Redbooks
ibm.com/redbooks

Implementation Guide for IBM Spectrum Virtualize for Public Cloud 8.5

For AWS and Azure Public Clouds

Vasfi Gucer

Sambasiva Andaluri

Pankaj Deshpande

Andrew Greenfield

John Nycz

Sushil Sharma

Jackson Shea

Saurabh Singh

Earl Springer

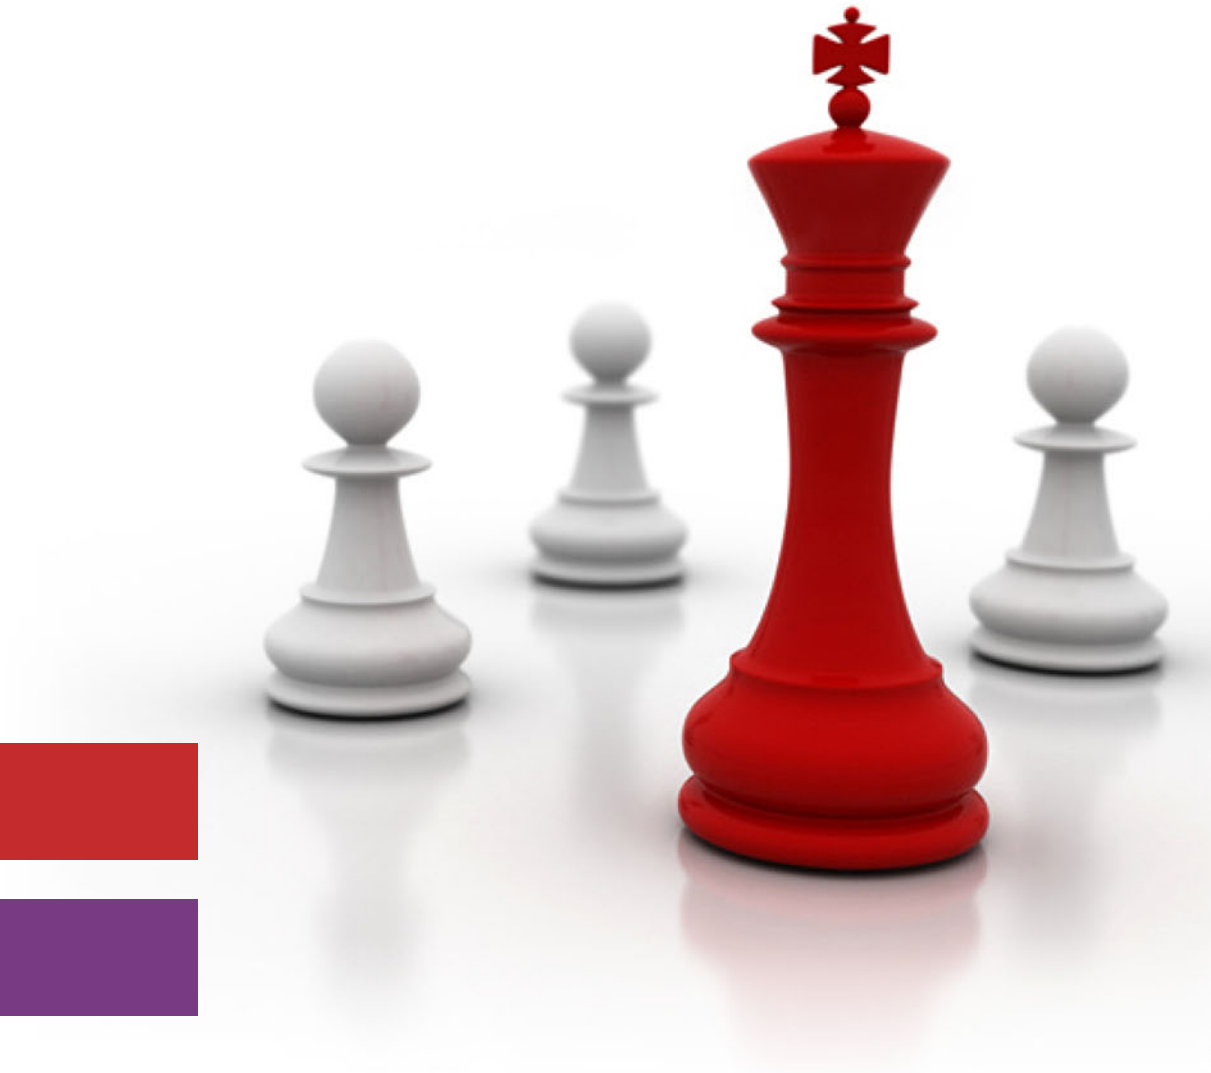

Cloud

Storage

IBM

Redpaper

IBM Redbooks

Implementation Guide for IBM Spectrum Virtualize for Public Cloud 8.5

July 2023

REDP-5671-00

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

First Edition (July 2023)

This edition applies to IBM Spectrum Virtualize for Public Cloud on Azure 8.5.X and IBM Spectrum Virtualize for Public Cloud on Azure 8.5.X.

Contents

Notices

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

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

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

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at https://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®
Aspera®
Db2®
Easy Tier®
FlashCopy®
HyperSwap®
IBM®
IBM Cloud®
IBM FlashSystem®
IBM Spectrum®
Redbooks®
Redbooks (logo)  ®
XIV®

The following terms are trademarks of other companies:

ITIL is a Registered Trade Mark of AXELOS Limited.

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

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.

Red Hat, OpenShift, are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

VMware, and the VMware logo are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

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

Preface

IBM Spectrum® Virtualize for Public Cloud supports clients in their IT architectural transformation and migration toward the cloud service model. It facilitates hybrid cloud strategy and cloud-native strategy, and it provides the benefits of having advanced storage features, functions, and capabilities usually found only in on-premises data centers. IBM Spectrum Virtualize for Public Cloud enhances the capabilities of traditional cloud storage offerings.

IBM Storage Virtualize software, as part of an IBM Storage IBM FlashSystem® or IBM SAN Volume Controller product running on-premises, can virtualize over 500 different storage systems from IBM® and other vendors. This wide range of storage support means that the storage solution can be used with almost any storage in a data center today and can integrate with its counterpart IBM Spectrum Virtualize for Public Cloud, which supports several public cloud environments, including Amazon Web Services (AWS) and Microsoft Azure to help facilitate an advanced hybrid cloud environment.

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

Authors

This paper was produced by a team of specialists from around the world.

Vasfi Gucer works as the Storage Team Leader on the IBM Redbooks® Team. He has more than 30 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 storage, cloud computing, and cloud storage technologies for the last 10 years. Vasfi is also an IBM Certified Senior IT Specialist, Project Management Professional (PMP), IT Infrastructure Library (ITIL) V2 Manager, and ITIL V3 Expert.

Sambasiva Andaluri (Sam) is an experienced Developer turned Solution Architect Leader with over 30 years of experience. For the past decade, he has been a pre-sales and post-sales solution architect for trading systems at Fidessa, a pre-sales solution architect at AWS, and a site reliability engineer for onboarding ISVs for Google Marketplace at an IBM Business Partner. He brings multifaceted experience to the table.

Pankaj Deshpande is an IBM Spectrum Virtualize for Public Cloud Architect working in IBM Pune, India. A key part of his role at IBM is in architecture, designing and developing products, and evolving technology and product strategy for IBM. His key technical contributions consist of storage technologies, storage security, cybersecurity, public and private cloud integration, storage for virtual machines (VMs), and Kubernetes. Pankaj also represents IBM on some standards bodies. He has designed and delivered innovative, quality system software products in the areas of storage, networking, security, cloud, and consumer segment. In a career spanning more than 20 years, he has worked on various designs involving scale-out, distributed, high-performance, and highly available (HA) architectures.

Andrew Greenfield is an IBM Global IBM XIV® and Flash Solution Engineer who is based in Phoenix, Arizona. He holds numerous technical certifications from Cisco, Microsoft, and IBM. Andrew brings over 25 years of data center experience inside the Fortune 100 to the team. He graduated magna cum laude honors from the University of Michigan, Ann Arbor. Andrew has written for and contributed to several IBM Redbooks publications.

John Nycz is an Advanced Subject Matter Expert for IBM Spectrum Virtualize and IBM FlashSystem. He has more than 10 years of experience in the areas of systems management, networking hardware, and software. John has been with IBM for more than 20 years, and has been a member of numerous development, project management, and support teams. In the last 7 years, he has been a member of the IBM Spectrum Virtualization Support Team.

Sushil Sharma is an Advisory Software Engineer working as the Development Lead at IBM India Systems Development Lab, Pune. He has 13 years of industry experience, and has been working on IBM FlashSystem and IBM SAN Volume Controller since 2015. He has worked and delivered features, including 3-Site Replication (Metro Mirror, IBM HyperSwap®, and GUI), iSCSI/iSER support on IBM SAN Volume Controller, and IBM Spectrum Virtualize for Public Cloud. He has contributed to Tech Sales enablement, Beta enablement, and L3 support enablement. He also created various customer collaterals, such as blueprints and technical white papers, about IBM SAN Volume Controller features. He holds a master's degree in Computer Application from Mumbai University, India.

Jackson Shea is a Level 2 certified IBM Information Technology Specialist/Architect who performs design and implementation engagements through IBM Lab Services. He has been with IBM since April 2010. He was a Lead Storage Administrator with a large health insurance consortium in the Pacific Northwest, and has been working with IBM equipment since 2002. He has over 12 years of experience with IBM Spectrum Virtualize (formerly known as the IBM SAN Volume Controller) and related technologies. Jackson is based in Portland, Oregon. He received his Bachelor of Science degree in Philosophy with minors in Communications and Chemistry from Lewis & Clark College. Jackson's professional focus is IBM Spectrum Virtualize, but he is conversant with storage area network design, implementation, extension, and storage encryption.

Saurabh Singh is an Advisory Software Engineer working as the development lead for IBM Spectrum Virtualize for Public Cloud at IBM Pune, India. He has almost 15 years of experience in IBM Spectrum Virtualize products and has worked on the development of various software features like Fibre Channel and Transparent Cloud Tiering (TCT). He holds a Bachelor of Technology degree in Computer Science and Engineering from Harcourt Butler Technological Institute, Kanpur, India.

Earl Springer is a Senior Partner Technical Specialist supporting IBM Storage hardware and software technologies for IBM Ecosystem. Working with IBM Value Add Distributors and IBM Business Partners across North America, Earl helps co-create opportunity-driven solutions, develops and delivers enablement, and helps IBM Business Partners understand the value and differentiators of IBM storage technologies. Originally working with IBM Canada, Earl now works with IBM US in Tampa, FL. He has an extensive technical background in software development, compute technologies including IBM Power, networking, virtualization, cloud, and storage technologies. Earl is a designated Subject Matter Expert for IBM Spectrum Virtualize, helped develop certification and education content for several hardware and software storage certification exams, and helped develop and deliver education for various IBM technical events, including IBM TechU.

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

Mandar Vaidya, Dhiraj Verma, Prasad Dasari, Viswanath Kuchnagari
IBM India

Also, special thanks to Guillaume Legmar from IBM France for his contributions to the project.

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 IBM Redbooks residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

  ibm.com/redbooks

- Send your comments in an email to:

  redbooks@us.ibm.com

- Mail your comments to:

  IBM Corporation, IBM Redbooks
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on LinkedIn:

  https://www.linkedin.com/groups/2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/subscribe

- Stay current on recent Redbooks publications with RSS Feeds:

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

1

Introduction

This chapter describes the IBM Storage Virtualize product that is implemented in a cloud environment, which is referred to as *IBM Spectrum Virtualize for Public Cloud*.

This chapter also provides a brief overview of the technology that is behind the product, and introduces the drivers and business values of IBM Spectrum Virtualize in the context of public cloud services. Finally, it describes from a high-level perspective how the solution works.

This publication describes IBM Spectrum Virtualize for Public Cloud 8.5.4

This chapter includes the following topics:

1.1 Introducing IBM Spectrum Virtualize for Public Cloud

Most businesses today are at some stage of digital transformation. Many are taking a customer-driven, digital-first approach to many aspects of their business, from business models to customer experiences, to processes and operations.

Businesses are starting to use artificial intelligence (AI), automation, and other digital technologies to leverage the power of data and drive intelligent workflows, faster and smarter decision-making, and real-time response to market disruptions with a goal of ultimately changing customer expectations and creating new business opportunities.

These businesses are deciding on IT architectures that will impact their operations over the next few years. Most of these organizations have realized the value that both public and private cloud services can deliver. The role of hybrid cloud as a critical part of digital transformation has matured and is considered the architecture of choice that helps drive innovation businesses seek allowing them to continually adapt to change.

One of the challenges for businesses is integrating public cloud services with existing IT infrastructure. Organizations want to achieve agility and flexibility without introducing new layers of complexity or requiring significant 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, service, application, or management layers as examples. 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 integrity and consistency.

IBM Spectrum Virtualize for Public Cloud supports clients in their IT architectural transformation and migration toward the cloud service model. It facilitates hybrid cloud strategy and cloud-native strategy, and it provides the benefits of advanced storage features, functions, and capabilities that usually are found only at on-premises data centers. IBM Spectrum Virtualize for Public Cloud enhances the capabilities of traditional cloud storage offerings.

IBM Storage Virtualize software, as part of an IBM FlashSystem or IBM SAN Volume Controller product running on-premises, can virtualize over 500 different storage systems from IBM and other vendors. This wide range of storage support means that the storage solution can be used with almost any storage in a data center today and can integrate with its counterpart IBM Spectrum Virtualize for Public Cloud, which supports several public cloud environments including Amazon Web Services (AWS) and Microsoft Azure to help facilitate an advanced hybrid cloud environment. For more information, see Chapter 3, “Solution architecture” on page 31, and Chapter 4, “Planning and preparing for IBM Spectrum Virtualize for Public Cloud” on page 55.

Important note on IBM Storage rebranding: IBM is rebranding its IBM Spectrum Storage family of products to as part of the IBM focus on simplification. Throughout 2023, IBM will change the IBM Spectrum branding to IBM Storage branding. Currently, IBM Spectrum Virtualize for IBM FlashSystem and IBM SAN Volume Controller are rebranded as IBM Storage Virtualize. At the time of writing, IBM Spectrum Virtualize for Public Cloud will be rebranded as IBM Storage for Public Cloud soon, and as of this release still retains its original branding.

1.2 IBM Spectrum Virtualize for Public Cloud

Designed for software-defined storage (SDS) environments, IBM Spectrum Virtualize for Public Cloud represents a solution for public cloud implementations, and includes technologies that complement and enhance the capabilities of public cloud offerings.

For example, traditional practices providing data replication by copying data from a storage array at one facility to a largely identical storage array at another facility is not an option for public cloud. Also, the use of conventional software to replicate data at the application layer imposes unnecessary loads on the application's servers. Use cases are analyzed in Chapter 2, "Typical use cases for IBM Spectrum Virtualize for Public Cloud" on page 13.

IBM Spectrum Virtualize for Public Cloud delivers a powerful solution for deploying IBM Storage Virtualize software within public clouds. This capability includes flexible licensing options and enables hybrid cloud solutions to transfer data between on-premises data centers by using any IBM Storage Virtualize based appliance, such as an IBM Storage IBM FlashSystem or IBM Storage SAN Volume Controller, and multiple public cloud environments.

With a deployment that is designed for the cloud, IBM Spectrum Virtualize for Public Cloud can be deployed in AWS or Azure cloud data centers around the world where, after defining the infrastructure, an installation template automatically installs and configures the storage.

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* as a technology that makes one set of resources resemble another set of resources, preferably with other characteristics.

Virtualization is a logical representation of resources that are not constrained by physical limitations while hiding any aspects of complexity. It also can add or integrate new functions with existing services and be nested or applied to multiple layers of a system.

The aggregation of volumes into storage pools enables better managed capacity, performance, and can provide multiple tiers 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 it is referred to as *block-level virtualization*, or the block aggregation layer. For clarification, the block-level volumes that are provided by the cloud provider are exposed as target volumes, and are seen by IBM Spectrum Virtualize for Public Cloud as a managed disk (MDisk).

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

The virtualization terminology is included into the wider concept of SDS, which is 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 storage systems or installed on any commodity x86 hardware and or hypervisor.

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

In general, SDS technologies apply the following concepts:

- A shared-nothing architecture (or sometimes, a partial or fully shared architecture) with no single point of failure (SPOF) and nondisruptive upgrades.
- Scale-up or scale-out mode: Each extra server running the SDS software acts as a building block that provides predictable increases in capacity, performance, and resiliency.
- Multiple classes of service: SDS can provide file-based, object-based, block-based, and auxiliary and storage support services. SDS solutions also can be integrated into a hybrid or composite SDS solution.
- High availability and disaster recovery (HADR): SDS solutions can tolerate various levels of availability and durability as SDS systems typically are self-healing and self-adjusting.
- Lower total cost of ownership (TCO): Lowers the TCO for those workloads that can use the services that are provided by SDS.

1.2.2 IBM Spectrum Virtualize for Public Cloud benefits

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 public cloud services. By using standard x86-based servers, IBM Spectrum Virtualize for Public Cloud can be added to cloud infrastructures to deliver more features, functions, and capabilities that enhance the storage offering that is available on the public cloud catalog.

The benefits of deploying IBM Spectrum Virtualize for Public Cloud are two-fold:

- Public cloud storage offering enhancement: IBM Spectrum Virtualize for Public Cloud enhances the public cloud storage offerings by increasing standard storage, and offering features, functions, and capabilities that help overcome specific limitations:
  - Snapshots: A volume's snapshots occur on high-tier storage with no options for a lower-end storage tier. By using IBM Spectrum Virtualize for Public Cloud, the administrator has more granular control, which enables a production volume to have a snapshot that is stored on lower-end and more cost-effective storage.
  - Volume size: Most cloud storage providers have a maximum volume size (typically a few terabytes) that can be provided by a few nodes. At the time of writing, IBM Spectrum Virtualize for Public Cloud allows for volumes that are sized up to 320 TB, which accommodates up to 20,000 host connections.
  - Native storage-based replication: Replication features are natively supported, but are typically limited to specific data center pairs, a predefined minimum recovery point objective (RPO), and are typically accessible only when the primary volume is down. IBM Spectrum Virtualize for Public Cloud provides greater flexibility in storage replication to allow for user-defined RPO and replication between any other on-premises system that is running IBM Storage Virtualize or another public cloud environment supporting IBM Spectrum Virtualize for Public Cloud.

- New features for public cloud storage offering: IBM Spectrum Virtualize for Public Cloud introduces to the public cloud catalog new storage capabilities. Those features are available on SAN Volume Controller and IBM Storage Virtualize, but are not available by default. The following extra features that are provided on public cloud are 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, replication consistency is achieved through storage-based replica peer cloud storage with primary storage on-premises. Because of standardization of the storage service model and inability to move its own storage to a cloud data center, the storage-based replica is achievable only by involving an SDS solution on-premises.

    In this sense, IBM Spectrum Virtualize for Public Cloud offers data replication between the IBM Storage IBM FlashSystem family, IBM Storage SAN Volume Controller and public cloud and extends replication to all types of supported virtualized storage on-premises.

    Working together, IBM Storage Virtualize-based appliances and IBM Spectrum Virtualize for Public Cloud support synchronous and asynchronous mirroring between the cloud and on-premises for more than 500 different storage systems from various vendors. In addition, they support other services, such as IBM FlashCopy®, Snapshots, IBM Easy Tier®, and IBM Safeguarded Copy.
  - Disaster recovery (DR) strategies between on-premises and public cloud data centers as alternative DR solutions.

    One of the reasons to replicate is to have an identical copy of the source data in another physical location from which to restart operations if there is data loss. IBM Spectrum Virtualize for Public Cloud enables DR for both virtual and physical environments, which adds new possibilities compared to the many software replicators in use today that primarily handle virtual infrastructure only.
  - Benefit from familiar, sophisticated storage functions in the cloud to implement reverse mirroring.

    IBM Spectrum Virtualize enables the possibility to reverse data replication to offload data from IBM Spectrum Virtualize for Public Cloud on AWS or IBM Spectrum Virtualize for Public Cloud on Azure back to on-premises or to another cloud provider.

IBM Storage Virtualize for on-premises and IBM Spectrum Virtualize for Public Cloud for public cloud storage provide a data strategy that is independent of the choice of infrastructure, which delivers tightly integrated features, functions, and capabilities along with consistent management across heterogeneous storage and cloud storage.

The software layer that is provided by IBM Storage Virtualize on-premises or IBM Storage Virtualize for Public Cloud 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 interactions.

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

1.2.3 IBM Spectrum Virtualize for Public Cloud features

IBM Spectrum Virtualize for Public Cloud helps make cloud storage volumes (block-level) more effective by providing key functions that are not natively available on the public cloud catalogs and are traditionally deployed within disk array systems in on-premises environments. For this reason, IBM Spectrum Virtualize for Public Cloud improves and expands the capabilities of the specific block storage offerings from both AWS and Azure.

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

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

Feature	Benefits
Single point of control for cloud storage resources.	Designed to increased management efficiency and help support application availability.
Pools the capacity of multiple storage volumes.	► Helps overcome volume size limitations. ► Helps manage storage as a resource to meet business requirements, and not only as a set of independent volumes. ► Helps an 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 a storage volume offering.
Clustered pairs of virtual servers that are configured as IBM Spectrum Virtualize for Public Cloud engines.	► Usage of standard Elastic Compute Cloud (EC2) compute nodes. ► Designed to avoid single point of hardware failures.
Manages tiered storage.	► Helps to balance performance needs against infrastructures costs in a tiered storage environment. ► AI-based algorithm automatically moves data to the right storage tier or class at the right time.
Easy-to-use IBM FlashSystem family management interface.	► A single common interface for storage configuration, management, and service tasks regardless of the configuration that is available from the public cloud portal. ► Helps administrators use storage assets and volumes more efficiently. ► IBM Storage Insights and IBM Storage Protect provide more capabilities to manage and monitor capacity and performance.
Dynamic data migration.	► Migrates data among volumes or LUNs without taking applications that use that data offline. ► Manages and scales storage capacity without disrupting applications.
Advanced network-based copy services.	► Copies 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.

Feature	Benefits
Thin provisioning, data reduction pools (DRPs), and snapshot replication.	► Reduces volume requirements by using storage only when data changes. ► Improves storage administrator productivity through automated on-demand storage provisioning. ► Supports thin-provisioning in standard and DRPs, which reduces capacity requirements by using storage only when data changes. ► Supports DRPs with native data reduction features, such as host unmap and reclaiming usable capacity. ► Supports deduplicated and or compressed volumes in DRPs for more capacity savings. Snapshots are available on lower-tier storage volumes.
IBM Storage Protect Snapshot application-aware snapshots.	► Performs near-instant, application-aware snapshot backups, with minimal performance impact for IBM Db2®, Oracle, SAP, Microsoft SQL Server, and Microsoft Exchange. ► Provides advanced, granular restoration of Microsoft Exchange data.
Native IP replication.	► Uses embedded WAN optimization technologies, such as data compression and acceleration. ► Reduces network costs and speeds replication cycles, improving accuracy of remote data.
IBM Storage Connect Cloud Storage Management.	Manages container storage in Kubernetes.
IBM Safeguarded Copy.	► Designed to protect critical data from cyberthreats, including ransomware, human error, and software corruption. ► Creates immutable point-in-time (PiT) copies of critical data on a schedule. ► Copies cannot be seen, modified, or deleted. ► Copies are used to restore primary data in as little as minutes but no more than hours depending on the amount of data.
Policy-Based Replication (PBR).	► Automatically deploy and manages replication to volume groups (VGs). ► Policies can be created and applied to the VG. Any volume that is assigned membership to the VG automatically inherits the policy.

Note: The IBM Spectrum Virtualize for Public Cloud “Upgrade capabilities from v8.3.1 to current 8.5.4” feature is not supported by this release.

Some of these features are planned for future releases and are prioritized for implementation based on customer feedback.

1.3 IBM Spectrum Virtualize for Public Cloud on AWS

The current release of IBM Spectrum Virtualize for Public Cloud available for AWS is v8.5.4. Block virtualization further uses public cloud infrastructure for various types of workload deployments whether it is new or traditional.

The following features are supported on the AWS infrastructure:

- Data replication with any IBM Storage Virtualize appliance-based product and between public clouds supporting IBM Spectrum Virtualize for Public Cloud.
- Volume Group Snapshots in the cloud.
- Common Management: IBM Storage Virtualize GUI.
- Deployment in any AWS region.
- Encryption at rest by using Amazon Elastic Block Store (EBS) encrypted volumes.
- Data redundancy with volume mirroring.
- Automated block-level storage tiering by using Easy Tier.
- Scales on demand by using thin-provisioned volumes and paying for AWS storage as you grow.
- Supports thin provisioning in both standard pools and DRPs, which reduces capacity requirements by using storage only when data changes.
- Supports DRPs with native data reduction features, such as host unmap and reclaiming usable capacity.
- Supports deduplicated and or compressed volumes within DRPs for more capacity savings.

The AWS infrastructure is an established platform for today's cloud computing needs. Deploying the IBM Spectrum Virtualize for Public Cloud platform further enriches the capabilities of the AWS cloud infrastructure.

Figure 1-1 shows the general layout of IBM Spectrum Virtualize for Public Cloud on AWS.

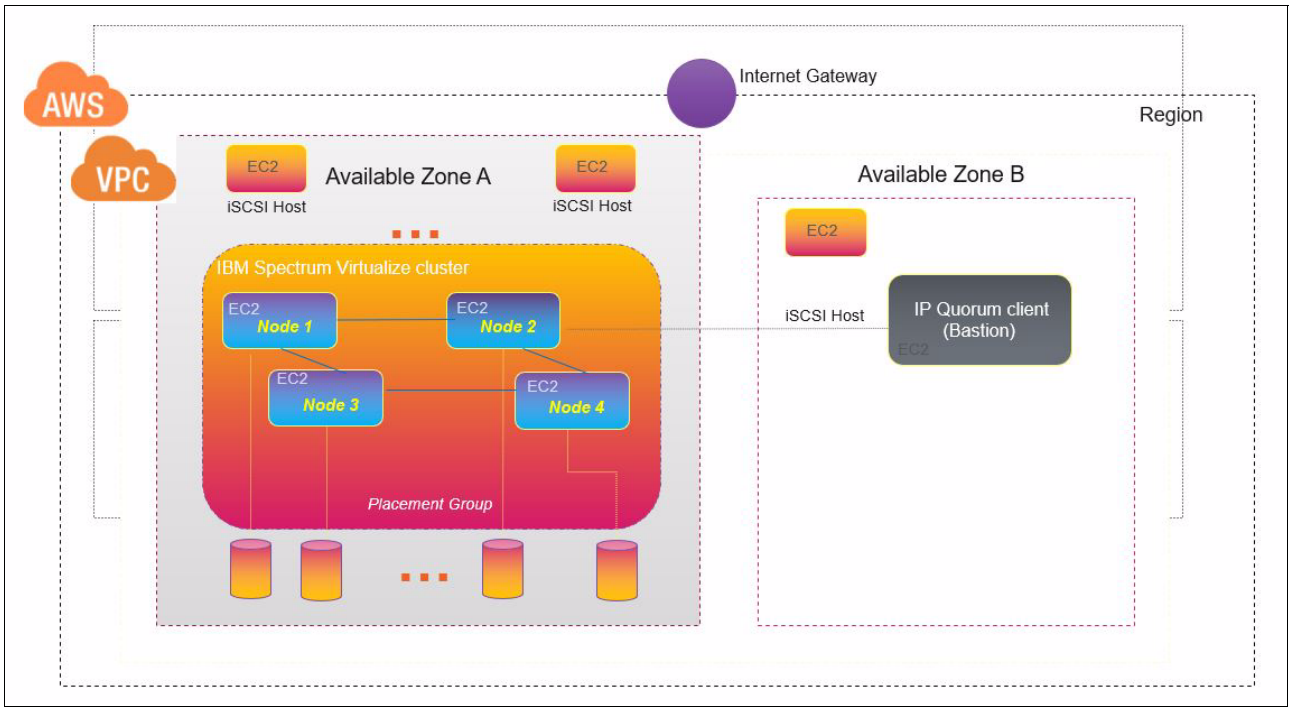

Figure 1-1 High-level architecture of IBM Spectrum Virtualize for Public Cloud on AWS

In AWS, the Amazon EBS storage is directly attached to the IBM Spectrum Virtualize node instances that comprise a single node pair (or I/O group) to provide a shared storage pool that is used by IBM Spectrum Virtualize for Public Cloud. At the time of writing, the following Amazon EBS types are supported:

- General-purpose solid-state drive (SSD) (gp2 or gp3)
- Provisioned IOPS SSD (io1)
- Throughput-optimized hard disk drive (HDD) (st1)
- Throughput-optimized HDD (sc1)

1.4 IBM Spectrum Virtualize for Public Cloud on Azure

IBM Spectrum Virtualize for Public Cloud 8.5.x is available for Azure. Block virtualization further uses public cloud infrastructure for various types of workload deployments, whether it is new or traditional.

The following features are supported on the Azure infrastructure:

- Data replication with any IBM Spectrum Virtualize appliance-based product and between public clouds providers.
- FlashCopy snapshots in the cloud.
- Common Management: IBM Spectrum Virtualize GUI.
- Deployment in any Azure region.
- Data redundancy with volume mirroring.
- Automated block-level storage tiering by using Easy Tier.

- Scale on demand by using thin-provisioning volumes and paying for Azure storage as you grow.
- Supports thin provisioning in both standard pools and DRPs, which reduces capacity requirements by using storage only when data changes.
- Supports DRPs with native data reduction features, such as host unmap and reclaiming usable capacity.
- Transparent Cloud Tiering (TCT) allows for fulls copies of data to be protected by logical air gaps, allowing for restoration on another IBM Storage Virtualize based appliance or even the cloud-only environment supporting IBM Spectrum Virtualize for Public Cloud.
- Supports deduplicated and or compressed volumes within DRPs for more capacity savings.
- Safeguarded Copy protects critical data by creating immutable PiT copies of selected volumes that cannot be seen, modified, or deleted, and can be used to recover from lost data in minutes to hours, including from ransomware.

The Azure infrastructure is an established platform for today's computing needs. By deploying the IBM Spectrum Virtualize for Public Cloud platform, IBM Spectrum Virtualize further enriches the capabilities of the cloud infrastructure.

Figure 1-2 shows the general layout of IBM Spectrum Virtualize for Public Cloud on Azure.

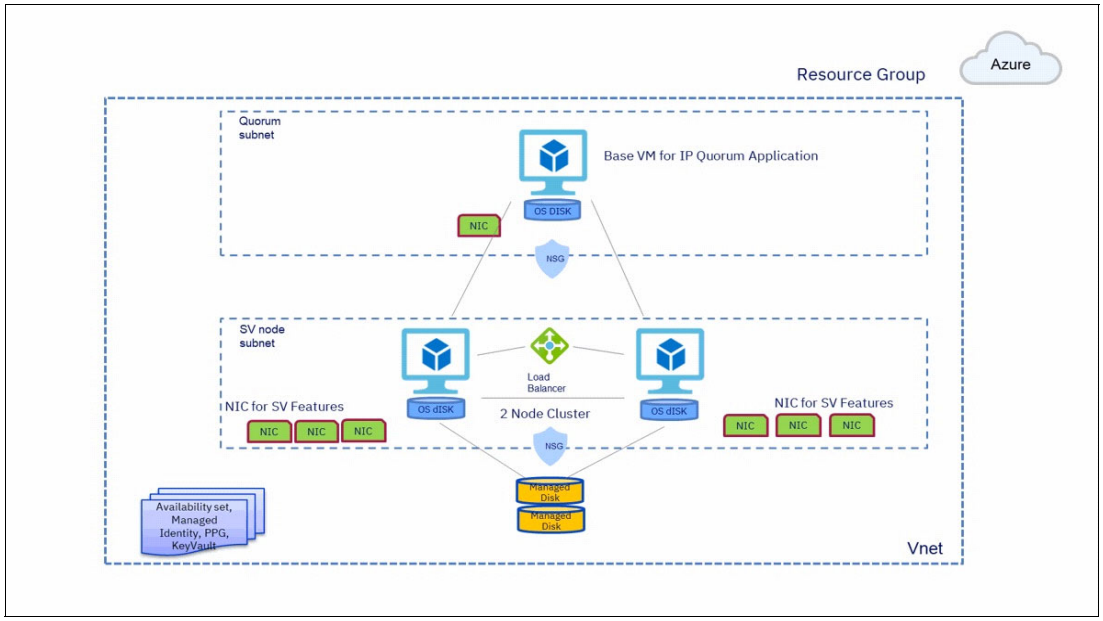

Figure 1-2 High-level architecture of IBM Spectrum Virtualize for Public Cloud on Azure

1.5 What is new with IBM Spectrum Virtualize for Public Cloud 8.5.4

Since the previous release of IBM Spectrum Virtualize for Public Cloud, IBM continues to enhance its capabilities. Here are some of the major key enhancements that are available with this updated version:

- Support for Safeguarded Copy 2.0 allowing for logical air-gapping of critical data from on-premises to AWS and or AWS to on-premises with simple recovery.
- Support for Snapshot (FlashCopy 2.0) helping cloud storage administrator to design strategy to make backups automatically and restore data when needed.
- Support for Multifactor Authentication providing simple authentication and single sign-on (SSO) during login to a IBM Spectrum Virtualize for Public Cloud cluster GUI.
- Support for PBR, which simplifies configuring, managing, and monitoring replication between two I/O Groups.
- Support for TCT allowing for backup volumes to not only be tiered restored from on-premises to AWS Simple Storage Service (S3)/Azure Blobs, but from AWS/Azure to on-premises and AWS to Azure.
- Support for Replication over IPsec, which provides secure replication between IBM Spectrum Virtualize for Public Cloud clusters with certification-based authentication. Provides an extra layer of security for data-in-flight for hybrid cloud and all-in-cloud user cases.
- Support for non-disruptive security and software fixes.

Announcement letter: For more information, see the announcement letter for IBM Spectrum Virtualize for Public Cloud 8.5, which is available at this web page.

2

Typical use cases for IBM Spectrum Virtualize for Public Cloud

This chapter describes four use cases for IBM Spectrum Virtualize for Public Cloud and includes the following topics:

- Deploying whole IT services in the public cloud
- Disaster recovery
- IBM FlashCopy in the public cloud
- Safeguarded Copy and Safeguarded snapshots
- Workload relocation into the public cloud

2.1 Deploying whole IT services in the public cloud

Companies are approaching and using public cloud services from multiple angles. Users that are rewriting and modernizing applications for cloud complement those users who are looking to move to cloud-only new services or to extend existing IT into a hybrid model to address quickly changing capacity and scalability requirements.

The delivery models for public cloud are available in the following general as-a-service categories:

- *Software as a Service*: SaaS provides the greatest level of abstraction in which the user interacts only with the software. IBM Storage Insights is such an example where clients are not at all involved with any of the back-end components.
- *Infrastructure as a Service (IaaS)*: In IaaS, server instances and even bare metal servers are provisioned on a subscription basis. IBM Cloud Classic Infrastructure is such an example. Network components also can be discretely subscribed, such as virtual private network (VPN) gateways.
- *Platform as a Service (PaaS)*: PaaS is the intermediate and the most common cloud environment. Microsoft Azure and Amazon Web Services (AWS), and Red Hat OpenShift are examples. In PaaS, virtualization is managed by the provider and abstracted from the user.

The workload deployment is composed of two major use cases, as shown in Figure 2-1:

- *Hybrid cloud*: The integration is between the off-premises public cloud services with an on-premises IT infrastructure.
- *Cloud-native*: The full application's stack is moved to cloud as SaaS, PaaS, IaaS, or as a combination of the three delivery models.

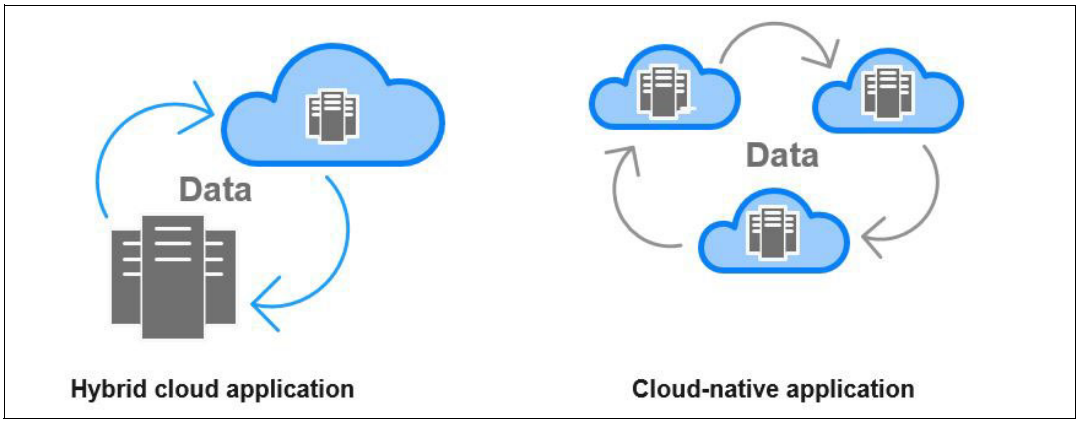

Figure 2-1 The two major deployment models for public cloud

Cloud-native implementations (that is, whole IT services that are 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, final architecture, and roles and responsibilities depend on SaaS, PaaS, or IaaS usage.

Within the IaaS domain, the transparency of cloud services is the highest because the user's visibility (and responsibility) into the application stack is much deeper compared to the other delivery models. Conversely, the *burden* for its deployment is higher because all the components must be designed from the server up.

At the time of this writing, IBM Spectrum Virtualize for Public Cloud is framed only within the IaaS cloud delivery model so that the user can interact with their storage environment as they did on-premises, which provides more granular control over performance.

2.1.1 Business justification

A stand-alone workload or an application, with few on-premises dependencies, relatively low-performance requirements, and that is not processing highly regulated data, represents a good fit for a cloud-native deployment. The drivers that motivate businesses toward cloud-native deployment are financial, such as decreasing capital expenditure (CapEx) and operating expenditure (OpEx), optimizing or eliminating resource management and controls against hidden or *shadow* IT resources. Other benefits are more flexibility and scalability, and streamlined flow in delivering IT service because of the global footprint of cloud data centers.

At its core, the cloud environment is highly focused on standardization and automation. Therefore, the full IBM Spectrum of features and customization that are available in a typical on-premises or outsourcing deployment might not be natively available in the cloud catalog.

Nevertheless, the client does not lose performance and capabilities when deploying a cloud-native application. In this context, the storage virtualization with IBM Spectrum Virtualize for Public Cloud enables the IT staff to maintain the technical capabilities and skills to deploy, run, and manage highly available (HA) 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 that the client uses in their traditional IT environment.

In a hybrid multicloud environment, the orchestration of the infrastructure requires multiple entities that are tightly integrated with each other and smartly respond to administrator or user needs, and that is where a software-defined environment (SDE) has an important role in the overall orchestration.

Integration between service delivery, management, orchestration, automation, and hardware systems is becoming a requirement to support the emergence of SDEs. For SDEs to provide their benefits, they must understand and manage all the components of the infrastructure, including storage, and that makes software-defined storage (SDS) more relevant and important.

The capability of collecting the information from storage systems and providing a simplified multicloud deployment across IBM Storage Systems is provided by IBM Spectrum Connect. IBM Spectrum Virtualize for Public Cloud on Microsoft Azure or Amazon AWS and IBM Spectrum Connect integrate vRealize Orchestrator with vRealize Automation, which takes the service around infrastructure beyond orchestration.

By integrating the Advanced Service Designer feature of vRealize Automation with vRealize Orchestrator, an organization can offer anything as a service (XaaS) to its users. By using the XaaS feature of vRealize Automation, IBM Spectrum Virtualize Storage System and IBM Spectrum Virtualize for Public Cloud on Azure or Amazon AWS can be delivered as SaaS in a multicloud environment, whether it is deployed in private cloud or a public cloud multicloud environment.

2.1.2 Highly available deployment models

The architecture is directly responsible for an application's reliability and availability if a component failure (hardware and software) occurs. When an application is fully hosted on cloud, the cloud data center becomes the primary site (production site). Cloud deployment does not automatically guarantee 100% uptime, that the backups are available by default, or that the application is automatically replicated between different sites.

These security, availability, and recovery features are often incorporated into the SaaS model. They might be partially provided in the PaaS model. However, in the IaaS model, they are *entirely* the customer's responsibility.

Having reliable cloud deployments means that the service provider must meet the required service-level agreement (SLA), which helps ensure service availability and uptime. Companies that use a public cloud IaaS can meet required SLAs by implementing HA solutions and duplicating the infrastructure in the same data center or in two data centers to maintain business continuity if there are failures.

If business continuity is not enough to reach the requirements of the SLA, disaster recovery (DR) implementations, which split the application among multiple cloud data centers (usually with a distance of at least 300 Km [186.4 miles]) prevent failure in a major disaster in the organization's main campus.

The HA deployment models for an application that is fully deployed on public cloud are summarized as follows:

- HA cloud deployment on a single primary site

  All the solution's components are duplicated (or more) within the same data center. This solution continues to function because there are not single points of failure (SPOFs), but it does not function if the data center is unavailable.

- HA cloud deployment on multiple sites

  The architecture is split among multiple cloud data centers from multiple cloud providers to mitigate the failure of an entire data center or provider, or spread globally to recover the solution if major disaster affects the campus.

Highly available cloud deployment on a single primary site

When fully moving an application to a cloud IaaS that is the primary site for service delivery, a reasonable approach is implementing at least a HA architecture. Each component (servers, network components, and storage) is redundant to avoid SPOF.

Within the single primary site deployment, storage is deployed as native cloud storage. By using the public cloud catalog storage, users can leverage the intrinsic availability (and SLAs) of the storage service, which is this case is Microsoft Azure managed disk (MDisk) or Elastic Block Store (EBS) in Amazon AWS.

When IBM Spectrum Virtualize for Public Cloud is deployed as clustered pair of Azure virtual machine (VM) or Amazon Elastic Compute Cloud (EC2) instances, it mediates between the Cloud Block Storage and the workload hosts. In the specific context of single-site deployment, IBM Spectrum Virtualize for Public Cloud supports extra features that enhance the public cloud block-storage offering.

At the storage level, IBM Spectrum Virtualize for Public Cloud resolves some limitations because of the standardized model of 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. It is a 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. Another advantage of IBM Spectrum Virtualize Public Cloud is that it integrates with our Storage Insights product to provide advance monitoring, reporting, and alerting by using data that is gathered from the IBM Spectrum Virtualize instances.

The benefits of an IBM Spectrum Virtualize for Public Cloud single site deployment are listed in Table 2-1.

Table 2-1 Benefits of IBM Spectrum Virtualize for Public Cloud single site deployment

Feature	Benefits
Single point of control for cloud storage resources	Designed to increase management efficiency and to help to support application availability.
Pools the capacity of multiple storage volumes	► 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 a storage volume offering.
Manages tiered storage	► 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 and classes.
Easy-to-use IBM FlashSystem family management interface	► Has a single interface for storage configuration, management, and service tasks regardless of the configuration that is available from the public cloud portal. ► Helps administrators use storage assets and volumes more efficiently. ► Has IBM Spectrum Control Insights and IBM Spectrum Protect for extra capabilities to manage capacity and performance.
Dynamic data migration	► Migrates data among volumes and LUNs without taking applications that use that data offline. ► Manages and scales storage capacity without disrupting applications.
Advanced network-based copy services	► Copy data across multiple storage systems with IBM FlashCopy. ► Replicate or migrate data across metropolitan and global distances as needed to create fault-tolerant storage solutions between multiple data centers.
Thin provisioning, snapshot replication, and data deduplication	► Snapshots reduce capacity requirements by using storage only when data changes. ► Improves storage administrator productivity through automated on-demand storage provisioning. ► Snapshots are available on lower-tier storage volumes. ► Data deduplication further enhances capacity savings when back-end storage is gathered into a data reduction pool (DRP).
IBM Spectrum Protect Snapshot application-aware snapshots	► Perform near-instant and application-aware snapshot backups, with minimal performance impact for IBM Db2, Oracle, SAP, VMware, Microsoft SQL Server, and Microsoft Exchange. ► Provide advanced and granular restoration of Microsoft Exchange data.

Feature	Benefits
Third-party native integration	Integration with VMware vRealize.
Safeguarded Copy	This new IBM Spectrum Virtualize function provides a valuable ransomware mitigation solution, especially when combined with an implementation of IBM Spectrum Virtualize for Public Cloud to achieve a physically isolated (air-gapped) cybervault.

Highly available cloud deployment on multiple sites

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 data center. The primary and secondary data centers can be deployed as:

- *Active-active*: The secondary site is always running and synchronously aligned with the primary site.[1]
- *Active-passive:* The secondary site is always running but asynchronously replicated (with a specific recovery point objective [RPO]) or running only for specific situations, such as acting as a recovery site or test environment. Storage is always active and available for data replication.

The active-passive configuration is usually the best fit for many cloud use cases, including DR, as described in 2.2, “Disaster recovery” on page 19. The ability to provision compute resources on demand in a few minutes with only the storage that is provisioned and aligned with a specific RPO is a huge driver for a cost-effective DR infrastructure, and lowers the total cost of ownership (TCO).

The replication among multiple cloud data centers is no different from the traditional approach, except for the number of available tools in the cloud. Although solutions that are based on hypervisor or application-layer replication, such as VMware, Veeam, and Zerto, are available in the public cloud, storage-based replication is still the preferable approach if the environment is heterogeneous (virtual servers, bare metal servers, multiple hypervisors, and so on).

Active-passive asynchronous mirroring that uses Global Mirror with Change Volumes (GMCV) provides a minimum RPO of 2 minutes (the Change Volume [CV] cycle period ranges is 1 minute - 1 day, and a best practice is setting the cycle period to be half of the RPO), and can replicate a heterogeneous environment.

IBM Spectrum Virtualize 8.5.2 introduces Policy-Based Replication (PBR) which streamlines and simplifies the replication configuration process, and it delivers improvements in efficiency, performance, manageability, and scalability for replication between IBM Spectrum Virtualize systems. The RPO can now be specified explicitly, and replication automatically attempts to stay as close to synchronous as the replication link and compute resources allow. If resources become constrained, PBR automatically adjusts the replication frequency to increasing intervals. If this interval exceeds the specified RPO alert threshold, an alert is sent to notify the storage administrator that replication has fallen outside of the RPO.

[1] At the time of writing, IBM Spectrum Virtualize HA multi-site topologies, such as IBM HyperSwap and Enhanced Stretch Cluster, are not supported by IBM Spectrum Virtualize Public Cloud.

2.2 Disaster recovery

Customers have adopted DR strategies to harness and secure proliferating data in their environment and infrastructure workloads in a cost-effective manner when a HA level of RPO of 0 is not a business requirement.

Technology is only one crucial piece of a DR solution, and not the one that always dictates the overall approach.

This section describes the DR approach and benefits of IBM Spectrum Virtualize for Public Cloud on Azure or Amazon AWS.

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

2.2.1 Business justification

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

Table 2-2 Drivers, challenges, and capabilities that are provided by IBM Spectrum Virtualize for Public Cloud

Adoption drivers	Challenges	IBM Spectrum Virtualize for IBM public cloud capabilities
The promise of reduced operational expenditure (OpEx) and CapEx	► Hidden costs ► Availability of data when needed	► Optimized for Cloud Block Storage ► IBM Easy Tier solution to optimize the most valuable storage usage, which maximizes Cloud Block Storage performance ► Thin provisioning to control the storage provisioning ► Snapshots feature for backup and DR solution ► HA clusters architecture
Bridging technologies from on-premises to cloud	► Disparate Infrastructure: How can my on-premises production data be readily available in the cloud in a disaster?	► Any to any replication ► Supporting over 400 different storage devices (on-premises), including iSCSI on-premises and when deployed in cloud

Adoption drivers	Challenges	IBM Spectrum Virtualize for IBM public cloud capabilities
Using the cloud for backup and DR	► Covering virtual and physical environments ► Solutions to meet a range of RPO/RTO needs	► A storage-based, serverless replication with options for low RPO/RTO: – Global Mirror for Asynchronous replication with an RPO close to “0” (not recommended for Public Cloud) – Metro Mirror for Synchronous replication (not supported for Public Cloud) – GMCV for Asynchronous replication with a tunable RPO (a best practice for Public Cloud deployments) – PBR for more streamlined configuration, improved performance and scalability, and RPO-based alerting and self-adjusting cycling rates

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

- Can be implemented at several locations in Microsoft Azure or Amazon AWS and installed by using their respective marketplaces.
- Deploys on an Azure VM or Amazon AWS EC2 instance.
- Offers data replication with the IBM FlashSystem family, IBM V9000, IBM SAN Volume Controller, or VersaStack, and other public cloud IBM Spectrum Virtualize instances (IBM Cloud®, Amazon AWS, or Microsoft Azure).
- Supports two node clusters in Microsoft Azure and up to four node clusters in AWS.
- Offers data services for Azure MDisks and Amazon EBS.
- Offers common management with the IBM Spectrum Virtualize GUI with full admin access and a dedicated instance.
- No incoming data transfer cost.
- Replicates between two Azure or AWS locations.
- Replicates between IBM Spectrum Virtualize on-premises and IBM Spectrum Virtualize for Public Cloud on Azure or Amazon AWS.

2.2.2 Two 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 DR for “Any to Any”.
- IBM Spectrum Virtualize for Public Cloud solution on Azure cloud DR, as shown in Figure 2-2 on page 21. The design for Cloud DR on IBM Spectrum Virtualize for Public Cloud in Amazon AWS is identical to that of Azure.

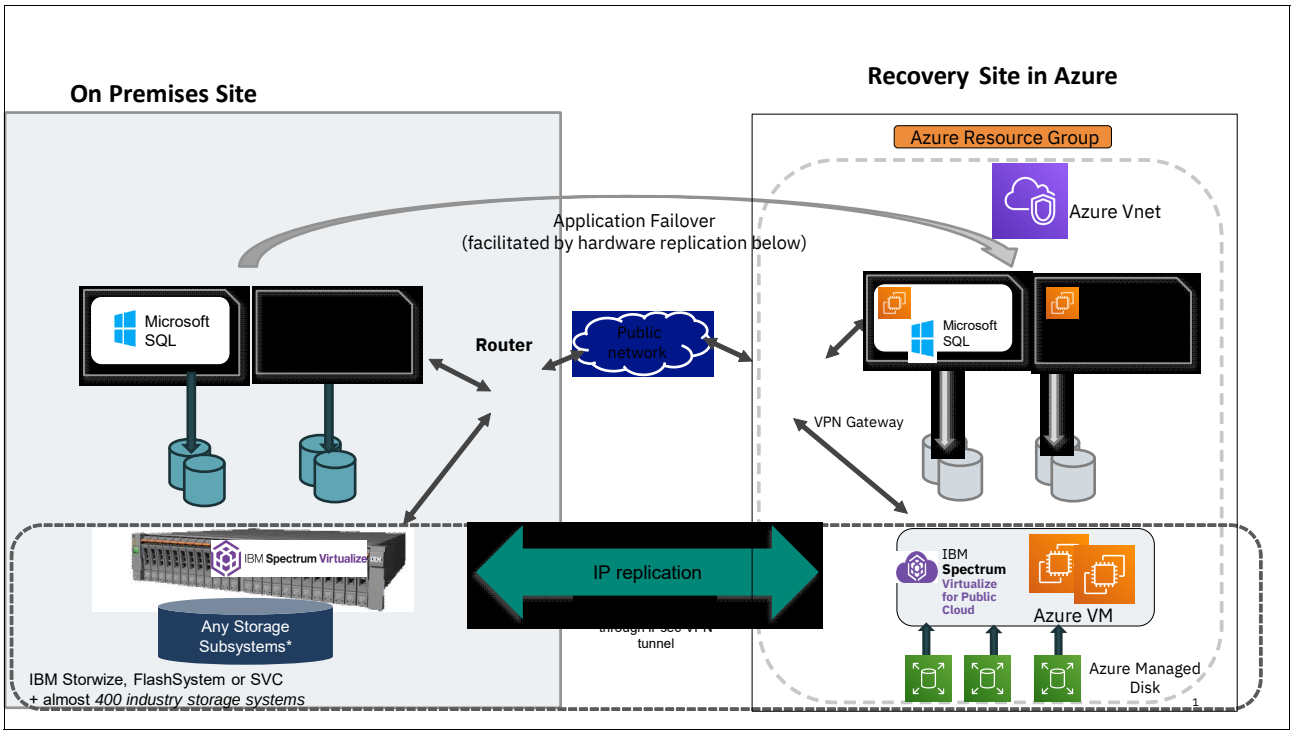

Figure 2-2 IBM Spectrum Virtualize for Public Cloud on Azure cloud DR solution

As shown in Figure 2-2, a customer can deploy a storage replication infrastructure in a public cloud by using IBM Spectrum Virtualize for Public Cloud.

This scenario includes the following scenarios:

- Primary storage is in the customer's physical data center. The customer has an on-premises IBM Spectrum Virtualize solution that is installed.
- Auxiliary storage sits on the DR site, which can be an IBM Spectrum Virtualize cluster running in the public cloud.
- The virtual IBM Spectrum Virtualize cluster manages the storage that is provided by Azure MDisks or Amazon AWS EBS volumes.

A replication partnership that uses GMCV or PBR is established between an on-premises IBM Spectrum Virtualize cluster or IBM FlashSystem solution and the virtual IBM Spectrum Virtualize on Azure or AWS cluster to provide DR.

When talking about DR, understand that IBM Spectrum Virtualize for Public Cloud is an important piece of a complex solution that has some prerequisites and best practices that must be considered.

2.3 IBM FlashCopy in the public cloud

The IBM FlashCopy function in IBM Spectrum Virtualize can perform a point-in-time (PiT) 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.

2.3.1 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
  - Auditing purposes and data mining
  - Quality assurance
- Rapidly creating copies of replication targets for testing data integrity.

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

2.3.2 FlashCopy mapping

The association between the source volume and the target volume is defined by a FlashCopy map. The FlashCopy map can have three different types (as defined in the GUI), four attributes, and seven different states.

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

- Snapshot

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

- Clone

  Sometimes referred to as *one time full copy*. A PiT copy of a volume with a 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, which is achieved with a copy and clean rate greater than zero and an autodelete flag; therefore, no cleanup of the map is necessary after the background copy is finished.

- Backup

  Sometimes referred to as an iterative incremental. A backup FlashCopy mapping consists of a PiT full copy of a source volume, plus periodic increments or “deltas” of data that changed between two points in time.

  This mapping is where the copy and clean rates are greater than zero, no autodelete flag is set, and you use an incremental flag to preserve the bitmaps between activations so that only the deltas since the last “backup” must be copied.

  It is named such as the most typical use case is with backup processes that cause heavy reads and so a full copy is made to insulate the primary volume against those heavy reads. Also, because backups occur periodical (typically daily), the incremental flag allows only the deltas between refresh to be copied.

The FlashCopy mapping has four property attributes (clean rate, copy rate, autodelete, and incremental) and seven different states. Users can perform the following tasks on a FlashCopy mapping:

- Create: Define a source and a target, and set the properties of the mapping.
- Prepare: The system must be prepared before a FlashCopy copy starts. It basically flushes the cache and makes it “transparent” for a short time so that 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 (by the system or 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, which 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 PiT 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. For scalability and performance reasons, FlashCopy source and target volumes and maps might need to be aligned in the same I/O group and possibly the same preferred node.

For more information, see 6.2.4 “FlashCopy planning considerations” of *IBM FlashSystem Best Practices and Performance Guidelines for IBM Spectrum Virtualize Version 8.4.2*, SG24-8508.

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

All FlashCopy operations occur on FlashCopy mappings. FlashCopy does not alter source volumes. Multiple operations can occur at the same time on multiple FlashCopy mappings by using consistency groups.

2.3.3 Consistency groups

To overcome the issue of dependent writes across volumes and create a consistent image of the client data, perform a FlashCopy operation on multiple volumes as an atomic operation. To accomplish this task, IBM Spectrum Virtualize supports the concept of consistency groups.

Consistency groups preserve PiT 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 only one mapping exists in the consistency group. If a FlashCopy mapping is not part of any consistency group, it is referred to as *stand-alone*.

2.3.4 Crash-consistent copy and host considerations

FlashCopy consistency groups do not provide application consistency. They ensure only that volume points-in-time are consistent between volumes.

Because FlashCopy is at the block level, you must understand the interaction between your application and the host operating system. From a logical standpoint, 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 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, which is required following a host crash. FlashCopy copies that are crash-consistent often can be used after the file system and application recovery procedures.

This concept is shown in Figure 2-3, where in-flight I/Os in cache buffers (if unflushed) are not in the volume; therefore, they are not be captured in the FlashCopy.

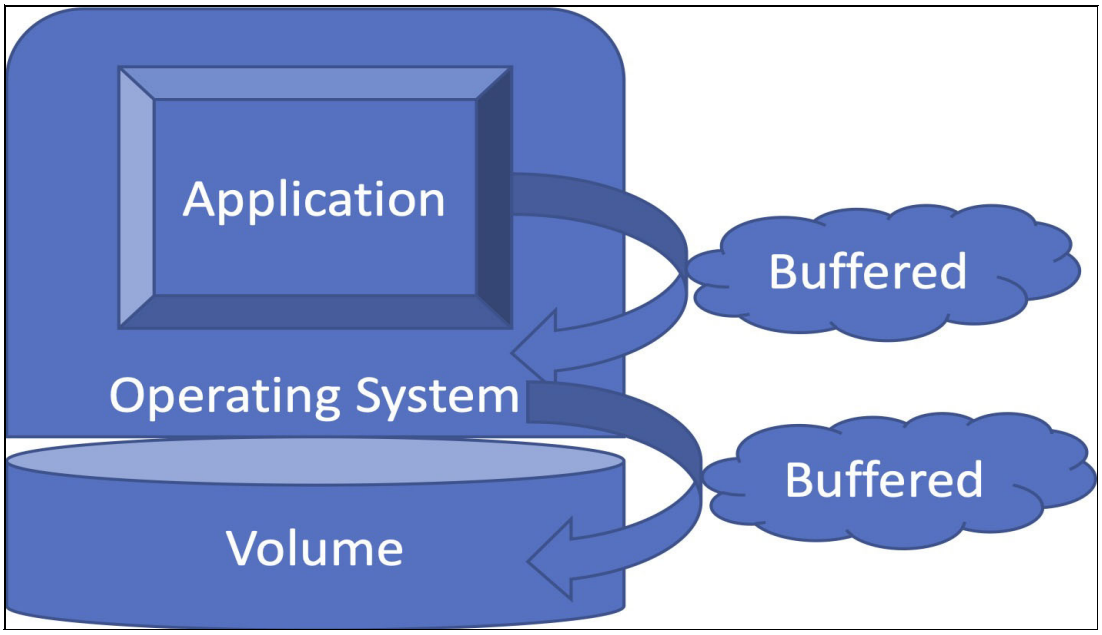

Figure 2-3 Buffered I/Os are lost if unflushed

Various operating systems and applications provide facilities to stop I/O operations and ensure that all data is flushed from the host cache. If these facilities are available, they can be used to prepare 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, Discard any data that is held on the host operating system (or application) caches for the target volumes. The simplest way to ensure that no data is held in these caches is to unmount the target volumes before the FlashCopy operation starts.

Note: From a practical perspective, 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 usually 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 opposed to flushing data from the application and backing database, which is a best practice because it is safer. However, this method can be used when facilities do not exist or your environment includes time sensitivity.

2.3.5 Volume Group Snapshots

In addition to PBR, IBM Spectrum Virtualize and IBM FlashSystem 8.5.2 introduced an improved way of making Copy-on-Write PiT copies. *Volume Group Snapshots* are more flexible and scalable than the snapshot presets (**`-cleanrate`** and **`-copyrate 0`**) for FlashCopy maps. By using the concept of volume groups (VGs) and snapshot policies, Volume Group Snapshots improve upon FlashCopy in the following ways:

- Simplicity: Creating or even specifying a target volume is no longer necessary. Simply create a VG and place volumes in it or create a volume or volumes directly in an existing group. When a snapshot is no longer needed, it can be deleted from the GUI with a single operation. You no longer need to delete the FlashCopy map and then the volume.
- Scheduling: When you add a function that did not exist previously, there is a now an internal scheduler that can create snapshots at a preset frequency, keeping them for a set amount of time for continuous protection.
- Scalability: Currently, on-premises IBM Spectrum Virtualize and IBM FlashSystem Volume Group Snapshots can scale up to 19,999 snapshots (on the IBM FlashSystem 9500 and SV3 SAN Volume Controller nodes) and 18 GiB of bitmap space.

Best practice: At the time of writing, one key limitation for Volume Group Snapshots exists, but it will be addressed: Snapshots can be recovered only to clones that are contained in a new VG. These clones (thin or "thick") are the volumes that you mount to a host. You cannot restore a snapshot back to the original volumes. The restoration to production function is in development.

2.4 Safeguarded Copy and Safeguarded snapshots

IBM Spectrum Virtualize 8.4.2.0 introduced the new Safeguarded child pool function to provide a powerful cyber resiliency solution to IBM Spectrum Virtualize and IBM FlashSystem users. Combining remote replication and FlashCopy with Safeguarded Copy, on-premises workloads can be protected from ransomware and other data corruption attacks with a truly air-gapped solution in the public cloud.

Volume Group Snapshots further enhances this function with Safeguarded snapshots by providing greater simplicity and flexibility.

2.4.1 Business justification

The regulatory and business justifications for this use case are clear and widely reported in the news of high profile cases of ransomware attacks crippling business processes and sometimes threatening lives as healthcare organizations were attacked in the midst of the COVID-19 global pandemic.

2.4.2 Solution design

As shown in Figure 2-4, on-premises primary volumes are replicated through an IP network to a IBM Spectrum Virtualize in Public Cloud instance on Azure. The design for Amazon AWS is identical in nature to the implementation on Azure in this regard.

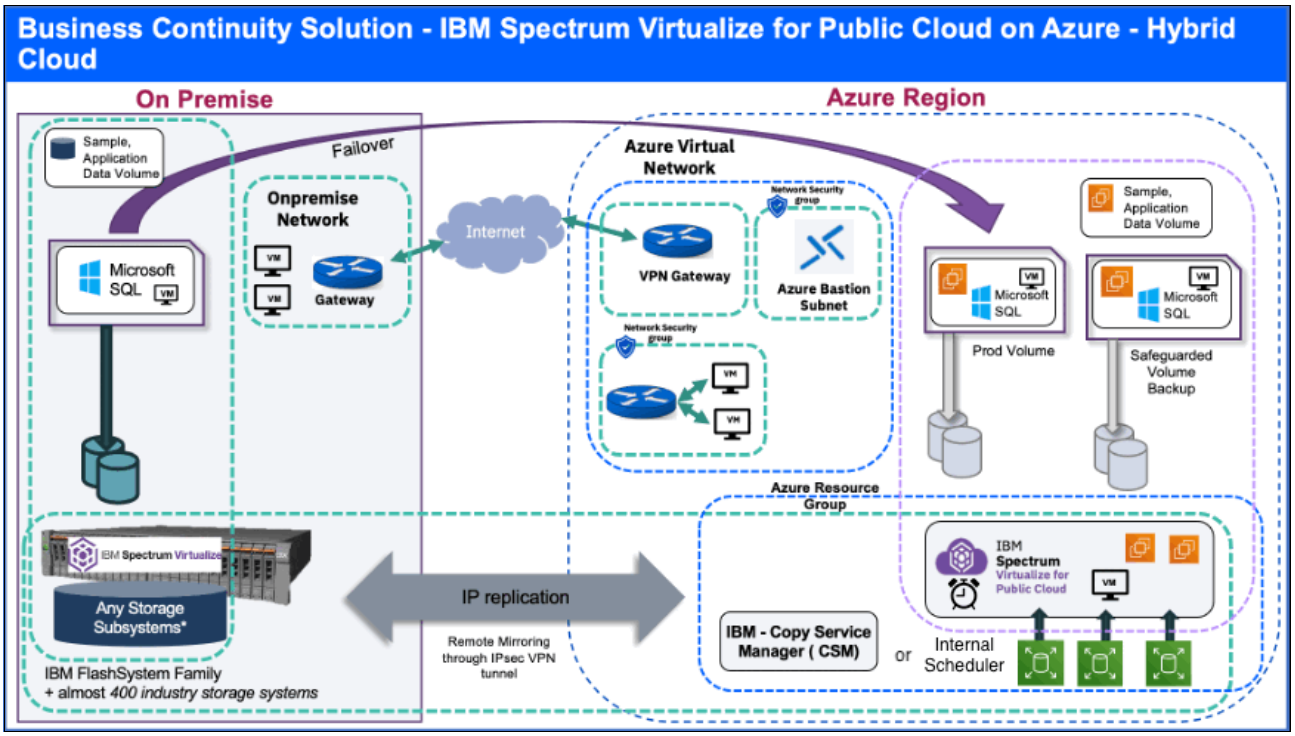

Figure 2-4 Safeguarded Copy

The pool from which those destination or auxiliary volumes are created was configured with a Safeguarded child pool, a VG was set up to contain those volumes, and a Safeguarded policy was assigned to the VG that governs the frequency and retention duration for the Safeguarded Copy copies.

Copy Services Manager (CSM) is installed on-premises, or ideally in Azure or AWS. It is configured to communicate with the IBM Spectrum Virtualize for Public Cloud instance to convert the policy into scheduled actions. CSM also provides an orchestration portal for managing recovery and restoration from the Safeguarded Copy copies to a recovery or original replication destination volume.

Because CSM is primarily a replication orchestration tool, it is positioned to manage the replication of the recovered or restored data back to the primary site. For more information, see *IBM FlashSystem Safeguarded Copy Implementation Guide*, REDP-5654.

2.4.3 Component summary

IBM Spectrum Virtualize 8.4.2.0 features the following components:

- Safeguarded Child Pool: A new feature that provides a region of a storage pool for making non-modifiable copies of volumes in that pool to guard against malicious or accidental data corruption.
- VG: A new container type that provides a way to group a set of volumes to which a Safeguarded policy is applied and acted upon in a crash consistent manner. When Safeguarded copies are taken for volumes in a VG, a consistency group is automatically created to keep those volumes crash consistent with one another.
- Safeguarded Policy: A new object type that governs the frequency and retention duration for Safeguarded copies. Three default policies and other custom policies can be created by using the **`mksafeguardedpolicy`** command-line interface (CLI) command. The three default policies are listed in Table 2-3.

Table 2-3 Default policies

Policy	Frequency	Retention
`predefinedsgpolicy0`	6 hour	7 days
`predefinedsgpolicy1`	1 week	30 days
`predefinedsgpolicy2`	1 month	365 days

- CSM: Application that has long existed as a replication and PiT copy orchestration tool for IBM storage (IBM Spectrum Virtualize, IBM DS8000, and IBM XIV). With version 6.2, CSM integrates with IBM Spectrum Virtualize 8.4.2.0 to periodically scan for VGs with volumes and a Safeguarded policy that is associated. Upon detection of such, CSM creates objects within its own framework (sessions, copy sets, and scheduled tasks) to run on the policy and create Safeguarded backups with the frequency that is stipulated in the policy.

  Moreover, it allows for the orchestration of recovery (create a copy of a Safeguarded Copy onto a new volume) and restoration (copy data back to the source volume from a Safeguarded Copy).

2.4.4 Safeguarded snapshots

Safeguarded snapshots provide several improvements over Safeguarded Copy:

- Flexibility: You no longer need a Safeguarded child pool. Because Volume Group Snapshots are no longer volumes and FlashCopy maps, volumes from any storage pool can be protected with immutable snapshots. Furthermore, recovery volumes no longer must come from the parent pool of the original volume. These clones of the Safeguarded snapshots can be created from any storage pool in the system.
- Scheduling and management: CSM is no longer required for scheduling and Safeguarded snapshot management. CSM and Copy Data Manager can both be used, but with the introduction of the internal scheduler, both the scheduling of backups and the creation and mapping of recovery volumes can now be done within the IBM Spectrum Virtualize interface.

2.5 Workload relocation into the public cloud

In this section, a use case for IBM Spectrum Virtualize for Public Cloud is described in which an entire workload segment is migrated from a customer's enterprise into the cloud. Although the process for relocating a workload into the cloud by using IBM Spectrum Virtualize can use only Remote Copy, other mechanisms are available that can accomplish this task.

2.5.1 Business justification

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

An ideal case regarding a hybrid cloud solution is the relocation of a specific segment of the environment that is well suited, such as development. Another might be a specific application group that does not require the regulatory isolation or low response time integration with on-premises applications.

Although performance might be a factor, do not assume that cloud deployments automatically create a diminished performance. Depending on the location of the cloud service data center and the intended audience for the migrated service, the performance can conceivably be superior to on-premises pre-migration.

In summary, moving a workload into the cloud might provide similar functions with better economies because of scaling physical resources in the cloud provider. Moreover, the cost of services in the cloud is structured, measurable, and predictable.

2.5.2 Data migration

Several methods are available for performing data migrations to the cloud, including the following 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

The first method was described in 2.3, "IBM FlashCopy in the public cloud" on page 21, and is essentially the same process as DR. The only difference is that instead of a persistent replication, after the initial synchronization is complete, the goal is to schedule the cutover of the application onto the compute nodes in the cloud environment that is attached to the IBM Spectrum Virtualize storage.

Host-side mirroring requires the server to have concurrent access to local and remote storage, which is not feasible. Also, because the object is to relocate the workload (compute and storage) into the cloud environment, that task is more easily accomplished by replicating the storage and after it is synchronized, bringing up the server in the cloud environment and making the suitable adjustments to the server for use in the cloud.

The second method is largely impractical because it requires the host to access source *and* target simultaneously.

Also, 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, but again, requires the target storage to be visible through iSCSI to the host.

The third method entails the use of third-party software and or hardware to move the data from one environment to another one. The general idea is that the target system includes an operating system and some empty storage that is provisioned to that acts 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; however, the process is no different between an on-premises to cloud migration as it is to an on-premises to on-premises migration.

Table 2-4 lists the migration methods.

Table 2-4 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

2.5.3 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. At the time of writing, in Azure and AWS, the VM compute nodes are available with storage that is provisioned to the VM compute instance by using an iSCSI connection.

2.5.4 Implementation considerations

The workload relocation into the public cloud use case includes the following implementation considerations:

- Naming conventions: This important consideration is in the manageability of a standard on-premises IBM Spectrum Virtualize environment. However, because of the many layers of virtualization in a cloud implementation, maintaining a consistent and meaningful naming convention for all objects, such as MDisks, volumes, FlashCopy mappings, Remote Copy relationships, hosts, and host clusters, is necessary.
- Monitoring integration: Integration into IBM Spectrum Control or some other performance monitoring framework is useful for maintaining metrics for reporting or troubleshooting. IBM Spectrum Control is well suited for managing IBM Spectrum Virtualize environments.
- Planning and scheduling: Regardless of the method that is chosen, gather as much information ahead of time as possible (file system information, application custodians, full impact analysis of related systems, and so on).
- Be sure to ensure a solid backout: If inter-related systems or other circumstances require rolling back the application servers to on-premises, plan the migration to ensure as little difficulty as possible in the roll-back, which might mean keeping zoning in the library (even if it is not in the active configuration), and not destroying source volumes for a specific period.

3

Solution architecture

This chapter provides a technical overview of an Amazon Web Services (AWS) and Microsoft Azure environment regarding architectural aspects of IBM Spectrum Virtualize for Public Cloud in AWS and Azure. It also explains various components of the AWS and Azure solution, and how they interact and interrelate with each other.

This chapter includes the following topics:

- IBM Storage Virtualize
- Amazon Web Services terminology
- Key components of the AWS solution
- AWS highly available infrastructure
- AWS security design considerations
- AWS solution architecture
- Azure terminology
- Key components of the Azure solution
- Azure highly available infrastructure
- Azure security design considerations
- Azure solution architecture: IBM Spectrum Virtualize that is used as storage for an all-in-cloud model

3.1 IBM Storage Virtualize

IBM Storage Virtualize (previously called *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 Storage Virtualize is a core software engine of established and IBM storage virtualization solutions, such as IBM SAN Volume Controller and all versions of the IBM FlashSystem family of products. This technology is now available in AWS and Azure, which provides increased flexibility in data center infrastructure and cloud systems. This section describes the components of IBM Storage Virtualize as they are deployed in the cloud.

3.1.1 Nodes

IBM Storage Virtualize software on AWS is installed on Elastic Compute Cloud (EC2) instances that are provisioned in AWS. Each EC2 is called a *node*.

IBM Storage Virtualize software on Azure is installed on Azure virtual machine (VM) instances that are provisioned in Azure cloud. Each Azure VM instance 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.

Note: At the time of writing, IBM Spectrum Virtualize for Public Cloud on AWS is limited to one or two I/O group and IBM Spectrum Virtualize for Public Cloud on Azure is limited to a single I/O group.

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 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 functions and fault tolerance, which are described next.

3.1.2 I/O groups

Each pair of IBM Storage Virtualize nodes is referred to as an *I/O group*. A specific volume always is presented to a host server or cluster 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 always are processed by the same node within the I/O group. This node is referred to as the *preferred node* for this specific volume. When 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.

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, servers that are connected to use multipath drivers must handle these failover situations.

If required, host servers can be mapped to more than one I/O group within the IBM 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.

Note: For more information about restrictions around Non-Disruptive Volume Move (NDVM), see this IBM Documentation web page.

3.1.3 Systems

At the time of writing, IBM Spectrum Virtualize for Public Cloud on AWS supports a clustered system that consists of one or two I/O group and IBM Spectrum Virtualize for Public Cloud on Azure supports a clustered system that consists of one I/O group.

Specific configuration limitations are set for the individual system. The AWS implementation is optimized around 20 Amazon Elastic Block Store (EBS) volumes and the largest single Amazon EBS volume on AWS is 16,384 GiB. Because of those limitations, the practical limit of managed disks (MDisks) that are managed is twenty 16 TB or 320 TB.

On Azure, the maximum managed disks (MDisks) that are supported is 992 TB per system. The Azure implementation is optimized around 31 Azure MDisks and the largest single Azure disk is 32 TB.

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 a disaster occurs. This method does not back up application data. Only the IBM 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 V8.5.2.x and V8.5.4.x Configuration Limits and Restrictions for IBM Spectrum Virtualize for Public Cloud.

3.1.4 MDisks

IBM Spectrum Virtualize for Public Cloud on AWS views the Amazon EBS volumes that are presented to the EC2 instance nodes by AWS as several disks or LUNs, which are known as *MDisks*. Because IBM 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, and you assume that the Amazon EBS volumes are suitably protected and redundant.

IBM Spectrum Virtualize for Public Cloud on Azure uses Azure MDisks as back-end storage. Azure MDisks are offered with two storage redundancy options: *zone-redundant storage (ZRS)*, and *locally redundant storage (LRS)*. ZRS provides higher availability for MDisks than LRS does. However, the write latency for LRS disks is better than ZRS disks because LRS disks synchronously write data to three copies in a single data center.

IBM Spectrum Virtualize for Public Cloud supports only LRS types of storage. Azure MDisks are presented to the Azure VM instance (IBM Spectrum Virtualize nodes) as disks or LUNs. IBM Spectrum Virtualize does not attempt to provide recovery from physical failures within the back-end controllers; an MDisk in Azure is provisioned from LRS storage.

LRS disk volumes are protected and include built-in redundancy. LRS replicates your disk data three times within a single data center in the selected region. LRS protects your data against server rack and drive failures.

The application servers do not see the MDisks. 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 extents that are 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 V8.5.2.x and V8.5.4.x Configuration Limits and Restrictions for IBM Spectrum Virtualize for Public Cloud.

MDisks that are presented to IBM 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. IBM 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.

3.1.5 Storage pools

A *storage pool* or *MDisk group* 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, although methods are available to address this issue with volume mirroring (see 3.1.7, “Volumes” on page 35). The size of the extent can be 16 MiB (mebibyte) - 8192 MiB, with the default being 1024 MiB.

It is a best 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.

For more information, see this IBM Documentation web page.

3.1.6 Child pools

A *child pool* is a subset of a pool that is created for administrative isolation. This feature can be useful when thin-provisioned volumes are used to prevent a single application from using all available space in the pool.

Another important usage of child pools was introduced in IBM Spectrum Virtualize 8.4.2: the Safeguarded child pool. This object enables the creation of unalterable snapshots of important volumes to guard against accidental or malicious corruption.

3.1.7 Volumes

Volumes are logical disks that are presented to the host or application servers by IBM 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.

For more information, see this IBM Documentation web page.

3.1.8 Hosts

Volumes can be mapped to a host to enable access for a specific server to a set of volumes. A *host* within the IBM Spectrum Virtualize is a collection of iSCSI-qualified names (IQNs) that are defined on the specific server. As a result, a node failover (when a node is restarted) can be handled without having a multipath driver that is installed on the iSCSI-attached server.

For more information, see this IBM Documentation web page.

3.1.9 Host clusters

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

For more information, see this IBM Documentation web page.

3.1.10 iSCSI

iSCSI 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 by using the 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 IP network, especially over a potentially unreliable IP network.

Every iSCSI node in the network must have the following iSCSI components:

- 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 the iSCSI name and location of an iSCSI node. The address consists of a hostname 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.

3.1.11 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. The IBM Spectrum Virtualize capability to manage in the background the destaging operations that are incurred by writes (in addition to still supporting full data integrity) helps the IBM Spectrum Virtualize capability to achieve good database performance.

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

Figure 3-1 on page 37 shows the separation of the upper and lower cache.

