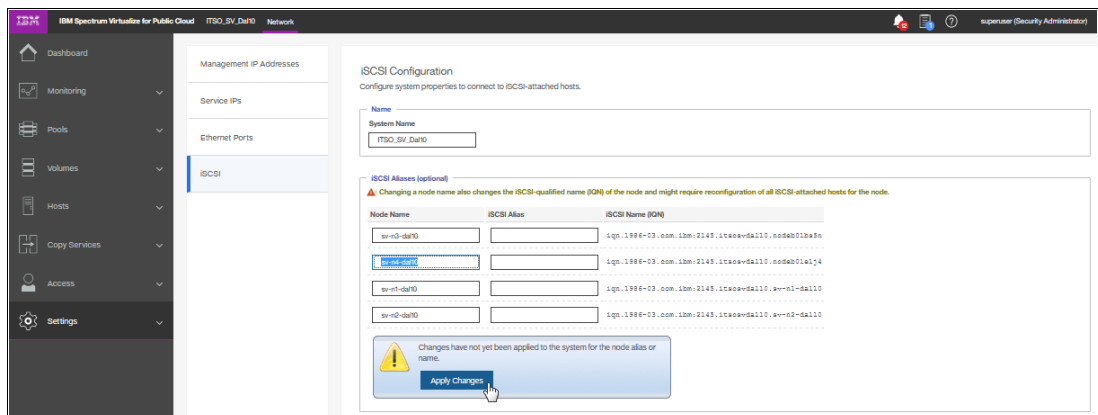


Figure 4-34 Changing the node name example

14. Configure your new node's port IPs by using the `cfgportip` command, as shown in Example 4-13.

Example 4-13 Configuring the port IP addresses

```

IBM_Spectrum_Virtualize:Cluster_10.183.120.10:superuser> svctask cfgportip -node 3
-ip 10.183.120.25 -gw 10.183.120.1 -mask 255.255.255.192 -storage yes 1
IBM_Spectrum_Virtualize:Cluster_10.183.120.10:superuser> svctask cfgportip -node 3
-ip 10.183.120.26 -gw 10.183.120.1 -mask 255.255.255.192 -storage yes 2
IBM_Spectrum_Virtualize:Cluster_10.183.120.10:superuser> svctask cfgportip -node 4
-ip 10.183.120.27 -gw 10.183.120.1 -mask 255.255.255.192 -storage yes 1
IBM_Spectrum_Virtualize:Cluster_10.183.120.10:superuser> svctask cfgportip -node 4
-ip 10.183.120.28 -gw 10.183.120.1 -mask 255.255.255.192 -storage yes 2
  
```

15. Authorize your back-end storage to the new nodes port IP addresses, as shown in Figure 4-35 and Figure 4-36 on page 102 by using IBM the Cloud IaaS web portal. Repeat the steps for all your MDisks that you want to make visible to the new nodes.

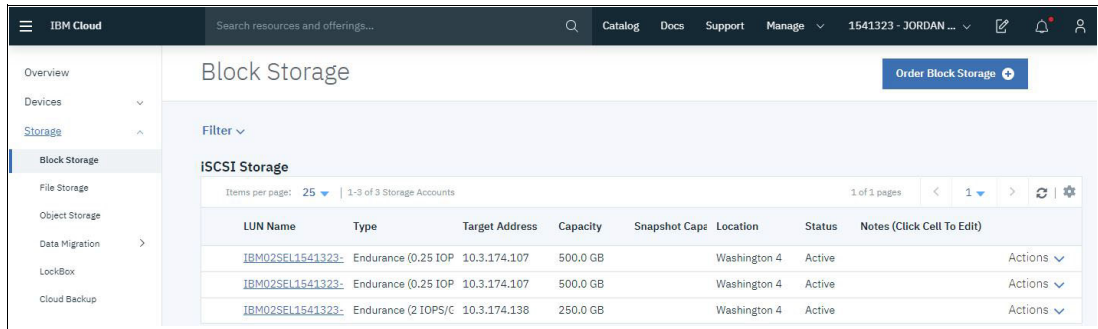


Figure 4-35 Authorize storage example

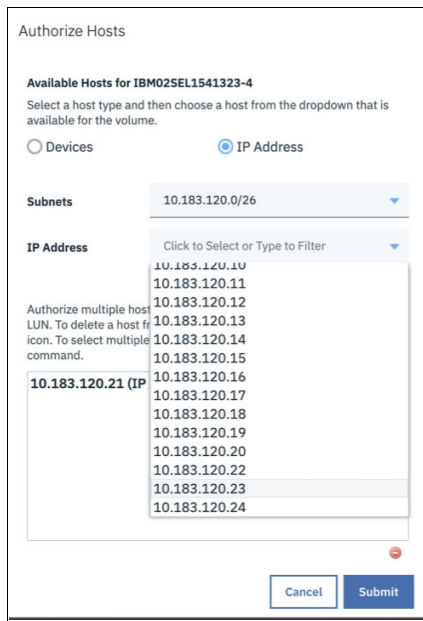


Figure 4-36 Authorize host example

Note: Each block storage LUN must be authorized by IP address. Each LUN can be authorized to a single IP address per node only. The IP addresses that are used must correlate to the same port number on each node in the cluster (that is, LUN 1 is authorized to the IP addresses that corresponds to port 1 on each node in the cluster).

- Validate that the port IP addresses are configured for each of the nodes and that the Storage Port IPv4 is listed as enabled for all IP addresses to be used to access storage. The Port IP address configuration is shown in Figure 4-37.

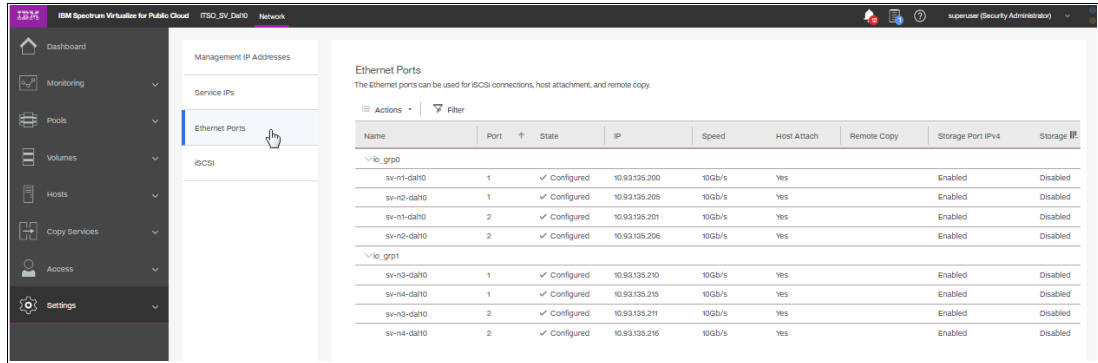


Figure 4-37 Port IP address configuration from GUI

17. Find the iSCSI Qualified Name (IQN) that the IBM Cloud assigned to the IP address when you authorized the IP address to the LUN. This information can be found in the block storage LUN details, as shown in Figure 4-38 on page 103.

Authorized Hosts							Authorize Host
Device Name	IP Address	Username	Password	Host IQN	Device Type		
	10.183.120.	IBM02SU1541323-I105805905	BzT2RVrnbX3GKpk5	iqn.2018-04.com.ibm:ibm02su1541323-i105805903	IP Address		Actions
	10.183.120.	IBM02SU1541323-I105805905	HweGA68sN7eeL7cT	iqn.2019-05.com.ibm:ibm02su1541323-i105805911	IP Address		Actions

Figure 4-38 Authorization details

18. When you expand the capacity of your IBM Spectrum Virtualize system from two nodes to four nodes, you have IBM Cloud storage that is managed by the two existing nodes. The two new nodes cannot access this storage until you synchronize your user name, password, and IQN in the IBM Spectrum Virtualize for Public Cloud software on each of the new nodes. You now must run through the CLI procedure as shown in Example 4-14.

The required credentials can be obtained from the IBM Cloud IaaS web portal. If you do not run this command, your MDisks are not accessible by the new nodes and they are in degraded state because a fundamental requirement (except in stretch clusters and HyperSwap configurations) of IBM Spectrum Virtualize is that all MDisks be visible to all nodes in the cluster.

Example 4-14 Adding storage access to the new nodes

```
IBM_Spectrum_Virtualize:Cluster_10.183.120.10:superuser>svctask chiscsiportauth
-src_ip 10.183.120.21 -iqn iqn.2018-04.com.ibm:ibm02su1541323-i105805905 -username
IBM02SU1541323-I105805905 -chapsecret MqZezxFbHMv7qdEQ
```

```
IBM_Spectrum_Virtualize:Cluster_10.183.120.10:superuser>svctask chiscsiportauth
-src_ip 10.183.120.23 -iqn iqn.2018-04.com.ibm:ibm02su1541323-i105805909 -username
IBM02SU1541323-I105805909 -chapsecret MrSL5DDey5vavQaU
```

```
IBM_Spectrum_Virtualize:Cluster_10.183.120.10:superuser>svctask
detectiscsistorageportcandidate -srcportid 2 -targetip 10.3.174.137
```


4.3 Configuring replication from on-prem IBM Spectrum Virtualize to IBM Spectrum Virtualize for IBM Cloud

In this section, we describe how to configure a replication from an on-prem solution that can be a Storwize or IBM SAN Volume Controller to IBM Spectrum Virtualize for IBM Cloud solution.

Our example uses a Storwize system in the on-prem data center and a four-nodes IBM Spectrum Virtualize for IBM Cloud as a DR Storage solution.

The scenario we are describing uses IBM Spectrum Virtualize Global Mirror with Change Volume (GM-CV) to replicate the data from the on-prem data center to IBM Cloud.

This implementation starts with the assumption that the IP connectivity between on-prem and IBM Cloud was established through a MPLS or VPN connection. Because several methods are available to implement the IP connectivity, this book does not consider that specific configuration. For more information, contact your IBM Cloud Technical Specialist.

To configure the GM-CV, complete the following steps:

1. Configure your IBM Spectrum Virtualize Private IP ports to be enabled for Remote Copy. This configuration is required on both sites, as shown in Figure 4-39.

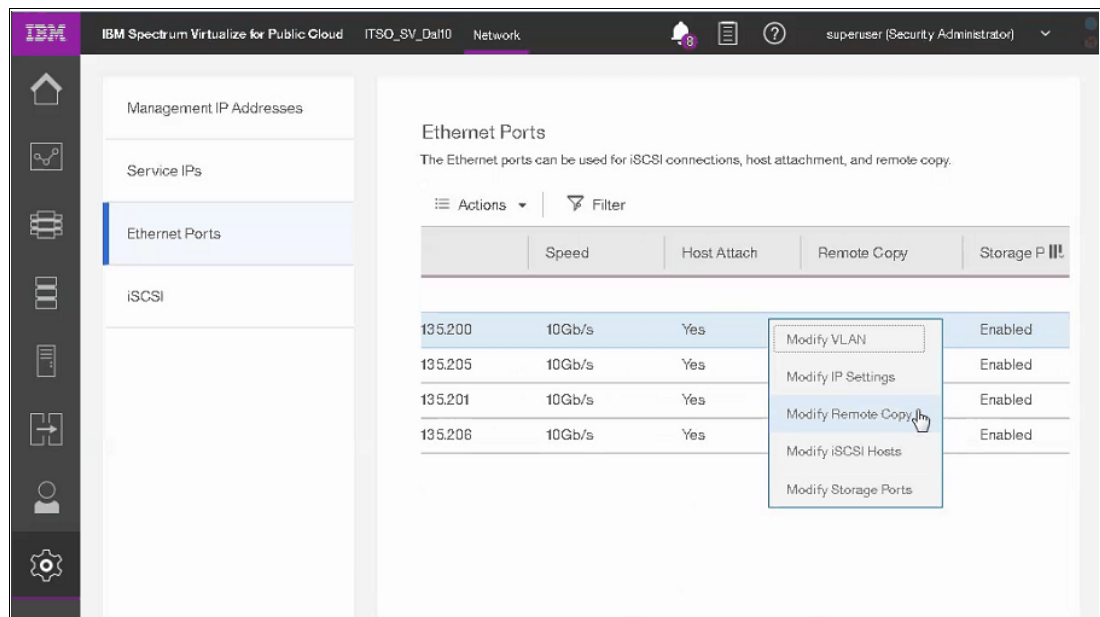


Figure 4-39 Remote Copy IP port

2. You are prompted to choose which copy group to use, as shown in Figure 4-40.




Figure 4-40 Group 1 configuration

- Repeat the previous steps for all of the IP ports that you want to configure. A configuration is created that is similar to the configuration that is shown in Figure 4-41.

Speed	Host Attach	Remote Copy	Storage Port IPv4
10Gb/s	Yes	Copy Group 1	Enabled
10Gb/s	Yes	Copy Group 1	Enabled
10Gb/s	Yes	Copy Group 1	Enabled
10Gb/s	Yes	Copy Group 1	Enabled
10Gb/s	Yes	Copy Group 1	Enabled
10Gb/s	Yes	Copy Group 1	Enabled
10Gb/s	Yes	Copy Group 1	Enabled
10Gb/s	Yes	Copy Group 1	Enabled

Figure 4-41 IBM Spectrum Virtualize for IBM Cloud configuration completed

- Run the same configuration for the on-prem Storwize Storage system or IBM SAN Volume Controller, as shown in Figure 4-42, Figure 4-43 on page 107, and Figure 4-44 on page 107.

Note: The on-prem solution has a different GUI because it is running on an older IBM Spectrum Virtualize software version than the version that is installed on the IBM Spectrum Virtualize in IBM Cloud. For more information about supported and interoperability versions, see the IBM interoperability matrix at [this web page](#).

Name	Port	State	IP	Speed	Host Attach	IPv4 Remote Copy
▼ io_grp0						
node1	1	✓ Configured	172.23.80.152	10Gb/s	Yes	Disabled
node2	1	✓ Configured	172.23.76.137	10Gb/s	Yes	Enabled
node1	2	✓ Configured	172.23.80.42	10Gb/s	Yes	Sup 1
node2	2	✓ Configured	172.23.66.218	10Gb/s	Yes	Sup 2

Figure 4-42 On-prem configuration example

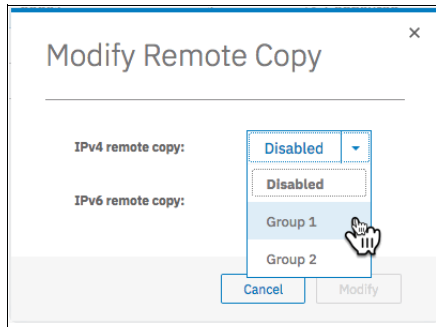


Figure 4-43 On-prem Copy Group example

Ethernet Ports
The Ethernet ports can be used for iSCSI connections, host attachment, and remote copy.

Actions | Filter

Speed	Host Attach	Remote Copy	Storage Port IPv4
10Gb/s	Yes	Copy Group 1	Enabled
10Gb/s	Yes	Copy Group 1	Enabled
10Gb/s	Yes	Copy Group 1	Enabled
10Gb/s	Yes	Copy Group 1	Enabled
10Gb/s	Yes	Copy Group 1	Enabled
10Gb/s	Yes	Copy Group 1	Enabled
10Gb/s	Yes	Copy Group 1	Enabled
10Gb/s	Yes	Copy Group 1	Enabled
10Gb/s	Yes	Copy Group 1	Enabled

Figure 4-44 On-prem configuration completed

5. Create a Cluster partnership between on-prem and IBM Spectrum Virtualize for IBM Cloud from the on-prem GUI, as shown in Figure 4-45.

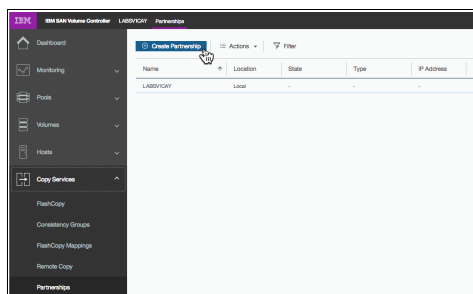


Figure 4-45 Creating partnership

- Complete the partnership creation from on-prem, as shown in Figure 4-46 and Figure 4-47.

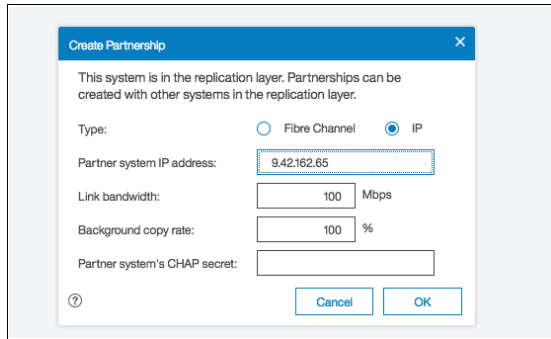


Figure 4-46 Insert IP address




Figure 4-47 Partnership partially configured

As you can see in Figure 4-47, the partnership is partially completed. You must complete the partnership on the IBM Spectrum Virtualize for IBM Cloud GUI, as shown in Figure 4-48 and Figure 4-49 on page 109.

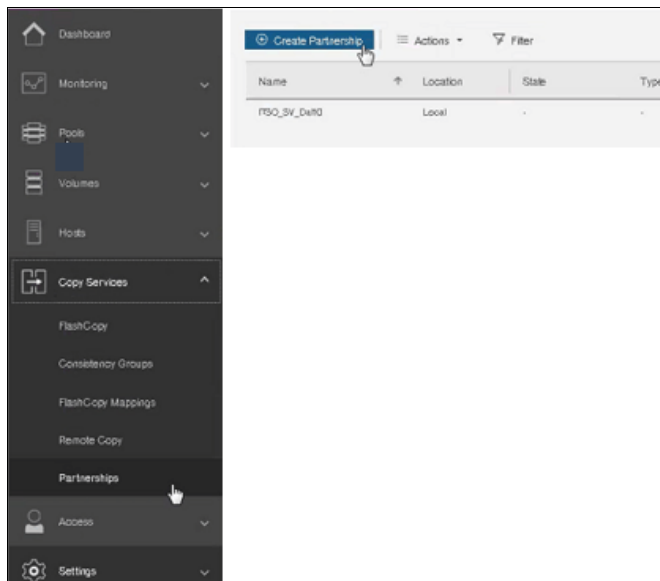


Figure 4-48 Creating partnership

Figure 4-49 Partnership example

When completed, your partnership is fully configured, as shown in Figure 4-50.

Name	Location	State	Type	IP Address
ITSO_SV_Dalt0	Local	-	-	-
ETECH_STORWIZE	Remote	✓ Fully Configured	IPv4	10.10.2.22

Figure 4-50 Fully configured

- In our example, we have an on-prem 100 GiB volume with its Change Volume (CV) that must be replicated to a 100 GiB volume in the IBM Cloud that is defined in our IBM Spectrum Virtualize for Public Cloud. The on-prem volumes are thin-provisioned, but this is not a specific requirement; instead, it is a choice. The CV can be thin-provisioned or fully provisioned, regardless of whether the master or auxiliary volume is thinly provisioned or space-efficient.

The CV only needs to store the changes accumulated during the cycle period and should therefore use as real capacity as possible (see Figure 4-51).

Name	State	Pool	Protocol Type	UID	Capacity
SVCPC_AWS	✓ Online	Pool1		600507607293A2FAA80000000000000...	100.00 GiB
SVCPC_AWS_CV	✓ Online	Pool1		600507607293A2FAA80000000000000...	100.00 GiB

Figure 4-51 On-prem volumes

8. Create a volume remote copy relationship for a GM-CV from on-prem, as shown in Figure 4-52.




Figure 4-52 Create relationship

9. Select the type of relationship, as shown in Figure 4-53.

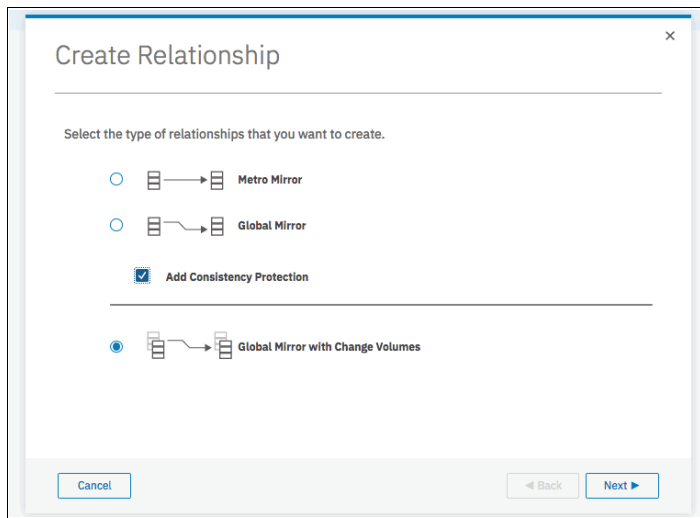


Figure 4-53 GM CV

10. Select the remote system (as shown in Figure 4-54), and select the volumes that must be in relationship (as shown in Figure 4-55 on page 111).

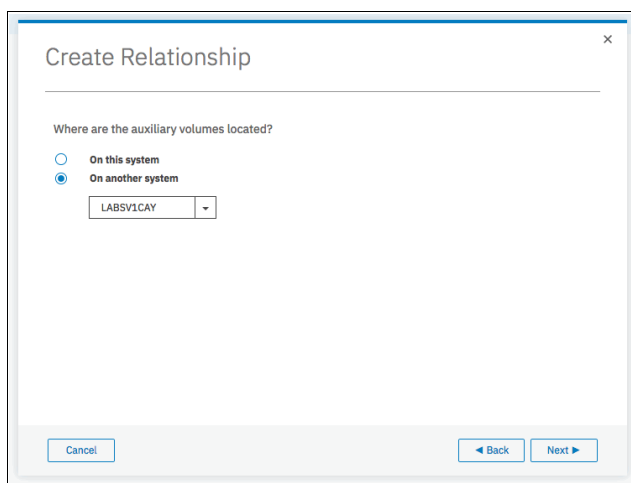


Figure 4-54 Remote system

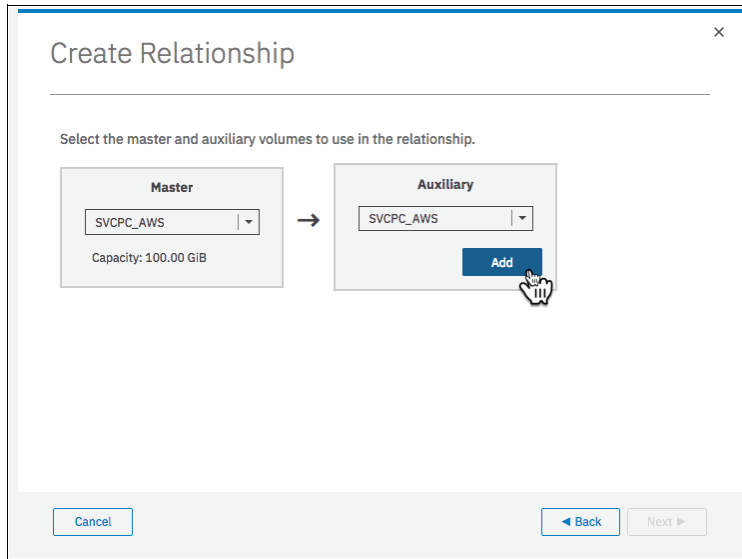


Figure 4-55 Master and auxiliary volumes example

In our example, we choose **No, do not add a master change volume** at this time, as shown in Figure 4-56.

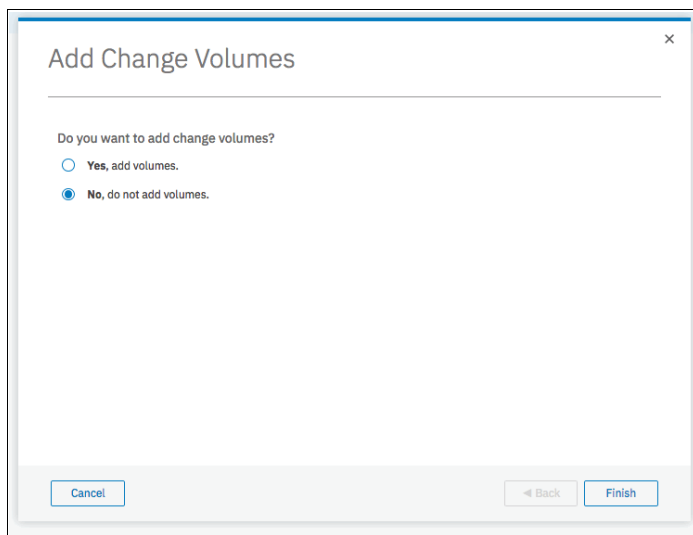


Figure 4-56 Do not add change volume

11. The volumes are added later. We choose **No, do not start copying**, as shown in Figure 4-57.

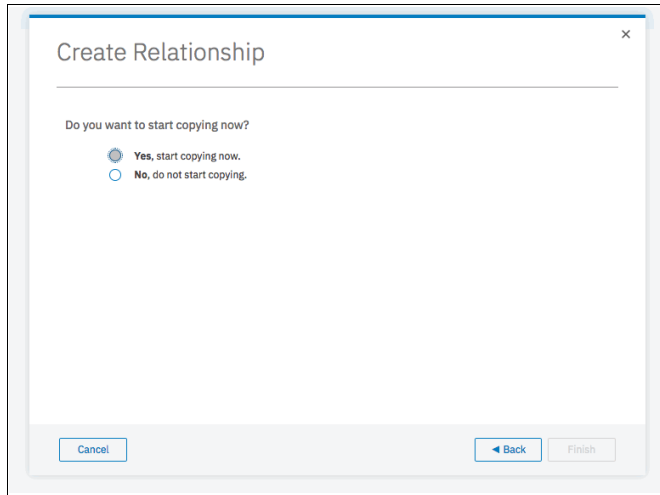


Figure 4-57 Do not start relationship example

12. Add the CV volumes to your relationship on both sides, as shown in Figure 4-58, Figure 4-59, and Figure 4-60 on page 113.

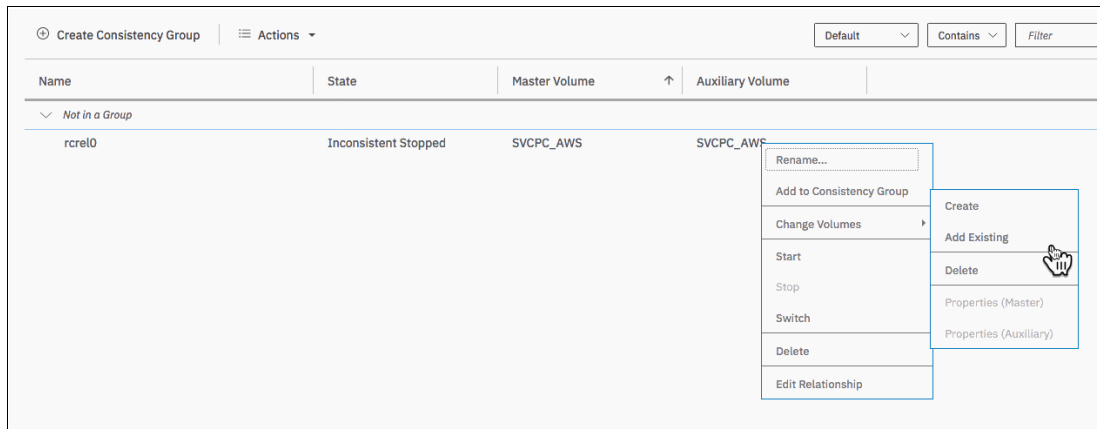


Figure 4-58 Add change volume from on-prem site

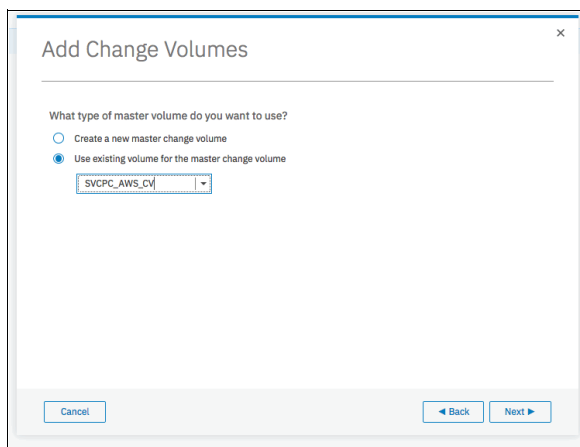


Figure 4-59 Choose the change volume from on-prem site

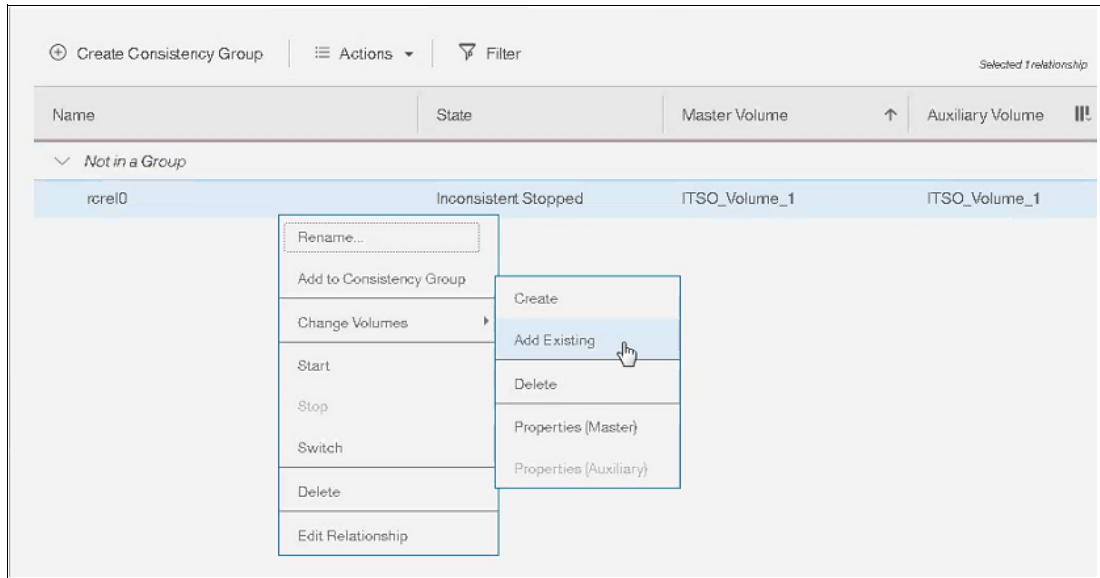


Figure 4-60 Add change volume to IBM Cloud site

13. Start your relationship from the on-prem site, as shown in Figure 4-61.

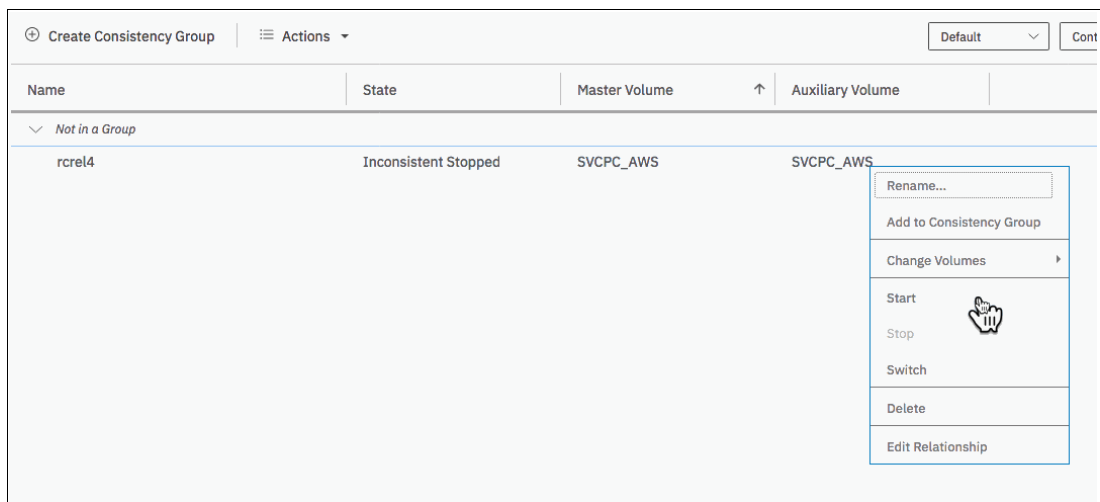


Figure 4-61 Start relationship

14. Create a GM consistency group and add your relationship to it, as shown in Figure 4-62 on page 114 and Figure 4-63 on page 114.

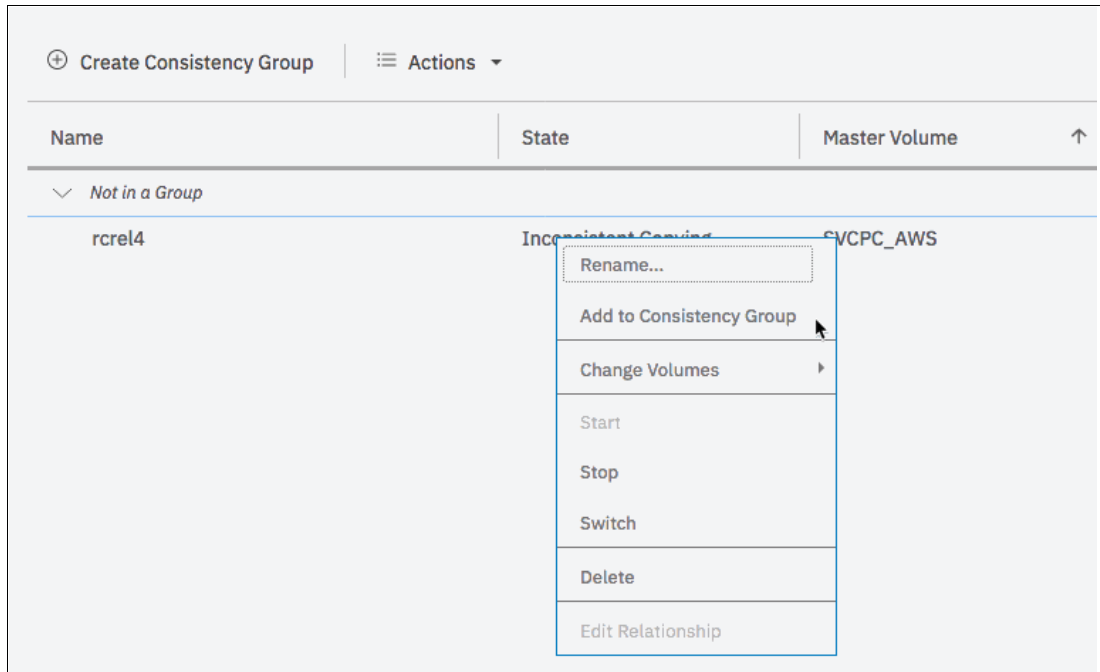


Figure 4-62 Add consistency group

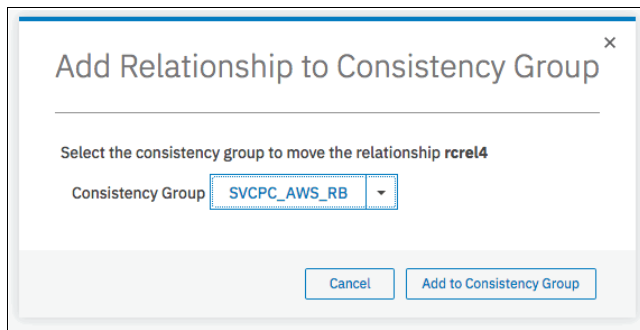


Figure 4-63 add relationship to a consistency group

Now you can see the status of your consistency group, as shown in Figure 4-64 and Figure 4-65 on page 115.

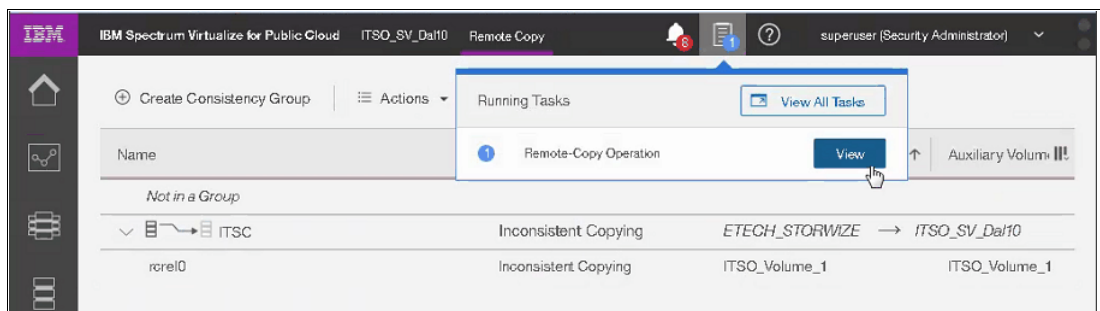


Figure 4-64 Consistency group status




Figure 4-65 Copying status

In our example, we show the status from IBM Spectrum Virtualize for IBM Cloud GUI.

When the copy approaches completion, the CV algorithm starts to prepare a freeze time in accordance with the cycling windows as defined in Figure 4-57 on page 112. When your copy reaches 100%, a FlashCopy is taken from the Auxiliary Volume to the Auxiliary-CV to be used if a real disaster or DR test occurs. At 100%, the status is “Consistent Copying”, as shown in Figure 4-66.

Name	State	Master Volume	Auxiliary Volume
<i>Not in a Group</i>			
ITSQ	Consistent Copying	ETECH_STORWIZE → ITSQ_SV_Dal10 Freeze Time: Nov 16, 2017, 7:41:54 PM	
rrel0	Consistent Copying	ITSQ_Volume_1	ITSQ_Volume_1

Figure 4-66 Consistency group status

This example shows how to configure a GM-CV relationship from an on-prem solution to an IBM Spectrum Virtualize for IBM Cloud solution.

It can be valuable to configure a snapshot (FlashCopy) of your GM-CV auxiliary volume to be used for DR Test or other purposes.

The steps that were shown in this example used the GUI, but they can also be run with the CLI.

For more information about how to manage Storwize or IBM Spectrum Virtualize or SVC Copy Functions, see the following publications:

- ▶ *Implementing the IBM Storwize V7000 with IBM Spectrum Virtualize V8.1, SG24-7938*
- ▶ *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 V8.1, SG24-7933*

4.4 Configuring Remote Support Proxy

The *Remote Support Proxy* (RSP) is a server that can be deployed to use the remote support assistance features that are offered in the IBM Spectrum Virtualize software. This section describes how to install the remote support proxy server and configure the proxy in IBM Spectrum Virtualize to enable remote support connections into the cluster.

Configuring the Remote Support Proxy Server

For the purposes of this IBM Redpaper publication, we assume that a separate virtual server is created in the environment that can access to the public network and the private network, including routes to the subnet in which IBM Spectrum Virtualize is running. Also, for the purposes of this publication, we assume that the virtual server deployed is RedHat Linux 7.x.

The first step is to obtain the remote support proxy software from your product support page. At the time of this writing, this code is under the Others category, as shown in Figure 4-67.



Figure 4-67 Downloading code on product support page

After the code is downloaded to the administrators notebook, you must upload the file to the server in which the proxy is to be installed. This process can be done by using the `scp` command. You also must install the `redhat-lsb` package if it is not installed.

When the file is uploaded to the server and all pre-requisite packages are installed, you can proceed with the installation, as shown in Example 4-16.

Example 4-16 Installing the Remote Support Proxy

```
[root@itso-dal10-sv-rsp ~]# chmod +x
supportcenter_proxy-installer-rpm-1.3.2.1-b1501.rhel7.x86_64.bin
[root@itso-dal10-sv-rsp ~]# ./supportcenter_proxy-installer-rpm-1.3.2.1-b1501.rhel7.x86_64.bin
Starting installer, please wait...
```

Tip: For the installation to succeed, ensure that the required packages are installed. On Red Hat systems, install the `redhat-lsb` package. On SUSE systems, install the `insserv` package. In both cases, install `bzip2`.

When the installer is started, you are presented with the International License Agreement for Non-Warranted Programs. To complete the installation, enter 1 to accept the license agreement and complete the installation.

When the installation completes, you must configure the proxy server to listen for connections. You can do this by editing the configuration file `supportcenter/proxy.conf`, which is in the `/etc` directory. The minimum modification required is to edit the fields `ListenInterface` and `ListenPort`. By default, the file includes “?” as the value for both of these fields.

To complete the configuration, specify the ListenInterface to be the interface name in Linux that can access the IBM Spectrum Virtualize clusters. This can be determined by using the **ifconfig** command, and identifying the interface that accesses the IBM Cloud private network. Also, set the ListenPort to the TCP port number to listen on for remote support requests. A sample configuration file is shown in Example 4-17.

Example 4-17 Sample Proxy Configuration

```
[root@itso-dal10-sv-rsp ~]# ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 10.93.4.91 netmask 255.255.255.192 broadcast 10.93.4.127
    inet6 fe80::490:fbff:fed6:7120 prefixlen 64 scopeid 0x20<link>
    ether 06:90:fb:d6:71:20 txqueuelen 1000 (Ethernet)
    RX packets 58690 bytes 59492454 (56.7 MiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 15492 bytes 2239603 (2.1 MiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

eth1: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 169.60.4.120 netmask 255.255.255.240 broadcast 169.60.4.127
    inet6 fe80::466:88ff:fe56:d4c prefixlen 64 scopeid 0x20<link>
    ether 06:66:88:56:0d:4c txqueuelen 1000 (Ethernet)
    RX packets 268 bytes 35536 (34.7 KiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 202 bytes 15628 (15.2 KiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

lo: flags=73<UP,LOOPBACK,RUNNING> mtu 65536
    inet 127.0.0.1 netmask 255.0.0.0
    inet6 ::1 prefixlen 128 scopeid 0x10<host>
    loop txqueuelen 1 (Local Loopback)
    RX packets 46 bytes 2693 (2.6 KiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 46 bytes 2693 (2.6 KiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

[root@itso-dal10-sv-rsp ~]# cat /etc/supportcenter/proxy.conf
# Configuration file for remote support proxy 1.3

# Mandatory configuration

# Network interface and port that the storage system will connect to
ListenInterface eth0
ListenPort 8988

#Remote support for SVC and Storwize systems on the following front servers
ServerAddress1 129.33.206.139
ServerPort1 443
ServerAddress2 204.146.30.139
ServerPort2 443

# Optional configuration

# Network interface (lo for local) for status queries
# StatusInterface ?
# StatusPort ?
```

```
# HTTP proxy for connecting to the Internet
# HTTPProxyHost ?
# HTTPProxyPort ?
# Optional authentication data for HTTP proxy
# HTTPProxyUser ?
# HTTPProxyPassword ?

# External logger (default is none)
# Logger /usr/share/supportcenter/syslog-logger

# Restricted user
# User nobody

# Log file
# LogFile /var/log/supportcenter_proxy.log

# Optional debug messages for troubleshooting
# DebugLog No

# Control IPv4/IPv6 usage
# UseIPv4 yes
# UseIPv6 yes
# UseIPv6LinkLocalAddress no
```

When the service is configured, the service must be started to allow the server to start listening for requests. Optionally, you can also configure the service to start on system start. To start the service, you can use the **service** or **systemctl** command.