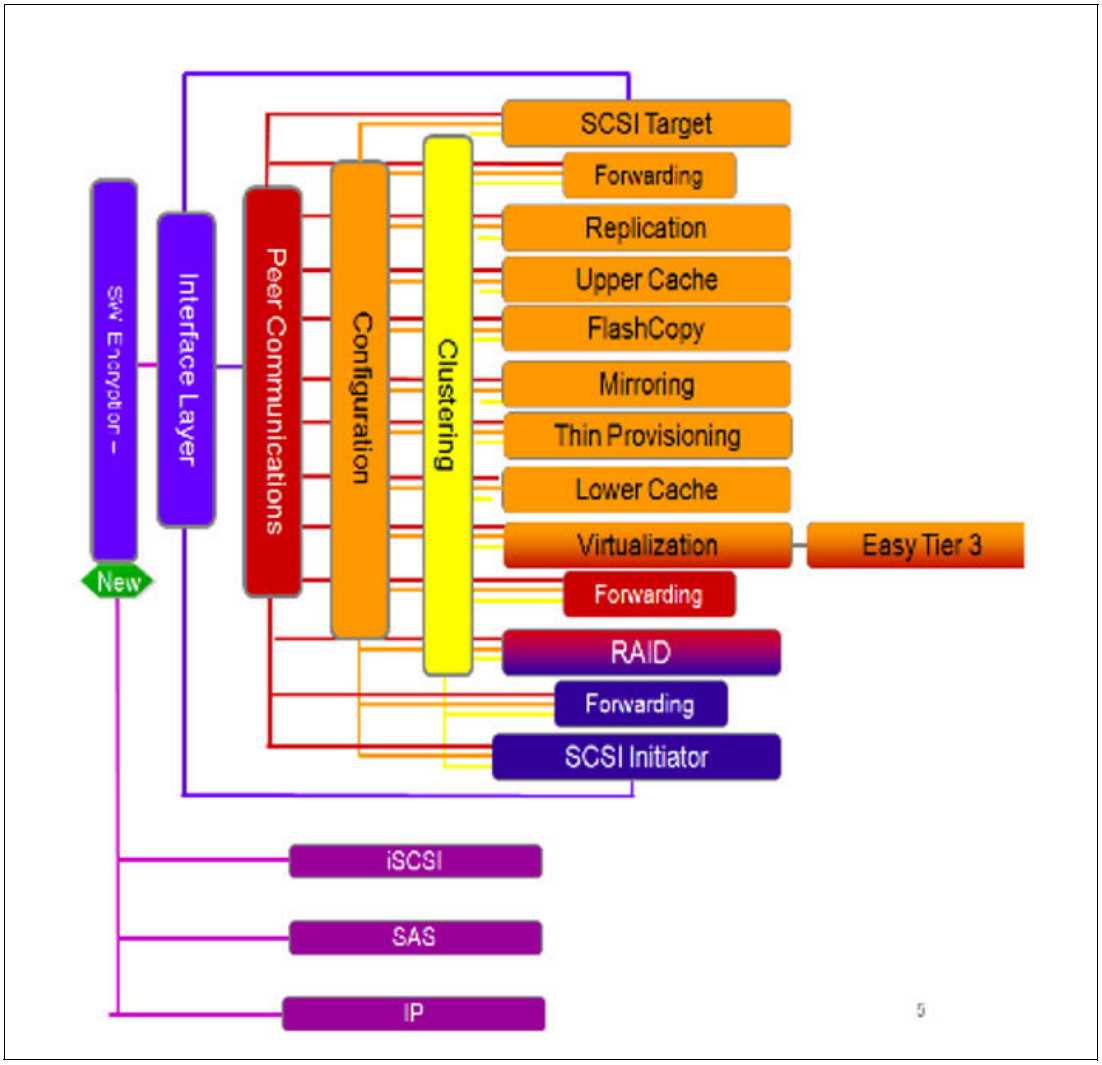

Figure 3-1 Separation of upper and lower cache

3.1.12 IBM Easy Tier

IBM Easy Tier is a performance function that automatically migrates or moves extents of 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, which 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. It 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 MDisks 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 by using the GUI. This load-balancing feature is not considered the same as the Easy Tier function, but it uses the same principles. Auto-balance evens the load for a pool across MDisks. Therefore, even the addition of new MDisks, or having MDisks of different sizes within a pool, does not adversely affect the performance.

The Easy Tier function can make it more suitable to use smaller storage pool extent sizes. The usage statistics file can be offloaded from the IBM 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 IBM Support web page.

3.1.13 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 FlashSystem family systems normally find each other in the network and can be selected from the GUI.

IP replication includes *Bridgeworks SANSlide* network optimization technology, which is available at no extra charge. Remote Mirror is a chargeable option, but the price does not change with IP replication. Existing Remote Mirror users can access the function at no extra 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 FlashSystem family, requires no separate appliances and so requires no extra 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.

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

IBM Spectrum Virtualize 8.5.2 introduces Policy-Based Replication (PBR), which provides simplified configuration and management of asynchronous replication between two systems. PBR uses volume groups (VGs) to deploy and manage automatically replication. This feature simplifies configuring, managing, and monitoring replication between two systems.

3.1.14 IBM 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 PiT 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 PiTs for the source volume, and to create multiple images from a source volume at a common PiT. 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 has multiple rollback points.

Most clients aim to integrate the FlashCopy feature for PiT copies and quick recovery of their applications and databases. An IBM solution is provided by IBM Storage Protect, which is described in What can IBM Storage Protect do for your business?

IBM Spectrum Virtualize 8.5.2 introduces Volume Group Snapshots, which enables snapshots of a VG. The snapshots can be recovered to thin clone type or clone type volumes in a new VG. Using the internal snapshot scheduler, a user automatically can schedule creation and deletion of either normal or Safeguarded snapshots.

3.2 Amazon Web Services terminology

AWS terminology is listed Table 3-1.

Table 3-1 AWS terminology

Item	Definition
EC2	A service that you can use to start VM instances in various operating systems.
EBS	Persistent block storage volumes that are used with Amazon EC2 (as opposed to the more common Simple Storage Service [S3]).
Availability zones	Distinct locations that are insulation from failures.
Virtual Private Cloud (VPC)	Virtual network (vNet) in your own logically isolated area within the AWS Cloud. It is populated by infrastructure, platform, and application services that share common security and interconnection.
CloudFormation Template	Creates and configures AWS resources and discovers dependencies.
Amazon Machine Images (AMIs)	Template that contains a software configuration (for example, an operating system, an application server, and applications).
S3	Storage for the internet. You can use Amazon S3 to store and retrieve any amount of data at any time, from anywhere on the web.

3.3 Key components of the AWS solution

The IBM Spectrum Virtualize for AWS infrastructure as a service (IaaS) solution includes the following building blocks:

- AWS EC2 instances
- Amazon EBS
- AWS networking elements, such as subnets and elastic network interfaces
- AWS constructs, such as VPC and placement group
- AWS security groups

AWS CloudFormation template-based publishing

IBM Spectrum Virtualize for Public Cloud is published in the AWS marketplace by using AMI-based CloudFormation to simplify provisioning of AWS resources. The IBM Spectrum Virtualize for Public Cloud installation provides two templates to provision and configure the required AWS services. One template installs the software on a new VPC and the other template is used for installations on an existing VPC. Both the templates provide the option of either public or private deployment.

AWS Virtual Private Cloud

The Amazon VPC service provisions a private, isolated section of the cloud where you can start AWS services and other resources in a vNet that you define. You have complete control over your virtual networking environment, including selection of an IP address range, creation of subnets, and configuration of route tables and network gateways. IBM Spectrum Virtualize for Public Cloud installation template in AWS Marketplace supports both creating a VPC or using an existing VPC.

3.4 AWS highly available infrastructure

An EC2 instance in the same IBM Spectrum Virtualize I/O group should spread across different underlying hardware to avoid a single point of failure (SPOF). As such, IBM Spectrum Virtualize for Public Cloud uses AWS Placement groups to spread EC2 instances on IBM Spectrum Virtualize nodes across different underlying hardware for high availability (HA).

Amazon EBS provides persistent block-level storage volumes for use with Amazon EC2 instances in the public cloud. Each Amazon EBS volume automatically is replicated within its availability zone to protect your environment from component failure through HA and durability.

3.5 AWS security design considerations

IBM Spectrum Virtualize nodes are protected from incoming traffic by using an AWS Security group, which acts as a virtual firewall. Inbound rules in a security group are automatically added during deployment from AWS Marketplace by allowing specific protocols and port ranges from specified source IP address ranges.

Note: For communication between two clusters from different VPC, you must do manual update of the inbound rules. For more information, see this IBM Documentation web page.

IBM Spectrum Virtualize for Public Cloud on AWS provides both public and private deployment. In private deployment, the default quorum or Azure Bastion host does not have a public IP address that is assigned to it, so by default it is not accessible from external network, which results in a more secure network. The default quorum or Azure Bastion host can be accessed securely by using AWS Session Manager. For more information about AWS Session Manager, see this Amazon Documentation web page.

3.6 AWS solution architecture

AWS delivers IaaS as VPCs, within which network, compute, and storage resources are housed. IBM Spectrum Virtualize nodes are built on EC2 instances and virtualize Amazon EBS volumes that are provisioned to those nodes, as shown in Figure 3-2. This configuration provides advanced capacity savings functions and replication and PiT copy services over a block virtualization layer through a user interface that is familiar to IBM Spectrum Virtualize users.

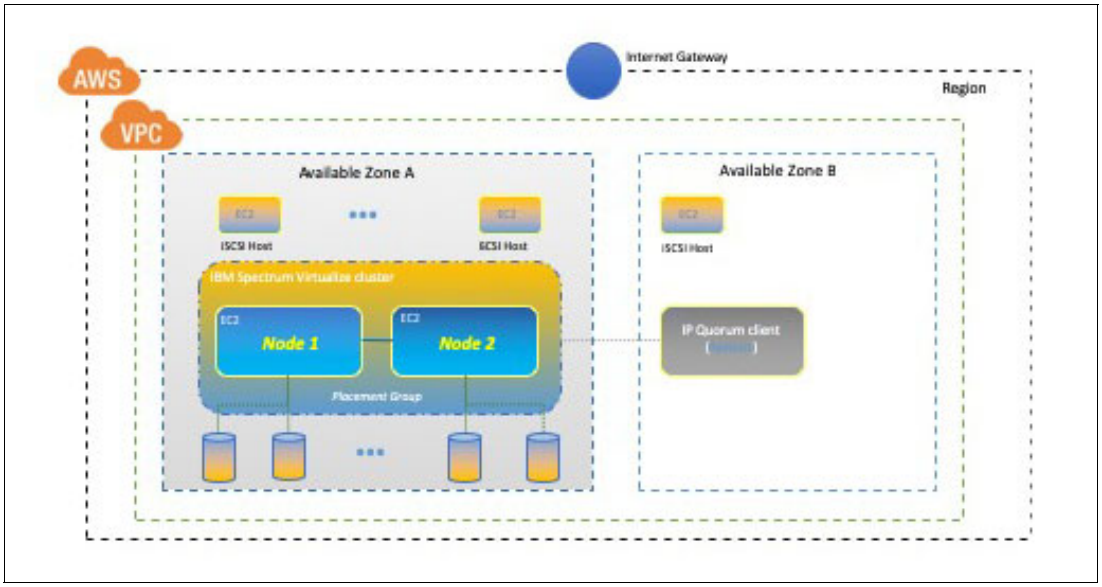

Figure 3-2 Architectural overview of high-level AWS components

3.6.1 Overview

As shown in Figure 3-2 on page 41, the IBM Spectrum Virtualize for Public Cloud environment is contained in a set of networks in a VPC; the IBM Spectrum Virtualize Cluster is on EC2 nodes; the Azure Bastion and initial IP quorum are on another smaller EC2 node; and the second IP quorum is in another availability zone for redundancy. As shown in Figure 3-3, those components are placed into the larger context of the solution as built for this document.

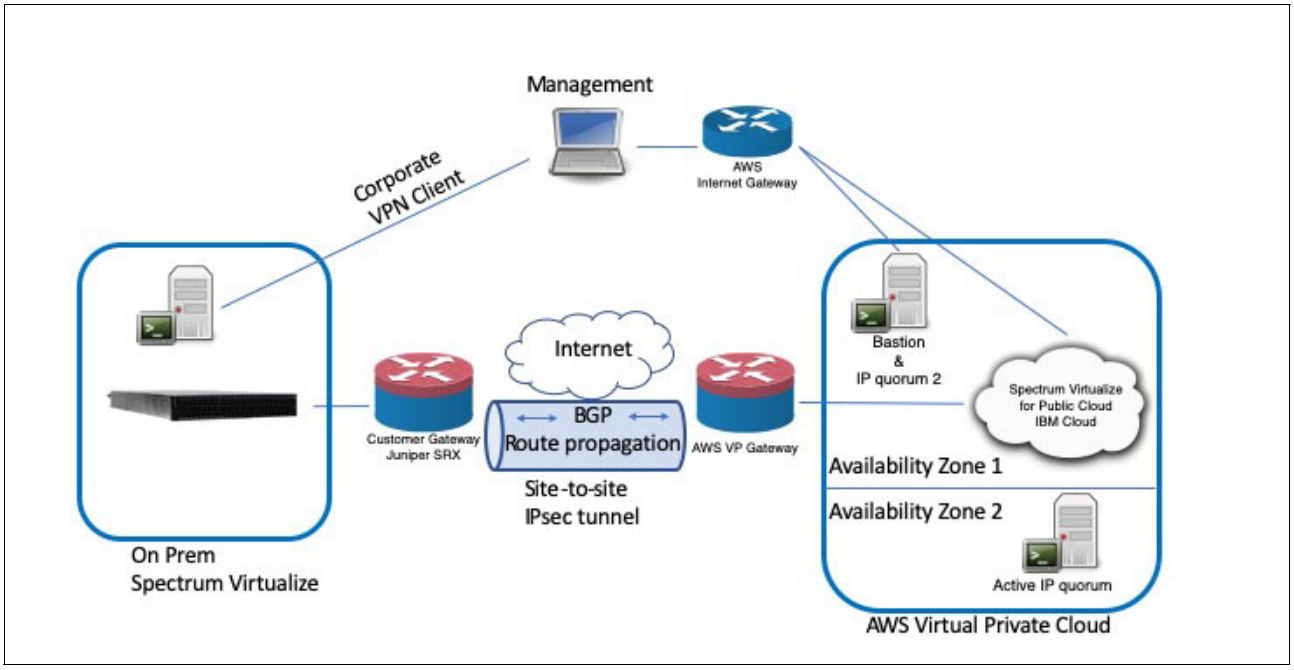

Figure 3-3 Connectivity between on-premises and Amazon AWS and management connections

In support of replication connectivity and the Transparent Cloud Tiering (TCT) function, internet access, and a site-to-site IPsec tunnel were added to the configuration. The site-to-site tunnel also provides an alternative method for managing the environment versus the use of Azure Bastion as a port-forwarder for the IBM Spectrum Virtualize GUI (which is not recommended from a security standpoint) or the complex process of setting up a *client virtual private network (VPN) Endpoint* in AWS (which requires configuring a certificate authority, and then, transferring certificates to the AWS Certificate Manager or integrating Active Directory for authentication).

3.6.2 Objective

The design of the solution for this publication was intended to show two key features of IBM Spectrum Virtualize for Public Cloud: Easy Tier and replication.

Easy Tier

Easy Tier with thin provisioning provides a compelling business justification for the use of IBM Spectrum Virtualize in Public Cloud, especially for AWS. With thin provisioning, capacity and performance can be extended for Amazon EBS volumes, which allows applications to achieve the same level of performance, but with lower cost than without IBM Spectrum Virtualize.

For the purposes of this example, two out of the four available Amazon EBS types are blended into a single Easy Tier pool. For the enterprise or nearline tier, *st1* is used and the flash tier is made up of *gp2 or gp3* volumes. The four available types of Amazon EBS volumes are compared in Table 3-2.

Table 3-2 Amazon EBS storage types

Items	Solid-state drives (SSDs)		Hard disk drives (HDDs)	
Volume type	General-purpose SSD (gp2 or gp3).	Provisioned IOPS SSD (io1).	Throughput-optimized HDD (st1).	Cold HDD (sc1).
Description	General-purpose SSD volume that balances price and performance for various workloads.	Highest-performance SSD volume for mission-critical, low-latency, or high-throughput workloads.	Low-cost HDD volume that is designed for frequently accessed, throughput-intensive workloads.	Lowest cost HDD volume that is designed for less frequently accessed workloads.
Use cases	► Recommended for most workloads. ► System boot volumes. ► Virtual desktops. ► Low-latency interactive apps. ► Development and test environments.	► Critical business applications that require sustained IOPS performance, or more than 16,000 IOPS or 250 MiBps of throughput per volume. ► Large database workloads, such as: – MongoDB – Cassandra – Microsoft SQL Server	► Streaming workloads requiring consistent, fast throughput at a low price. ► Big data. ► Data warehouses. ► Log processing. ► Cannot be a boot volume.	► Throughput-oriented storage for large volumes of data that is infrequently accessed. ► Scenarios where the lowest storage cost is important. ► Cannot be a boot volume.
API name	gp2 or gp3.	io1.	st1.	sc1.

Replication

Data mobility is firmly established as a foundational use case for block storage virtualization and the cornerstone of the IBM Spectrum Virtualize product. Replication between on-premises and AWS is the use case for the network diagram that is shown in Figure 3-3 on page 42.

3.6.3 Considerations

Consider the following important points when this solution is implemented:

- Use only the first three categories (and not the *sc1* class storage because of the high latency of that storage class).
- Because of the nature of Amazon EBS storage provisioning to EC2 instances, any single Amazon EBS volume is provisioned only to one of the two IBM Spectrum Virtualize nodes. However, the forwarding layer within the IBM Spectrum Virtualize software uses the custom tags that are attached to the Amazon EBS volumes to provide seamless handling of failover events.

- To ensure seamless handling of failover events, a 20 Amazon EBS volume limit per I/O group or IBM Spectrum Virtualize node pair is enforced.
- The *Cloud Formation Template (CFT)* that governs the installation of IBM Spectrum Virtualize for Public Cloud on AWS implements the Azure Bastion host and IP quorum server in the same availability zone for installations that are requesting a new VPC, and installations into an existing VPC allows the installer to choose a different availability zone for the Azure Bastion and IP quorum server. In installations that include the creation of a VPC, it is a best practice that an extra subnet is provisioned in a different availability zone and a secondary IP quorum server is started in that subnet and made the active quorum device.

3.7 Azure terminology

The Azure terminology is listed in Table 3-3.

Table 3-3 Azure terminology

Azure term	Explanation
Azure Resource Group (Azure RG)	A container that holds all the related resources for an Azure solution. For more information, see this web page.
Azure Key Vault	This vault is used to store the credentials in Azure and the store temporary key for cluster configuration, which are deleted after the cluster configuration is complete. For more information, see this web page.
Azure Network Security Group (Azure NSG)	This group contains security rules that define rules for inbound network traffic, and outbound network traffic between Azure resources. For more information, see this web page.
Azure Virtual Network (Azure vNet)	Azure vNet facilitates secured communication between different types of Azure resources, such as Azure VMs, external network, and the internet. For more information, see this web page.
Subnet	A range of IP addresses in a vNet. You can divide a vNet into multiple subnets for organization and security. Each network interface card (NIC) in a VM is connected to one subnet in one vNet. For more information, see this web page.
Proximity placement group	An Azure construct that is used to specify proximity and placement input to deployment when many logically grouped Azure VMs are deployed. When specified, Azure ensures that associated compute resources are physically close to each other. For more information, see this web page.
Availability set	A logical grouping of VMs that allows Azure to understand how your application is built to provide for redundancy and availability. For more information, see this web page.
Fault domain	A set of hardware components that share a SPOF.
Update domain	This group of resources can be updated and brought down for maintenance and patching (for example, system patches and software updates), possibly concurrently. Typically, only one update domain is brought down at a time for patching.

3.8 Key components of the Azure solution

Figure 3-4 shows the high-level architecture of IBM Spectrum Virtualize for Public Cloud on Microsoft Azure.

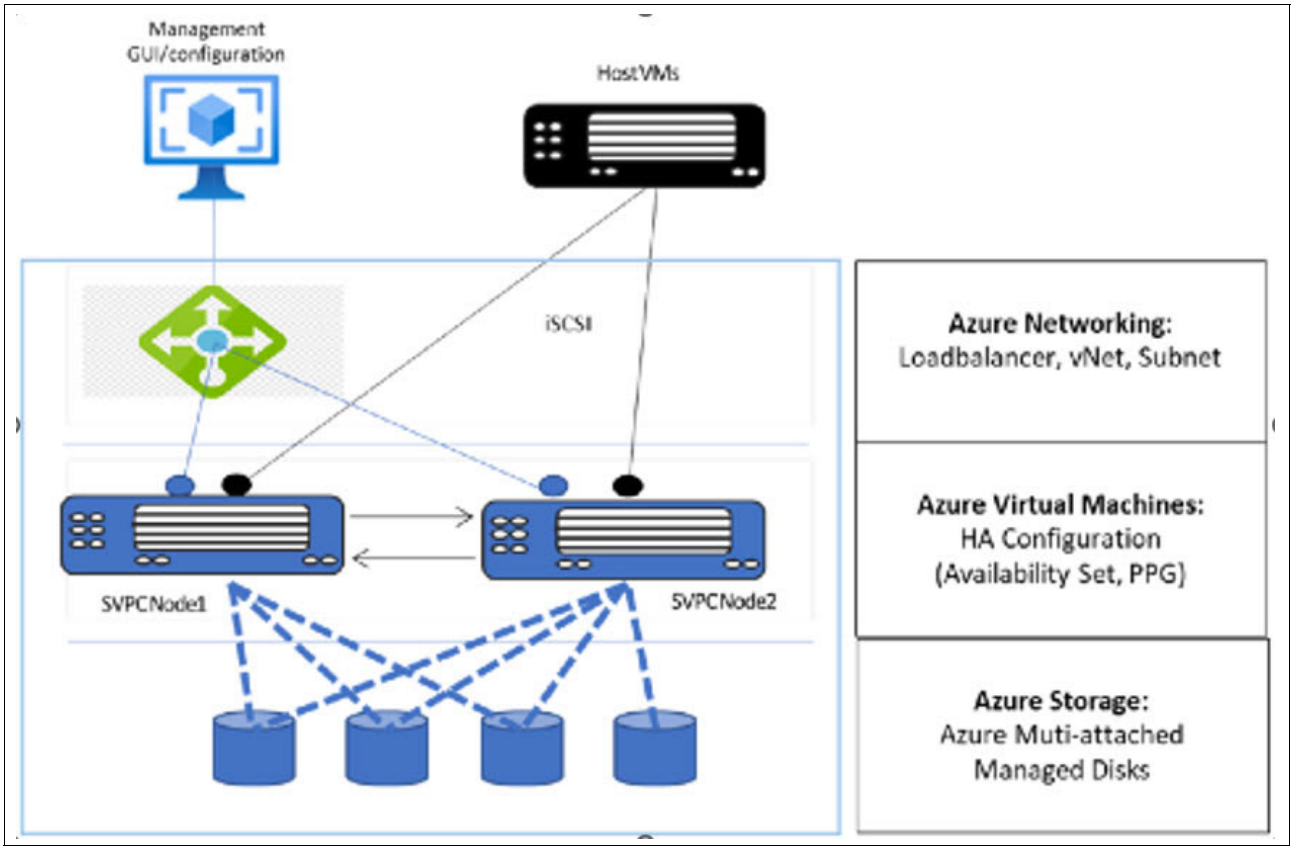

Figure 3-4 Components of the Azure solution

The IBM Spectrum Virtualize for Azure IaaS solution includes the following building blocks:

- Azure VMs
- Azure MDisks
- Azure networking elements, such as vNET, subnets, load-balancer, and vNICs
- Azure constructs, such as RG, availability set, and proximity
- Azure policies for NSGs

Azure Resource Manager template-based publishing

IBM Spectrum Virtualize for Public Cloud is published in the Azure commercial marketplace as an Azure Application with solution template. It employs a nested Azure Resource Manager (ARM) template, which is an infrastructure as code (IaC) primitive in Azure. The user must manage the solution and the transactions are billed through a separate plan. ARM template-based deployment automates deployment of several VMs along with a customized network infrastructure, security groups, key vault management, and MDisks in a customer's Azure subscription.

Resource group and vNets

An *RG* in Azure is a unit of management, so each IBM Spectrum Virtualize cluster is deployed in a fresh RG. All resources that are part of the cluster are held in their respective RGs.

Although multiple clusters cannot be deployed in the same RG, multiple RGs can share a vNet across different clusters, as shown in Figure 3-5.

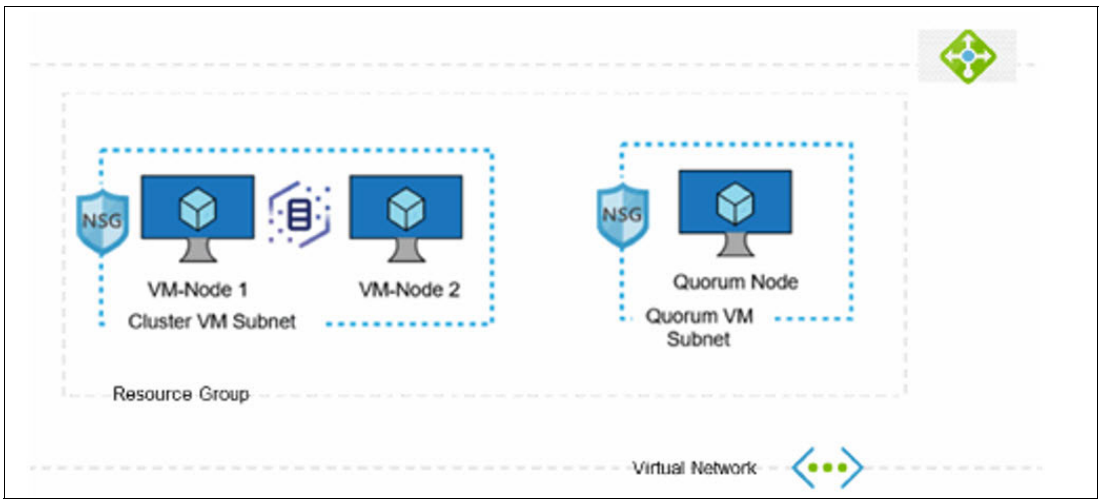

Figure 3-5 Resource group and vNets

A *vNet* is the building block for a customer's private network in Azure. A vNet facilitates various Azure resources to securely communicate with each other, with external resources over internet, and on-premises networks. Azure vNet is a vNet environment that provides the basis for provisioning resources and service in the Azure cloud.

IBM Spectrum Virtualize can be deployed in a new or an existing vNet. For nodes and quorum, a separate subnet is required. Azure vNet subnets are defined by the IP address block that is assigned to it.

Note: Multiple RGs can share a vNet across different clusters.

3.9 Azure highly available infrastructure

IBM Spectrum Virtualize nodes from an I/O group must be isolated in terms of different fault domains and different update domains so that cloud does not bring them down for any maintenance activity simultaneously. Figure 3-6 shows such a HA structure.

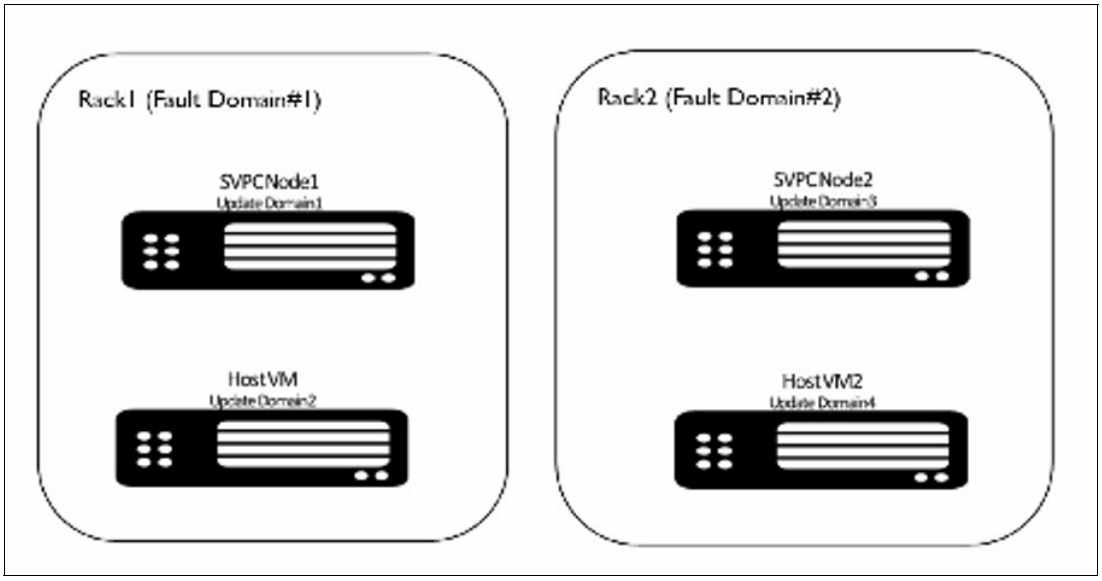

Figure 3-6 Highly available infrastructure

The IBM Spectrum Virtualize design uses an availability set construct to help Azure understand how the solution is designed to provide for redundancy and availability. The Azure Availability Set construct enables you to express logical groupings of IBM Spectrum Virtualize VMs and define the HA requirement.

Note: Consider the following points:

- Update domains indicate groups of VMs and underlying physical hardware that are brought down concurrently for Azure maintenance. With each node VM from an I/O group in a different fault domain, Azure ensures HA at the controller level.
- Fault domains define the group of VMs that share a common power source and network switch. IBM Spectrum Virtualize node VMs that are configured within an availability set are separated across up to three fault domains. Therefore, if a fault occurs, not all nodes in an I/O group go down because they are in different fault domains.

Controller node proximity

A proximity placement group's defined logical grouping ensures that Azure places the node compute resources close to each other.

IBM Spectrum Virtualize Node VMs must be separated by low network latency. Placing node VMs in an availability zone helps to reduce the physical distance between the instances. However, it might still result in network latency that is not acceptable for cluster functioning. Therefore, to get node VMs as close as possible with the lowest possible latency, use a proximity placement group.

Note: IBM Spectrum Virtualize nodes in a cluster all are added to the same proximity placement group.

Disk types and multi-attached shared disk support

IBM Spectrum Virtualize VMs interact with an attached MDisk to serve I/Os. When one VM from an I/O group goes down, I/O services that are provided by that VM are required to fail over to another surviving VM, which requires that MDisks are reattached to another partner VM quickly. This reattachment is best served by using the Azure multi-attach support for certain types of disks.

The multi-attach feature (see Figure 3-7) allows two or more node VMs to connect simultaneously to the same disk so that when one node VM goes down, the disk is still reachable from other VMs instantaneously.

Note: Microsoft provides the multi-attach (shared disk capability) support for standard and premium SSDs.

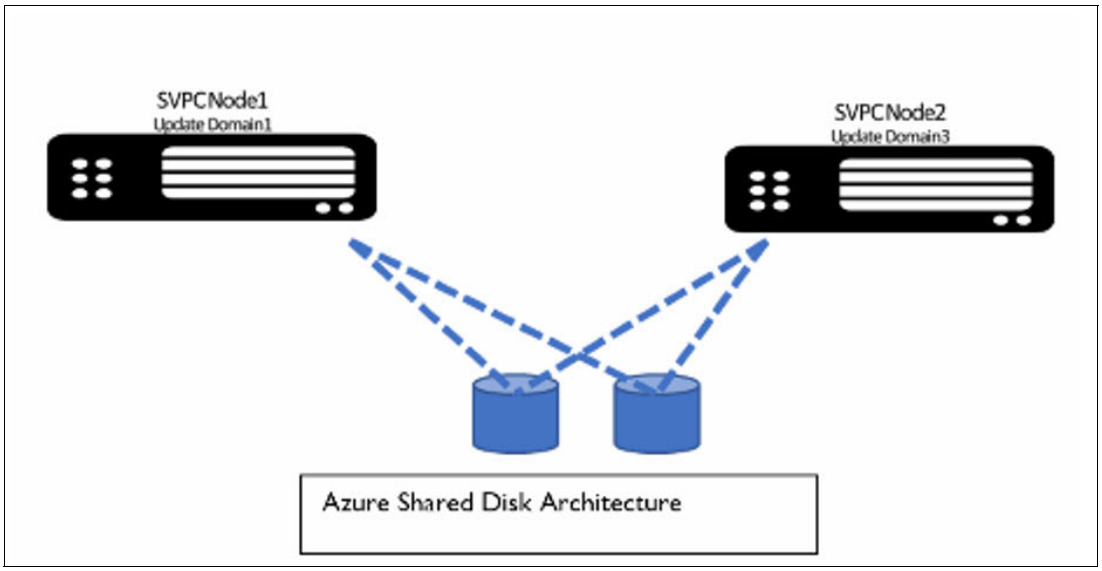

Figure 3-7 Multi-attach shared disk support

In this design, multi-attached disks are simultaneously connected with multiple nodes in the specific I/O group. The IBM Spectrum Virtualize requirement is to attach Azure MDisks to two VMs from same I/O group (see Figure 3-7).

Note: The disk that is required to be attached to IBM Spectrum Virtualize must include multi-attach capable disks with the number of mount points (num_share) equal to 2 or more.

IBM Spectrum Virtualize block storage allows for the creation of raw storage volumes to which server-based operating systems can connect. Table 3-4 on page 49 lists the Azure disk types that are required for different use cases and workloads that are best served by it.

Table 3-4 Disk types that are required for different use cases and workloads

Disk type	Block storage workload	Use cases
Standard SSDs	► Large data processing workloads ► Enterprise applications ► File storage and VM file system	► Hybrid cloud replication use case ► Hybrid cloud disaster recovery (DR) use case ► Low-end and all-in-cloud ► infrastructure
Premium SSDs	► High-performance database storage ► Mission-critical enterprise applications (Oracle, SAP, Microsoft Exchange, and Microsoft SharePoint	High-end and all-In-cloud infrastructure

Load balancer for management traffic failover

In Azure, IP address reassignment through IP failover from one VM to another VM can take much time. To solve this issue, as a best practice, use an Azure Load Balancer based design.

An IBM Spectrum Virtualize design uses the services of the Azure Load Balancer for management IP address failover. The Load Balancer redirects traffic from a received front-end IP address to a suitable back-end IP address of the node, which is running management service.

If the node fails, the Load Balancer service switches the traffic to another back-end IP address from the surviving node, as shown in Figure 3-8.

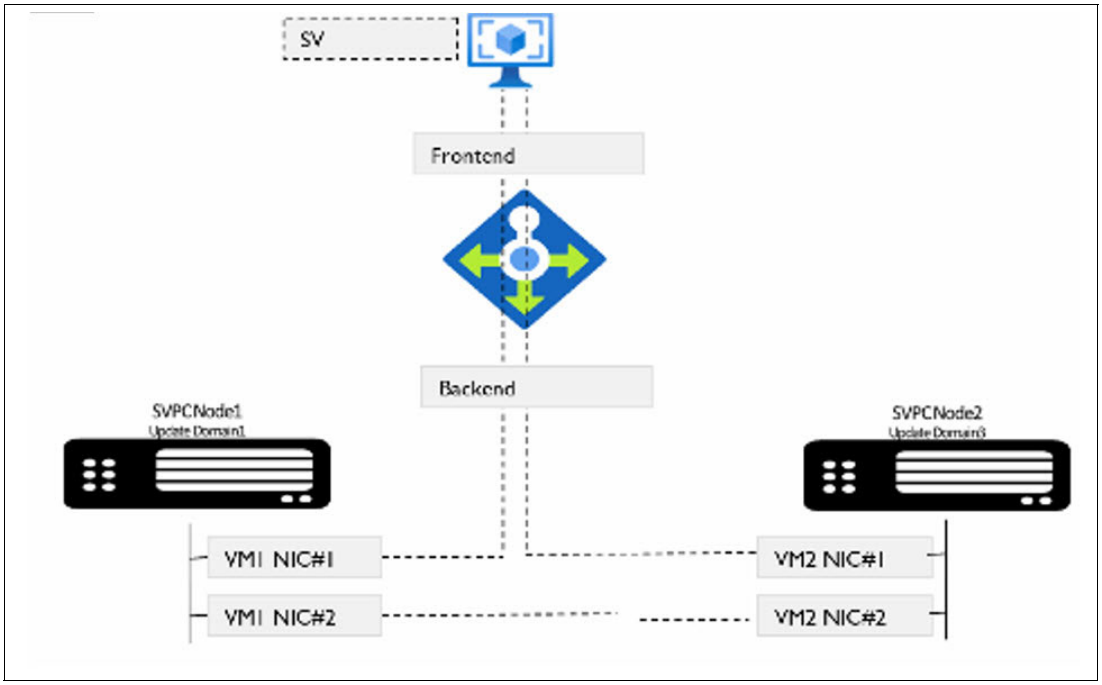

Figure 3-8 Load Balancer for management traffic failover

Multipath for data traffic failover and better performance

For I/O traffic failover, the IBM Spectrum Virtualize design relies on host multipathing technology by configuring redundant iSCSI I/O paths from a host to the dual controller. When a node fails in Azure, the original path (with a target IP address) is not presented by the partner node after a failover, so that path is stopped until the failed node is repaired. For this reason, you must use multi-paths from the host server to different nodes in the I/O group. Having multiple links to each of the node helps to protect the server against link failures in the network.

Note: Consider the following points:

- The IBM Spectrum Virtualize design does not use load balancer technology for data traffic. Instead, it relies on multipathing technology.
- Because the individual iSCSI connection from VM can be throttled in the cloud environment, increase the iSCSI sessions from the host to realize the high throughput performance.

3.10 Azure security design considerations

Cloud features share the responsibility model of security between the cloud provider, IBM Spectrum Virtualize, and customers.

An Azure cloud provides security for compute, network, and physical resources in cloud data centers, and various virtual resources, such as VM, vNet, MDisk storage, and various associated IaaS services. Azure also provides a list of tools and services to manage the security of Azure resources. For more information about Azure security benchmark documentation and guidelines, see this web page.

Consider the following aspects of security of IBM Spectrum Virtualize for Public Cloud for Azure:

- The core IBM Spectrum Virtualize software provides several basic common security features.
- IBM Spectrum Virtualize for Public Cloud provides other areas of security (see Figure 3-9 on page 51) that factor into a specifically Azure public cloud environment:
  - Network security (all-in-cloud private deployment and restrictive network security rules).
  - Azure access security (Azure Bastion access and hybrid cloud access).
  - Operating system security (operating system patch update management and Azure Update Manager).

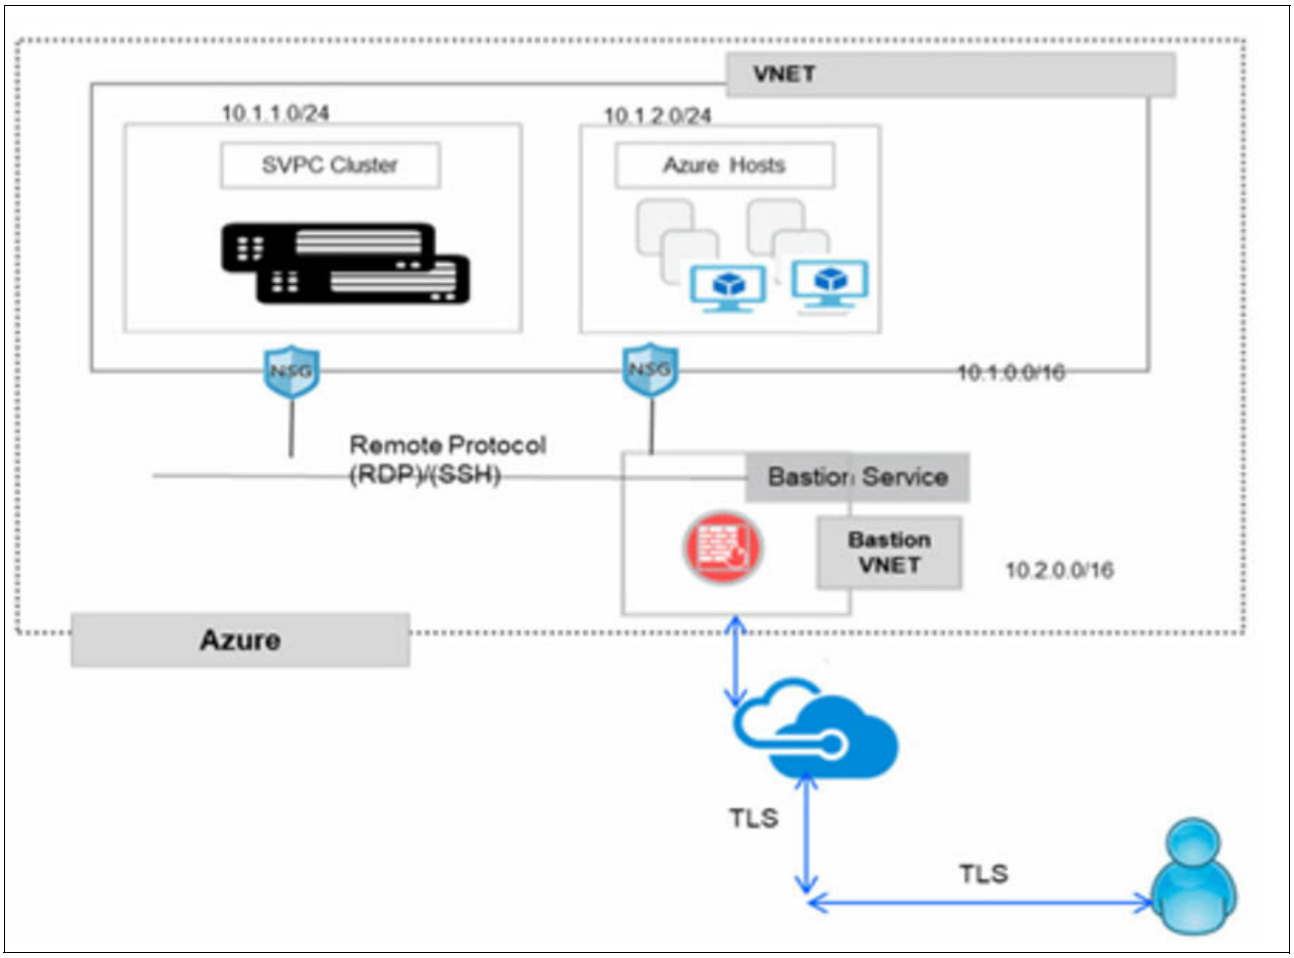

Figure 3-9 Security design aspects

Restrictive network security rules

When an IBM Spectrum Virtualize network is deployed by using the Azure Resource Management template, NSGs are created and applied at the subnet and network interface levels. NSGs also automatically are configured for each subnet or interface to protect the vNet. Default NSG rules tighten security to achieve the goal of network security.

For more information, see this Microsoft Documentation web page.

To bolster security at the network level, consider the following points:

- Default NSG rules that are provided with IBM Spectrum Virtualize deployment use only private IP addresses and provide access to or from only private IP addresses.
- By default (unless assigned), no public endpoints exist that access IBM Spectrum Virtualize configuration IP addresses or data traffic IP addresses.
- To communicate with some specific services (for example, entitlement), some outbound access rules are enabled.
- NSG rules that contain tight security rules are defined to create a secured design.

Azure Bastion access

Azure Bastion is a fully managed service that provides secure and seamless Remote Desktop Protocol (RDP) and Secure Shell (SSH) Protocol access to VMs without any exposure through public IP addresses. By using Azure Bastion, private and time-bound access can be provided through SSH and RDP. Azure Bastion acts like a gateway proxy, which allows users to connect to resources in a private subnet.

Operating system patch update interface and Azure Update Manager integration

Figure 3-10 shows the construction of a node. In Azure, IBM Spectrum Virtualize is packaged in a container and runs in an Azure VM. RHEL OS is used as an operating system in the VM. Managing patching and updates is an important security requirement, especially in a public cloud. The following methods are available that can be used to meet this security requirement:

- Manually by using IBM Spectrum Virtualize CLIs
- Automatically by using Azure Update Manager

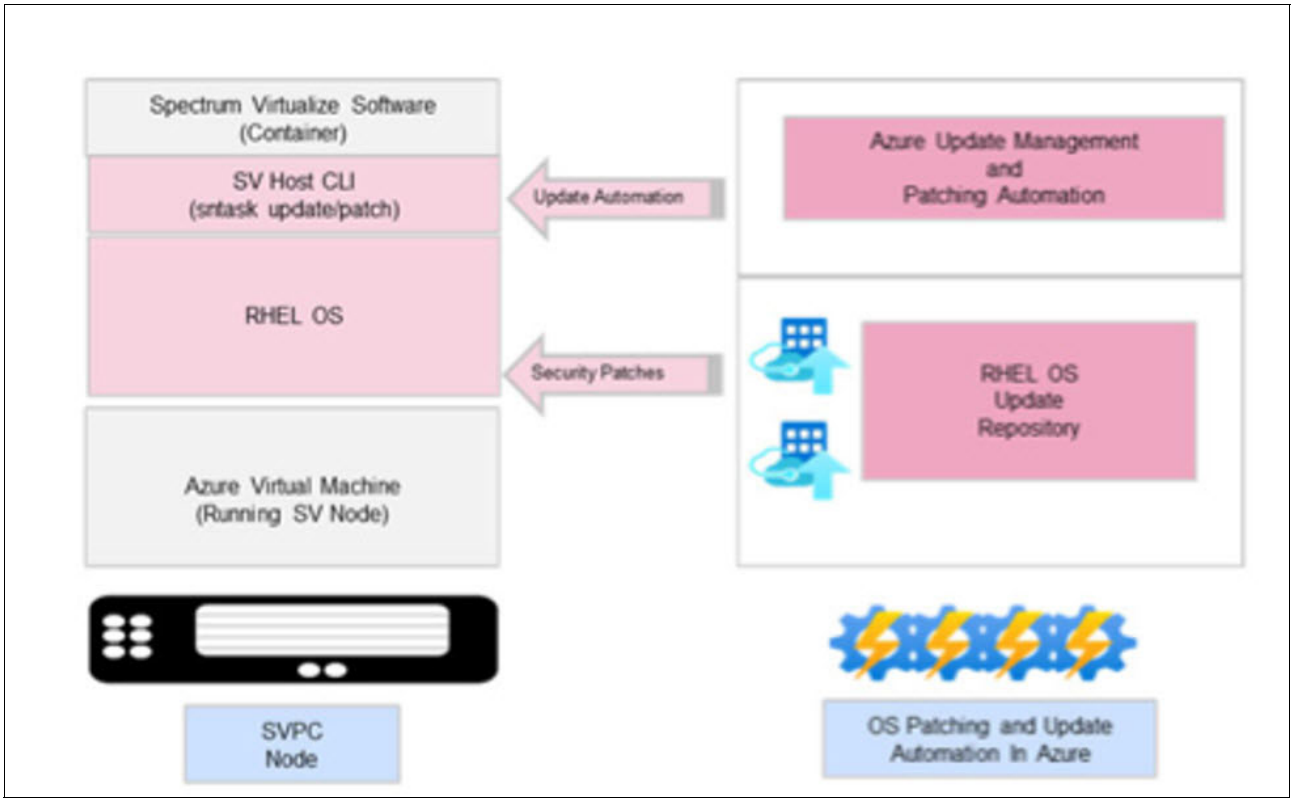

Figure 3-10 Operating system patch update interface and Azure Update Manager integration

Manual patching by using IBM Spectrum Virtualize CLIs

IBM Spectrum Virtualize provides service node task (**`sntask`**) commands for managing hosts in public cloud offerings. In Azure solution, these **`sntask`** commands and service node information (**`sninfo`**) commands are stored and run as terminal commands on the RHEL OS of the VM instances.

One of the options of these CLI options (**`sntask applysecuritypatch`**) can be used to manually apply security patches to RHEL OS. An Azure Red Hat repository is used to retrieve and apply security patches to the node VMs.

Automated patching by using Azure Update Manager

Azure Update Manager is an important security service that uses Azure Automation to manage operating system updates in Linux VMs. The following process is used to configure automated patching for all IBM Spectrum Virtualize nodes:

1. Create an Azure automation account (for more information, see this web page).
2. Configure this automation account with a Log Analytics workspace (for more information, see this web page).
3. Configure the required roles and permissions for automation.
4. Create run books and configure IBM Spectrum Virtualize scripts under them.
5. Add IBM Spectrum Virtualize node VMs under the update management.
6. Configure patching schedules and policies.

After the configuration is done, Azure Update Manager automatically evaluates Azure VMs to maintain security compliance regarding released Critical and Security updates.

3.11 Azure solution architecture: IBM Spectrum Virtualize that is used as storage for an all-in-cloud model

In the all-in-cloud model, as shown in Figure 3-11, IBM Spectrum Virtualize for Public Cloud is deployed as a clustered pair of Azure Compute VM instances. It arbitrates between the Cloud Block Storage and the workload hosts. In the specific context of all-in-cloud deployment, IBM Spectrum Virtualize for Public Cloud supports extra features that enhance the public cloud block-storage offering.

In this use case, a clustered system of IBM Spectrum Virtualize nodes (realized by using memory rich, network rich Azure VMs) presents volumes to the hosts (Virtualized Hosts - Realized as Azure VMs) running customer workloads. Hosts (VMs) access these volumes and read/write data. Most of the advanced functions that are provided by IBM Spectrum Virtualize are defined on volumes. These volumes are created from Azure MDisks. All data transfer occurs through the IBM Spectrum Virtualize node, which is described as symmetric virtualization.

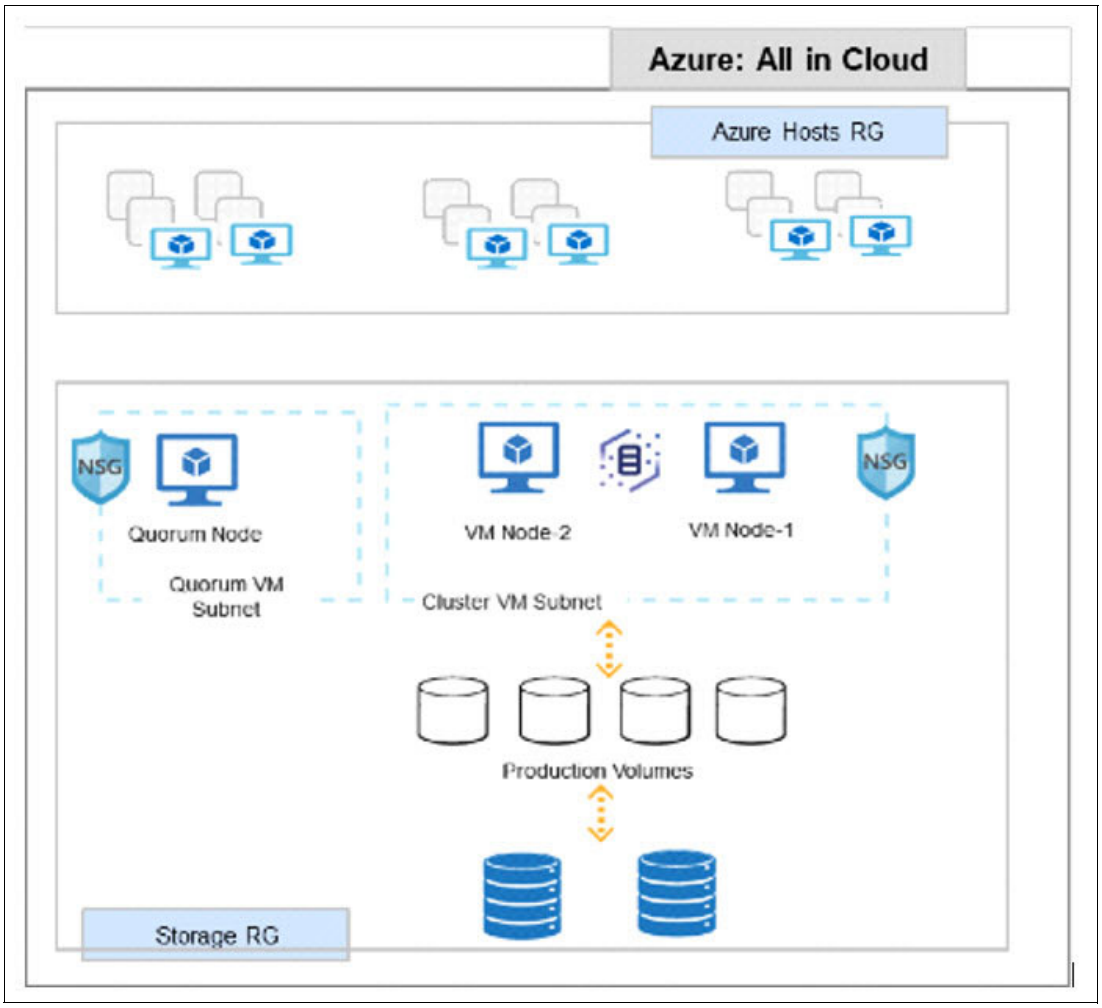

Figure 3-11 IBM Spectrum Virtualize as storage for an all-in-cloud model

4

Planning and preparing for IBM Spectrum Virtualize for Public Cloud

This chapter describes the planning considerations for successfully implementing IBM Spectrum Virtualize for Public Cloud on both the Amazon Web Services (AWS) and Microsoft Azure cloud platforms.

This chapter includes the following topics:

- Introduction
- General planning introduction
- Planning for Amazon Web Services
- Planning for Microsoft Azure

4.1 Introduction

This chapter describes the planning and preparation steps to provision network, server, and storage components on AWS and Microsoft Azure in support of IBM Spectrum Virtualize for Public Cloud.

Background information about the cloud networking architecture, compute, and storage offerings also are described to help the reader who is unfamiliar with the placement of IBM Spectrum Virtualize for Public Cloud within the larger context of an application environment.

4.2 General planning introduction

Both AWS and Microsoft Azure offer secure cloud services with storage and resources for deploying IBM Spectrum Virtualize for Public Cloud software. The IBM Spectrum Virtualize for Public Cloud installation is available within both the AWS and Azure marketplaces.

The IBM Spectrum Virtualize for Public Cloud installation uses the Cloud Formation Template (CFT) in AWS and the Azure Resource Manager (ARM) template to simplify provisioning of all needed resources.

When the installation templates are used from either the AWS or Azure marketplace, the user is prompted to provide information (such as customer ID) for an entitlement check. The process provisions resources that are based on the settings that are defined within the installation template of the IBM Spectrum Virtualize for Public Cloud software stack.

4.2.1 Prerequisites for IBM Spectrum Virtualize for Public Cloud

The IBM Spectrum Virtualize for Public Cloud on AWS software is a Bring Your Own License (BYOL) offering from IBM. During its deployment on either AWS or Azure, the installation template verifies the proof of entitlement that indicates that a valid license was acquired from IBM. If the proof of entitlement is not present or valid, the installation fails. To obtain the license and proof of entitlement for the software, complete the following steps:

1. Go to the IBM Passport Advantage website to obtain a license and proof of entitlement for the software.
2. At the website, follow the directions to enter your IBM Customer Number and the maximum number of terabytes of virtual storage to provision with your systems.

Note: The proof of entitlement and the IBM Customer Number are provided by email to the person who acquired the license.

4.2.2 Prerequisites for AWS

Before you install IBM Spectrum Virtualize for Public Cloud software from the AWS Marketplace, ensure that you complete the following tasks on the AWS site:

1. Sign up for AWS.
2. Create an AWS Identity and Access Management (IAM) administrator profile.
3. Assign the appropriate rules for installation and usage.
4. Create a key pair.

You can use the default AWS administrator profile to install the IBM Spectrum Virtualize for Public Cloud software or you can create an installer user profile that includes only the required permissions for deploying the software. Creating a second user for only monitoring is a best practice. To create those two extra users, complete the following steps:

1. Create a suitable user profile for the installer and the monitoring user.
2. Create one user for installation and one for monitoring, and assign the suitable user profile.

Creating an AWS user profile

To create an installer user profile, complete the following steps:

1. Log on to the AWS Management Console with the AWS default administrator profile.
2. Select **Services** in the upper left and click **IAM** to open the IAM console.
3. In the **Navigation** pane, select **Policies** → **Create policy**.
4. Click the **JSON** tab and add the JSON content that is shown in Example 4-1.

Example 4-1 User profile for an installer user

```
{
   "Version": "2012-10-17",
   "Statement": [
       {
           "Sid": "VisualEditor0",
           "Effect": "Allow",
           "Action": [
               "aws-marketplace:ViewSubscriptions",
               "aws-marketplace:Unsubscribe",
               "aws-marketplace:StartBuild",
               "aws-marketplace:ListBuilds",
               "aws-marketplace:Subscribe",
               "iam:CreateInstanceProfile",
               "cloudformation:CreateUploadBucket",
               "sns:DeleteTopic",
               "iam:RemoveRoleFromInstanceProfile",
               "iam:CreateRole",
               "cloudformation:UpdateTerminationProtection",
               "s3:CreateBucket",
               "sns:ListTopics",
               "sns:Unsubscribe",
               "iam:PutRolePolicy",
               "iam:AddRoleToInstanceProfile",
               "iam:PassRole",
               "cloudformation:DescribeStackEvents",
               "ssm:DescribeParameters",
               "iam:DeleteRolePolicy",
               "cloudformation:UpdateStack",
               "sns:Subscribe",
               "s3:DeleteObject",
               "s3:DeleteBucket",
               "cloudformation:ListStackResources",
               "iam:DeleteInstanceProfile",
               "iam:GetRole",
               "cloudformation:ListStacks",
               "iam:GetInstanceProfile",
```

```
                "sns:GetTopicAttributes",
                "cloudformation:DescribeStackResources",
                "sns:CreateTopic",
                "iam:ListRoles",
                "iam:DeleteRole",
                "ssm:GetParameters",
                "iam:ListInstanceProfiles",
                "cloudformation:GetTemplateSummary",
                "cloudformation:DescribeStacks",
                "s3:PutObject",
                "s3:GetObject",
                "cloudformation:GetStackPolicy",
                "s3:ListAllMyBuckets",
                "cloudformation:CreateStack",
                "cloudformation:GetTemplate",
                "cloudformation:DeleteStack",
                "ec2:*",
                "cloudformation:ListChangeSets"
            ],
            "Resource": "*"
        }
    ]
}
```

5. Click **Review Policy** and add a name for the policy, such as `SV_install_policy`.
6. Click **Create Policy**.

To create a user profile with limited permissions, repeat steps 1 - 6 but use the JSON content that is shown in Example 4-2 when you create the customized policy.

Example 4-2 User profile for a monitoring user

```
{
   "Version": "2012-10-17",
   "Statement": [
       {
           "Sid": "VisualEditor0",
           "Effect": "Allow",
           "Action": [
               "ec2:RebootInstances",
               "iam:GetRole",
               "ec2:Describe*",
               "ec2:StartInstances",
               "iam:ListRoleTags",
               "iam:ListAttachedRolePolicies",
               "iam:ListRoles",
               "iam:ListPolicies",
               "ec2:StopInstances",
               "iam:ListRolePolicies",
               "iam:ListInstanceProfiles",
               "iam:GetRolePolicy",
               "ec2:Get*"
           ],
           "Resource": "*"
       }
   ]
}
```

7. Click **Review Policy** and add a name for the policy, such as `SV_monitor_policy`. Click **Create Policy**.

Creating an AWS user and assigning the appropriate profile

To create two users for installation and monitoring, complete the following steps twice:

1. Log on to the AWS Management Console with the AWS default administrator profile.
2. Select **Services** in the upper left and click IAM to open the IAM console.
3. In the **Navigation** pane, select **Users** → **Add user**.
4. Enter a name and password, and ensure that you select AWS Management Console access for the **Access type**. Click **Next: Permissions**.
5. Select **Attach existing policies directly** and select the policy that you created in “Creating an AWS user profile” on page 57. Click **Next: Tags**.
6. Ensure that you add a tag that includes the email address for the installer user profile.

A successful login that uses the new user ID requires the login link that is provided by the AWS email, which must be manually sent during the creation process.

Because these steps might change, see IBM Documentation for updated information.

4.2.3 Prerequisites for Microsoft Azure

Before you install IBM Spectrum Virtualize for Public Cloud software from the Microsoft Azure Marketplace, complete the following tasks:

1. Sign up for Microsoft Azure.

   To use the Microsoft Azure services that IBM Spectrum Virtualize for Public Cloud installation provisions, such as your virtual machine (VM) instances, you must have a valid Azure account.

2. Choose the proper Azure profile.

   The Azure default administrator profile can set credentials for purchasing, setting up, and configuring Azure resources that are necessary for an IBM Spectrum Virtualize for Public Cloud deployment in an Azure cloud environment.

   The Azure default administrator profile can be used to install the IBM Spectrum Virtualize for Public Cloud software or an installer user profile can be created that includes only the required permissions for deploying the software.

   In addition, a separate user profile is a best practice for users that are completing daily work on the system. These users have limited permissions, which restrict them from specific actions that are based on Azure policies.

   For more information about these user profiles and security group settings, see 4.4.1, “Planning security access control on Microsoft Azure” on page 67.

3. Create a Secure Shell (SSH) key pair.

   Azure encrypts login information that uses public-key cryptography for security. You can create VMs in Azure that use SSH keys for authentication.

Microsoft Azure services that are used for IBM Spectrum Virtualize for Public Cloud installation templates

Microsoft Azure provides several services that the IBM Spectrum Virtualize for Public Cloud installation template uses to ensure that the required components are configured in the Azure cloud.

The ARM templates service is used in an IBM Spectrum Virtualize for Public Cloud installation. These ARM templates create and manage a collection of related resources that are required for the IBM Spectrum Virtualize for Public Cloud environment. One template is used for new instances and the other template is used for instances.

The following resources are installed and provisioned with these templates:

- Azure Compute Service (VMs)

  The Azure Compute Service starts VM instances with various operating systems. You can choose from virtual hard disk (VHD) or import your own VM images. As a part of IBM Spectrum Virtualize for Public Cloud, three VMs are provisioned and deployed: Two VMs are deployed as IBM Spectrum Virtualize for Public Cloud nodes, and the third is used for IP quorum management.

- Azure Virtual Network (Azure vNet)

  vNet provides an isolated and highly secure environment in which to run your VMs and other resources. vNet is the fundamental building block for your private network in Azure that enables many types of Azure resources, such as Azure VMs, to securely communicate with the internet and on-premises networks.

  The IBM Spectrum Virtualize for Public Cloud installation template in Azure Marketplace supports creating a vNet or using an existing vNet.

- Azure Network Security Group (NSG)

  An Azure NSG is a static set of rules that protects each network. An NSG contains security rules that allow or deny network traffic to or from an Azure resource. The NSG settings for the inbound or the outbound network traffic are configured during the installation of templates from the Azure Marketplace.

- Standard Load Balancer

  Azure Standard Load Balancer is used for management IP address failover, and it is provisioned as part of IBM Spectrum Virtualize for Public Cloud installation. The Load Balancer service is configured per resource group (RG).

4.3 Planning for Amazon Web Services

The section describes important planning considerations for a successful implementation of IBM Spectrum Virtualize for Public Cloud on the AWS cloud platform.

4.3.1 Requirements and limitations

The installation is available on AWS Marketplace and uses AWS CloudFormation service to simplify provisioning of AWS resources. The IBM Spectrum Virtualize for Public Cloud installation provides two templates to provision and configure the required AWS services. One template installs the software on a new virtual private cloud (VPC) and the other template is used for installations on an existing VPC.

When the installation template is started from AWS Marketplace, the user is prompted to provide information, such as a customer ID for the entitlement check. For more information, see IBM Documentation.

IBM Spectrum Virtualize for Public Cloud on AWS provides the following set of storage functions, which are similar to the functions that are provided by IBM Spectrum Virtualize for on-premises installations:

- IP-based Copy Services are available for the following types of replication:
  - Global Mirror
  - Metro Mirror
  - Global Mirror with Change Volumes (GMCV)
- Replication is possible between components:
  - On-premises SAN Volume Controller, IBM FlashSystem, or IBM Spectrum Virtualize as software only on Bare Metal Servers to AWS Cloud
  - IBM Spectrum Virtualize for Public Cloud on AWS instances that are deployed into two different availability zones
- FlashCopy, IBM Easy Tier, and thin provisioning are supported by IBM Spectrum Virtualize for Public Cloud on AWS.

Consider the following scalability-related limitations:

- Two or four nodes per cluster only.
- IPv4 only (no IPv6).
- A total of 20 Amazon Elastic Block Store (EBS) volumes per I/O group.
- The general maximum Amazon EBS Volume size is 16 TiB.

Multiple availability zones are supported only through a separate Global Mirror instance of IBM Spectrum Virtualize for Public Cloud.

The following features are unsupported in this release:

- Stretched Cluster
- HyperSwap
- Real-time Compression
- IBM Spectrum Virtualize native encryption
- Hot Spare Node (not applicable to cloud)
- DRAID and encrypted DRAID (not applicable to cloud)
- N_Port ID virtualization (NPIV) (not applicable to cloud)
- SCSI Unmap for host and back end

Note: IBM Spectrum Virtualize for Public Cloud on AWS is configured by the AWS time server by using underlying operating system methods. Changing the time server or setting a static time is not a best practice and might cause difficulties. For more information about the AWS time server, see Setting the Time for Your Linux Instance.

Consider the following points:

- Configuration requirement: The Amazon EBS volumes and the Elastic Compute Cloud (EC2) instance to which it attaches must be in the same Availability Zone.
- Configuration best practice: The EC2 instance that is acting as the IP quorum should be in another Availability Zone in the same VPC.

For more information, see IBM Documentation.

4.3.2 Amazon Web Services resources

Multiple resources are required for IBM Spectrum Virtualize for Public Cloud on AWS. Each IBM Spectrum Virtualize node requires one EC2 server instance. A single EC2 instance is required for the IP quorum device. Amazon EBS storage devices are used as IBM Spectrum Virtualize managed disks (MDisks).

After the CloudFormation template completes the installation, the result is a fully configured 2-node cluster in a private network with two MDisks. In addition, the installation process performs the following tasks:

- Validate the entitlement for an IBM Spectrum Virtualize for Public Cloud purchase (the client provides a customer number).
- Configure all IP addresses (cluster IP address, service IP addresses, node IP addresses, and iSCSI port IP addresses).
- Configure NTP and DNS with AWS internal servers.

The security rules are automatically configured according to AWS requirements:

- For access to services that are provided by the IBM Spectrum Virtualize for Public Cloud system (web GUI, SSH, IPsec, and iSCSI).
- For IBM Spectrum Virtualize inter-node communication, including network failover.
- For IBM Spectrum Virtualize for Public Cloud to manage Amazon EBS.
- One IP quorum client is configured on a third EC2 instance (an Azure Bastion host).

At least 11 IP addresses are required for a single IBM Spectrum Virtualize for Public Cloud on AWS installation in a VPC:

- Two node IP addresses per node
- Two port IP addresses per node
- One service IP address per node
- One IBM Spectrum Virtualize for Public Cloud on AWS cluster IP address

4.3.3 Amazon EC2 instances

The network bandwidth, the number of virtual central processing units (vCPUs), and the amount of memory are determined by instance type. The AWS instances (C5.4xlarge, C5.9xlarge, and C5.18xlarge) are the available options in the first release. The technical specifications are shown on Table 4-1. Dedicated Hosts mode is not supported for the first release.

Table 4-1 Amazon AWS EC2 on-demand resources

EC2 instance	vCPU	Memory (GiB)	Dedicated Amazon EBS bandwidth (Mbps)	Network performance (Gbps)
c5.4xlarge	16	32	3.500	Up to 10
c5.9xlarge	36	72	7.000	10
c5.18xlarge	72	144	14.000	25
c5.large (quorum only)	2	4	Up to 4,750	Up to 10

Note: These specifications were valid at the time of writing (May 2019).

For more information about available Amazon EC2 instances and pricing, see the following resources:

- Amazon EC2 Instances Types
- Amazon EC2 Pricing

Note: Selection of the EC2 instances should be done carefully because no EC2 node exchange or upgrade process is in place. A mix of different EC2 instances in the same IBM Spectrum Virtualize cluster is not supported in the first release. Migration to a different hardware platform can be done by replication to a new cluster.

4.3.4 AWS Elastic Block Stores

All Amazon EBS volume types are designed for 99.999% availability. They fall into the following categories:

- Solid-state drive (SSD)-based volumes that are optimized for transactional workloads with a small I/O size. The dominant performance attribute is IOPS.
- Hard disk drive (HDD)-based volumes that are optimized for streaming workloads, which are measured in MiBps.

Different volume types are available for Amazon EBS. They differ in performance characteristics, as listed in Table 4-2.

Table 4-2 EBC volume types

Item	SSDs		HDDs	
Volume type	General-purpose SSD	Provisioned IOPS SSD	Throughput-optimized HDD	Cold HDD
API name	gp2 or gp3	io1	st1	sc1
Max IOPS / volume	16,000	64,000	500	250
Max throughput / volume in MiBps	250 (gp2) 500 (gp3)	1,000	500	250

The following Amazon EBS volume types are recommended:

- General-purpose SSD (gp2 or gp3)
- Provisioned IOPS SSD (io1)
- Throughput-optimized(st1)

Note: The Cold HDD class of storage is *not* recommended for use with IBM Spectrum Virtualize for Public Cloud on AWS.

All volume types appear in the “Enterprise Disks” tier in IBM Spectrum Virtualize for Public Cloud and the tier level should be adapted afterward according to their capabilities.

General-purpose SSD (gp2) volumes are the default volume type for Amazon EBS volumes that are created from the console. The gp2 volumes have a throughput limit of 128 MiBps - 250 MiBps, depending on volume size. These volumes earn I/O credits at the baseline performance rate of 3 IOPS per GiB of volume size. For example, a 100 GiB gp2 volume has a baseline performance of 300 IOPS. When a volume below 1 TiB size requires more than the baseline performance I/O level, it draws on I/O credits in the credit balance to burst to the required performance level.

For more about the Amazon EBS volume types, the I/O credits, and pricing, see the following resources:

- Amazon EBS Volume Types
- Amazon EBS Pricing

4.3.5 AWS cost estimation

THe AWS cost depends on the following factors:

- Bandwidth
- Virtual CPUs and memory
- Storage capacity and performance
- Duration of usage

Amazon Total Cost of Ownership (TCO) helps with AWS cost estimation.

4.3.6 Network and security

AWS uses a *shared responsibility model* where AWS provides a global secure infrastructure and services. AWS customers are responsible for protecting the confidentiality, integrity, and availability of their data in the cloud, and fulfillment of specific business requirements. Careful planning is required for the network environment to fulfill future scalability and performance requirements.

Important: Involve the customer network architect and the AWS architect in the early phases of planning to ensure a successful implementation.

For more information about how AWS keeps your data safe and how AWS meets compliance requirements, see AWS Cloud Security.

For more information about Amazon Security best practices, see *AWS Security Best Practices*.

4.3.7 Data security

Data security and protection can be achieved by encrypting the following types of data:

- Data in motion
- Data in use
- Data at rest

At the time of writing, native encryption is not supported by IBM Spectrum Virtualize for Public Cloud. However, Amazon EBS volumes can be ordered as encrypted or non-encrypted. Use of unencrypted and Amazon EBS volumes that are encrypted with the `(default) aws/ebs` Master key are supported. Amazon EBS volumes that are encrypted with the AWS `Adminkey` are not supported. These `Adminkey` encrypted volumes generate an error and fail when a user attempts to add them to a pool.

The data-at-rest encryption occurs on the servers that host EC2 instances running the IBM Spectrum Virtualize nodes, which provide encryption of data as it moves between EC2 instances and the encrypted Amazon EBS volumes. For more information, see the following resources:

- Amazon EBS features
- AWS IAM

AWS has a full range of security infrastructure as a service (IaaS). Software installation and system management must be integrated under that infrastructure consistently and securely. Three different user types with the suitable IAM roles are available that are required for successful installation and system management, as described in the following sections.

Installer

To install IBM Spectrum Virtualize for Public Cloud from AWS Marketplace, manually predefine an Installer user profile in the AWS IAM service. Any user can be used as the Installer if the user meets the minimum privileges of the Installer profile. Six service-related permissions are needed to deploy an IBM Spectrum Virtualize for Public Cloud cluster on the AWS cloud.

User

A user profile can be defined based on your own IT security policy. However, it is a best practice to limit the permissions of these users to actions that they complete as part of their daily work.

IP quorum management

IP quorum management requires permissions to access quorum-related actions.

For more information about planning an installation on AWS, see IBM Documentation.

4.3.8 Storage performance optimization for AWS

IBM Spectrum Virtualize for Public Cloud assigns workloads to MDisks according to their physical capabilities. Those capabilities must be set manually for external MDisks, such as Amazon EBS volumes.

Easy Tier is a feature that can be used to optimize EBS block storage performance. Easy Tier relocates hot extents to realize optimal performance. By setting the suitable performance tier for each EBS volume, IBM Spectrum Virtualize can use the back-end MDisk according to its capabilities, and not underdrive or overdrive the volume.

By default, all MDisks appear in IBM Spectrum Virtualize for Public Cloud on AWS as “Enterprise” tier with the “default” **easytierload** (medium), as shown in Example 4-4 on page 66.

Example 4-3 Verifying the current MDisk tier settings

```
IBM_IBM Spectrum_Virtualize:REDBOOKS-SV4PC:superuser>lsmdisk 0 | grep tier
tier tier_enterprise
easy_tier_load
IBM_IBM Spectrum_Virtualize:REDBOOKS-SV4PC:superuser>I
```

The assignment to the suitable Easy Tier level that is shown in Table 4-3 is a best practice and must be adjusted manually by using the **chmdisk -tier -easytierload** command.

Table 4-3 AWS assignment to an Easy Tier level

Drive	AWS volume type	Easy Tier level	easytierload command flag
Provisioned IOPS SSD	io1	tier0_flash	**high**
General-purpose SSD	gp2 or gp3	tier1_flash	**low**
Throughput-optimized HDD	st1	tier_nearline	**low**

As shown in Example 4-4, you assign an Amazon EBS gp2 volume to the suitable IBM Spectrum Virtualize for Public Cloud storage tier. Finally, you verify the tier level by running the **lsmdisk** command.

Example 4-4 Changing the MDisk tier level settings

```
IBM_IBM Spectrum_Virtualize:REDBOOKS-SV4PC:superuser>chmdisk -tier tier1_flash
-easytierload low 0
IBM_IBM Spectrum_Virtualize:REDBOOKS-SV4PC:superuser>lsmdisk 0 | grep tier
tier tier1_flash
easy_tier_load low
IBM_IBM Spectrum_Virtualize:REDBOOKS-SV4PC:superuser>
```

4.3.9 Planning for data reduction pools on AWS

Consider the following points regarding the use of data reduction pools (DRPs).

- Minimum pool size (see Table 4-4).
- Pool extent size: 4 GB to accommodate a 512 TB pool.
- Maximum utilization of pool: 85% to allow maximum efficiency in space reclamation.

Table 4-4 DRP extent size, minimum, and maximum

Extent Size (in gigabytes)	Minimum (In terabytes)	Maximum per I/O group
1 GB or smaller	1.1 TB	128 TB
2 GB	2.1 TB	256 TB
4 GB	4.2 TB	512 TB
8 GB	8.5 TB	1024 TB

The minimum size indicates the space that is needed for metadata management, the maximum is a reflection of the 128,000 extent per volume limit (for more information, see 3.6, “AWS solution architecture” on page 41).

Be especially aware of the minimum size requirement when creating a minimum base configuration that provisions 2x512GB Amazon EBS volumes at time of cluster creation. This requirement is below the minimum for a DRP with 1 GB extents, and it is below the default 4 GB extent.

4.4 Planning for Microsoft Azure

The section details important planning considerations for a successful implementation of IBM Spectrum Virtualize for Public Cloud on the Microsoft Azure cloud platform.

4.4.1 Planning security access control on Microsoft Azure

To deploy and manage IBM Spectrum Virtualize for Public Cloud in Microsoft Azure, the user must have a set of specific permissions to access the Microsoft Azure resources.

Microsoft Azure provides role-based access control (RBAC) to manage access to Microsoft Azure resources. RBAC allows fine-grained access control to your resources that are hosted in Azure. For more information, see What is Azure role-based access control (Azure RBAC)? in the Microsoft Azure documentation.

When a subscription is created in Microsoft Azure, a default administrator role is created. The permissions on the default administrator role allow access that is related to subscriptions, management groups, and all RGs that are configured within the subscription. You can use this role to install and manage your IBM Spectrum Virtualize for Public Cloud deployment; however, separate roles provide granular access control and better protection of your resources.

Several user roles can be created to manage resources in your IBM Spectrum Virtualize for Public Cloud deployment. Each of these user roles feature specific permissions that allow or deny access to resources within your deployment. These roles divide actions among several users, which minimize unauthorized access to resources within your environment.

The deployment template also automatically creates roles that provide access between the IBM Spectrum Virtualize for Public Cloud nodes and other objects within your configuration.

Table 4-5 lists the roles that are required for the deployment and management of IBM Spectrum Virtualize for Public Cloud software in Microsoft Azure.

Table 4-5 Required user roles for IBM Spectrum Virtualize for Public Cloud deployments

User role	Description	Tasks allowed by this role
Installer	This role is assigned to the user who is deploying the IBM Spectrum Virtualize for Public Cloud cluster. This user is responsible for installing the cluster through the deployment template, and has permissions that include creating VMs, provisioning virtual networks (vNet), and attaching Azure disks.	Installation step: Deploy IBM Spectrum Virtualize for Public Cloud.
Management	This role provides permissions to complete day-to-day operations on the IBM Spectrum Virtualize for Public Cloud cluster after it is deployed in Microsoft Azure. This role can run system setup and any related configuration tasks in IBM Spectrum Virtualize for Public Cloud management interfaces.	Postinstallation steps: 1. Complete system setup. 2. Configuring pools and assigning storage. 3. Configure volumes. 4. Configure hosts and mappings.
Azure Bastion	This role can create a separate user role to manage all Azure Bastion connections between your public and private network, or include these permissions as part of the user role or the installer user role. An Azure Bastion host allows public networks to access a private vNet.	Postinstallation step: Create an Azure Bastion host.