To have the service start on system start, you can use the **chkconfig** command. Both of these processes are shown in Example 4-18.

Example 4-18 Starting the service

```
[root@itso-dal10-sv-rsp ~]# service supportcenter_proxy start
Starting IBM remote support proxy: [ OK ]
[root@itso-dal10-sv-rsp ~]# chkconfig supportcenter_proxy on
```

When the service is started, you are ready to configure IBM Spectrum Virtualize to use the proxy to initiate remote support requests.



Typical use cases for IBM Spectrum Virtualize for Public Cloud

In this chapter, we describe four use cases for IBM Spectrum Virtualize for Public Cloud. This chapter includes the following topics:

- ▶ 5.1, “Whole IT services deployed in the Public Cloud” on page 120
- ▶ 5.2, “Disaster recovery” on page 124
- ▶ 5.3, “IBM FlashCopy in the Public Cloud” on page 128
- ▶ 5.4, “Workload relocation into the Public Cloud” on page 131

5.1 Whole IT services deployed in the Public Cloud

Companies are approaching and leveraging Public Cloud services from multiple angles: users that are rewriting and modernizing applications for cloud complement those looking to move to cloud only new services or to extend the existing IT into a hybrid model to quickly address changing capacity and scalability requirements. There are different delivery models for Public Cloud, such as SaaS, PaaS, and IaaS. In Chapter 1, “Introduction” on page 1, the overarching Public Cloud application and workload deployment can be seen as composed of two major use cases:

- ▶ Hybrid cloud: The integration between the off-premise Public Cloud services with an existing on-premises IT environment.
- ▶ Cloud-native: The full application’s stack is moved to cloud as SaaS, PaaS, IaaS or as a combination of the three delivery models.

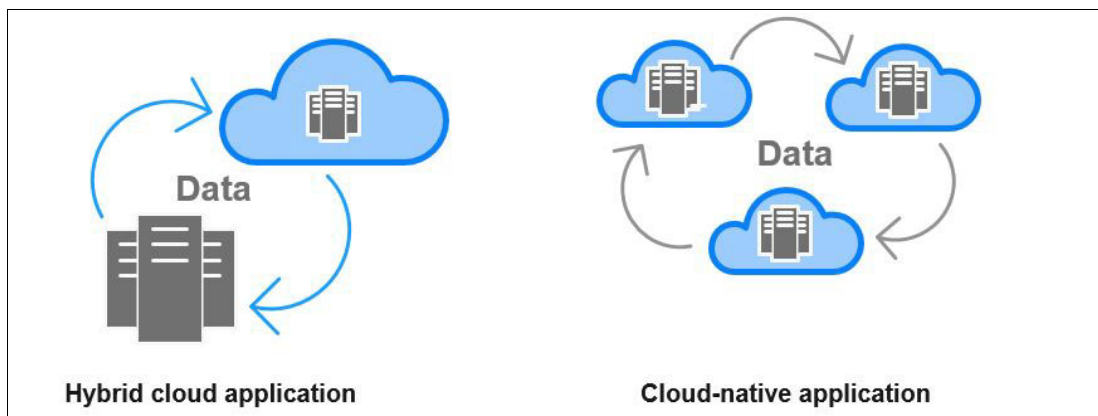


Figure 5-1 The two major deployment models for Public Cloud

Cloud-native implementations (aka whole IT services deployed in the Public Cloud) are composed of several use cases, all with the lowest common denominator of having a full application deployment in the Public Cloud data centers. The technical details and final architecture will depend, along with roles and responsibilities, on SaaS, PaaS or IaaS usage. Within the IaaS domain the transparency of cloud services is the highest, as the user’s visibility (and responsibility) into the application stack is much deeper compared to the other delivery models. On the other side, the *burden* for its deployment is higher as all the components have to be designed from the server up. IBM Spectrum Virtualize for Public Cloud, at the time of writing, is framed only within IaaS cloud delivery model, allowing the user to interact with their storage environment as they did on-prem, which provides more granular control over performance.

5.1.1 Business justification

A workload or an application that often is a good fit for a cloud-native deployment has the following characteristics:

- ▶ Is stand-alone, with few on-prem dependencies
- ▶ Relatively undemanding of I/O performance (low-latency/response-time and high IOPS)
- ▶ Is not processing highly regulated data

The drivers that motivate businesses towards cloud-native deployment span from capex and opex reduction, better resource management and controls against shadow IT, more flexibility and scalability along with drastically improved capillarity in delivering IT service due to the global footprint of the cloud data centers.

The cloud environment is highly focused on standardization and automation at its core. Because of this focus, the full spectrum of features and customization that are available in a typical on-premise or outsourcing deployment might not be natively available in the cloud catalog.

Nevertheless, the client does not lose the performances and capabilities when deploying a cloud-native application. In this sense, the storage virtualization with IBM Spectrum Virtualize for Public Cloud allows the IT staff to maintain the technical capabilities and skills to deploy, run, and manage highly available and highly reliable cloud-native applications in a Public Cloud.

In this context, the IBM Spectrum Virtualize for Public Cloud acts as a bridge between the standardized cloud delivery model and the enterprise assets the client leverages in their traditional IT environment.

5.1.2 Highly available deployment models

The architecture is directly responsible for the application's reliability and availability if a component failure (hardware and software) occurs. When the application is fully hosted on cloud, the cloud data center becomes the primary site (production site).

Cloud deployment does not guarantee 100% uptime, or that the backups are available by default or even that the application is automatically replicated between different sites. These security, availability, and recovery features are likely not client responsibility if the service is delivered in the SaaS model. It is partially the user's responsibility in PaaS, but is entirely the client's design responsibility in the case of the IaaS model.

Having reliable cloud deployments means meeting the required Service Level Agreement (SLA), a guaranteed service availability, and uptime. Companies using Public Cloud IaaS can meet required SLAs either by implementing highly available solutions and duplicating the infrastructure in the same data center or in two or more in-campus data centers (for example IBM Dallas10 and Dallas09) to maintain business continuity in case of failures. If business continuity is not enough to reach the desired SLA then disaster recovery (DR) implementations, splitting the application into multiple cloud data centers (usually with a distance of at least 300 km [186.4 miles]) prevents failure in case of a major disaster in the campus-area.

The following highly available deployment models are available for an application that is fully deployed on Public Cloud:

- ▶ On a single primary site

All the solution's components are duplicated (or more) within the same data center. This solution tolerate only the failure of single components, but not the data center unavailability.

- ▶ On multi-site

The architecture splits among multiple cloud data centers within the same campus to tolerate the failure of an entire datacenter or spread globally to recover the solution in case of a major disaster affecting the campus area.

Highly available cloud deployment on a single primary site

When fully moving an application to cloud IaaS as a primary site for service delivery, a reasonable approach is implementing at least a highly available architecture. By having each component (servers, network components and storage) redundant without any Single Point of Failures (SPoF).

Within the single site deployment, storage is usually deployed as native cloud storage. Leveraging Public Cloud catalog storage, users can use the intrinsic availability (and SLAs) of the storage service, whether as object storage (for example IBM Cloud Object Storage), file, or block storage. IBM Cloud block storage (delivered as Endurance or Performance format) is natively highly available (as multiple 9s).

A typical use case of the IBM Cloud highly available architecture is a VMware environment where physical hosts are N+1 with datastores that are hosted on the cloud block storage and shared simultaneously among multiple hosts.

IBM Spectrum Virtualize for Public Cloud, when deployed as clustered pair of Intel bare-metal servers, mediates the cloud block storage to the workload hosts. In the specific context of single site deployment, IBM Spectrum Virtualize for Public Cloud supports more features, which enhances the Public Cloud block-storage offering. This is true at the storage level where IBM Spectrum Virtualize for Public Cloud resolves some limitations because of the standardized model of the Public Cloud providers: a maximum number of LUNs per host, a maximum volume size, and poor granularity in the choice of tiers for storage snapshots.

IBM Spectrum Virtualize for Public Cloud also provides a new view for the storage management other than the cloud portal, which gives an high level view of the storage infrastructure and some limited specific operations at the volume level (such as volume size, IOPS tuning and snapshot space increase). What is not provided is a holistic view of the storage from the application perspective. More detailed reasons are highlighted in Table 5-1.

Table 5-1 Benefits of IBM Spectrum Virtualize for Public Cloud on single site deployment

Feature	Benefits
Single point of control for cloud storage resources	<ul style="list-style-type: none"> ▶ Designed to increase management efficiency and to help to support application availability
Pools the capacity of multiple storage volumes	<ul style="list-style-type: none"> ▶ Helps to overcome volume size limitations ▶ Helps to manage storage as a resource to meet business requirements, and not just as a set of independent volumes ▶ Helps administrator to better deploy storage as required beyond traditional “islands” ▶ Can help to increase the use of storage assets ▶ Insulate applications from maintenance or changes to storage volume offering
Manage tiered storage	<ul style="list-style-type: none"> ▶ Helps to balance performance needs against infrastructures costs in a tiered storage environment ▶ Automated policy-driven control to put data in the right place at the right time automatically among different storage tiers/classes

Feature	Benefits
Easy-to-use IBM Storwize family management interface	<ul style="list-style-type: none"> ▶ Single interface for storage configuration, management, and service tasks regardless the configuration available from Public Cloud portal ▶ Helps administrators use storage assets/volumes more efficiently ▶ IBM Spectrum Control Insights and IBM Spectrum Protect for additional capabilities to manage capacity and performance
Dynamic data migration	<ul style="list-style-type: none"> ▶ Migrate data among volumes/LUNs without taking applications that use that data offline ▶ Manage and scale storage capacity without disrupting applications
Advanced network-based copy services	<ul style="list-style-type: none"> ▶ Copy data across multiple storage systems with IBM FlashCopy ▶ Copy data across metropolitan and global distances as needed to create high-availability storage solutions between multiple data centers
Thin provisioning and snapshot replication	<ul style="list-style-type: none"> ▶ Reduce volume requirements by using storage only when data changes ▶ Improve storage administrator productivity through automated on-demand storage provisioning ▶ Snapshots available on lower tier storage volumes
IBM Spectrum Protect Snapshot application-aware snapshots	<ul style="list-style-type: none"> ▶ 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
Third parties native integration	Integration with VMware vRealize

Highly available cloud deployment on multi-site

When the application architecture spans over multiple data centers, it can tolerate the failure of the entire primary data center by switching to the secondary allowing it to tolerate a major disaster affecting a wide area. The primary and secondary data centers can be deployed as:

- ▶ **Active-active:** The secondary site is always up and running and synchronously aligned with the primary.
- ▶ **Active-passive:** The secondary site is either always up but asynchronously replicated (with a specific RPO) or up only for specific situation acting as a recovery site or test environment. Storage is, of course, always active and available for data replication.

The active-passive is usually the best fit for many cloud use cases including the DR as shown in 5.2, “Disaster recovery” on page 124. The possibility to provision compute resources on-demand in a few minutes, having just the storage always provisioned and aligned with a specific RPO represent a huge driver for a cost effective DR infrastructure and lowers the TCO.

The replication among multiple cloud data centers is no different from the traditional approach, except for the number of available tools in cloud. The considerations that are described in 1.3.1, “Hybrid scenario: on-premises to IBM Cloud” on page 12 for a hybrid environment are still applicable.

Solutions that are based on hypervisor or application-layer replication, such as VMware, Veeam, and Zerto are available in the Public Cloud if the environment is heterogeneous (such as virtual servers, bare metals, and multiple hypervisor). Storage based replication is still the preferable approach.

Storage-based replication is available in almost every cloud provider. IBM Cloud for example allows for block-level replication on IBM Endurance storage. A volume can be replicated to another cloud site with a minimum recovery point objective (RPO) of 60 minutes. The replication features are specific to the cloud offering and are not editable nor tunable. For example, the remote copy of Endurance Storage is not accessible until unavailability of the primary is declared or its replica is limited to specific data centers pairs.

For this reason, the model does not fit all clients' requests. Some of these deltas are covered by an application or hypervisor level replica that is limited to a specific environment and specific infrastructure to replicate (for example VMware with vReplicator and DR startup managed by SRM).

However, asynchronous mirroring that uses Global Mirror with Change Volumes (GMCV) allows for a minimum RPO of 2 minutes (the change volume cycle period ranges from 1 minute to 1 day and we recommend setting the cycle period to be half of the RPO) and is capable of replicating a heterogeneous environment.

Also, Spectrum Virtualize supports several third-party integrations, such as VMware Site Recovery Manager (SRM), to automate failover at the application layer while the storage replica is used. SRM also automates the taking of storage snapshots with FlashCopy for testing purposes.

5.2 Disaster recovery

In 2018 and later years, as customers harness and secure proliferating data in their environment, infrastructure workloads will see the highest adoption increases.

Technology is just one crucial piece of a disaster recovery (DR) solution, and not the one that dictates the overall approach.

In this section we talk about the IBM Spectrum Virtualize for IBM Cloud DR approach and benefits. In addition in Appendix A, "Guidelines for disaster recovery solution in the Public Cloud" on page 157 we cover the suggested practices and some considerations you should take into account when creating a DR solution.

Disaster recovery strategy is the predominant aspect of the overall resiliency solution because it determines what classes of physical events the solution can address, and sets the requirements in terms of distance, and sets constraints on technology.

Considering the cloud space, most cloud providers offer a redundant infrastructure, with the following several layers of resiliency:

- ▶ Local: A physical and logical segregation zone (for example, availability zone or availability set) within a Cloud Service Provider (CSP) location (physical datacenter) that is independent from other zones for what pertains to power supply, cooling, and networking.
- ▶ Site: CSPs group multiple *sites* in a so-called *region*. Using different sites within the same region offers a better level of protection in cases of limited natural disaster, compared to two different zones on a single site because sites in the same region are usually in close proximity (tens of kilometers or miles).

- ▶ Region: Using two sites in two regions in the same geography represents the top level of protection against natural disasters because sites are generally over 400 km (248 miles) apart.
- ▶ Geography: Selecting two sites in different *geographies* extends the level of protection against natural disasters, as geographies are generally over 1500 km (1000 miles) apart, which represents the latest option in terms of wider protection requirements.

IBM Cloud operates over 60 datacenters in six regions and 18 availability zones around the world, as shown in Figure 5-2.

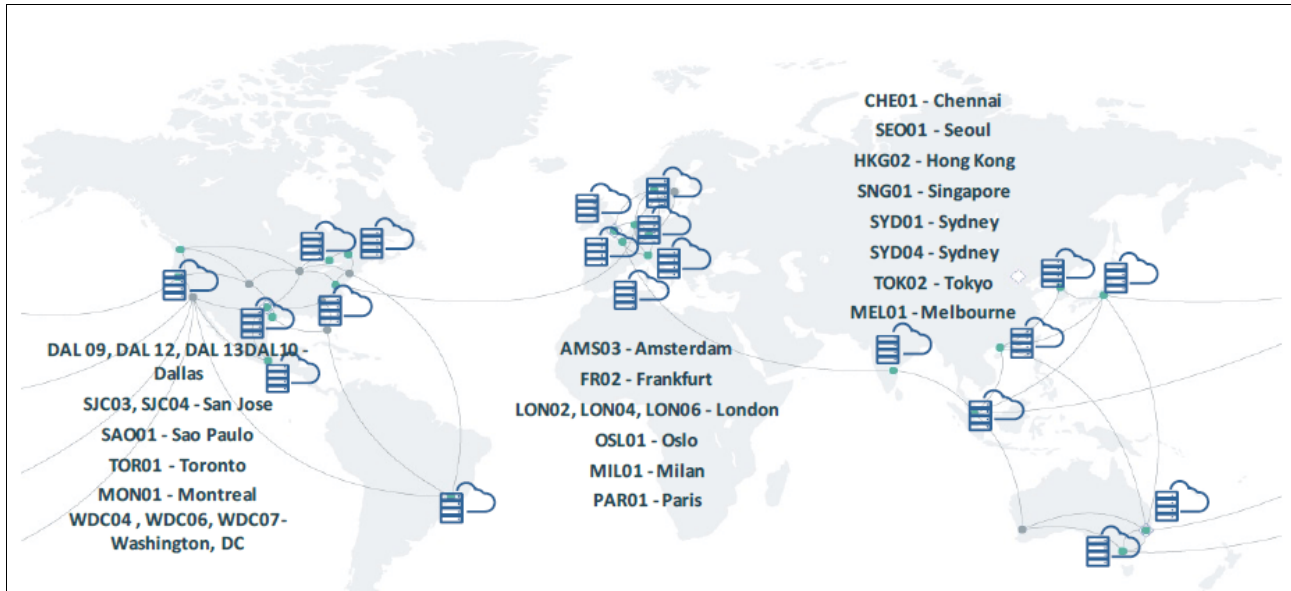


Figure 5-2 IBM Cloud with more being added all the time

For more information about data cloud centers, see this IBM Cloud [web page](#).

5.2.1 Business justification

Table 5-2 lists the drivers and the challenges of having a DR solution on cloud and the capabilities IBM Spectrum Virtualize for Public Cloud provides in these areas.

Table 5-2 Drivers and challenges and capabilities IBM Spectrum Virtualize for Public Cloud provide

Adoption drivers	Challenges	Spectrum Virtualize for IBM Public capabilities
The promise of reduced Opex and Capex	<ul style="list-style-type: none"> ▶ Hidden costs ▶ Availability of data when needed 	<ul style="list-style-type: none"> ▶ Optimized for Cloud Block Storage ▶ EasyTier solution to optimize the most valuable storage usage maximizing Cloud Block Storage performance ▶ Thin Provisioning to control the storage provisioning ▶ Snapshots feature for backup and DR solution ▶ High availability clusters architecture

Adoption drivers	Challenges	Spectrum Virtualize for IBM Public capabilities
Bridging technologies from on premises to cloud	<ul style="list-style-type: none"> ▶ Disparate Infrastructure – How can my on-premises production data be readily available in the cloud in case of a disaster? 	<ul style="list-style-type: none"> ▶ Any to any replication ▶ Supporting over 400 different storage devices (on prem) including iSCSI even on-premises than when deployed in cloud
Leveraging the cloud for Backup and Disaster Recovery	<ul style="list-style-type: none"> ▶ Covering virtual and physical environments ▶ Solutions to meet a range of RPO/RTO needs 	<ul style="list-style-type: none"> ▶ A storage based, serverless replication with options for low RPO/RTO named: <ul style="list-style-type: none"> – Global Mirror for Asynchronous replication with a RPO close to “0” – Metro Mirror for Synchronous replication – Global Mirror with Change Volumes for Asynchronous replication with a tunable RPO – Supports virtualized and bare-metal applications (unlike VM-based solutions)

At the time of this writing, IBM Spectrum Virtualize for Public Cloud includes the following DR-related features:

- ▶ Can be implemented in over 60 data centers in 19 countries. For more information, see this [web page](#).
- ▶ Was first available on IBM Cloud; Amazon Web Services Marketplace availability announced for June 25, 2019.
- ▶ Is deployed on IBM Cloud Infrastructure.
- ▶ Offers data replication with Storwize family, V9000, IBM SAN Volume Controller, FlashSystem 9100 or VersaStack and Public Cloud.
- ▶ Supports 2,4,6 or 8 node clusters in IBM Cloud.
- ▶ Offers data services for IBM Cloud Block Storage.
- ▶ Offers common management with IBM Spectrum Virtualize GUI with full admin access and dedicated instance.
- ▶ No incoming data transfer cost.
- ▶ No bandwidth cost within IBM Cloud.
- ▶ Replicates between IBM Cloud data centers.

5.2.2 Common DR scenarios with IBM Spectrum Virtualize for Public Cloud

The following most common scenarios can be implemented with IBM Spectrum Virtualize for Public Cloud:

- ▶ IBM Spectrum Virtualize Hybrid Cloud disaster recovery for “Any to Any”, Physical and Virtualized applications as shown in Figure 5-3 on page 127.