4.4.2 Installer user role permissions

Before IBM Spectrum Virtualize for Public Cloud is installed, ensure that an installer role is created with the correct permissions. If permissions are not assigned, actions that are required for successful installation of the IBM Spectrum Virtualize for Public Cloud fail.

You can use the Azure default administrator profile to install the IBM Spectrum Virtualize for Public Cloud software, or you can create an installer user profile that includes only the required permissions for deploying the software. When you create permissions in Microsoft Azure, you can select specific permissions in the Azure portal or add permissions in JSON format.

Make sure that an installer user profile includes the following permissions:

```
"Microsoft.Compute/virtualMachines/read",
"Microsoft.Resources/deployments/validate/action",
"Microsoft.Resources/deployments/write",
"Microsoft.Resources/tags/write",
"Microsoft.Compute/proximityPlacementGroups/write",
"Microsoft.Compute/availabilitySets/write",
"Microsoft.Network/virtualNetworks/write",
"Microsoft.Network/virtualNetworks/subnets/write",
"Microsoft.Network/networkSecurityGroups/write",
"Microsoft.ManagedIdentity/userAssignedIdentities/write",
"Microsoft.Authorization/roleDefinitions/write",
"Microsoft.Authorization/roleAssignments/write",
"Microsoft.KeyVault/vaults/write",
"Microsoft.Compute/disks/write",
"Microsoft.Network/networkInterfaces/write",
"Microsoft.Compute/virtualMachines/write",
"Microsoft.Compute/virtualMachines/extensions/write",
"Microsoft.Resources/deployments/operationstatuses/read",
"Microsoft.ManagedIdentity/userAssignedIdentities/read",
"Microsoft.Network/virtualNetworks/read",
"Microsoft.Network/networkSecurityGroups/read",
"Microsoft.Resources/deployments/read",
"Microsoft.KeyVault/vaults/read",
"Microsoft.Network/networkSecurityGroups/join/action",
"Microsoft.Network/virtualNetworks/subnets/read",
"Microsoft.Network/virtualNetworks/subnets/join/action",
"Microsoft.Network/networkInterfaces/join/action",
"Microsoft.ManagedIdentity/userAssignedIdentities/assign/action",
"Microsoft.Compute/virtualMachines/extensions/read",
"Microsoft.Network/LoadBalancers/*",
"Microsoft.Resources/subscriptions/resourceGroups/read",
"Microsoft.Resources/subscriptions/resourceGroups/write",
"Microsoft.Resources/deploymentScripts/read",
"Microsoft.Resources/deploymentScripts/write",
"Microsoft.Resources/deployments/operations/read",
"Microsoft.Resources/subscriptions/resourcegroups/resources/read",
"Microsoft.Network/virtualNetworks/joinLoadBalancer/action",
"Microsoft.Resources/resources/read",
"Microsoft.Resources/deployments/cancel/action",
"Microsoft.Resources/subscriptions/resources/read",
"Microsoft.Compute/disks/read",
"Microsoft.Network/networkInterfaces/read",
```

```
"Microsoft.Network/networkInterfaces/ipconfigurations/read",
"Microsoft.Compute/proximityPlacementGroups/read",
"Microsoft.Compute/availabilitySets/read",
"Microsoft.Network/bastionHosts/read",
"Microsoft.Network/virtualNetworks/Azure BastionHosts/action",
"Microsoft.Network/virtualNetworks/bastionHosts/default/action",
"Microsoft.Network/virtualNetworks/subnets/virtualMachines/read",
"Microsoft.Network/virtualNetworks/virtualMachines/read",
"Microsoft.Compute/virtualMachines/start/action",
"Microsoft.Compute/virtualMachines/restart/action",
"Microsoft.Compute/virtualMachines/deallocate/action",
"Microsoft.Network/publicIPAddresses/write",
"Microsoft.Network/bastionHosts/write",
"Microsoft.Network/publicIPAddresses/read",
"Microsoft.Network/publicIPAddresses/join/action",
"Microsoft.Network/publicIPAddresses/dnsAliases/read",
"Microsoft.ContainerInstance/containerGroups/write",
"Microsoft.ContainerInstance/containerGroups/read",
"Microsoft.ContainerInstance/containerGroups/delete",
"Microsoft.ContainerInstance/containerGroups/start/action",
"Microsoft.ContainerInstance/containerGroups/stop/action",
"Microsoft.ContainerInstance/containerGroups/restart/action",
"Microsoft.Storage/storageAccounts/read",
"Microsoft.Storage/storageAccounts/write",
"Microsoft.Storage/storageAccounts/listkeys/action",
"Microsoft.Storage/storageAccounts/delete"
```

4.4.3 Management user role permissions

To create and manage IBM Spectrum Virtualize for Public Cloud operations, another user role can be created to complete these management tasks. The `SV_Cloud_User_Role` provides permissions to a user that completes the daily configuration and management tasks of your IBM Spectrum Virtualize for Public Cloud cluster.

The `SV_Cloud_User_Role` can be defined with the following permissions:

```
"Microsoft.Compute/virtualMachines/read",
"Microsoft.Resources/deployments/validate/action",
"Microsoft.Resources/deployments/operationstatuses/read",
"Microsoft.ManagedIdentity/userAssignedIdentities/read",
"Microsoft.Compute/virtualMachines/start/action",
"Microsoft.Compute/virtualMachines/restart/action",
"Microsoft.Compute/virtualMachines/deallocate/action",
"Microsoft.Network/virtualNetworks/read",
"Microsoft.Network/networkSecurityGroups/read",
"Microsoft.Resources/deployments/read",
"Microsoft.KeyVault/vaults/read",
"Microsoft.Network/networkSecurityGroups/join/action",
"Microsoft.Network/virtualNetworks/subnets/read",
"Microsoft.Network/virtualNetworks/subnets/join/action",
"Microsoft.Network/networkInterfaces/join/action",
"Microsoft.ManagedIdentity/userAssignedIdentities/assign/action",
"Microsoft.Compute/virtualMachines/extensions/read",
"Microsoft.Network/LoadBalancers/*",
"Microsoft.Resources/subscriptions/resourceGroups/read",
```

```
"Microsoft.Resources/deploymentScripts/read",
"Microsoft.Resources/deployments/operations/read",
"Microsoft.Resources/subscriptions/resourcegroups/resources/read",
"Microsoft.Network/virtualNetworks/joinLoadBalancer/action",
"Microsoft.Resources/resources/read",
"Microsoft.Resources/deployments/cancel/action",
"Microsoft.Resources/subscriptions/resources/read",
"Microsoft.Compute/disks/read",
"Microsoft.Network/networkInterfaces/read",
"Microsoft.Network/networkInterfaces/ipconfigurations/read",
"Microsoft.Compute/proximityPlacementGroups/read",
"Microsoft.Compute/availabilitySets/read",
"Microsoft.Network/bastionHosts/read",
"Microsoft.Network/virtualNetworks/Azure BastionHosts/action",
"Microsoft.Network/virtualNetworks/bastionHosts/default/action",
"Microsoft.Network/virtualNetworks/subnets/virtualMachines/read",
"Microsoft.Network/virtualNetworks/virtualMachines/read",
"Microsoft.Compute/virtualMachines/start/action",
"Microsoft.Compute/virtualMachines/restart/action",
"Microsoft.Compute/virtualMachines/deallocate/action",
"Microsoft.Network/publicIPAddresses/write",
"Microsoft.Network/bastionHosts/write",
"Microsoft.Network/publicIPAddresses/read",
"Microsoft.Network/publicIPAddresses/join/action",
"Microsoft.Network/publicIPAddresses/dnsAliases/read"
```

4.4.4 Azure Bastion user role permissions

If you plan to use an Azure Bastion Service to connect to your deployment, you can create another user role or add the necessary permissions to the installer user profile. The `Azure Bastion_User_Role` can be defined with the following permissions:

```
"Microsoft.Network/bastionHosts/read",
"Microsoft.Network/virtualNetworks/Azure BastionHosts/action",
"Microsoft.Network/virtualNetworks/bastionHosts/default/action",
"Microsoft.Network/networkInterfaces/read",
"Microsoft.Network/networkInterfaces/ipconfigurations/read",
"Microsoft.Network/virtualNetworks/read",
"Microsoft.Network/virtualNetworks/subnets/virtualMachines/read",
"Microsoft.Network/virtualNetworks/virtualMachines/read",
"Microsoft.Network/publicIPAddresses/write",
"Microsoft.Network/bastionHosts/write",
"Microsoft.Network/publicIPAddresses/read",
"Microsoft.Network/publicIPAddresses/join/action",
"Microsoft.Network/publicIPAddresses/dnsAliases/read"
```

For more information, see Configure Azure Bastion from VM settings in the Azure documentation.

Note: Permissions to the Azure Bastion Service are controlled by Microsoft Azure and are subject to change. For more information, see the Azure documentation.

4.4.5 Planning networking for Microsoft Azure

As part of the basic deployment of IBM Spectrum Virtualize for Public Cloud in Microsoft Azure, the installation template provisions and configures all necessary networking objects for a private vNet. If you are implementing your environment exclusively in Microsoft Azure as an all-in-cloud deployment, more networking objects are needed, such as an Azure Bastion server and Windows VM to access the management interfaces for IBM Spectrum Virtualize for Public Cloud. If you plan to extend your configuration to include access to a public network over the internet, such as in a hybrid-cloud use case, more network planning and configuration are required.

4.4.6 Network considerations for basic deployment

IBM Spectrum Virtualize for Public Cloud deployment allows you to create your configuration within a vNet or create a vNet during deployment.

If you select an existing vNet, ensure that the RG that contains the IBM Spectrum Virtualize for Public Cloud installation does not contain any other resources. As part of the all-in-cloud deployment of IBM Spectrum Virtualize for Public Cloud in Microsoft Azure, the following networking objects are configured based on your values in the deployment template:

- Two VMs for nodes

  During installation, you can select the size and type of D-Series VM to use as redundant nodes for your IBM Spectrum Virtualize for Public Cloud cluster. As part of deployment, each VM includes a dedicated network interface card (NIC). For more information, see 4.4.8, “Planning an Azure virtual machine” on page 73.

- One VM for a quorum node

  During installation, a single B-Series VM is deployed for the quorum node. A dedicated NIC also is deployed for the quorum node. The quorum node is used to maintain redundancy if one of the cluster nodes is unavailable.

- A private subnet for the quorum node

  The template deploys a private subnet that includes IP addresses and ranges for the quorum node.

- Standard load balancer

  As part of the installation, an Azure load balancer is automatically provisioned for your configuration. The standard load balancer connects the cluster IP address to new back-end IP address on IBM Spectrum Virtualize for Public Cloud nodes. The load balancer is used for the management IP address only and the traffic is routed on all the ports.

- NSGs

  In addition, NSGs are automatically configured for each subnet to protect the vNet. An NSG filters network traffic to and from resources in the vNet. It uses specific security rules that allow or deny inbound and outbound traffic to and from the vNet.

  As part of basic deployment, specific rules are defined automatically that protect the private network and limit connections to and from public networks. If you plan to extend your configuration to connect to public networks, these rules must be updated.

  For more information, see network security groups in the Azure documentation.

The rules that are included in NSGs during a basic deployment are listed in Table 4-6.

Table 4-6 Default rules for Network Security Groups

Type	Protocol	Port range	Use
SSH	TCP	22	SSH traffic to a node instance
Port Node Discovery	UDP	21451 - 21452	Node discovery traffic
Customer TCP rule	TCP	21450	Node-to-node communication traffic
Customer TCP rule	TCP	3260	iSCSI target discovery, login, and IP replication traffic
Customer TCP rule	TCP	3265	IP replication traffic
Customer TCP rule	TCP	8443	Redirects for port 443
HTTPS	TCP	443	HTTPS inbound traffic
Customer TCP rule	TCP	1260	IP quorum traffic
Deny All vNet Inbound	Any	Any	Deny all except allowed
Allow vNet Inbound	Any	Any	Default rule
Deny All Inbound	Any	Any	Azure default rule
Allow vNet Outbound	Any	Any	Azure default rule
Allow Internet Outbound	Any	Any	Azure default rule
Deny All Outbound	Any	Any	Azure default rule

4.4.7 Network considerations for cross-public network deployments

If you are extending your basic deployment to access private networks across the internet, more planning and configuration steps are necessary.

Ensure that your current network infrastructure provides the acceptable security and performance for cloud-based environments. In general, before implementing a hybrid-cloud or multi-cloud solution with an on-premises data center, storage administrators must work with their networking administrators to determine whether a virtual private network (VPN) connection exists between the on-premises data center and Microsoft Azure. If not, a VPN or IPsec tunnel must be configured. Also, the current infrastructure must be updated to include a supported VPN device. Microsoft Azure supports many types of VPN and IPsec providers.

Ensure that the following prerequisites are completed before creating a site-to-site VPN between the on-premises site and the site in the Microsoft Azure cloud:

- The on-premises network uses one of the validated VPN devices that Microsoft Azure supports. Microsoft Azure provides more information about each provider and supported features. For more information, see Validated VPN devices and device configuration guides.

  If your current VPN device is not a validated device, contact your VPN service provider for assistance in setting up your on-premises network with Microsoft Azure.
- The on-premises VPN device is configured with an external public IPv4 address.
- The security requirements are determined for cryptographic keys that are generated for VPN connections for hybrid environments. Microsoft Azure provides a list of validated VPN devices and a list of IPsec or internet key exchange parameters for virtual private network gateways (VPNGW) for each of these devices. For more information, see the following topics at the Microsoft Azure documentation web page:
  - About cryptographic requirements and Azure VPN gateways
  - About VPN devices and IPsec or internet key exchange parameters for Site-to-Site VPN Gateway connections
- Coordinate with your network administrator to ensure that you have accurate IP address ranges for your on-premises network. During configuration of a site-to-site VPN connection, you must specify the IP address range prefixes that Microsoft Azure uses to route to your on-premises environment. In addition, all internal subnets for your on-premises network cannot overlap with any of the vNet subnets to which you are connecting.

For more information, see Extend an on-premises network using VPN in the Microsoft Azure documentation.

In addition to planning your on-premises network, you must change the security rules that are configured as part of basic deployment to allow connections to and from your on-premises VPN connection.

For more information about customizing network security rules, see Create, change, or delete an NSG.

4.4.8 Planning an Azure virtual machine

Determine the size of the Azure VMs to match your expected workloads. IBM Spectrum Virtualize for Public Cloud deployment supports the following types of D-Series VMs for node operations:

- Standard_D16s_v3
- Standard_D32s_v3
- Standard_D64s_v3

Another fixed size B-Series (B1ms VM size) is provisioned for quorum management.

IBM Spectrum Virtualize for Public Cloud nodes

As part of the deployment of IBM Spectrum Virtualize for Public Cloud software, you can choose from three types of D-Series VMs to use for your IBM Spectrum Virtualize for Public Cloud nodes instances. These VMs are deployed as pairs to provide node failover and redundancy.

Also, a basic load balancer is provisioned during installation to manage the I/O to each node. Each supported VM type features the following properties:

- VM Size

  Indicates the total amount of capacity that is allocated to the VM.

- vCPU

  Indicates the amount of vCPU per VM. One or more vCPUs are assigned to every VM within a cloud environment. Each vCPU is seen as a single physical CPU core by the VM's operating system.

- Family

  Defines the usability for the VM type. For example, general-purpose VM sizes provide balanced CPU-to-memory ratio, which ideal for testing and development, small to medium databases, and low to medium traffic web servers.

- RAM

  Indicates the amount of memory that is allocated to the VM.

- Data Disk

  Indicates the number of virtual HDDs that are attached to the VM.

- Max IOPS

  Indicates the maximum I/O operations per second that the VM can process.

- Temporary Storage

  Indicates the size (in GiB) of temporary storage that provides short-term storage for applications and processes. All VMs include a temporary drive.

Table 4-7 compares the properties of each D-Series type VM that can be selected for your IBM Spectrum Virtualize for Public Cloud nodes.

Table 4-7 Properties of each D-Series type VM

VM Size	Family	vCPUs	RAM (GiB)	Data disks	Max IOPS	Temporary storage (GiB)
Standard_D16s_v3	General-purpose	16	64	32	25600	128
Standard_D32s_v3	General-purpose	32	128	32	51200	256
Standard_D64s_v3	General-purpose	64	256	32	80000	512

For more information about the supported D-Series VMs, such as pricing information, see D-Series comparison in the Azure documentation.

Quorum node

As part of a deployment, a single, fixed-size B1ms VM is provisioned for quorum management for the IBM Spectrum Virtualize for Public Cloud cluster. This VM hosts the IP quorum application that determines which node handles I/O if the connection between the nodes is lost. This VM type is ideal for this use because it provides low-cost option for low or moderate workloads.

Table 4-8 lists the properties for the VM that is provisioned for the quorum node.

Table 4-8 Properties of B-Series type VM

VM Size	Family	vCPUs	RAM (GiB)	Data disks	Max IOPS	Temporary storage (GiB)
B1ms	General-purpose	1	2	2	640	4

4.4.9 Planning Azure managed disks

Azure managed disks (MDisks) are block-level storage volumes that provide disk-based data storage for the IBM Spectrum Virtualize for Public Cloud deployment. You can select the type of Azure MDisk that are used with IBM Spectrum Virtualize for Public Cloud for storage provisioning.

A minimum of two Azure MDisks are required for initial cluster creation from Azure Marketplace. After the installation, disks can be created in the same RG where the cluster was created by using the Azure portal or the Azure command-line interface (CLI). The cluster automatically detects such disks, which can be used later for creating storage pools. The cluster can support up to 31 MDisks.

For more information about Azure MDisks pricing, see this web page.

4.4.10 Attaching MDisks

IBM Spectrum Virtualize for Public Cloud supports only a two-node cluster in Microsoft Azure. You can attach Azure MDisks to both nodes in the cluster by using the Azure shared disks feature. Attaching an MDisk to multiple VMs allows you to deploy new or migrate existing clustered applications to Azure.

Azure shared disks include a `maxShares` property value that signifies the maximum VMs that can be attached to a MDisk simultaneously. Enable the Azure shared disks feature on the disk by using the Azure portal or Azure CLI. When enabling the feature, set `maxShares=2`. If the `maxShares` value is less than 2, it cannot be attached to IBM Spectrum Virtualize for Public Cloud node cluster.

Note: When the Azure shared disks feature is used on the premium SSD or standard SSD MDisks, each extra mount of the MDisk is charged per month based on the disk size.

4.4.11 MDisk support

IBM Spectrum Virtualize for Public Cloud supports following types of Azure MDisks:

- Standard SSD

  This type of Azure MDisk can be used for web servers, lightly used enterprise applications, and testing. It includes the following features:

  - Disk Type: Standard SSD
  - Max Disk Size (GiB): 32,767
  - Max throughput (MBps): 750
  - Max IOPS: 6,000

- Premium SSD

  This type of Azure MDisk can be used for the production and performance-sensitive workloads. It includes the following features:

  - Disk Type: Premium SSD
  - Max Disk Size (GiB): 32,767
  - Max throughput (MBps): 900
  - Max IOPS: 20,000

4.4.12 Planning deployment access

The basic deployment of IBM Spectrum Virtualize for Public Cloud creates a private vNet exclusively. To access the private network from outside of that network, the planning and configuration of an access method for your environment is required.

Depending on the use case of your deployment, different methods can be used to access the private vNet on which your IBM Spectrum Virtualize for Public Cloud cluster is deployed. Access methods can overlap different use cases.

The tables in this section describe different supported access methods, their corresponding use cases, and the permissions that are required for users who are configuring resources.

Table 4-9 lists a use case of a basic deployment in a cloud model in which all setup and configuration occurs with the basic deployment of IBM Spectrum Virtualize for Public Cloud.

Table 4-9 Access method for a basic deployment in a cloud use case

Access methods	Required for use case?	Azure permissions	Description
Windows VM	Yes	Installer user role permissions.	A Windows host must be deployed to enable access to the management GUI and management interfaces for IBM Spectrum Virtualize for Public Cloud deployment.
Azure Bastion service	Yes	For more information, see Permissions for Azure Bastion service in the Azure documentation.	An Azure Bastion host can be created on the quorum node in your deployment or with the Azure Bastion Service.

Table 4-10 lists the use case of an all-in-cloud with replication model in which resources are in two private vNets that are separated by a public network. Data replication is configured between the two private vNets. You have the choice of two access methods: vNet peering and a VPN connection.

Table 4-10 Access method for an all-in-cloud with replication use case

Access methods	Required for use case?	Azure permissions	Description
Windows VM	Yes	Installer user role permissions.	A Windows host must be deployed to enable access to the management GUI and management interfaces for IBM Spectrum Virtualize for Public Cloud deployment.
vNet peering	Yes	For more information, see Permissions for vNet Peering in the Azure documentation.	vNet peering creates secure connections between different vNets within Microsoft Azure. In this use case, vNet peering configuration is simplified, and extends the private network between the two endpoints.
Site-to-site VPN	No	For more information, see Create a Site-to-Site connection in the Azure documentation.	As an alternative, you also can configure a peer-to-peer VPN connection that secures data that is replicated between the separate private vNets. You might want to use a VPN if your security policy requires more protection.

Table 4-11 lists a use case of a hybrid cloud with replication model in which data is replicated from an on-premises data center to an IBM Spectrum Virtualize for Public Cloud cluster in Microsoft Azure. If the production site becomes unavailable, the IBM Spectrum Virtualize for Public Cloud cluster can act as a recovery site.

Table 4-11 Access method for Hybrid cloud with replication use case

Access method	Required for use case?	Azure permissions	Description
Site-to-site VPN	Yes	For more information, see Extend an on-premises network using VPNCreate a site-to-site VPN in the Microsoft Azure documentation.	A VPN connection between the on-premises data center and the recovery site in Microsoft Azure is required to protect data transfer across the public network.

4.4.13 Storage performance optimization

IBM Spectrum Virtualize for Public Cloud assigns MDisk characteristics according to their physical capabilities. Those capabilities must be set manually for external MDisks, such as Azure MDisks.

IBM Easy Tier is a solution that you can use to optimize the most valuable storage usage and maximize Cloud Block Storage performance. Those settings are used on Easy Tier for hot extent relocation and optimal performance. By selecting the suitable tier for the Azure MDisk, IBM Spectrum Virtualize can use the MDisk according to its capabilities, and not under-drive or overdrive the volume.

By default, all MDisks appear in IBM Spectrum Virtualize for Public Cloud on Microsoft Azure as `Enterprise` tier with the default `easytierload` (medium). The assignment to the suitable Easy Tier level that is listed in Table 4-12 is a best practice and must be adjusted manually by using the **`chmdisk -tier -easytierload`** command.

Table 4-12 Microsoft Azure MDisks assignment to an Easy Tier level

Drive	Easy Tier level	easytierload command
Premium SSD	`tier0_flash`	**`high`**
Standard SSD	`tier1_flash`	**`low`**

After assigning an Azure MDisk to the suitable IBM Spectrum Virtualize for Public Cloud storage tier, run the **`lsmdisk`** command to verify the change in the tier level.

4.4.14 Planning for data reduction pools

A deduplicated volume or volume copy can be created in a DRP. When you implement deduplication, you must consider specific requirements in the storage environment.

Deduplication can be configured with volumes that use different capacity-saving methods, such as thin-provisioning. Deduplicated volumes must be created in DRPs for added capacity savings.

Deduplication is a type of data reduction that eliminates duplicate copies of data. User data deduplication occurs within a DRP and only between volumes or volume copies that are marked as deduplicated.

You can migrate any type of volume from a standard pool to a DRP. You can use volume mirroring to migrate data from a volume in a regular storage pool to a deduplicated volume in a DRP. To create a deduplicated volume copy of a volume in a standard pool in a DRP, the following options can be used:

- Add Volume Copy page in the management GUI
- The **`addvdiskcopy`** command

Note: Nodes must have a minimum of 32 GB of memory to support deduplication.

Consider the following points:

- Avoid GMCV to or from a deduplicated volume.
- You can use the Data Reduction Estimation Tool (DRET) to estimate how much capacity you might save if a standard volume that a host can access was a deduplicated volume. The tool scans target workloads on all attached storage arrays, consolidates these results, and generates an estimate of potential data reduction savings for the entire system.

  For more information about DRET, see this IBM Support web page.

  For more information about Compestimator, see this IBM Support web page.
- To ensure that your intended use of deduplicated volumes includes adequate performance for your application, see *IBM SAN Volume Controller Best Practices and Performance Guidelines for IBM Spectrum Virtualize Version 8.4.2*, SG24-8509.
- For more information about DRPs, see *Introduction and Implementation of Data Reduction Pools and Deduplication*, SG24-8430.

The usage of DRPs and deduplication on IBM Spectrum Virtualize for Public Cloud in Microsoft Azure is supported on all three VM types (see Table 4-7 on page 74).

5

Implementing IBM Spectrum Virtualize for Public Cloud on Amazon Web Services

This chapter describes how to implement an IBM Spectrum Virtualize for Public Cloud on Amazon Web Services (AWS) cloud environment.

This chapter includes the following topics:

5.1 Implementing IBM Spectrum Virtualize for Public Cloud on Amazon Web Services

This section describes implementing IBM Spectrum Virtualize for Public Cloud on AWS. The IBM Spectrum Virtualize for Public Cloud on AWS implementation starts from the assumption that the required IBM Storage Virtualize Public Cloud licenses were purchased and you can access IBM Passport Advantage.

Designed for software-defined environments (SDEs), IBM Spectrum Virtualize for Public Cloud on AWS represents a solution for public cloud implementations. It also includes technologies that complement and enhance public cloud offering capabilities.

IBM Spectrum Virtualize for Public Cloud on AWS provides for the deployment of IBM Spectrum Virtualize software in public clouds inside AWS. IBM Spectrum Virtualize for Public Cloud on AWS includes a monthly license to deploy and uses IBM Spectrum Virtualize for Public Cloud on AWS to enable hybrid cloud solutions, which offer the ability to have storage as a service in a multicloud environment.

Table 5-1 lists the IBM Spectrum Virtualize for Public Cloud on AWS product that is available on Amazon Cloud.

Table 5-1 IBM Spectrum Virtualize for Public Cloud on AWS at a glance

Items	On AWS
Storage supported	Amazon Elastic Block Store (EBS)
Licensing approach	Simple, flat cost per managed terabyte and monthly licensing
Platform	IBM Spectrum Virtualize for Public Cloud on AWS installed on an Elastic Compute Cloud (EC2) instance

5.1.1 Installing Storage Virtualize for Public Cloud on AWS

The IBM Spectrum Virtualize for Public Cloud on AWS installation uses AWS CloudFormation templates that simplify provisioning and management on AWS. These templates are available on AWS Marketplace and simplify the provisioning and installation process.

Note: Before installing IBM Spectrum Virtualize for Public Cloud on AWS, ensure that you have a valid IBM Customer Number with licenses, and have a local key pair (`ssh_key`) ready or define it inside AWS before the installation begins. This key is used to access the Azure Bastion host, any other EC2 instances that are created, and other key-based authentication.

Ensure that all prerequisites are complete before you install IBM Spectrum Virtualize for Public Cloud from AWS Marketplace. For more information, see AWS Marketplace.

Deployment video: As part of this IBM Redbooks publication, the authors also created an IBM Spectrum Virtualize for Public Cloud on AWS deployment video. You can see all the following steps in the video.

To install the IBM Spectrum Virtualize for Public Cloud software, complete the following steps:

1. Go to the IBM Spectrum Virtualize for Public Cloud BYOL Marketplace Offering (or search for “IBM Spectrum Virtualize” on the AWS marketplace).

   Log in with your AWS account, as shown in Figure 5-1.

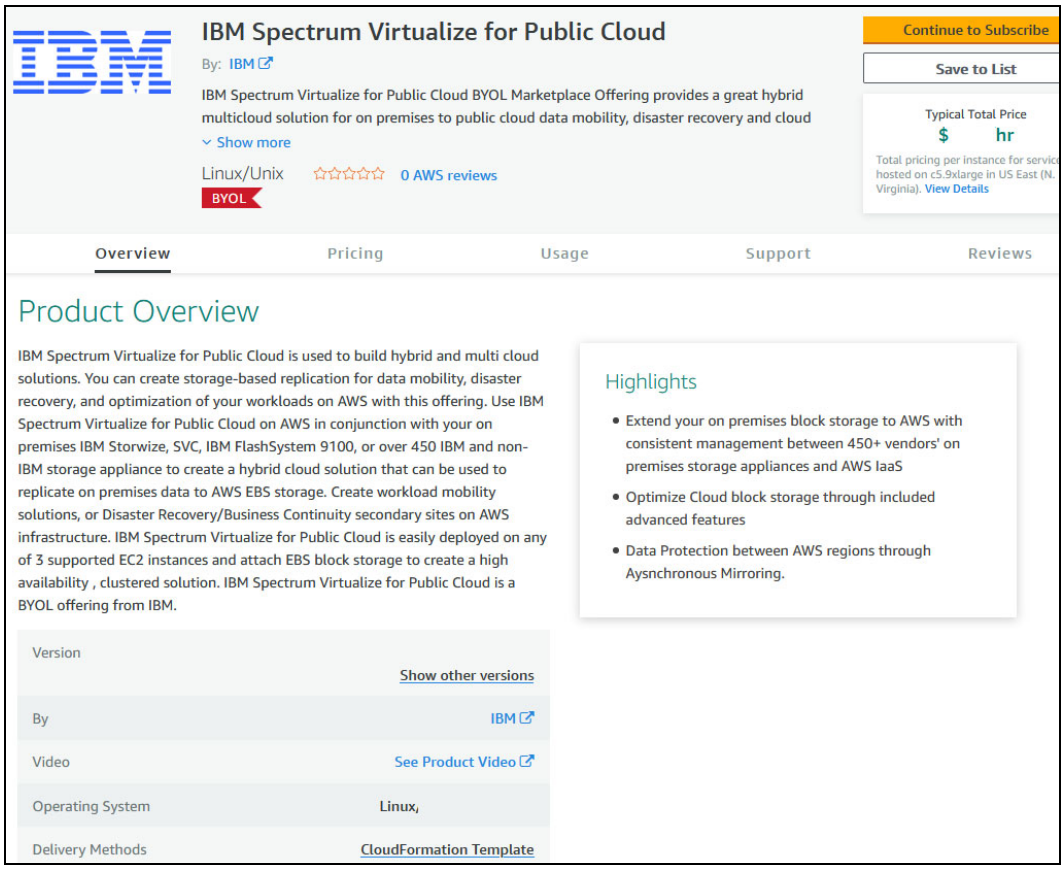

Figure 5-1 IBM Spectrum Virtualize for Public Cloud: AWS Marketplace page

2. Scroll down from the **Overview** section to **Pricing** or select it from the menu items at the top of the Marketplace page. Enter or validate the following information for your installation:

   – Region
   – Fulfillment Option (by using a virtual private cloud [VPC] or a new VPC)
   – EC2 Instance type (the default is c5.9xlarge)

AWS Marketplace provides a dynamic pricing display that is based on your selections. If you are satisfied with your selections, click **Continue to Subscribe** in the upper right of the page and follow the instructions, as shown in Figure 5-2.

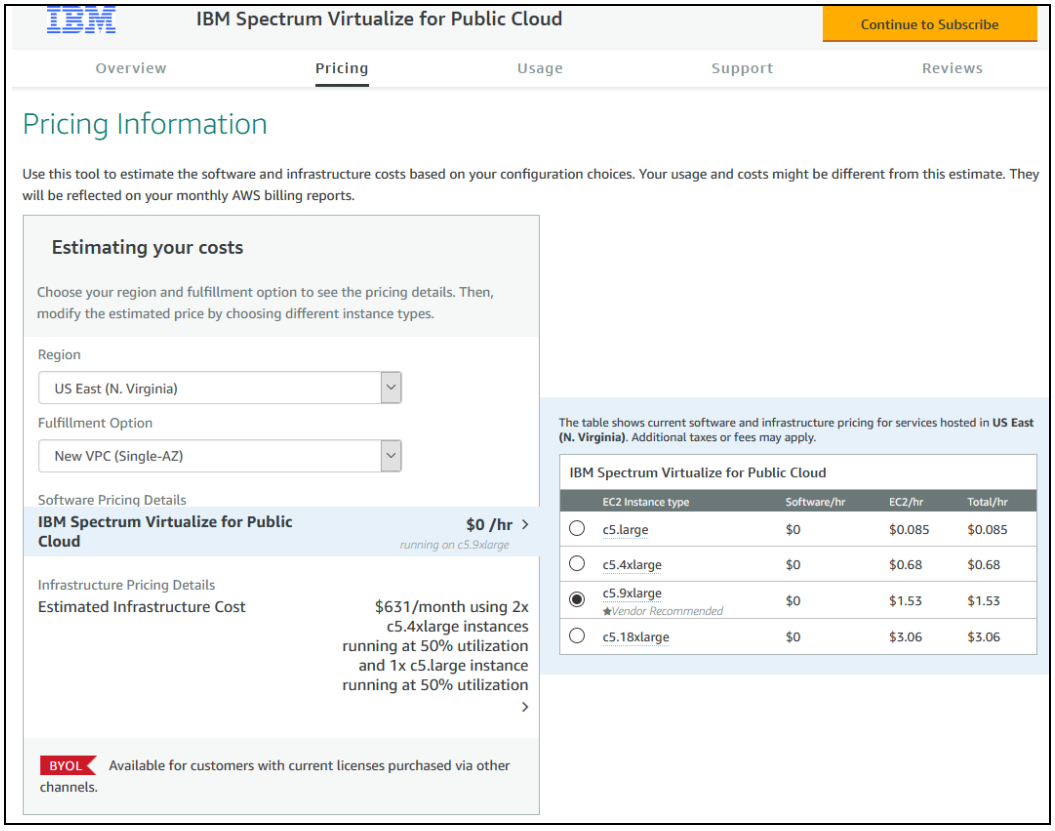

Figure 5-2 IBM Spectrum Virtualize for Public Cloud on AWS: Marketplace Pricing Summary

3. The Terms and Conditions window opens and shows the Product information, as shown in Figure 5-3. After you are satisfied with the results, click **Continue to Configuration**.

Figure 5-3 IBM Spectrum Virtualize for Public Cloud: AWS Marketplace Terms and Conditions

4. Specify whether a new VPC is requested or if you want to deploy into an existing VPC. If the deployment into a new VPC option is selected, it reminds you that only a single Availability Zone is provided by default. Therefore, the Azure Bastion host and the initial IP quorum are placed in the same Availability Zone as the IBM Spectrum Virtualize nodes.

   We elaborate on the significance of this setup and the steps to remediate it by adding another subnet in a different Availability Zone, starting a private network-only EC2 instance into that new Availability Zone, and installing the IP quorum application on that server. The software version also can be changed if others are available, and the region of AWS for the deployment.

   In our example, select the existing VPC option. After all the options are finalized, click **Continue to Launch** (see Figure 5-4).

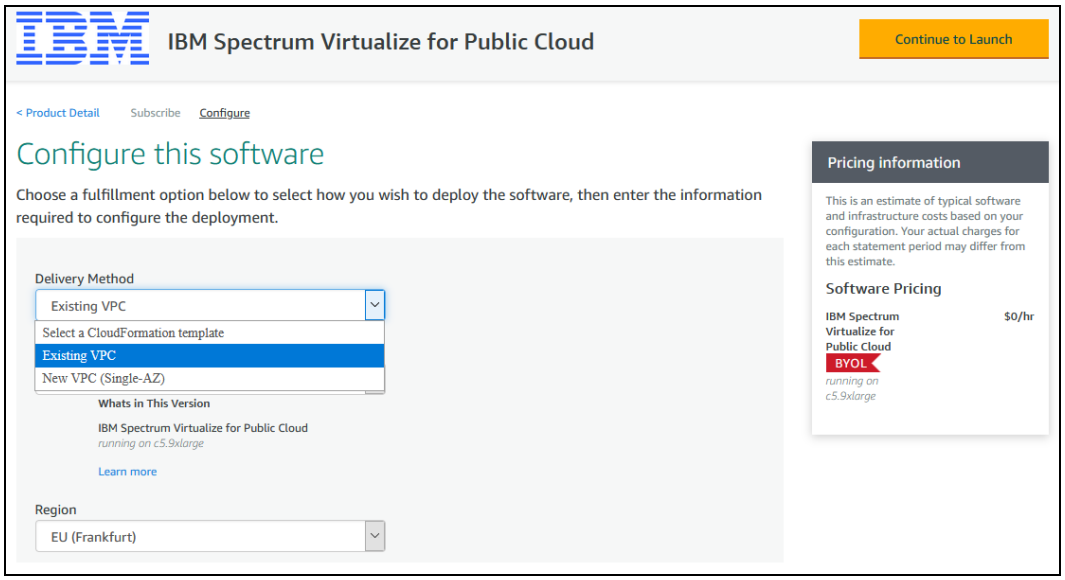

Figure 5-4 IBM Spectrum Virtualize for Public Cloud: Select VPC destination and region

5. The CloudFormation template opens, which automates the rest of the installation process after some key parameters are entered. Because the default action is to launch the CloudFormation process, click **Launch** (see Figure 5-5).

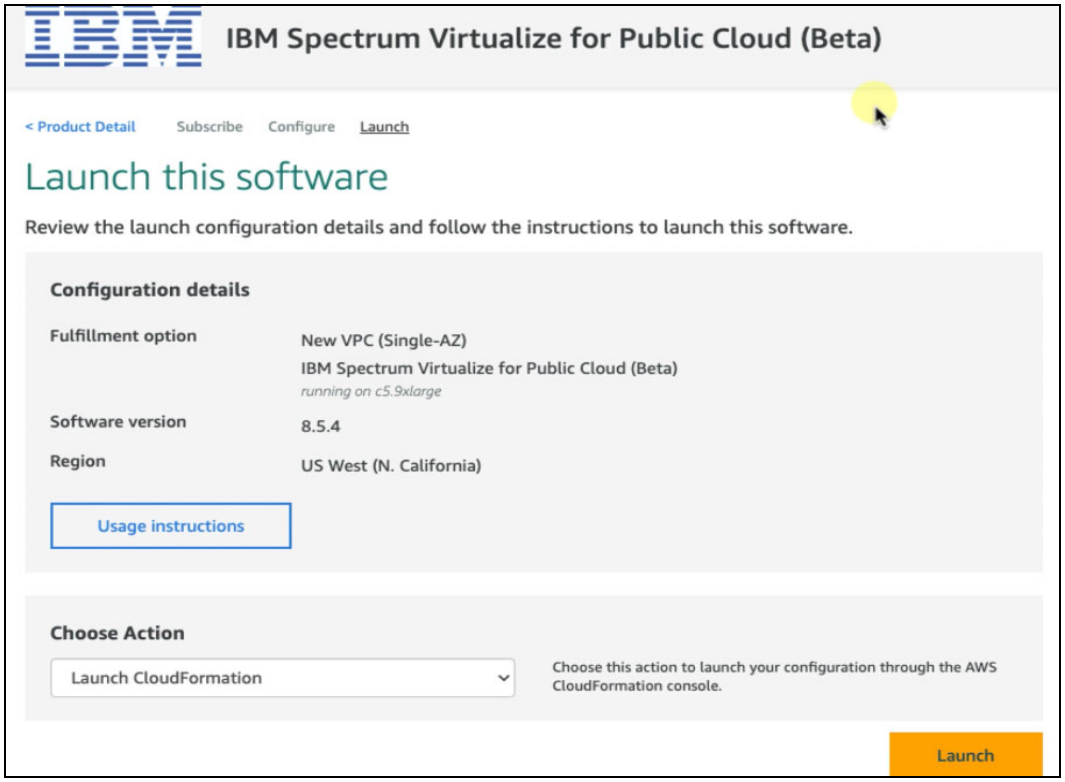

Figure 5-5 IBM Spectrum Virtualize for Public Cloud about to launch the CloudFormation template to deploy the IBM Spectrum Virtualize for Public Cloud stack

6. Create the stack. For this page, use the default settings and *do not change the Amazon Simple Storage Service (S3) URL*. This template location is provided by IBM Spectrum Virtualize for Public Cloud and contains critical information for installation automation. Click **Next** (see Figure 5-6 on page 87).

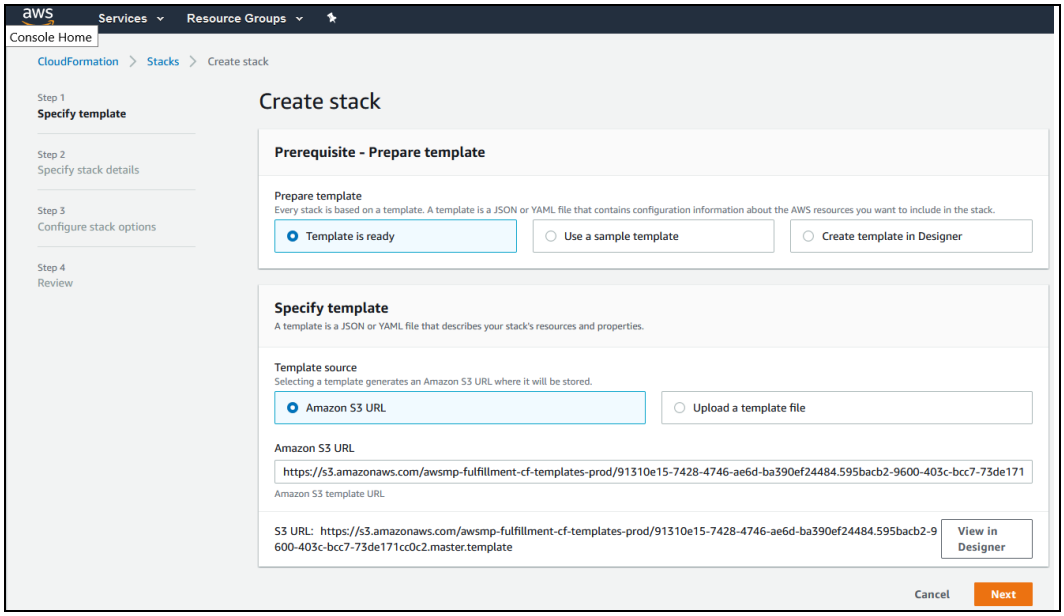

Figure 5-6 Starting the IBM Spectrum Virtualize for Public Cloud: CloudFormation stack creation process

7. Enter the stack name that is the basis of the IBM Spectrum Virtualize for Public Cloud cluster or system name. Specify the node performance type and I/O groups. Continue with each section (see Figure 5-7).

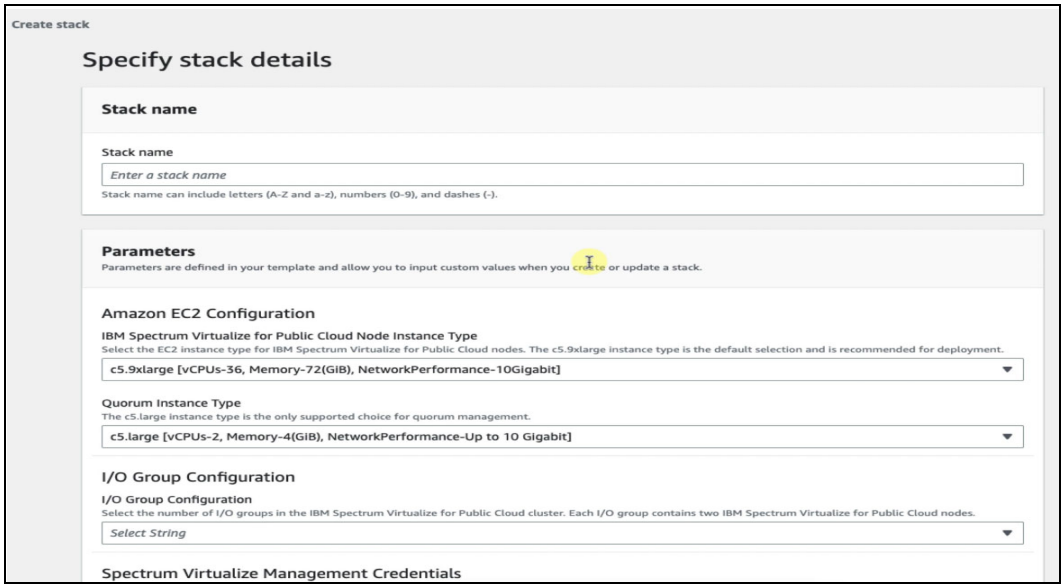

Figure 5-7 IBM Spectrum Virtualize for Public Cloud Stack details: Name, Node Types, and I/O groups

8. Enter your customer number and email for the deployment report. Both of these items are critical to successful deployment and access. The deployment fails on the Quorum node if a customer number does not validate the confirmed licenses. (see Figure 5-8).

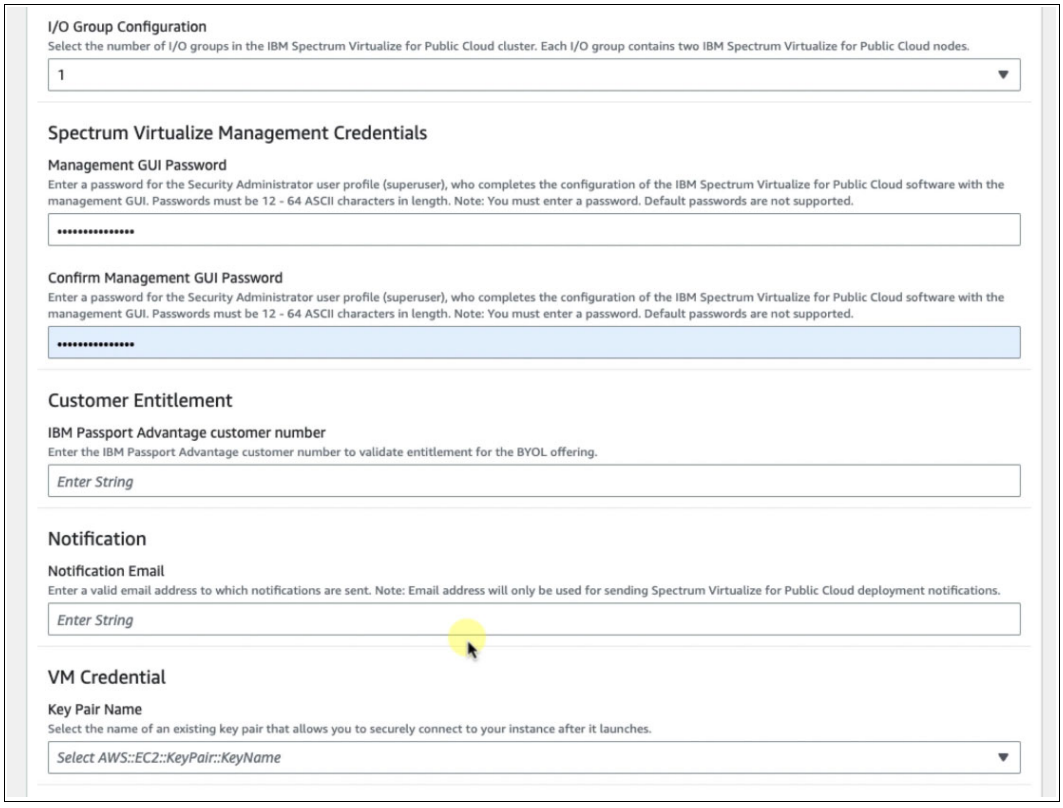

Figure 5-8 IBM Spectrum Virtualize for Public Cloud: Stack Customer Number, Notification Email, and Key Pair

9. Continue with the network settings, including the IP access filter, as needed. Usually, the defaults are suitable for most new deployments, but check with your local on-premises network technicians for any virtual private network (VPN) considerations, and adjust as needed. For this example, we did not restrict which IP addresses were allowed and used a public IP address, so we used 0.0.0.0, as shown in Figure 5-9.

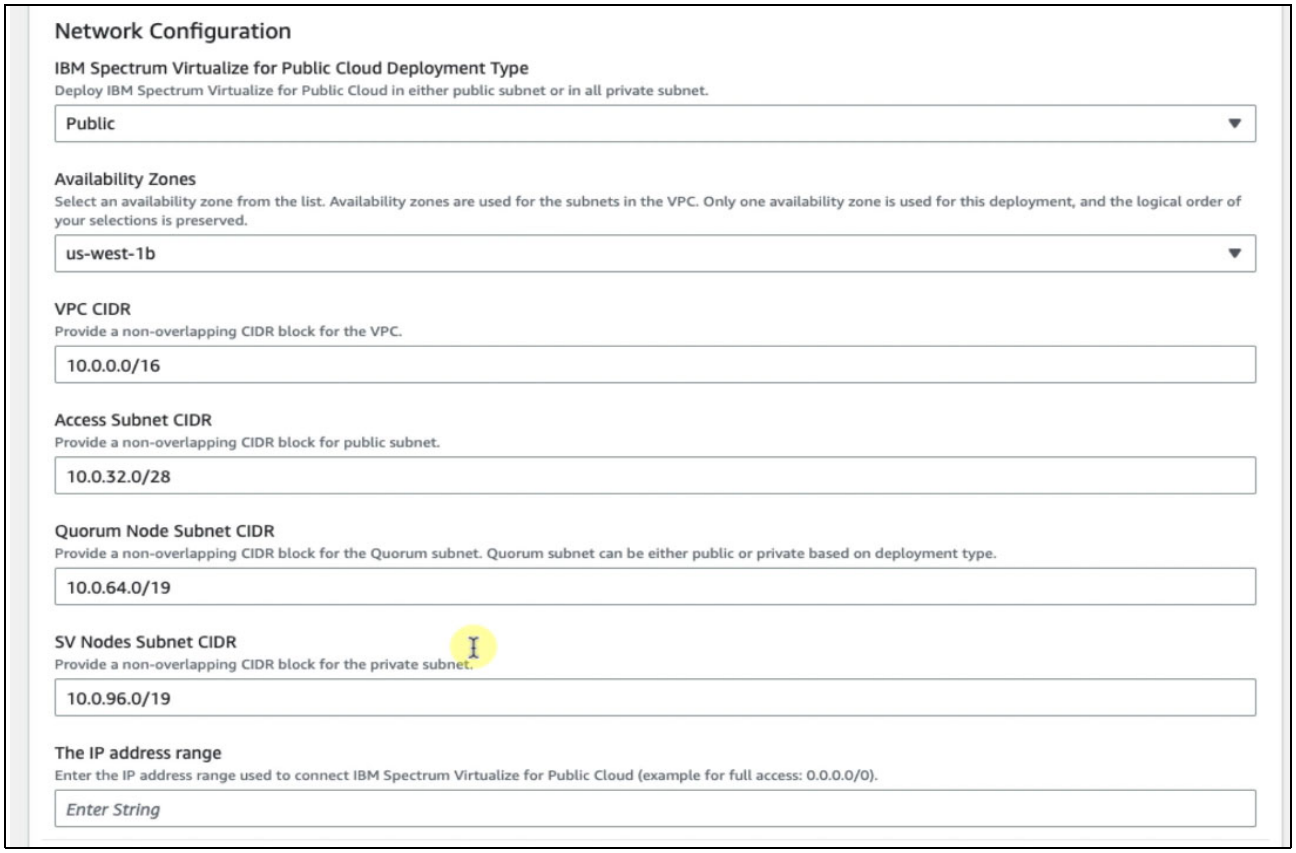

Figure 5-9 IBM Spectrum Virtualize for Public Cloud: Network and IP address filters configuration

10. Continuing with the configuration page that is shown in Figure 5-10, specify the size of the two Amazon EBS gp2 volumes, which will be configured in the IBM Spectrum Virtualize for Public Cloud pool as part of the cluster creation. Accept the license and click **Next**.

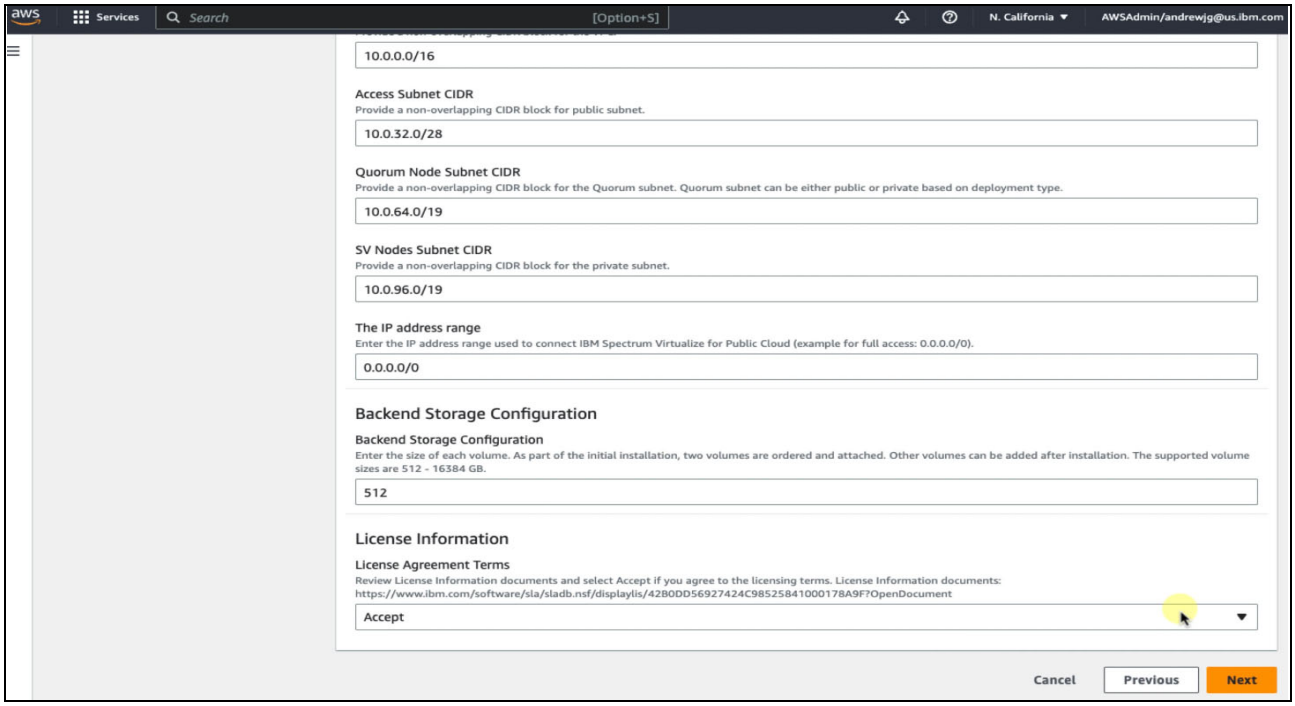

Figure 5-10 IBM Spectrum Virtualize for Public Cloud: Stack Back-end EBS storage capacity selection and license acceptance

11. The review and acknowledgment pages are shown next. Review your prior selections and edit them if necessary. The **Rollback** option helps clean up the deployed items if there is a rollback failure; however, if deeper analysis is required for support, clear it and click **Preserve**, as shown in Figure 5-11. The rest of the page is advanced options that often are not needed, as shown in Figure 5-12. Click **Next**.

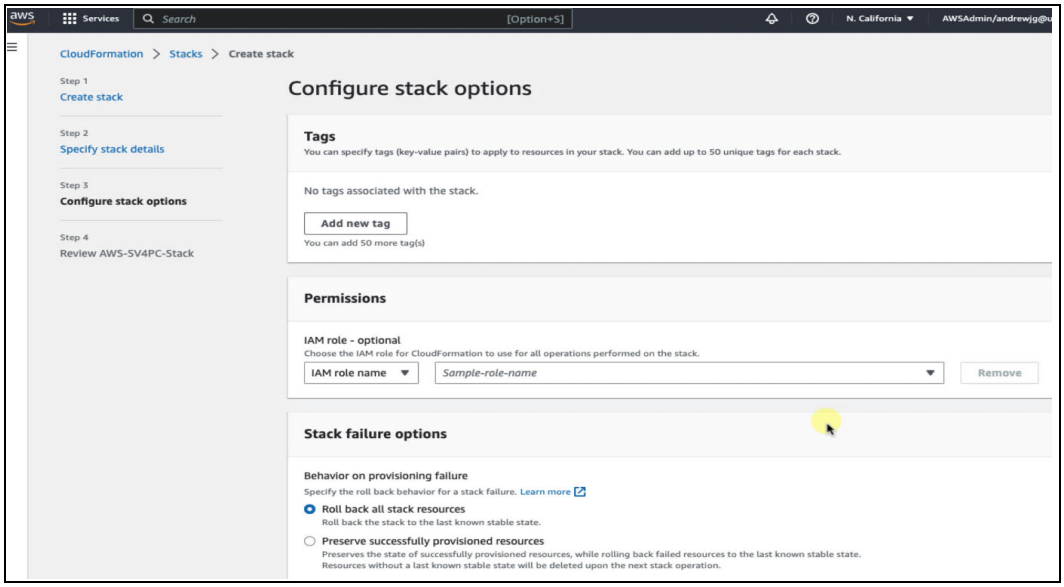

Figure 5-11 IBM Spectrum Virtualize for Public Cloud: Review CloudFormation Stack Optional Tags and Rollback

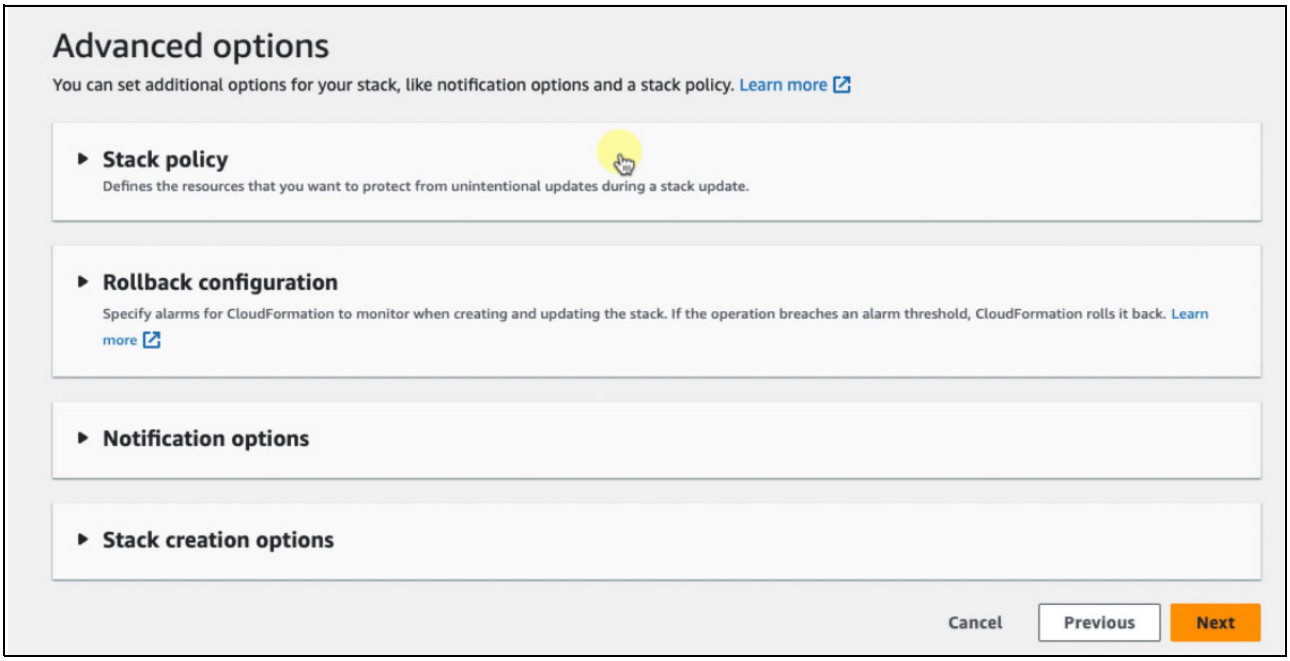

Figure 5-12 IBM Spectrum Virtualize for Public Cloud: CloudFormation Stack Optional Advanced options

12. After you review your entire stack selections, click **Create Stack**, as shown in Figure 5-13 and Figure 5-14.

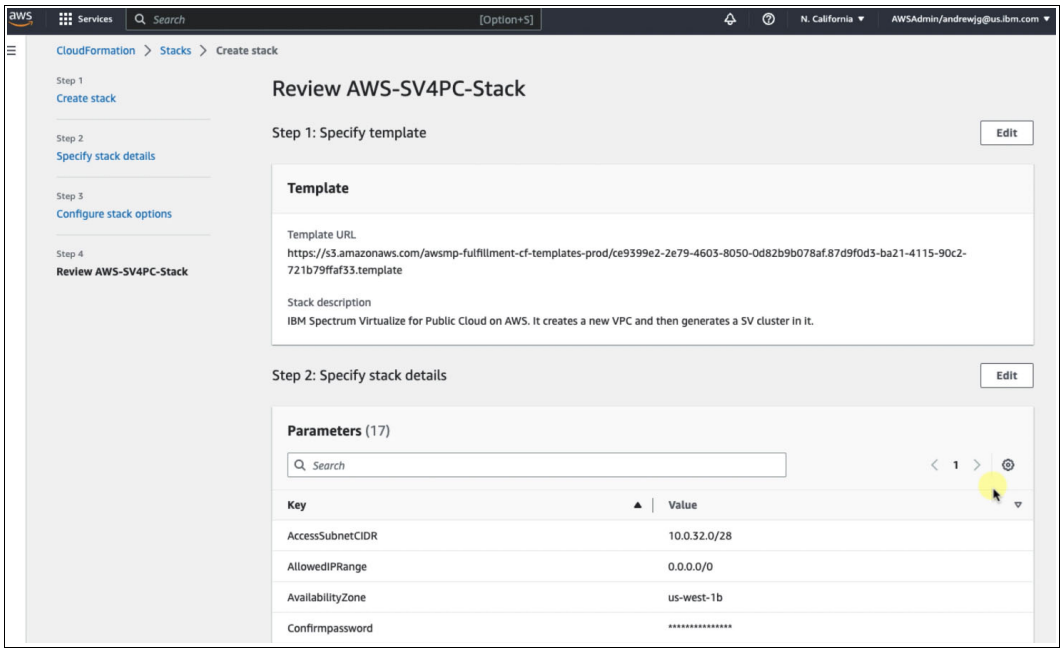

Figure 5-13 IBM Spectrum Virtualize for Public Cloud: Stack Review selections and finalization

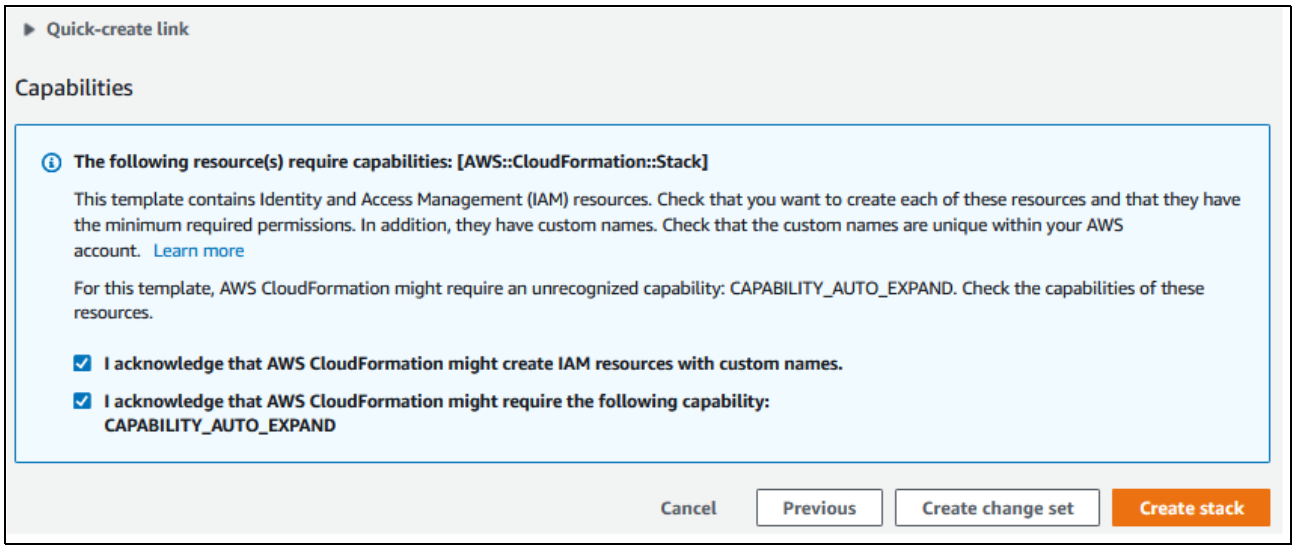

Figure 5-14 IBM Spectrum Virtualize for Public Cloud: CloudFormation Stack Final acknowledgment checkboxes

13. The stack creation process takes about 20 minutes for new VPCs and 15 minutes for existing VPCs. Progress can be monitored by going to the AWS console and selecting **CloudFormation** → **Stacks**, and then clicking the **Events** tab. After the stack and associated WorkloadStack reaches `CREATE_COMPLETE`, the environment is ready for interaction, as shown in Figure 5-15. Afterward, you can use Secure Shell (SSH) to access the Quorum Node for further configuration or enable GUI access.

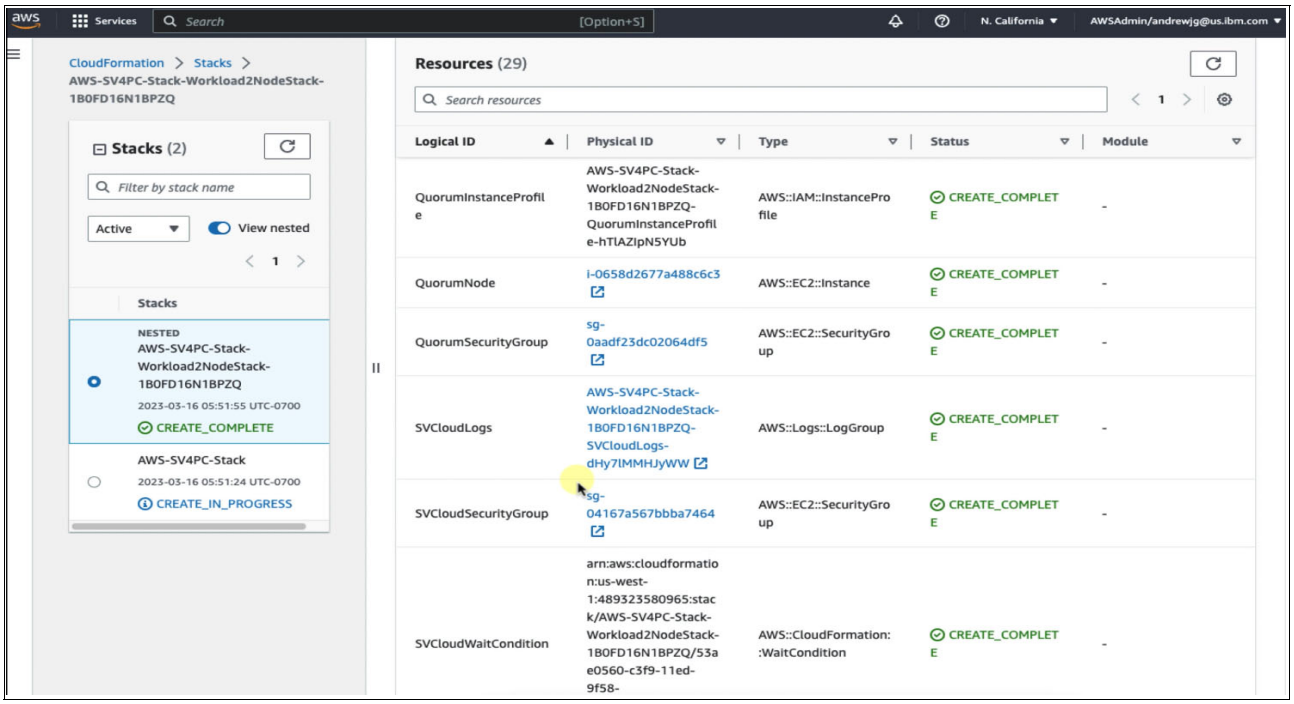

Figure 5-15 IBM Spectrum Virtualize for Public Cloud: CloudFormation Stack Creation complete

14. In this same view, you can view important IP address information by clicking the **Outputs** tab, as shown in Figure 5-16.

Names	IP address	Descriptions
IBMSVClusterIP	172.16.1.61	IBM Spectrum Virtualize Cloud Cluster IP
IBMSVNode1Port1NodeIP	172.16.1.104	IBM Spectrum Virtualize Node 1 Port 1
IBMSVNode1Port2NodeIP	172.16.1.134	IBM Spectrum Virtualize Node1 Port 2
IBMSVNode1PortIP1	172.16.1.91	IBM Spectrum Virtualize Node 1 Port IP 1
IBMSVNode1PortIP2	172.16.1.154	IBM Spectrum Virtualize Node 1 Port IP 2
IBMSVNode1ServiceIP	172.16.1.181	IBM Spectrum Virtualize Node1 Service IP
IBMSVNode2Port1NodeIP	172.16.1.241	IBM Spectrum Virtualize Node 2 Port 1
IBMSVNode2Port2NodeIP	172.16.1.40	IBM Spectrum Virtualize Node 2 Port 2
IBMSVNode2PortIP1	172.16.1.36	IBM Spectrum Virtualize Node 2 Port IP 1
IBMSVNode2PortIP2	172.16.1.236	IBM Spectrum Virtualize Node 2 Port IP 2
IBMSVNode2ServiceIP	172.16.1.193	IBM Spectrum Virtualize Node 2 Service
IBMSVNode3Port1NodeIP	172.16.1.20	IBM Spectrum Virtualize Node 3 Port 1
IBMSVNode3Port2NodeIP	172.16.1.73	IBM Spectrum Virtualize Node 3 Port 2
IBMSVNode3PortIP1	172.16.1.198	IBM Spectrum Virtualize Node 3 Port IP 1
IBMSVNode3PortIP2	172.16.1.77	IBM Spectrum Virtualize Node 3 Port IP 2
IBMSVNode3ServiceIP	172.16.1.211	IBM Spectrum Virtualize Node 3 Service
IBMSVNode4Port1NodeIP	172.16.1.59	IBM Spectrum Virtualize Node 4 Port 1
IBMSVNode4Port2NodeIP	172.16.1.173	IBM Spectrum Virtualize Node 4 Port 2
IBMSVNode4PortIP1	172.16.1.107	IBM Spectrum Virtualize Node 4 Port IP 1
IBMSVNode4PortIP2	172.16.1.94	IBM Spectrum Virtualize Node 4 Port IP 2
IBMSVNode4ServiceIP	172.16.1.119	IBM Spectrum Virtualize Node 4 Service
IBMSVQuorumClientEC2IP	172.16.2.42	IBM Spectrum Virtualize Quorum Client EC2 Private IP
IBMSVVersion	8.3.1.1	IBM SV Cloud version
StackUpdateTemplate	https://svcloud-beta.s3-us-west-2.amazonaws.com/8311/Beta/sv-cloud-node-	Template to add or remove IBM SV Cluster I/O group

Figure 5-16 Output of CloudFormation auto-provisioning after AWS finishes stack creation

5.2 Logging in to IBM Spectrum Virtualize for Public Cloud on Amazon Web Services

After the IBM Spectrum Virtualize for Public Cloud CloudFormation is created, you can log in to IBM Spectrum Virtualize for Public Cloud to provision the software array for volumes, hosts, and replication. Depending on your earlier IP address choices, you take a few extra steps by using the Azure Bastion service. Because this service is the only (external and optional) exposed address, it features the following functions:

- SSH jump host
- GUI proxy
- Cloud Call Home gateway
- SMTP gateway (optional)

- Remote Support Proxy (RSP) server (optional)
- Storage Insights DataCollector host (optional)

By using SSH after the IBM Spectrum Virtualize for Public Cloud deployment, you can enable the IBM Spectrum Virtualize for Public Cloud GUI through HTTPS, and directly by accessing the command-line interface (CLI) on the newly deployed system.

5.2.1 Using SSH to access the Azure Bastion host

Use the AWS console to access the list of EC2 instances and look for an instance that starts with your stack name and ends in `IBM-SV-QuorumNode`. Above that instance are the four IBM Spectrum Virtualize nodes. Select the QuorumNode instance or Azure Bastion host and look for the IPv4 Public IP in the **Description** tab, as shown in Figure 5-17.

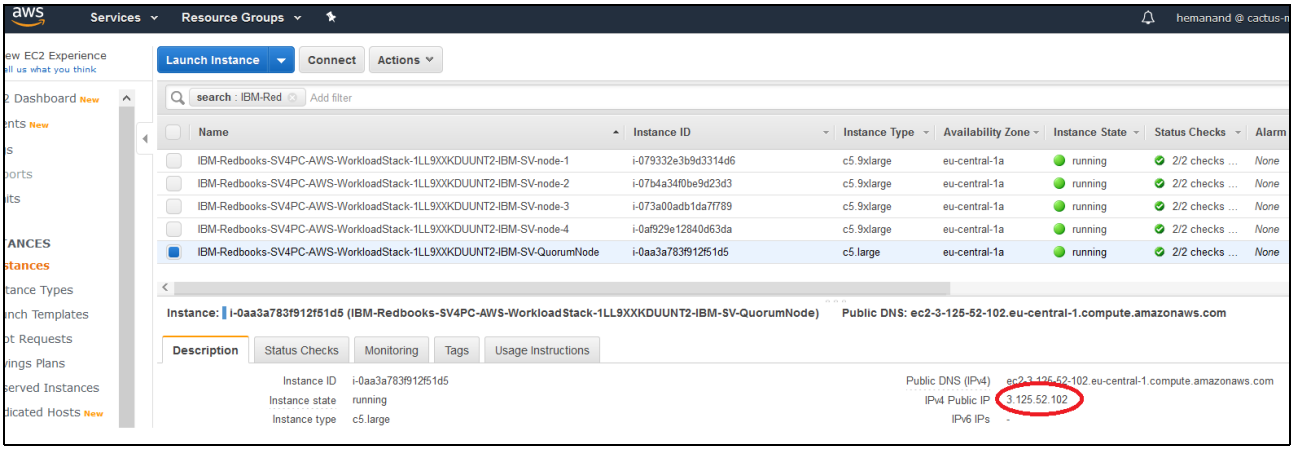

Figure 5-17 Public IP address of QuorumNode (Azure Bastion host)

One of the fastest ways to connect to any node is by using the built-in AWS connection method, which is named Session Manager. Click **Connect**, as shown in Figure 5-18.

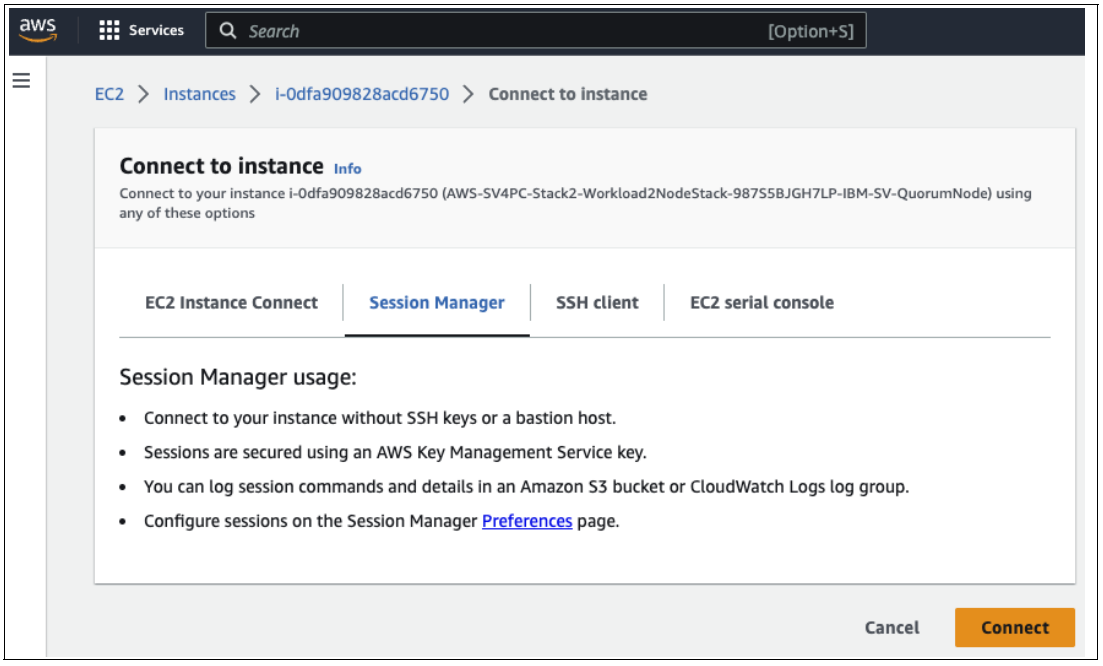

Figure 5-18 AWS Session Manager tab: connect to a node directly without using SSH keys or a client

By using your AWS session on your web browser, log in to the Quorum node, issue a (one-time) command to enable the Management GUI, and then connect by using that method, as shown in Figure 5-19.

Figure 5-19 AWS Session Manager showing sudo and enabling GUI without a KeyPair

You also can toggle this setting, which further reduces attack vectors. As an alternative to using an AWS Browser session, note the Node IP address and use a local SSH client and KeyPair. Make sure to change permissions on your KeyPair, and then run **ssh** by using the `centos` user and specifying the KeyPair and IP address to access the Azure Bastion host, as shown in Figure 5-20.

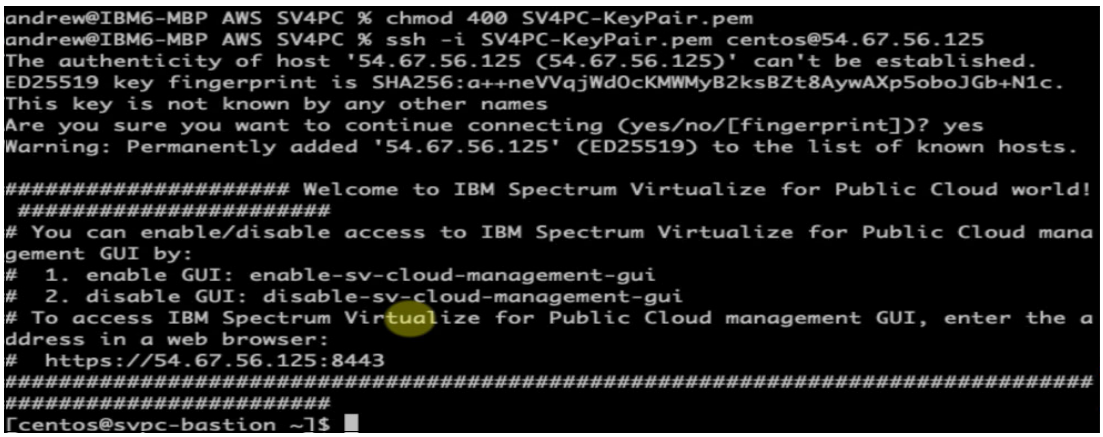

Figure 5-20 Running SSH to access the Azure Bastion host

5.2.2 Configuring the Azure Bastion host

To configure the Azure Bastion host, complete the tasks that are described in the following sections.

Enabling IBM Spectrum Virtualize for Public Cloud Management GUI access

Run **ssh** to access the Azure Bastion host by using the `ssh-key` that you specified during the installation, as shown in Example 5-1 and Figure 5-20.

Example 5-1 SSH connection to the Azure Bastion host to enable GUI access

```
[centos@svpc-bastion~]$ enable-sv-cloud-management-gui
```

Note: Port forwarding of port 8443, which is needed for GUI access, is disabled by default. Enable it for added security.

Configuring the Remote Support Proxy server

An *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 RSP server and configure the proxy in IBM Spectrum Virtualize for Public Cloud to enable remote support connections into the cluster.

For the purposes of this publication, assume that a separate virtual server is created in the environment that can access the public network and the private network, including routes to the subnet in which IBM Spectrum Virtualize for Public Cloud is running. Also, for this example, assume that the virtual server that is deployed is Red Hat Linux 7.x. (At the time of writing, it was the default OS).

Complete the following steps:

1. Download the RSP software from your product support page. At the time of this writing, this code is under the Others category, as shown in Figure 5-21.

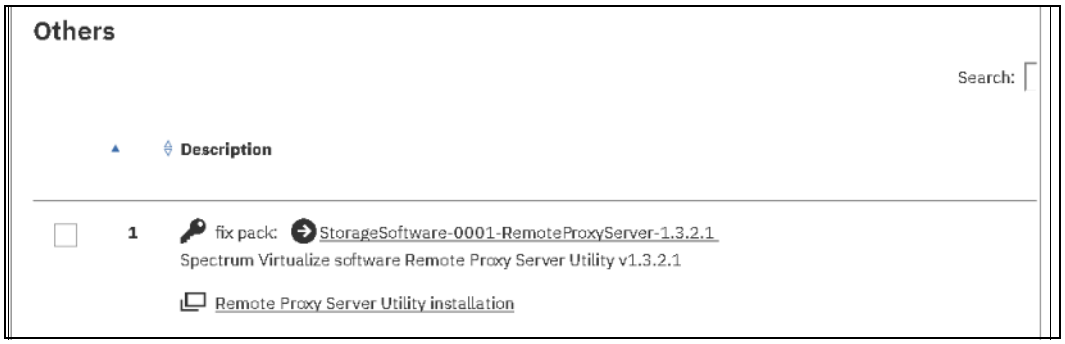

Figure 5-21 Downloading code from the product support page

2. After the code is downloaded to the administrator's notebook, you must upload the file to the server in which the proxy will be installed. To do so, run 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 prerequisite packages are installed, you can proceed with the installation, as shown in Example 5-2.

Example 5-2 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 packages `redhat-lsb` and `bzip2`. On SUSE systems, install the package `insserv`.

3. When the installer is started, you see the International License Agreement for Non-Warranted Programs. To complete the installation, enter `1` to accept the license agreement and complete the installation.
4. When the installation completes, you must configure the proxy server to listen for connections. This configuration is done by editing the `supportcenter/proxy.conf` configuration file, which is in the `/etc` directory. The minimum modification that is required is to edit the fields `ListenInterface` and `ListenPort`. By default, the file has "?" as the value for both.

5. To complete the configuration, specify `ListenInterface` with the interface name in Linux that can access the IBM Spectrum Virtualize clusters. You can discover this name by running the **`ifconfig`** command and identifying the interface that accesses the AWS Cloud private network. Also, set `ListenPort` to the TCP port number to listen on for remote support requests. A sample configuration file is shown in Example 5-3.

Tip: Consider the internal address of the Azure Bastion host (in Example 5-3, it is `10.0.32.86`). This address is available from the same AWS console view from where we retrieved the public IP, but it is useful to find it from `ifconfig` on the server. The internal IP is used for several configuration items on the IBM Spectrum Virtualize system. Also, make a note the port that is specified for `ListenPort` of the remote proxy because it is needed later in EasySetup for Support Proxy.

Example 5-3 Sample proxy configuration

```
[centos@svpc-bastion ~]$ ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
        inet 10.0.32.86 netmask 255.255.224.0  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

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 SAN Volume Controller 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
```

6. When the service is configured, you must start the service so that the server can start listening for requests. Optionally, you can also configure the service to start on system start. To start the service, run the **service** or **systemctl** command. To make the service start on system start, run the **chkconfig** command. Both of these processes are shown in Example 5-4.

Example 5-4 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 starts, you are ready to configure IBM Spectrum Virtualize for Public Cloud to use the proxy to start remote support requests.

5.2.3 Logging in to the IBM Spectrum Virtualize for Public Cloud cluster and completing the installation

After the deployment of the IBM Spectrum Virtualize for Public Cloud stack is complete, you can log in to the IBM Spectrum Virtualize for Public Cloud on AWS cluster through the Web GUI (Figure 5-22) by using the Azure Bastion public IP address as a proxy after the **`enable-sv-cloud-management-gui`** command is run on the Azure Bastion host. Complete the following steps:

1. With the proxy enabled, open a browser to the Azure Bastion public IP and append the port ID (:8443) to access the IBM Spectrum Virtualize WebGUI. In our example, it was `https://3.125.52.102:8443`.

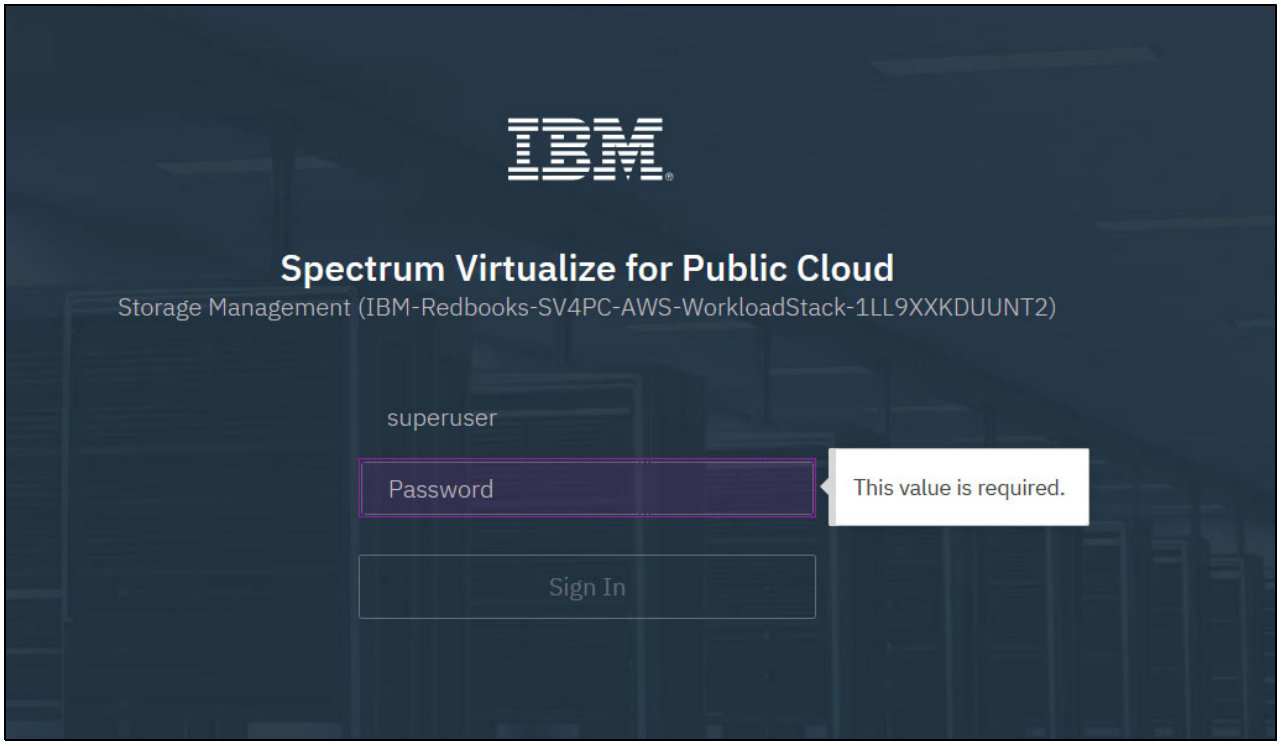

Figure 5-22 Logging in to the IBM Spectrum Virtualize for Public Cloud Management Web GUI

2. You are redirected to the Welcome window. Click **Next** (see Figure 5-23).

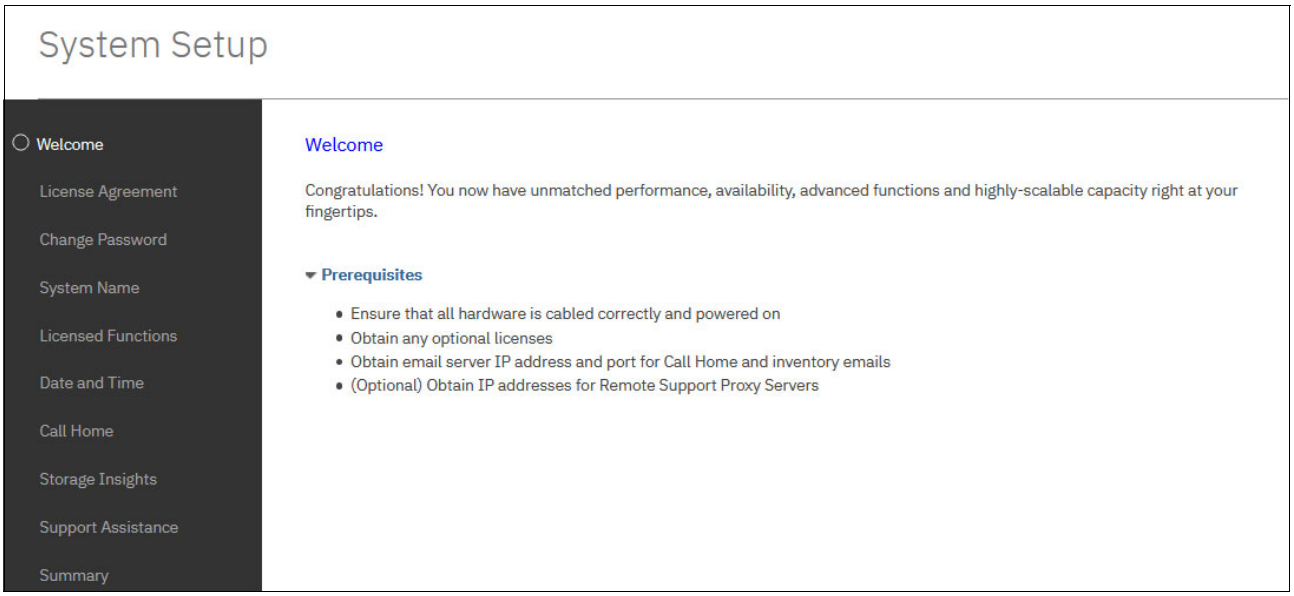

Figure 5-23 EasySetup: Welcome window

3. You are redirected to the Change Password window, as shown in Figure 5-24. Change your password, and then, click **Apply** and then, **Next** to open the next window.

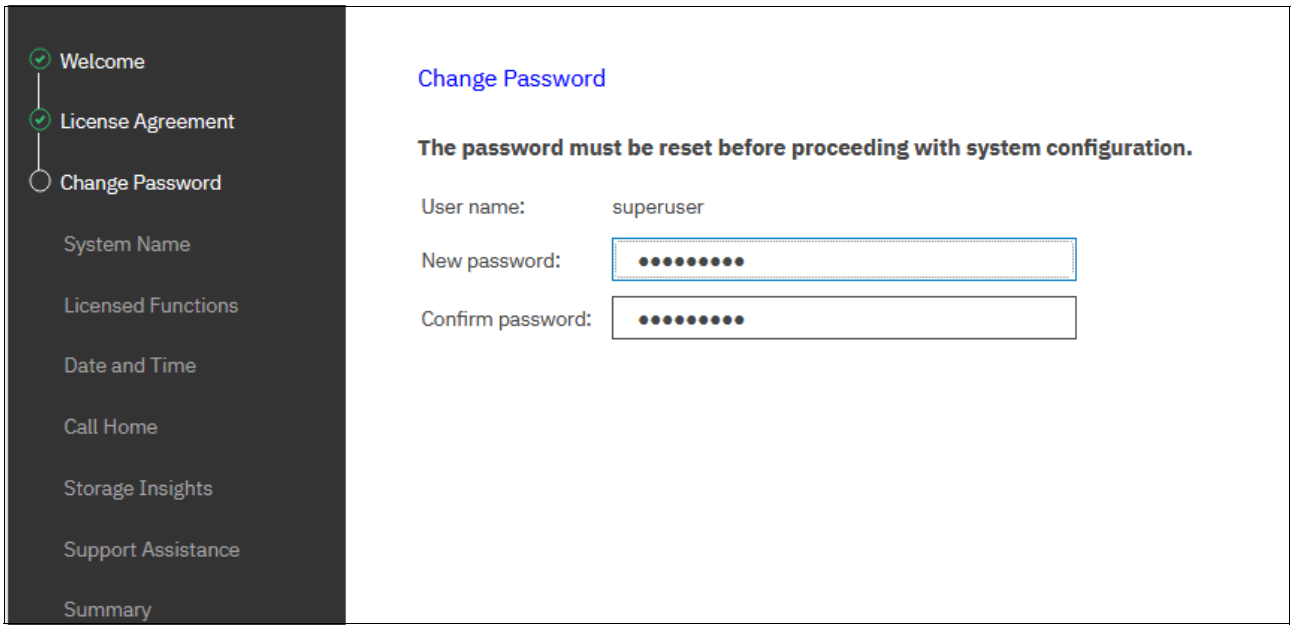

Figure 5-24 Easy Setup: Change Password window

4. You can change your IBM Spectrum Virtualize for Public Cloud (system) cluster name, which defaults to the stack ID name and WorkloadStack (stack unique identifier). As a best practice, trim the unique identifier at the end, as shown in Figure 5-25. Click **Apply** and **Next**.

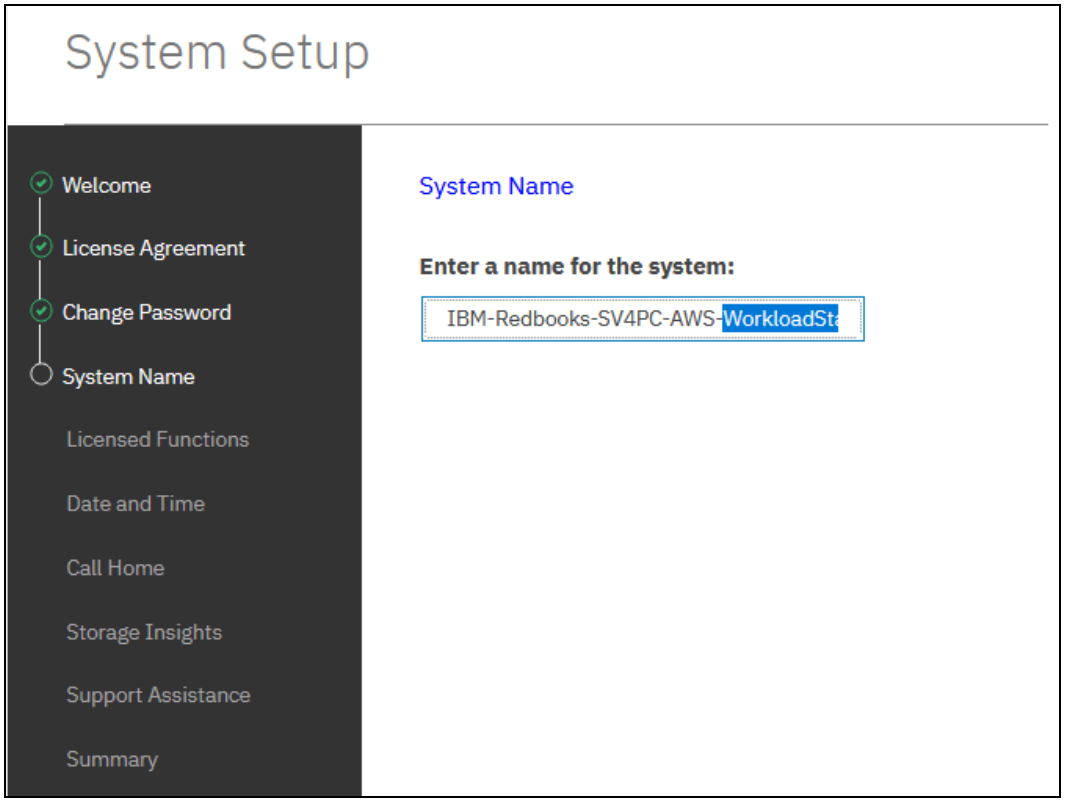

Figure 5-25 EasySetup: Trimming the system name

5. Enter your capacity license in accordance with your IBM agreement, as shown in Figure 5-26 on page 103. Click **Apply** and **Next** to open the next window.

Note: An IBM Spectrum Virtualize for Public Cloud license uses simple TiB values instead of Storage Capacity Units. This feature keeps the licensing model simple and still realizes economic benefits through thin provisioning and IBM Easy Tier. It also allows for overallocation of the Amazon EBS volumes that are purchased and for the use of fewer expensive high-performance Amazon EBS volumes and cheaper low-performance volumes.

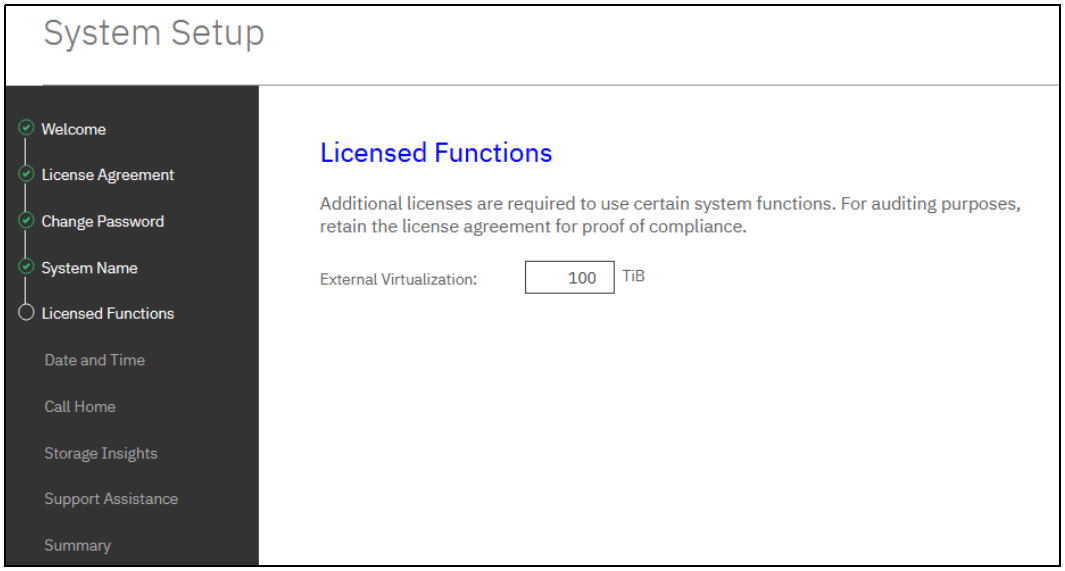

Figure 5-26 EasySetup: Licensed Functions

6. You do not need to set the date and time because that function is controlled by AWS. IBM Spectrum Virtualize for Public Cloud is configured by the existing AWS time server by using underlying operating system methods; however, it is beneficial to update the time zone for your business requirements.

Note: Changing the time server or setting a static time *is not recommended* and might cause difficulties.

For more information about the AWS time server, see Setting the Time for Your Linux Instance.

Ensure that the time zone is set. *For ease of troubleshooting across multiple time zones, it is a best practice to use GMT or UTC+0*, as shown in Figure 5-27.

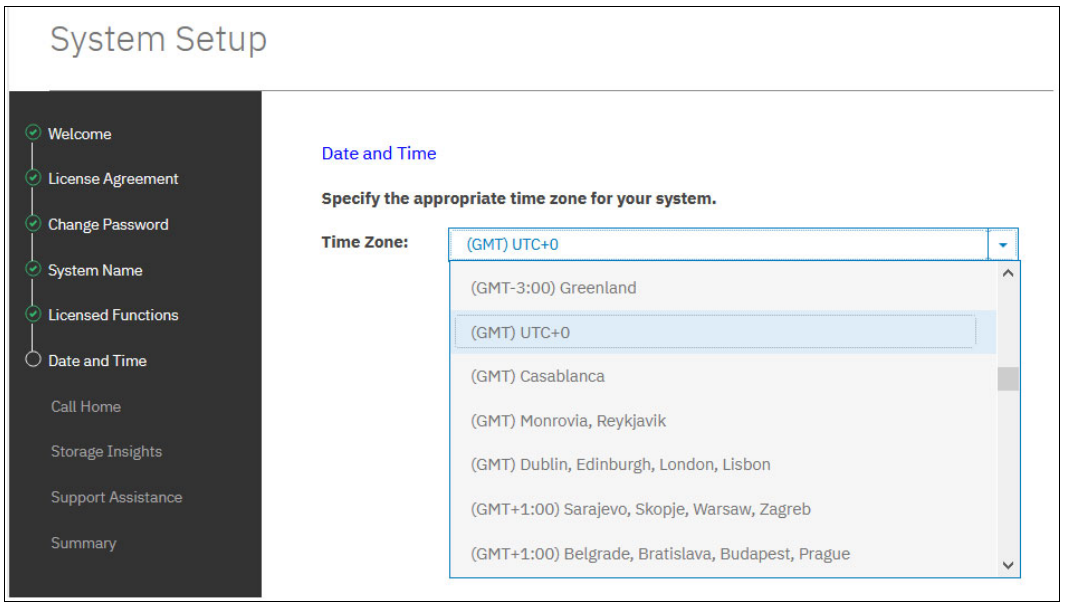

Figure 5-27 EasySetup: Time Zone

7. IBM Spectrum Virtualize for Public Cloud on AWS is preconfigured with an IBM Cloud Call Home feature that uses the Azure Bastion host as a gateway. When the EasySetup process enters the Call Home configuration, Cloud Call Home verifies the connection to the support center, as shown in Figure 5-28.

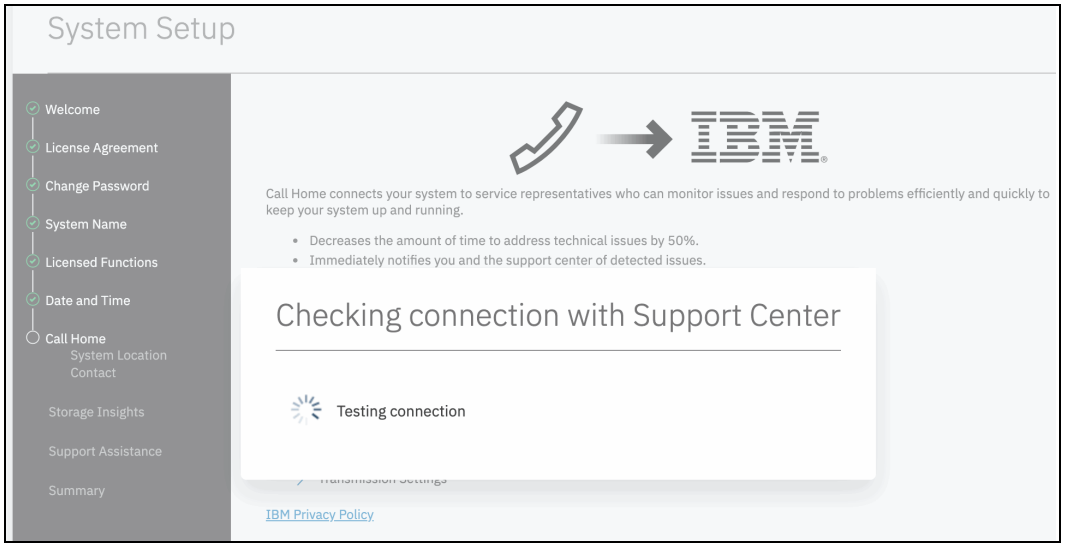

Figure 5-28 EasySetup: Cloud Call Home verification

This verification should succeed, as shown in Figure 5-29, which is the System Location window.

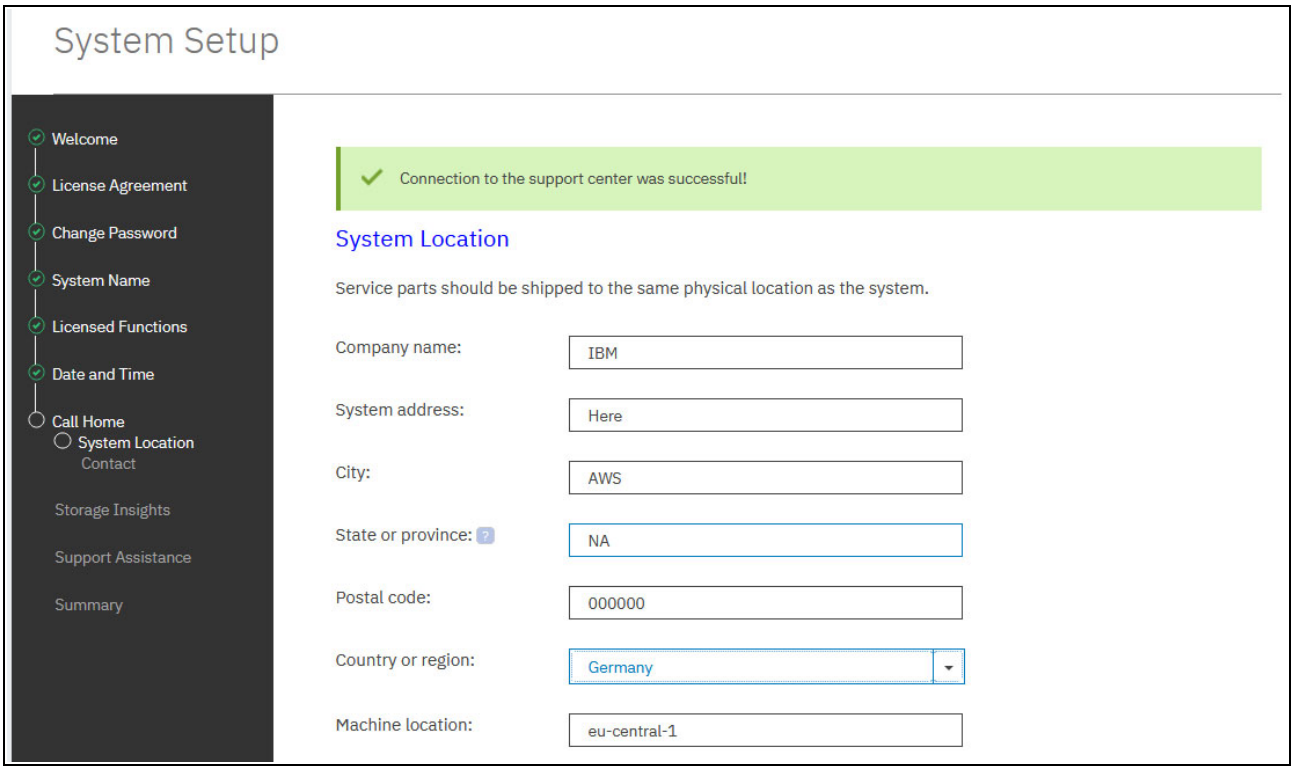

Figure 5-29 EasySetup: Successful Cloud Call Home and System Location information

8. Finish the Call Home configuration by entering the contact information, as shown in Figure 5-30 on page 105.

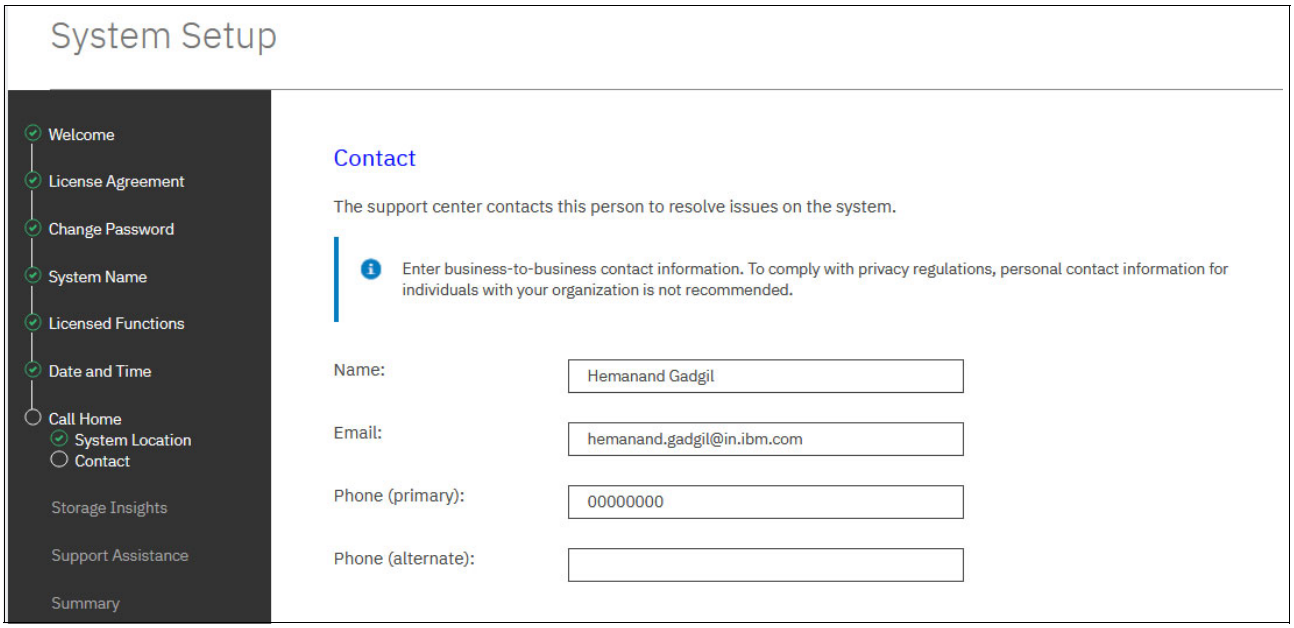

Figure 5-30 EasySetup: Contact information

9. The IBM Storage Insights configuration must be completed or turned off (this process is *not* done during EasySetup). The process requires registering for a no-charge account. Figure 5-31 shows the IBM Storage Insights configuration window. Skip this step for now.

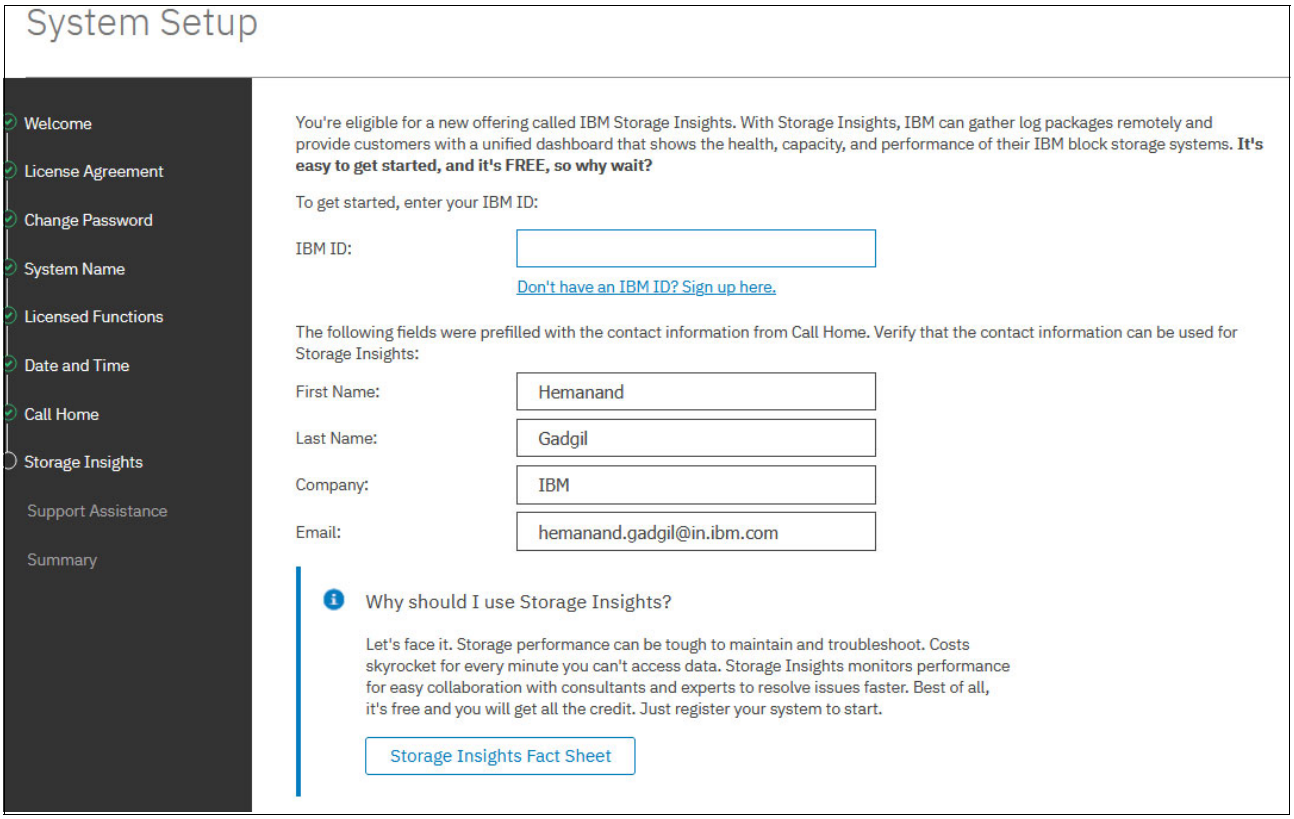

Figure 5-31 EasySetup: IBM Storage Insights

10. Configure your direct or (optional RSP), depending on your configuration, as shown in Figure 5-32.

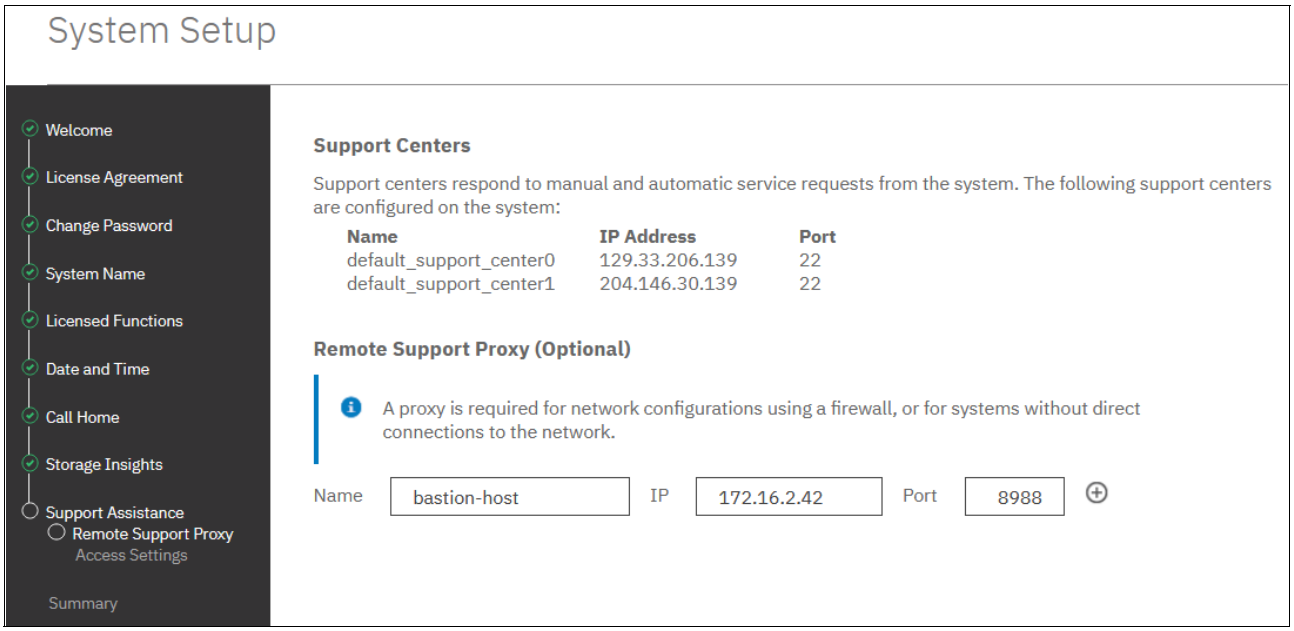

Figure 5-32 EasySetup: Direct or by using an optional Remote Support Proxy

Note: This step assumes that you deployed an RSP. Again, the Azure Bastion host is a logical choice. Note the internal IP address of the Azure Bastion host and the `ListenPort` that was specified in 5.2.2, “Configuring the Azure Bastion host” on page 96.

Figure 5-33 on page 107 shows a summary of your configuration. Your cluster setup is complete.

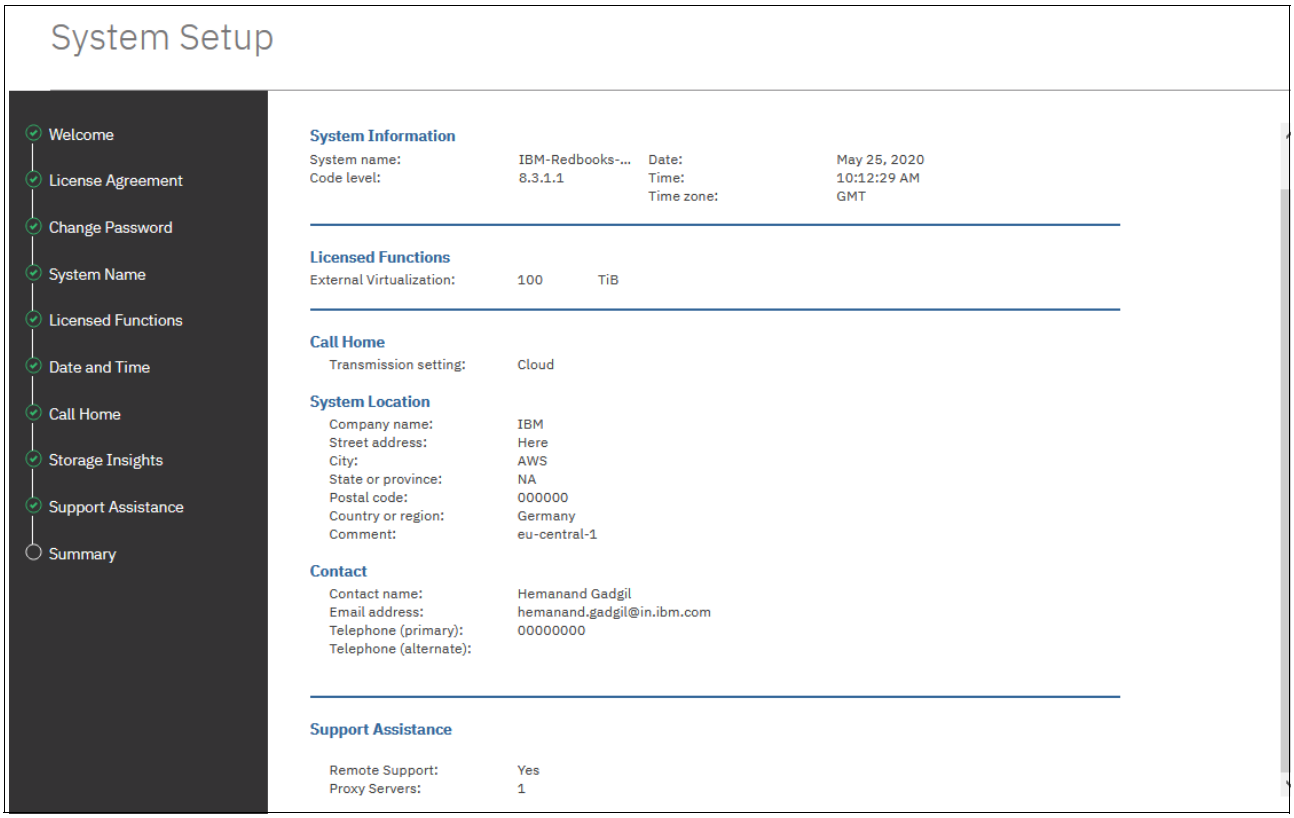

Figure 5-33 EasySetup: Summary

Note: Call Home is set up with Cloud Call Home. However, email notification is useful for event notification and can be set up after the EasySetup process is complete. The Azure Bastion host runs an SMTP service and can be used as the email gateway.

5.3 Configuring the Cloud Quorum

IP quorum applications are used in Ethernet networks to resolve failure scenarios when half the nodes on the system become unavailable. These applications determine which nodes can continue processing host operations and avoids a split-brain scenario in which both halves attempt to service independently I/O, which causes corruption.

As part of the installation of IBM Spectrum Virtualize for Public Cloud on AWS, an Azure Bastion host is provisioned and the IP quorum application is installed and configured on this instance. This Azure Bastion host operates as the IP quorum and the network gateway for the configuration.

Note: An IP quorum is configured during the installation. You can configure an extra IP quorum if you want to enhance fault tolerance by putting the active quorum (in a different Availability Zone for installations) into new, extra VPCs.

Strict requirements for the IP network that uses IP quorum applications must be met. All IP quorum applications must be reconfigured and redeployed to hosts when specific aspects of the system configuration change. These aspects include adding or removing a node from the system or when node service IP addresses are changed.

Other examples include changing the system certificate or experiencing an Ethernet connectivity issue. Such a 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 Figure 5-34.

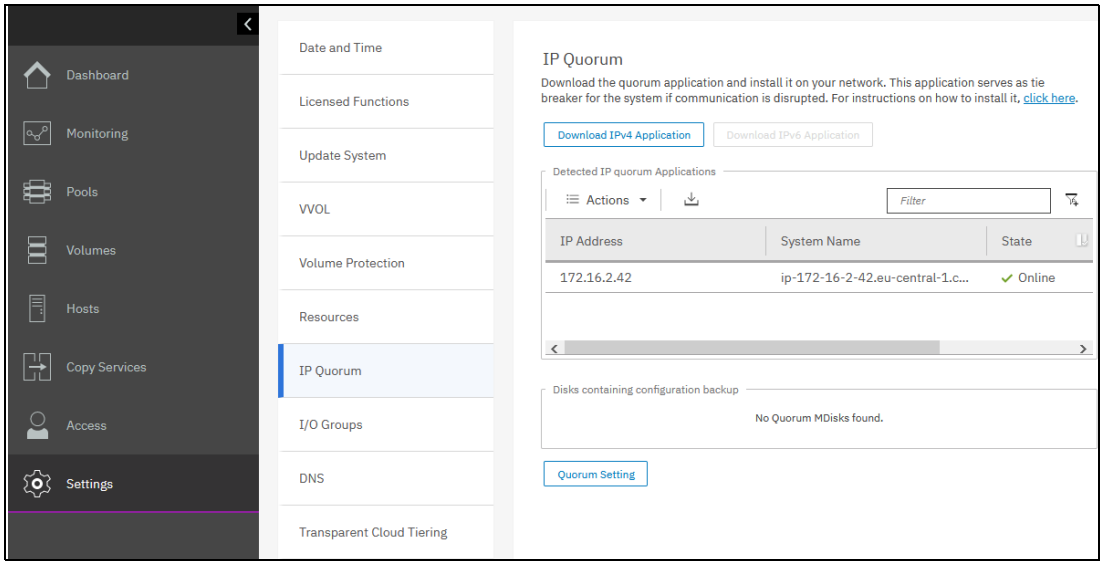

Figure 5-34 IP quorum example from the GUI

Even with IP quorum applications on an EC2 instance, quorum disks are required on each node in the system to contain backups of the configuration and recovery information. On EC2 instances where IBM Spectrum Virtualize connectivity with its nontraditional back-end storage connectivity, the quorum disks cannot be on external storage or internal disk as in SAN Controller Volume or IBM FlashSystem systems. Therefore, they are automatically allocated on the EC2 instance boot device for each IBM Spectrum Virtualize node.

The IBM Spectrum Virtualize command **`lsquorum`** shows only the IP quorum.

The maximum number of IP quorum applications that can be deployed is five. Applications can be deployed on multiple hosts to provide redundancy.

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

- 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 the possible security implications of exposing the service IP addresses because this connectivity also can 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 MBps is guaranteed for node-to-quorum traffic.

For more information about the IP quorum configuration, see IBM Documentation.

Note: The current Cloud Formation Template (CFT) for new VPCs deploys the Azure Bastion host (which houses the initial IP quorum device) into the same Availability Zone as the IBM Spectrum Virtualize nodes. If deploying into an existing VPC, it is possible to place that Azure Bastion host on a subnet that is in a different Availability Zone from the IBM Spectrum Virtualize nodes.

However, if you are deploying in to a new VPC that is created as part of the IBM Spectrum Virtualize for Public Cloud deployment process, it is a best practice that you create a subnet in that VPC that belongs to a *different* Availability Zone. Then, start a new secure EC2 instance by using only a private interface in that new subnet with no direct access from the internet. Next, you deploy an IP quorum application on that server and restart the one on the Azure Bastion host so that the secure, redundant IP quorum is the active quorum device.

In summary, deploying a second IP quorum server with a new VPC includes the following steps:

1. Create a subnet within the VPC in a different Availability Zone than the IBM Spectrum Virtualize nodes and Azure Bastion host.
2. Start a new EC2 instance. You can use the Amazon Linux Amazon Machine Images (AMIs) 2018.03.0 image from the quick start because it has Java preinstalled. The default type of t2.micro is suitable but do *not* select Review and Launch.
3. Click **Next: Configure Instance Details** and select the correct VPC and subnet that you created in step 1. Leave **Public IP** disabled for added security and use an existing security group (same as the Azure Bastion host).
4. Click **Review and Launch** to review the configuration and then, click **Launch**.
5. Select the key pair that was used during the creation of the cluster because the key pair is needed to access the new EC2 instance.
6. After the instance is provisioned, run the **scp** command on the private key that is used to access the Azure Bastion host over *to* the Azure Bastion host.
7. Run **ssh** to access the Azure Bastion host and run **scp** to transfer the `ip_quorum.jar` file from the Azure Bastion host over to the new EC2 instance by using the private key:

   ```
   scp -i ~/.ssh/privkey.pem /usr/local/bin/ip_quorum.jar ec2-user@{new EC2 IP}:
   ```

8. Run **ssh** to access the new EC2 instance and test the ip_quorum service:

   ```
   java -jar ~/ip_quorum.jar
   ```

9. Set up the quorum as a service or install a cronjob to ensure that it is always running.
10. Exit the new EC2 instance and restart the `ip_quorum` service on the Azure Bastion host:

    ```
    systemctl restart ip-quorum
    ```

5.4 Expanding from a 2-node to a 4-node cluster in AWS

IBM Spectrum Virtualize for Public Cloud software in AWS supports 2-node and 4-node cluster configurations. You can expand a 2-node cluster to four nodes by adding nodes to a stack in AWS.

5.4.1 Prerequisites

Before you expand a 2-node cluster to a 4-node cluster, you must ensure that both the nodes that are added to the configuration and existing nodes in the cluster are updated to the latest version of the IBM Spectrum Virtualize for Public Cloud software. For information, see IBM Documentation.

To expand a 2-node cluster to a 4-node cluster in AWS, complete the following steps:

1. Log on to the AWS Management Console with the AWS default administrator profile or the installer profile.
2. Select **CloudFormation** → **Stacks**. Select the existing 2-node cluster configuration. It is displayed as a nested workload with the following name format:

   `[stack-name]-workstack-{resource id}`

   The stack-name is specified when the cluster is created with the AWS CloudFormation template (see Figure 5-35).

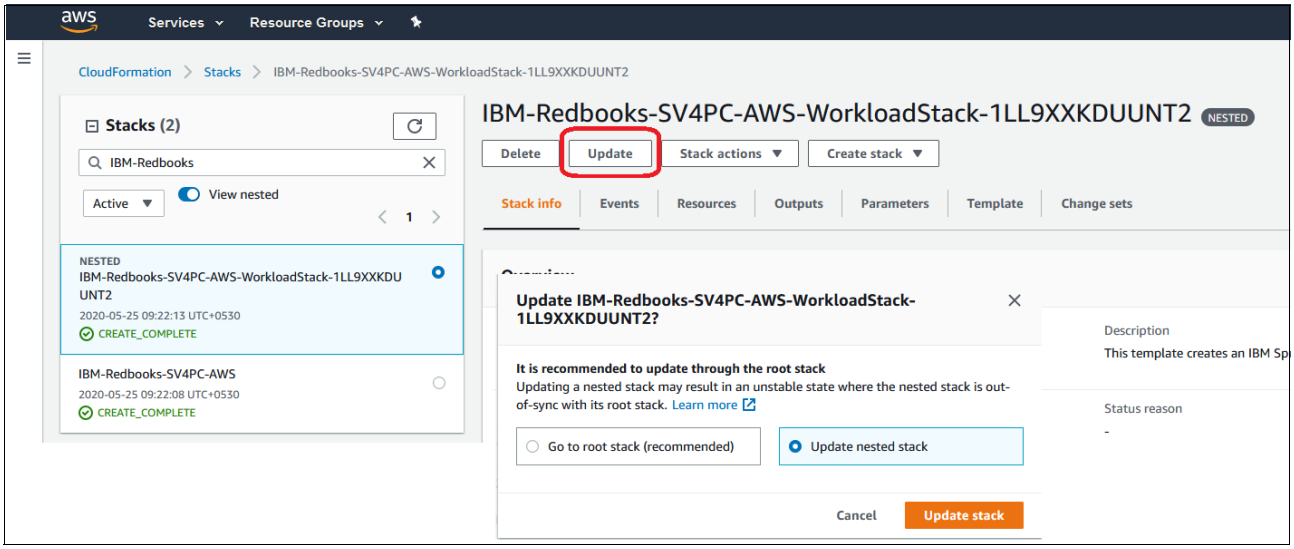

Figure 5-35 Expanding the 2-node cluster to 4-node cluster

3. Click **Update**.
4. Select **Updated nested stack** and click **Update stack**.
5. On the Update stack page, select the following options:
   a. In the Prerequisite-Prepare template section, select **Replace current template**.
   b. In the Specify template section, select **Amazon S3 URL**.
   c. In the Amazon S3 URL field, enter the URL that is displayed in the StackUpdateTemplate field.
   d. Click **Next**.

This information was included in the summary and email notification when the IBM Spectrum Virtualize for Public Cloud node instances were first installed in AWS. This information is also included on the Output tab when the node instances were first installed in AWS (see Figure 5-36).

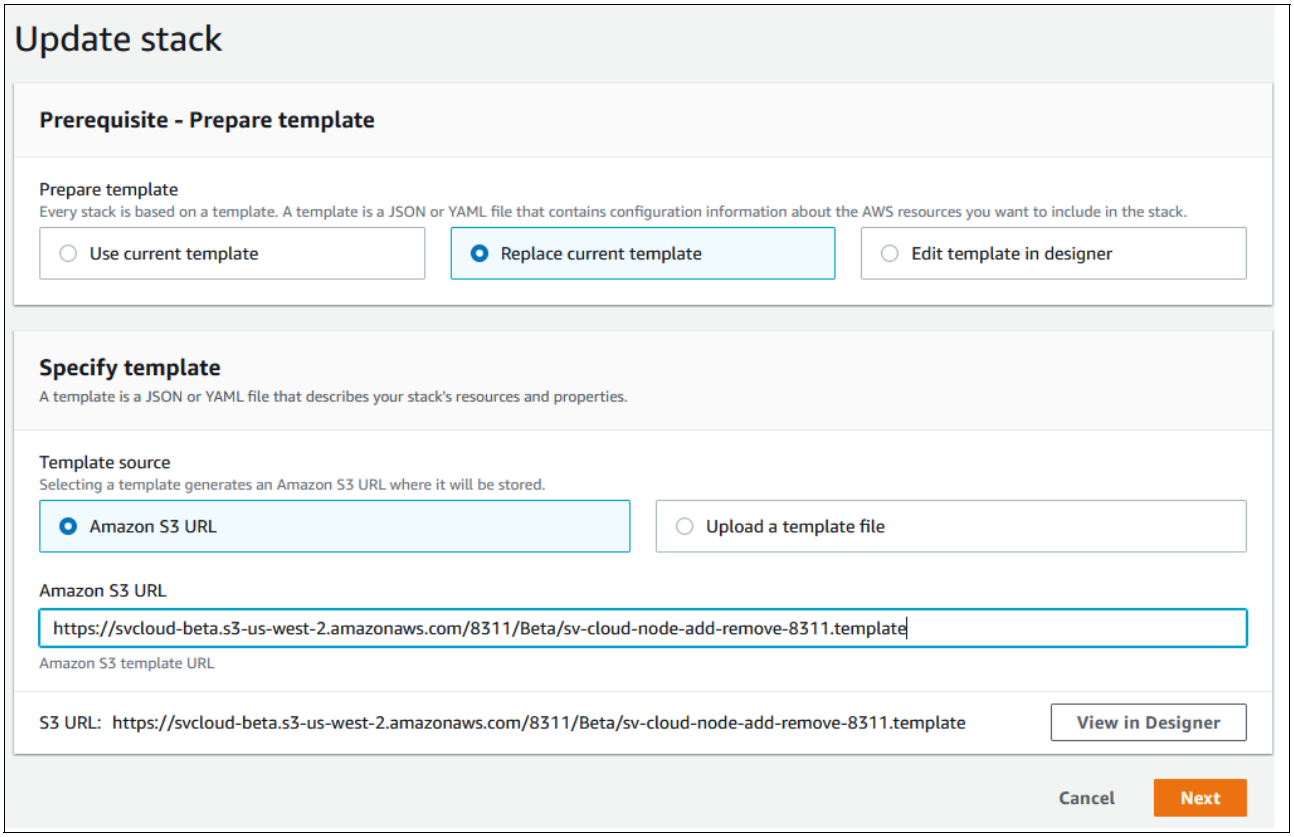

Figure 5-36 Update stack URL

6. Click **Next**.

7. On the Specify stack details page, keep the values that are configured for the existing configuration. Review the Amazon EC2 Configuration section and confirm the node instance type for the new I/O group is correct. Click **Next**, as shown in Figure 5-37.

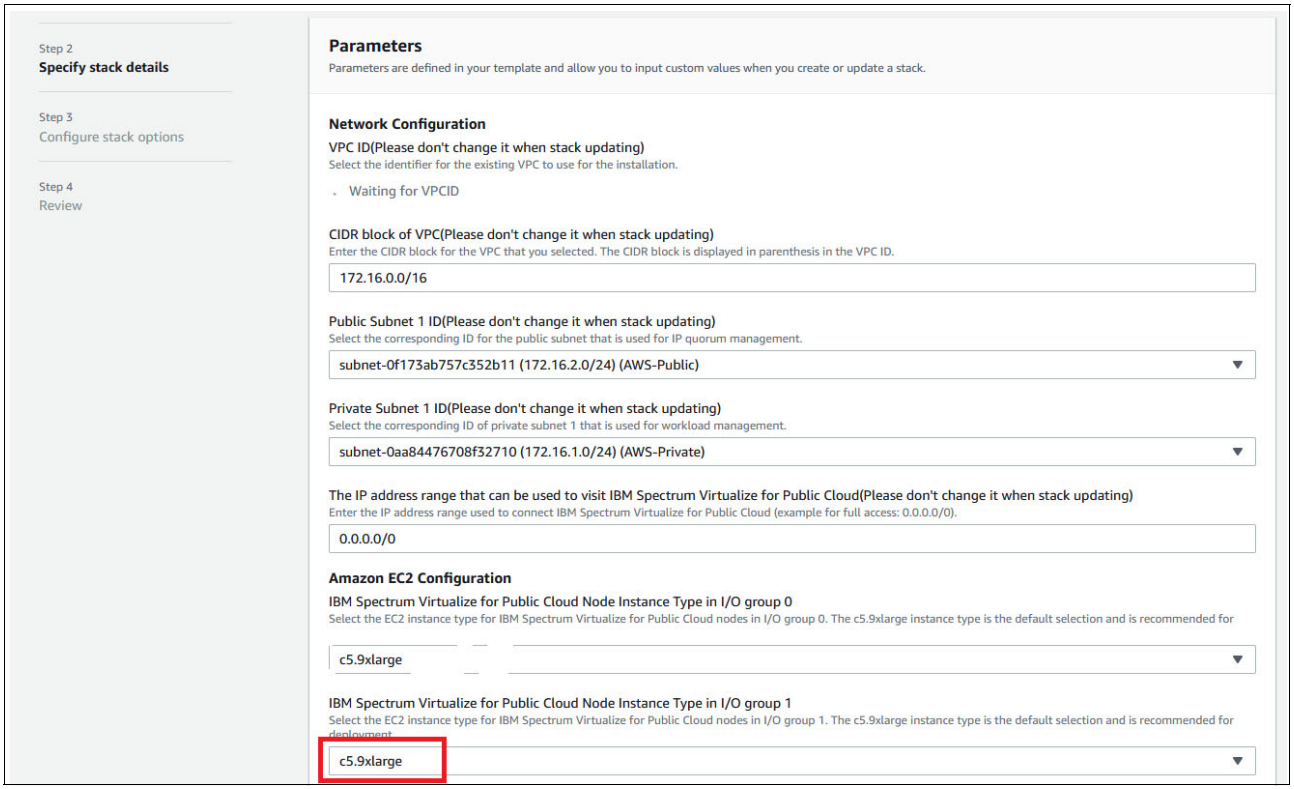

Figure 5-37 Selecting the I/O group configuration to expand from a 2 nodes to a 4 nodes

8. On the Configure stack options page, keep the values that are configured. Click **Next**.
9. On the Review page, review the options. Click **Next**.
10. On the Change set preview page, review the changed resources. Several resources are modified and two more EC2 instances are added for the nodes. After you verify these changes, ensure that **I acknowledge that AWS CloudFormation might create IAM resource** is selected.
11. Click **Update stack**. Verify that the status of the nested stack changes to `Update_In_Progress`.

12. After the stack is listed as `UPDATE_COMPLETE`, review the details that are listed for the updated nested stack on the **CloudFormation** → **Stacks** page. It includes the configuration of the existing cluster and the new nodes (see Figure 5-38).

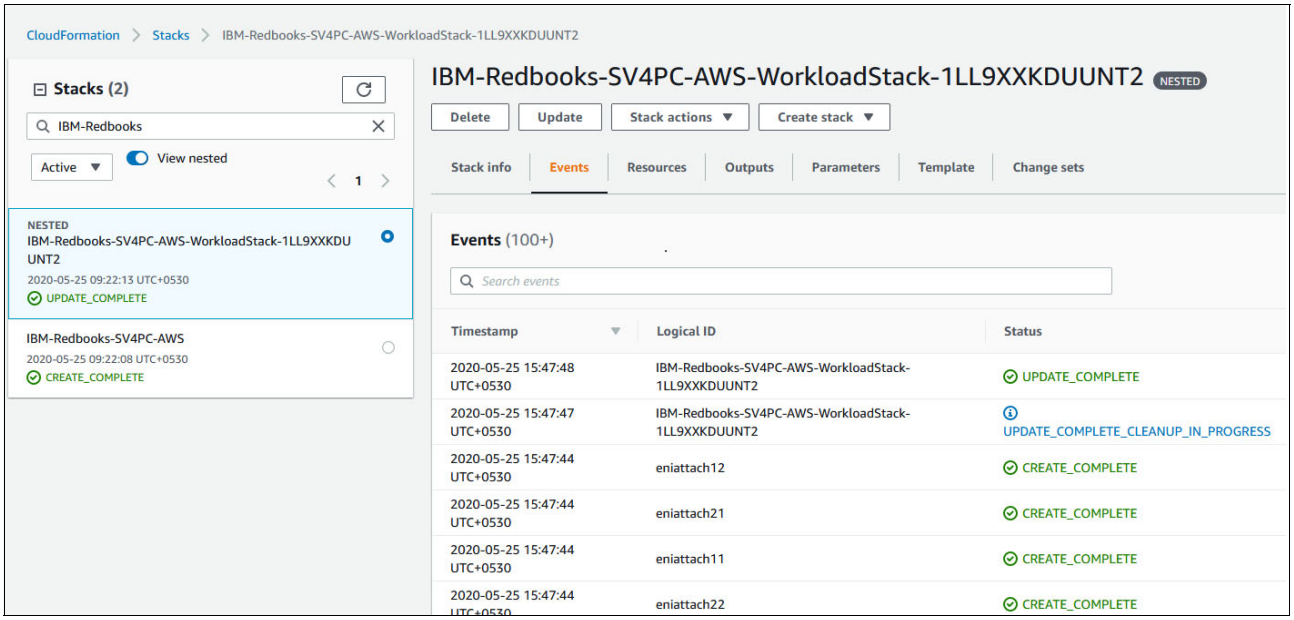

Figure 5-38 Completing the 2-node to 4-node update task

The output with all the IP addresses of newly deployed nodes is displayed in Output tab with node 3 and node 4 added, as shown in Figure 5-39.

Names	IP address	Descriptions
IBMSVNode1Port1NodeIP	172.16.1.104	IBM Spectrum Virtualize Node 1 Port 1 Node IP
IBMSVNode1Port2NodeIP	172.16.1.134	IBM Spectrum Virtualize Node1 Port 2 Node IP
IBMSVNode1PortIP1	172.16.1.91	IBM Spectrum Virtualize Node 1 Port IP 1
IBMSVNode1PortIP2	172.16.1.154	IBM Spectrum Virtualize Node 1 Port IP 2
IBMSVNode1ServiceIP	172.16.1.181	IBM Spectrum Virtualize Node1 Service IP
IBMSVNode2Port1NodeIP	172.16.1.241	IBM Spectrum Virtualize Node 2 Port 1 Node IP
IBMSVNode2Port2NodeIP	172.16.1.40	IBM Spectrum Virtualize Node 2 Port 2 Node IP
IBMSVNode2PortIP1	172.16.1.36	IBM Spectrum Virtualize Node 2 Port IP 1
IBMSVNode2PortIP2	172.16.1.236	IBM Spectrum Virtualize Node 2 Port IP 2
IBMSVNode2ServiceIP	172.16.1.193	IBM Spectrum Virtualize Node 2 Service IP
IBMSVNode3Port1NodeIP	**172.16.1.20**	**IBM Spectrum Virtualize Node 3 Port 1 Node IP**
IBMSVNode3Port2NodeIP	**172.16.1.73**	**IBM Spectrum Virtualize Node 3 Port 2 Node IP**
IBMSVNode3PortIP1	**172.16.1.198**	**IBM Spectrum Virtualize Node 3 Port IP 1**
IBMSVNode3PortIP2	**172.16.1.77**	**IBM Spectrum Virtualize Node 3 Port IP 2**
IBMSVNode3ServiceIP	**172.16.1.211**	**IBM Spectrum Virtualize Node 3 Service IP**
IBMSVNode4Port1NodeIP	**172.16.1.59**	**IBM Spectrum Virtualize Node 4 Port 1 Node IP**
IBMSVNode4Port2NodeIP	**172.16.1.173**	**IBM Spectrum Virtualize Node 4 Port 2 Node IP**
IBMSVNode4PortIP1	**172.16.1.107**	**IBM Spectrum Virtualize Node 4 Port IP 1**
IBMSVNode4PortIP2	**172.16.1.94**	**IBM Spectrum Virtualize Node 4 Port IP 2**
IBMSVNode4ServiceIP	**172.16.1.119**	**IBM Spectrum Virtualize Node 4 Service IP**

Figure 5-39 IP addresses of newly deployed nodes

5.5 Shrinking the IBM Spectrum Virtualize for Public Cloud node configuration from four nodes to two nodes in Amazon Web Services

IBM Spectrum Virtualize for Public Cloud software inside AWS supports 2-node and 4-node cluster configurations. You can shrink a four-node cluster to two nodes by removing nodes in the stack in AWS.

For more information about prerequisites and restrictions, see IBM Documentation.

Move the volumes to the remaining I/O group by using the procedure that is described in IBM Documentation.

After the prerequisites are completed, follow steps 1 - 6 as described in 5.4, “Expanding from a 2-node to a 4-node cluster in AWS” on page 110. Then, complete the following steps:

1. On the Specify stack details page, keep the values that are configured for the configuration. Review the Amazon EC2 Configuration section and ensure that the I/O group that is being removed is set to **None** (see Figure 5-40). For example, in this procedure, `iogrp0` is being removed and the IBM Spectrum Virtualize node instance type for I/O group 0 must be set to **None**. Click **Next**.

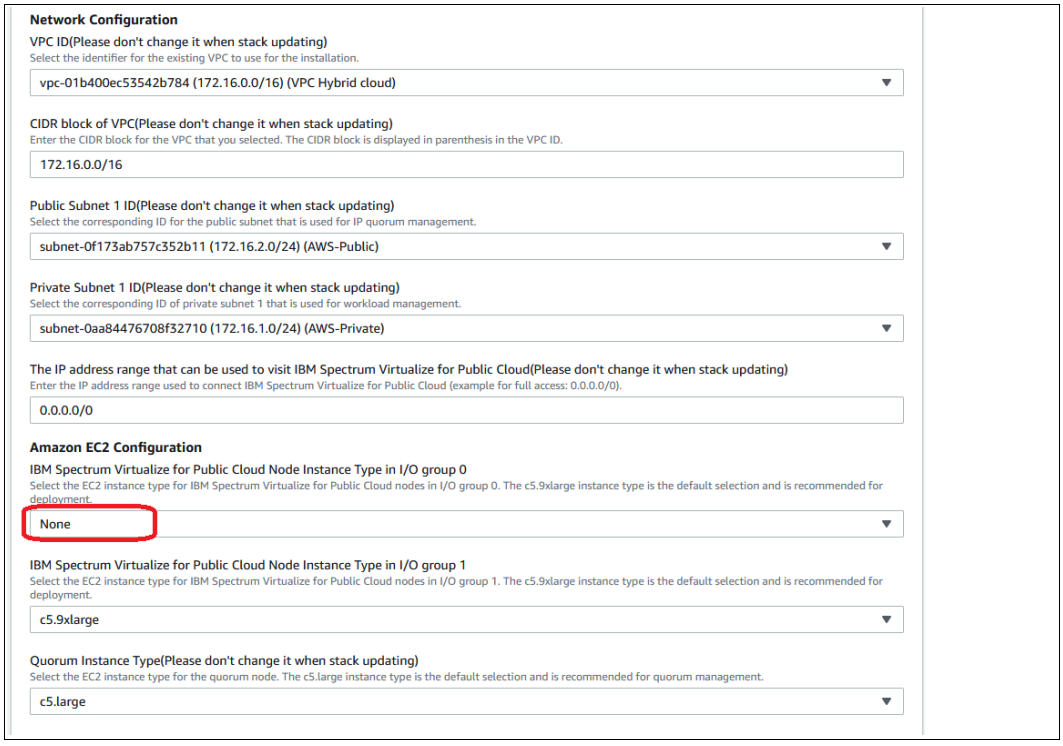

Figure 5-40 Selecting I/O group None to shrink the cluster

2. On the Configure stack options page, keep the configured values. Click **Next**.
3. On the Review page, review the options. Click **Next**.

4. On the Change set preview page, review the changed resources. Several resources are modified and two EC2 instances are removed for the two nodes that are being deleted from the cluster. After verifying these changes, ensure that **I acknowledge that AWS CloudFormation might create IAM resource** is selected.
5. Click **Update stack**. Verify that the status of the nested stack changes to `Update_In_Progress`.
6. After the stack is listed as `UPDATE_COMPLETE`, review the details that are listed for the updated nested stack on the **CloudFormation** → **Stacks** page (see Figure 5-41).

Names	IP address	Descriptions
IBMSVNode3Port1NodeIP	172.16.1.20	IBM Spectrum Virtualize Node 3 Port 1 Node IP
IBMSVNode3Port2NodeIP	172.16.1.73	IBM Spectrum Virtualize Node 3 Port 2 Node IP
IBMSVNode3PortIP1	172.16.1.198	IBM Spectrum Virtualize Node 3 Port IP 1
IBMSVNode3PortIP2	172.16.1.77	IBM Spectrum Virtualize Node 3 Port IP 2
IBMSVNode3ServiceIP	172.16.1.211	IBM Spectrum Virtualize Node 3 Service IP
IBMSVNode4Port1NodeIP	172.16.1.59	IBM Spectrum Virtualize Node 4 Port 1 Node IP
IBMSVNode4Port2NodeIP	172.16.1.173	IBM Spectrum Virtualize Node 4 Port 2 Node IP
IBMSVNode4PortIP1	172.16.1.107	IBM Spectrum Virtualize Node 4 Port IP 1
IBMSVNode4PortIP2	172.16.1.94	IBM Spectrum Virtualize Node 4 Port IP 2
IBMSVNode4ServiceIP	172.16.1.119	IBM Spectrum Virtualize Node 4 Service IP

Figure 5-41 Details that are listed for the updated nested stack

5.6 Configuring IBM Spectrum Virtualize for Public Cloud back-end storage and pools

IBM Spectrum Virtualize for Public Cloud on AWS uses the back-end storage that is provided by Amazon EBS as external managed disks (MDisks). As part of the initial default installation, two gp2 Amazon EBS volumes are allocated and put into a pool on the IBM Spectrum Virtualize cluster (see Figure 5-42).

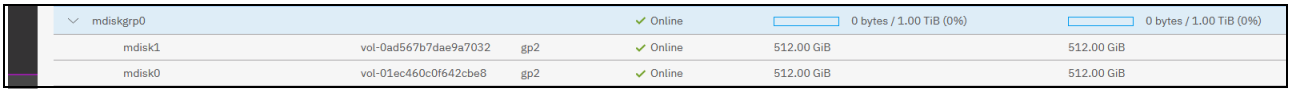

Figure 5-42 Default Amazon EBS gp2 volumes that are specified during CloudFormation template configuration

If more or different storage is wanted, complete the following steps:

1. To order back-end storage, log in to the AWS Console.
2. Click **Services** in the upper left of the browser window. Then, click **EC2**.
3. Under Resources, click **Volumes**. In the window that opens, you can create volumes and view current volumes.

Note: The AWS CloudFormation template provides two gp2 Amazon EBS volumes of a size that is specified during the CloudFormation template configuration for use with your IBM Spectrum Virtualize cluster.

Either before adding Amazon EBS volumes to a storage pool or as a part of the assignment process, be sure to follow the recommendation for correctly aligning the Amazon EBS volume type to IBM Spectrum Virtualize performance expectations in accordance with Table 4-2 on page 63.

4. To create a volume, click **Create Volume** in upper left of the window.
5. Select the volume type and size of the volume that is required, as shown in Figure 5-43.

Note: When you create an Amazon EBS volume, ensure that you choose the same Availability Zone as the existing IBM Spectrum Virtualize for Public Cloud on AWS instance.

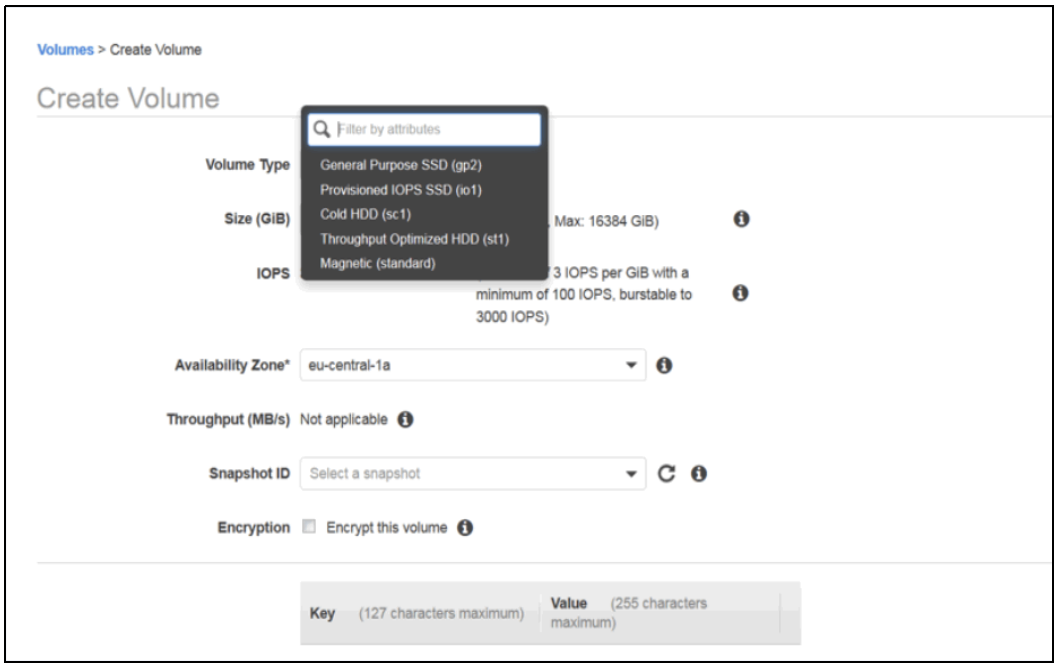

Figure 5-43 Amazon EBS: Create Volume on the AWS Console

Volumes that are created are viewable on the AWS Console in the **EBS volumes** section, and they should include a status of `Available`.

As shown in Figure 5-44, two pools are created on IBM Spectrum Virtualize for Public Cloud on AWS and each pool features one MDisk that is assigned, which is the Amazon EBS external storage that is purchased on AWS Cloud.

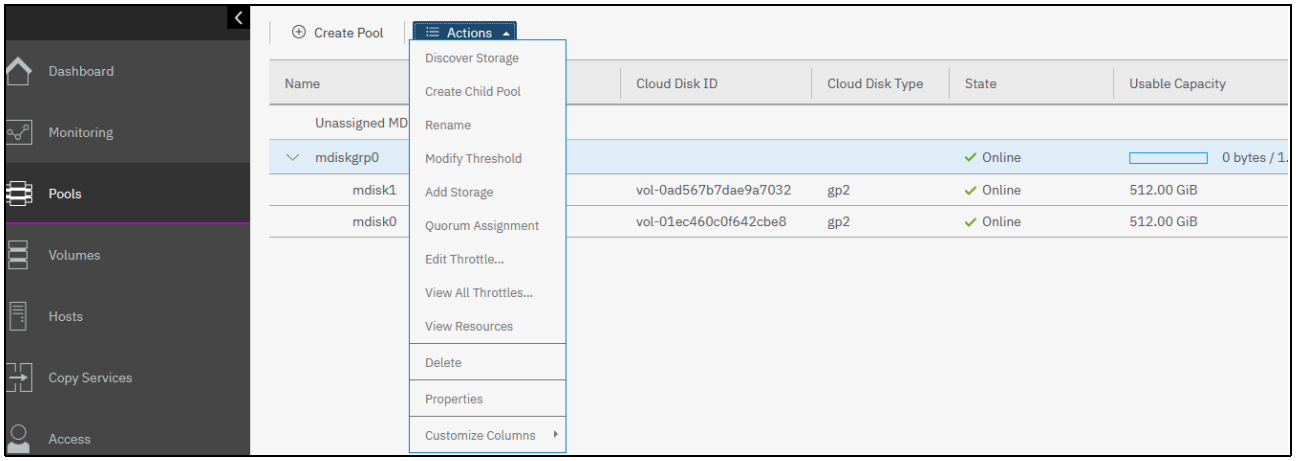

Figure 5-44 Pool creation

6. To create a pool inside IBM Spectrum Virtualize for Public Cloud on AWS, log in to the IBM Spectrum Virtualize for Public Cloud on AWS Management GUI and select **Pools** → **Create Pool**.

   This procedure also can be done through the IBM Spectrum Virtualize for Public Cloud CLI:

   a. Run the **mkmdiskgrp** command to create the pool.

   b. Run **lsmdisk** command to list the new MDisks.

   c. Run the **addmdisk** command to add the new MDisks to the new pool.

   A sample is shown in Example 5-5.

   Example 5-5 Creating a pool inside IBM Spectrum Virtualize for Public Cloud on AWS by using the CLI

```
detectmdisk
lsmdiskcandidate
mkmdiskgrp -name pool2 -ext 32
addmdisk -mdisk mdsk4:mdsk5:mdsk6:mdsk7 pool2
```

7. After the pool is created, select **Action** → **Discover Storage**. The Amazon EBS volumes that were purchased on AWS Cloud and have available space and unused are visible under Unassigned MDisk. To cross-verify that the correct volume is added to the pool, check to see whether the Amazon EBS Volume ID is the same volume ID that is seen on the AWS Cloud console.
8. Add storage in the form of MDisks to the pool. Only 16 MDisks can be used per I/O group.
9. [Optional] IBM Spectrum Virtualize for Public Cloud on AWS beginning with Version 8.3.1 supports data reduction pools (DRPs). To use data reduction technologies on IBM Spectrum Virtualize for Public Cloud, you must create a DRP, create volumes inside that DRP, and map these new volumes to hosts that support SCSI **unmap** commands. For more information about DRPs, see IBM Documentation.
10. Create a VDisk and assign the volume for host access by using iSCSI.