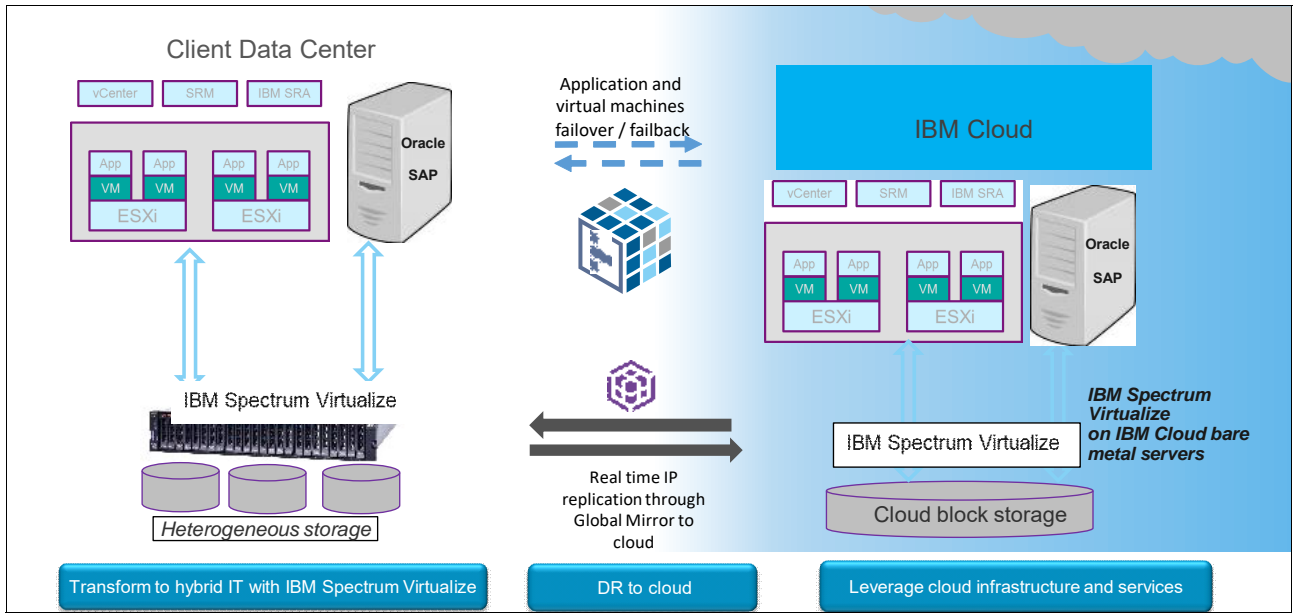


Figure 5-3 Hybrid scenarios

- ▶ IBM Spectrum Virtualize for Public Cloud DR solution with VMware Site Recovery Manager (SRM) as shown in Figure 5-4.

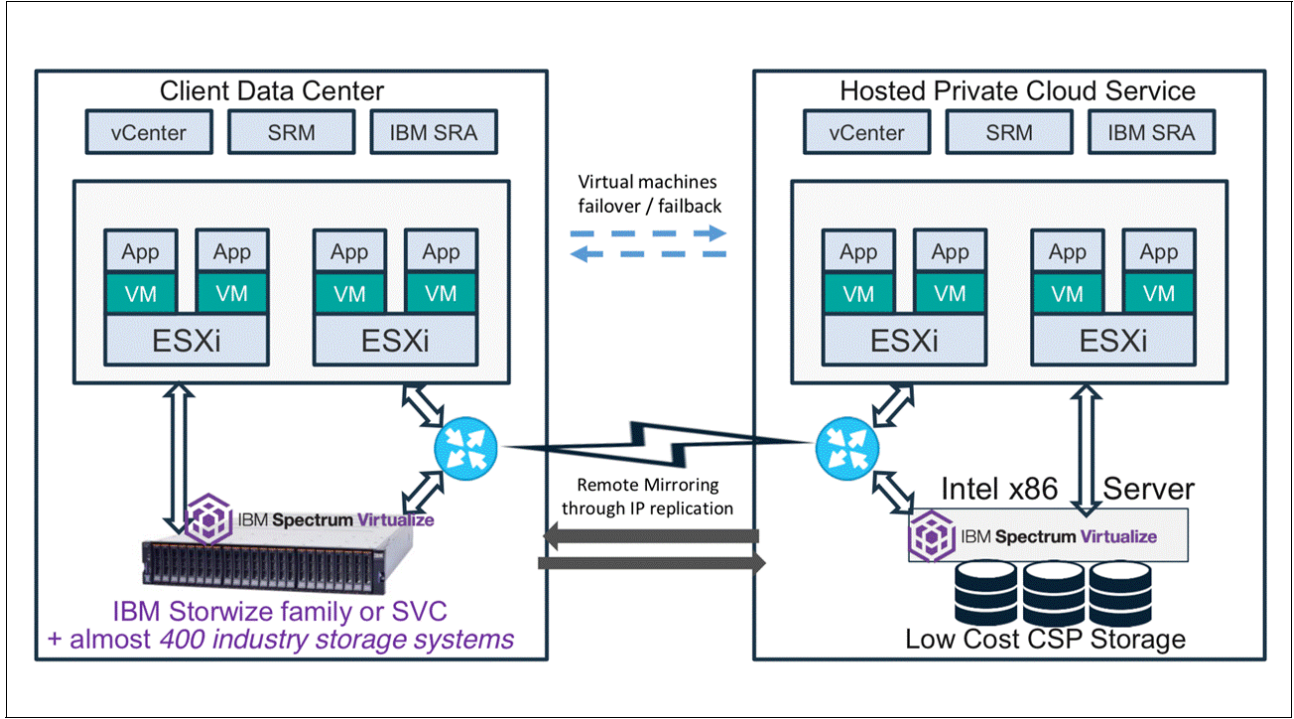


Figure 5-4 VMware with Site Recovery Manager scenario

As shown in Figure 5-4, a customer can deploy a storage replication infrastructure in a Public Cloud with the IBM Spectrum Virtualize for Public Cloud.

The following are the details of this scenario:

- ▶ Primary storage sits on customer's physical data center. Customer has on-premises SVC cluster installed or IBM Storwize solution.
- ▶ Secondary storage sits on the DR site which includes a virtual IBM Spectrum Virtualize cluster running in the Public Cloud.
- ▶ The virtual IBM Spectrum Virtualize cluster manages the storage provided by Cloud Service Provider (CSP).

A replication partnership that uses Global Mirror with Changed Volumes is established between on-premises IBM Spectrum Virtualize cluster or Storwize solution and the virtual IBM Spectrum Virtualize cluster to provide disaster recovery.

When talking about disaster recovery it is important to mention that IBM Spectrum Virtualize for Public Cloud is an important piece of a more complex solution that has some prerequisites considerations, and recommended best practices that need to be applied.

Note: Refer to Appendix A, "Guidelines for disaster recovery solution in the Public Cloud" on page 157 we cover preferred practices when designing a resiliency solution, and considerations for using the cloud space as a possible alternative site.

Also, to see an example of a simple implementation of a DR solution, including IBM Storwize and IBM Spectrum Virtualize for Public Cloud, see 4.3, "Configuring replication from on-prem IBM Spectrum Virtualize to IBM Spectrum Virtualize for IBM Cloud" on page 105.

5.3 IBM FlashCopy in the Public Cloud

The IBM FlashCopy function in IBM Spectrum Virtualize can perform a *point-in-time copy* of one or more volumes. You can use FlashCopy to help you solve critical and challenging business needs that require duplication of data of your source volume. Volumes can remain online and active while you create consistent copies of the data sets. Because the copy is performed at the block level, it operates below the host operating system and its cache. Therefore, the copy is not apparent to the host unless it is mapped.

5.3.1 Business justification

The business applications for FlashCopy are wide-ranging. Common use cases for FlashCopy include, but are not limited to, the following examples:

- ▶ Rapidly creating consistent backups of dynamically changing data.
- ▶ Rapidly creating consistent copies of production data to facilitate data movement or migration between hosts.
- ▶ Rapidly creating copies of production data sets for application development and testing.
- ▶ Rapidly creating copies of production data sets for auditing purposes and data mining.
- ▶ Rapidly creating copies of production data sets for quality assurance.
- ▶ Rapidly creating copies of replication targets for testing data integrity

Regardless of your business needs, FlashCopy within the IBM Spectrum Virtualize is flexible and offers a broad feature set, which makes it applicable to many scenarios.

5.3.2 FlashCopy mapping

The association between the source volume and the target volume is defined by a FlashCopy mapping. The Flashcopy mapping can have three different types, four attributes, and seven different states.

FlashCopy in the GUI can be one of the following types:

- ▶ Snapshot

Sometimes referred to as *nocopy*, a snapshot is a point-in-time copy of a volume without background copy of the data from the source volume to the target. Only the changed blocks on the source volume are copied. The target copy cannot be used without an active link to the source. This is achieved by setting the copy and clean rate to zero.

- ▶ Clone

Sometimes referred to as *full copy*, a clone is a point-in-time copy of a volume with background copy of the data from the source volume to the target. All blocks from the source volume are copied to the target volume. The target copy becomes a usable independent volume. This is achieved with a copy and clean rate greater than zero and an autodelete flag so no cleanup is necessary once the background copy is finished.

- ▶ Backup

Sometimes referred to as *incremental*, a backup FlashCopy mapping consists of a point-in-time full copy of a source volume, plus periodic increments or “deltas” of data that have changed between two points in time. This is a mapping where the copy and clean rates are greater than zero, no autodelete flag is set and the use of an incremental flag to preserve the bitmaps between activations so that only the deltas since the last “backup” need be copied.

The FlashCopy mapping has 4 property attributes (*clean rate*, *copy rate*, *autodelete*, *incremental*) described later in this chapter and 7 different states described later in this chapter as well. The *actions* users can perform on a FlashCopy mapping are:

- ▶ Create: Define a source, a target and set the properties of the mapping
- ▶ Prepare: The system needs to be prepared before a FlashCopy copy starts. It basically flushes the cache and makes it “transparent” for a short time, so no data is lost.
- ▶ Start: The FlashCopy mapping is started and the copy begins immediately. The target volume is immediately accessible.
- ▶ Stop: The FlashCopy mapping is stopped (either by the system or by the user). Depending on the state of the mapping, the target volume is usable or not.
- ▶ Modify: Some properties of the FlashCopy mapping can be modified after creation.
- ▶ Delete: Delete the FlashCopy mapping. This does not delete any of the volumes (source or target) from the mapping.

The source and target volumes must be the same size. The minimum granularity that IBM Spectrum Virtualize supports for FlashCopy is an entire volume. It is not possible to use FlashCopy to copy only part of a volume.

Important: As with any point-in-time copy technology, you are bound by operating system and application requirements for interdependent data and the restriction to an entire volume.

The source and target volumes must belong to the same IBM Spectrum Virtualize system, but they do not have to be in the same I/O group or storage pool.

Volumes that are members of a FlashCopy mapping cannot have their size increased or decreased while they are members of the FlashCopy mapping.

All FlashCopy operations occur on FlashCopy mappings. FlashCopy does not alter the source volumes. However, multiple operations can occur at the same time on multiple FlashCopy mappings because of the use of consistency groups.

5.3.3 Consistency groups

To overcome the issue of dependent writes across volumes and to create a consistent image of the client data, a FlashCopy operation must be performed on multiple volumes as an atomic operation. To accomplish this method, the IBM Spectrum Virtualize supports the concept of consistency groups. Consistency groups address the requirement to preserve point-in-time data consistency across multiple volumes for applications that include related data that spans multiple volumes. For these volumes, consistency groups maintain the integrity of the FlashCopy by ensuring that *dependent writes* are run in the application's intended sequence.

FlashCopy mappings can be part of a consistency group, even if there is only one mapping in the consistency group. If a FlashCopy mapping is not part of any consistency group it is referred as *stand-alone*.

5.3.4 Crash consistent copy and hosts considerations

FlashCopy consistency groups do not provide application consistency. It only ensures volume points-in-time are consistent between them.

Because FlashCopy is at the block level, it is necessary to understand the interaction between your application and the host operating system. From a logical standpoint, it is easiest to think of these objects as "layers" that sit on top of one another. The application is the topmost layer, and beneath it is the operating system layer.

Both of these layers have various levels and methods of caching data to provide better speed. Because the IBM SAN Volume Controller and, therefore, FlashCopy sit below these layers, *they are unaware of the cache at the application or operating system layers*.

To ensure the integrity of the copy that is made, it is necessary to flush the host operating system and application cache for any outstanding reads or writes before the FlashCopy operation is performed. Failing to flush the host operating system and application cache produces what is referred to as a *crash consistent copy*.

The resulting copy requires the same type of recovery procedure, such as log replay and file system checks, that is required following a host crash. FlashCopies that are crash consistent often can be used following file system and application recovery procedures.

Various operating systems and applications provide facilities to stop I/O operations and ensure that all data is flushed from host cache. If these facilities are available, they can be used to prepare for a FlashCopy operation. When this type of facility is unavailable, the host cache must be flushed manually by quiescing the application and unmounting the file system or drives.

The target volumes are overwritten with a complete image of the source volumes. Before the FlashCopy mappings are started, it is important that any data that is held on the host operating system (or application) caches for the target volumes is discarded. The easiest way to ensure that no data is held in these caches is to unmount the target Volumes before the FlashCopy operation starts.

Preferred practice: From a practical standpoint, when you have an application that is backed by a database and you want to make a FlashCopy of that application's data, it is sufficient in most cases to use the write-suspend method that is available in most modern databases, because the database maintains strict control over I/O.

This method is as opposed to flushing data from both the application and the backing database, which is always the suggested method because it is safer. However, this method can be used when facilities do not exist or your environment includes time sensitivity.

IBM Spectrum Protect Snapshot

IBM FlashCopy is not application aware and a third-party tool is needed to link the application to the FlashCopy operations.

IBM Spectrum Protect Snapshot protects data with integrated, application-aware snapshot backup and restore capabilities using FlashCopy technologies in the IBM Spectrum Virtualize.

You can protect data that is stored by IBM DB2 SAP, Oracle, Microsoft Exchange, and Microsoft SQL Server applications. You can create and manage volume-level snapshots for file systems and custom applications.

In addition, it enables you to manage frequent, near-instant, nondisruptive, application-aware backups and restores using integrated application and VMware snapshot technologies. IBM Spectrum Protect Snapshot can be widely used in both IBM and non-IBM storage systems.

For more information on IBM Spectrum Protect Snapshot, see https://www.ibm.com/support/knowledgecenter/en/SSERFV_8.1.0

5.4 Workload relocation into the Public Cloud

In this section, yet another use case for IBM Spectrum Virtualize for Public Cloud is illustrated wherein an entire workload segment is migrated from a client's enterprise into the cloud. While the process for relocating a workload into the cloud via IBM Spectrum Virtualize can certainly simply entail Remote Copy, there are other mechanisms through which this can be accomplished, making this a topic worth discussing.

5.4.1 Business justification

All the drivers that motivate businesses to utilize virtualization technologies makes deploying services into the cloud even more compelling because the cost of idle resources are further absorbed by the cloud provider. However, certain limitations in regulatory or process controls may prevent a business from moving all workloads and application services into the cloud.

An ideal case with regards to a hybrid cloud solution would be the relocation of a specific segment of the environment that is particularly well suited, such as development. Another might be a specific application group that doesn't require either the regulatory isolation or low response time integration with on-premises applications.

While performance may or may not be a factor, it should not be assumed that cloud deployments will automatically provide diminished performance. Depending on the location of the cloud service data center and the intended audience for the migrated service, the performance could conceivably be superior than on-premises pre-migration.

In summary, moving a workload into the cloud may provide similar functionality with better economies due to scaling of physical resources in the cloud provider. Moreover, the cost of services in the cloud are structured and easily measurable and predictable.

5.4.2 Data migration

There are multiple methods for performing data migrations into the cloud just as there are for on-premises migrations. Let us discuss three general approaches:

- ▶ IBM Spectrum Virtualize Remote Copy
- ▶ Host-side mirroring (Storage vMotion or IBM AIX® Logical Volume Manager mirroring)
- ▶ Appliance-based data transfer, such as IBM Aspera® or IBM Transparent Data Migration Facility (TDMF)

The first method has already been discussed in previous sections and is essentially the same process as disaster recovery. The only difference being that instead of a persistent replication, once the initial synchronization is completed, the goal is to schedule the cutover of the application onto the compute nodes in the cloud environment attached to the IBM Spectrum Virtualize storage. This method is likely the preferred method for bare-metal Linux or Microsoft Windows environments.

Host side mirroring would require the server to have concurrent access to both local and remote storage which is not feasible. Also, because the object is to relocate the workload (both compute and storage) into the cloud environment, that will more easily be accomplished by replicating the storage and once synchronized, bringing up the server in the cloud environment and making the appropriate adjustments to the server for use in the cloud.

The second method is largely impractical as it requires the host to be able to access both source and target simultaneously and the practical impediments to creating an iSCSI (the only connection method currently available for IBM Spectrum Virtualize in the Public Cloud) connection from on-premises host systems into the cloud are beyond the scope of this use case. Traditional VMware Storage vMotion is similar to this but again, would require the target storage to be visible via iSCSI to the existing

The third method entails the use of third party software or hardware to move the data from one environment to another. The general idea is that the target system would have an operating system and some empty storage provisioned to it that would act as a landing pad for data that is on the source system. Going into detail about these methods is also outside the scope of this document, but suffice it to say that the process would be no different between an on-premises to cloud migration as it would be to an on-premises to on-premises migration.

VMware environments, however, do have some interesting options that uses either a combination of the first two methods or something similar to the second and the third method.

The first of the two options is a creative migration method that involves setting up a pair (or multiple pairs) of *transit datastores* that remain in a remote copy relationship (see Figure 5-5 on page 133). After these are sync (or if they are set up from scratch, they can be created with the sync flag and then assigned to the VMware clusters) selected VMware guests can be storage vMotioned into these datastores.

After that process is complete, a batch of guests can be scheduled for cutover with Site Recovery Manager. When cut over, then the guest can be storage vMotioned out of the cloud transit datastore into a permanent datastore in the IBM Cloud environment.

For ESX clusters on vSphere 5.1 or higher, there are circumstances under which vMotion without shared storage is possible, with the appropriate licensing. As long as the conditions are met for the two vSphere clusters, a guest can be moved from an ESX host from one cluster to another, and the data will move to a datastore that is visible to the target ESX host. This falls somewhere between the second and the third migration methods as it employs something similar to mirroring but really leverages VMware as a migration appliance. For more information, see the [VMware Documentation](#).

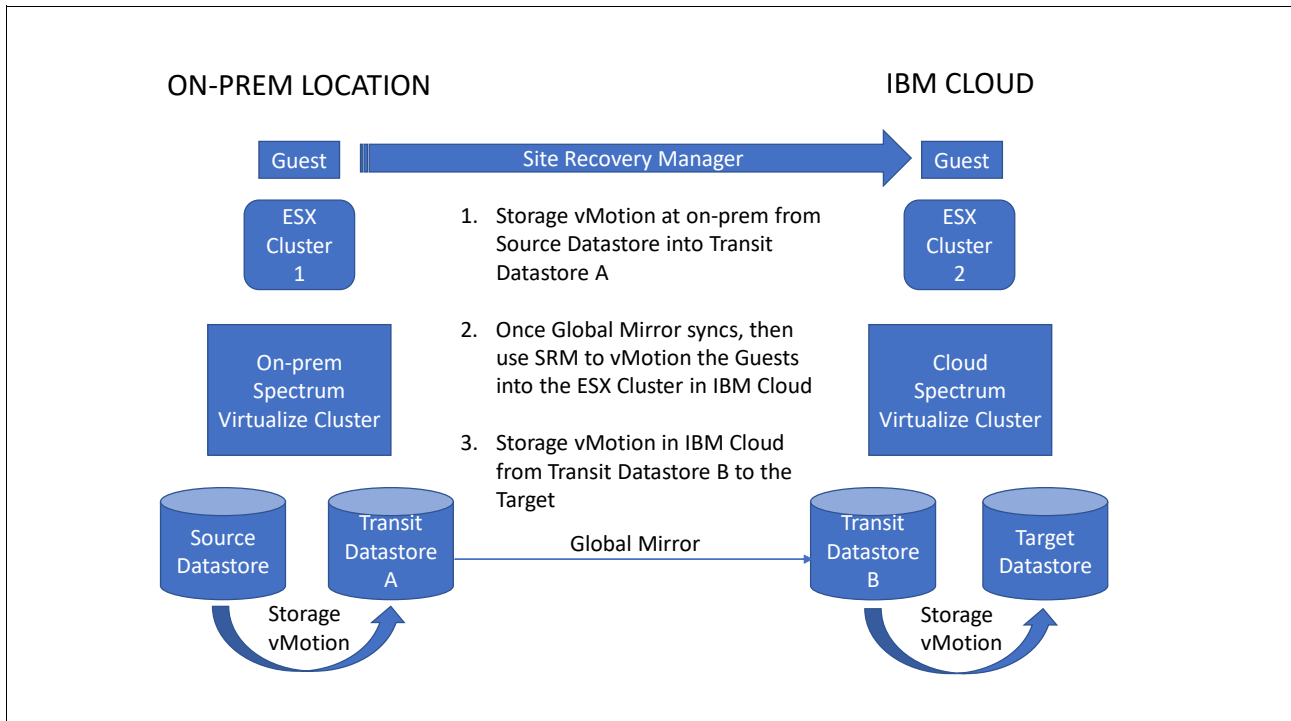


Figure 5-5 On-prem to Cloud VMware migration via “transit datastores”

Table 5-3 lists the migration methods.

Table 5-3 Migration methods

Migration method	Best-suited operating system	Pros versus cons
Remote Copy	Stand alone Windows, Linux, or VMWare (any version)	Simple versus limited scope
Host Mirror	VMWare vSphere 5.1 or higher	Simple versus limited scope
Appliance	N/A	Flexible versus cost and complexity

5.4.3 Host provisioning

In addition to the replication of data, it is necessary for compute nodes and networking to be provisioned within the Cloud provider upon which to run the relocated workload. Currently, in the IBM Cloud, bare-metal and virtual servers are available. Within the bare-metal options, Intel processor based machines and Power8 machines with OpenPOWER provide high performance on Linux-based platforms. As Spectrum Virtualize in the Public Cloud matures and expands into other Cloud Service Providers, other platforms might become available.