5.6.1 Configuring an IBM Spectrum Virtualize for Public Cloud volume for host access by using a pool

In this section, you create a volume by using the pool that was created with the Amazon EBS volumes or MDisks. Volumes can be fully allocated or thinly provisioned (space-efficient). The default pre-allocation that is indicated by the CLI in Example 5-6 is 2% (specified by the real size [rsize]). You have 98% of the capacity for the volumes that is available in the pool for other volumes until this volume claims it.

Example 5-6 Thinly provisioned (space-efficient) volume creation by using the CLI

```
svctask mkvdisk -autoexpand -grainsize 256 -mdiskgrp 2 -name thin-test -rsize 2%
-size 32212254720 -unit b -warning 80%
```

Figure 5-45 shows thinly provisioned (space-efficient) volume creation by using the GUI.

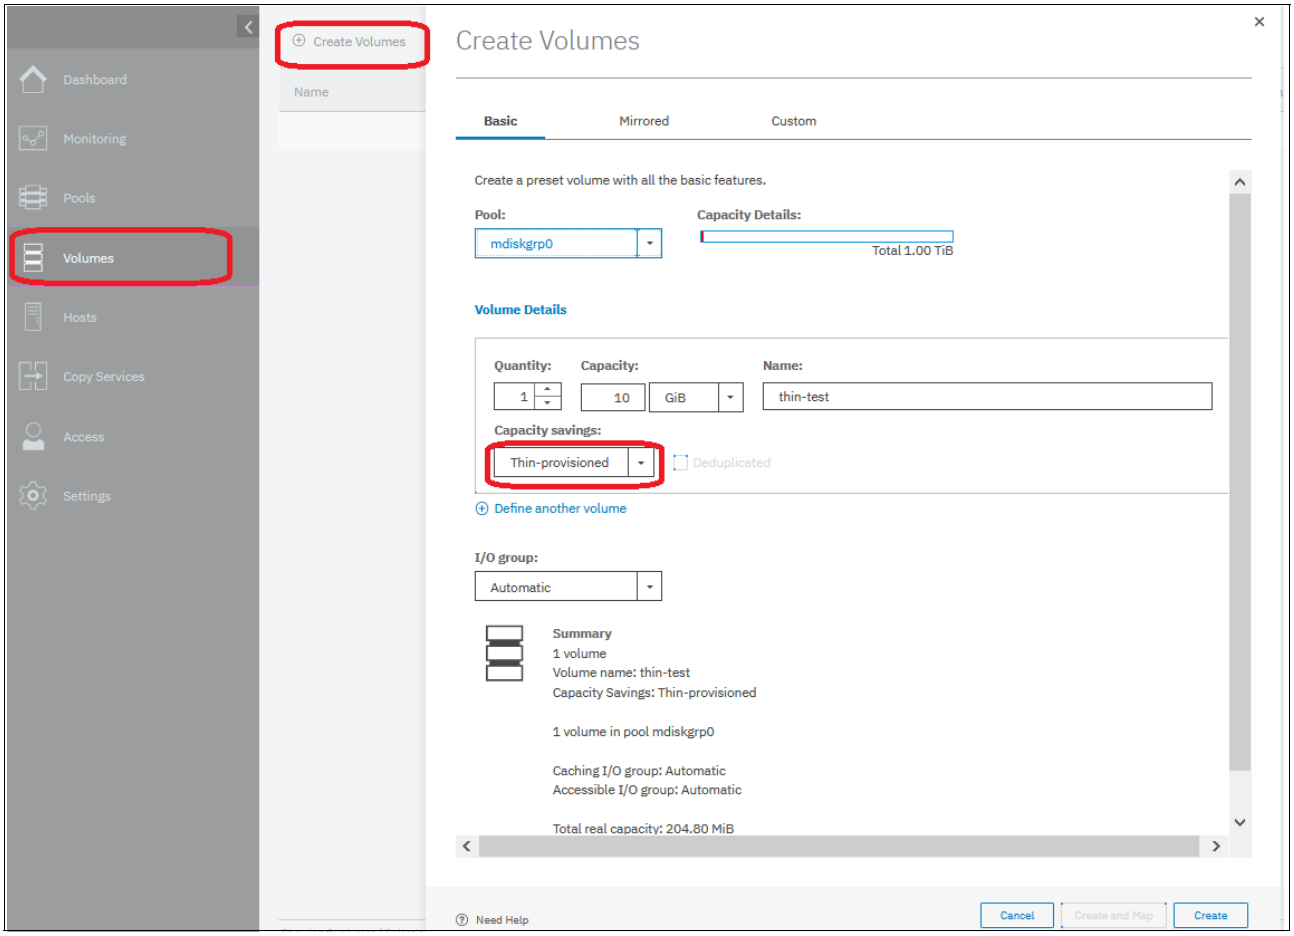

Figure 5-45 Thinly provisioned (space-efficient) volume creation with the GUI

Thinly provisioned volumes allow users to over-provision the Amazon EBS volumes, which reduce the overall operational cost in AWS.

5.6.2 Configuring the host and volume mapping

To use the volume that you created, you must map it to a host object. The host object represents a single Bare Metal Server on your cloud account and its iSCSI-qualified name (IQN), which is similar to a worldwide port name (WWPN) for an FC host.

To create a host object, you must find and collect its specific IQN. The procedure to collect the IQN varies with each operating system. For more information about the required steps, see the specific operating system's documentation.

When you create your host object and map your volume, depending on what operating system you using, you must install the iSCSI initiator and run some specific operations to use your mapped volumes with the hosts.

Linux host

Install the Linux software iSCSI initiator. The initiator software on RHEL systems is packaged as `iscsi-initiator-utils`, and the suggested version is 6.2.0.873-35 or later. The initiator software on SUSE Linux Enterprise Server systems is packaged as `open-iscsi`, and the suggested version is 6.2.0.873-33.2 or later.

According to IBM Documentation, set the IQN, target discovery, and authentication. Then, enable multipathing for the Linux hosts.

After creating the host object and mapping VDisks to it, on the IBM Spectrum Virtualize cluster, scan for the disks on the host by using the specific iSCSI command as with an anon-premises IBM Spectrum Virtualize Cluster.

Check the **`multipath`** output (run **`multipath -ll`**) to ensure that your VDisks are attached correctly through the multipath tool. A typical output of a VDisk is shown in Example 5-7.

Example 5-7 Linux multipath -ll output example

```
mpathch (3600507680181820bc800000000000009) dm-1 IBM     ,2145
size=500G features='1 queue_if_no_path' hwhandler='0' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 26:0:0:5 sdf  8:80   active ready running
| `- 27:0:0:5 sdl  8:176  active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 28:0:0:5 sdr  65:16  active ready running
  `- 29:0:0:5 sdx  65:112 active ready running
```

Windows host

The software iSCSI initiator is built in to the operating system on Windows 2008 and later. Access the iSCSI initiator from the Control Panel or search from the **Start** menu. An example is shown in Figure 5-46.

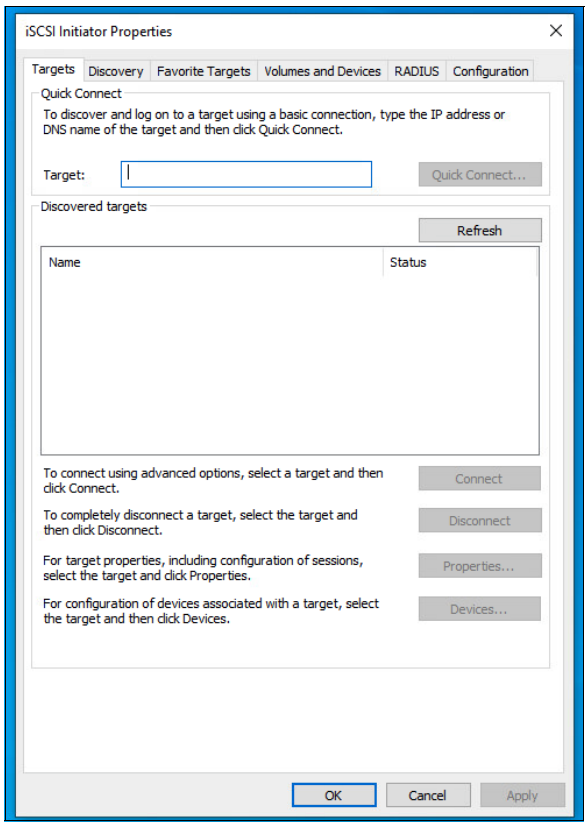

Figure 5-46 Windows iSCSI window to configure and find an IQN and storage

Discover the iSCSI target by using Send Targets or by using iSNS. For more information, see IBM Documentation.

Connect to the discovered targets, as described in IBM Documentation.

Now, the mapped volumes are visible to Windows disk management services. The system volumes can be initialized, formatted, and mounted. You can view the details of the discovered disks by using the Windows Command Prompt. An example output is shown in Example 5-8.

Example 5-8 Windows OS DISKPART command example

```
DISKPART> list disk
 Disk ###  Status      Size     Free     Dyn  Gpt
 --------  ----------  -------  -------  ---  ---
 Disk 0    Online       149 GB    78 GB   *
 Disk 1    Online       149 GB    78 GB   *
 Disk 2    Online       565 MB   565 MB
 Disk 3    Online       337 MB   337 MB
DISKPART> select disk 2
Disk 2 is now the selected disk.
DISKPART> detail disk
IBM      2145             SCSI Disk Device
```

```
Disk ID: 00000000
Type   : iSCSI
Bus    : 0
Target : 2
LUN ID : 0
There are no volumes.
```

5.7 Configuring a site-to-site virtual private network IPsec tunnel for hybrid cloud connectivity in AWS Cloud

This section describes how to configure hybrid cloud connectivity between the AWS Cloud and the on-premises environment. This section also describes the lab setup and the steps to configure the site-to-site IPsec tunnel for communication between AWS Cloud and the on-premises site.

The VPN IPsec site-to-site tunnel creates a secure communication network between the AWS Cloud infrastructure and on-premises infrastructure. Network communication between the private subnets is controlled by the access control list (ACL) that is populated when you create the VPN IPsec site-to-site tunnel.

AWS configuration for the VPN IPsec tunnel

Complete the following steps at the VPC level in AWS Cloud to establish the IPsec tunnel:

1. Create a customer gateway by logging in to the AWS console with resource provisioning privileges. Select **Services** at the upper left, and then, **VPC**. Select **Virtual Private Network (VPN)** in the pane on the left. Then, click the customer gateways and enter the required details.
2. Create the virtual private gateways by clicking the **Virtual private gateways** section in the VPC and configure the required details.
3. Attach a virtual private gateway to the VPC.
4. Create a site-to-site VPN connection in AWS Console by selecting the virtual private gateway and customer gateway parameters. Attach the virtual private gateway to the VPC in AWS.
5. After the site-to-site connection is complete, a configuration file is generated for the end-to-end point. This step creates two tunnels in the VPC. The same configuration file is used for the configuration at the other end of the tunnel.

5.8 Configuring replication from an on-premises IBM Spectrum Virtualize array to IBM Spectrum Virtualize for Public Cloud on AWS

This section describes how to configure replication from an on-premises IBM FlashSystem solution or SAN Volume Controller system to an IBM Spectrum Virtualize for Public Cloud on AWS solution.

Our example uses an IBM FlashSystem array in an on-premises data center and a 2-node IBM Spectrum Virtualize for Public Cloud on AWS as a disaster recovery (DR) storage solution.

This scenario uses IBM Spectrum Virtualize *volume groups (VGs)* to replicate the data from the on-premises data center to AWS Cloud.

This implementation starts with the assumption that the IP network connectivity between the on-premises data center IBM FlashSystem array and AWS Cloud is established through a *Multiprotocol Label Switching* (MPLS) or VPN connection. Because several methods are available to implement the IP network connectivity, this section does not consider that specific configuration. For more information, contact your organizations's network technical specialist.

To configure the replication, complete the following steps:

1. Go to **Copy Services** and click **Create Partnership**. This configuration is required on both sites, as shown in Figure 5-47.

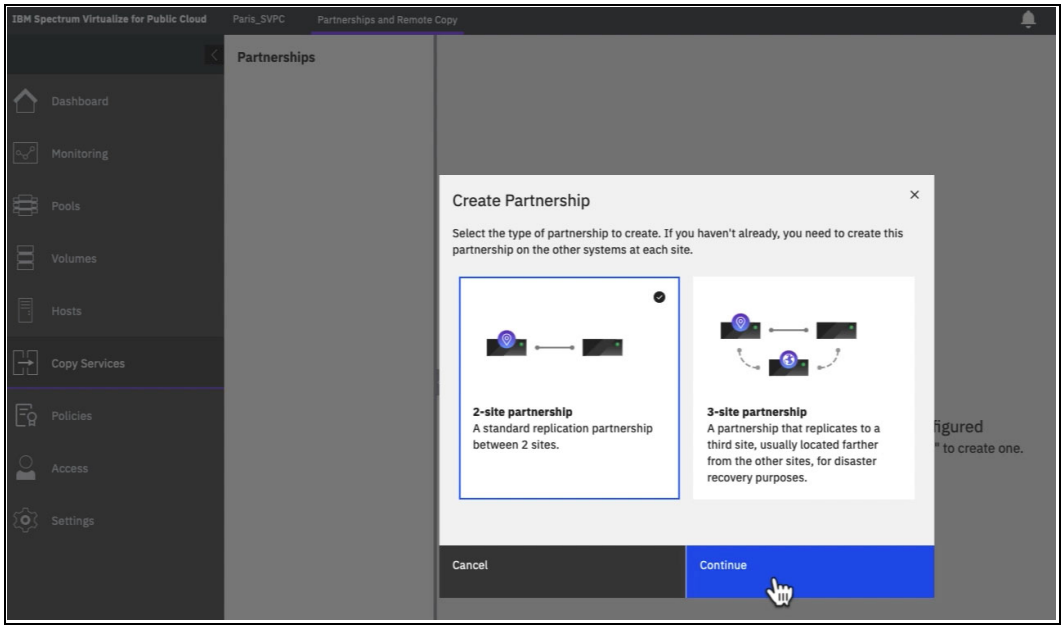

Figure 5-47 Setting up a partnership between IBM Spectrum Virtualize for Public Cloud and an on-premises IBM FlashSystem

2. Input the IP address of the target system, as shown in Figure 5-48.

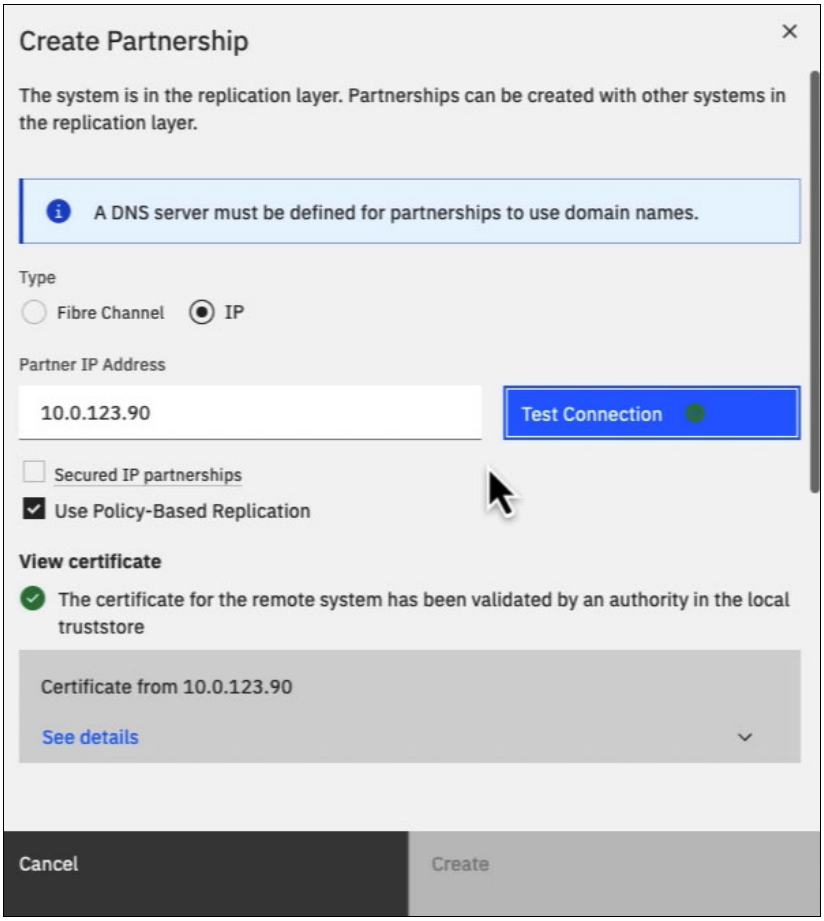

Figure 5-48 Entering an IP address and testing a Replication Partnership configuration example

3. Scroll down and choose the bandwidth and portset, as shown in Figure 5-49.

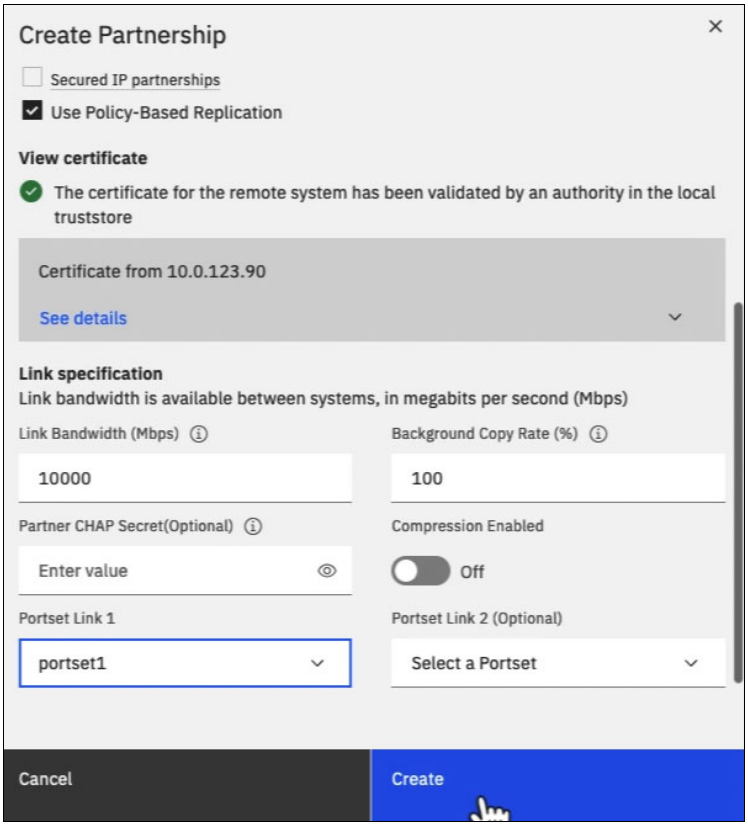

Figure 5-49 Partnership bandwidth and portset choices

4. Repeat steps 1 on page 122 - 3 on page 124 for the other Virtualize instance, pointing the IP address to the other instance. Link the pools so that volumes in the source pool are replicated to the target replication pool system, which can be done as a child or parent pool. Optionally, a provisioning policy can be used.

   A sample is shown in Figure 5-50.

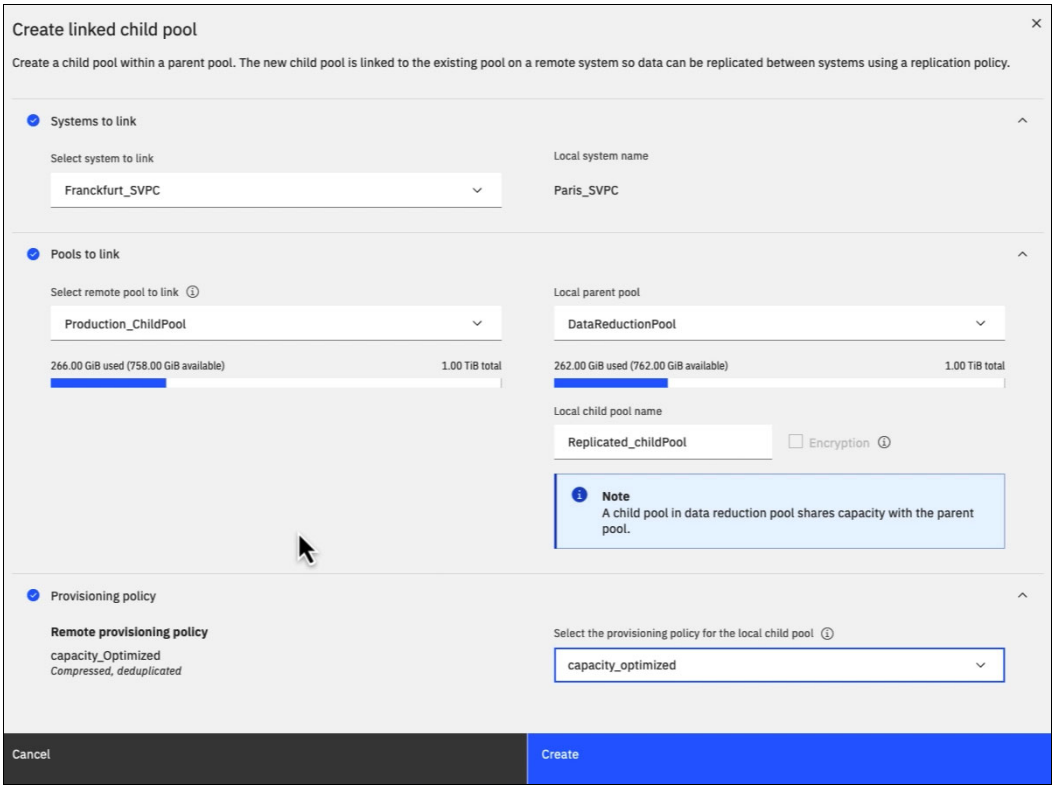

Figure 5-50 Linking the primary and target pools by using child pools replication example

5. Following the checklist, create the replication policy, as shown in Figure 5-51.

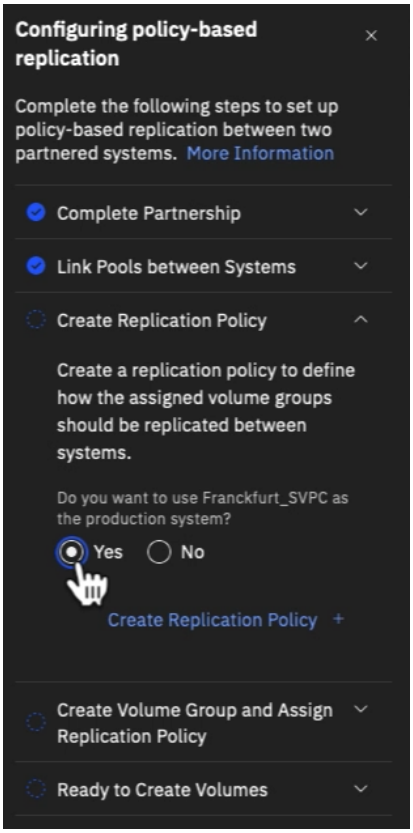

Figure 5-51 Replication checklist: Create Replication Policy

6. Create the replication policy and configure the recovery point objective (RPO) value based on business needs, as shown in Figure 5-52.

Note: Know what versions of IBM Spectrum Virtualize software are supported. For supported and interoperability versions, see IBM Spectrum Virtualize Family of Products Inter-System Metro Mirror and Global Mirror Compatibility Cross Reference.

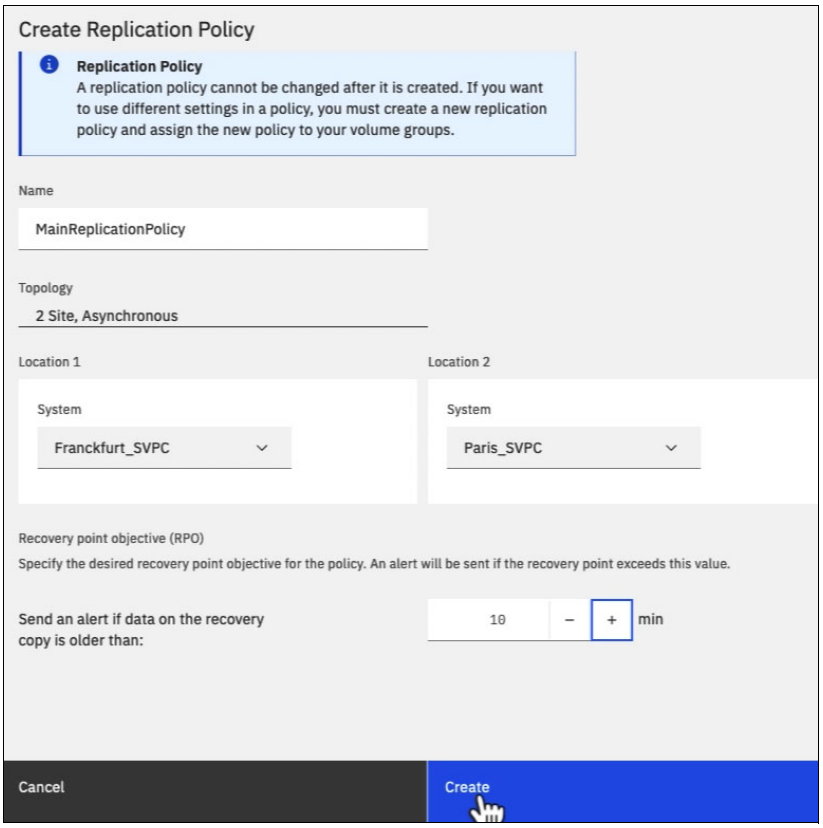

Figure 5-52 Replication Policy Creation

7. Create the VG on the source or primary Virtualize system, as shown in Figure 5-53.

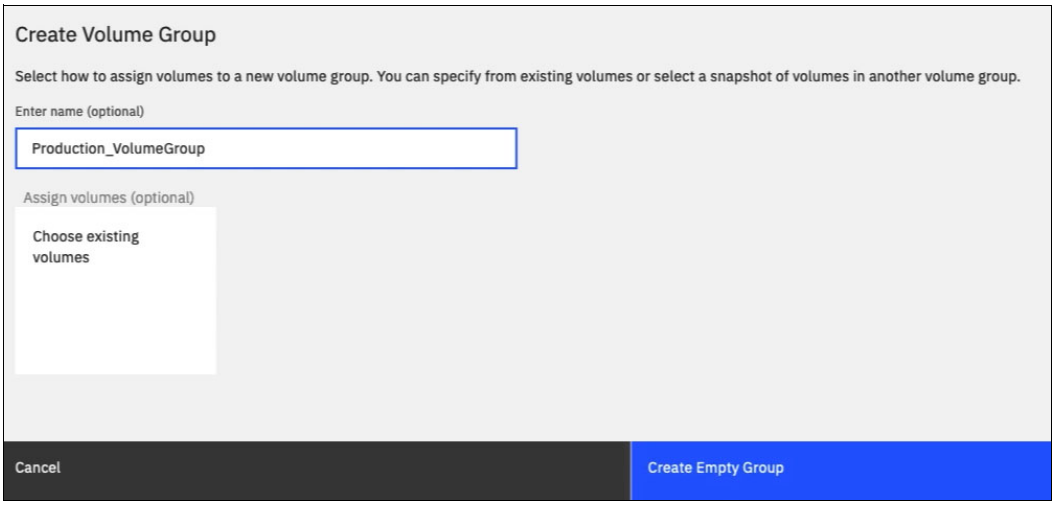

Figure 5-53 Creating a volume group on a IBM Spectrum Virtualize for Public Cloud source example

8. Assign the replication policy, as shown in Figure 5-54.

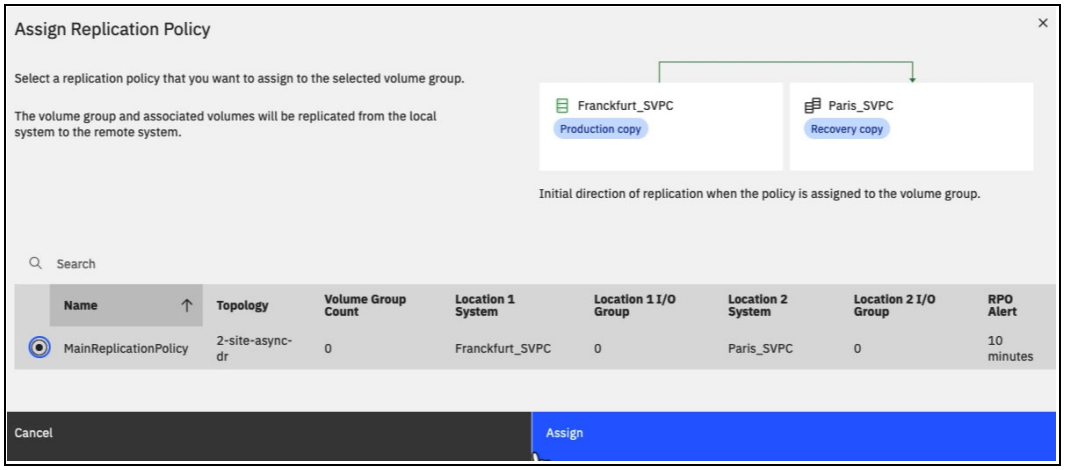

Figure 5-54 Assigning a replication policy example

As shown in Figure 5-55, you completed the various steps of partnership, that is, linking the source and target pools, and creating a replication policy and its RPO. Now, you are ready to create volumes (or add existing source volumes into this VG, which will be replicated).

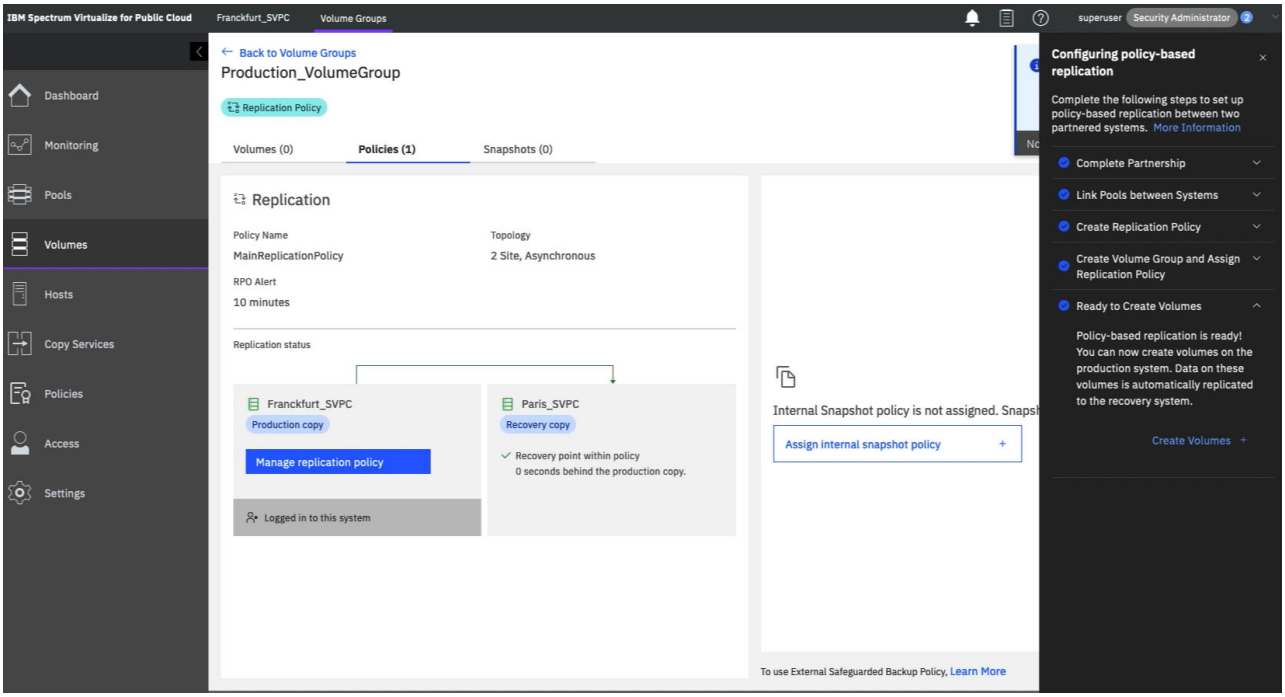

Figure 5-55 Replication policy, volume group, and linked pools complete

9. Complete the partnership configuration in the IBM Spectrum Virtualize for Public Cloud on AWS on the AWS side by providing an on-premises cluster IP address.

   Your partnership is fully configured, as shown in Figure 5-56.

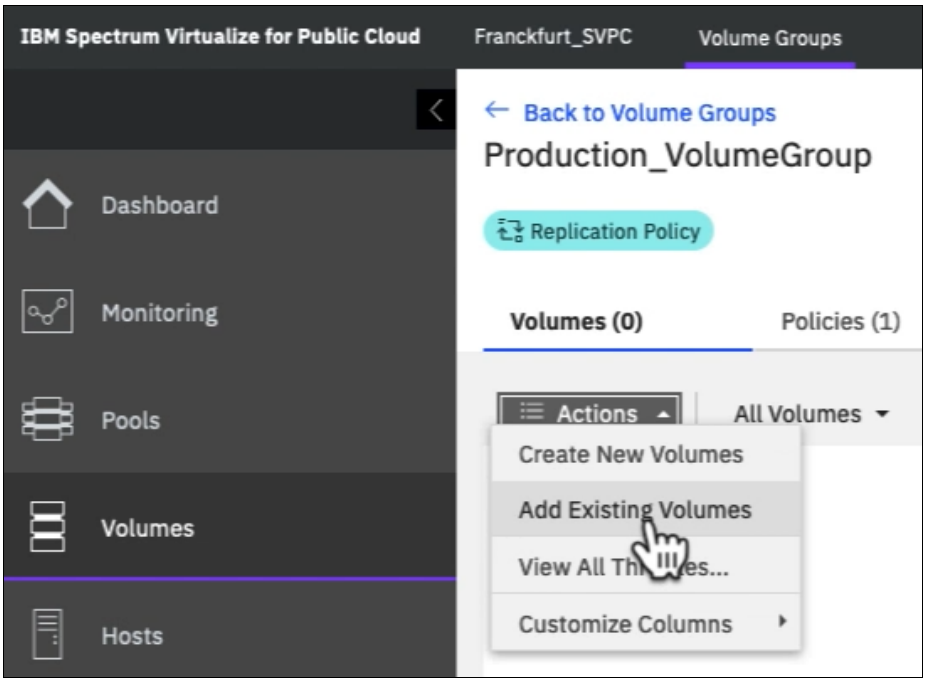

Figure 5-56 Adding or creating volumes in the volume group for replication

Tip: The connection might take a few seconds to synchronize, but double-clicking **Partnership** shows the confirmed status of the partnership.

10. In our example, four 100 GiB volumes are added to the VG for replication, as shown in Figure 5-57.

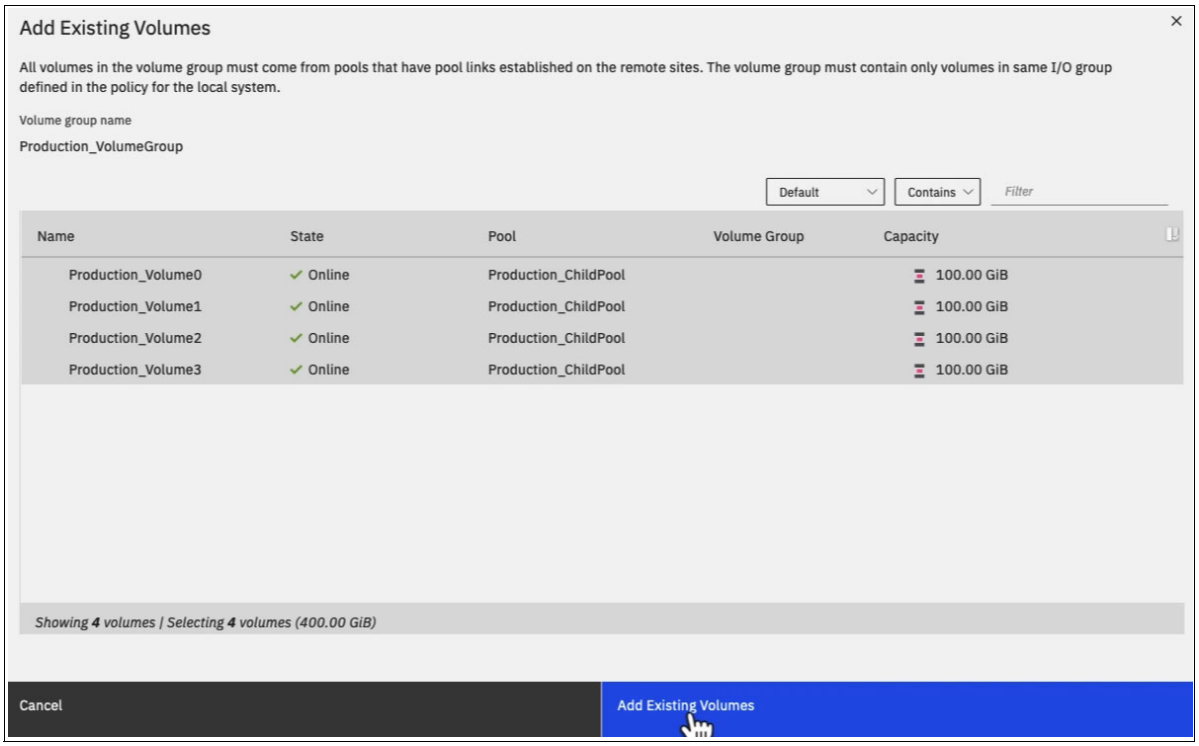

Figure 5-57 Adding existing volumes to a volume group example

11. After adding the volumes, immediately the replication begins, as shown in Figure 5-58.

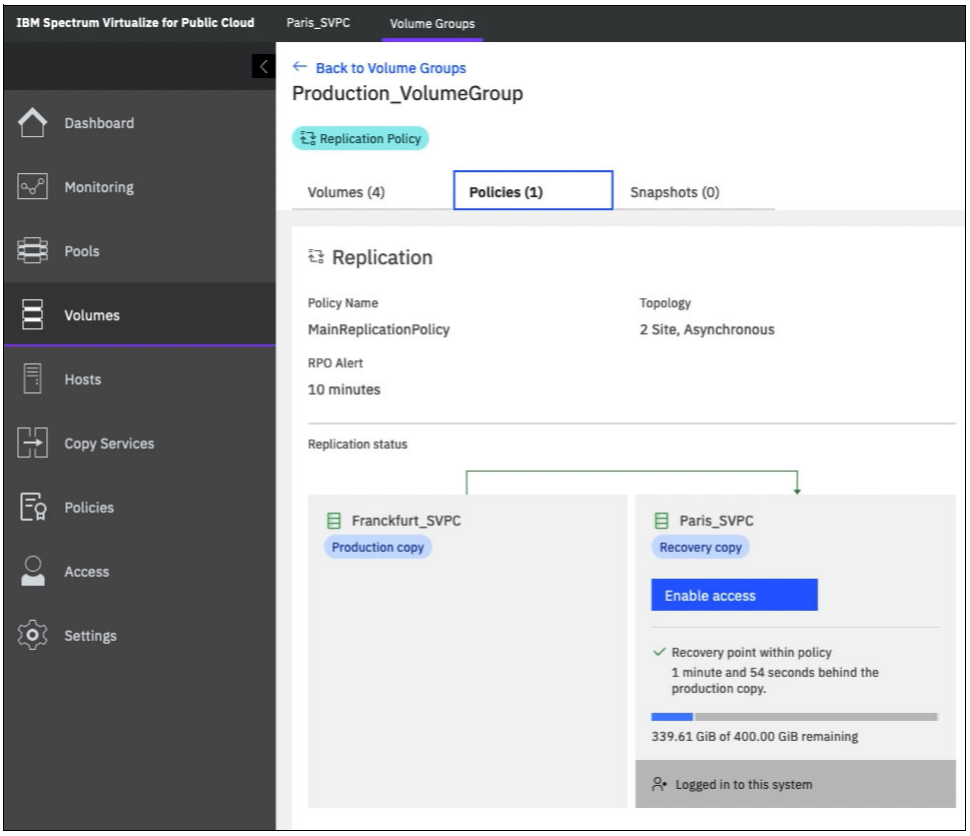

Figure 5-58 Volumes immediately begin replicating to the target array after being added

6

Implementing IBM Spectrum Virtualize for Public Cloud on Microsoft Azure

This chapter describes how to implement an IBM Spectrum Virtualize for Public Cloud on Microsoft Azure environment and includes the following topics:

- Installing IBM Spectrum Virtualize for Public Cloud on Azure
- Logging in to IBM Spectrum Virtualize for Public Cloud on Azure
- Finishing the setup of IBM Spectrum Virtualize for Public Cloud on Azure
- Configuring the IBM Spectrum Virtualize for Public Cloud Azure cloud quorum
- Configuring the back-end storage
- Adding more back-end storage
- Configuring a site-to-site virtual private network gateway for hybrid cloud connectivity in Azure cloud
- Configuring replication from on-premises IBM Spectrum Virtualize to IBM Spectrum Virtualize for Public Cloud on Azure

6.1 Installing IBM Spectrum Virtualize for Public Cloud on Azure

Deployment video: As part of this IBM Redbooks publication, the authors created an IBM Spectrum Virtualize for Public Cloud on Azure deployment video.

The IBM Spectrum Virtualize for Public Cloud software is a Bring Your Own License (BYOL) offering in Azure Marketplace. During the installation, the template verifies proof of entitlement to ensure that a valid license is purchased from IBM. If the proof of entitlement is not present, the installation fails.

As described in 4.4, “Planning for Microsoft Azure” on page 67 to obtain the license and proof of entitlement that is needed for the software, complete the following steps:

1. See the IBM Passport Advantage web page to obtain a license and proof of entitlement for the software.
2. On the web page, follow the directions to enter your IBM Customer Number and the maximum number of terabytes of virtual storage that you want to provision your systems.

The IBM Spectrum Virtualize for Public Cloud installation uses Azure Resource Manager (ARM) templates that simplify provisioning and management on Azure. These templates are available on Azure Marketplace and simplify the provisioning and installation process.

Ensure that all prerequisites are complete before you install the IBM Spectrum Virtualize for Public Cloud software from the Azure Marketplace.

Important: Before installing IBM Spectrum Virtualize for Public Cloud on Azure, ensure that the following tasks are complete:

- A valid Azure Account is created if such an account does not exist.
- Any extra Azure cloud profiles for users are created.
- Resource providers are registered for your subscription.

Azure Marketplace provides user with guided windows to collect deployment-specific information and ease the deployment user experience. These dynamic windows provide pricing estimates to users that are based on the user’s resource selections.

To install the IBM Spectrum Virtualize for Public Cloud software on Azure, complete the following steps:

1. Go to the Azure Portal and log in by using the installer or administrator user profile for your Azure Account.
2. At the portal, search for “IBM Spectrum Virtualize for Public Cloud” in the search bar to see the offering from Azure Market place. This search result displays the available product, as shown in Figure 6-1. Review the Plans and the Support sections and then, click **Create**.

Figure 6-1 IBM Spectrum Virtualize for Public Cloud on Azure Marketplace[1]

[1] Microsoft Azure screen captures in this chapter are used with permission from Microsoft.

3. Enter the basic information about your deployment on Azure (see Figure 6-2):
   - Subscription to use for deployment.
   - Create a resource group (RG) to group all resources.
   - Region that hosts the resources.
   - Project Tag to uniquely identify the deployed resources.
   - (Optional) Enable Rollback support if a deployment failure occurs.

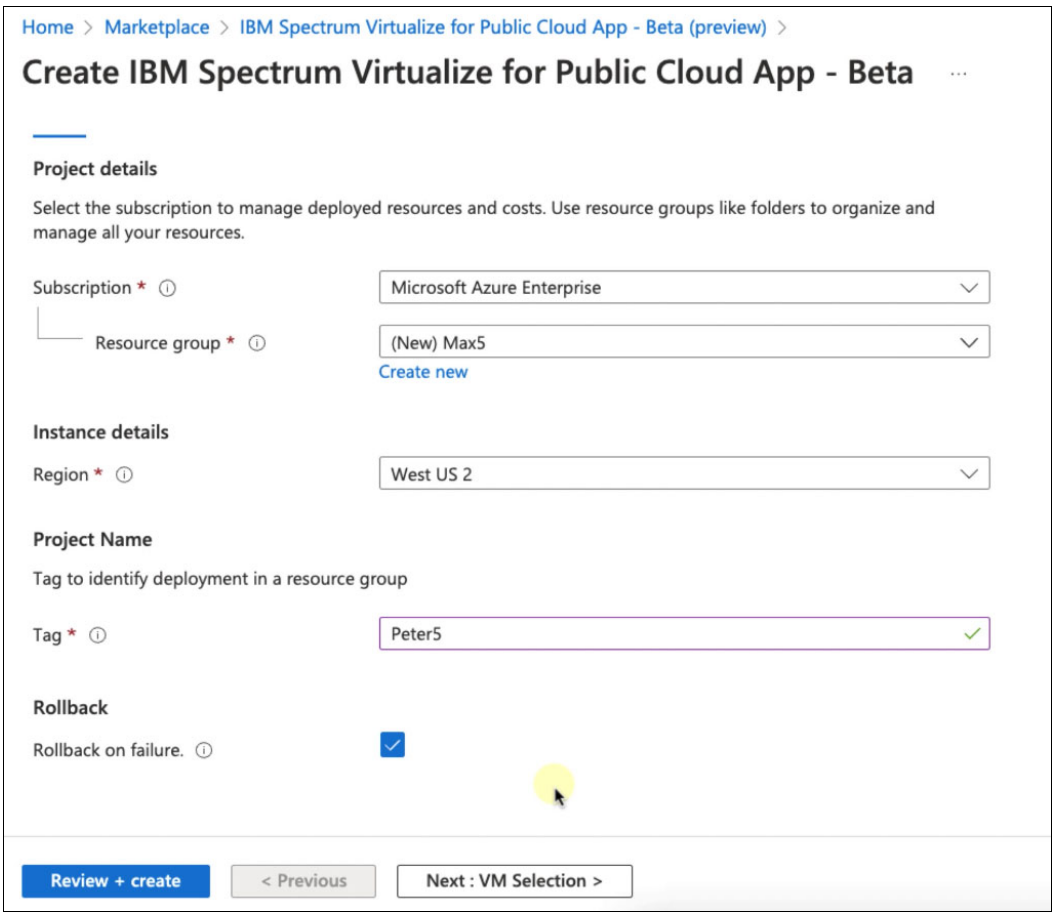

Figure 6-2 First configuration window for IBM Spectrum Virtualize for Public Cloud deployment on Azure

Note: IBM Spectrum Virtualize requires an empty RG for any deployment.

4. Select the size of virtual machine (VM) in VM Settings tab. IBM Spectrum Virtualize for Public Cloud on Azure classifies VMs in two categories:
   - IBM Spectrum Virtualize for Public Cloud Node
   - Quorum Node

   The IBM Spectrum Virtualize for Public Cloud Node supports three VMs from Microsoft Azure D-series V3 VMs. By default, the `D16s_v3 VM` type is selected. For more information about the VM size that is supported by IBM Spectrum Virtualize, see 4.4, “Planning for Microsoft Azure” on page 67.

These VMs are used to run the IBM Spectrum Virtualize software inside a container and facilitate the IBM Spectrum Virtualize features on Azure cloud. The Quorum Node supports a fixed size `B-Series-B1 VM`. It is used to run the quorum management service and maintain the quorum state for the two nodes that are being deployed.

Figure 6-3 shows the information that is available during the VM selection.

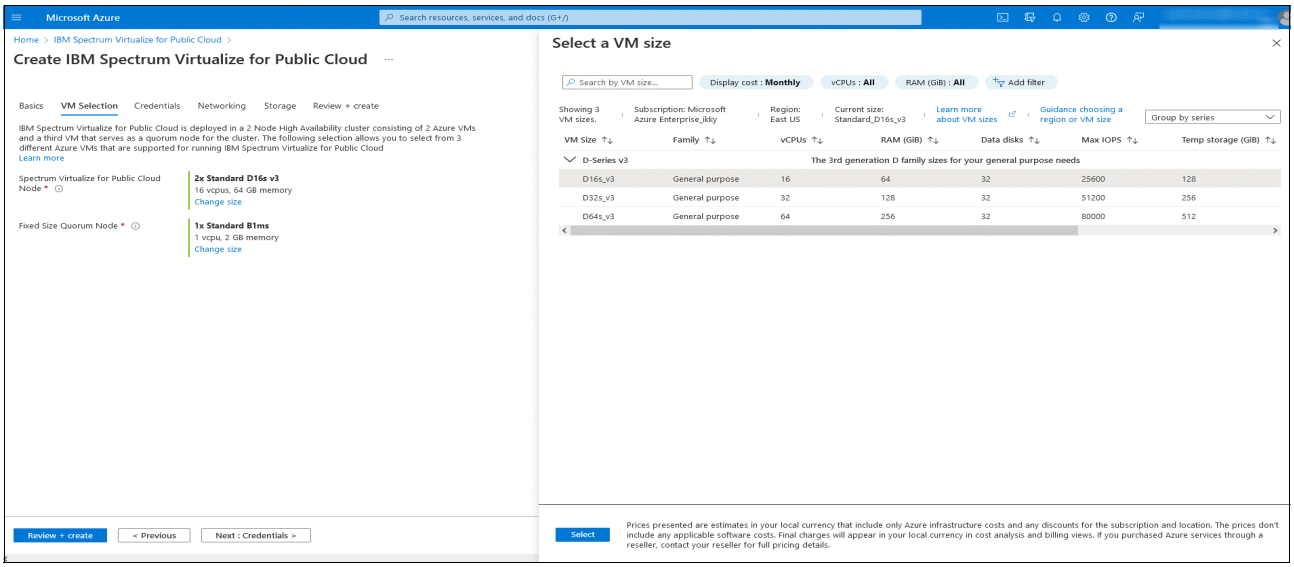

Figure 6-3 IBM Spectrum Virtualize for Public Cloud deployment: VM options with details

5. The Credentials window (see Figure 6-4) shows all the credentials that are related to IBM Spectrum Virtualize deployment; namely, the password for IBM Spectrum Virtualize cluster, Customer entitlement check, Secure Shell (SSH) public key for access to VM and the notification email address.

   Note: Enter a customer number for a user who is entitled to the IBM Spectrum Virtualize License on Azure. An invalid number results in a failed deployment.

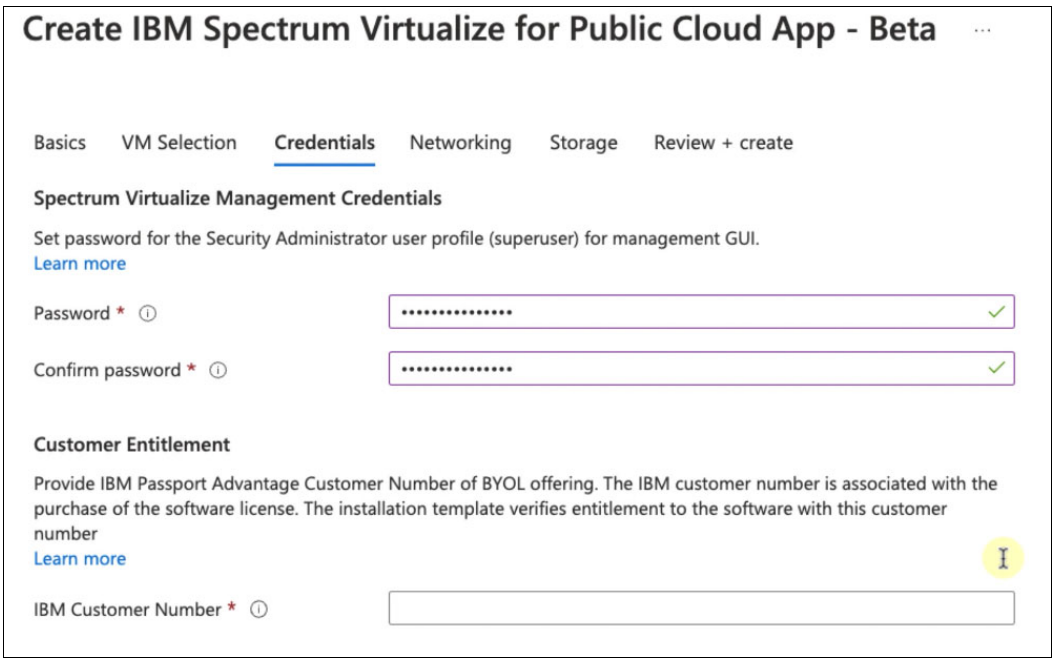

Figure 6-4 IBM Spectrum Virtualize for Public Cloud: Credentials and Customer account

6. In Figure 6-5 on page 139, choose one of the following options:
   - Create an SSH KeyPair.
   - Use an existing KeyPair from the Azure account.
   - Use your own created KeyPair.

   This KeyPair is needed to access the Nodes through secured communications. When finished, click **Next** to access the Networking items.

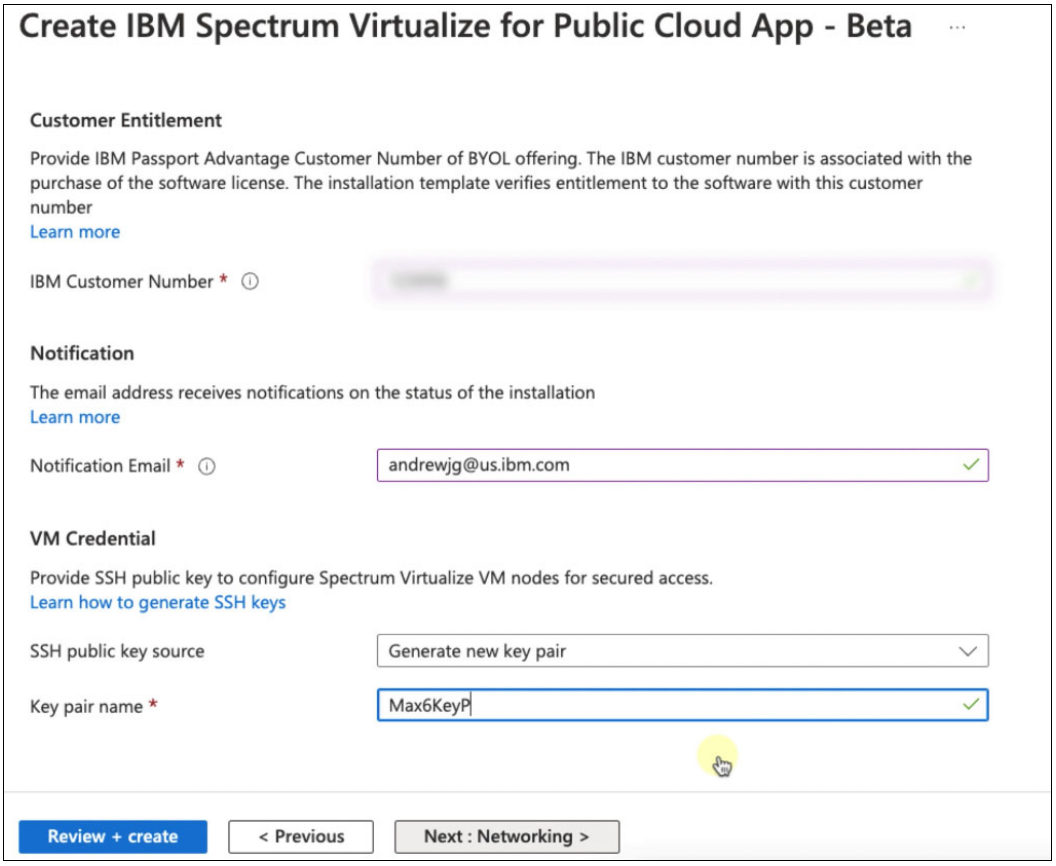

Figure 6-5 IBM Spectrum Virtualize for Public Cloud: email notification and VM nodes SSH KeyPair creation

7. In the **Networking** tab, select the **Azure Virtual Network** for this deployment. Users can select a vNet or create a virtual network (vNet) for deployment. As part of extra security and resource management, the subnets for Quorum and IBM Spectrum Virtualize nodes are separated. If the user plans to use an existing subnet, two subnets are needed in the existing VN: one for quorum, and one for the cluster nodes with minimum 255 IP addresses.

Figure 6-6 shows the **Network** tab during deployment.

Notes: Consider the following points:

- The same subnet cannot be used for node and quorum deployments.
- The public IP can be added to the quorum node after the deployment process is complete.

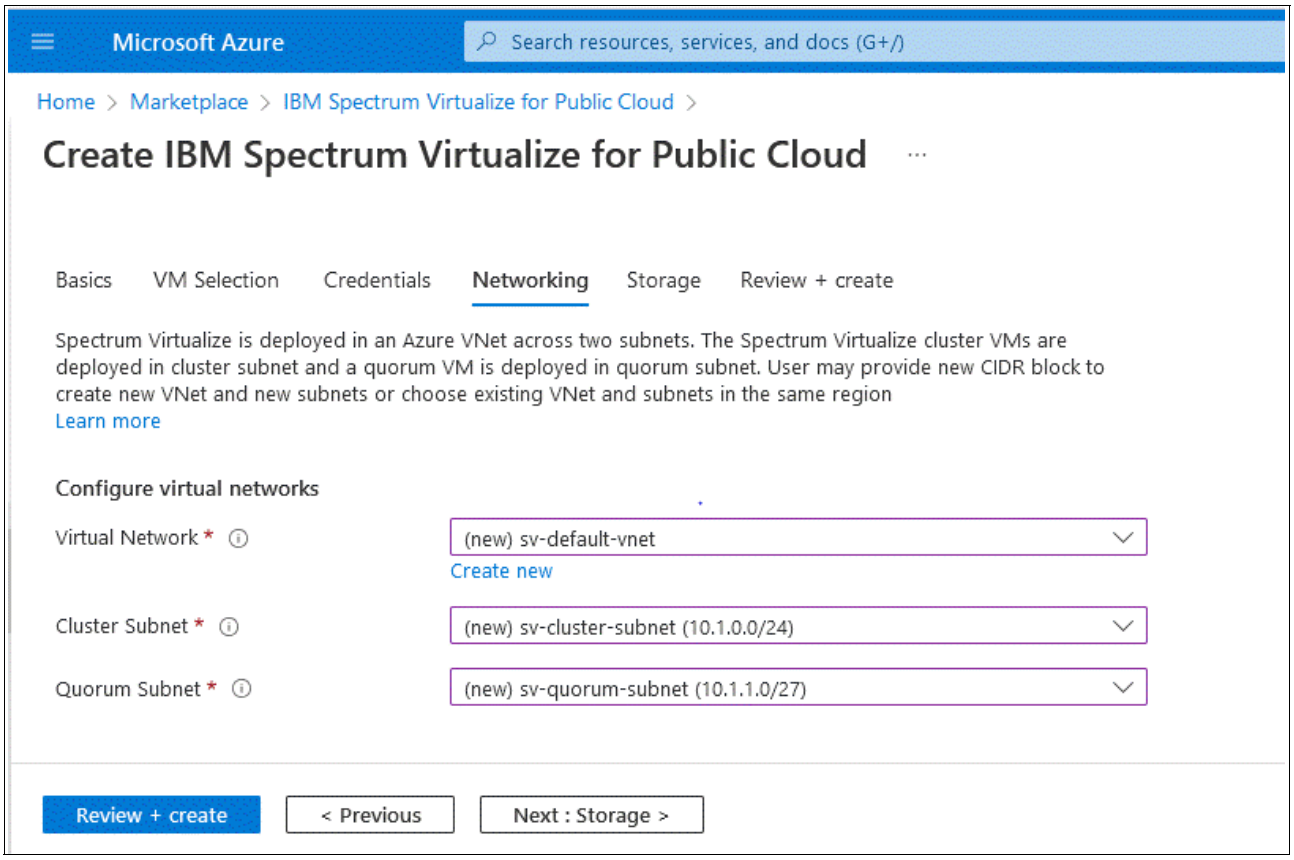

Figure 6-6 Network selection during deployment

8. In the **Storage** tab, users select the default back-end storage that will be attached to the IBM Spectrum Virtualize nodes as part of the deployment. By default, two storage disks with similar configurations that are provided by the user are provisioned and attached to the IBM Spectrum Virtualize nodes. The user can choose from various supported types of Azure disks. The back-end storage for Azure deployment is shown in Figure 6-7 on page 141.

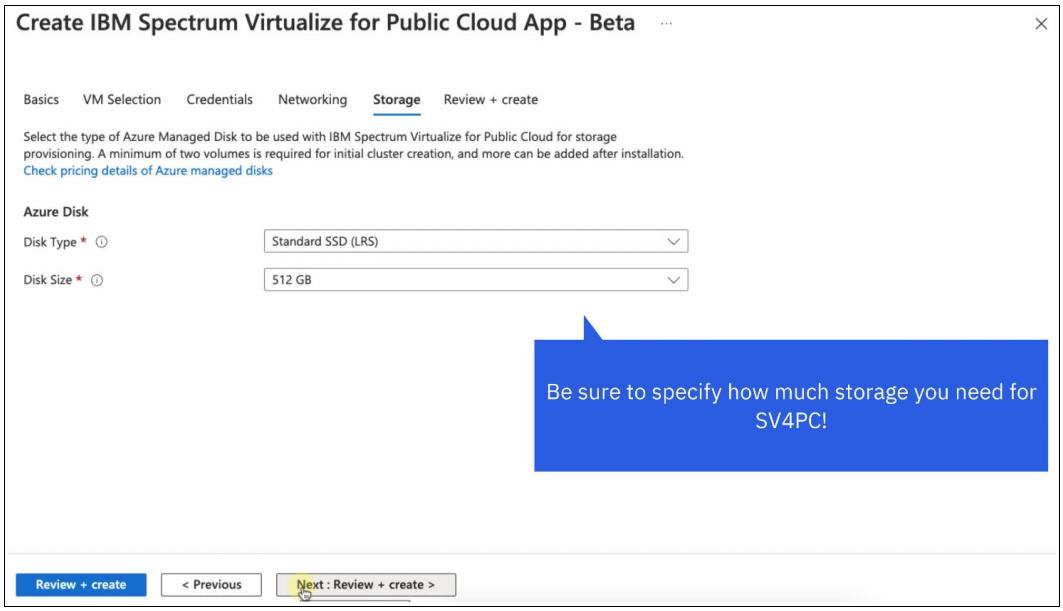

Figure 6-7 IBM Spectrum Virtualize for Public Cloud: Azure Managed Disk back-end storage selection

9. After all selections are made by the user, a basic validation runs to verify the information. In the **Review and Create** tab, an option is available to validate the availability of Azure resources in the user-selected region and ensure that the deployment does not fail because of a lack of Azure resources.

The successful validation window before the deployment is triggered is shown in Figure 6-8.

Note: By clicking **Create**, the user agrees to the terms of deployment and costs that are related to the configuration.

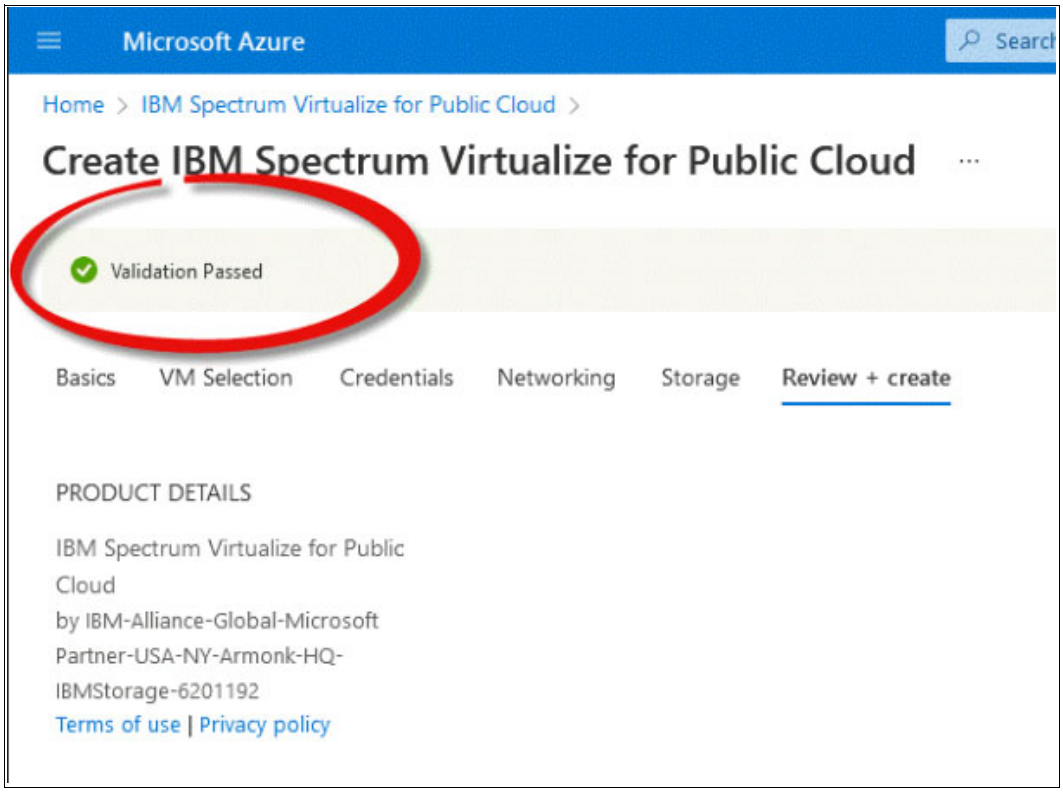

Figure 6-8 Validation Passed message

10. The deployment process takes approximately 30 minutes to complete. Progress can be monitored by reviewing the Deployment window, which is displayed after you click **Create** (see Figure 6-9).

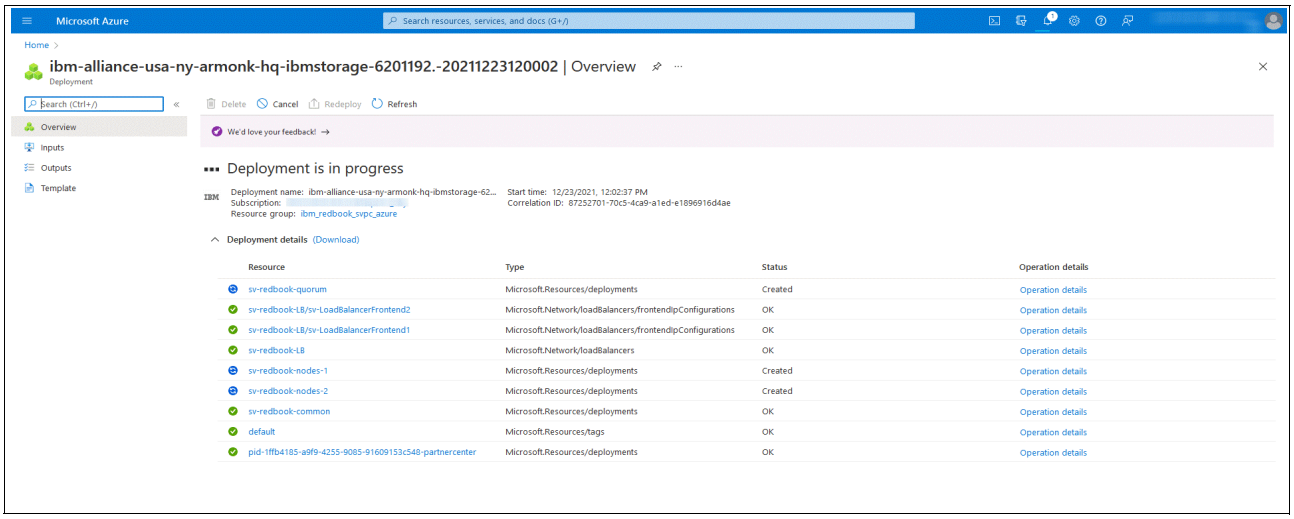

Figure 6-9 Deployment in progress

Alternatively, the user can also monitor the deployment progress in the Deployment sections that are under the RG that is used for deployment, as shown in Figure 6-10.

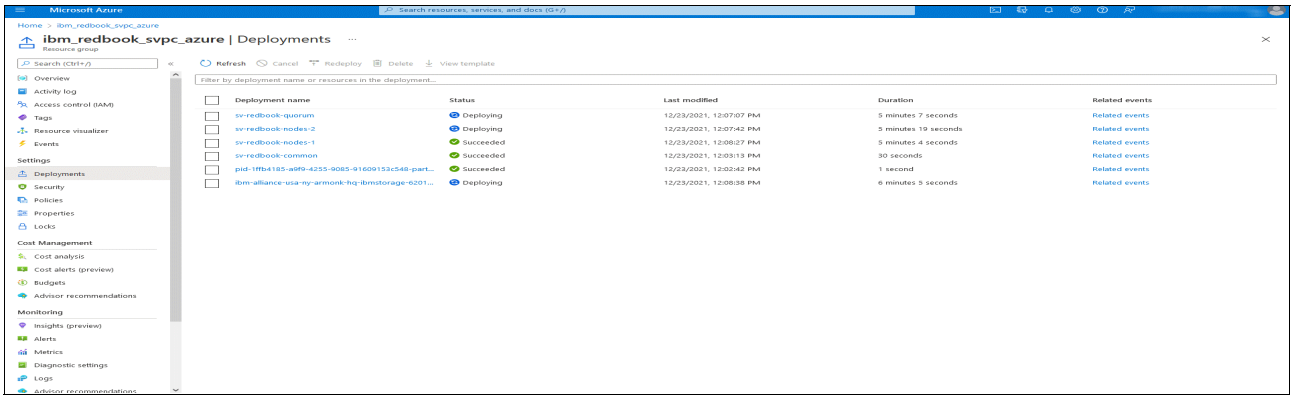

Figure 6-10 Deployment progress under a resource group

Figure 6-11 shows the successful deployment for IBM Spectrum Virtualize for Public Cloud on Azure.

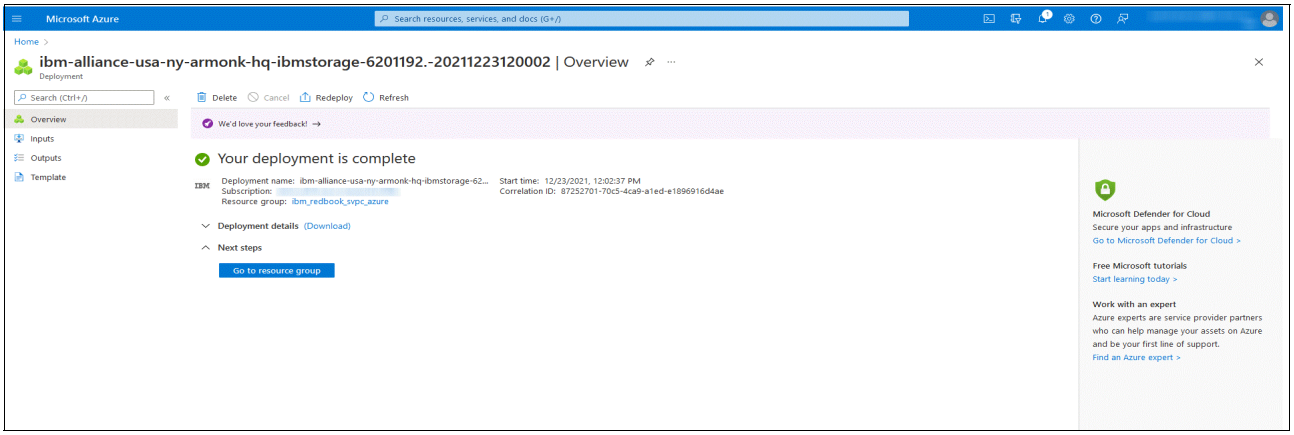

Figure 6-11 Successful deployment

11. Click the **Outputs** tab to check the various IP addresses that are deployed, as shown in Figure 6-12.

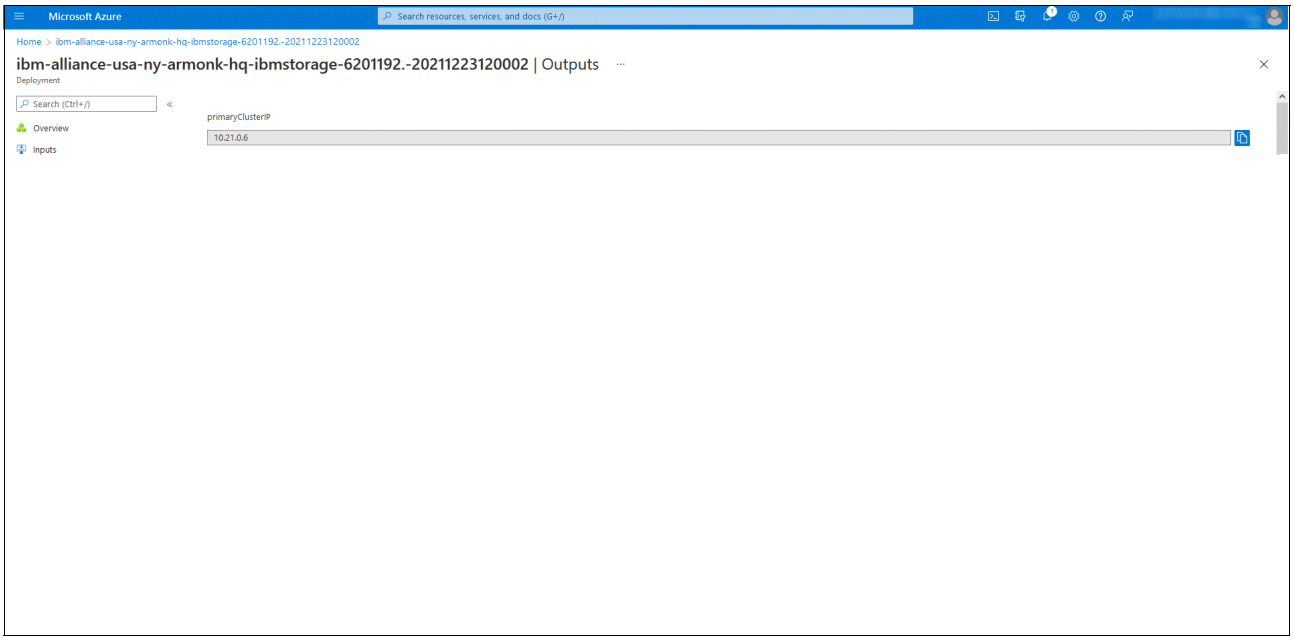

Figure 6-12 Deployment output window

After successful deployment, a window opens in which the successful deployment is shown. An email also is sent to the address that was provided by the user during the setup process. This email includes more information about the deployment (see Figure 6-13).

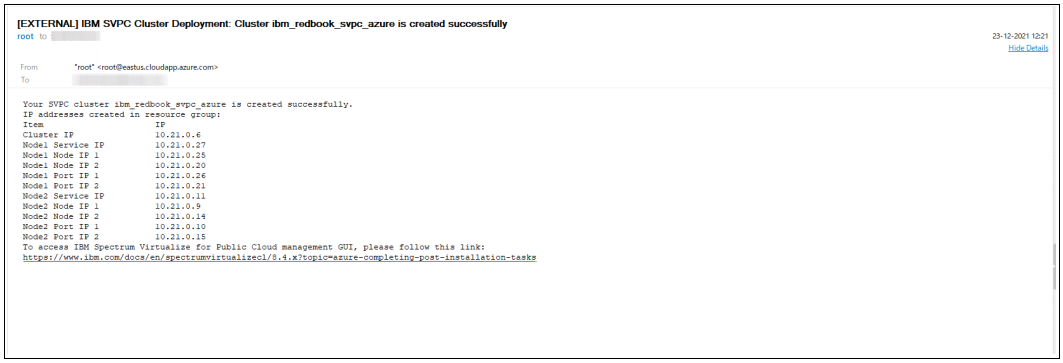

Figure 6-13 Sample email of successful deployment

6.2 Logging in to IBM Spectrum Virtualize for Public Cloud on Azure

IBM Spectrum Virtualize for Public Cloud on Azure supports both public and private deployments in the Azure cloud. Depending on the network options that are selected during deployment, an optional public IP address is assigned to the quorum node. Depending on your business requirements, the IBM Spectrum Virtualize for Public Cloud deployment might be isolated from all internet traffic and controlled by Azure Network Security Group (Azure NSG) rules.

Note: To allow traffic from specific port or IP address, NSG rules for VM network interface can be modified by using the Azure portal.

Because all traffic passes through a private network connection, network access to the Azure VM can be enabled by using the following methods:

- Azure Bastion Service
- Azure VPN gateway

IBM Spectrum Virtualize for Public Cloud on Azure deploys all resources within a new or existing vNet with an optional public IP address to the quorum node. Therefore, to access the resources that are provisioned by the deployment, a user has several options to connect to the IBM Spectrum Virtualize for Public Cloud instance. In addition to the SSH method, the Azure Bastion Service, Remote Desktop Protocol (RDP), or Azure Virtual Private Network (VPN) gateway are available. The simplest method is to use the Azure Bastion Service because it is built in to the HTML5 session of the authenticated Azure account.

6.2.1 Configuring the Azure Bastion Service

Azure Bastion is a service with which you connect to a VM in a VPN by using your web browser and Azure Portal. The Azure Bastion is a fully managed platform as a service (PaaS) that is provisioned inside the vNet. For more information, see this web page.

Complete the following steps to configure an Azure Bastion Service:

1. Create a subnet for Azure Bastion service under the vNet to be used for the Azure Bastion service. The user can choose to use a vNet that was used for the deployment of IBM Spectrum Virtualize for Public Cloud on Azure or configure a vNet.

   Note: If a new network is used, the user must configure vNet peering between the networks that use Azure Bastion Service.

The name of subnet to be created must be `AzureBastionSubnet`. The IPv4 CIDR Block range for IP addresses that are assigned to the subnet is unlimited. Figure 6-14 shows a subnet that was created under a vNet to be used for the Azure Bastion Service.

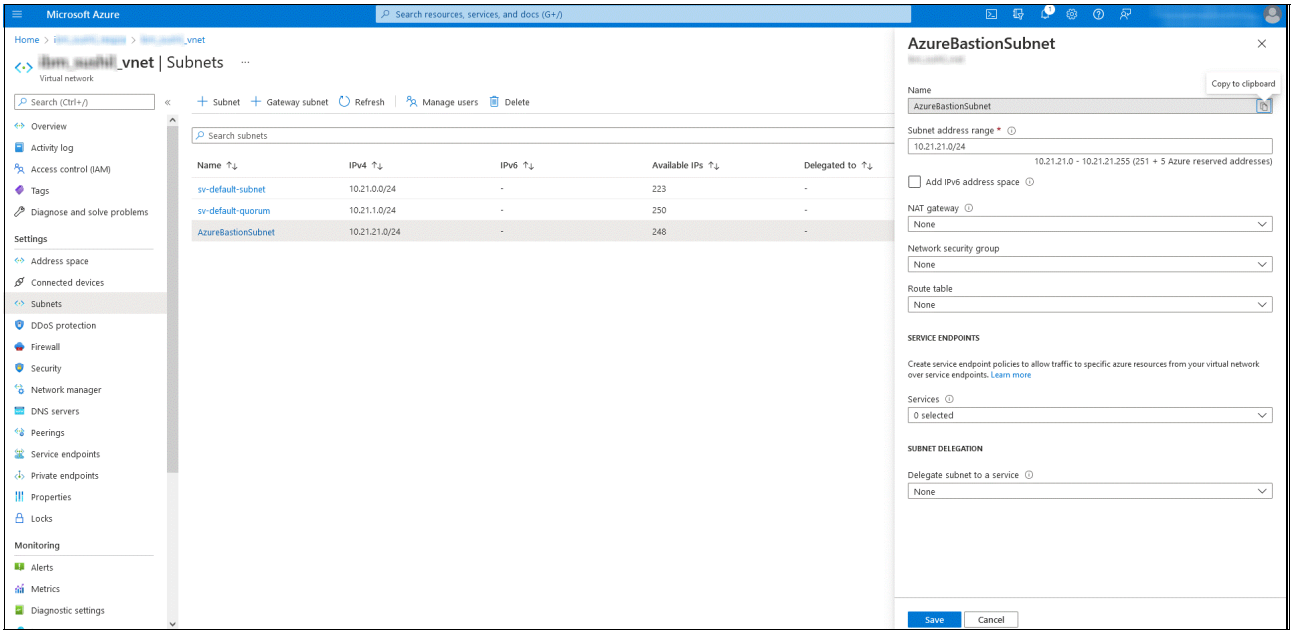

Figure 6-14 Azure Bastion Subnet in an existing virtual network

2. Create the Azure Bastion Service resource by searching for “Azure Bastion” in the **Create a resource** window on the Azure Portal (see Figure 6-15).

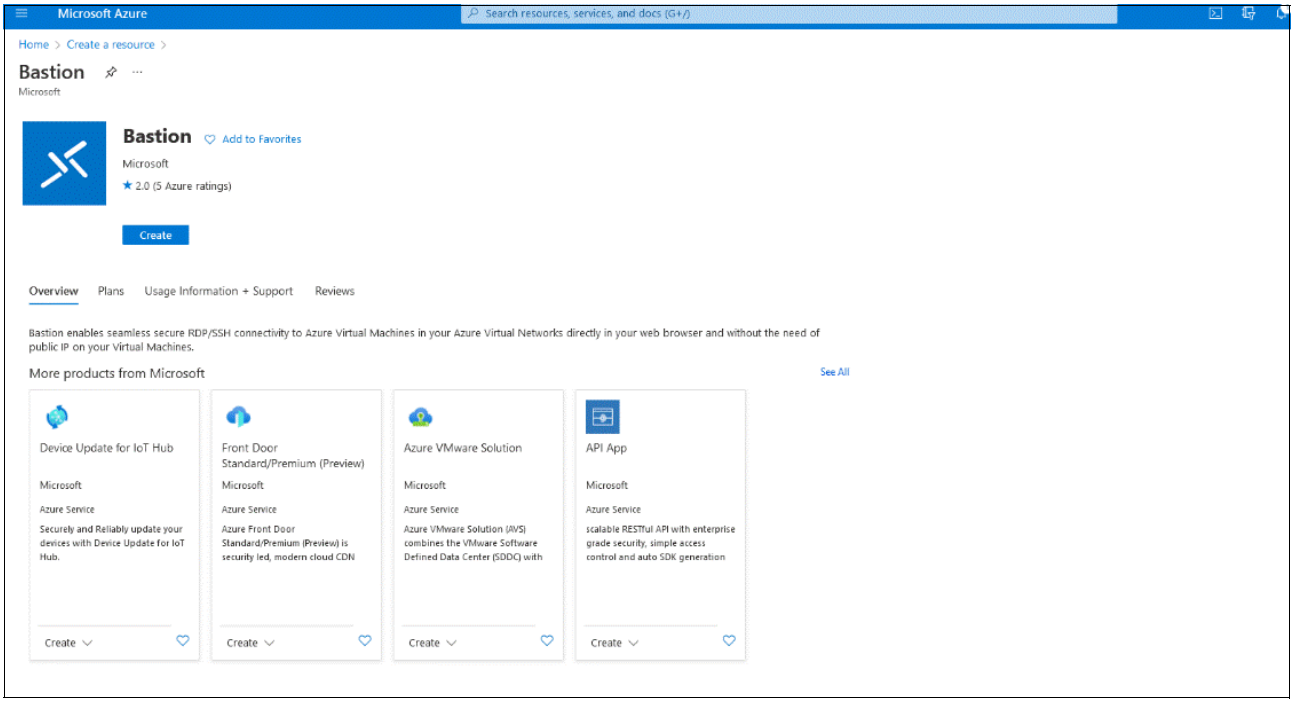

Figure 6-15 Azure Bastion Service on Azure Portal

3. Click **Create**, and then enter the values for the Azure Bastion Service.

Note: Select the vNet that is used in step 1 to create the Azure Bastion Subnet as the vNet for your Azure Bastion Service. For more information abut Azure Bastion Service, see this web page.

6.2.2 Connecting to an IBM Spectrum Virtualize for Public Cloud VM by using Azure Bastion Service

To connect the Azure VM in all-in-cloud deployment, use the Azure Bastion Service as configured in 6.2.1, “Configuring the Azure Bastion Service” on page 145.

Complete Perform the following steps to connect the VM by using the SSH key that was used during the deployment:

1. Log in to Azure Portal.
2. Find the VM that is to be connected under the newly deployed RG, as shown in Figure 6-16.

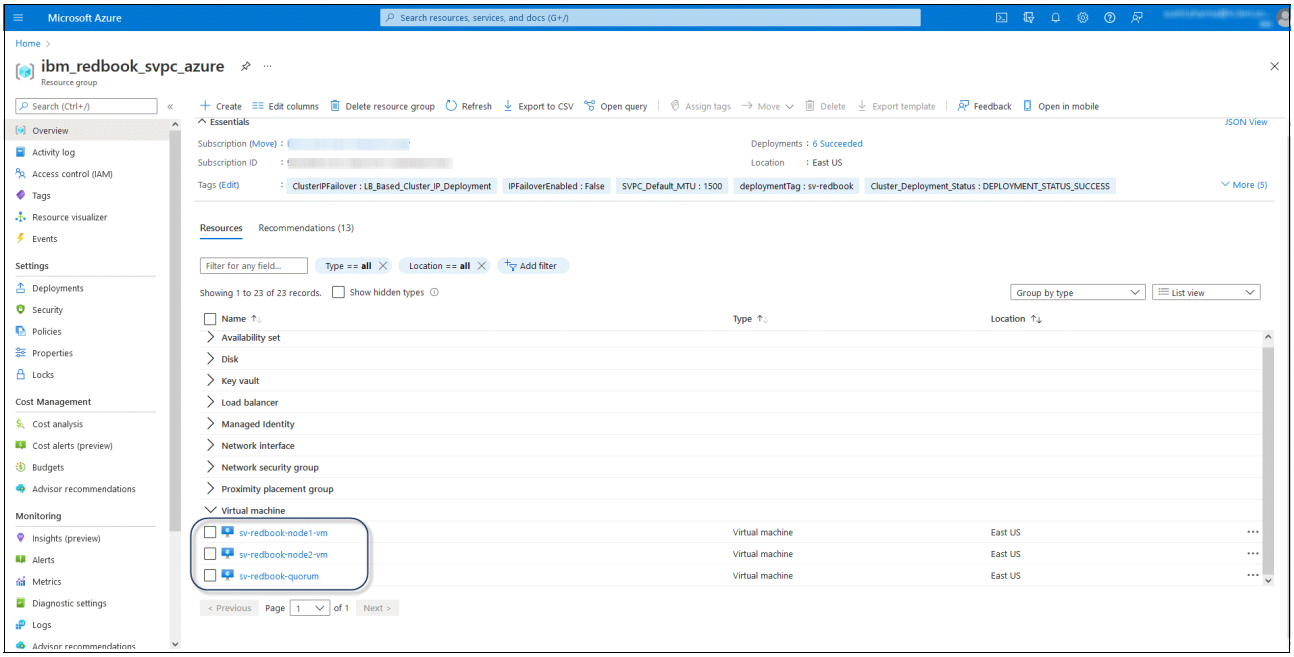

Figure 6-16 VMs in a resource group

3. Go to the Quorum VM to connect and select the **Connect** option, as shown in Figure 6-17.

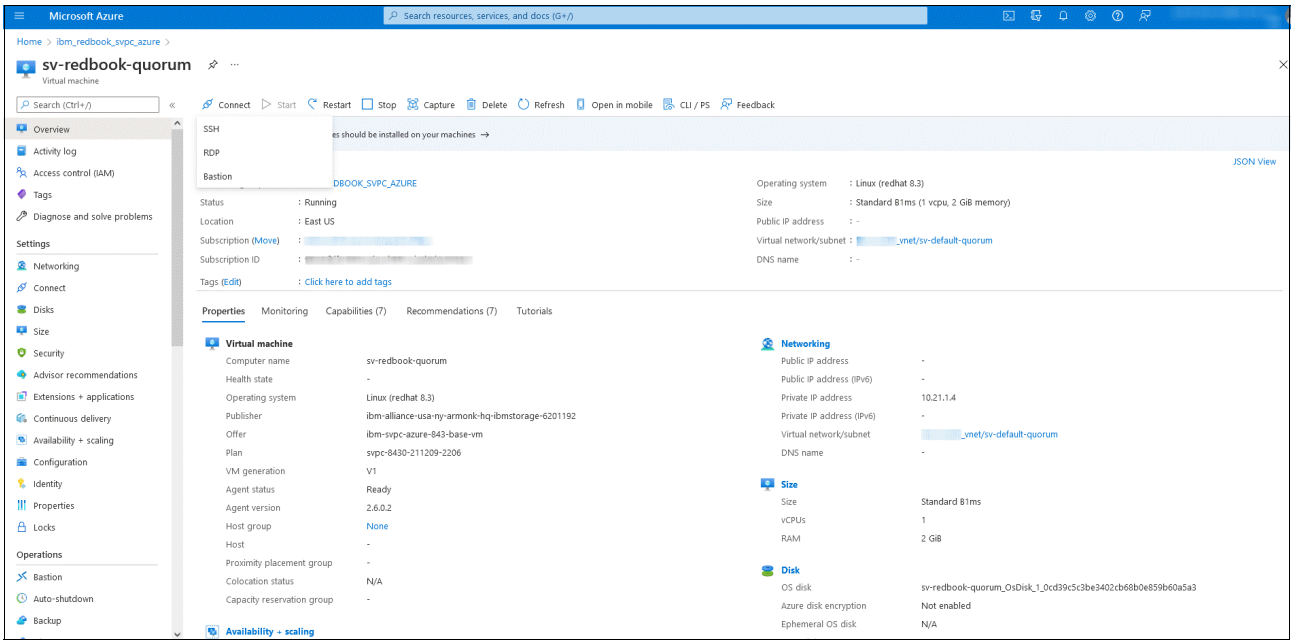

Figure 6-17 Virtual machine Connect option

4. Select the **Azure Bastion** option to connect and then, select the configured Azure Bastion service. The Connect page in the Azure portal features options that are used to provide the connection credentials, such as username to use for logging in and the password and SSH key information. For the purposes of this document, the connection is made to quorum node by using the SSH key file.

> **Note:** The SSH key must be from same set of public and private key pairs that were used during the IBM Spectrum Virtualize for Public Cloud deployment.

Figure 6-18 shows the Connection window on Azure portal for quorum node. Similarly, a connection can be made to the two VM nodes by using the same SSH key and username as `sv-cloud`.

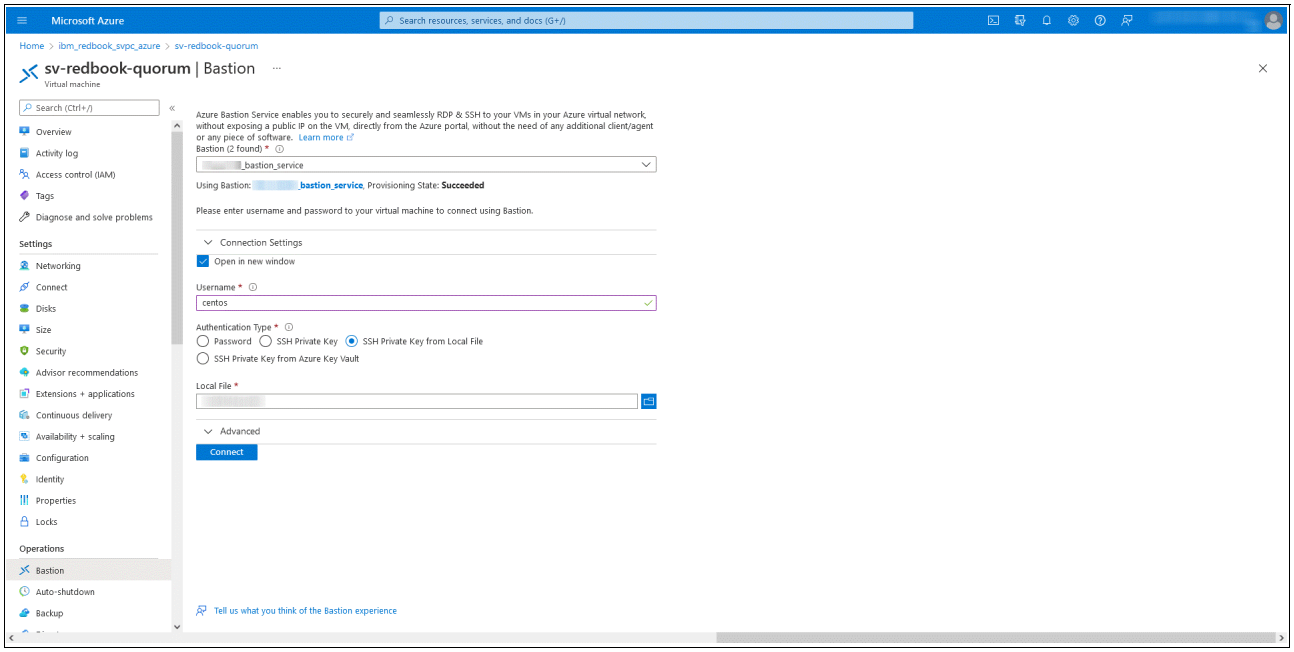

Figure 6-18 Connection page for a virtual machine that uses Azure Bastion Service

5. Click **Connect** to open the connection to the VM from the Azure portal that uses the Azure Bastion Service. A successful connection to the quorum node is shown in Figure 6-19.

> **Note:** All connections that use Azure Bastion are done by using an HTML5 web browser that is logged in to the Azure cloud account.

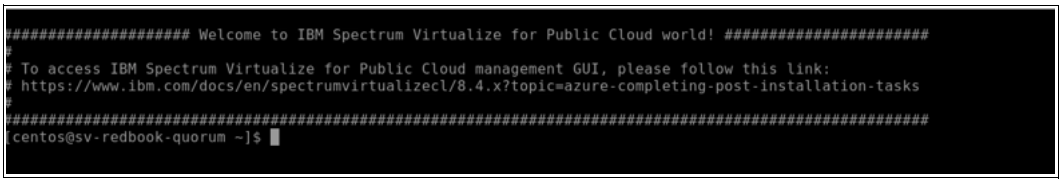

Figure 6-19 Connecting to a Quorum Node that uses Azure Bastion Service through a web browser

6.3 Finishing the setup of IBM Spectrum Virtualize for Public Cloud on Azure

After IBM Spectrum Virtualize for Public Cloud is deployed in the Azure cloud, you can log in to IBM Spectrum Virtualize for Public Cloud to set up and provision the software array for volumes, hosts, and replication. Depending on your earlier IP network choices, you might need to take a few extra steps by using the Azure Bastion service. Because this service is the only (external and optional) exposed address, it features the following functions:

- SSH jump host
- GUI proxy
- Cloud Call Home gateway
- SMTP gateway (optional)
- Remote Support Proxy (RSP) server (optional)
- Storage Insights DataCollector host (optional)

By using a management VM after IBM Spectrum Virtualize for Public Cloud deploys, you can quickly enable the IBM Spectrum Virtualize for Public Cloud Management GUI by using an Azure browser window, HTTPS, or a command-line interface (CLI) in the newly deployed IBM Spectrum Virtualize for Public Cloud.

6.3.1 Creating or using an existing cloud VM to access IBM Spectrum Virtualize for Public Cloud through Azure Bastion Service

Use the Azure console to create or use an existing VM in the same RG to access the IBM Spectrum Virtualize for Public Cloud. Selecting a Windows Server image is convenient, but any modern OS with a web browser works.

Figure 6-20 shows the Azure new VM creation window.

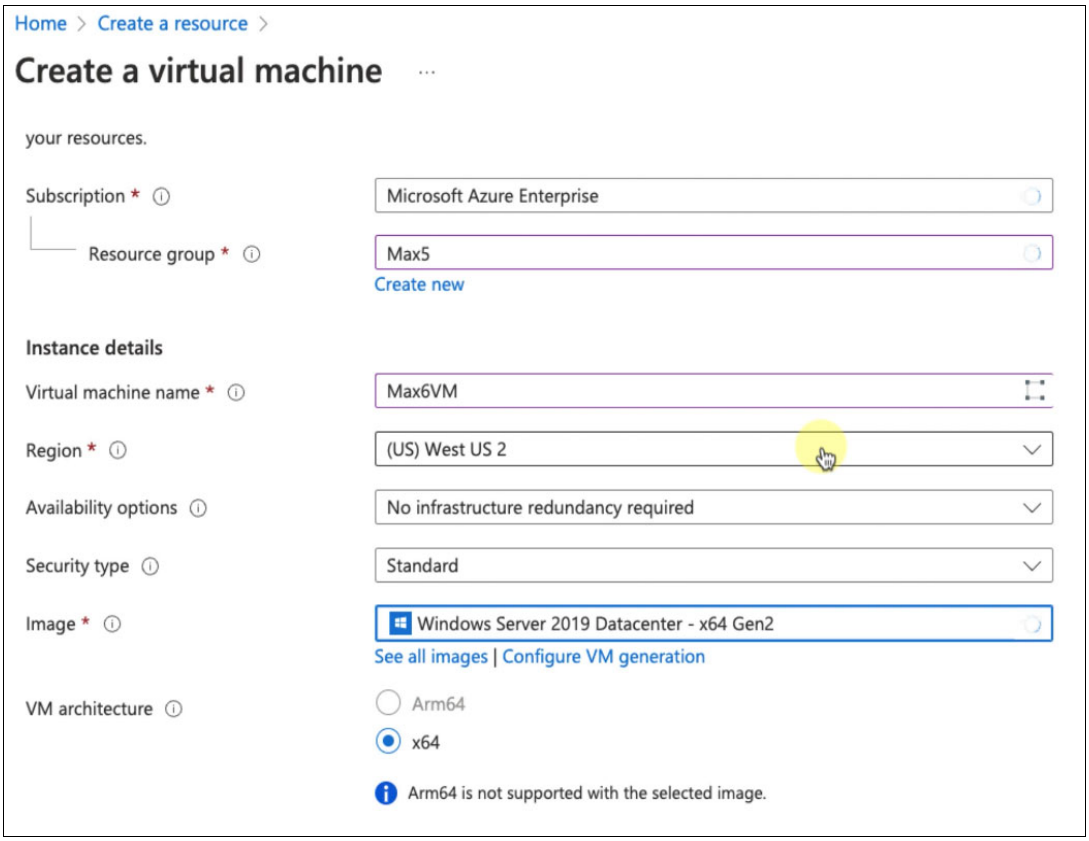

Figure 6-20 Creating a VM in the same RG to access IBM Spectrum Virtualize for Public Cloud (Azure Bastion host)

Complete the following steps:

1. Continue with the required user credentials, as shown in Figure 6-21.

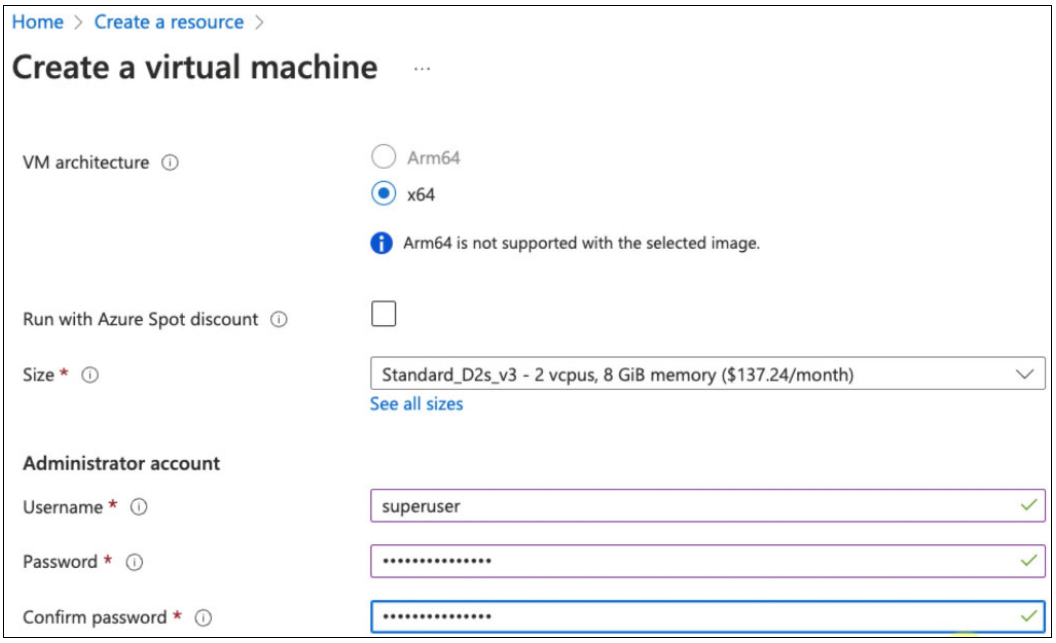

Figure 6-21 Azure virtual machine creation: credentials

2. Enable the inbound ports so that the VM may communicate with the IBM Spectrum Virtualize for Public Cloud, as shown in Figure 6-22.

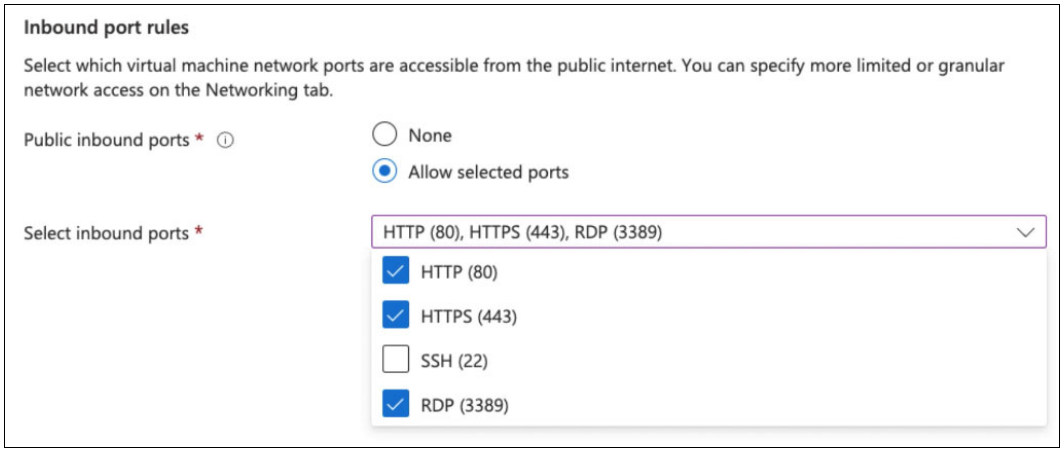

Figure 6-22 Azure VM creation: enable inbound ports for communications

3. In the **Network** tab, there are more options to restrict access so that this VM complies with your business guidelines. There are various options for the disk types. In this example, use the defaults, which are for management functions, as shown in Figure 6-23 on page 153.

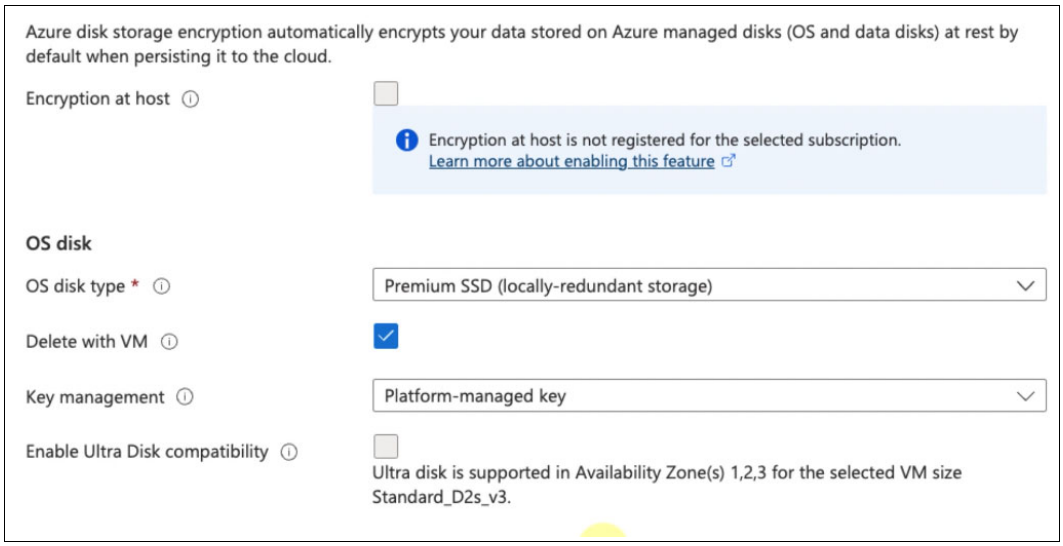

Figure 6-23 Azure VM creation: Defaults for OS disk and VM (host) encryption

4. The **Networking** tab contains key items for the VM to communicate with IBM Spectrum Virtualize for Public Cloud instance. If you create the instance inside the same RG, the defaults are sufficient. If you create the instance outside the RG, adjust for your business requirements and ensure cross-RG communication by assigning a public IP to the VM. Also, you can harden various network flows by using the “NIC network security group” field, as shown in Figure 6-24. When done, click **Next**.

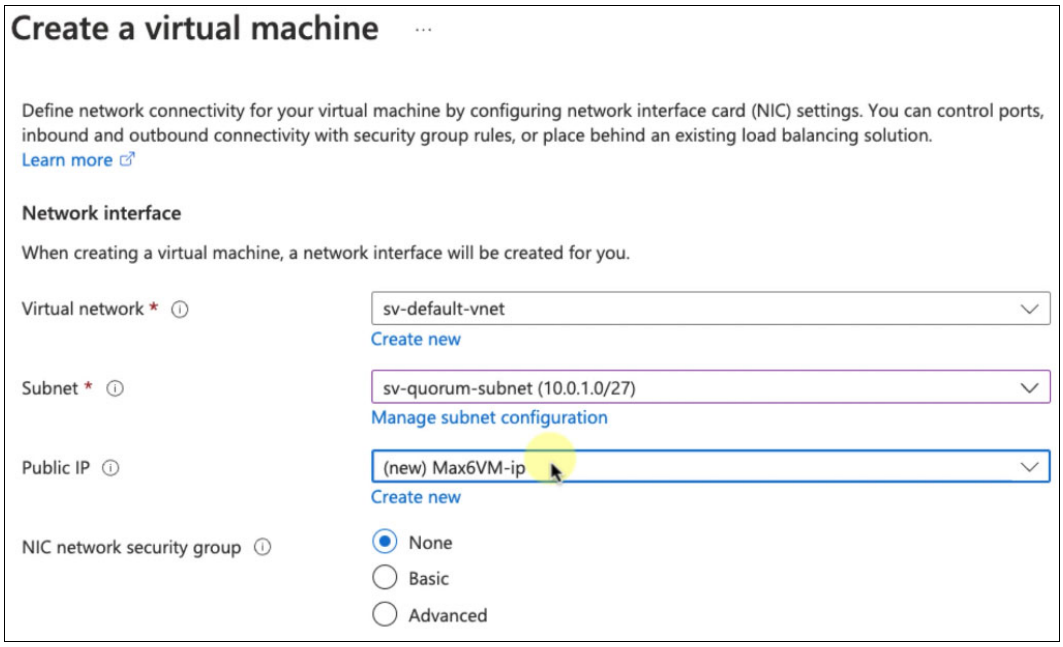

Figure 6-24 Azure VM creation: Network settings

5. There are more VM options, but none of them are essential. If you create a VM in the same RG, the extra windows are optional. If you use another existing VM or host, ensure that the VM or host has access to the same network as the IBM Spectrum Virtualize for Public Cloud Quorum VM.

The extra windows are shown in Figure 6-25 and Figure 6-26.

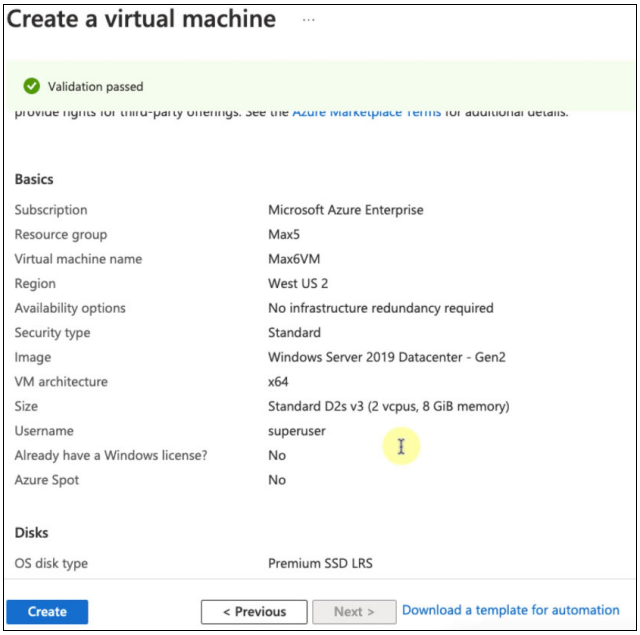

Figure 6-25 Azure VM creation: Review items before deployment

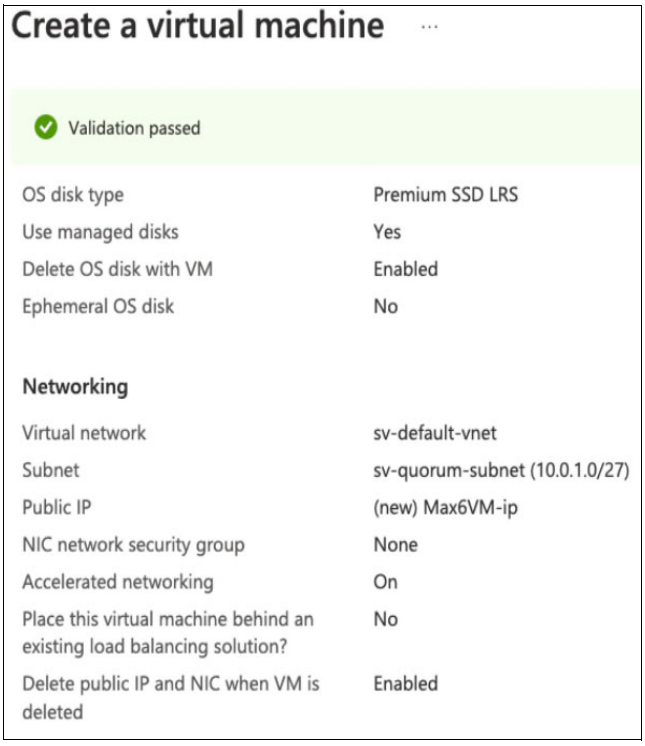

Figure 6-26 Azure VM creation: Final items before deployment

6. After deploying your VM, connect to it by using the Azure Bastion service, as shown in Figure 6-27 on page 155.

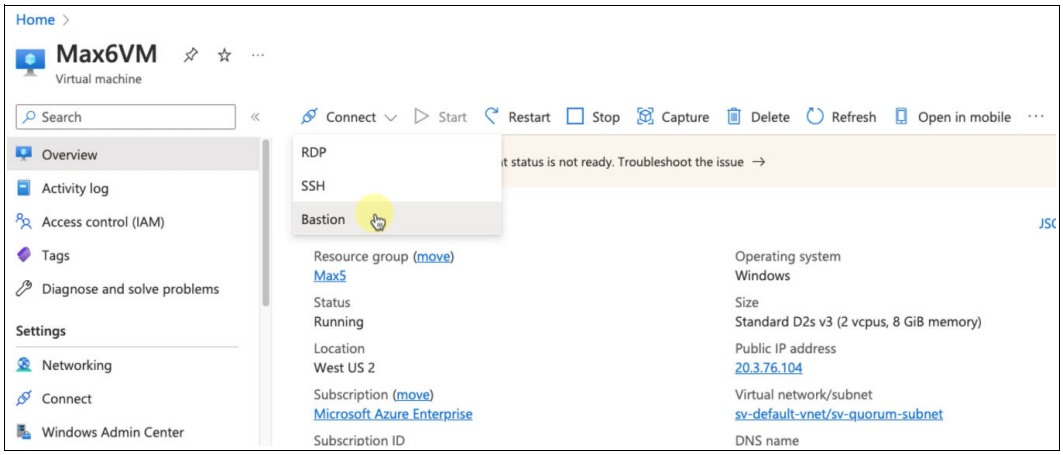

Figure 6-27 Azure connecting to the new VM through Azure Bastion

7. If you do not already have an Azure Bastion service, Azure creates one for you. Figure 6-28 shows key items to specify to ensure connectivity to IBM Spectrum Virtualize for Public Cloud.

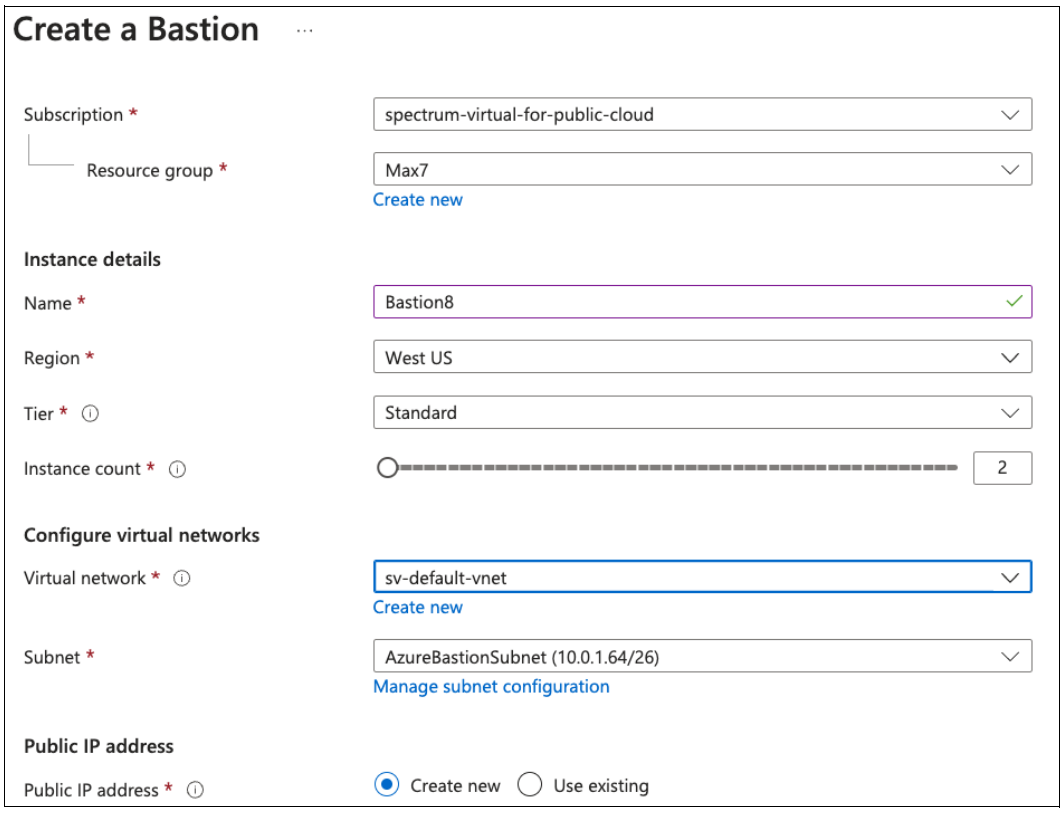

Figure 6-28 Creating an Azure Bastion Service instance for a resource group to use for IBM Spectrum Virtualize for Public Cloud Management

As shown in Figure 6-28 on page 155, your Azure Bastion instance may have a public IP address that you use for testing. For any production deployment, fewer public attack vectors are a best practice, and using your Azure account and the internal Azure Bastion method is the most secure, followed by SSH.

8. Figure 6-29 shows the features for your Bastion instance. After you select the Bastion features (as shown in the figure), click **Review + create**.

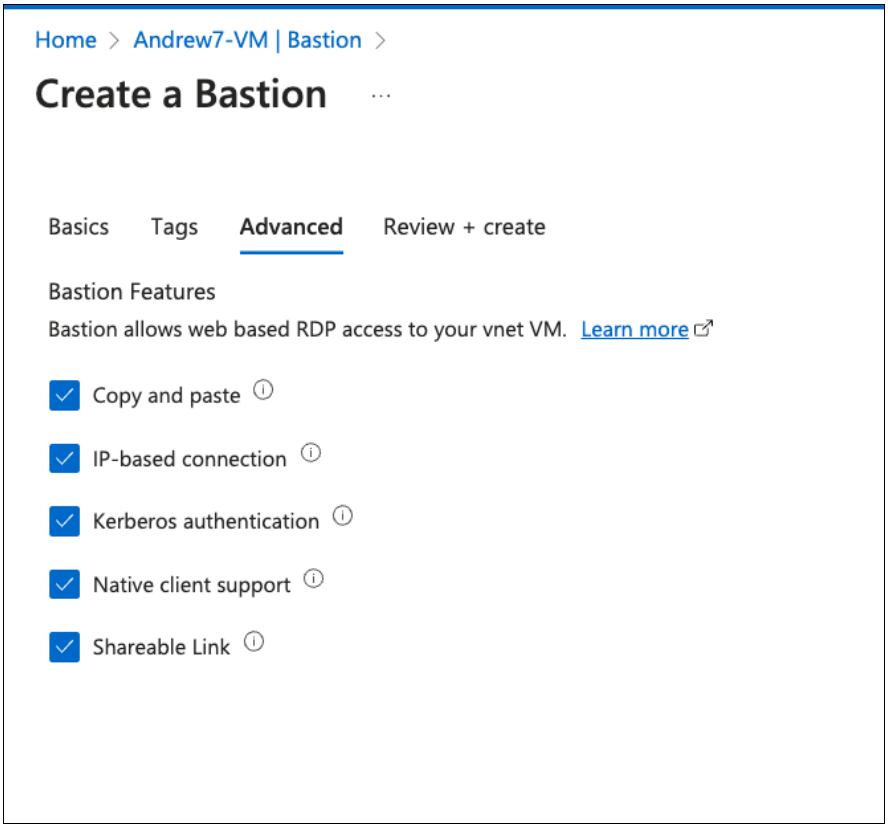

Figure 6-29 Optional but useful Azure Bastion features for Azure IBM Spectrum Virtualize for Public Cloud

9. In Figure 6-30 on page 157, click **Deploy Bastion** to deploy Azure Bastion into the same resource group as an IBM Spectrum Virtualize for Public Cloud for Management VM.

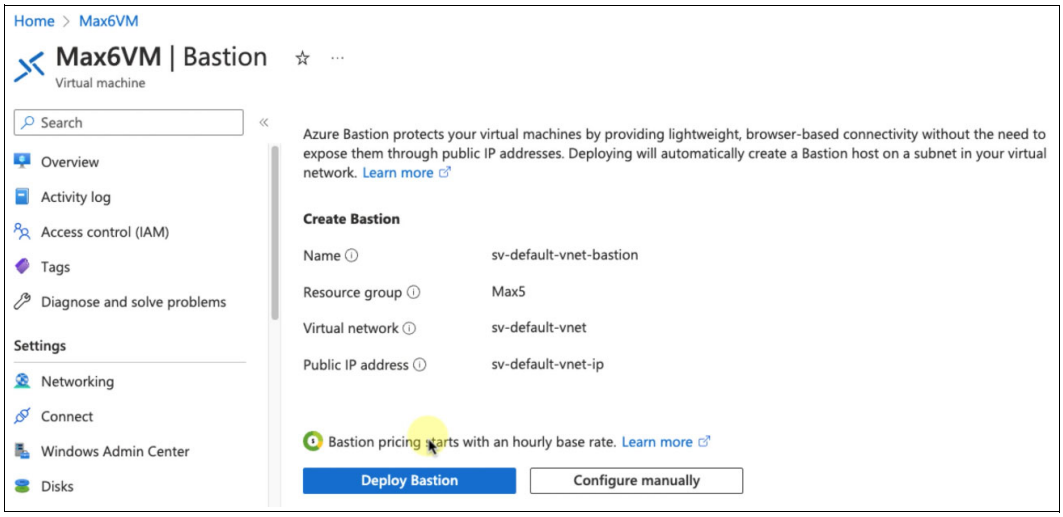

Figure 6-30 Azure deploying Azure Bastion into the same RG as an IBM Spectrum Virtualize for Public Cloud for Management VM

10. Connect to the VM by using the Azure Bastion service with the Azure Bastion service deployed in the RG, Azure presents a window in which you enter the Management credentials that you specified during the VM deployment, as shown in Figure 6-31. Ensure that your browser has not entered the wrong values.

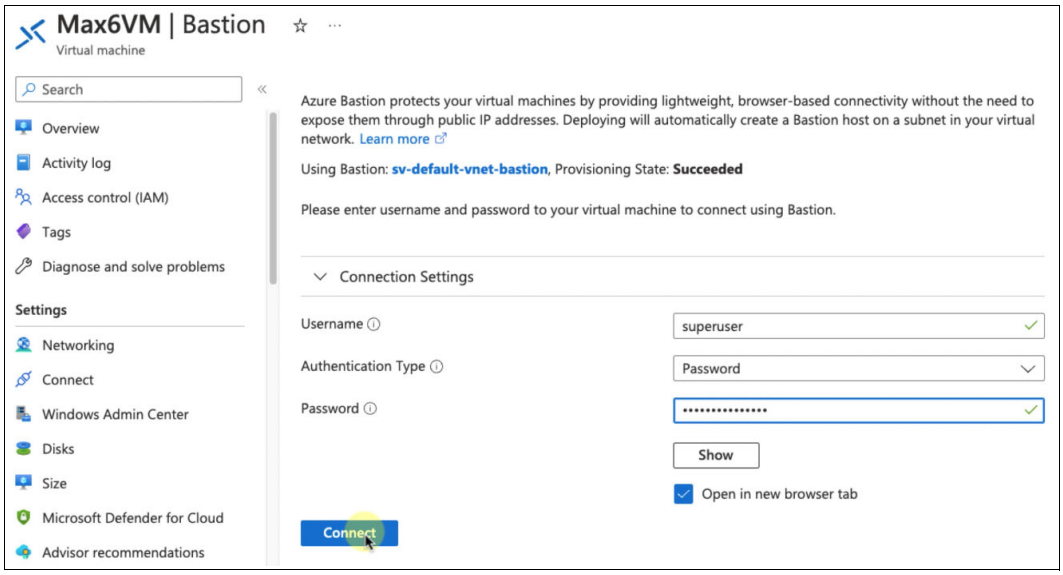

Figure 6-31 Connecting to Azure Bastion by using the specified credentials

As we are using a Windows VM for this example, you must allow its internal (OS) firewall to connect, as shown in Figure 6-32.

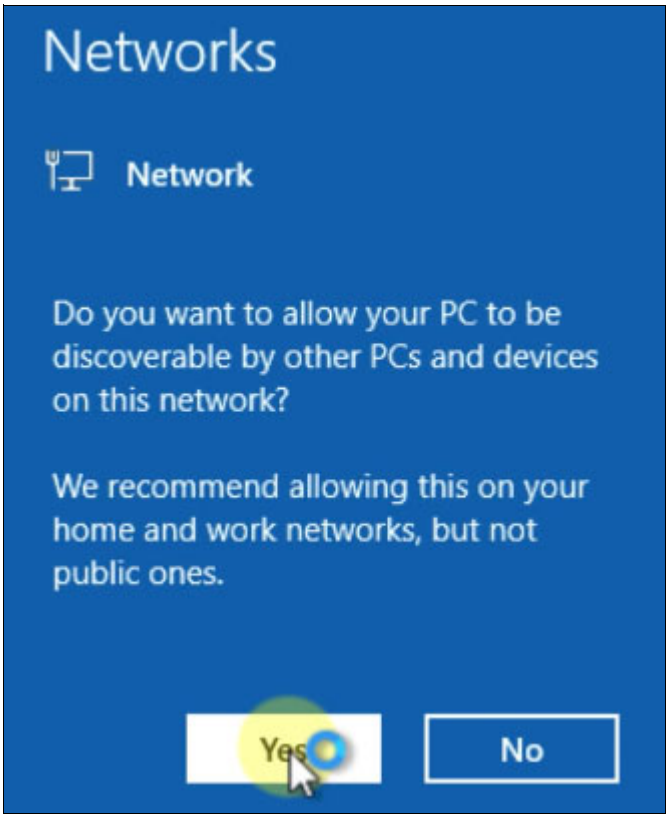

Figure 6-32 VM initial setup for OS firewall items: enable network connections

11. After you enable the VM's local web browser, use the previously supplied cluster IP address that Azure sent a copy to the email address that is specified in the IBM Spectrum Virtualize for Public Cloud deployment, and use that IP address to connect. Because the default installation uses a self-signed certificate (that can be updated later), proceed through various warnings, as shown in Figure 6-33 on page 159.

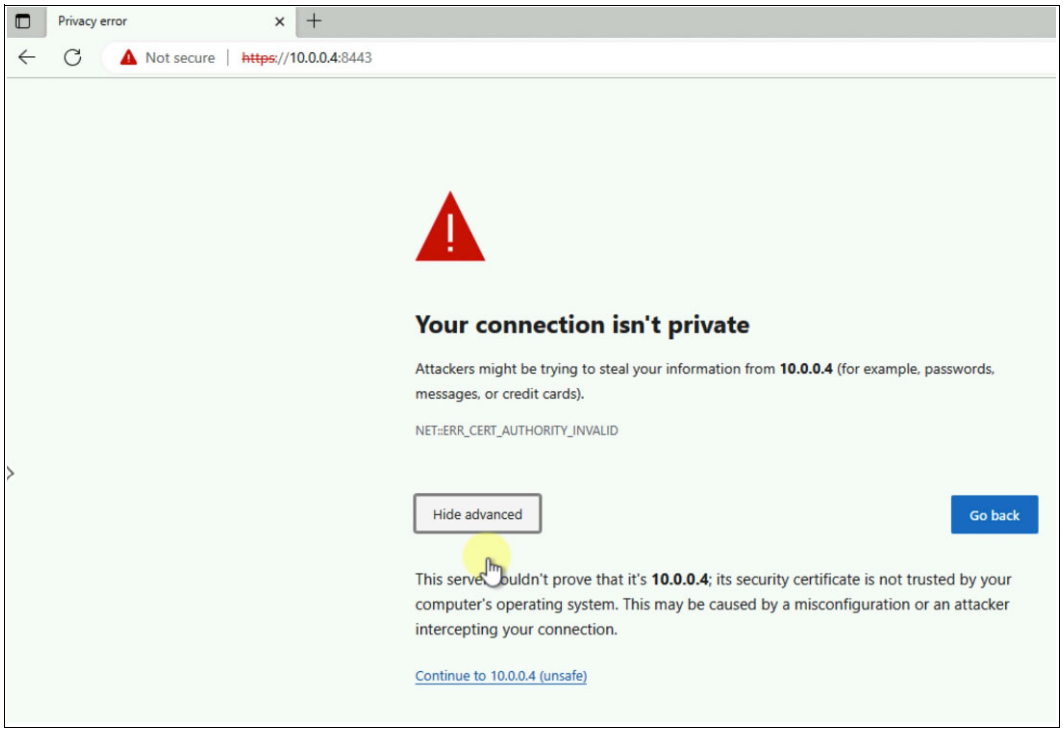

Figure 6-33 Expected browser warning when first connecting to IBM Spectrum Virtualize for Public Cloud Quorum by using a self-signed certificate

As shown in Figure 6-34, after you click to continue, you can access IBM Spectrum Virtualize for Public Cloud by using the IBM Spectrum Virtualize for Public Cloud deployment credentials that are specified. As a reminder, the URL to access the Quorum node VM management GUI is: `https://<cluster IP> :8443`.

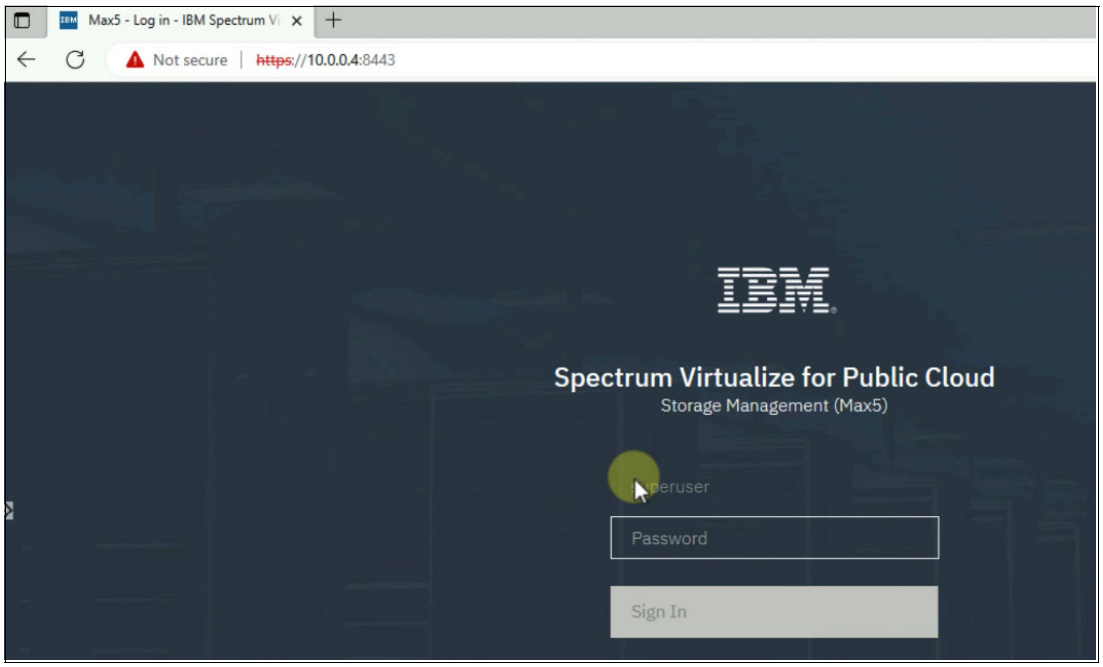

Figure 6-34 IBM Spectrum Virtualize for Public Cloud Management GUI: ready to log in with IBM Spectrum Virtualize for Public Cloud credentials

An alternative method to connect is by using a local SSH client and a previously specified KeyPair. Change the permissions on your KeyPair, and then run **ssh by** using the `centos` user, specifying KeyPair and IP address to access the Azure Bastion host. The CLI commands that you use with the `centos` user are shown in Figure 6-35.

```
andrew@IBM6-MBP AWS SV4PC % chmod 400 SV4PC-KeyPair.pem
andrew@IBM6-MBP AWS SV4PC % ssh -i SV4PC-KeyPair.pem centos@54.67.56.125
The authenticity of host '54.67.56.125 (54.67.56.125)' can't be established.
ED25519 key fingerprint is SHA256:a++neVVqjWdOcKMWMyB2ksBZt8AywAXp5oboJGb+N1c.
This key is not known by any other names
Are you sure you want to continue connecting (yes/no/[fingerprint])? yes
Warning: Permanently added '54.67.56.125' (ED25519) to the list of known hosts.

###################### Welcome to IBM Spectrum Virtualize for Public Cloud world!
 ######################
# You can enable/disable access to IBM Spectrum Virtualize for Public Cloud mana
gement GUI by:
#   1. enable GUI: enable-sv-cloud-management-gui
#   2. disable GUI: disable-sv-cloud-management-gui
# To access IBM Spectrum Virtualize for Public Cloud management GUI, enter the a
ddress in a web browser:
#   https://54.67.56.125:8443
################################################################################
########################
[centos@svpc-bastion ~]$
```

Figure 6-35 Running SSH to access the Azure Bastion host

After you are logged in to the Quorum node, a user can issue the (one-time) command to enable or disable the IBM Spectrum Virtualize for Public Cloud Management GUI.

6.3.2 Configuring IBM Spectrum Virtualize for Public Cloud EasySetup and completing the installation

Using the preconfigured VM, earlier, log in to the IBM Spectrum Virtualize for Public Cloud on Azure cluster through the WebGUI, as shown in Figure 6-36.

To complete the installation, complete the following steps:

1. Use the Cluster IP address that was provided to you by email to connect the WebGUI for IBM Spectrum Virtualize (see Figure 6-36).

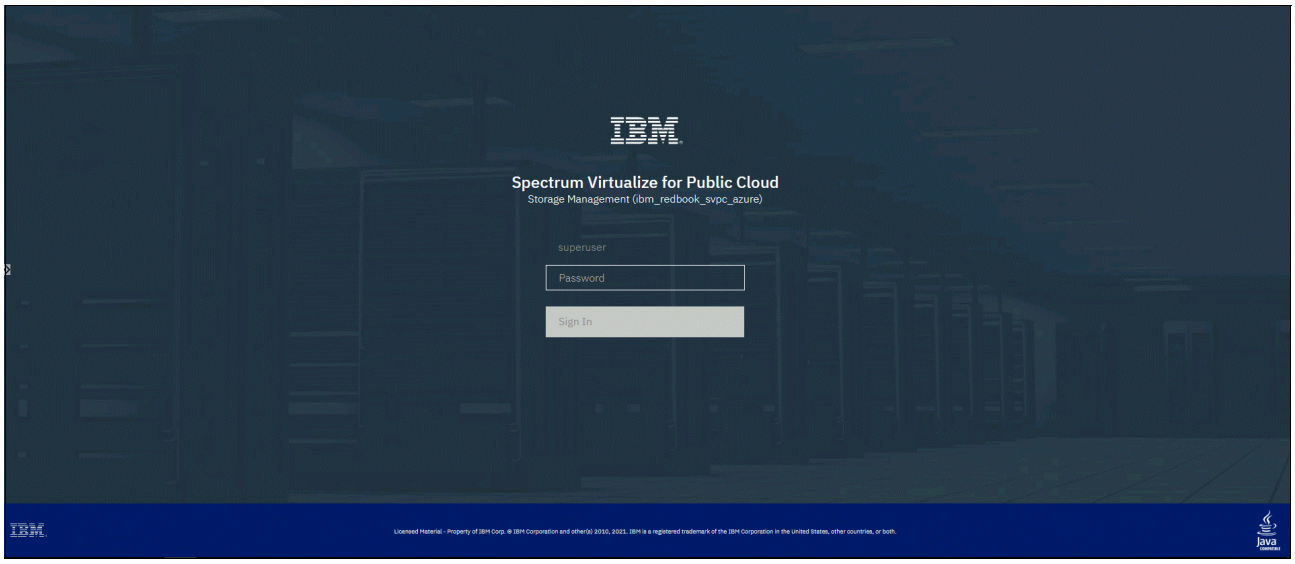

Figure 6-36 WebGUI interface for IBM Spectrum Virtualize for Public Cloud on Azure

2. You are directed to the IBM Spectrum Virtualize for Public Cloud setup window. Click **Next** (see Figure 6-37).

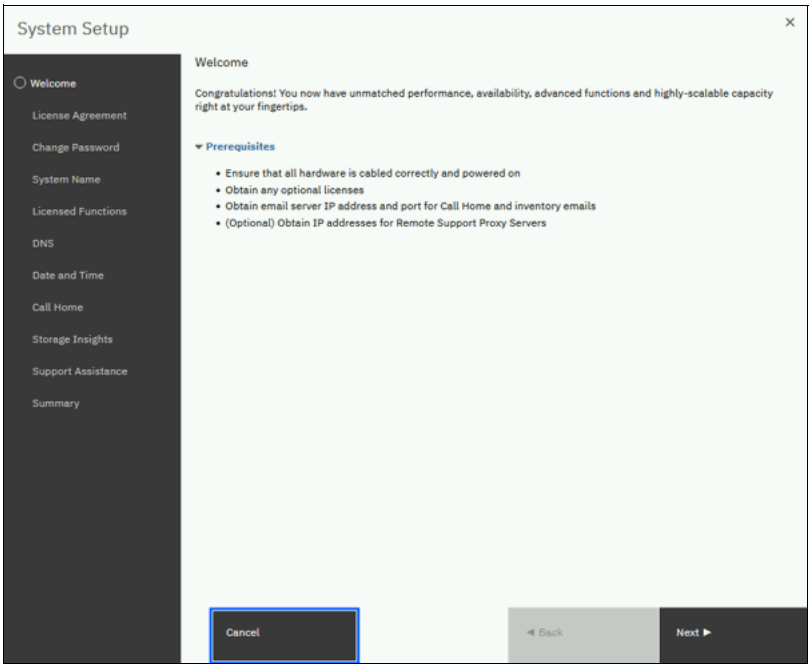

Figure 6-37 Welcome window

3. Accept the License Agreement. Read the license agreement and select **I agree** and then click **Next**, as shown in Figure 6-38.

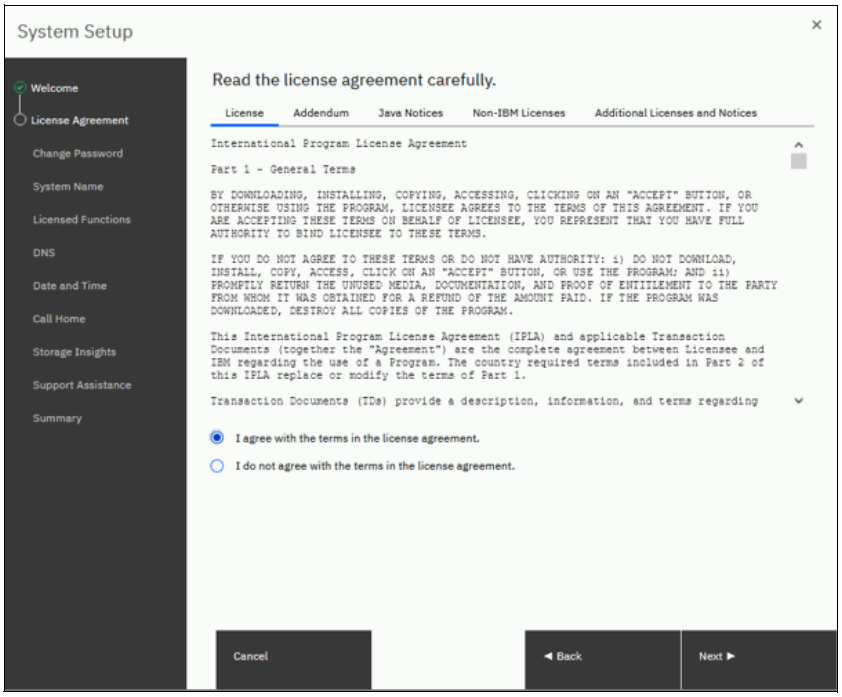

Figure 6-38 License Agreement window

4. You are directed to the Change Password window, as shown in Figure 6-39. Change your IBM Spectrum Virtualize for Public Cloud `superuser` password, and then click **Apply** and **Next**.

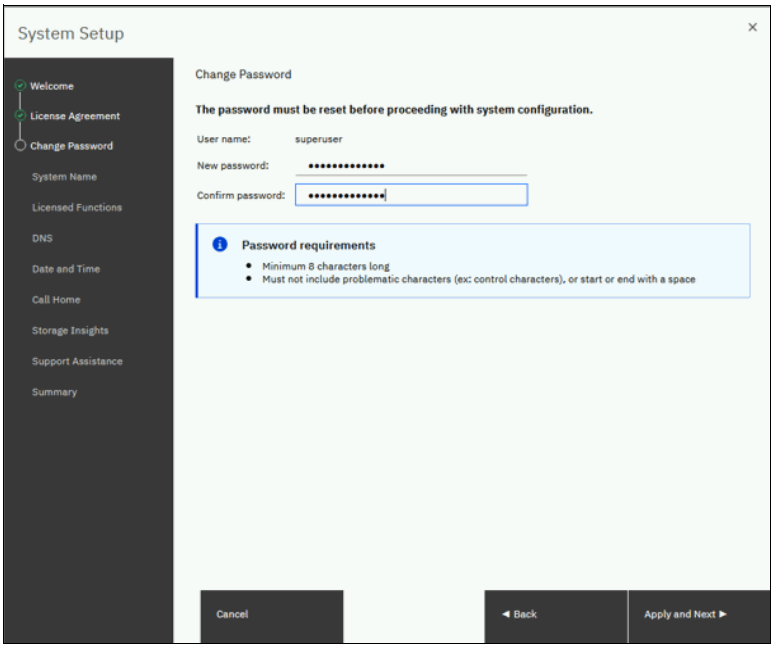

Figure 6-39 IBM Spectrum Virtualize for Public Cloud: Superuser change password window

5. Change your IBM Spectrum Virtualize for Public Cloud cluster name (if needed), which defaults to the RG name, as shown in Figure 6-40. Click **Apply** and then **Next**.

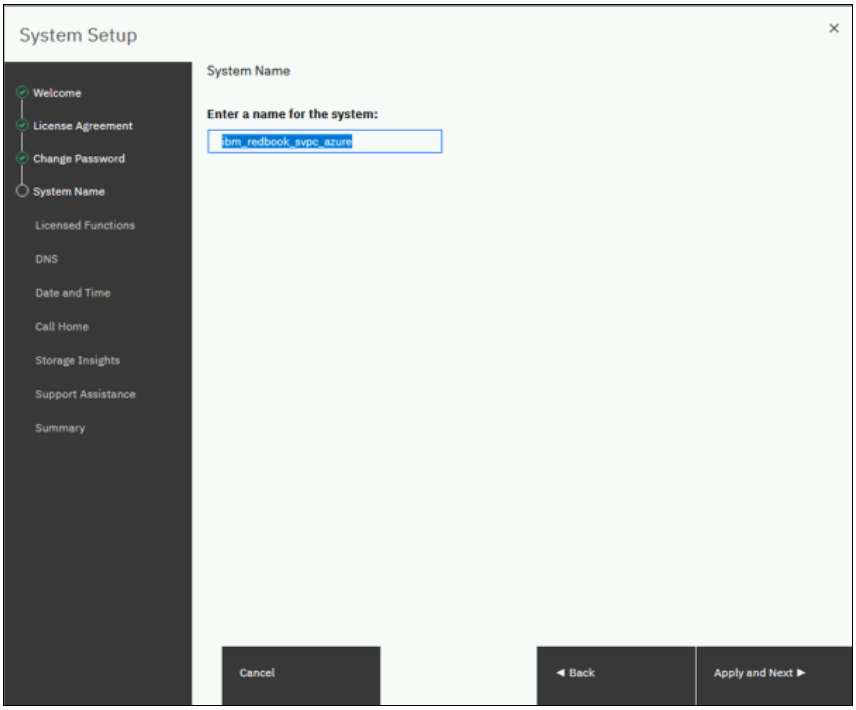

Figure 6-40 Change Cluster Name window

6. Enter your capacity license in accordance with your IBM agreement, as shown in Figure 6-41. Click **Apply** and **Next**.

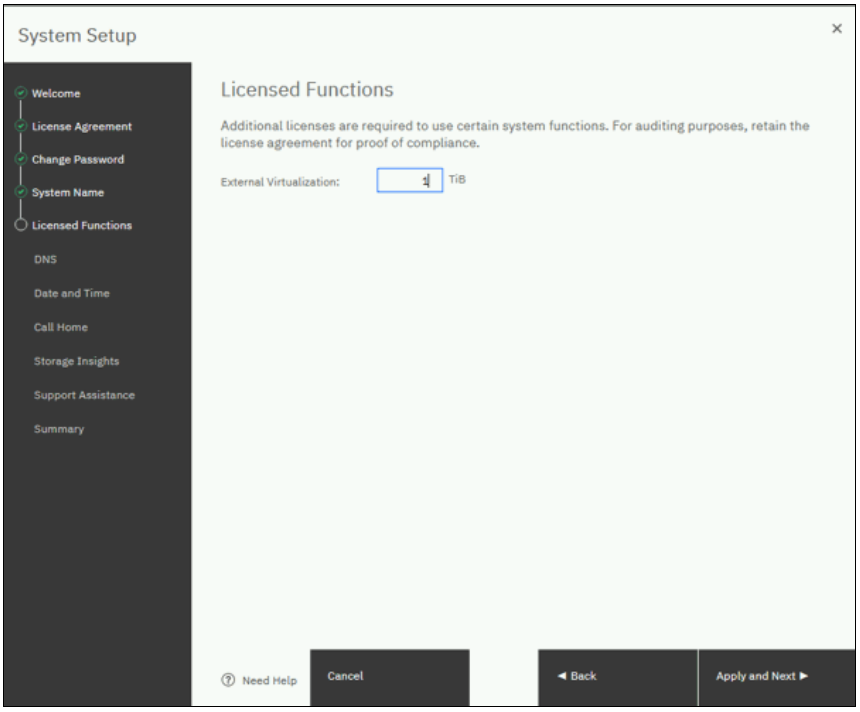

Figure 6-41 Capacity License window

7. Enter (optional) DNS servers to manage any resources that are on an external network, as shown in Figure 6-42. Because we are using only internal resources in our example, skip this step and click **Next**.

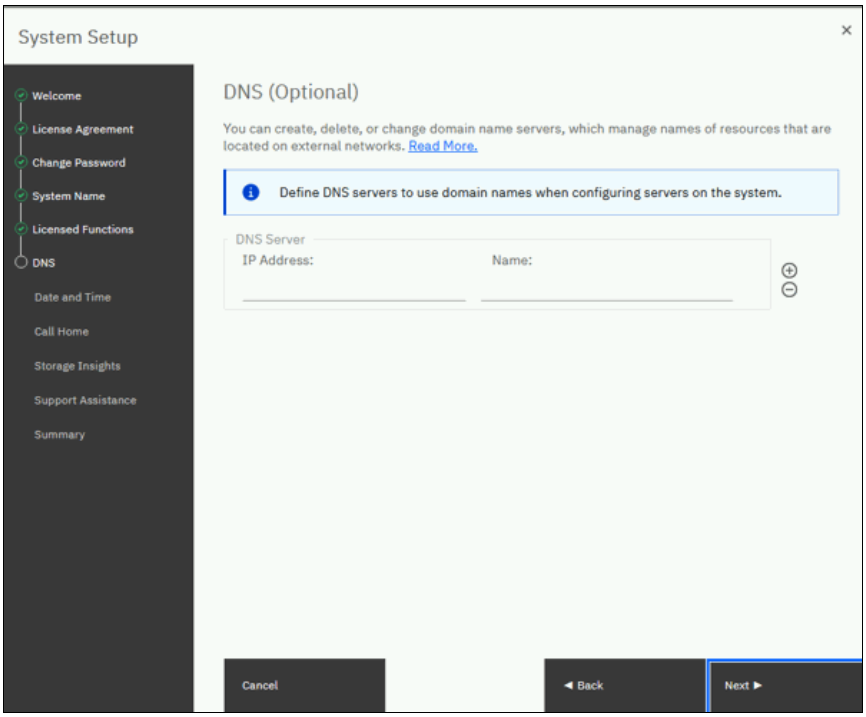

Figure 6-42 DNS Setting window

8. You do not need to set the date and time because this information is controlled by Azure. IBM Spectrum Virtualize for Public Cloud on Azure is configured by the Azure time server by using underlying operating system methods.

> **Important:** Changing the time server or setting a static time is *not* recommended because it can cause errors.

For more information about the Azure time sync, see *Time sync for Linux VMs in Azure - Azure Virtual Machines*.

9. Ensure that the time zone is set. For troubleshooting across multiple time zones, it is a best practice to use Greenwich Mean Time or Coordinated Universal Time+0, as shown in Figure 6-43.

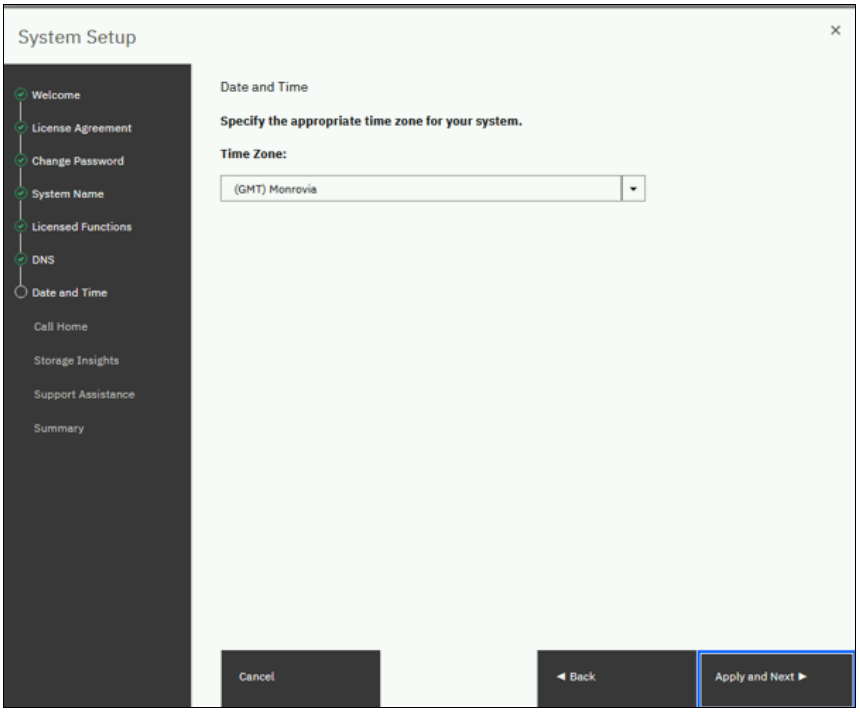

Figure 6-43 Time zone setting

10. IBM Spectrum Virtualize for Public Cloud on Azure is preconfigured with Cloud Call Home because Azure VMs can send data to IBM Support Call Home servers. When the EasySetup process enters the Call Home configuration, Cloud Call Home verifies the connection to the support center by using Cloud Service and no proxy, as shown in Figure 6-44.

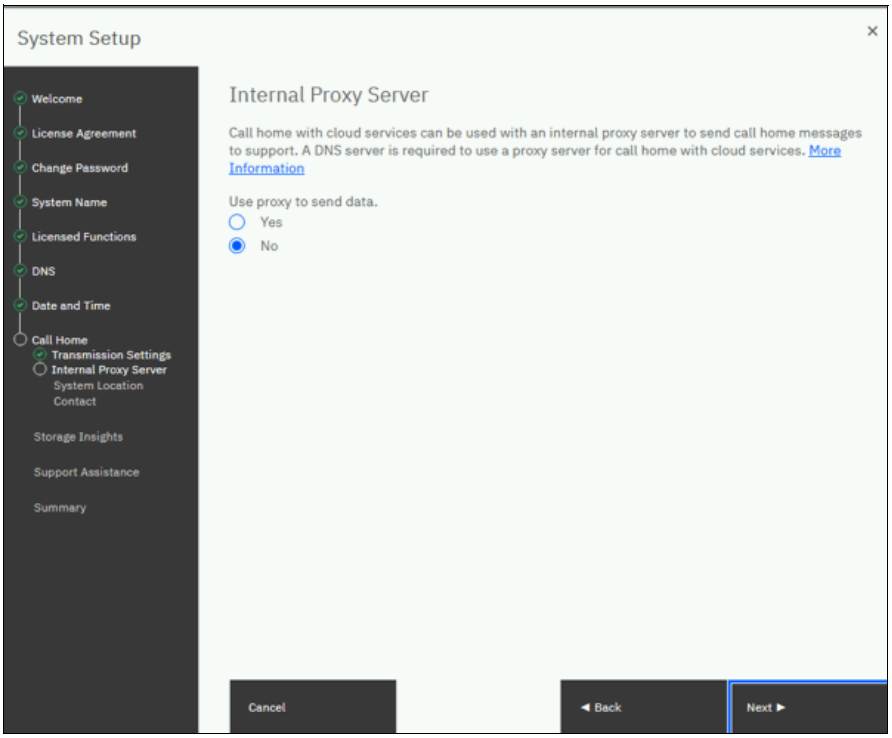

Figure 6-44 Call Home Setting by using Cloud Services

This verification should succeed, as shown in Figure 6-45, which is the System Location window.

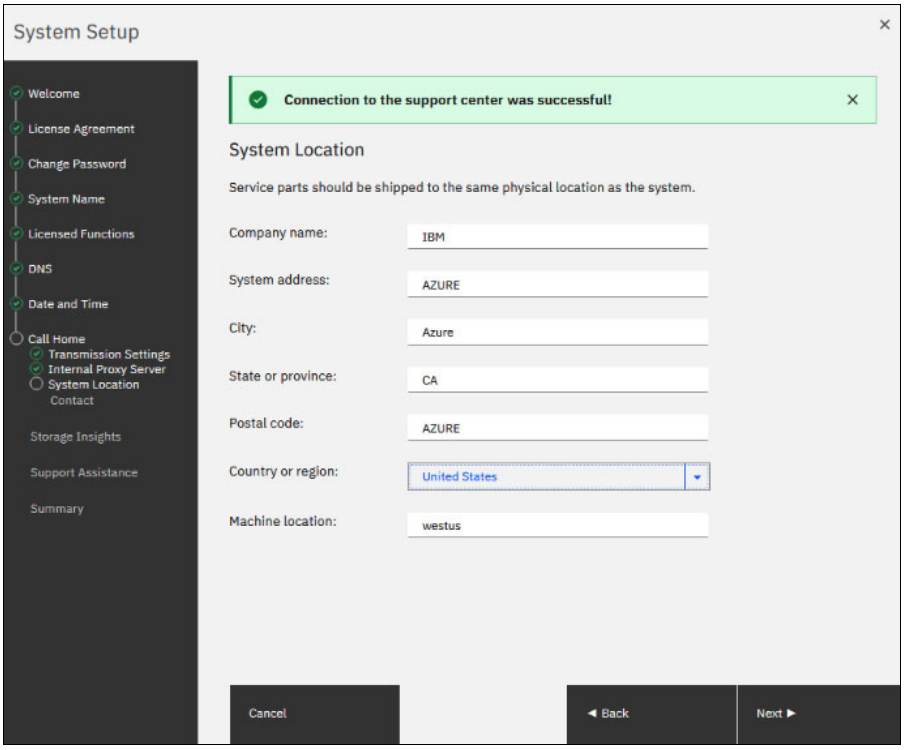

Figure 6-45 Successful verification of Call Home

11. Complete the Call Home configuration by entering the contact information, as shown in Figure 6-46.

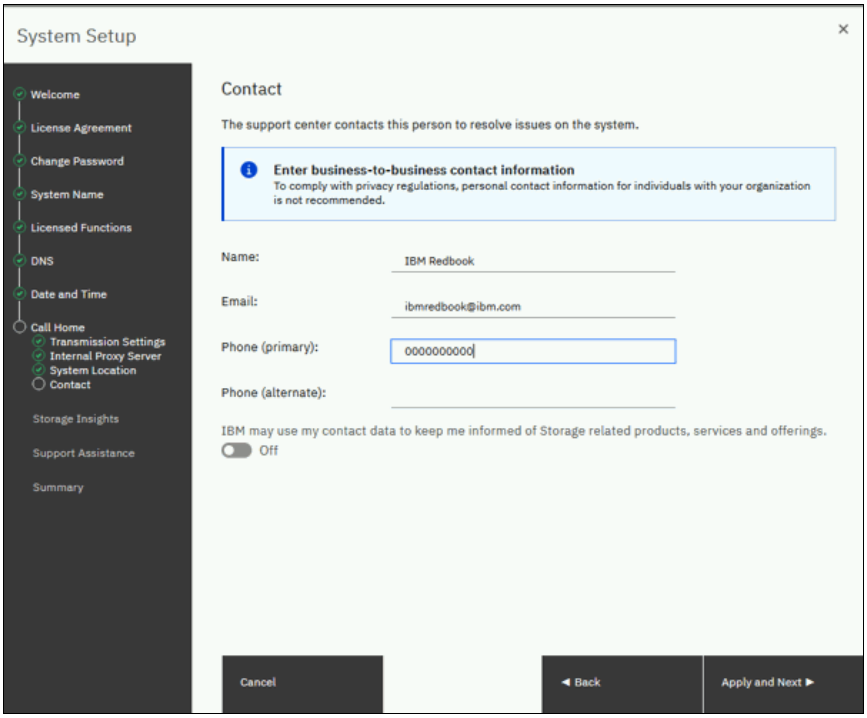

Figure 6-46 Call Home contact information

12. IBM Storage Insights is configured next. To do so, you must register a no-charge account, as shown in Figure 6-47.

    You can register by using your existing IBMid, or skip this step. Click **Next**.