5.4.4 Implementation considerations

The following list describes implementation considerations for the workload relocation into the Public Cloud use case:

- ▶ Naming conventions

This is an important consideration in the manageability of a standard on-premises IBM Spectrum Virtualize environment, but given the multiple layers of virtualization in a cloud implementation, maintaining a consistent and meaningful naming convention for all objects (managed disks, volumes, FlashCopy mappings, Remote Copy relationships, hosts and host clusters).

- ▶ Monitoring integration

Integration into IBM Spectrum Control or some other performance monitoring framework will be useful for maintaining metrics for reporting or troubleshooting. IBM Spectrum Control is just particularly well suited for managing IBM Spectrum Virtualize environments.

- ▶ Planning and Scheduling

Regardless of the method chosen, gather as much information ahead of time as possible: Filesystem information, application custodians, full impact analysis of related systems, and so forth.

- ▶ Be sure to ensure solid backout

In the event that inter-related systems or other circumstances require rolling back the application server(s) to on-prem, plan the migration in such a way as to ensure as little difficulty as possible in the roll-back. This may mean keeping zoning in the library (even if not in the active configuration), not destroying source volumes for a certain waiting period.



Supporting the solution

This chapter provides guidance about how support for this solution developed. This solution is consists of two basic support segments: IBM Cloud and IBM Storage support teams. Because of this, it is important to understand who to contact if a problem occurs.

This chapter includes the following topics:

- ▶ 6.1, “Who to contact for support” on page 136
- ▶ 6.2, “Working with IBM Cloud Support” on page 137
- ▶ 6.3, “Working with IBM Spectrum Virtualize Support” on page 138

6.1 Who to contact for support

The IBM Spectrum Virtualize for Public Cloud solution consists of several components, much like the traditional storage offerings. However, when deployed in the public cloud, IBM Spectrum Virtualize is simply an application running in the cloud, as shown in Figure 6-1.




Figure 6-1 IBM Spectrum Virtualize for Public Cloud solution components

In this solution, the cloud provider is responsible for providing the infrastructure, network components and storage, and support and assistance for this portion. The cloud user or any involved third party will be responsible for deploying and configure the solution from the network layer up to the OS and the software installed. IBM Systems support is responsible for providing support and assistance with the IBM Spectrum Virtualize application.

As per current state, the solution consists of multiple parties with different roles and responsibilities. For this reason is good practice, in such cross-functional projects and processes, to clarify roles and responsibilities with a workflow definition for handling tasks and problems when they arise. In this sense a responsibility assignment matrix, also known as *RACI matrix*, describes the participation by various roles in completing tasks or deliverables for a project or business process, splitting as (R) Responsible, (A) Accountable, (C) Consulted and (I) Informed.

The RACI matrix will be specific for each solution deployment: how the cloud service is provided, who is the final user, who are multiple parties involved and so forth. To assist in creating a workflow for handling problems when they arise, we created Table 6-1 as an example.

Table 6-1 Simplified workflow definition based on RACI matrix

Situation	Client	Cloud Provider	Spectrum Virtualize
SV error 2030	Informed	Consulted	Responsible
mdisk is offline	Informed	Responsible	Accountable

Situation	Client	Cloud Provider	Spectrum Virtualize
Network port is down	Informed	Responsible	Consulted
Configuration question	Responsible	Consulted	Accountable

In the situations where the cloud provider is responsible or accountable, the client should collect as much detail about the problem that is known and open a ticket with the cloud provider. In the situations where IBM Spectrum Virtualize is responsible or accountable, the client should collect as much detail about the problem and diagnostic data surrounding the event and raise a PMR with IBM. In the situations where the client is responsible, it is up to the client to be as detailed as possible in any requests or questions raised to the cloud provider or IBM Spectrum Virtualize or any other third party involved in the support.

6.2 Working with IBM Cloud Support

The IBM Cloud IaaS portal (that is, Softlayer) is the system that provisions the infrastructure, network, operating systems, and back end storage that will be used in this solution. As such, IBM Cloud support is responsible for assisting in resolving problems and answering questions for products and services acquired through the IBM Cloud IaaS portal. The IBM Cloud support team is engaged in the ticketing system provided by the IBM Cloud IaaS portal. Figure 6-2 shows where in the portal a ticket can be opened for the IBM Cloud systems.



Figure 6-2 Add a ticket in IBM Cloud IaaS portal

In this page, you can review support documentation, review tickets, and create tickets, as shown in Figure 6-3.

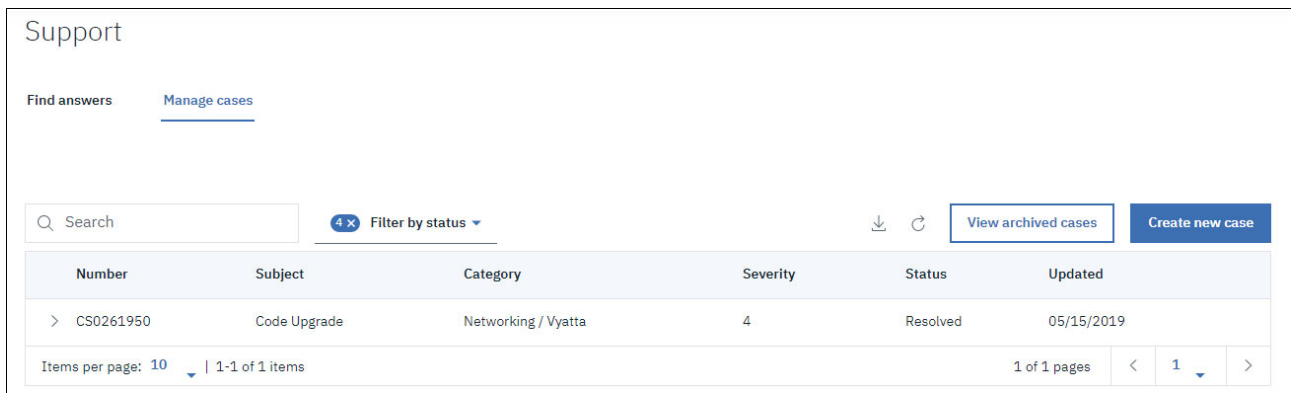


Figure 6-3 Adding a ticket

After a ticket is generated, a representative from IBM Cloud support reviews and updates the ticket. An email is sent to the master account and to all the customer representative accounts that are assigned to the ticket or entitled to receive it, as shown in Figure 6-3. The accuracy of email addresses must be verified in advance for the correct functionality of email notifications.

6.3 Working with IBM Spectrum Virtualize Support

Support engagement for the IBM Spectrum Virtualize for Public Cloud component of the solution is the same as it is for all of the other IBM Spectrum Virtualize based solutions. IBM Support can be engaged by using one of the following methods:

- ▶ Visit the IBM Service requests and PMRs [web page](#) to open a PMR
- ▶ By phone: 1-800-IBM-SERV
- ▶ IBM Call Home

After you receive a Problem Management Record (PMR) or ticket number, you can begin working with support to troubleshoot the problem. You might be asked to collect diagnostic data or to open a remote support session for an IBM Support representative to dial in to the system and investigate.

6.3.1 Email notifications and the Call Home function

The Call Home function of IBM Spectrum Virtualize uses the email notification being sent to the specific IBM Support center, therefore the configuration is similar to sending emails to the specific person or system owner.

Complete the following steps to configure email notifications and emphasizes what is specific to Call Home:

1. Prepare your contact information that you want to use for the email notification and verify the accuracy of the data. From the GUI's left menu, select **Settings** → **Notifications** (see Figure 6-4).

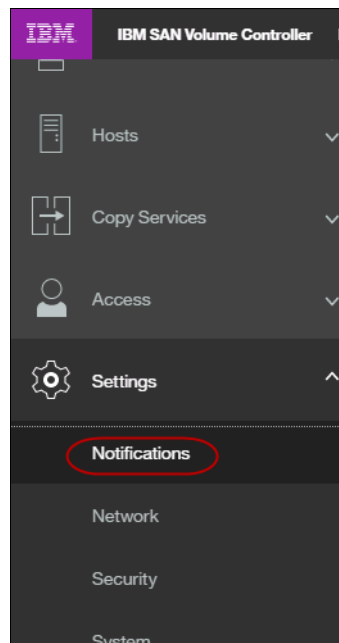


Figure 6-4 Notifications menu

2. Select **Email** and then, click **Enable Notifications** (see Figure 6-5 on page 139).

For the correct functionality of email notifications, ask your network administrator if Simple Mail Transfer Protocol (SMTP) is enabled on the network and is not blocked by firewalls or the foreign destination "@de.ibm.com" is not blocked.

Be sure to test the accessibility to the SMTP server by using the `telnet` command (port 25 for a non-secured connection, port 465 for Secure Sockets Layer (SSL)-encrypted communication) using any server in the same network segment.

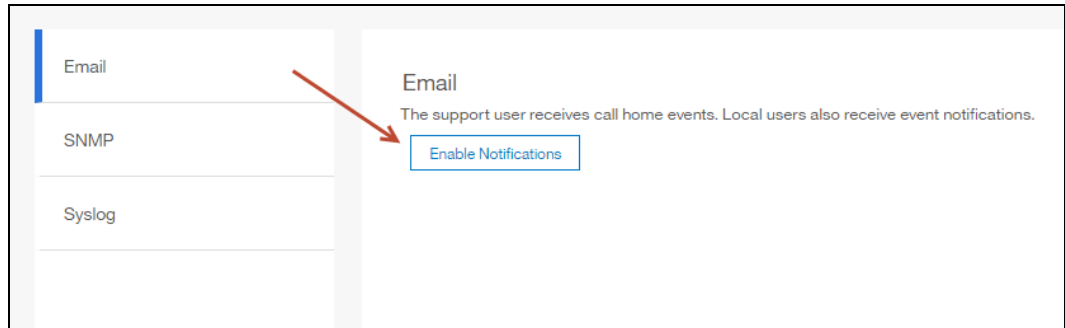


Figure 6-5 Configuration of email notifications

3. After clicking **Next** on the welcome panel, enter the information about the location of the system (see Figure 6-6) and contact information of IBM Spectrum Virtualize administrator (see Figure 6-7 on page 140) to be contactable by IBM Support. *Always* keep this information current.

A screenshot of a dialog box titled 'Email Event Notifications'. The dialog has a blue header bar with a close button (X) in the top right. On the left is a dark sidebar with a list of tabs: 'Welcome' (selected with a green checkmark), 'System Location', 'Contact', 'Email Servers', and 'Summary'. The main area is titled 'System Location' and contains the text 'Service parts should be shipped to the same physical location as the system.' Below this are several input fields: 'Company name:' with 'IBM ITSO', 'System address:' with '120 Holger way', 'City:' with 'San Jose', 'State or province:' with 'CA', 'Postal code:' with '95134', 'Country or region:' with a dropdown menu showing 'United States', and 'Machine location:' with 'Bid#3 2ndFloor Back Room ITSO racks'. At the bottom left is a 'Cancel' button, and at the bottom right are 'Back' and 'Next' buttons.

Figure 6-6 Location of the device

Figure 6-7 shows the contact information of the owner.

Email Event Notifications

✓ Welcome
✓ System Location
○ Contact
Email Servers
Summary

Contact

The support center contacts this person to resolve issues on the system.

Name: System Administrator

Email: name@company.com

Phone (primary): +1 234 567 8900

Phone (alternate):

Cancel Back Apply and Next

Figure 6-7 Contact information

4. Configure the IP address of your company SMTP server, as shown in Figure 6-8. When the correct SMTP server is provided, you can test the connectivity by using the **Ping** option to its IP address. You can configure additional SMTP servers by clicking the **+** at the end of the entry line.

Email Event Notifications

✓ Welcome
✓ System Location
✓ Contact
○ Email Servers
Summary

Email Servers

Call home and event notifications are routed through this email server.

Server IP: 10.18.228.118 Port: 25

Ping

Call home regularly sends emails to the support center that describes your system hardware and critical configuration information. Object names and other potentially sensitive information, such as IP addresses, are not sent.

Cancel Back Apply and Next

Figure 6-8 Configure email servers and inventory reporting

- The summary window opens. Verify it and click **Finish**. You are returned to the Email Settings window where you can verify email addresses of IBM Support (callhome0@de.ibm.com) and optionally add local users who also must receive notifications (see Figure 6-9).

The default support email address callhome0@de.ibm.com is predefined by the system to receive Error Events and Inventory, we recommend not changing these settings.

You can modify or add local users by using **Edit** mode after the initial configuration was saved.

The **Inventory Reporting** function is enabled by default for Call Home. Rather than reporting a problem, an email is sent to IBM that describes your system hardware and critical configuration information. Object names and other information, such as IP addresses, are not included. By default the inventory email is sent on a weekly basis, allowing an IBM Cloud service to analyze and inform you if the hardware or software that you are using requires an update because of any known issues.

Figure 6-9 shows the configured email notification and Call Home settings.

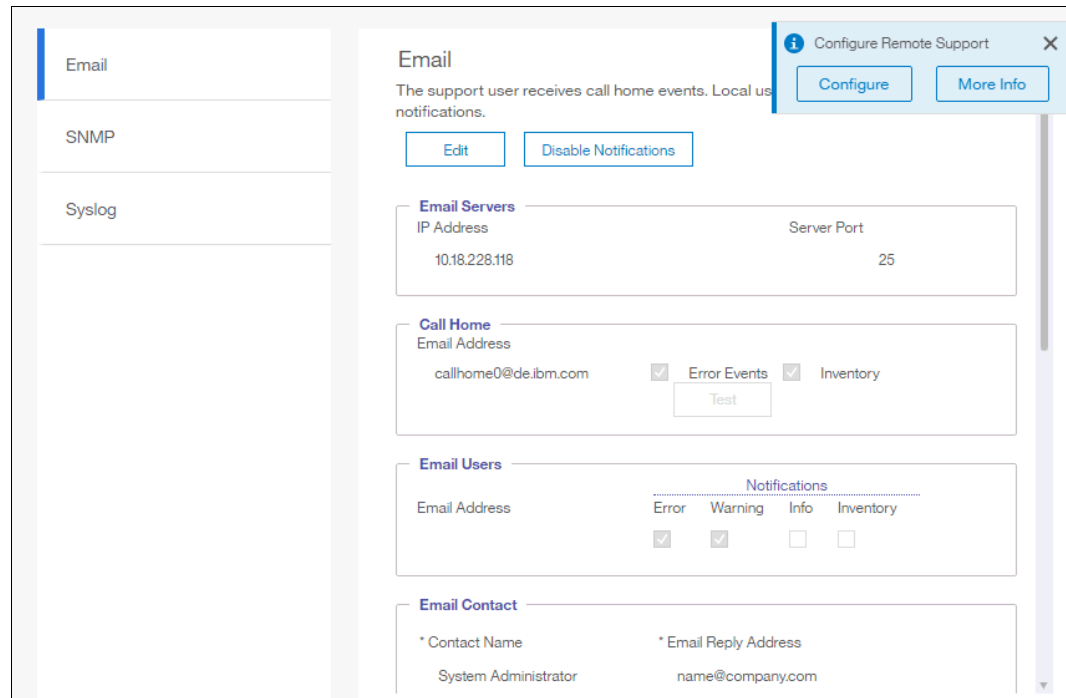


Figure 6-9 Setting email recipients and alert types

- After completing the configuration wizard we can test the email function. To do so, enter **Edit** mode, as shown in Figure 6-10. In the same window, email recipients can be defined or any contact and location information can be changed as needed.




Figure 6-10 Entering edit mode

We strongly suggest that you keep the sending inventory option enabled to IBM Support; however, it might not be of interest to local users, although inventory content can serve as a basis for inventory and asset management.

7. In **Edit** mode, we can change any of the previously configured settings. After these parameters are edited, recipients are added, or the connection is tested, the configuration can be saved so that any changes take effect (see Figure 6-11).

The screenshot shows the 'Email' configuration page. At the top, there are 'Save' and 'Cancel' buttons. Below them is the 'Email Servers' section with fields for 'IP Address' (10.18.228.118) and 'Server Port' (25). The 'Call Home' section shows an 'Email Address' field with 'callhome0@de.ibm.com', checkboxes for 'Error Events' and 'Inventory', and a 'Test' button circled in red. The 'Email Users' section has a table with columns for 'Email Address', 'Error', 'Warning', 'Info', and 'Inventory', and a 'Test' button. The 'Error' and 'Warning' columns have checked boxes, while 'Info' and 'Inventory' have unchecked boxes.

Email Address	Notifications				
	Error	Warning	Info	Inventory	
DaniR@teach.com	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	+ -

Figure 6-11 Saving modified configuration

Note: The Test button appear for new email users after first saving and then edit again.

6.3.2 Disabling and enabling notifications

At any time, you can temporarily or permanently disable email notifications, as shown in Figure 6-12 on page 143. This is good practice when performing activities in your environment which could generate expected errors on your IBM Spectrum Virtualize, such as SAN re-configuration / replacement activities. After the planned activities, remember to re-enable the email notification function. The same results can be achieved with the CLI `svctask stopmail` and `svctask startmail` commands.

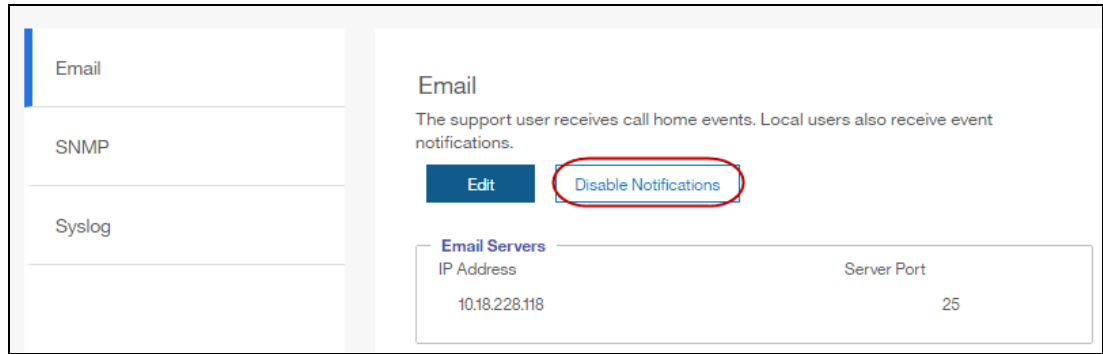


Figure 6-12 Disabling or enabling email notifications

6.3.3 Collecting Diagnostic Data for IBM Spectrum Virtualize

Occasionally, if a problem occurs and the IBM Support Center is contacted, they most likely ask you to provide the support package. You can collect and upload this package from the **Settings** → **Support** menu.

Collecting information by using the GUI

To collect information by using the GUI, complete the following steps:

1. Click **Settings** → **Support** and the **Support Package** tab (see Figure 6-13). Then, click the **Upload Support Package** button.

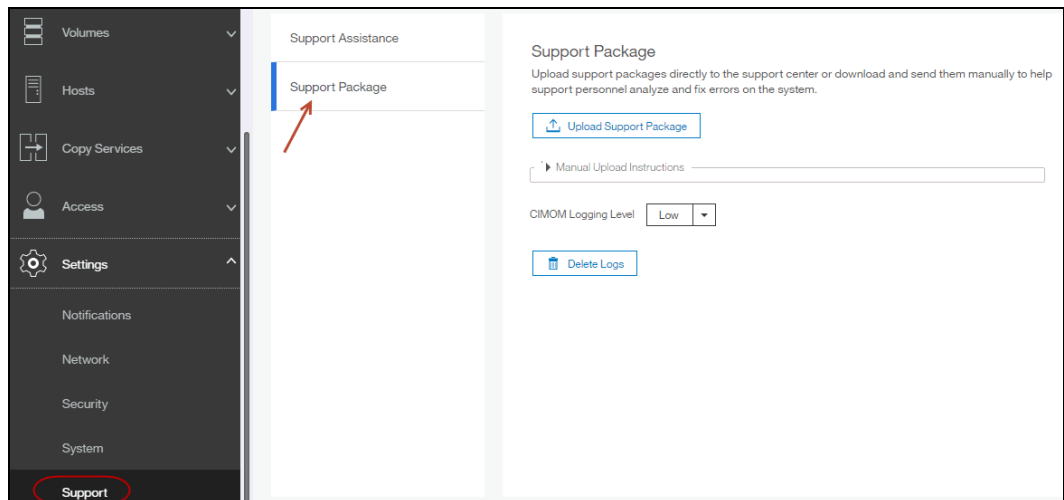


Figure 6-13 Support Package option

Assuming the problem encountered was an unexpected node restart that has logged a 2030 error, we collect the default logs plus the most recent statesave from each node to capture the most relevant data for support.

Note: When a node unexpectedly restarts, it first dumps its current statesave information before it restarts to recover from an error condition. This statesave is critical for support to analyze what occurred. Collecting a snap type 4 creates statesaves at the time of the collection, which is not useful for understanding the restart event.