> **Note:** For more information about IBM Storage Insights, see “Capacity monitoring in IBM Spectrum Control and IBM Storage Insights” on page 237.

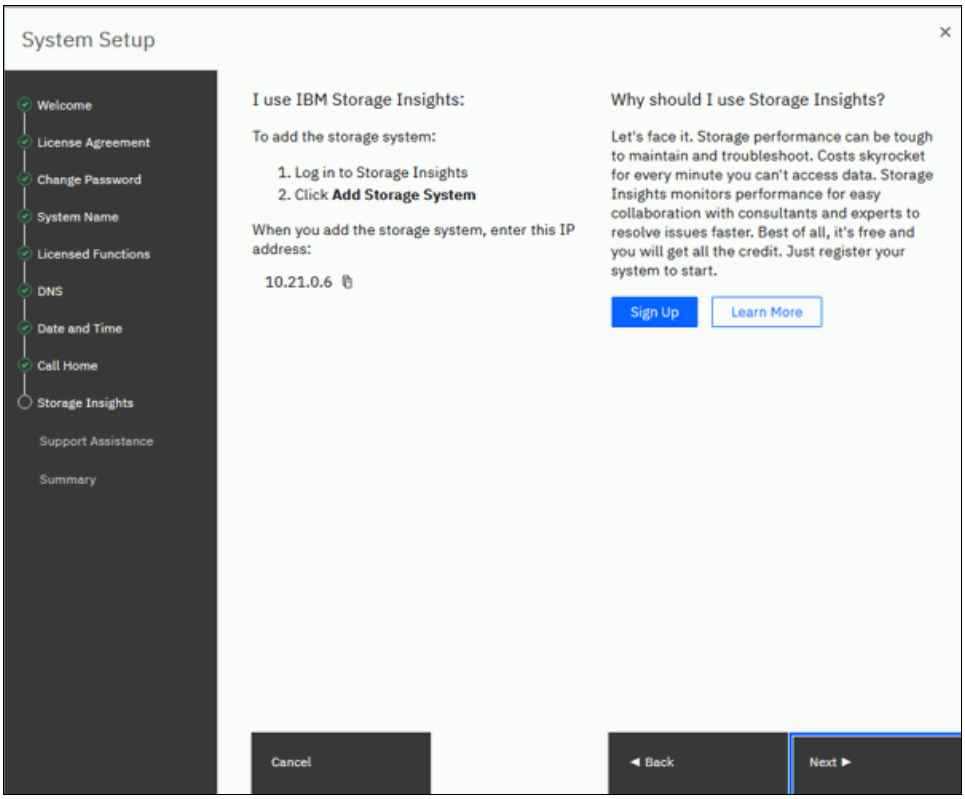

Figure 6-47 Storage Insights window

13. You can (optionally) configure an RSP, as shown in Figure 6-48. For our example, we skip this step.

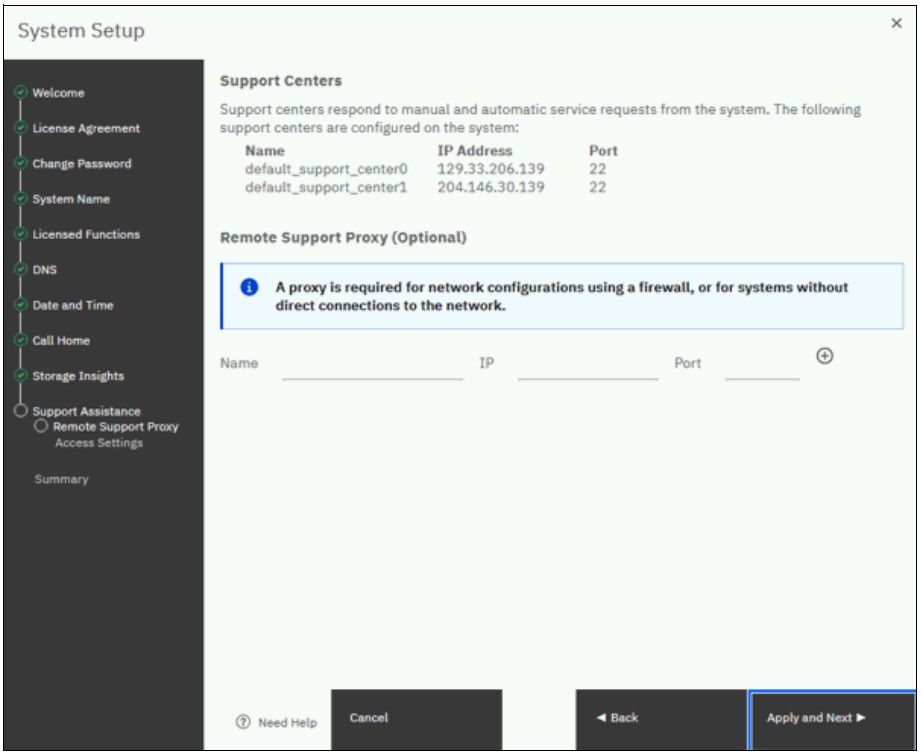

Figure 6-48 Remote Support Assistance Config window

14. Figure 6-49 shows a summary of your cluster configuration. After careful review, click **Finish** to complete the setup.

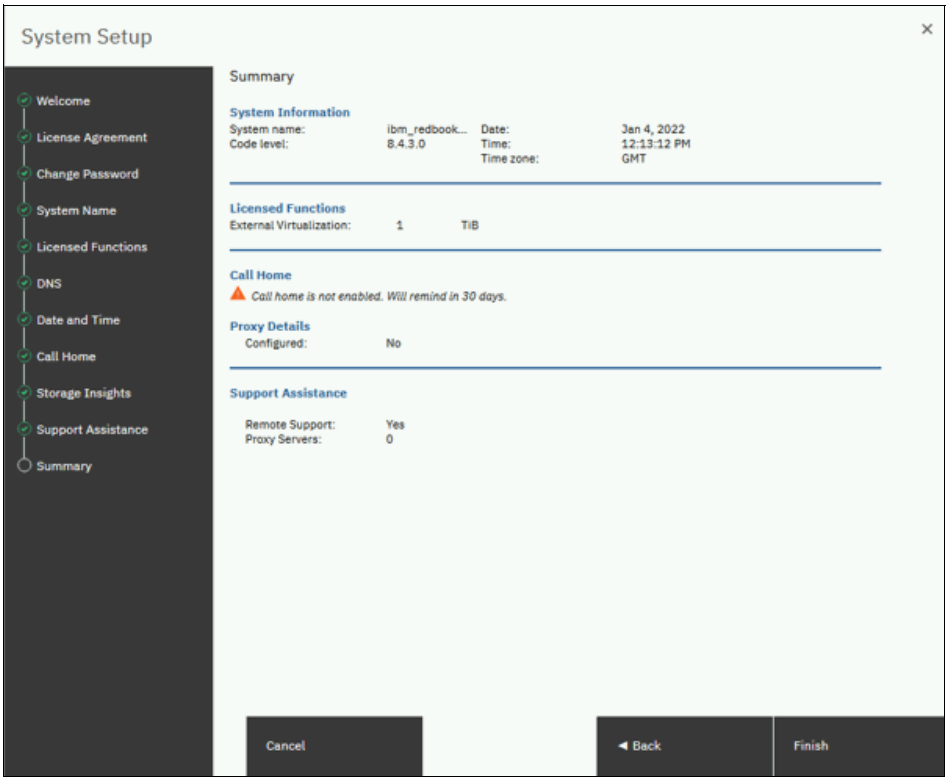

Figure 6-49 Summary window

6.4 Configuring the IBM Spectrum Virtualize for Public Cloud Azure cloud quorum

IP quorum applications are used in Ethernet networks to resolve failure scenarios when half of the nodes on the system become unavailable. These applications determine which nodes can continue processing host operations and avoids a split-brain scenario where both halves attempt to service independently I/O, which causes corruption.

As part of the installation of IBM Spectrum Virtualize for Public Cloud on Azure, a quorum host is provisioned, and the IP quorum application is installed and configured on this instance. This quorum host operates as the IP quorum for the configuration.

Note: An IP quorum is configured during the installation. You configure an extra IP quorum only if you want to enhance the fault tolerance by putting the active quorum in a different Availability Zone for installations into new vNets.

Strict requirements exist on the IP network for the use of IP quorum applications. All IP quorum applications must be reconfigured and redeployed to hosts when specific aspects of the system configuration change. These aspects include adding or removing a node from the system or when node service IP addresses are changed.

Other examples include changing the system certificate or experiencing an Ethernet connectivity issue. Such an 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 Figure 6-50.

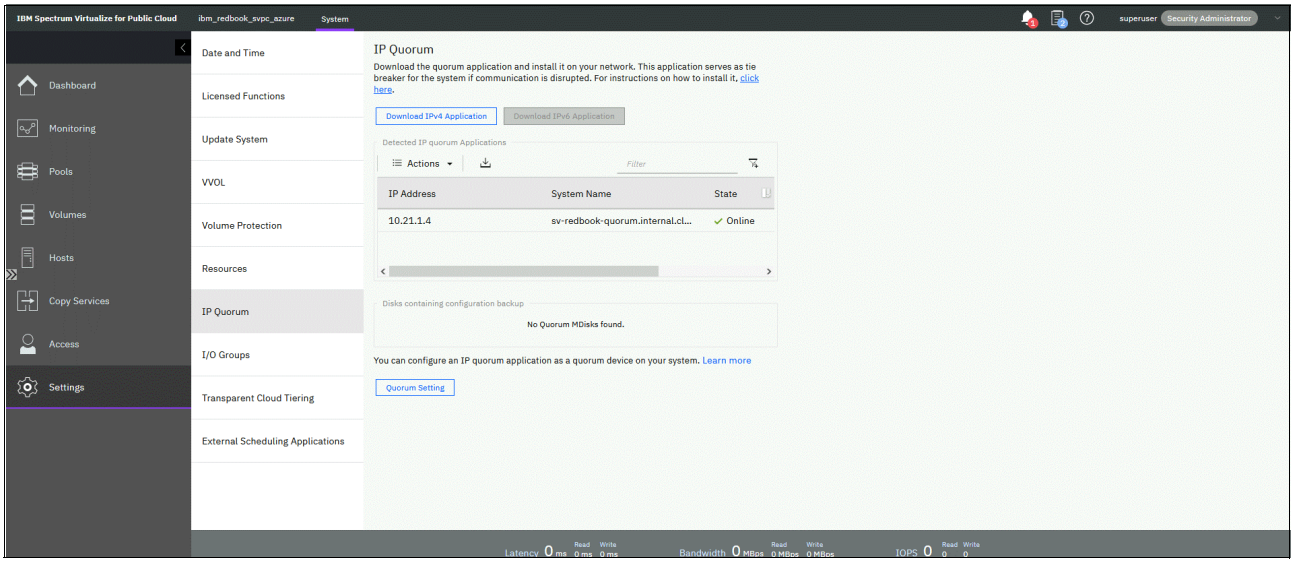

Figure 6-50 IP Quorum Config page

Even with IP quorum applications on an Azure VM instance, quorum disks are required on each node in the system to contain backups of the configuration and recovery information. On Azure VM instances where IBM Spectrum Virtualize connectivity with its nontraditional back-end storage connectivity, the quorum disks cannot be on external storage or internal disk as in IBM SAN Controller Volume or IBM FlashSystem systems. Therefore, they are automatically allocated on the VM instance boot device for each IBM Spectrum Virtualize node.

The IBM Spectrum Virtualize command **`lsquorum`** shows only the IP quorum. A maximum of 5 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:

- The servers that are running an IP quorum application are connected to the service IP addresses of all nodes.
- The network manages the possible security implications of exposing the service IP addresses because this connectivity also can 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 does not exceed 80 milliseconds (ms), which means 40 ms in each direction.
- A minimum bandwidth of 2 MBps is ensured for node-to-quorum traffic.

6.5 Configuring the back-end storage

IBM Spectrum Virtualize for Public Cloud on Azure uses the back-end storage that is provided by Azure Managed Disk (MDisk) as external Virtualize MDisks. As part of the initial default installation, two Azure MDisks are allocated to the IBM Spectrum Virtualize cluster (see Figure 6-51).

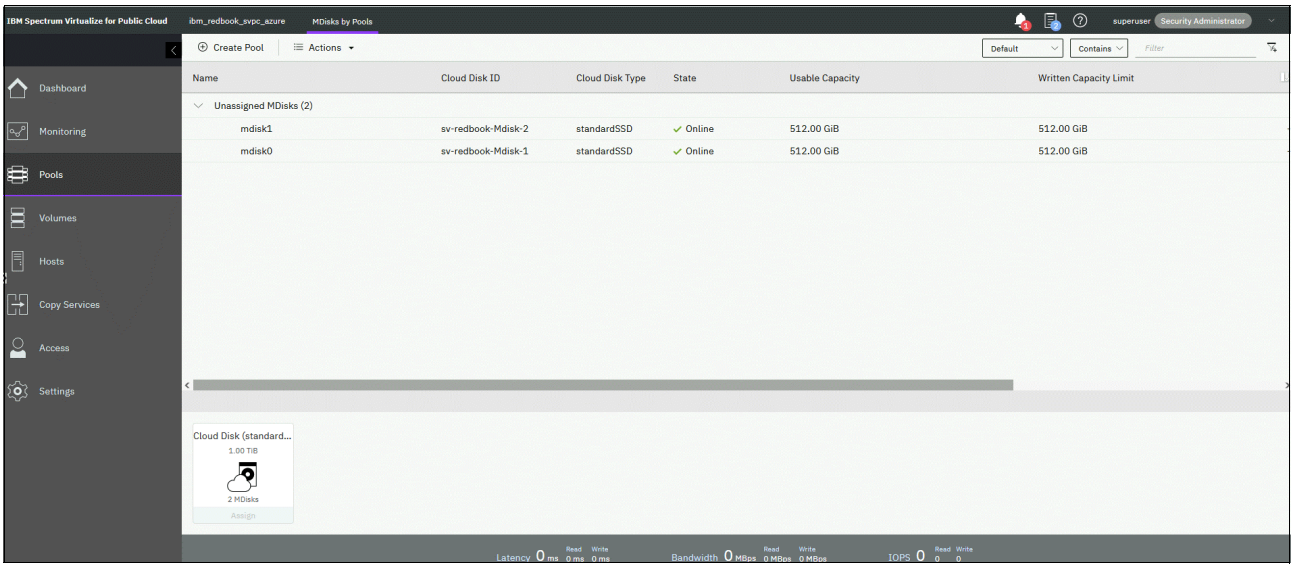

Figure 6-51 Default storage added as part of the initial deployment

To create a pool from the Azure MDisks, complete the following steps:

1. Click **CreatePool**. Enter information for the new pool, such as the Pool Name, Extent Size, and Data Reduction Capabilities for the pool, as shown in Figure 6-52.

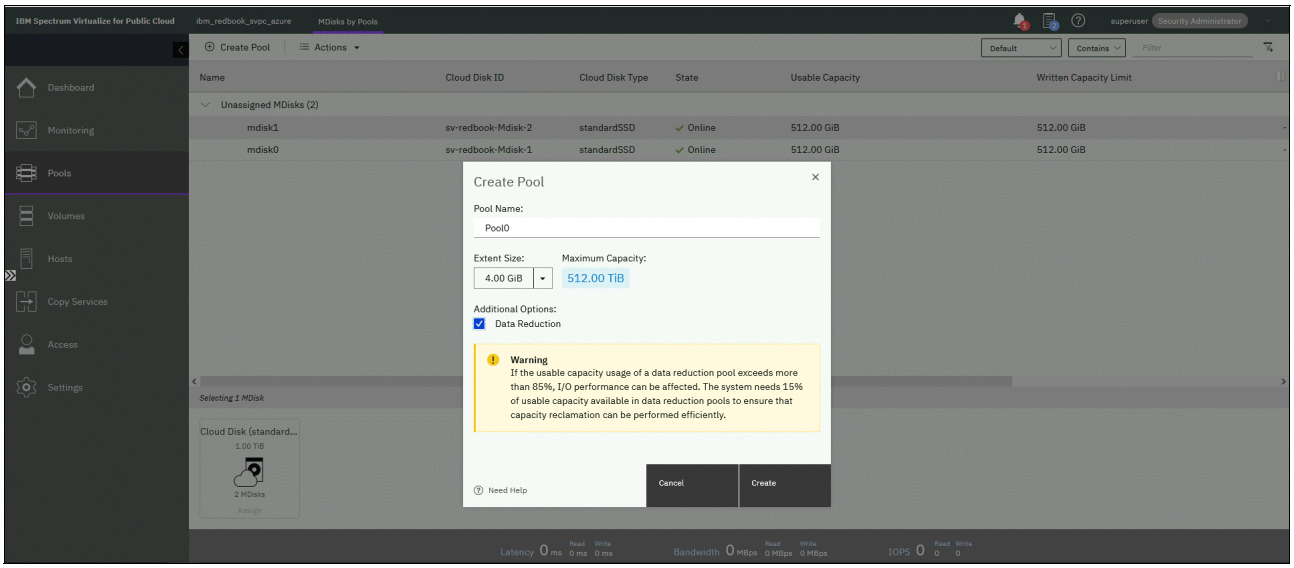

Figure 6-52 Create Pool

2. Complete the following step to add the Unmanaged Disk to the created pool (see Figure 6-55 on page 174):

   a. Right-click **Unmanaged Disk** and then, click **Assign**.
   b. Select the pool from the Pool list.
   c. Specify the suitable Tier.
   d. Click **Assign**.

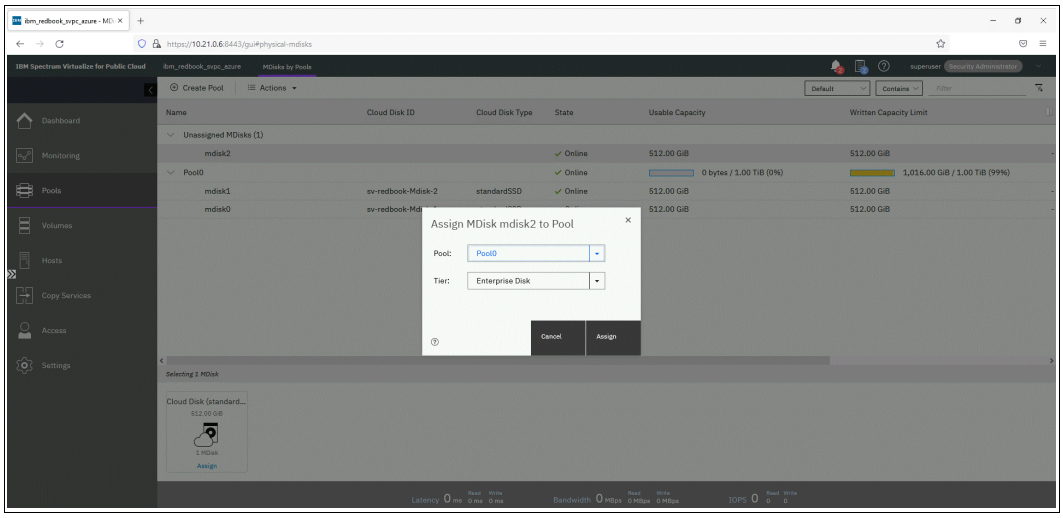

Figure 6-53 Assigning an unmanaged disk to a pool

6.6 Adding more back-end storage

After the initial deployment, adding more storage and different tiers is possible.

To add new Azure MDisks to IBM Spectrum Virtualize for Public Cloud, complete the following steps:

1. Open the Azure portal.
2. Select **Create Resource** and search for “Managed Disk”, as shown in Figure 6-54.

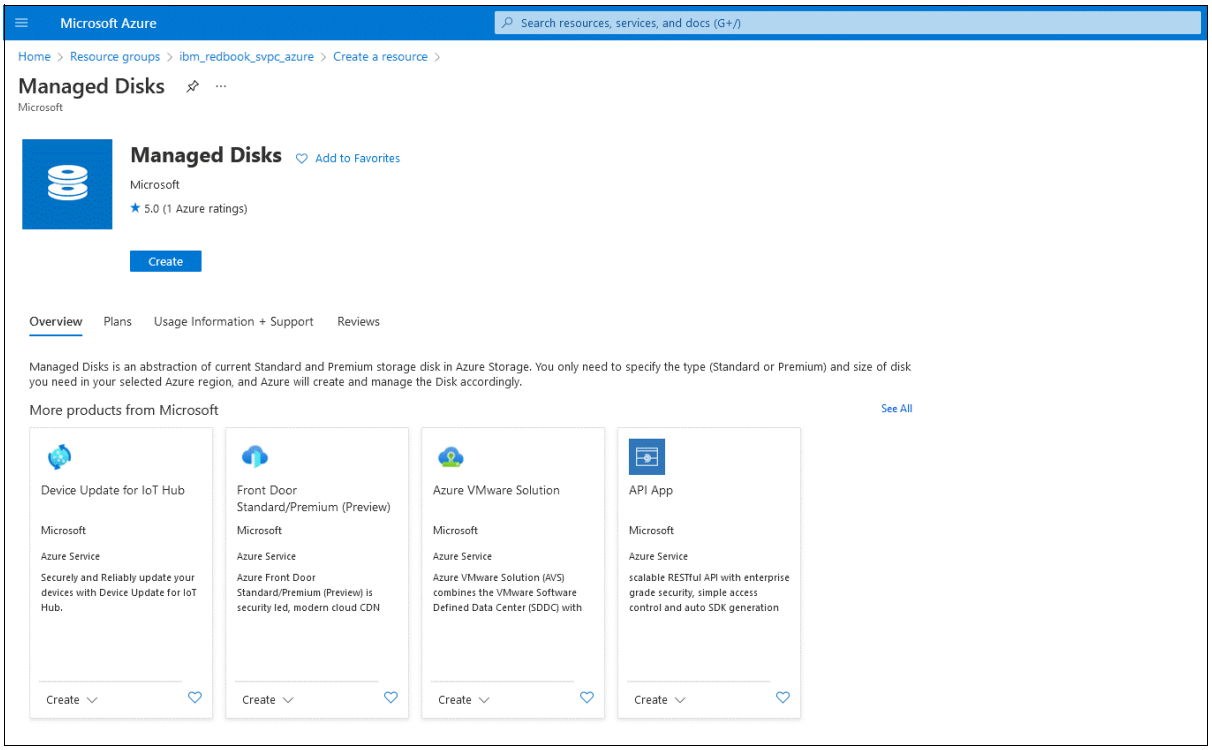

Figure 6-54 MDisk resource on the Azure portal

3. Enter the basic information for the new MDisk, as shown in Figure 6-55.

> **Note:** Use the correct IBM Spectrum Virtualize for Public Cloud RG to add the new Azure MDisks.

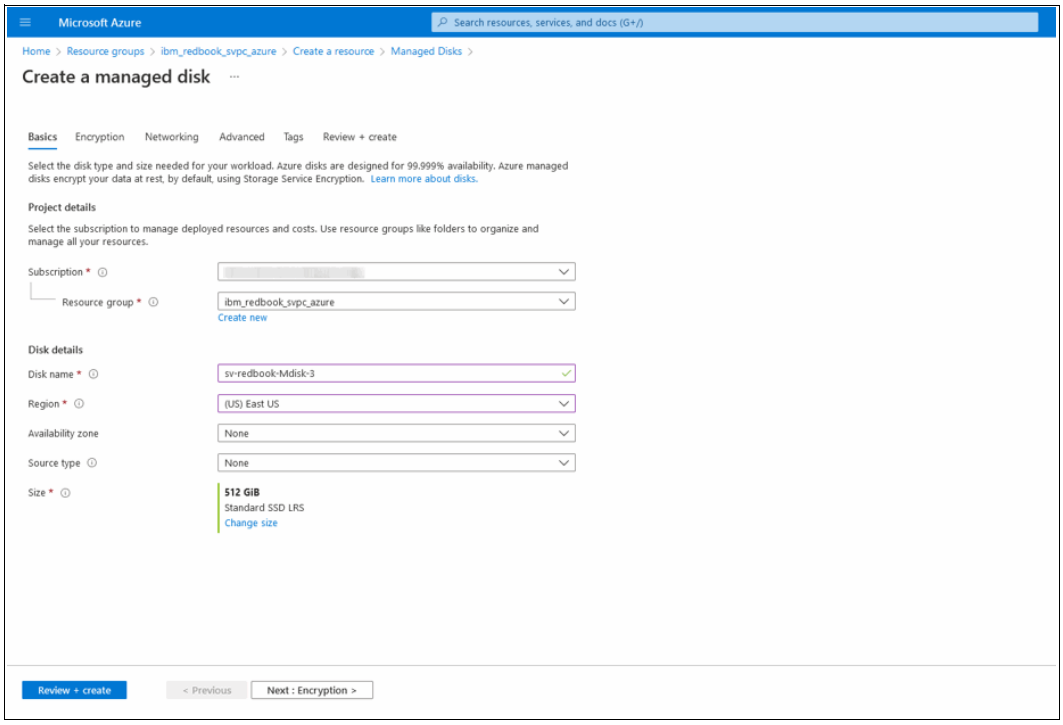

Figure 6-55 Basic disk configuration information in the Create a managed disk window

4. When selecting the size of disk, select the type of disk and the capacity, as shown in Figure 6-56. Click **Next**.

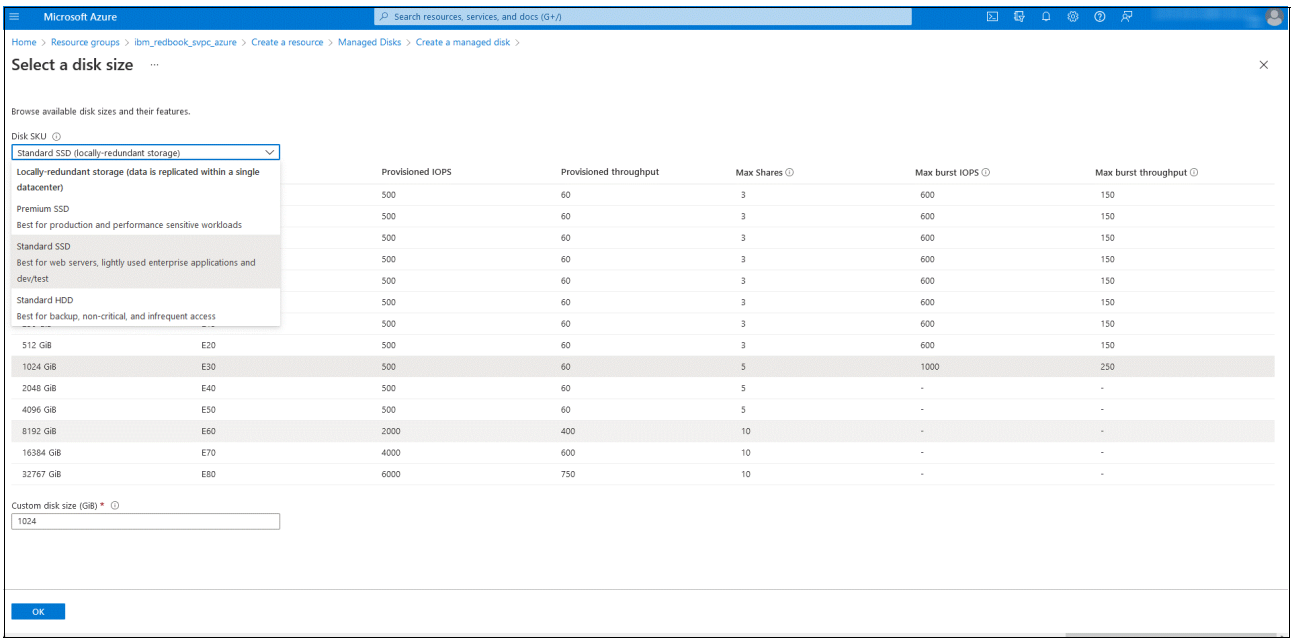

Figure 6-56 Selecting a size for an Azure MDisk

5. Select the (optional) encryption policy that is provided by Azure cloud Infrastructure, as shown in Figure 6-57. Click **Next**.

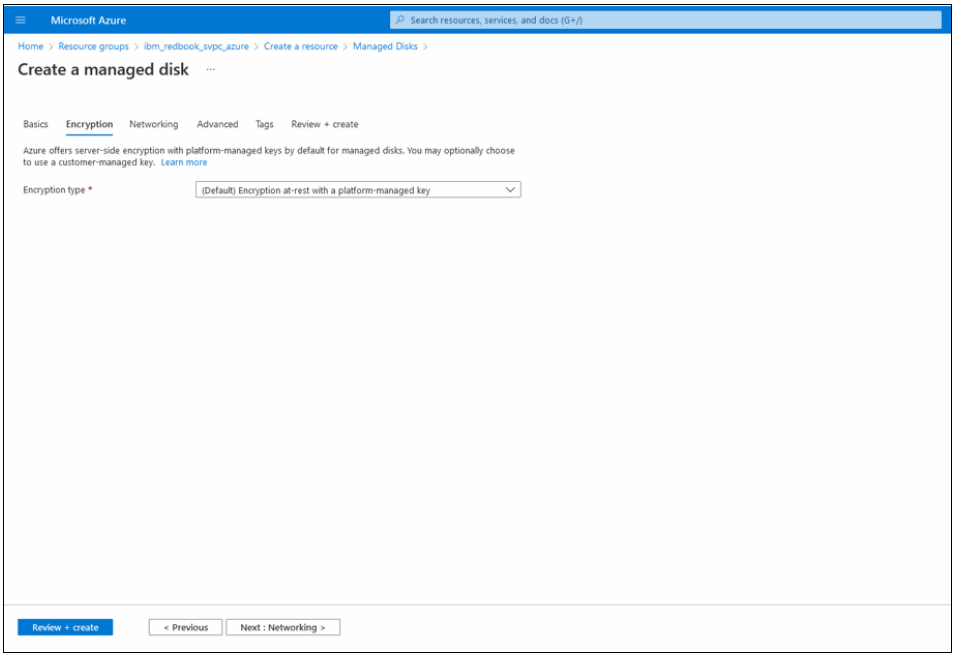

Figure 6-57 Encryption setting for Azure MDisk

6. Select the networking access for the MDisk, as shown in Figure 6-58, and then click **Next**.

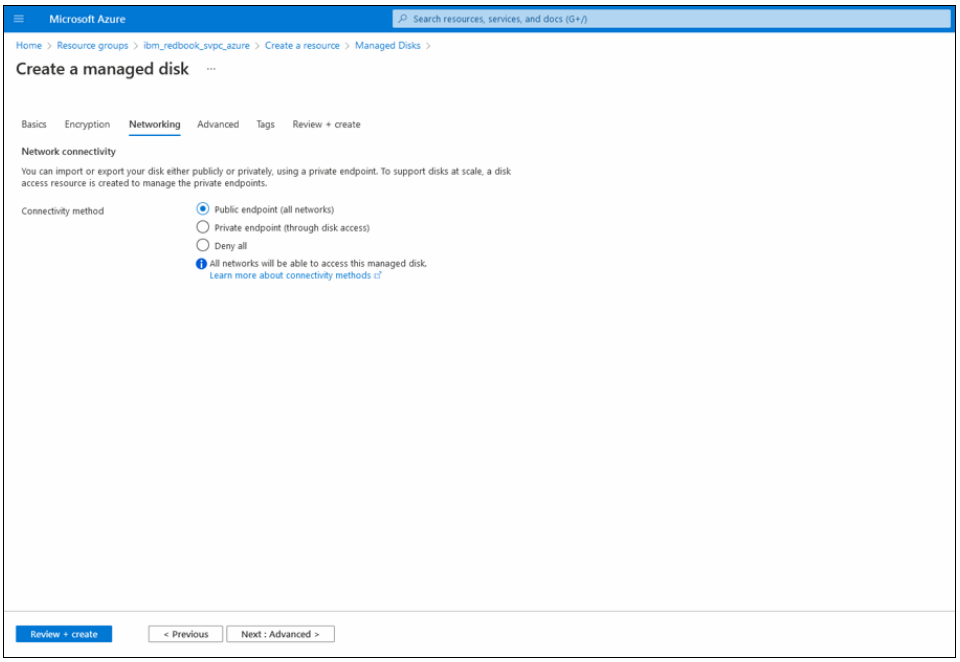

Figure 6-58 Network setting for more Azure MDisks

7. Enable **Shared Disk** as `Yes` and select the **Max share** as `2`, as shown in Figure 6-59. Click **Next**.

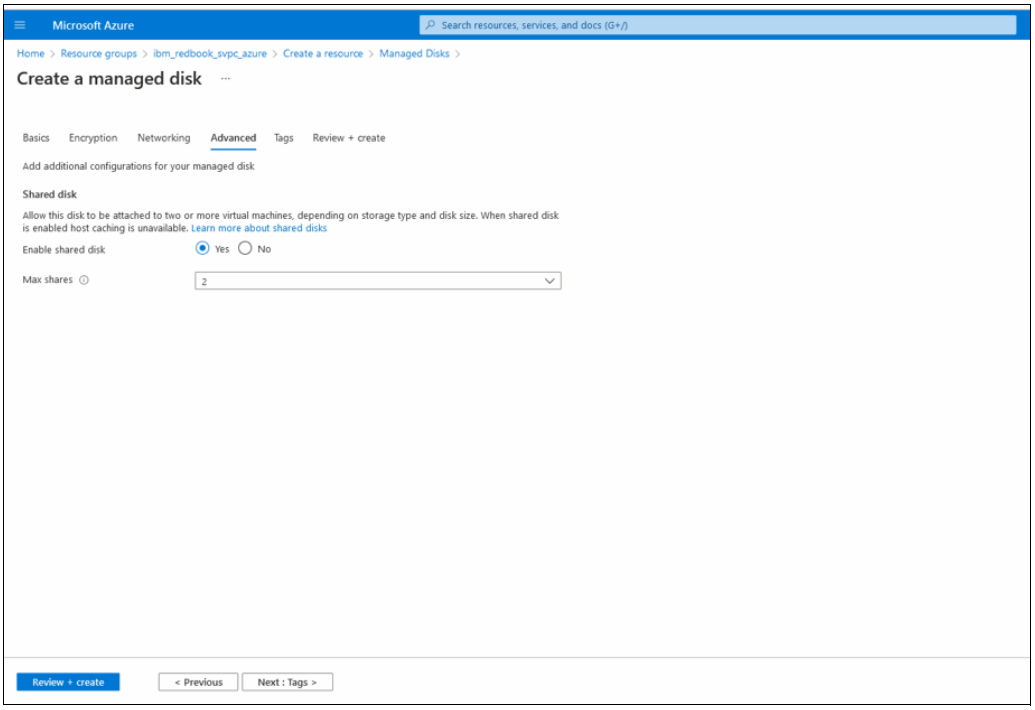

Figure 6-59 Disk Share setting for an Azure MDisk

8. Enter any optional Tags, and click **Next** to review and create the Azure MDisk, as shown in Figure 6-60.

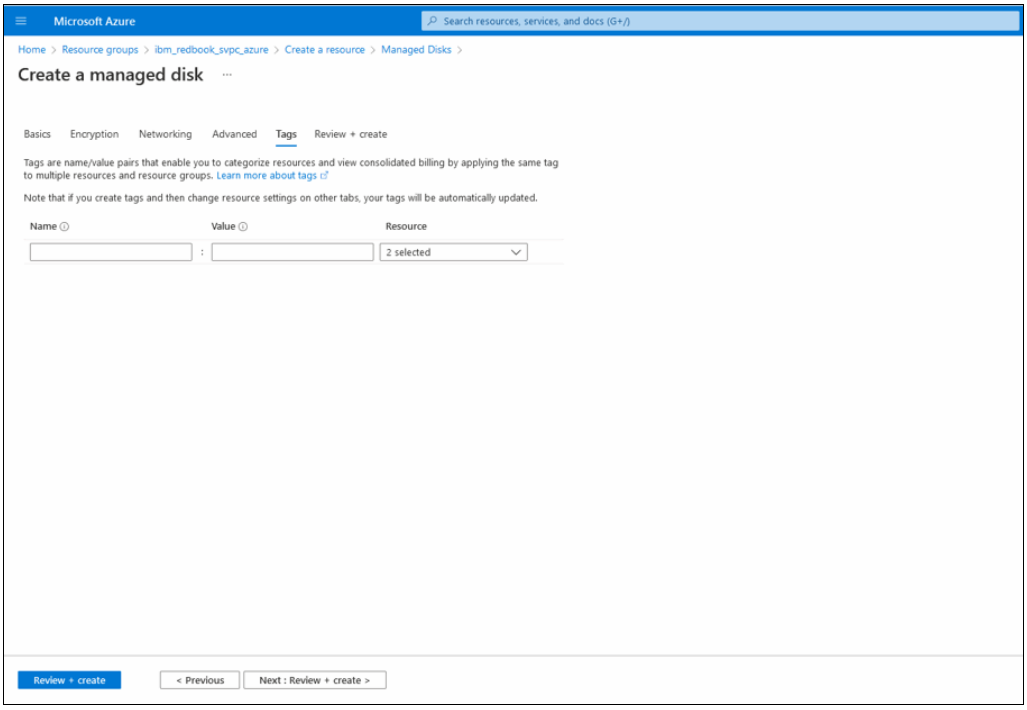

Figure 6-60 Tags setting for an Azure MDisk

9. Review the values that are provided. After the validation is successful, click **Create** to create the Azure MDisk, as shown in Figure 6-61.

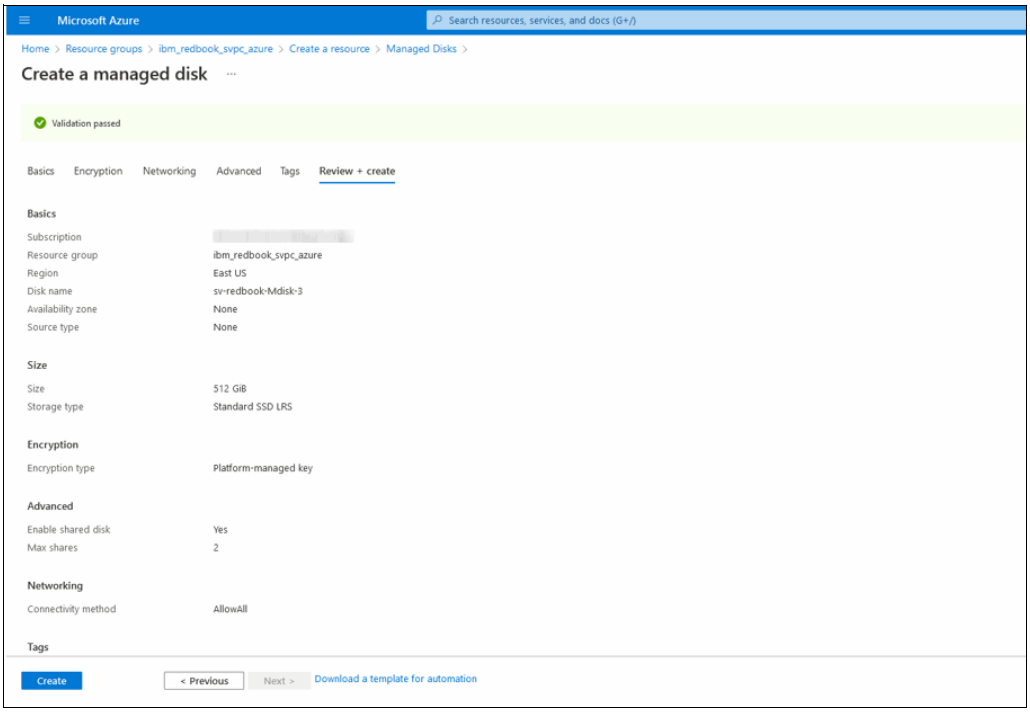

Figure 6-61 Azure MDisk validation and review

A successful deployment message is shown after the new Azure MDisk is provisioned, as shown in Figure 6-62.

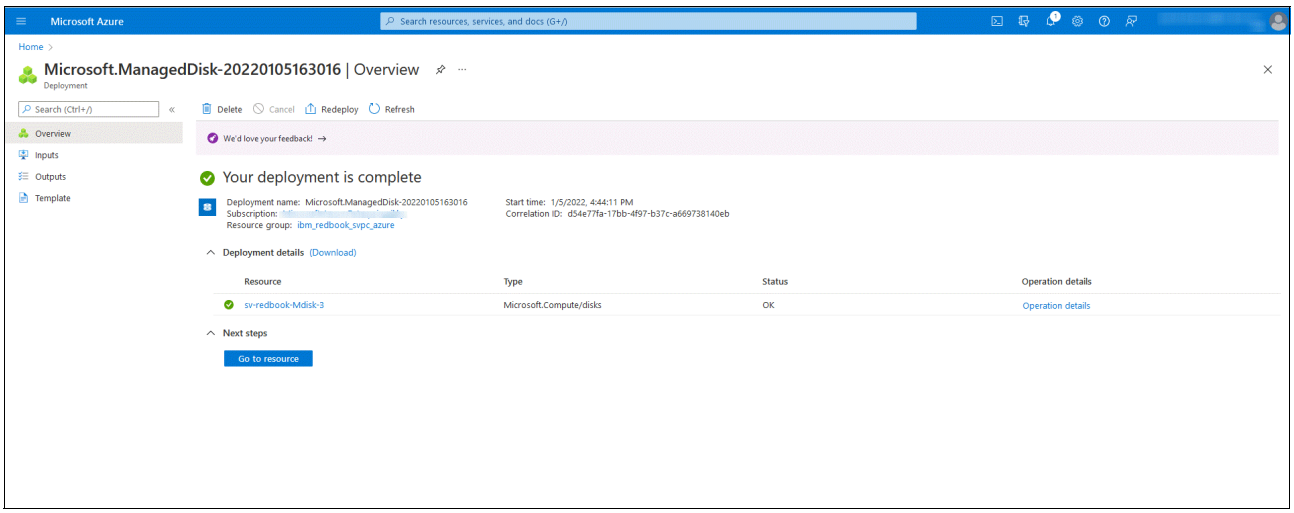

Figure 6-62 Successful deployment of an Azure MDisk

The disk is now available in the Azure RG, as shown in Figure 6-63.

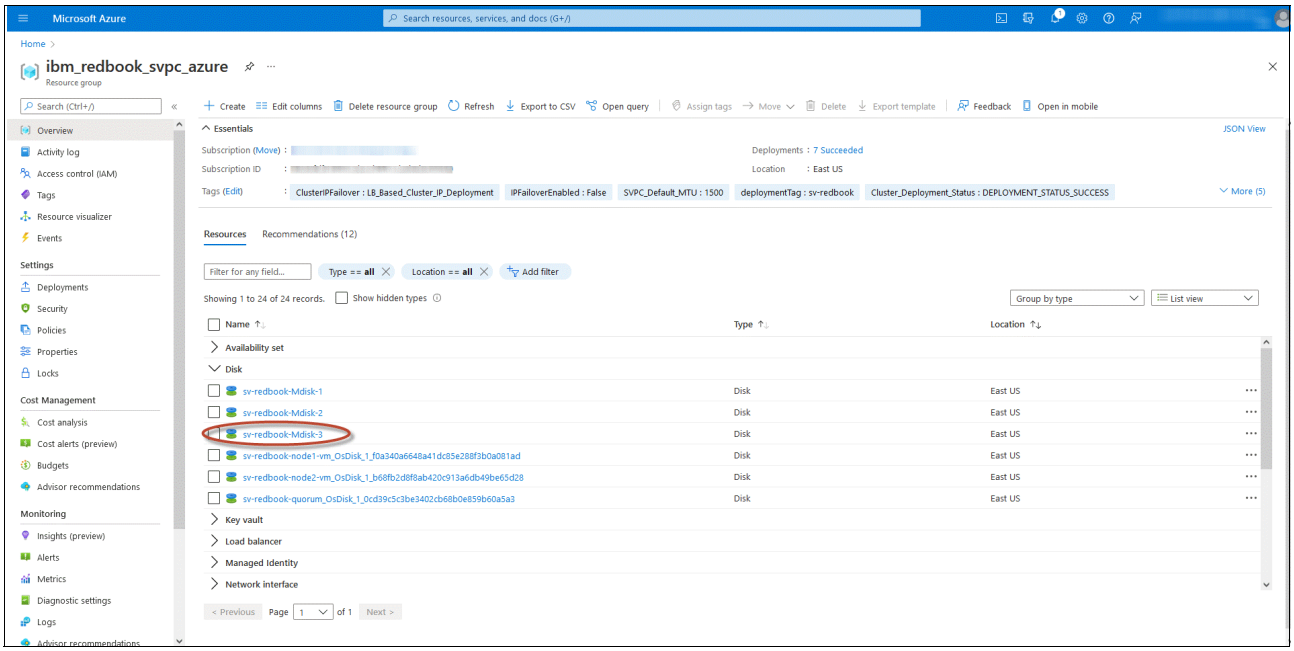

Figure 6-63 An Azure MDisk in an Azure resource group

10. In the IBM Spectrum Virtualize GUI, select **Pools** → **MdiskByPools**, as shown in Figure 6-64.

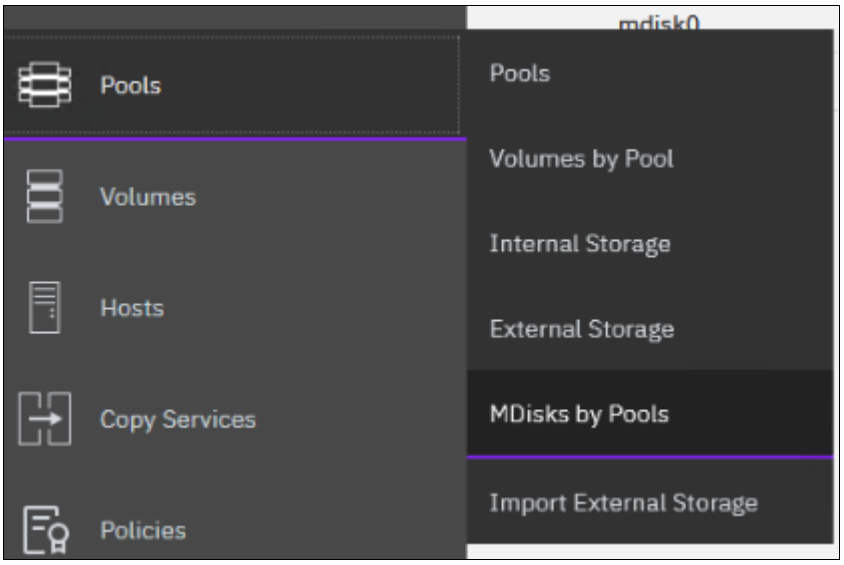

Figure 6-64 IBM Spectrum Virtualize for Public Cloud menu item to list all storage items: MDisks

11. Select **Actions** → **Discover Storage**, as shown in Figure 6-65 on page 179.

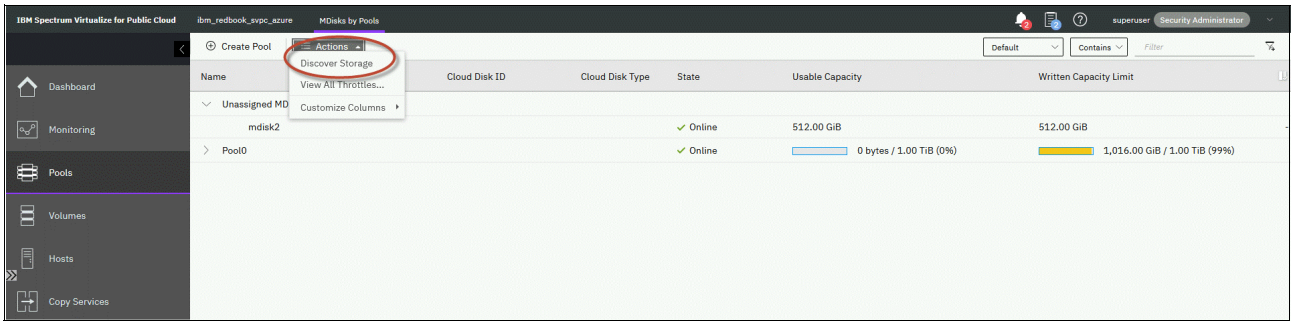

Figure 6-65 Discover Storage

The newly created disk is available under `UnManagedDisk`, as shown in Figure 6-66.

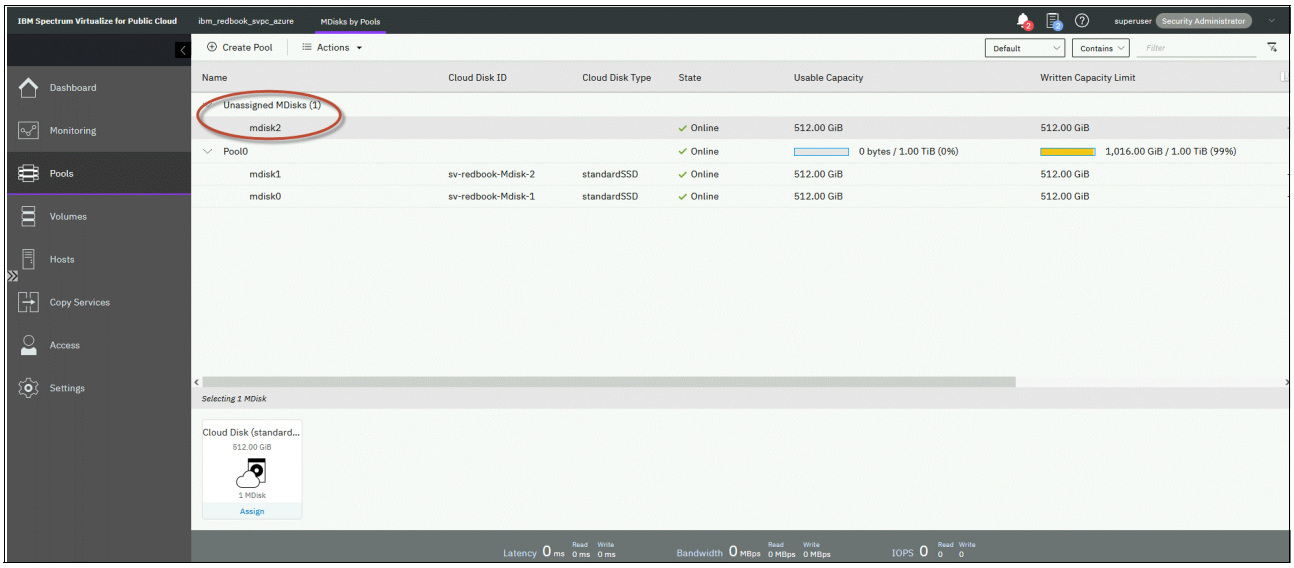

Figure 6-66 A newly created Azure MDisk seen as a unmanaged disk

We can add the disk to the existing pool by following the steps that are described in 6.5, “Configuring the back-end storage” on page 171.

Now, you can create the VDisk and assign the volume for host access by using ISCSI.

6.6.1 Configuring an IBM Spectrum Virtualize volume

In this section, you create a volume by using the pool that was created with the Azure MDisks. Volumes can be fully allocated or thinly provisioned (which are more space-efficient).

You can see both the CLI and GUI methods below. The default for thin-provisioned volumes that is indicated by the CLI, as shown in Example 6-1, is 2% (specified by the real size [`rsize`]). You have 98% of the capacity for the volumes that is available in the pool for other volumes until this volume claims it.

Example 6-1 Thinly provisioned (space-efficient) volume creation by using the CLI svctask

```
svctask mkvdisk -autoexpand -cache readwrite -iogrp io_grp0 -mdiskgrp 0 -name
svpc-azure-thin-vol -rsize 2% -size 32212254720 -unit b
```

Figure 6-67 shows thinly provisioned (space-efficient) volume creation by using the GUI.

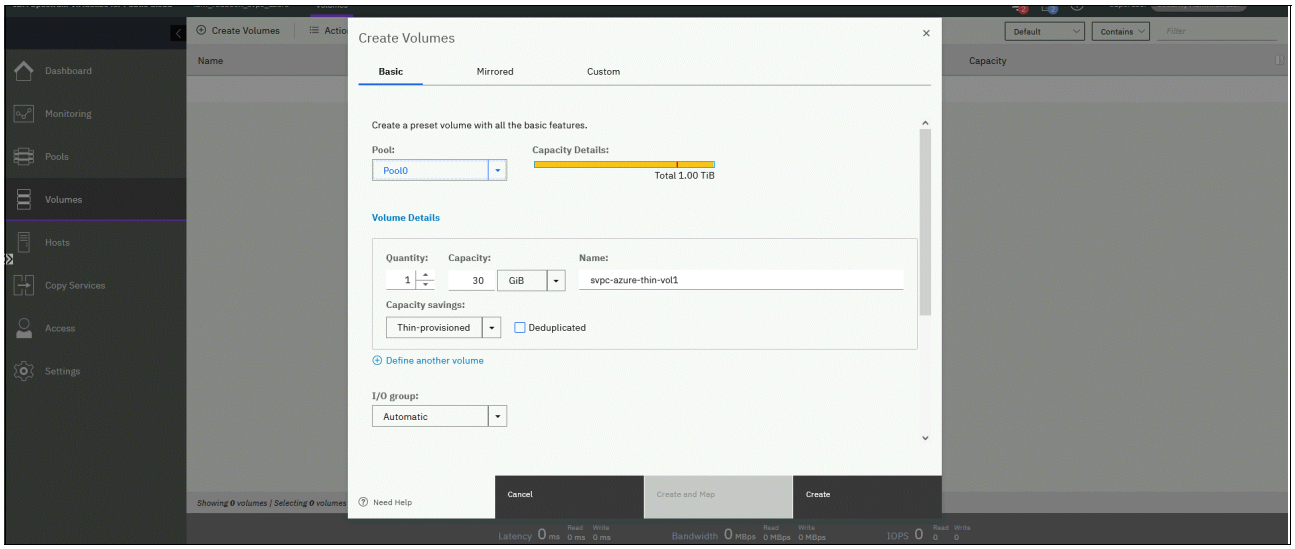

Figure 6-67 Thinly provisioned volume

Thinly provisioned volumes allow users to over-provision the Azure MDisk, which reduces the overall operational cost in Azure. Users can also use the Deduplication feature by checking the Deduplication option when the volume is created.

6.6.2 Configuring host and volume mapping

To use the volume that you created, you must map it to a host object. The host object represents a single server on your cloud account and its iSCSI-qualified name (IQN), which is similar to a worldwide port name (WWPN) for an FC host. To create a host object, you must collect its IQN. The place and the procedure to collect the IQN from can vary with each operating system. For the suitable steps for an operating system, see the operating system's documentation.

When you create your host object and map your volume, depending on what operating system you use, you must install the iSCSI initiator and run some specific operations to use your mapped volumes with the hosts.

Linux host

Install the Linux software iSCSI initiator. The initiator software on RHEL systems is packaged as `iscsi-initiator-utils`, and the suggested version is 6.2.1.2-1 or later. The initiator software on SUSE Linux Enterprise Server systems is packaged as `open-iscsi`. According to IBM Documentation, set the IQN, target discovery, and authentication, and enable multipathing for the Linux hosts.

After creating the host object and mapping VDisks to it, on the IBM Spectrum Virtualize cluster, scan for the disks on the host by using the specific iSCSI command, as is done for an on-premises IBM Spectrum Virtualize Cluster.

Check the multipath output (run **`multipath -ll`**) to ensure that your VDisks are attached correctly through the multipath tool.

A typical output of a VDisk should resemble the output that is shown in Example 6-2 on page 181.

Example 6-2 Multipath output

```
multipath -ll
mpathah (36005076072a06dc4f00000000000000c) dm-1 IBM,2145
size=30G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 9:0:0:1 sdi 8:128 active ready running
| `- 8:0:0:1 sdj 8:144 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 6:0:0:1 sde 8:64  active ready running
  `- 7:0:0:1 sdf 8:80  active ready running
mpathag (36005076072a06dc4f00000000000000d) dm-0 IBM,2145
size=1.0G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 8:0:0:0 sdg 8:96  active ready running
| `- 9:0:0:0 sdh 8:112 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 6:0:0:0 sdc 8:32  active ready running
  `- 7:0:0:0 sdd 8:48  active ready running
```

Windows host

The software iSCSI initiator is built into the system on Windows 2008 and later. Access the iSCSI initiator from the Control Panel or search from the **Start** menu.

An example is shown in Figure 6-68.

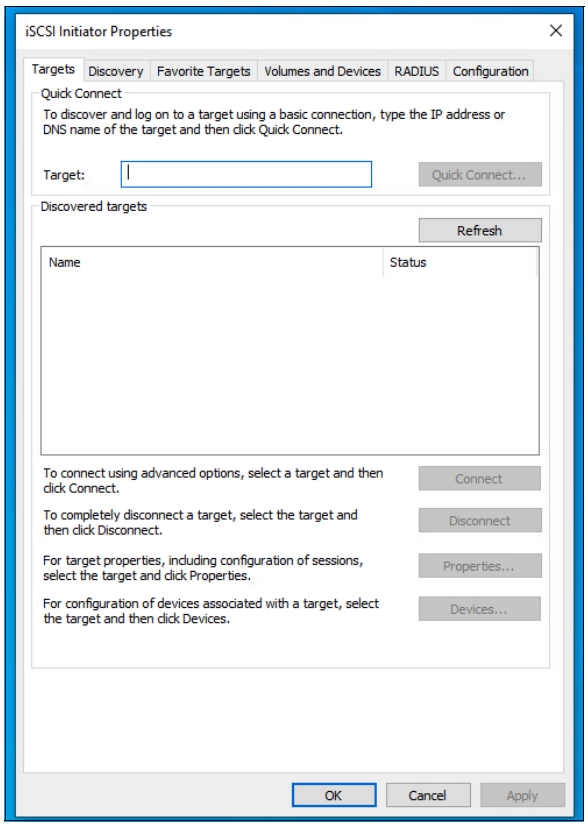

Figure 6-68 Windows iSCSI window to configure and find the IQN and storage

Discover the iSCSI target by using Send Targets or by using iSNS. For more information, see IBM Documentation.

Connect to the discovered targets, as described in IBM Documentation.

Now, the mapped volumes are visible to Windows disk management services. The system volumes can be initialized, formatted, and mounted. You can view the details of the discovered disks by using the Windows Command Prompt. An example output is shown in Example 6-3.

Example 6-3 Windows OS Diskpart command example

```
DISKPART> list disk
 Disk ###  Status      Size     Free     Dyn  Gpt
 --------  ----------  -------  -------  ---  ---
 Disk 0    Online       149 GB    78 GB   *
 Disk 1    Online       149 GB    78 GB   *
 Disk 2    Online       565 MB   565 MB
 Disk 3    Online       337 MB   337 MB
DISKPART> select disk 2
Disk 2 is now the selected disk.
DISKPART> detail disk
IBM      2145             SCSI Disk Device
Disk ID: 00000000
Type   : iSCSI
Bus    : 0
Target : 2
LUN ID : 0
There are no volumes.
```

6.7 Configuring a site-to-site virtual private network gateway for hybrid cloud connectivity in Azure cloud

This section describes how to configure hybrid cloud connectivity between Azure cloud and the on-premises environment. This section also describes the lab setup and the steps to configure the site-to-site IPsec tunnel for communication between Azure cloud and the on-premises site. A video example is at YouTube and IBM MediaCenter.

The virtual private network gateway (VPNGW) IPsec site-to-site tunnel creates a secure communication network between the Azure cloud infrastructure and on-premises infrastructure. Network communication between the private subnets is controlled by the access control list (ACL) that is populated when you create the VPN IPsec site-to-site tunnel.

6.7.1 Azure configuration for a VPNGW IPsec tunnel

This section describes the steps that are required at the vNet level in Azure cloud to establish the IPsec tunnel. For more information, see this Azure tutorial web page.

Complete the following steps:

1. Create a VPNGW:
   a. Log in to Azure console with administrator privileges.
   b. Select **Create a Resource**, and then search for “Virtual Network Gateway”.
   c. Enter the required information for the Virtual Network Gateway and associate it with the vNet in Azure to be used for hybrid connectivity.
2. Create a local network gateway:
   a. Log in to the Azure console with administrator privileges.
   b. Select **Create a Resource**, and then search for “Local Network Gateway”.
   c. Enter the required information for the Local Network Gateway for hybrid connectivity.
3. Create a VPN device.

   A site-to-site connection requires a VPN device for connection to on-premises setup. Follow the Azure documentation to create the VPN device for your on-premises setup.
4. Create a VPN connection:
   a. Select the Virtual Network Gateway that was created in Step 1.
   b. Select **Connections** and then, click **Add** to create a connection.
   c. Enter the local gateway that was created in Step 2.

      When this process is complete, a VPN connection is established between your on-premises and Azure cloud network.
5. Verify the connection by connecting to a VM in cloud or on-premises.

6.8 Configuring replication from on-premises IBM Spectrum Virtualize to IBM Spectrum Virtualize for Public Cloud on Azure

This section describes how to configure replication from an on-premises solution that can be an IBM FlashSystem or IBM SAN Volume Controller system to an IBM Spectrum Virtualize for Public Cloud on Azure solution.

Our example uses a IBM FlashSystem system in the on-premises data center and a 2-node IBM Spectrum Virtualize for Public Cloud cluster on Azure as a disaster recovery (DR) storage solution.

This scenario uses IBM Spectrum Virtualize *volume groups (VGs)* to replicate the data from the on-premises data center to Amazon Web Services (AWS) Cloud.

This implementation starts with the assumption that the IP network connectivity between the on-premises data center IBM FlashSystem array and Azure cloud is established through a *Multiprotocol Label Switching* (MPLS) or VPN connection. Because several methods are available to implement the IP network connectivity, this section does not consider that specific configuration. For more information, contact your organization’s network technical specialist.

To configure the replication, complete the following steps:

1. Go to **Copy Services** and click **Create Partnership**. This configuration is required on both sites, as shown in Figure 6-69.

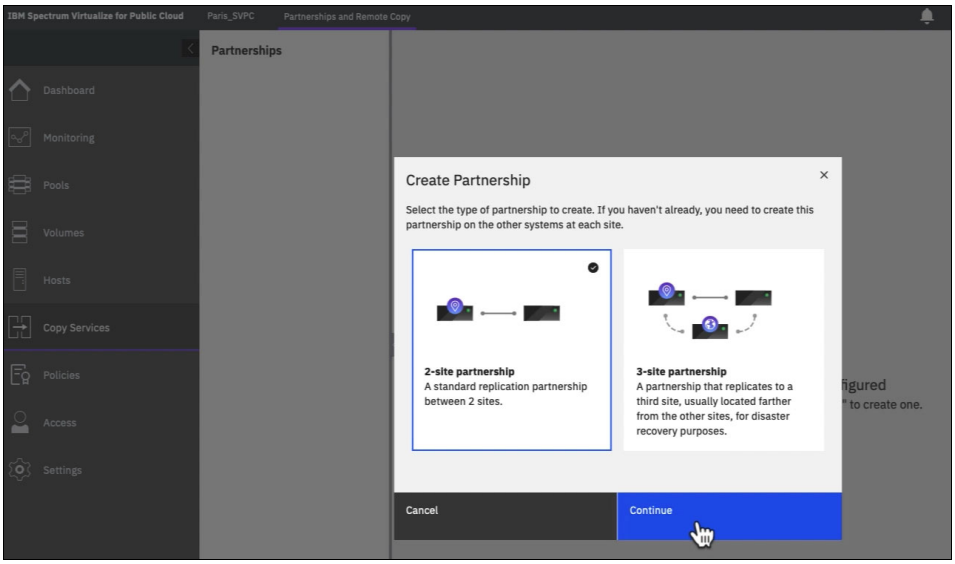

Figure 6-69 Setting up a partnership between IBM Spectrum Virtualize for Public Cloud and an on-premises IBM FlashSystem

2. Input the IP address of the target system, as shown in Figure 6-70.

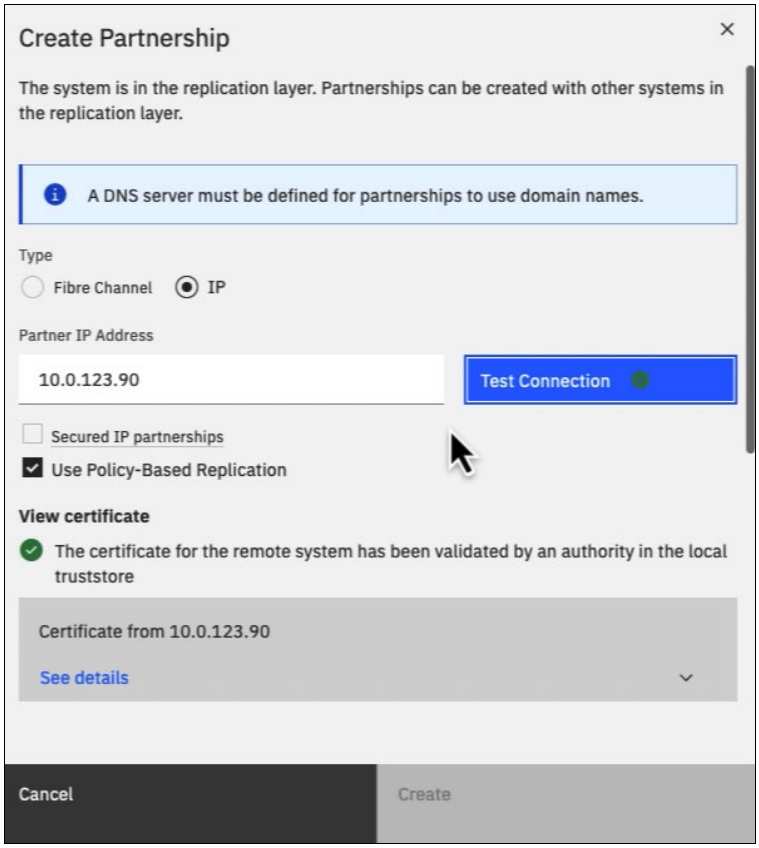

Figure 6-70 Entering the IP address and testing the replication partnership configuration example

3. Scroll down and choose the bandwidth and portset, as shown in Figure 6-71.

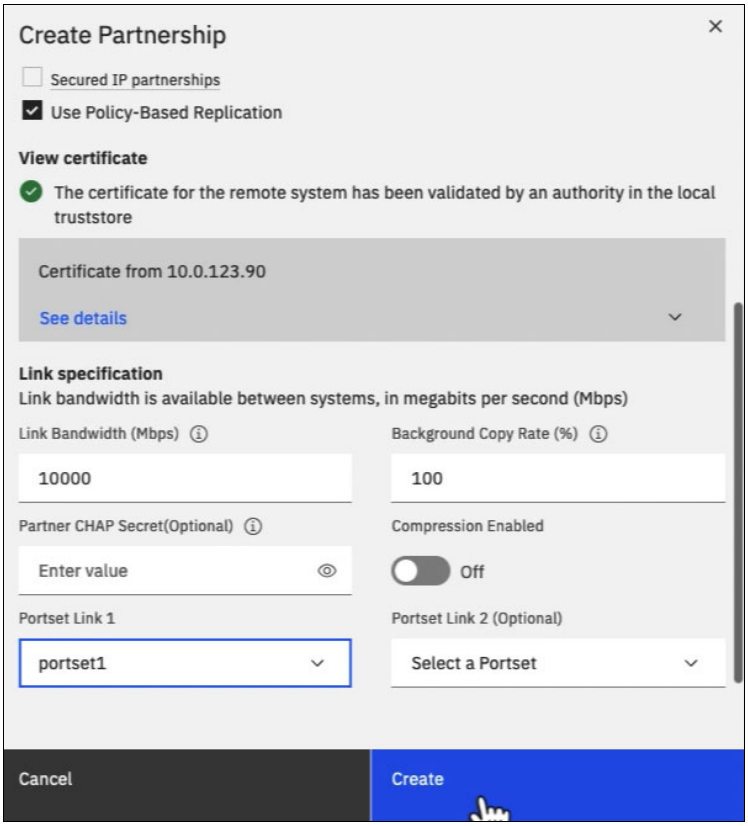

Figure 6-71 Choosing the partnership bandwidth and portset

4. Repeat steps 1 on page 184 to 3 for the other instance, but point the IP address to the other instance. Link the pools so that volumes in the source pool are replicated to the target replication pool system. You can use a child or parent pool. Optionally, you can use a provisioning policy.

A sample is shown in Figure 6-72.

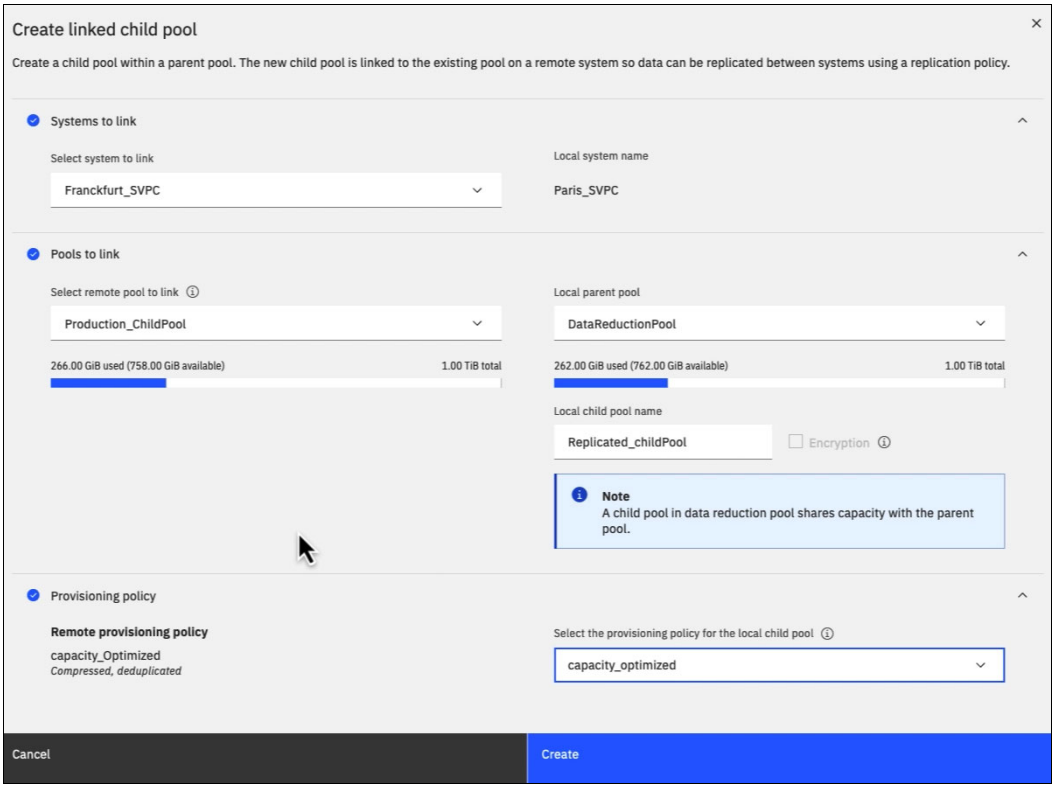

Figure 6-72 Linking the primary and target pools by using child pools replication example

5. Follow the checklist to create the replication policy, as shown in Figure 6-73.

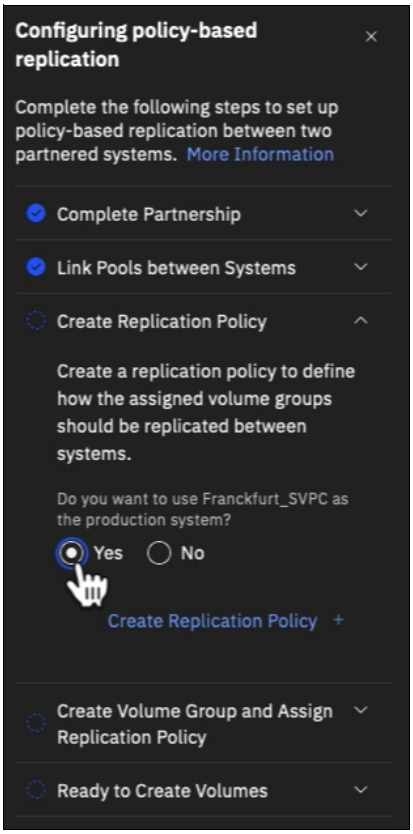

Figure 6-73 Replication checklist: Create Replication Policy

6. Create the replication policy and configure the recovery point objective (RPO) value based on your business needs, as shown in Figure 6-74.

Note: Know what versions of IBM Spectrum Virtualize software are supported. For supported and interoperability versions, see IBM Spectrum Virtualize Family of Products Inter-System Metro Mirror and Global Mirror Compatibility Cross Reference.

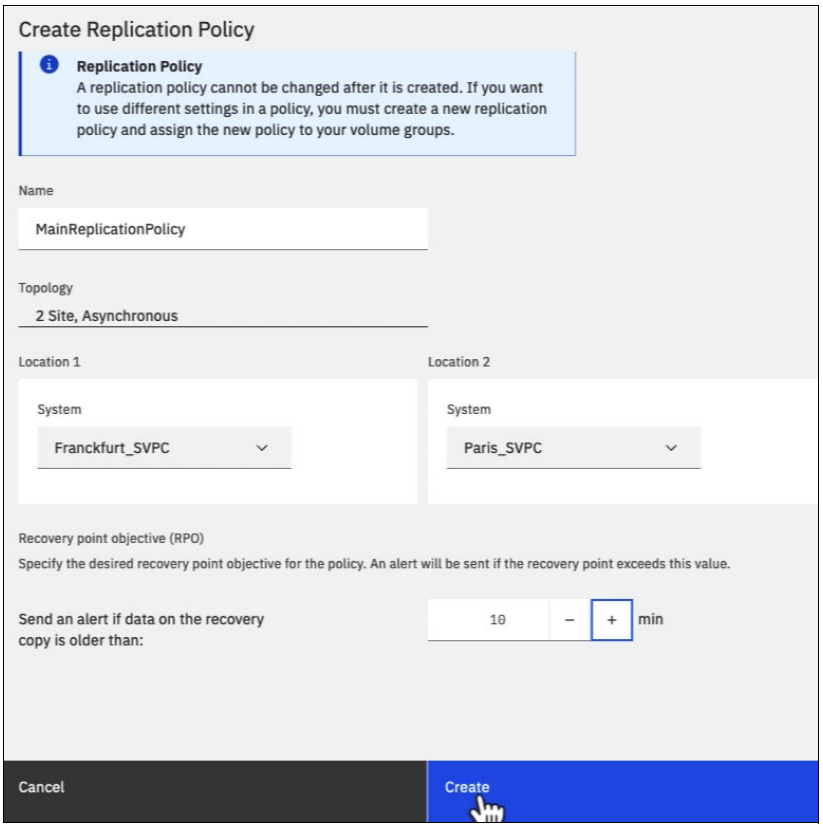

Figure 6-74 Replication policy creation

7. Create the VG on the source or primary IBM Spectrum Virtualize for Public Cloud system, as shown in Figure 6-75 on page 189.

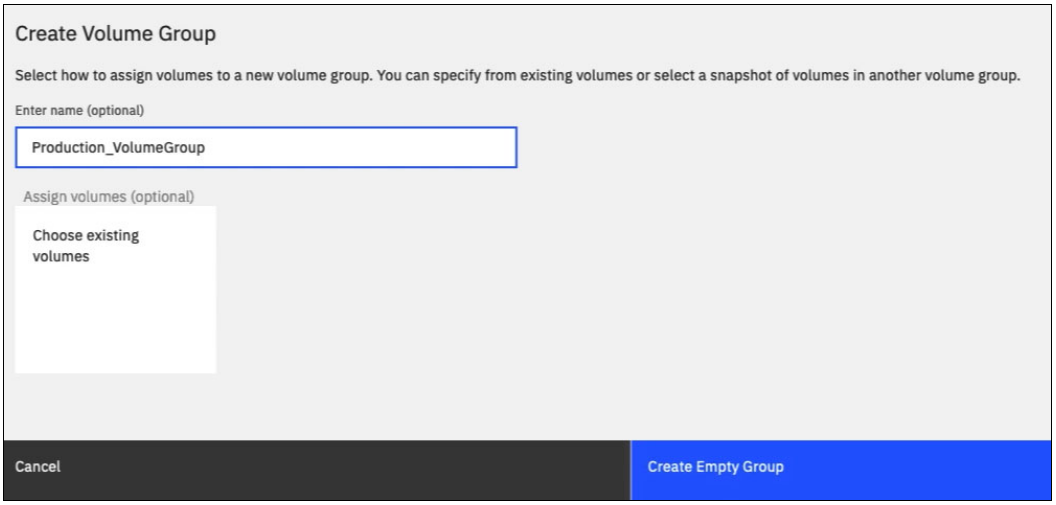

Figure 6-75 Creating a volume group on the source IBM Spectrum Virtualize for Public Cloud system example

8. Assign the replication policy, as shown in Figure 6-76.

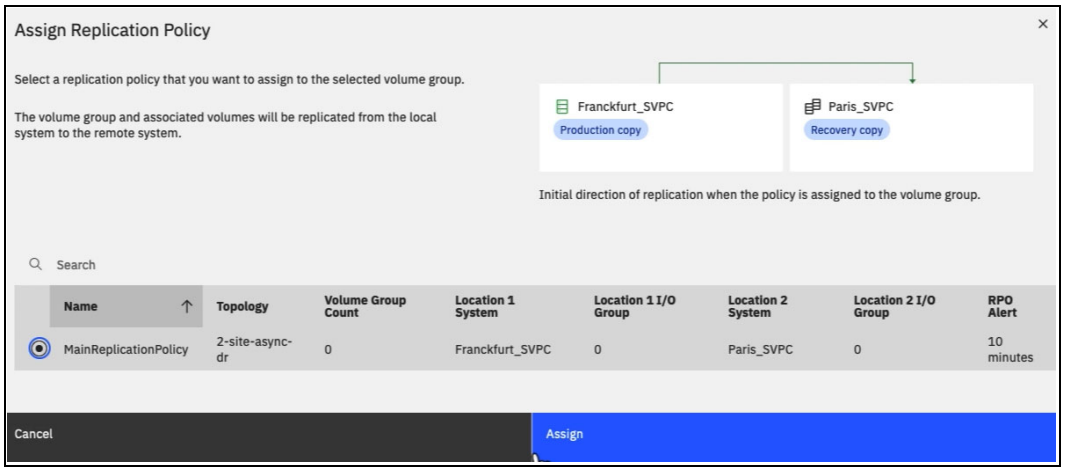

Figure 6-76 Assigning a replication policy example

As shown in Figure 6-77, you completed the various steps of partnership, linking the source and target pools, and creating a replication policy and its RPO. Now, you can create volumes (or add existing source volumes into this volume group, which will be replicated).