- From the Upload Support Package panel, we are given four options of data collection. Because you were contacted by IBM Support because your system called home or you manually opened a call with IBM Support, you receive a *PMR number*. Enter that PMR number into PMR field and select the snap type, which is often referred to as an *option 1, 2, 3, 4 snap*, as requested by IBM Support (see Figure 6-14). In our case, we enter our PMR number, select snap type 3 (option 3), because this choice automatically collects the statesave created at the time the node restarted. Click **Upload**.

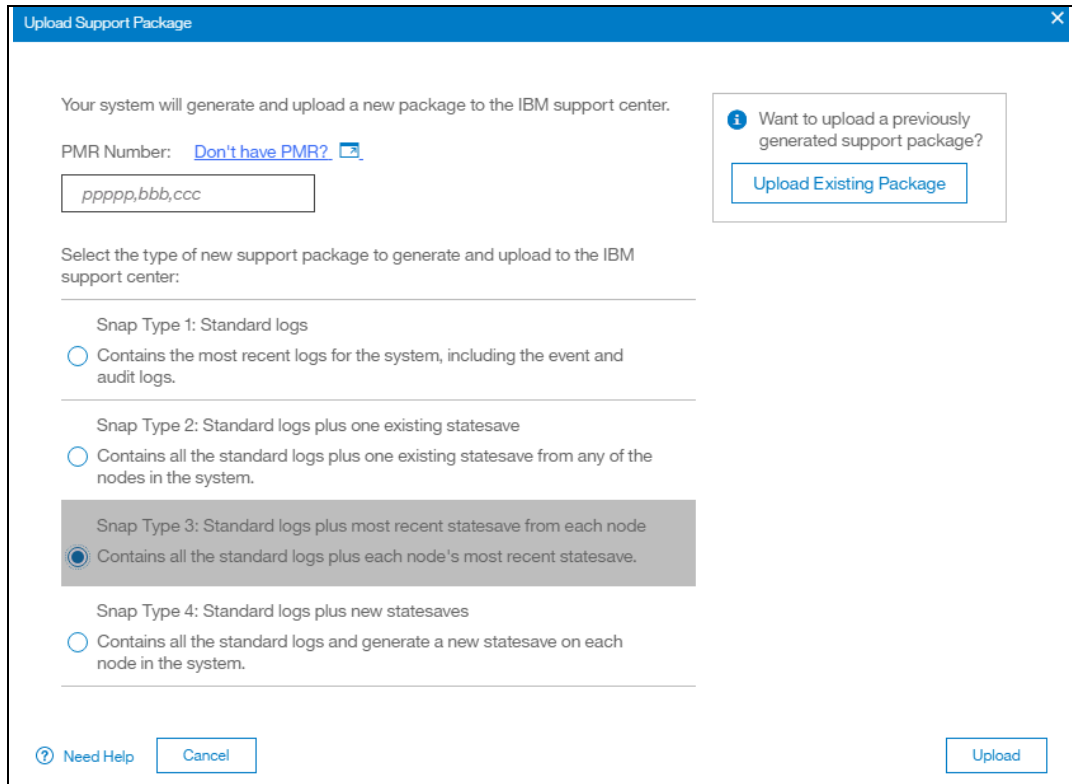


Figure 6-14 Upload Support Package window

The procedure to create the snap on an IBM Spectrum Virtualize system, including the latest statesave from each node, starts. This process might take a few minutes (see Figure 6-15).

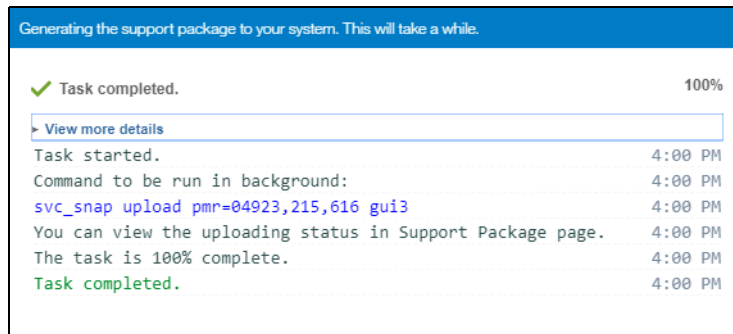


Figure 6-15 Task detail window

Collecting logs by using the CLI

Complete the following steps to use the CLI to collect and upload a support package as requested by IBM Support:

1. Log in to the CLI and to run the `svc_snap` command that matches the type of snap requested by IBM Support:
 - Standard logs (type 1):
`svc_snap upload pmr=ppppp,bbb,ccc gui1`
 - Standard logs plus one existing statesave (type 2):
`svc_snap upload pmr=ppppp,bbb,ccc gui2`
 - Standard logs plus most recent statesave from each node (type 3):
`svc_snap upload pmr=ppppp,bbb,ccc gui3`
 - Standard logs plus new statesaves:
`svc_livedump -nodes all -yes`
`svc_snap upload pmr=ppppp,bbb,ccc gui3`
2. We collect the type 3 (option 3) and have it automatically uploaded to the PMR number that is provided by IBM Support, as shown in Example 6-1.

Example 6-1 The `svc_snap` command

```
ssh superuser@10.18.228.64
Password:
IBM_2145:ITS0 DH8_B:superuser>>svc_snap upload pmr=04923,215,616 gui3
```

3. If you do not want to automatically upload the snap to IBM, do not specify the ‘upload pmr=ppppp,bbb,ccc’ part of the commands. In this case, when the snap creation completes, it creates a file named in the following format:
`/dumps/snap.<panel_id>.YYMMDD.hhmmss.tgz`
It takes a few minutes for the snap file to complete (longer if statesaves are included).
4. The generated file can then be retrieved from the GUI under the **Settings** → **Support** → **Manual Upload Instructions** twisty → **Download Support Package**. Click **Download Existing Package**, as shown in Figure 6-16.

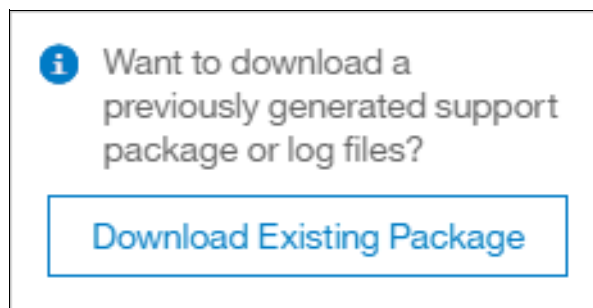


Figure 6-16 Downloaded Existing Package

5. A new panel opens. Click in the Filter box and enter snap (hit Enter). A list of snap files is shown (see Figure 6-17). Locate the exact name of the snap that was generated by the `svc_snap` command that was issued earlier. Click to select that file and then, click **Download**.

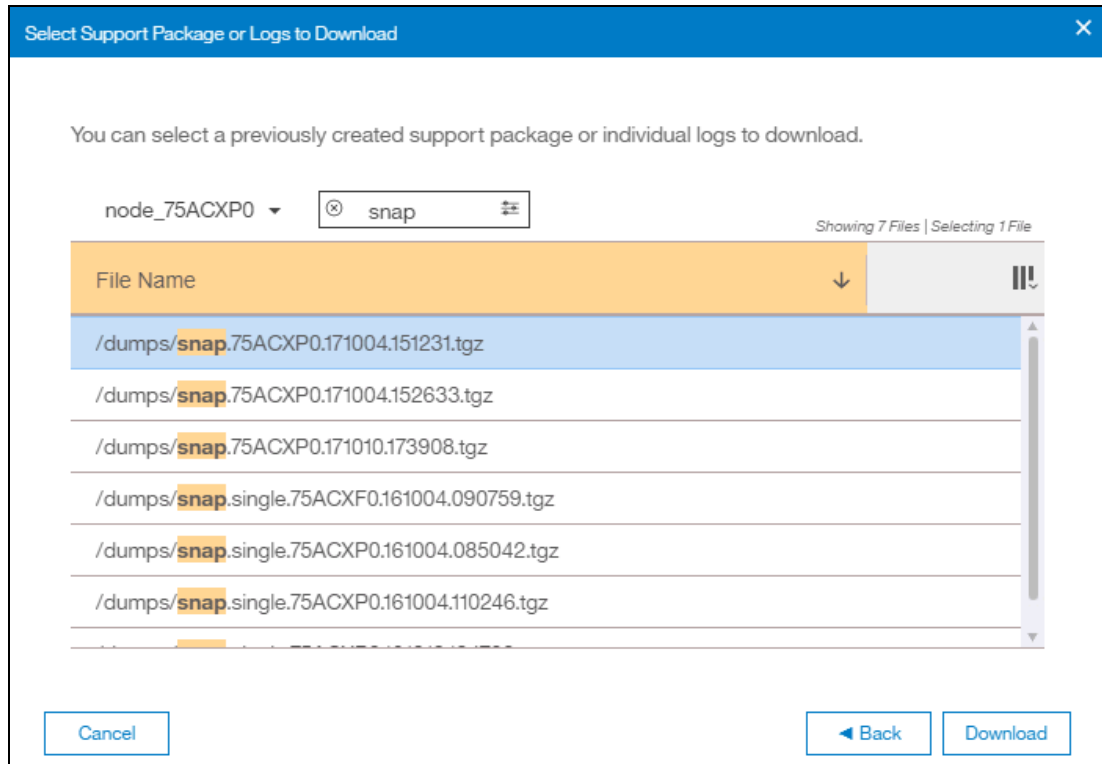


Figure 6-17 Filtering on snap to download

6. Save the file to a folder of your choice on your workstation.

6.3.4 Uploading files to the Support Center

If you chose not to have IBM Spectrum Virtualize upload the support package automatically, it might still be uploaded for analysis by using the Enhanced Customer Data Repository (ECuRep). Any uploads are associated with a specific problem management report (PMR). The PMR is also known as a *service request* and is a mandatory requirement when uploading.

To upload information, complete the following steps:

1. Using a browser, navigate to [the ECuRep Secure Upload web page](#) (see Figure 6-18).

Figure 6-18 ECuRep details

2. Complete the following required fields:

- PMR number (mandatory) as provided by IBM Support for your specific case. This should be in the format of ppppp,bbb,ccc, for example, 04923,215,616 using a comma (,) as a separator.
- Upload is for (mandatory). Select **Hardware** from the drop-down menu.

Although completing the Email address field is not required, we suggest entering your email address to be automatically notified of a successful or unsuccessful upload.

3. When completed, click **Continue**. The Input window opens (see Figure 6-19).

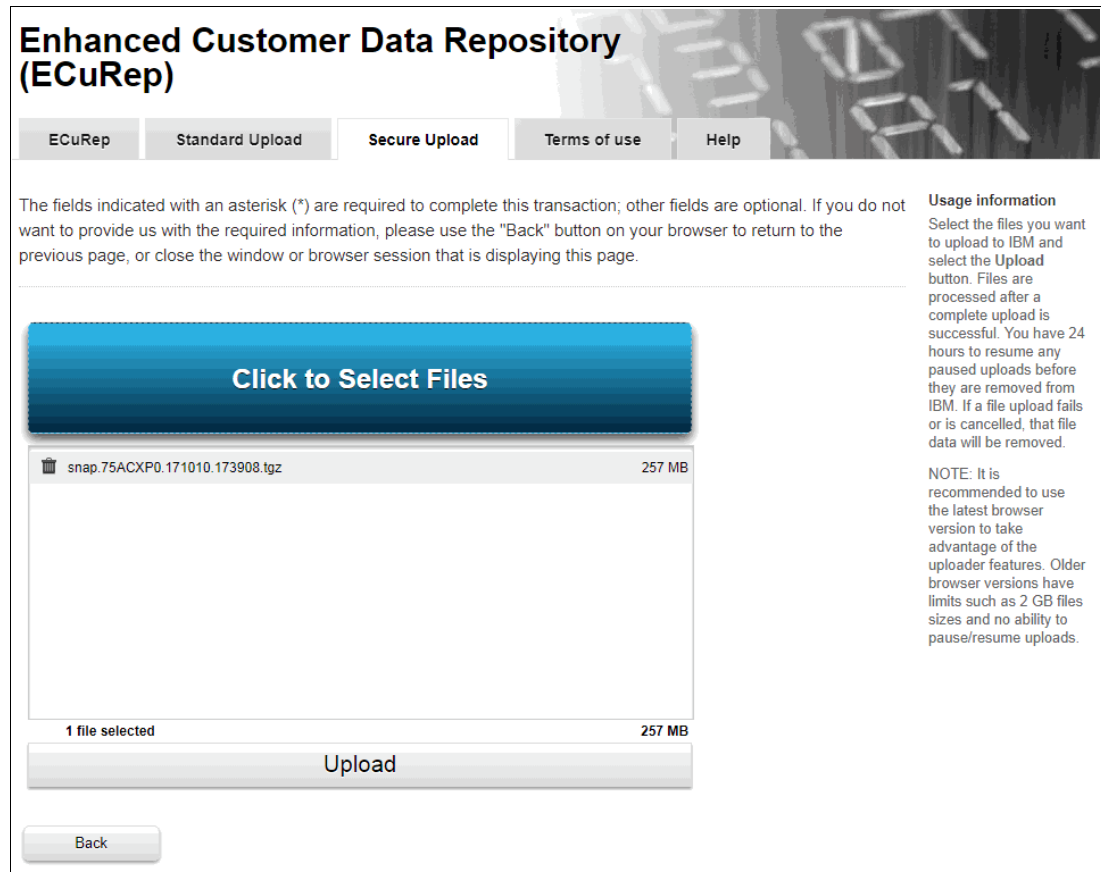


Figure 6-19 ECuRep File upload

4. After the files are selected, click **Upload** to continue, and follow the directions.

6.3.5 Service Assistant Tool

The Service Assistant Tool (SAT) is a web-based GUI that is used to service individual node canisters, primarily when a node has a fault and is in a service state. A node is not an active part of a clustered system while it is in service state.

Typically, an IBM Spectrum Virtualize cluster is initially configured with the following IP addresses:

- ▶ One service IP address for each IBM node.
- ▶ One cluster management IP address, which is set when the cluster is created.

The SAT is available even when the management GUI is not accessible. The following information and tasks can be accomplished with the Service Assistance Tool:

- ▶ Status information about the connections and the nodes
- ▶ Basic configuration information, such as configuring IP addresses
- ▶ Service tasks, such as restarting the Common Information Model (CIM) object manager (CIMOM) and updating the WWNN

- ▶ Details about node error codes
- ▶ Details about the hardware such as IP address and Media Access Control (MAC) addresses.

The SAT GUI is available by using a service assistant IP address that is configured on each node. It can also be accessed through the cluster IP addresses by appending /service to the cluster management IP.

If the clustered system is down, the only method of communicating with the nodes is through the SAT IP address directly. Each node can have a single service IP address on Ethernet port 1 and should be configured for all nodes of the cluster, including any Hot Spare Nodes.

To open the SAT GUI, enter one of the following URLs into any web browser:

- ▶ `http(s)://<cluster IP address of your cluster>/service`
- ▶ `http(s)://<service IP address of a node>/service`

Complete the following steps to access the SAT:

1. When you are accessing SAT by using `<cluster IP address>/service`, the configuration node canister SAT GUI login window opens. Enter the Superuser Password, as shown in Figure 6-20.

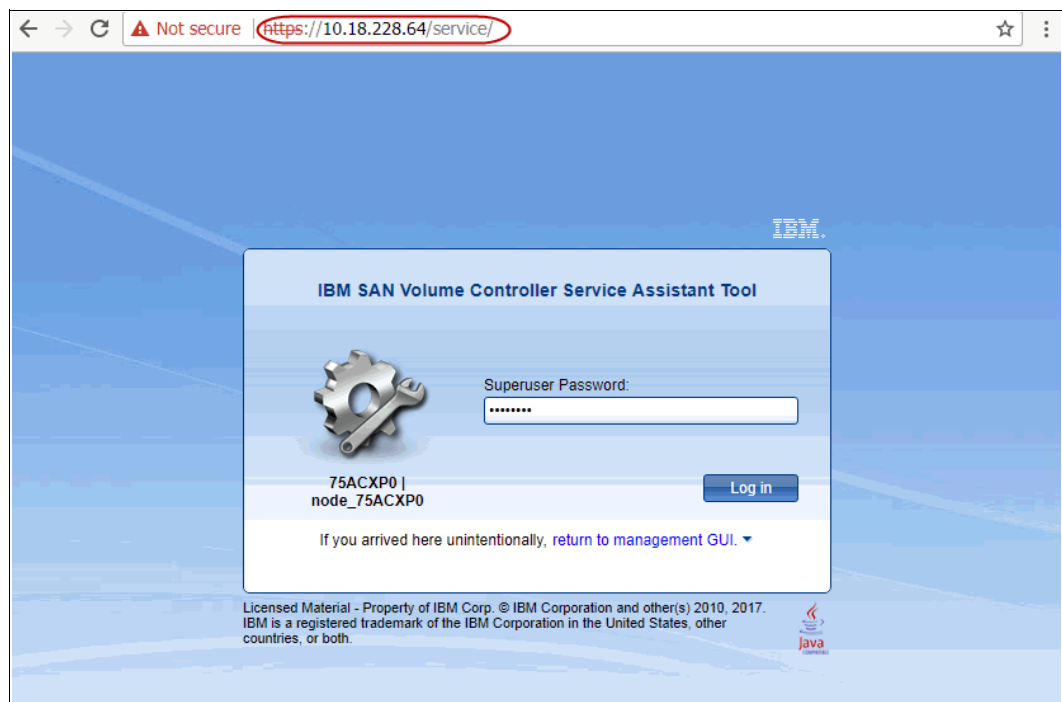


Figure 6-20 Service Assistant Tool Login GUI

- After you are logged in, you see the Service Assistant Home panel, as shown in Figure 6-21. The SAT can view the status and run service actions on other nodes, in addition to the node that the user is logged into.



Figure 6-21 Service Assistant Tool GUI

- The current selected Spectrum Virtualize node is displayed in the upper left corner of the GUI. In Figure 6-21, this is node ID 1. Select the desired node in the Change Node section of the window. You see the details in the upper left change to reflect the selected node.

Note: The SAT GUI provides access to service procedures and shows the status of the nodes. It is advised that these procedures should only be carried out if directed to do so by IBM Support.

For more information about how to use the SA Tool, see [this website](#).

6.3.6 Remote Support Assistance

Introduced with V8.1, Remote Support Assistance allows IBM Support to remotely connect to the Spectrum Virtualize via a secure tunnel to perform analysis, log collection or software updates. The tunnel can be enabled ad-hoc by the client or enable a permanent connection if desired.

Note: Clients who have purchased Enterprise Class Support (ECS) are entitled to IBM Support using Remote Support Assistance to quickly connect and diagnose problems. However, because IBM Support might choose to use this feature on non-ECS systems at their discretion, we recommend configuring and testing the connection on all systems.

If you are enabling Remote Support Assistance, then ensure that the following prerequisites are met:

- ▶ Ensure that call home is configured with a valid email server.
- ▶ Ensure that a valid service IP address is configured on each node on the IBM Spectrum Virtualize.
- ▶ If your IBM Spectrum Virtualize is behind a firewall or if you want to route traffic from multiple storage systems to the same place, you must configure a Remote Support Proxy server. Before you configure remote support assistance, the proxy server must be installed and configured separately. During the set-up for support assistance, specify the IP address and the port number for the proxy server on the remote support centers panel.
- ▶ If you do not have firewall restrictions and the IBM Spectrum Virtualize nodes are directly connected to the Internet, request your network administrator to allow connections to 129.33.206.139 and 204.146.30.139 on Port 22.
- ▶ Both uploading support packages and downloading software require direct connections to the Internet. A DNS server must be defined on your IBM Spectrum Virtualize for both of these functions to work.
- ▶ To ensure that support packages are uploaded correctly, configure the firewall to allow connections to the following IP addresses on port 443: 129.42.56.189, 129.42.54.189, and 129.42.60.189.
- ▶ To ensure that software is downloaded correctly, configure the firewall to allow connections to the following IP addresses on port 22: 170.225.15.105, 170.225.15.104, 170.225.15.107, 129.35.224.105, 129.35.224.104, and 129.35.224.107.

Figure 6-22 shows a pop-up that appears in the GUI after updating to V8.1, to prompt you to configure your IBM Spectrum Virtualize for Remote Support, you might select not to enable it, open a tunnel when needed or to open a permanent tunnel to IBM.




Figure 6-22 pop-up prompt to configure Remote Support Assistance

From the pop-up prompt, we can choose to configure or learn some more about the feature or close the pop-up by clicking the X. Figure 6-23 shows how we can find the Setup Remote Support Assistance if the pop-up is closed.

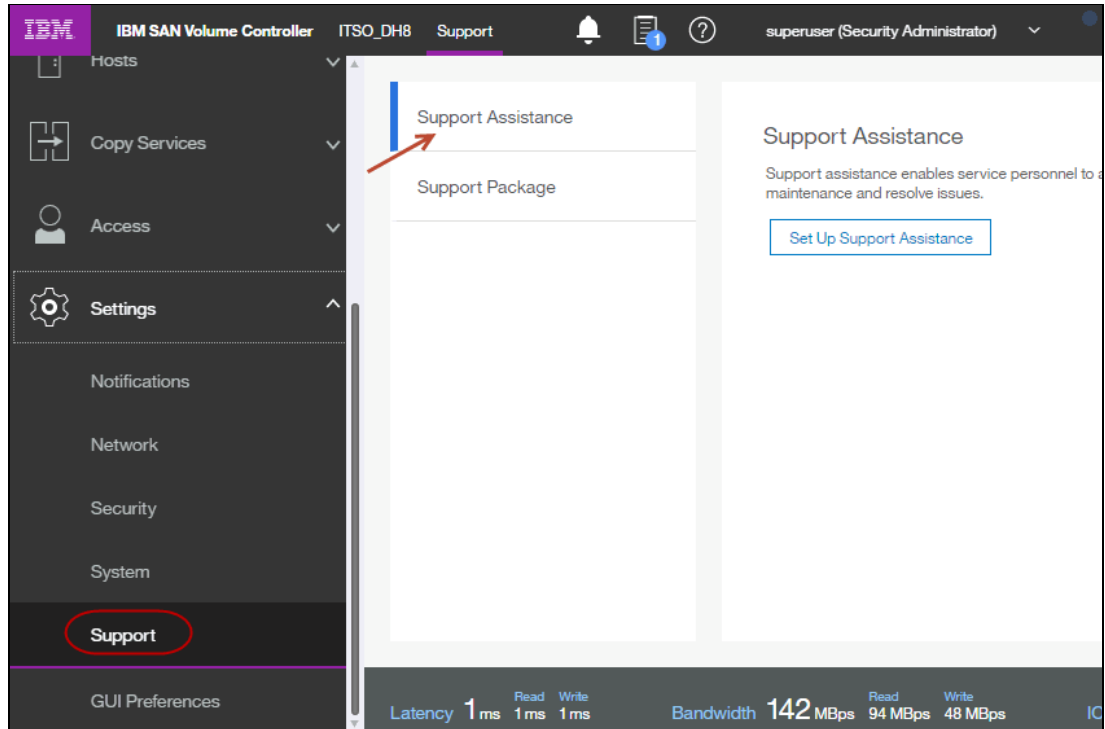


Figure 6-23 Remote Support Assistance menu

Choosing to set up support assistance opens a wizard to guide us through the configuration. Figure 6-24 on page 153 shows the first wizard panel, where we can choose not to enable remote assistance by selecting **I want support personnel to work on-site only** or enable remote assistance by choosing **I want support personnel to access my system both on-site and remotely**. We chose to enable remote assistance and click **Next**.

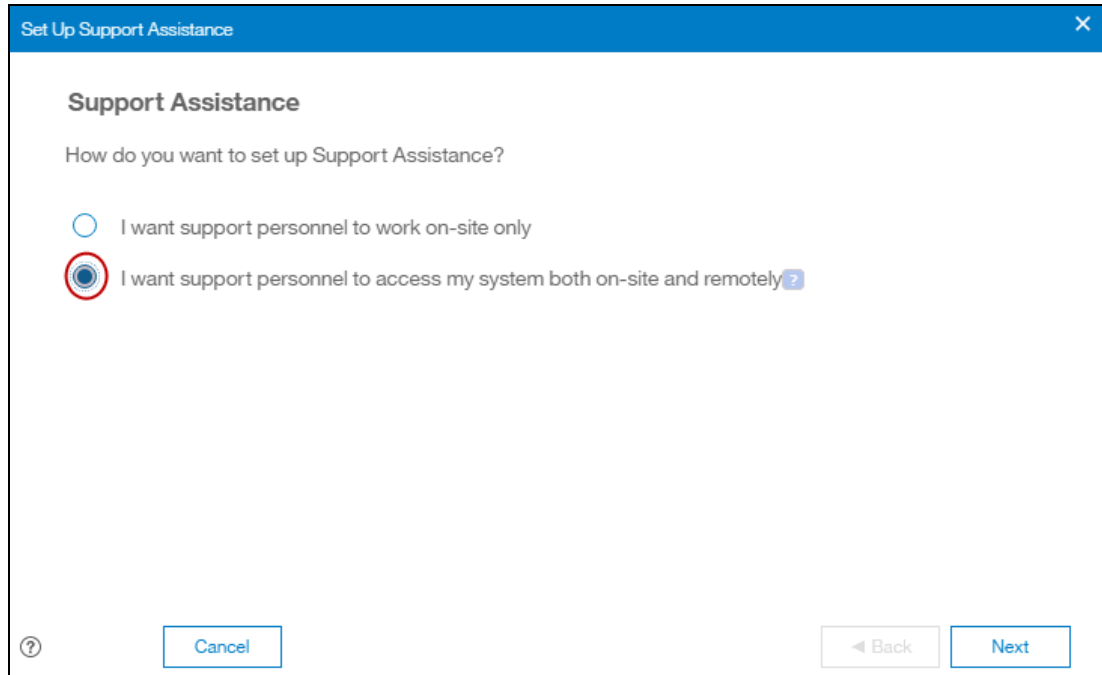


Figure 6-24 Remote Support wizard enable or disable

The next panel, shown in Figure 6-25, lists the IBM Support centers IP addresses and SSH port that will need to be open in your firewall, here we can also define a Remote Support Assistance Proxy if we have multiple IBM Spectrum Virtualize clusters in the same cloud, enabling firewall configuration only being required for the Proxy Server and not every storage system. We do not have a proxy server and leave the field blank. Click **Next**.

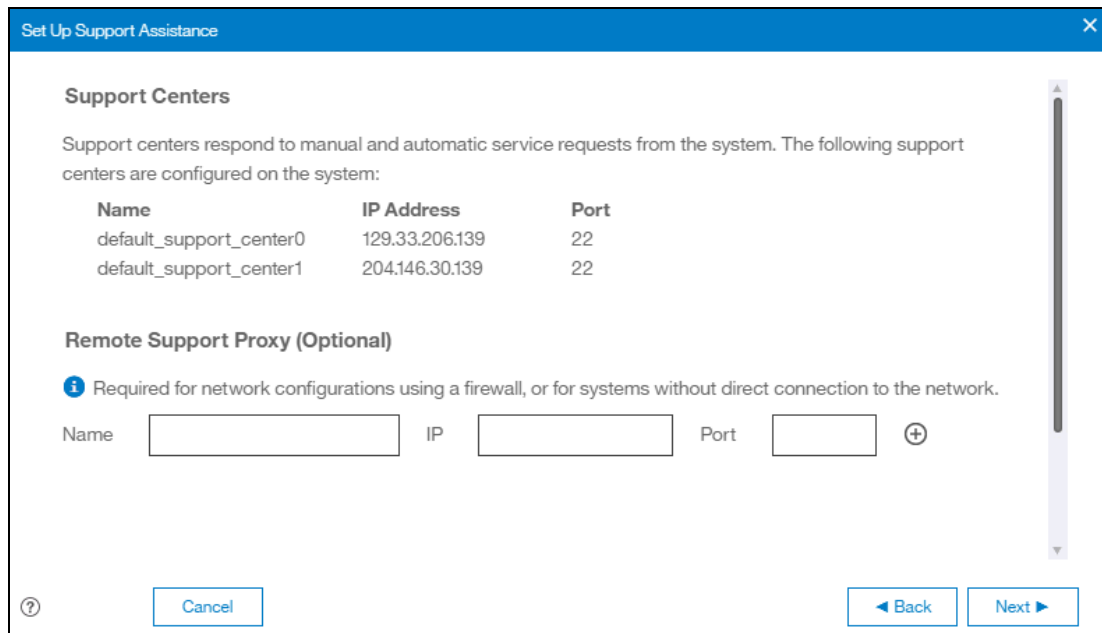


Figure 6-25 Remote Support wizard proxy setup

The next panel asks if we want to open a tunnel to IBM permanently, allowing IBM to connect to your IBM Spectrum Virtualize cluster at any time, or **On Permission Only** which requires a storage administrator to log on to the GUI and enable the tunnel when required. We select this **Permission** option as shown in Figure 6-26 and click **Finish**.

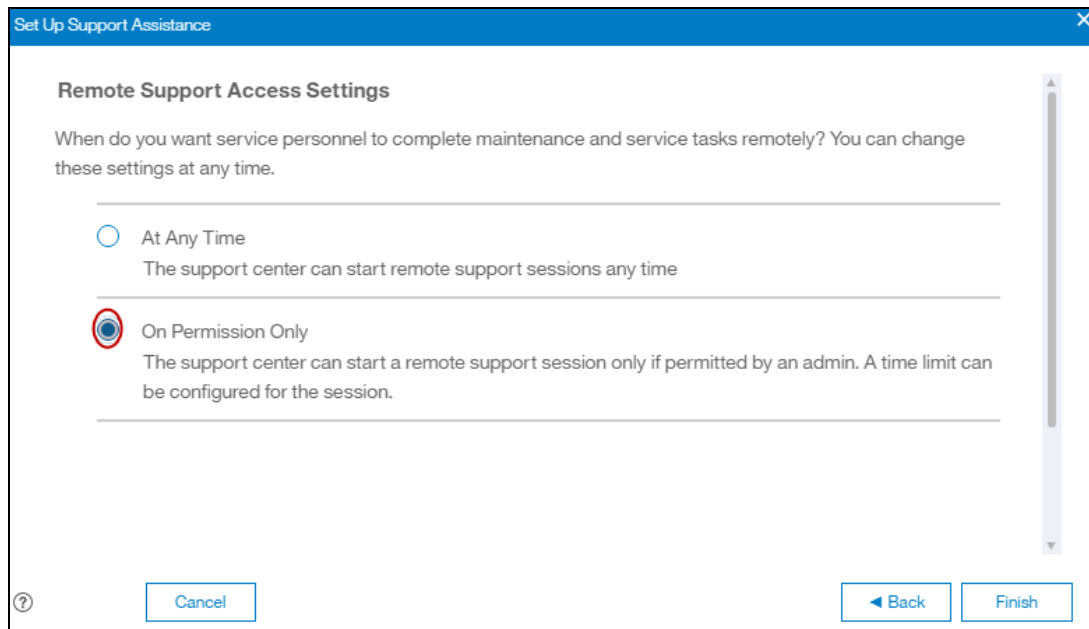


Figure 6-26 Remote Support wizard access choice

When we have completed the remote support set-up, we can view the current status of any remote connection, start a new session, test the connection to IBM, and reconfigure the setup. As shown in Figure 6-27, we successfully tested the connection. Now, click **Start New Session** to open a tunnel for IBM Support to connect.

The screenshot displays the 'Support Assistance' interface. At the top, it states: 'Support assistance enables service personnel to access the system to perform maintenance and resolve issues.' Below this, there are three main sections:

- Local Support: Session Ready**
IBM Service Status: **Not Connected**
- Remote Support**
Contains a 'Start New Session' button and a 'Test Connection' button. The 'Test Connection' button is highlighted in dark blue, and next to it, it says 'Last test: **Successful Connection** on 10/9/17 11:35:50 PM'.
- Support Users and History**
Monitor Users: **0**
Privileged Users: **0**

Below these sections, there is a 'Token Age: 0 days' label and a 'Generate New Token' button. Further down, it shows 'Remote Support Access Setting: **On Permission Only**' and 'Proxy Servers Configured: **0**'. At the bottom, there is a 'Reconfigure Settings' button.

Figure 6-27 Remote Support Status and session management

A pop-up asks how long we would like the tunnel to remain open if there is no activity, by setting a time-out value. Then, as shown in Figure 6-28, the connection establishes and is waiting for IBM Support to connect.

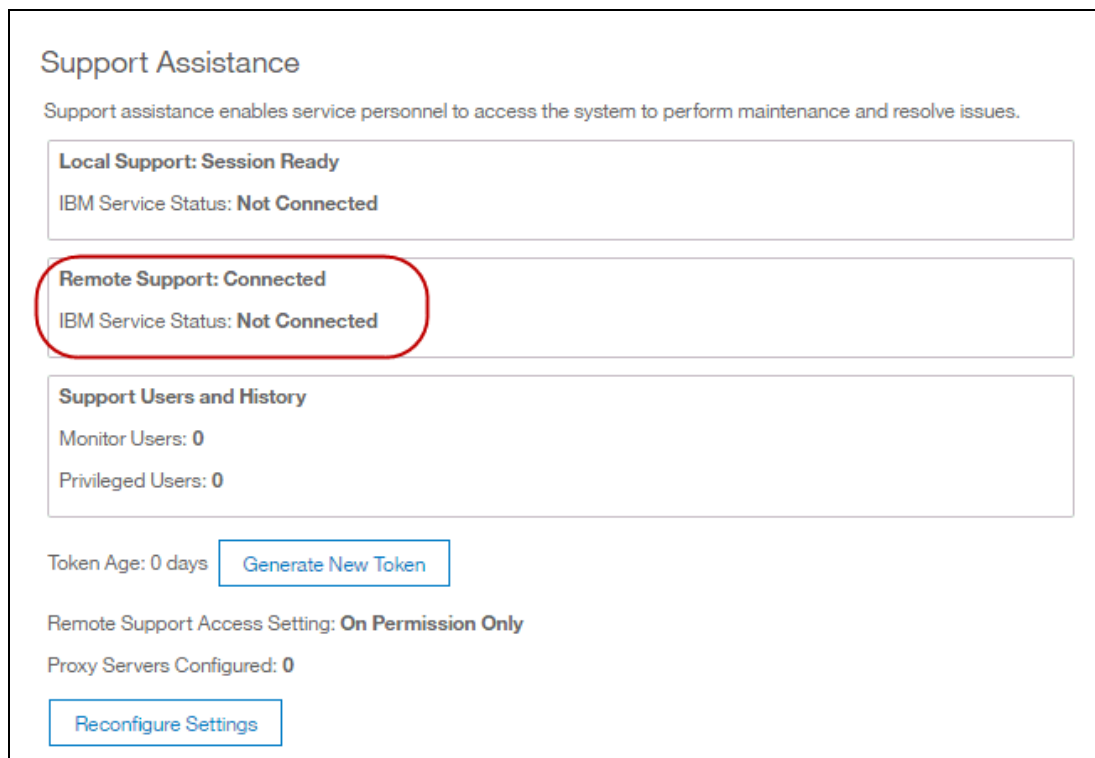


Figure 6-28 Remote Assistance tunnel connected



A

Guidelines for disaster recovery solution in the Public Cloud

In this appendix, we briefly describe the recommended practices when designing a resiliency solution and considerations when using the cloud space as a possible alternative site.

This appendix includes the following topics:

- ▶ “Plan and design for the worst case scenario” on page 158
- ▶ “Recovery tiers” on page 158
- ▶ “Common pitfalls” on page 165
- ▶ “Networking aspects” on page 166
- ▶ “Full versus partial failover” on page 168
- ▶ “Network virtualization” on page 170

Plan and design for the worst case scenario

When designing a resiliency solution the key is in the *resiliency* word. Your solution has to work in the worst possible scenario and provide enough performance, not only to achieve your recovery objectives, but also to be able to sustain and run your production from the alternate site for an extended period of time as in the case of an actual disaster that destroys the primary site.

Your design should not only consider a single plan, but several possible alternatives to better adapt and be able to work in a degraded or impacted ecosystem.

Of course, costs are always an important consideration but nevertheless you should have a clear view of the *minimal* acceptable conditions and the other more optional elements that may be eliminated when cost is a challenge, but with a full appreciation for the consequences of the omission of those elements.

Your design should not be limited to the IT only. Your IT is dependant on the following factors that can be impacted as well:

- ▶ Key personnel availability
- ▶ External networks services
- ▶ Dependencies on critical providers
- ▶ Road conditions (when for example planning to physical transfers of personnel, backup)
- ▶ Disaster recovery (DR) resources availability when required

Recovery tiers

The DR solution can have *recovery tiers* with a different set of appropriate *Recovery Time Objective (RTO)* and *Recovery Point Objective (RPO)* requirements for each wave or tier.

The following is an example to better explain this tiered approach:

- ▶ P0 Class: With an immediate restart or take-over requirement (< 2 hours), which requires active resources at the DR site kept aligned with an application-aware data replication to guarantee a possible take-over of the operations.

Note: P0 class is accomplished with synchronous replication or using a federated application like Active Directory, Oracle Real Application Clusters (RAC), VMWare MetroStorage Cluster, or Live Partition Mobility (LPM). It may be acceptable and might lend itself to tape restoration instead of storage replication.

- ▶ P1 Class: With a near-immediate restart (4-12 hours), which can be implemented with dedicated stand-by resources at the DR site associated to a technological data replication. Although this RTO window might appear to be too wide, you have to consider that your restart will be in an emergency and in this condition, the best you can achieve is the equivalent of restarting from power failure at your on-premises. P1 might be asynchronous replication with up to a 5 minute possible delay (GMCV with 150 sec cycle period).
- ▶ P2 Class: With a restart within one day, you could leverage shared or re-usable assets at the alternate site. Since the time is short, usually this spare capacity must be already there and cannot be acquired at the time of the disaster (ATOD).
- ▶ P3 Class: With a restart after two days, you can use additional compute power that can be freed-up or provisioned on-demand at the alternate site.

Note: Depending on business need, P2 and P3 classes could have either the same RPO or longer RPOs depending on business process or needs. For instance, if this is back office documentation that is only refreshed on a quarterly basis, or web servers, then an RPO of a day or a week might be acceptable and might lend itself to tape restoration instead of storage replication.

The previous list is just a general definition; you might have more tiers or different combinations of requirements.

Whatever the recovery tier is, four basic things are essential to provide a usable DR solution:

- ▶ A copy of your production data in the DR site according to the RPO requirement.
- ▶ Alternate IT resources (compute) to restart the IT operation in emergency according to the RTO requirement.
- ▶ An alternate connection to your production network (WAN).
- ▶ Ability to operate and run production at the alternate DR location.

Design for production

One of the cost elements impacting a resiliency solution is the fact that the alternate site has to run the production in emergency for an extended period of time.

The bigger the emergency, the more services (hence resources) you need to operate during the crisis.

Compute power could be acquired on-demand leveraging the cloud service provider's (CSP) ability to scale the infrastructure.

On the storage side, you can opt to replicate and run the DR test simulation with low performance and using low cost storage, but that performance might not be enough to sustain actual production. Thus, you might need to move your production data over a different tier of storage. You have to evaluate what this means for your cloud provider. This means having a clear picture of the effort in migrating the data to the new tier and the associated monetary and time (which might impact your RTO).

Concurrent on-demand DR provisioning requests

We have seen how components could be requested on-demand in the CSP space, but how might this be impacted by the concurrent requests the CSP might have from other clients in your same situation, as in the case of a metro area disaster? This could exhaust its local scalability in that site or region.

That is, cloud does not have an infinite scale.

Note: Although conceptually, a cloud infrastructure is designed to be able to scale infinitely through multiple sites, a single given site is finite and an individual client's data isn't going to be instantaneously replicated to all sites. Therefore, the limitations of the specific site chosen for DR must be considered.

Another thing to remember is that cloud is a different business than Disaster Recovery (DR).

Most DR providers actively use resource syndication to contain the costs. Resource syndication means the same resource is used to provide service to different customers that might not be exposed to the same concurrent event.

In other words, the compute resource that you use for your DR solution might also be used by other customers of the DR provider, not sharing the same risk situation with you.

DR providers have consolidated this in years of experience in planning and running the DR business, so their site planning is such that it offers a *reasonable* guarantee that their customers in the same risk zone will have their contract honored with the availability of the agreed upon resources.

Cloud providers do not syndicate, because this way of operating is not within their business model. *CSP planning algorithms do not take disaster recovery concurrent requests into consideration when performing a planning for a site*, thus in case of concurrent on-demand provisioning requests, they will provide resources simply based on the timeline of requests up to when they have exhausted the resources in that specific site. At this point, the CSP cannot fulfill requests at that particular site in a reasonable time.

Tip: Be careful to search for CSP SLA on provisioning and read all the caveats.

DR test

Disaster recovery testing is what you have to do to ensure your ability to be resilient, hoping not to have to do it in real life.

There are two types of DR testing you can execute to verify your resiliency capability:

- ▶ DR simulation
- ▶ Switch-over

DR simulation

DR simulations are mainly done to verify and audit the emergency runbooks and check the RTO and RPOs provided by the in an environment that, as closely as possible, resembles a real emergency.

This means introducing disruptions on the replication network connections before interrupting the communications among the sites, simulating the sudden lost of the primary site.

You can only do this if your data replication solution is resilient, and such a disruption does not have a negative impact on the production, which would continue running on the primary site. For instance, if you suddenly disrupted the partnership in a IBM Spectrum Virtualize Metro Mirror relationship, any write activity to your master volumes would suspend. Similarly, if you simulated a site failure with Enhanced Stretch Cluster by isolating the primary site from the secondary site *and* the quorum, the primary site would go offline, by design, until connectivity is restored. So, in IBM Spectrum Virtualize, the only likely scenario appropriate for this test would be Global Mirror with Change Volumes (GMCV). Refer to section “Plan and design for the worst case scenario” on page 158 for more information.

DR simulation deploys a duplication of your production environment at the DR site to used to perform validation. This environment is cleaned at the end of the simulation and changes to the DR test environment are discarded, as the real production has continued on the primary site.

Thus, it is important that network streams flowing to the DR test environment are copies of the real production. Usually these network flows are intercepted and duplicated at the primary site (that receives the real flow), and sent to the DR site which operates on a separate network from production.

Switch-over

Opposite to the DR simulation, in the switch-over test, production is moved from the primary site to the DR site to verify and audit the ability to run and sustain production operations for an extended period.

In this option we do not test the DR runbooks, as this would imply the emergency restart, and would have an impact on the production.

Similarly, because this move impacts the production, you should be careful in performing this switch to minimize all those impacts through scheduling, communication and planning.

Production operations are shut down at the primary site and started up at the DR site. Replication is reversed once the switchover is complete and production is running at the DR site.

Production runs in the DR site for whatever period you have chosen before returning to the Primary site. During this period, your production site performs as the DR site until the next switch-over.

Updates to the production running at the DR site is kept and replicated back to the Primary site.

DR test frequency

DR test frequency depends on many factors, but in general you should have a DR test *at least once per year*. This is the bare minimal frequency accepted by auditors to prove your ability to recover.

In setting the frequency of your cloud DR testing you should strive for more frequent than once a year, depending on multiple elements specific to the cloud, such as the following concepts:

- ▶ How dynamic is your production environment? The more your production environment changes, the more you need to exercise a DR test to verify that changes introduced have not affected your ability to recover,
- ▶ How dynamic is your CSP infrastructure? As an example, if you have chosen to provision some of the DR resources *on demand*, you should consider that you have no control on the type of resources and technologies you will be able to find at the moment of a disaster. In fact, the CSP may have changed the underlying technologies (for example, the servers) or configurations (for example, network topologies) or service levels (for example, time to provision new resources) since your last test. So, it is suggested to increase the testing frequency in such a case, especially if you believe that your applications may be sensitive to changes in the underlying infrastructure.

Of course, each test carries a cost associated with the use of the DR site, effort and so on, so leveraging automation, tools and software products like IBM Cloud Resiliency Orchestrator, operations required to run a DR test can be simplified, allowing the associated costs and effort and test more frequently. More automation will also enhance overall recoverability and improve the recovery time objective (RTO).

On-premise to DR cloud considerations

There are some considerations when planning to use a cloud service provider to implement a DR solution for a production running on-premises:

Compatibility with the cloud provider

Doing DR on a cloud provider from on-premises has the same requirements as a migration project to the cloud. You must verify that the workload which is going to be deployed at the cloud provider can run on that infrastructure.

Constraints on the cloud provider

Despite all the abstractions the cloud provider can do, in the end you might still be facing some constraints in your design. An example is the limitation you might have to boot a virtual machine from a replicated disk (LUN), when that VM runs in the in cloud provider managed hypervisor infrastructure (Public Cloud). Other examples are in the areas of LAN and WAN, where the Software Defined Network (SDN) provides some flexibility, but not the full flexibility you have in your own infrastructure.

Cloud resource scaling

Having a DR on cloud might provide the false expectation that cloud scalability is infinite. It is not particularly in the cloud data center where you have decided to replicate your data.

Provisioning time for on-demand resources

Cloud providers are not created equal, so you need to evaluate different offerings and the SLAs associated with the provisioning time of on-demand resources, as they impact the RTO of your DR solution. Another important point of consideration is the commitment offered by the cloud provider (or the SLA), and what limitations there might be. (For example, does it guarantee the SLA regardless of the number of concurrent requests?).

Resource pricing

On-demand resources are cheaper, as they are usually billed by usage, which means you only pay when you use the resource. It is important to evaluate all the caveats associated with the usage billing, and how that might impact your total cost of ownership (TCO) with hidden and unplanned costs - as in emergency the use of on-demand resources which might incur higher costs compared to what you have originally planned.

Cloud networking (LAN)

Having a DR on cloud might also imply that you need to rebuild your network structure to adapt to the network structures and constraints of the cloud provider. In some cases, you will be forced to use Software Defined Network (SDN) and Network Functions Virtualization (NFV) to reproduce your complex enterprise network layout. Careful planning and evaluation of the performance and scalability limitations is essential to assure your reprovisioned compute can work properly.

Cloud networking (WAN)

All the external networks connected to your primary (failed) site must be re-routed to the alternate cloud site. Different options are available and you need to evaluate and plan for options that best fit your requirements. Consider the charges associated with the use of the network (download and upload charges) when using the cloud provider resources, as they might be substantial.

Cloud to DR cloud considerations

Here are some considerations on top of what has been listed for on-premises to DR cloud, when planning to use a cloud service provider to implement a DR solution for a production running on cloud:

Cloud provider compatibility amongst different providers

Different providers have different technologies. Like in the on-premises to cloud topology, moving your workload from one provider to another is equivalent and has the same complexity.

Particular attention has to be placed in investigating the different methodologies, tools and interfaces of the two cloud providers. In fact, they might strongly impact the portability of workloads among different providers and could create a true lock in. In the following we cover some of the technical domains to which particular attention should be paid.

Monitoring

Often native tools have been used to implement the monitoring of the cloud environment. Hence, moving a workload to a different CSP may involve changing monitoring frameworks, which may involve the installation of different clients on the host systems or integration back to on-premise notification systems or event correlation. However, the native monitoring tools are often limited to fully address enterprise level monitoring requirements. For example, they work quite well to monitor the cloud resources, but may be lacking in the areas of middleware or databases deployed on the VMs. So, we suggest adopting additional tools that augment the native monitoring tools. In our case we have introduced a cloud monitoring tool that natively integrates with the CSP's monitoring via APIs. There are also other tools with similar capabilities. The introduction of such a tool reduces CSP dependency since the tool can natively integrate with multiple cloud monitoring frameworks and facilitate DR processes.

APIs

Any scripts or applications that interchange data with the CSP native APIs are impacted in the case of a relocation of the workloads and this has to be taken into account.

Network

Networking is one of the most dynamic elements between different CSPs. Key features that have to be taken into account include the possibility of bringing the original IP when moving workloads to a cloud. Other functionalities to look for are the ability to interconnect different subscriptions environments and to interconnect virtual networks between different cloud sites.

Provisioning

The interface (graphical or APIs) to provision resources on a cloud varies from CSP to CSP. We suggest introducing a brokerage tool (such as the *IBM Cloud Brokerage*) to consistently manage resources provisioning and cost management across CSPs at the DR site.

Backup

Generally CSPs provide a mechanism for performing basic backups. To implement a more complete and sophisticated backup functionality, consider implementing an additional product or framework. One example is Azure IaaS VM backup. On AWS, tools are available in the marketplace. It is strongly suggested to evaluate the trade off between using a CSP native backup capability versus an independent backup software that can be utilized on any CSP. In fact, using the native tools may generate a lock in or at least make it difficult to manage your workload on a different provider. This could be overcome using a tool, such as *IBM Spectrum Protect*, to be installed on cloud and to store backup data on the CSP storage resources.

When implementing a DR cloud to cloud on the same CSP, of course, the solution is simplified as the same technology is leveraged in both the primary and the secondary site. However, these considerations remains valid in a more broad context to avoid lock-ins.

Other considerations might come from the following topics:

Resource pricing when different providers are used

It is important to evaluate and understand all the caveats associated with the usage billing of the two providers, and how that might impact your TCO with hidden and unplanned costs.

Cloud networking: LAN and WAN

The two different CSP might have a different approaches to Software Defined Network (SDN) and Network Functions Virtualization (NFV) and reproducing what has been implemented on one CSP to another might not be possible, presenting challenges or hidden costs that ought to be discovered up-front.

Provisioning time and SLA

How the two providers differ on provisioning time and SLAs has to be investigated and determined. In particular, if you plan to use on-demand (for example pay-per-use) resources, you should carefully evaluate the SLA that the CSP commits to on provisioning time, as that will impact the whole RTO.

Cloud provider compatibility with same provider

When implementing a DR cloud to cloud on the same CSP, of course, the solution is simplified given that the same technology in the primary and secondary site. However, these considerations remain valid in a more broad context to avoid lock-ins.

The use of the same provider for DR is the easiest choice, however it could not be possible in all cases. Reasons could be, for example, availability of a secondary site in a given region and SLAs provided by the CSP to support the DR.

Cloud to DR on-premises considerations

A company that wants to maintain an extra flexibility in its cloud strategy might decide to maintain a copy of the company data on premises. The main aspects to consider in planning to use a traditional site as a DR site for a production running on clouds are:

Compatibility with the cloud provider

Doing DR from a cloud provider to on-premises has challenges on the compatibility side. You might need to mimic the hypervisor infrastructure you have used on the cloud to avoid hypervisor migrations.

Dependencies from the cloud provider

You should also be mindful of the possible dependencies on the services or features that your applications have started using on the cloud provider (such as DNS, monitoring, logging) as they might not be available in your on-premises.

Operating system licenses

On the CSP, operating systems are usually part of the running fee for the compute you use on the CSP infrastructure. When you run on your own on-premises infrastructure you, most likely, have to reprovision all the licenses and you are not entitled to use the CPS's licenses.

Reprovisioning operating system and subsystem (for example database management system) licenses might also impact your recovery time objective and constitute a substantial cost.

Common pitfalls

Based on our experience, here are some common pitfalls when implementing a DR solution on cloud:

Plan and design for the best scenarios

As already said, it is important that your solution and your testing methodology are planned and designed to mimic what the real condition might be. Designing for the best case reduces the costs, but also entails some risks.

Some examples:

- ▶ Designing the solution to use low cost or low performance components in the DR site (such as storage, server). While these systems may be able to prove the concept of a disaster recovery, they will likely not be able to run your full volume production workload for an extended period of time in case of an emergency.
- ▶ Reuse your decommissioned production components (such as storage, server) due to technology change in the DR solution. If you have changed your production technology it was probably because these old component were not a good fit for your requirements any longer.

Single-points-of-failure

Single-points-of-failure (SPOF) can be anywhere in a solution, not only on the technology side. Your solution might dependent on people, vendors, providers, and other external dependencies. It is essential that you identify clearly in advance most of your SPOF or at least have a plan to mitigate your dependencies. Also be prepared to discover SPOF during the first sessions of your DR test.

Among SPOFs, Provider Risk is a condition to consider in your DR plan. When you have both Production and DR on the same provider your risk has increased and must be carefully considered.

Only have a Plan A

Risks cannot be eliminated; they can only be mitigated. The more risk you mitigate with a single solution, the more the solution would cost. Thus any solutions have a residual risk.

Having a Plan B with a different recovery time objective might help you to mitigate additional risks, while not adding too much costs.

Your Plan B recovery time objective might not be the one expected, meaning you will need more time to restart your operations, and thus you will suffer greater business impact from the emergency, but restarting later is far better than not restarting at all.

A possible Plan B in a two sites topology might be to have a periodic third site backup of the data, at least located at a distance sufficient to avoid the risk of being affected by the same events that affect the primary and principal DR site.

Poor DR testing methodology

An untested DR solution is like running blind in a woods; there is a fair chance that you will come to an abrupt and painfully injurious stop.

Testing is essential to guarantee that you have a valid solution, also but the validity of the solution is dependent on establishing the correct conditions under which you are testing. Invalid conditions will invalidate the solution no matter how rigorously you have tested it.

If you plan to perform a DR test, by doing an orderly shutdown of operations at the primary site and an orderly start at the DR site, you can be sure that your DR site environment is valid and capable of supporting a workload that you are able to dynamically relocate. But is this what will happen during an emergency? You must build in processes and budget time into your RTO to account for resolving inconsistent application states and reconciling interdependent applications.

You should design your tests to mimic as much as possible the possible emergency conditions, by simulating a so called *rolling disaster condition*, where your IT service is impacted progressively by the emergency. This is the best way to test your solution and have a reasonable understanding whether it is resilient (ability to resist to stress conditions).

Networking aspects

In this section, we briefly illustrate key aspects in network design that might impact the DR solution.

Five networks

In a DR solution we can basically classify the networks in five types:

- ▶ Management and monitoring
- ▶ Replication
- ▶ DR test
- ▶ Failover (or emergency)
- ▶ Fallback

In cloud, the concept of these networks continue to be valid but may need to be implemented differently depending on the CSP specific network services and policies.

Management and monitoring

This is the least demanding connectivity in terms of bandwidth, as it is only designed to sustain:

- ▶ Administrator access to remote systems
- ▶ Monitoring, logging and support activities (such as SNMP, traps, logging and so on)
- ▶ Antivirus definitions (active, real time, in-line antivirus scanning can be a latency sensitive, intensive process) and security patches

You can normally create network tunnels within the replication network to sustain these functions.

Replication

Replication is the connectivity required to duplicate data from the primary to the DR site. It is mainly determined by the synchronization method:

- ▶ Online full synchronization

In the online synchronization, a full copy the entire volume (including empty spaces) is transmitted over the network to the DR site, together with delta updates that are captured and sent. In the online synchronization, this might not be a one-shot operation.

Sometimes, you need to periodically perform a full-refresh of the replicated data, or to perform a bulk-transfer of a good portion of your source data and iteratively replicate the deltas until they are sufficiently small.

- ▶ Offline synchronization

In the offline synchronization, the large part of the data is made available at the DR site through alternative means (disk images), and thus only delta updates that are captured and sent will traverse the network. This allows you to reduce the network requirement for replication, but you should carefully and methodically determine the frequency of the periodic refreshes, as this might impact your RTO/RPO.

Another consideration is that replication network has to work in two ways:

- ▶ To DR site: This is the normal flow when production is executed on-prem and DR site is there just for protection.
- ▶ From DR site: This is the flow you have in emergency where production is running at the DR site and your on-premises (or an alternate on-premise) is the target of replication.

The safest way to size your bandwidth requirement for is to size for online synchronization, and not just daily updates.

In the cloud, the replication network is generally provided by the CSP. For simplicity let us assume that both the production and DR sites are in the same CSP's cloud. In such a case the replication network is part of the site to site connectivity that the CSP makes available to its clients. It is important to verify the costs (if any) and performance of such a connectivity while planning for a DR solution.

DR test

As we have seen previously, two types of DR testing exist, and their network requirements of are different:

- ▶ DR simulation: Run while the production operations continue on-premises, DR test users access the DR simulation environment by pointing to different servers reprovisioned at the DR site. If the original IP addresses of the recovered servers are needed, they'll have to be NATed to these different IP addresses to avoid duplicated IP addresses in production or, even worse, real transaction are executed toward the DR servers instead of the real production servers.

Another option could be to perform policy based routing based on the source address of the test users.

- ▶ Switch-over: Run by moving the production to the DR site. From a networking perspective this is equivalent to the real emergency, as the entire production networks have to be routed to the DR site, while replication network flow is reversed (from DR site to on-premises), to bring updated data back to on-premise, which now works as the DR site.

In the cloud, the test network is provisioned as part of the DR cloud resources and is often subject to limitations as mentioned before. If the DR is implemented on the same CSP, it would be easier to configure a network to simulate or switch the production environment.

Failover or emergency

The entire production network has to be routed to the DR site, while replication network flow is reversed (from DR Site to on-premise), to bring updated data back to on-premises.

During the emergency the network functions (routing and security) are also transferred from on-premises to the DR site and if the primary site is physical on-premises, it might need to be virtualized to adapt to the target CSP requirements. A pre-check and maintenance of this physical-to-virtual network functions is a key success factor during the emergency and has the same importance as the Data Replicator or the Server Reprovisioning technique.

Apply all the considerations for cloud to DR cloud to this failover (or emergency) network as well.

Fallback

Fallback presents the same challenges as the DR simulation by switch-over seen before, as you keep having the production running at DR site, while intercepting and sending updates back to the fallback site.

The quantity of data that flows back to the fallback sites depends on the emergency that happened.

For short-term emergencies, where the original site is unavailable for a period of time, but that has left servers and storage intact, a delta-resynch might fit the need to bring the operation back to the on-premise.

For other emergencies that have forced a change to the site server or storage in the original on-premises might require a full resynchronization of data, and so the fallback will happen ordinarily at the most convenient time after the synch-point has been achieved.

Full versus partial failover

The *partial failover* is a feature that potentially increases the availability of the customer services by allowing the customer to run and recover some of their services in DR, while all the others are still running production at the customer data center.

This requires the extension of the production site network(s) to DR site. That can happen on Layer 2 (like: L2TPv3, and Cisco Overlay Transport Virtualization) or in Layer 3 (like: Cisco Locator and ID Separation Protocol (LISP), Virtual Private LAN Services (VPLS) or Software Defined Network technique). Additional considerations may require the full control of the landing DR hypervisor, so techniques like Virtual Extensible LAN overlays included in solutions like VMware NSX (Network Virtualization and Security Platform).

From a network perspective this extension requires additional planning because of possible impacts on security and performances, at least because the servers that were previously in the local LAN, are now placed over a longer distance (WAN).

This impact must be evaluated in two directions: from a server to user perspective and from a server to server perspective, especially for latency-sensitive applications or services.

User-to-DR server

Consider the scenario where only a server (server A) has been failed over and is currently running at the DR data center in a partial failover situation.

The user-to-server A session, will reach the server A DR-VLAN through the existing customer routing and security running on-premise. From the customer-managed routing perspective, server A is still in the source-VLAN; that is, the fact that source-VLAN is extended in the DR data center is transparent to the customer managed router.

So, the customer managed router will place an Address Resolution Protocol (ARP) frame for server A on the source-VLAN. This broadcast frame will be seen and picked up by the network extension device or function at the customer location, encapsulated in the L2TPv3 tunnel, and sent to the network extension device or function at the DR site. The frame is then placed DR-VLAN where server A is running.

Note: If the customer managed router has a static ARP table for server A, this will not work and the static entry needs to be replaced.

While this is applicable to both internet and directly connected customers, the network design must consider the additional latency due to the processing of the L2TPv3 tunnel, and the distance.

Where the additional latency may be predictable within a range for directed connected clients with network latency SLAs (the overall latency does not come only from the link latency), the internet connected customers might experience a different situation due to the unpredictable latency of the public internet.

The network design must also consider possible maximum transmission unit (MTU) adjustments or additional packet loss situations that the application running in the production site might demand in terms of requirements.

Server-to-DR servers

Consider the scenario where two or more servers have been failed over and are currently running at the DR data center in a partial failover scenario (server A on DR-VLAN1 and server B on DR-VLAN2), while the rest of the servers keep on running on-premises.

In this case even the main concepts are similar to what we have seen in previous example, since the customer routing and security services are still running on-premises. Each frame has an additional latency to be considered, as all the routing and security functions are performed on-premises.

The more servers in a partial F/O mode, the more the impact of latency on the customer routing and security functions.

Network virtualization

Network virtualization (NV) is the ability to create logical, virtual networks that are decoupled from the network hardware. This ensures that the network can better support virtual environments.

NV abstracts networking connectivity and services that have traditionally been delivered via hardware into a logical virtual network that is decoupled from and runs independently on top of a physical network in a hypervisor.

NV can deliver a virtual network within a virtualized infrastructure that is separate from other network resources.

NV is available on cloud service providers, as they largely adopted network virtualization to provide a secure multi-tenant environment, sharing the same physical hardware component among all their customers. Additional considerations should be taken into account to understand at what level the customer has control over the NV function.

Some NV functions require the full control of the landing DR hypervisor, so techniques like VXLAN (or Network Overlays) may not be applicable or be subject to restrictions.

Network function virtualization

With network virtualization, a set of network functions can be virtualized as well, these typically are services like:

- ▶ Customer firewalls
- ▶ Load balancing solutions
- ▶ WAN optimization
- ▶ Intrusion Detection Systems (IDS)

The virtualization of those function cloud be the only essential step in being able to exploit cloud service providers, where it is not possible to recreate them using traditional appliances or specialised hardware.

In a DR solution, network function virtualization might bring in the additional challenges on how to transform the backup of the configuration and definition of the physical appliance on-premises to adapt to what is available from the CSP space, both natively or from its marketplace.

Another challenge might be on the performance side, as the network function virtualization runs over a standardized compute and not on specialized hardware like an appliance on-premises.

Bring Your Own IP

Network virtualization enabled the *Bring Your Own IP* (BYOIP), which is a function that allows you to define and use, at the DR site, subnets having a user-defined address space.

You might want to exploit this feature to maintain your on-premises IP in the DR site. When doing so, consider that pros and cons exist.

Pros

BYOIP includes the following advantages:

- ▶ System administrators are more comfortable with operating in an environment that mimics their production site
- ▶ Hard coded IP address in systems works like on-premises
- ▶ Domain Name Servers does not require re-convergence

Cons

BYOIP includes the following disadvantages:

- ▶ Presents more challenges in network extensions
- ▶ Presents more challenges when handling Partial Fail-Over conditions

Related publications

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

IBM Redbooks

The following IBM Redbooks publications provide more information about the topic in this document. Note that some publications that are referenced in this list might be available in softcopy only:

- ▶ *Implementation guide for IBM Spectrum Virtualize for Public Cloud on AWS*, REDP-5534
- ▶ *Implementing the IBM Storwize V7000 with IBM Spectrum Virtualize V8.1*, SG24-7938-06
- ▶ *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 V8.1*, SG24-7933

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

Online resources

The following websites are also relevant as further information sources:

- ▶ Solution for integrating the FlashCopy feature for point in time copies and quick recovery of applications and databases:
<http://www.ibm.com/support/docview.wss?uid=ssg1S4000935>
- ▶ Information about the total storage capacity that is manageable per system regarding the selection of extents:
<https://www.ibm.com/support/docview.wss?uid=ssg1S1010644>
- ▶ Information about the maximum configurations that apply to the system, I/O group, and nodes:
<https://www.ibm.com/support/docview.wss?uid=ssg1S1010644>
- ▶ IBM Systems Journal, Vol. 42, No. 2, 2003:
http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=5386853
- ▶ IBM Storage Advisor Tool:
<https://www.ibm.com/us-en/marketplace/data-protection-and-recovery>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



REDP-5466-01

ISBN 0738457809

Printed in U.S.A.

Get connected