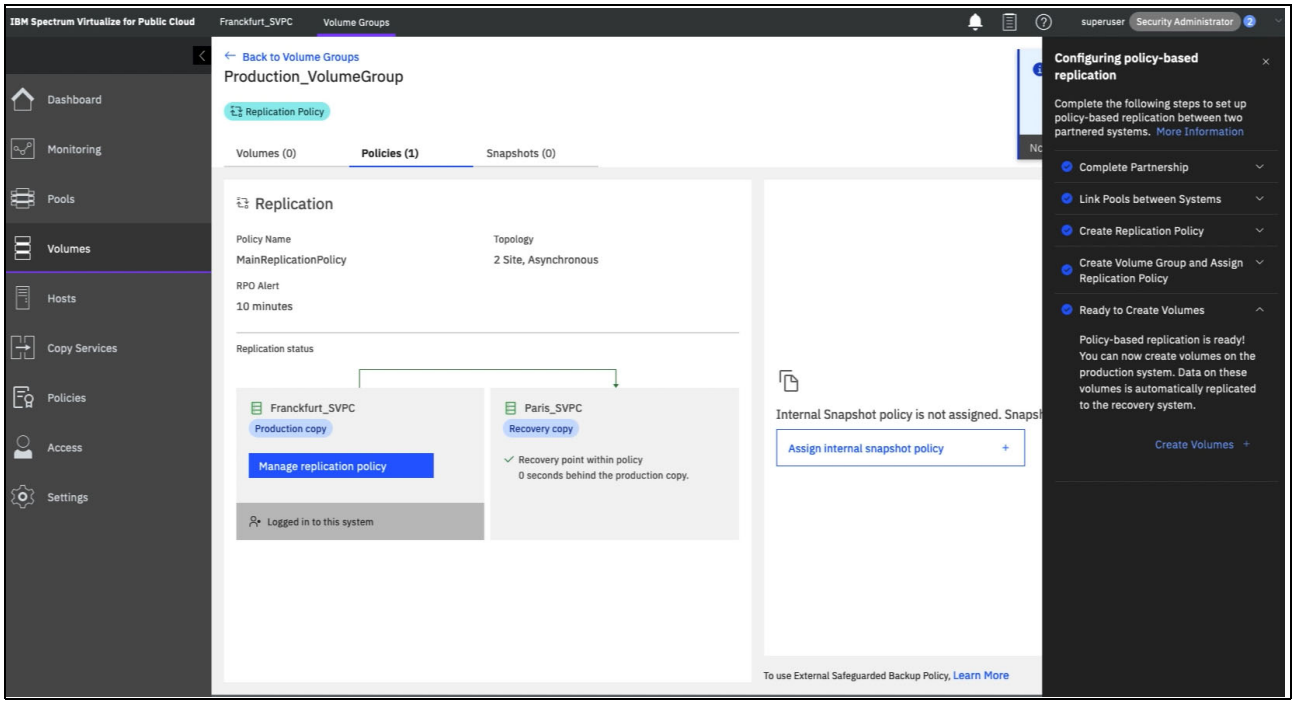

Figure 6-77 Replication policy, volume group, and linked pools completed

9. Complete the partnership configuration in the IBM Spectrum Virtualize for Public Cloud on the Azure side by providing the on-premises cluster IP address.

   Your partnership is fully configured, as shown in Figure 6-78.

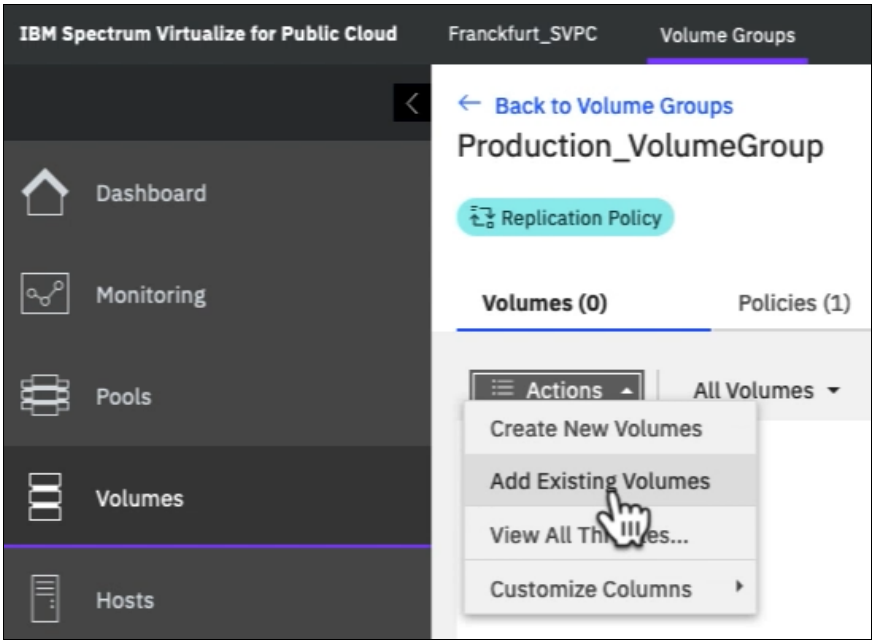

Figure 6-78 Adding or creating volumes in the volume group for replication

Note: The connection might take a few seconds to synchronize, but double-clicking **Partnership** shows the confirmed status of the partnership quicker.

10. In our example, four 100 GiB volumes are added to the VG for replication, as shown in Figure 6-79.

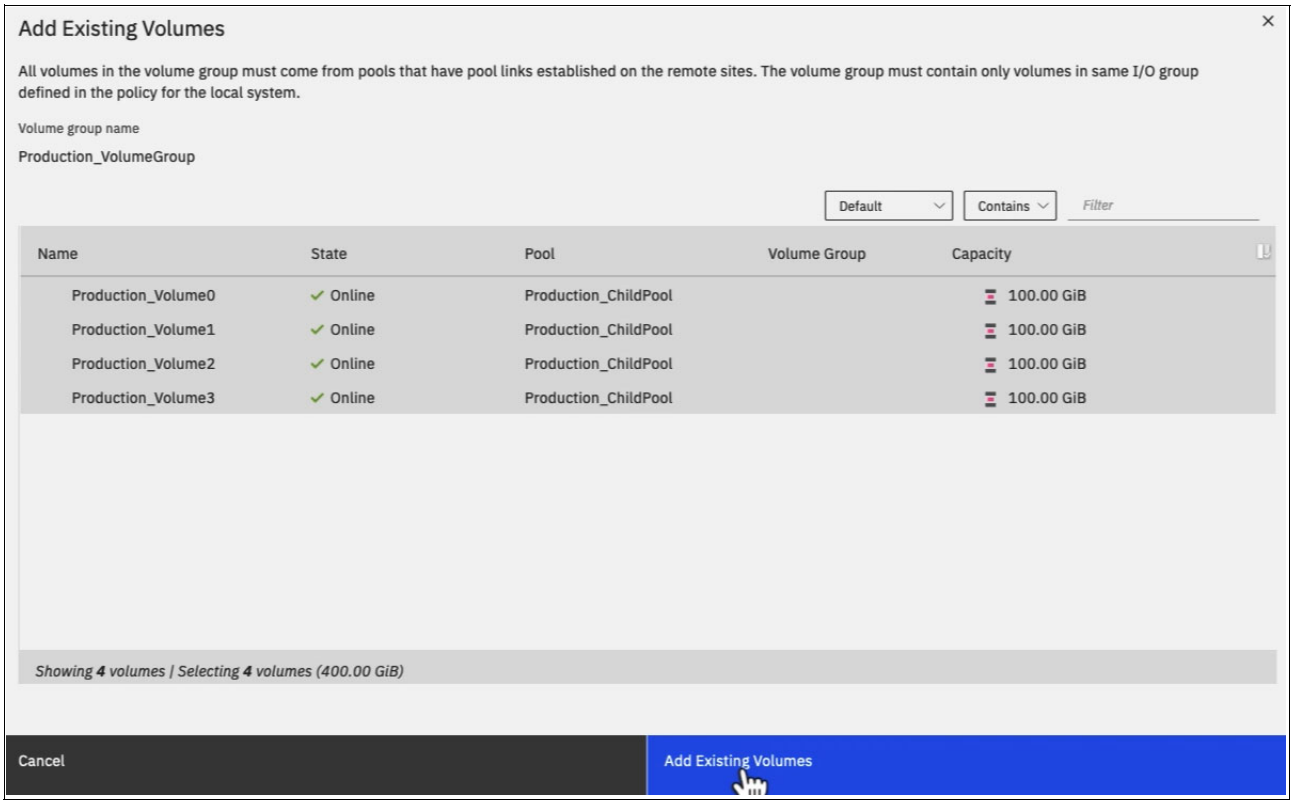

Figure 6-79 Add existing volumes to a volume group example

11. After adding the volumes, the replication begins immediately, as shown in Figure 6-80.

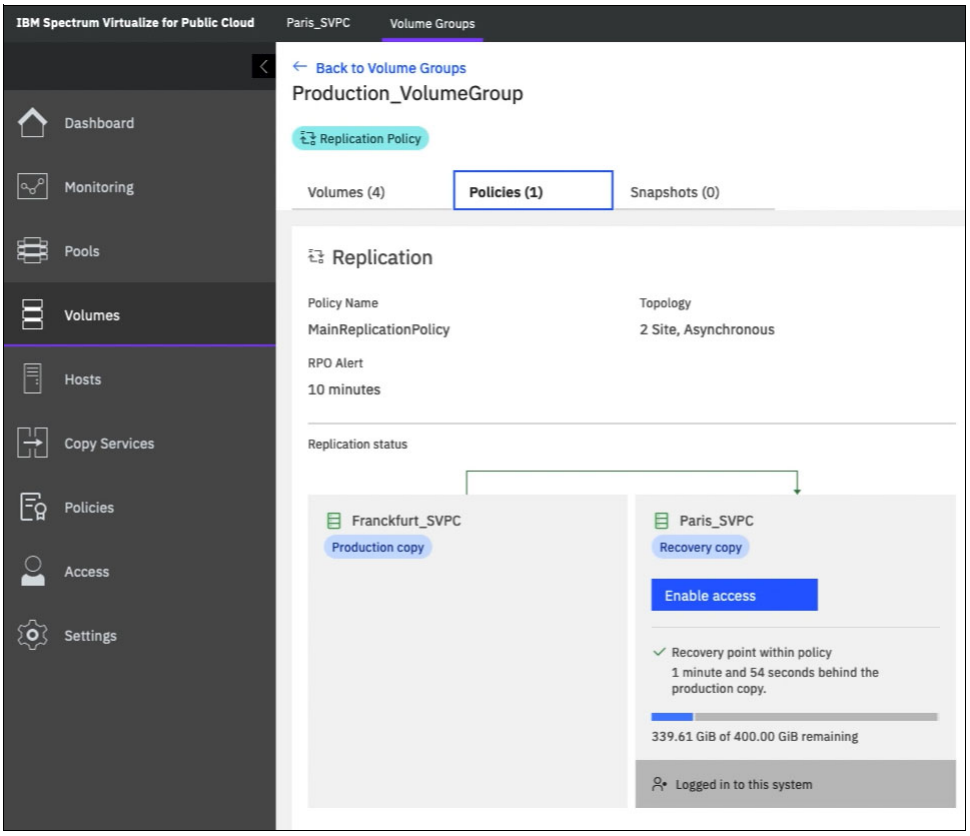

Figure 6-80 Volumes immediately begin replicating to a target array after they are added

7

Implementing key features

This chapter describes how to implement Safeguarded snapshots and Policy-Based Replication (PBR) on IBM Spectrum Virtualize for Public Cloud on Amazon Web Services (AWS) and Microsoft Azure environments.

This chapter includes the following topics:

- Configuring Safeguarded snapshots
- Configuring secured Policy-Based Replication

7.1 Configuring Safeguarded snapshots

This section describes how to configure Safeguarded snapshots by using an Internal Snapshot policy. The section also describes how to recover Safeguarded snapshots.

Complete the following steps:

1. For our example implementation, we create one volume group (VG) (`production_volume_group`) and two volumes (`production_source1` and `production_source2`) inside the VG, as shown in Figure 7-1.

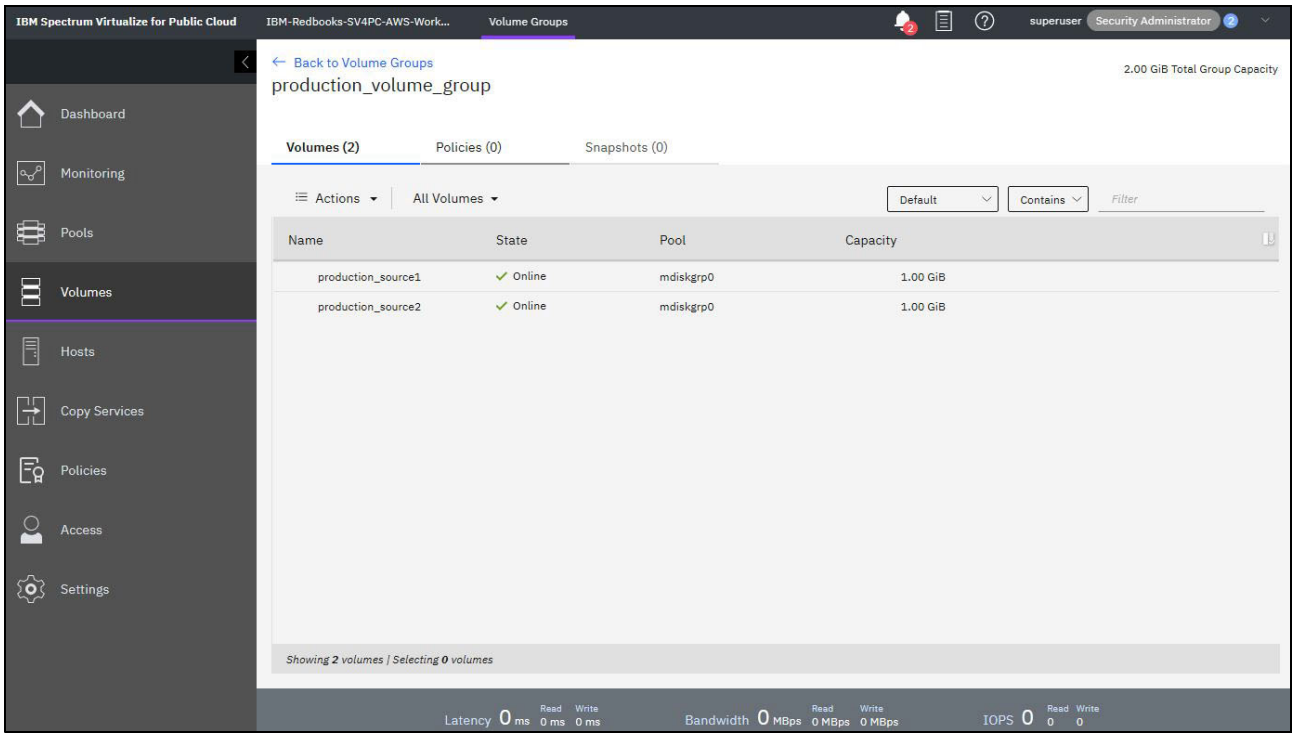

Figure 7-1 Source Safeguarded volume group

2. Safeguarded snapshot can be created by either using predefined policies, as shown in Figure 7-2 on page 195, or by using user-defined policy.

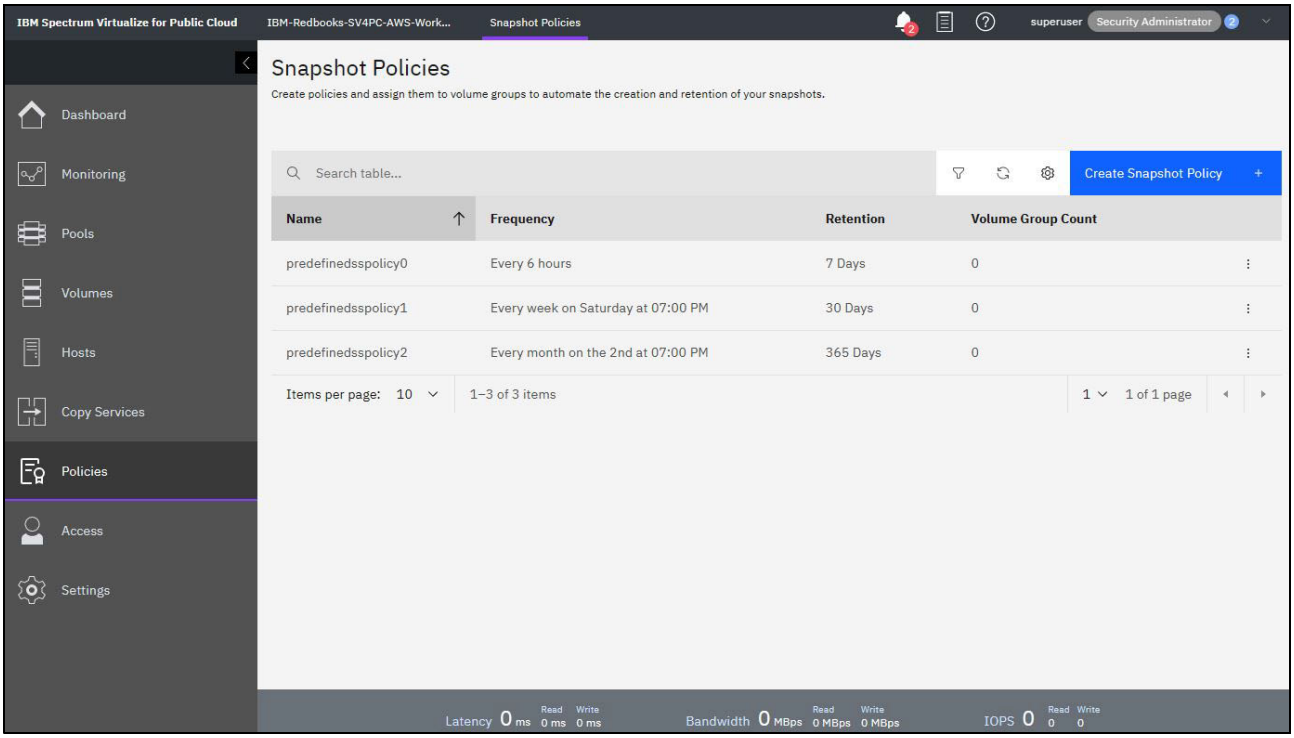

Figure 7-2 Predefined snapshot policies

3. To create a user-defined policy, select **Policies** → **Snapshot Policies**, click **Create Snapshot Policy**, and enter the required details. Figure 7-3 shows the creation of a user-defined policy `userdefinedsspolicy0` for creating snapshots every hour with a retention period of 1 day.

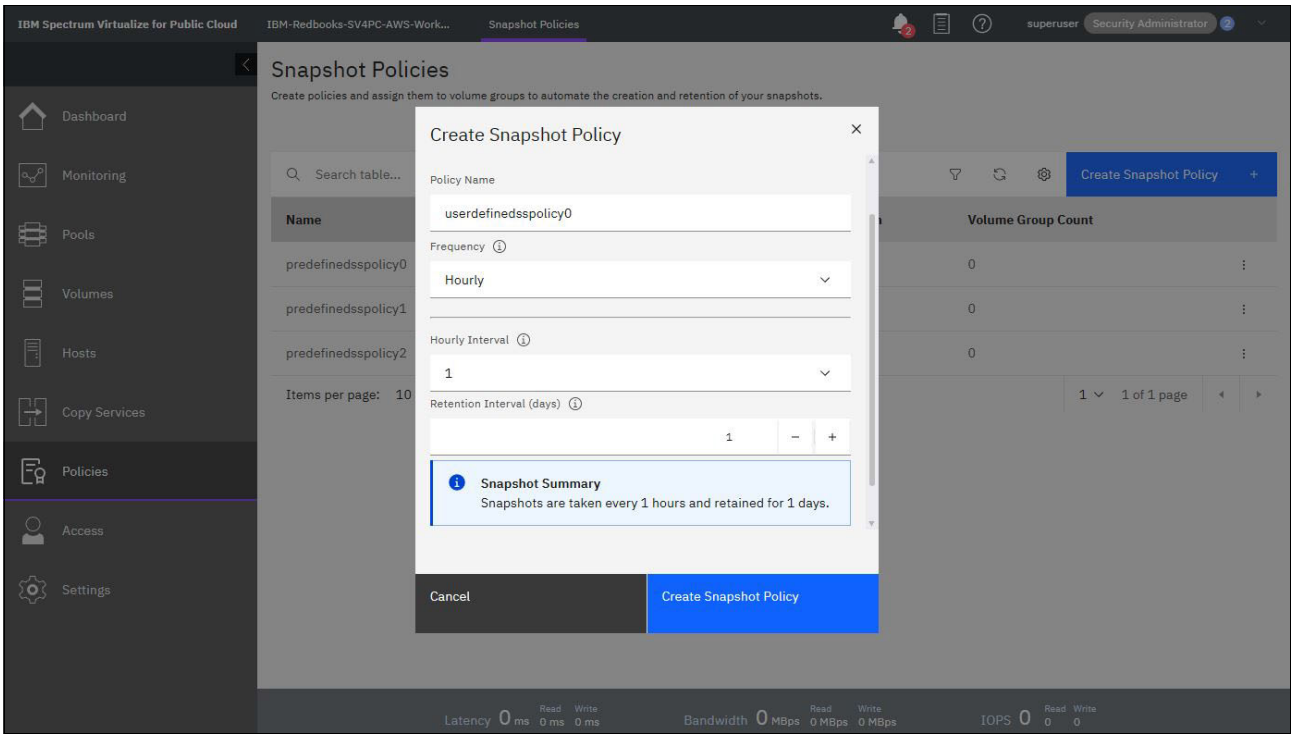

Figure 7-3 Creating a user-defined snapshot policy

4. Assign the internal snapshot policy to the VG by selecting **Volumes** → **Volume Group**, and click the VG to which the policy is to be assigned. For our example, VG `production_volume_group` is selected, as shown in Figure 7-1 on page 194.
5. Go to the Policies tab and click **Assign internal snapshot policy**. A window opens and lists all the predefined policies and user-defined policies. Select any one of the policies and select the **Safeguarded** checkbox. You can optionally specify the start date and time. Click **Assign Policy** to assign the selected policy to the VG.

   In our example, we selected user-defined policy `userdefinedsspolicy0`, as shown in Figure 7-4.

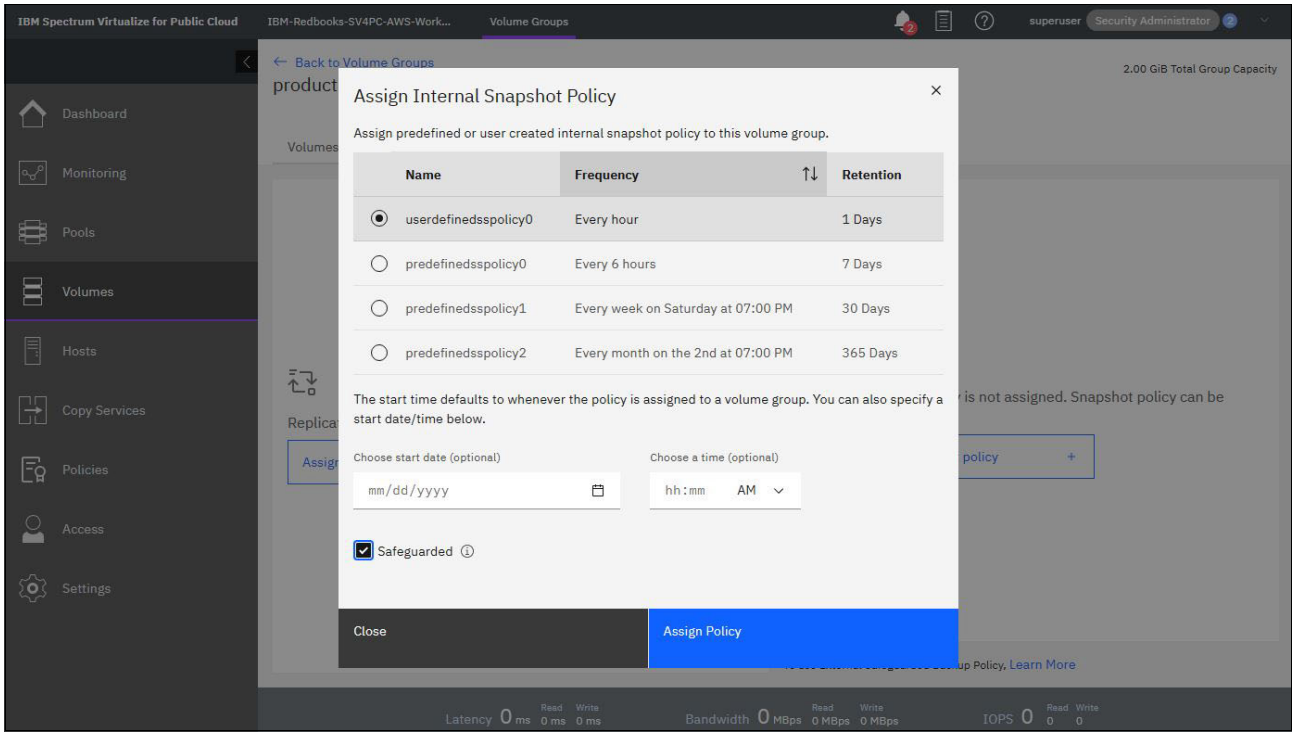

Figure 7-4 Assigning an internal snapshot policy

6. The **Policies** tab shows the Safeguarded Snapshot Policy that is associated with the VG, the retention days for the snapshots, and the next snapshot creation time, as shown in Figure 7-5 on page 197.

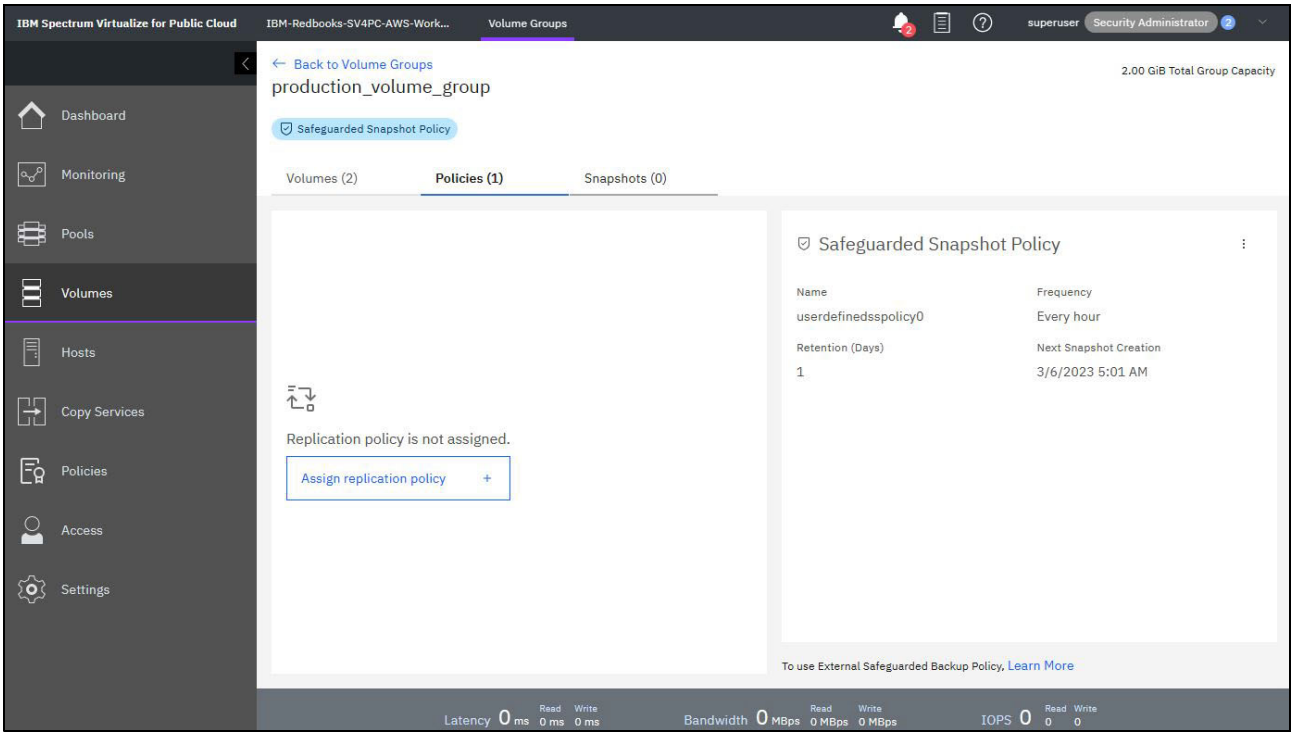

Figure 7-5 Safeguarded Snapshot Policy

7. Figure 7-6 shows a Safeguarded snapshot that is generated automatically by the assigned policy.

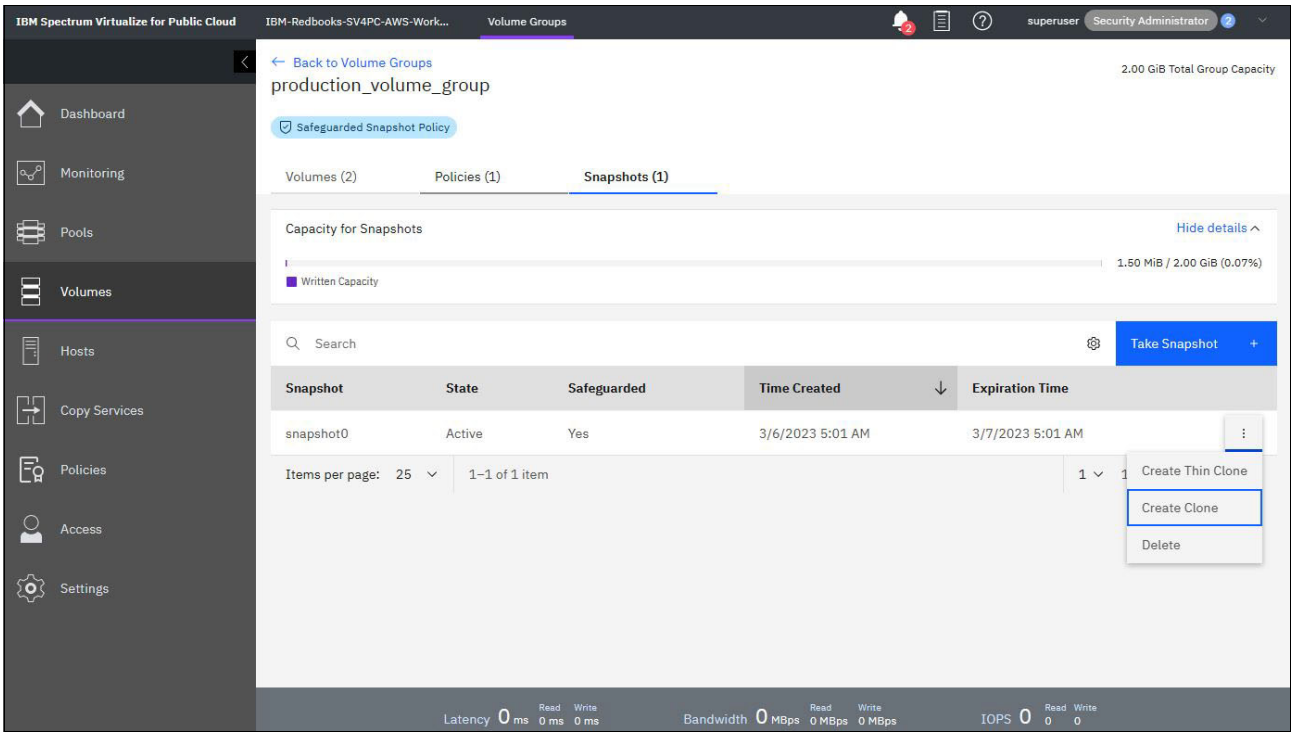

Figure 7-6 Snapshot creation by the assigned policy

8. To recover the Safeguarded snapshot, click either **Create Thin Clone** or **Create Clone**, as shown in Figure 7-7.

   Figure 7-7 shows the recovery of a Safeguarded snapshot by using the **Create Clone** option. For our example, we named the new VG as `recovery_volume_group`.

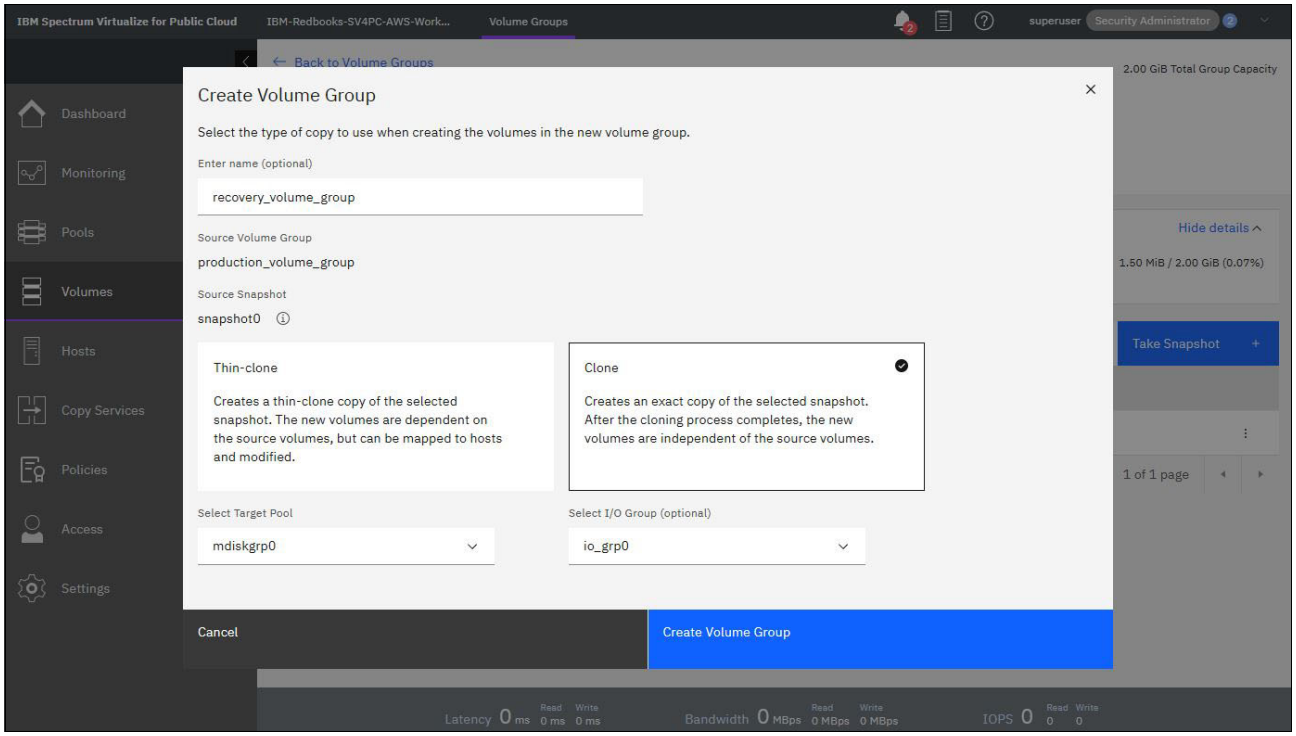

Figure 7-7 Creating a clone

9. After clicking **Create Volume Group** in Figure 7-7, a new VG is created. For our example, a new VG `recovery_volume_group` with two volumes (`production_source1-0` and `production_source2-0`) are created, as shown in Figure 7-8 on page 199.

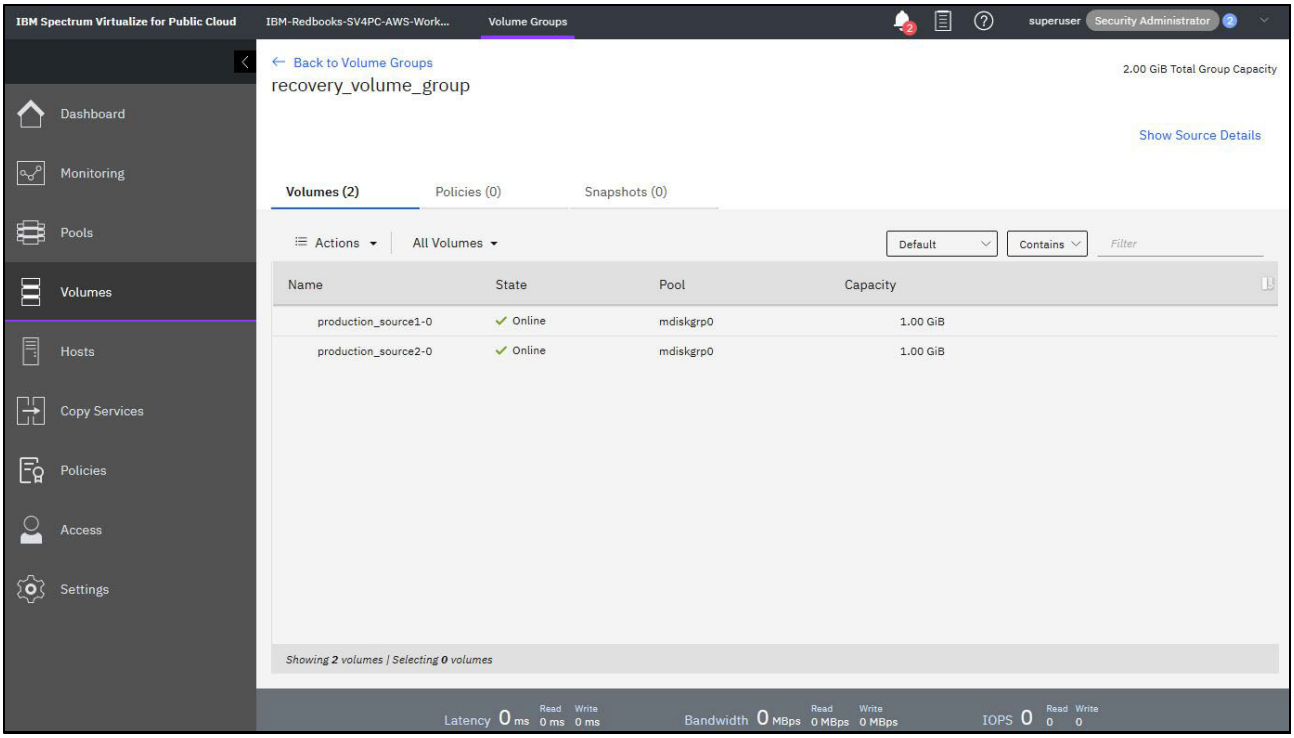

Figure 7-8 New volume group after a recovery

Safeguarded snapshots automatically are deleted according to the retention period that is specified in the assigned policy. Alternatively, it can be deleted by a superuser or Security Administrator by using the **Delete** option, as shown in Figure 7-6 on page 197.

7.2 Configuring secured Policy-Based Replication

This section describes how to configure secured PBR:

- Between AWS and AWS cloud
- Between Azure and Azure cloud
- Between AWS or Azure cloud and on-premises solutions like IBM FlashSystem or IBM SAN Volume Controller system

7.2.1 Pre-configuration requirements

This section discusses the pre-configuration requirements for implementing secured PBR.

Pre-configuration steps that are required for AWS to AWS secured policy-based replication

If the production and recovery clusters are in different virtual private clouds (VPCs), then VPC peering is required to setup PBR.

Complete the following steps:

1. To configure VPC peering, go to the AWS console, find the VPC of the production cluster, and create a peering connection with the recovery cluster VPC, as shown in Figure 7-9.

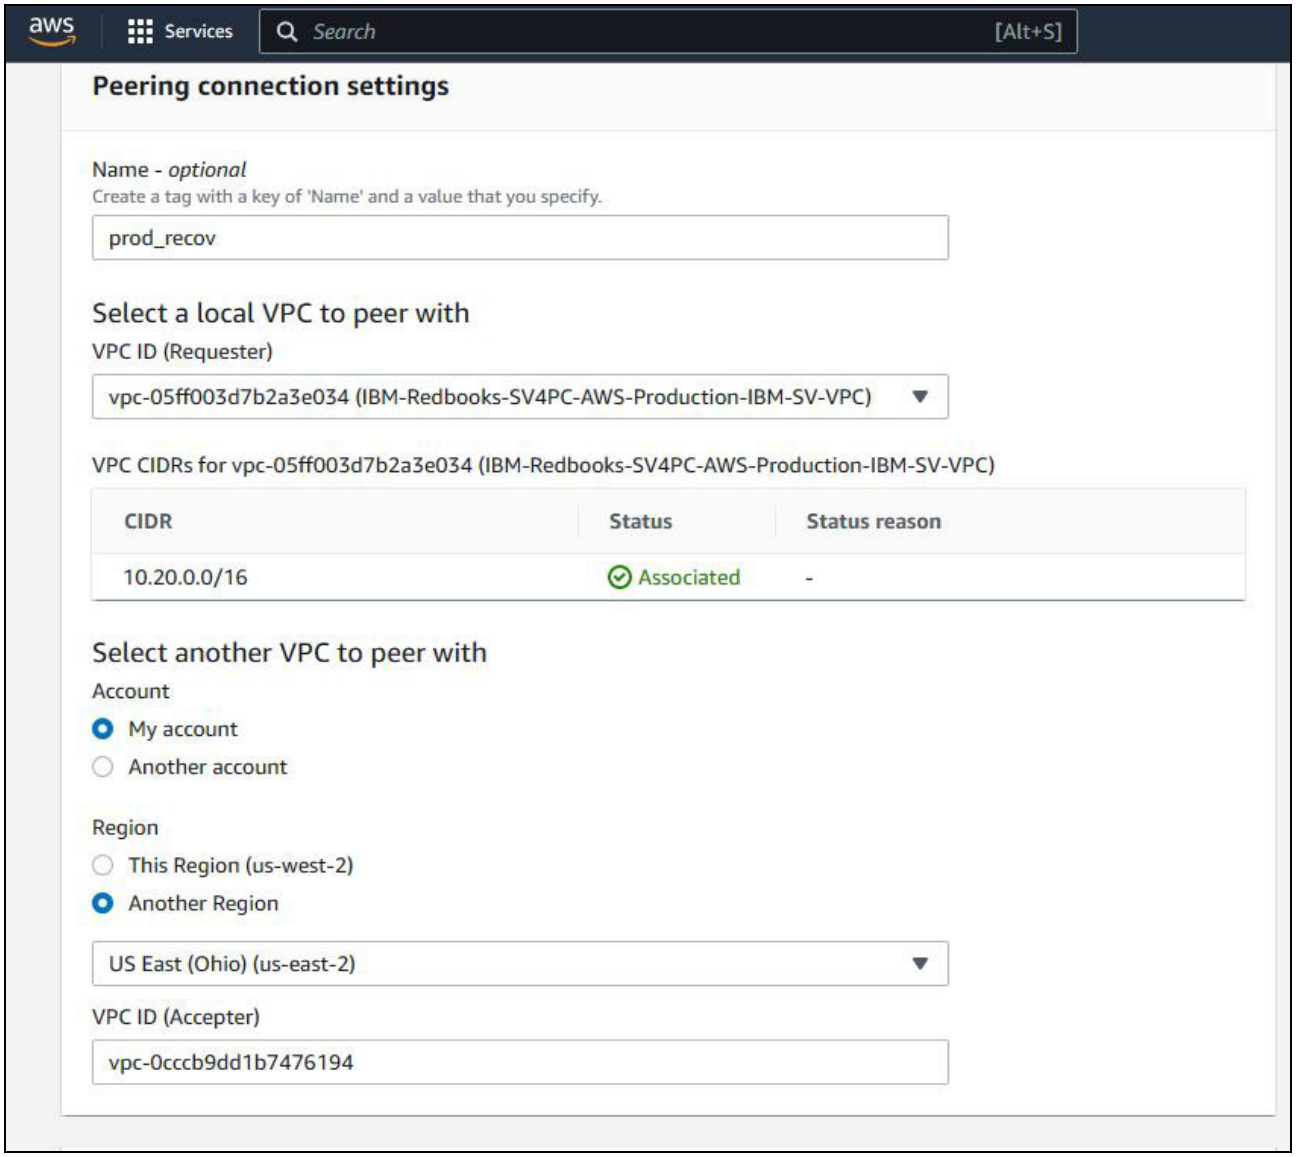

Figure 7-9 AWS: create VPC peering

2. Go to the recovery cluster region Peering connections list and accept the peering connection request, as shown in Figure 7-10 on page 201.

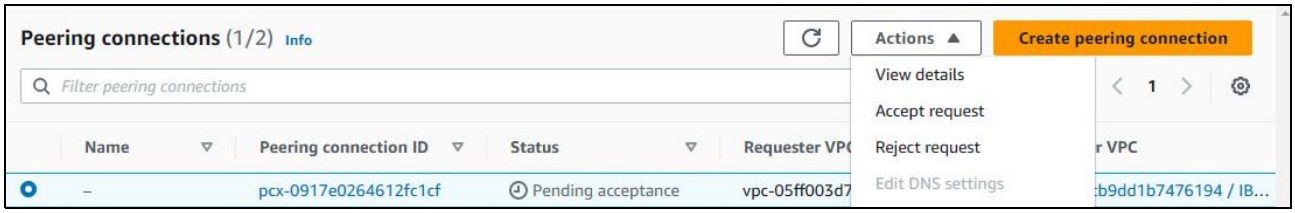

Figure 7-10 AWS: accept VPC peering

Tip: Ensure that the two VPCs do not have an overlapping CIDR range or the peering connection will not be created.

3. After peering is set up, modify the private route table that is associated with IBM Spectrum Virtualize nodes subnet of the production cluster on the AWS console. Add the route to the recovery cluster VPC CIDR range by using the VPC peering connection in the route table, as shown in Figure 7-11. Perform a similar step at the recovery cluster route table.

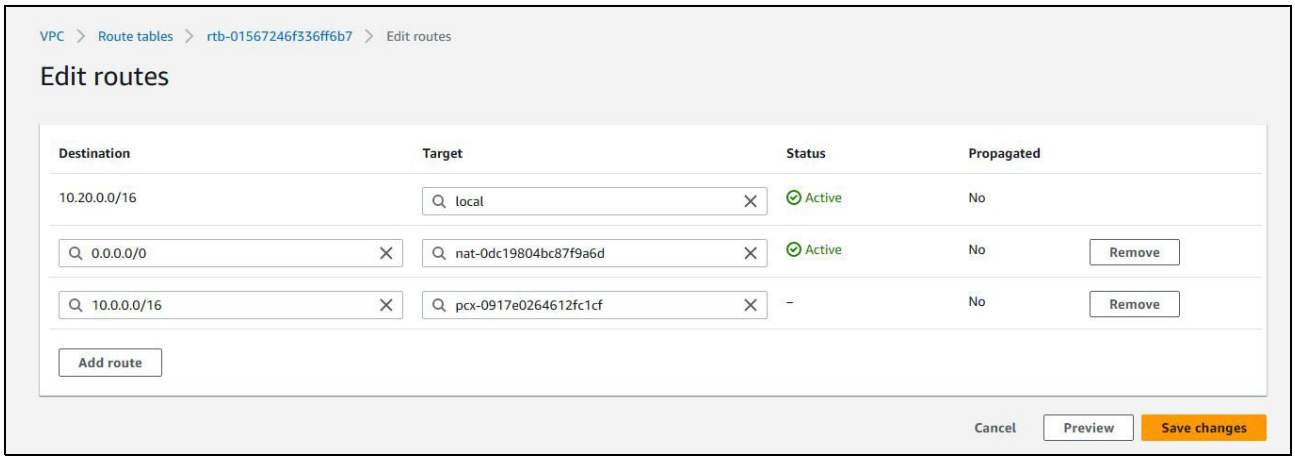

Figure 7-11 Editing routes to the route table

4. After the route table is modified, the production cluster can ping the recovery cluster by using the **`svctask ping`** command if *all ICMP-IPv4* protocol traffic is allowed in the inbound rules of the IBM Spectrum Virtualize node security group. Example 7-1 shows a production cluster pinging a recovery cluster by using a replication port IP address that is configured according to step 2 on page 206.

Example 7-1 Sample svctask ping command

```
#> svctask ping -srcip4 10.20.102.2  10.0.119.174
PING 10.0.119.174 (10.0.119.174) from 10.20.102.2 : 56(84) bytes of data.
64 bytes from 10.0.119.174: icmp_seq=1 ttl=64 time=50.7 ms
64 bytes from 10.0.119.174: icmp_seq=2 ttl=64 time=50.8 ms
64 bytes from 10.0.119.174: icmp_seq=3 ttl=64 time=50.7 ms
64 bytes from 10.0.119.174: icmp_seq=4 ttl=64 time=50.7 ms
64 bytes from 10.0.119.174: icmp_seq=5 ttl=64 time=50.7 ms

--- 10.0.119.174 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4007ms
rtt min/avg/max/mdev = 50.704/50.736/50.769/0.247 ms
```

If the production and recovery cluster are in different VPCs, then specific ports must be opened to set up PBR. For more information about opening ports, see this IBM Documentation web page. Open TCP ports 3260, 3265, 7443, and 443, and UDP ports 500 and 4500 in the inbound rules of IBM Spectrum Virtualize node security group of both the clusters for configuring a secured PBR.

Pre-configuration that is required for Azure to Azure secured policy-based replication

Complete the following steps:

1. Go to the production cluster resource group (RG) on the Azure console and search for the virtual network (vNet) resource. Click it and select **Peerings** under **Settings** on the left of the page. Click **Add** and enter the names of the local and remote peering links, as shown in Figure 7-12. Select **remote vNet** and click **Add**.

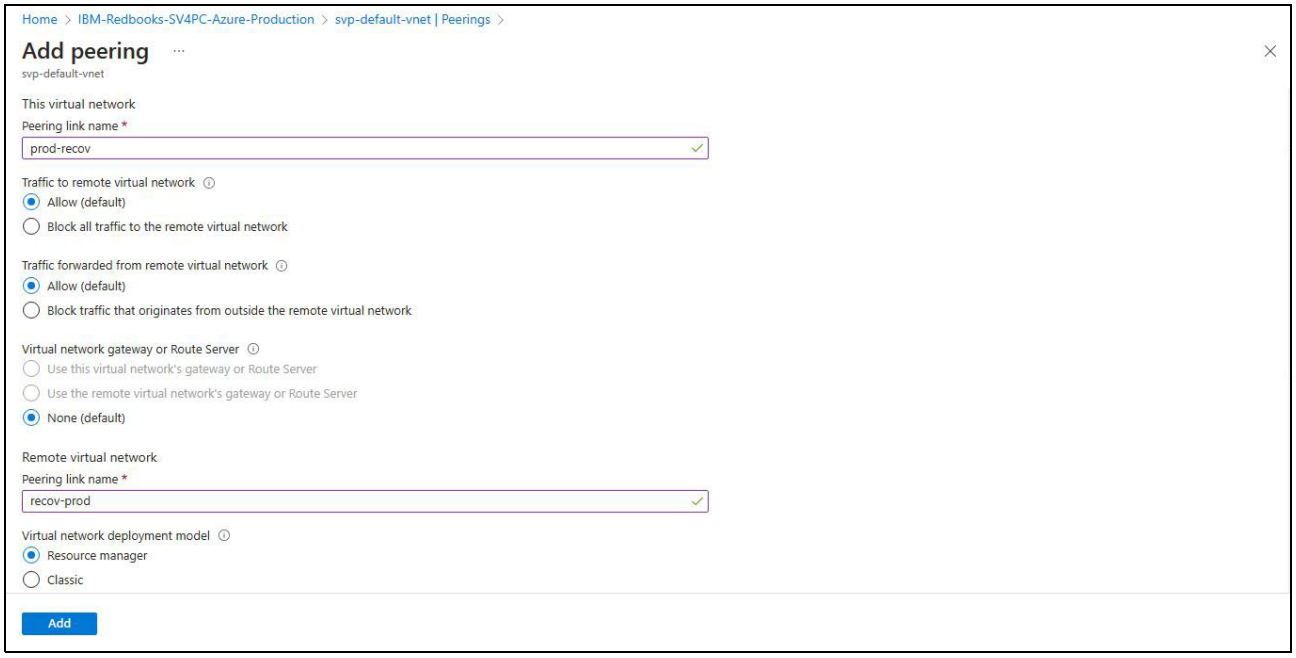

Figure 7-12 Azure Virtual Network peering

2. Now, if the network security group (NSG) of the IBM Spectrum Virtualize cluster subnet allows inbound ICMP protocol packets, then both clusters can ping each other.

   For IBM Spectrum Virtualize for Public Cloud on Azure 8.5.2, first update the subject alternative name in the certificate before creating the IP partnership between clusters. To do this task, select **Settings** → **Security** → **Secure Communications**, click **Update Certificate**, and update the Subject Alternative Name field with a local cluster IP address, for example `IP:10.13.0.4`, as shown in Figure 7-13 on page 203. Click **Update**.

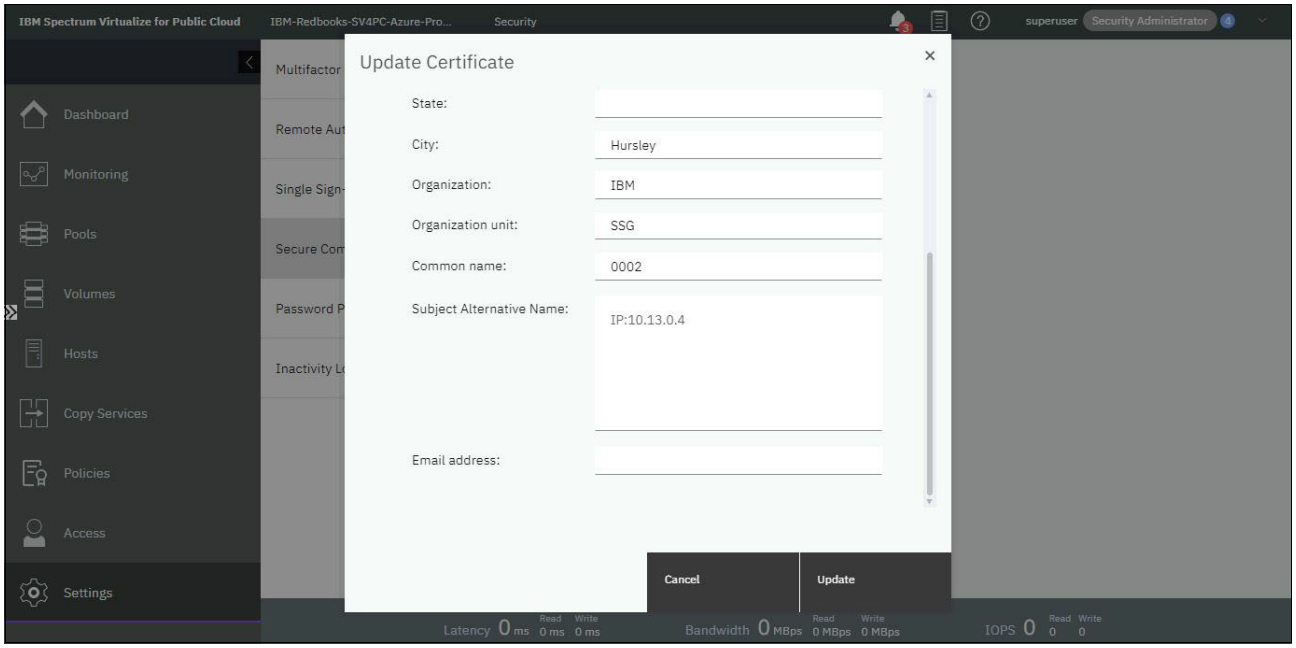

Figure 7-13 Update Certificate

> **Note:** When the certificate update occurs, the web browser loses connection to the system before the process completes. To proceed, close the browser and wait approximately 2 minutes before accessing the management GUI.
>
> **Note:** After the certificate update, reinstall the IP quorum application.

Here is a sample procedure for reinstalling the IP quorum application on the default quorum server that is deployed during Azure Marketplace installation:

1. In the management GUI, select **Settings** → **System** → **IP Quorum** and download the version of the IP quorum Java application. You also can use the command-line interface (CLI) to enter the **mkquorumapp** command to generate an IP quorum Java application. The application is stored in the `dumps` directory of the system with the file name `ip_quorum.jar`.

   ```
   IBM_IBM Spectrum_Virtualize:IBM-Redbooks:superuser>mkquorumapp
   ```

2. Transfer the IP quorum application from the system to a separate directory on the server or host that is to run the IP quorum application by running the following command:

   ```
   # scp -i key.pem  superuser@cluster:/dumps/ip_quorum.jar
   ```

3. Replace the existing IP quorum application on the host by running the following command:

   ```
   # cp ip_quorum.jar /usr/local/bin/ip_quorum.jar
   ```

4. Restart the quorum service by running the following command:

   ```
   #  systemctl restart quorum.service
   ```

Pre-configuration that is required from AWS or Azure to on-premises secured policy-based replication

IP network connectivity between the on-premises data center and AWS or Azure cloud must be established through a virtual private network (VPN) connection before configuring secured PBR. This section describes how to configure hybrid cloud connectivity between the AWS or Azure cloud and on-premises solutions like IBM FlashSystem or IBM SAN Volume Controller.

AWS configuration for the VPN IPsec tunnel

This section describes the steps to configure the site-to-site IPsec tunnel for communication between AWS Cloud and the on-premises site. The VPN IPsec site-to-site tunnel creates a secure communication network between the AWS Cloud infrastructure and on-premises infrastructure. Network communication between the private subnets is controlled by the access control list (ACL) that is populated when you create the VPN IPsec site-to-site tunnel.

To establish the IPsec tunnel, complete the following steps at the VPC level in AWS Cloud:

1. Create a customer gateway by logging in to the AWS console with resource provisioning privileges. Select **Services** at the upper left, and then select **VPC**. Select **Virtual Private Network (VPN)** in the pane at the left. Then, click the customer gateways and enter the required details.
2. Create the virtual private gateways by clicking **Virtual private gateways** in the VPC and configuring the required details.
3. Attach a virtual private gateway to the VPC.
4. Create a site-to-site VPN connection in AWS Console by selecting the virtual private gateway and customer gateway parameters. Attach the virtual private gateway to the VPC in AWS.
5. After the site-to-site connection is complete, a sample configuration file can be downloaded. This step creates two tunnels in the VPC. The same configuration file is used for the configuration at the other end of the tunnel.
6. After completing these steps, modify the private route table that is associated with the IBM Spectrum Virtualize nodes subnet of the cloud cluster on the AWS console. Add a route to the on-premises cluster IP range by using the virtual private gateway in the route table, as shown in Figure 7-14.

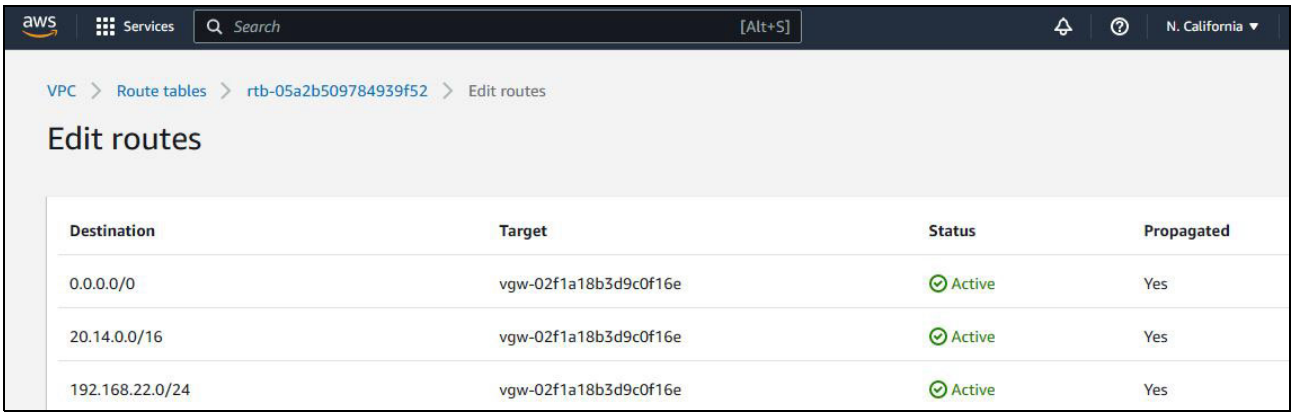

Figure 7-14 Edit routes

7. Modify the IBM Spectrum Virtualize node NSG to allow inbound traffic to the required ports from the remote on-premises cluster, as set in step 4 on page 201.

Azure configuration for a VPN IPsec tunnel

This section describes the steps that are required at the vNet level in Azure cloud for establishing the IPsec tunnel. For more information, see this Azure tutorial web page.

Complete the following steps:

1. Create a virtual private network gateway (VPNGW):
   a. Log in to the Azure console with administrator privileges.
   b. Select **Create a Resource**, and then search for “Virtual Network Gateway”.
   c. Enter the required information for the Virtual Network Gateway and associate it with the vNet in Azure to be used for hybrid connectivity.
2. Create a local network gateway:
   a. Log in to the Azure console with administrator privileges.
   b. Select **Create a Resource**, and then search for “Local Network Gateway”.
   c. Enter the required information for the Local Network Gateway for hybrid connectivity.
3. Create a VPN device.

   A site-to-site connection requires a VPN device for connection to the on-premises setup. Follow the Azure documentation to create the VPN device for your on-premises setup.
4. Create a VPN connection:
   a. Select the Virtual Network Gateway that was created in step 1.
   b. Select **Connections**, and then click **Add** to create a connection.
   c. Enter the local gateway that was created in step 2.

   When this process is complete, a VPN connection is established between your on-premises environment and the Azure cloud network.
5. Verify the connection by connecting to a virtual machine (VM) in cloud or on-premises.

   For IBM Spectrum Virtualize for Public Cloud on Azure 8.5.2, the certificate must be updated, as mentioned in step 2 on page 202.

7.2.2 Secured policy-based replication steps

For our example, we created one production cluster in the AWS Oregon region and one recovery cluster in the AWS Ohio region. The deployment type is public for both clusters.

To configure the PBR, complete the following steps:

1. Configure your IBM Spectrum Virtualize Private IP network ports so that they are enabled for remote copy. This configuration is required on both sites, as shown in Figure 7-15.

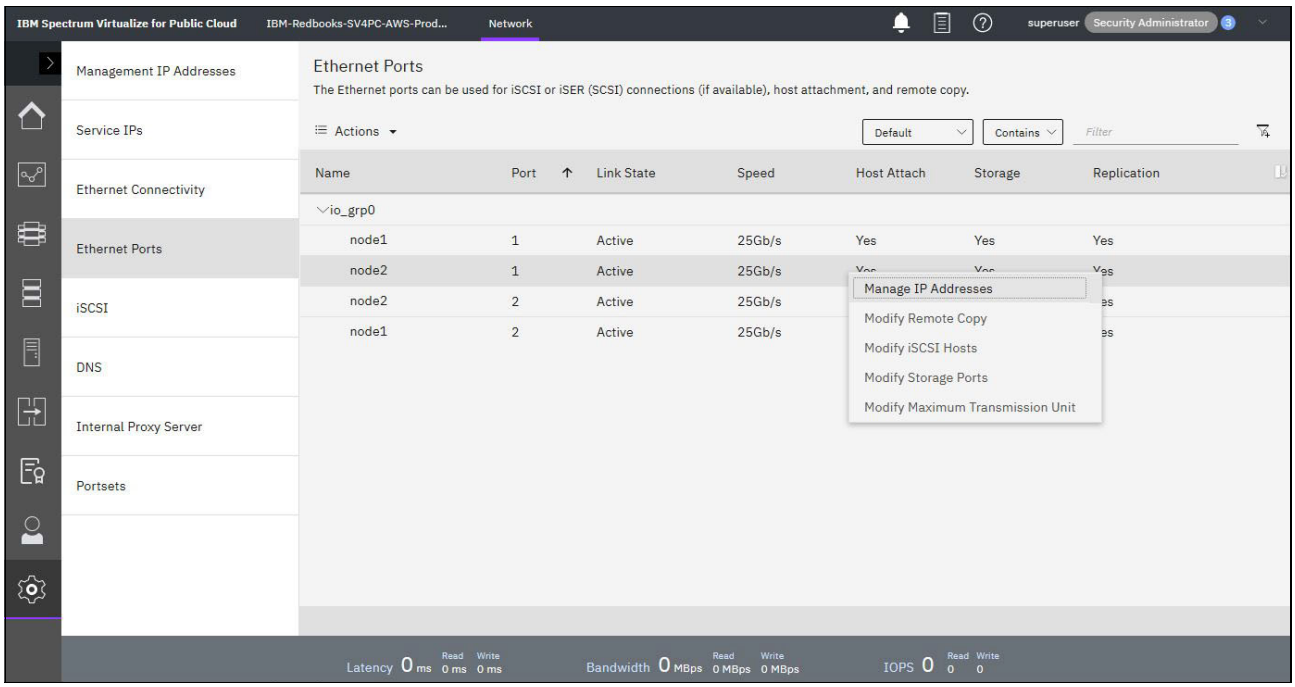

Figure 7-15 Enabling remote copy

2. Enable the IP portset for remote copy. When you use the new feature of multiple portsets and multiple IP addresses, you must enable the IP address for the portset. To do so, use the CLI to enable the IP address, as shown in Example 7-2. The share IP address option is used in this example because multiple IP assignments on the same Ethernet port by using the **mkip** command are not supported on IBM Spectrum Virtualize for Public Cloud. This process must be done on at both clusters to set up the IP address partnership.

> **Note:** The share IP address option is not available on the IBM Spectrum Virtualize for Public Cloud GUI.

Example 7-2 Using the same IP address for the replication portset

```
#> svctask mkip -gw 10.21.0.1 -ip 10.21.0.21 -node 1 -port 2 -portset 2 -prefix 24 -shareip
IP address, id [10], successfully
created
```

3. On the production cluster, select **Copy Services** → **Partnerships and Remote Copy**. Click **Create Partnership** and select **2-site partnership**, and then click **Continue**. Select **IP** for the partnership type. Enter a cluster IP address for recovery cluster into the Partner IP Address field and click **Test Connection**. Select the **Secured IP partnerships** and **Use Policy-Based Replication** check boxes, as shown in Figure 7-16.

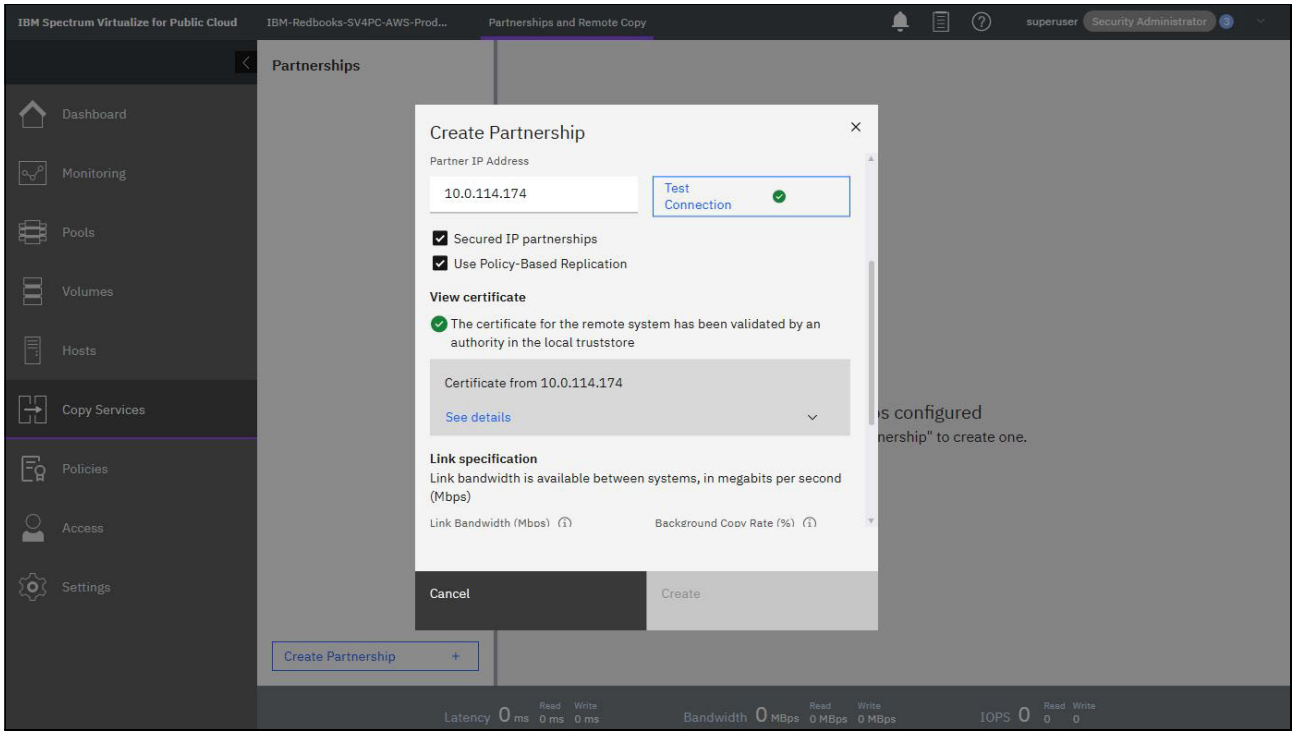

Figure 7-16 Create Partnership

4. Enter values into the Link Bandwidth and Background Copy Rate fields. To enable compression on the partnership link, select **Compression Enabled**. Click **Create**.

   This process creates a partial partnership. A replication checklist, as shown in Figure 7-17, lists the remaining steps to complete configuring PBR.

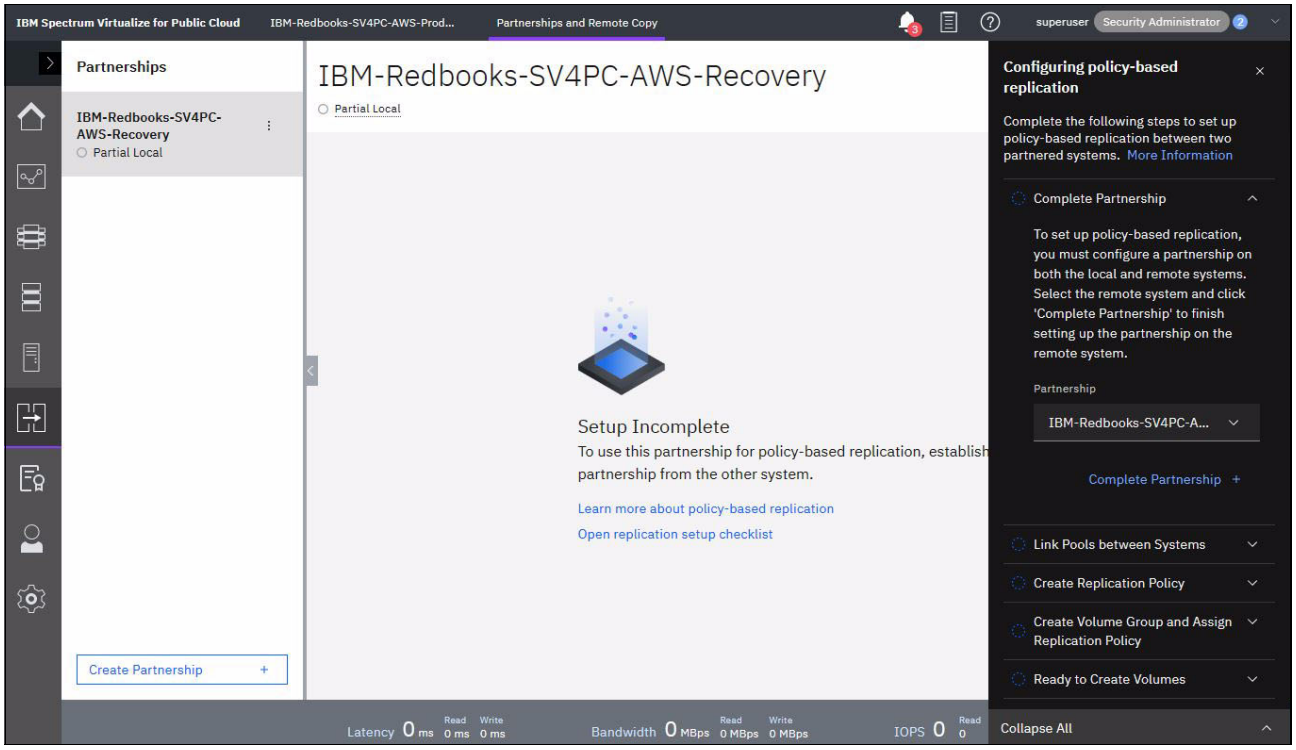

Figure 7-17 Partial partnership with a replication checklist

5. To complete the partnership, click **Complete Partnership** in the replication checklist if your browser can access the internal cluster IP address of the recovery cluster. In our example, we manually opened the recovery cluster GUI by using a quorum server public IP address and repeated the same partnership creation steps.

   Now, the partnership is created, as shown in Figure 7-18 on page 209.

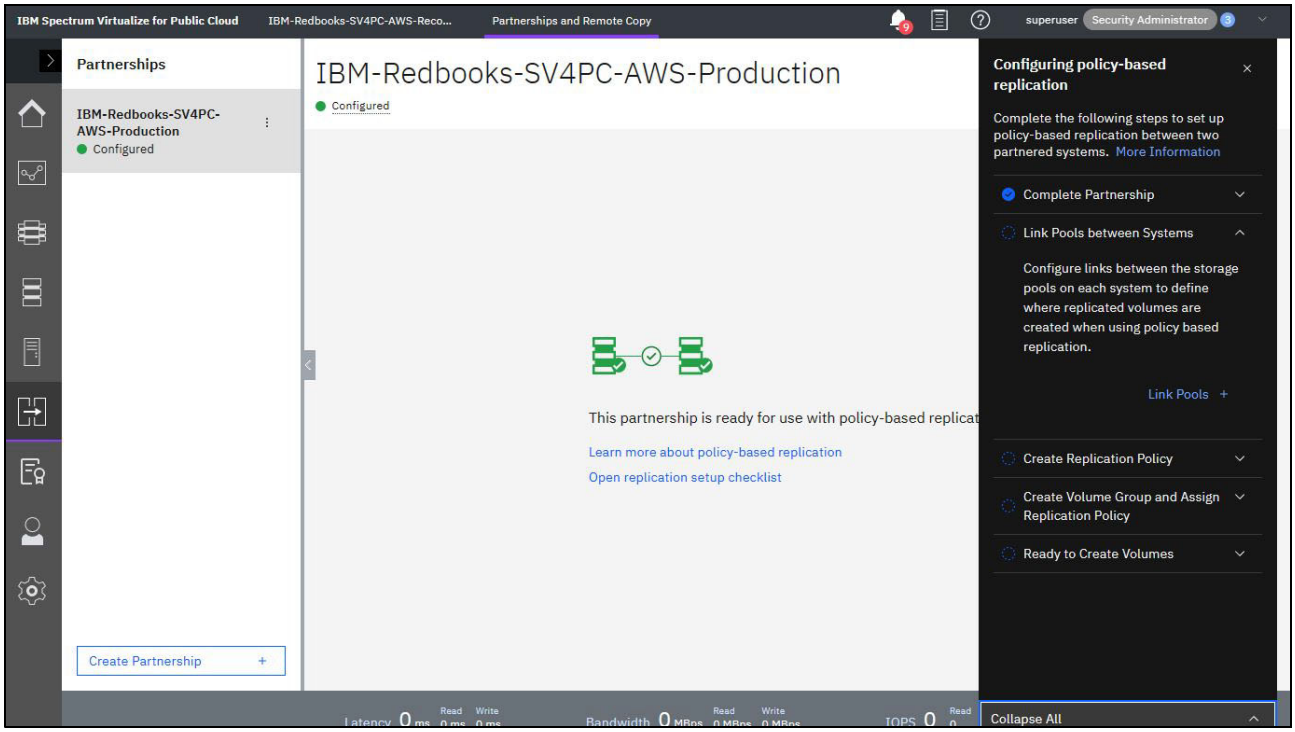

Figure 7-18 Fully configured partnership with replication checklist

You can either continue to use the GUI on this system to configure pool links between the systems or you can close this tab and return to the GUI of the first system.

Linking pools between systems

The next item in the replication setup checklist is linking the pools between systems. Before you link the pools between the systems, assign a provisioning policy to the pool to be linked to on the recovery cluster GUI by completing the following steps:

1. Go to **Pools** and right-click the pool to be linked and select **Assign Provisioning Policy**. For our example, we selected `capacity_optimized provisioning` policy, as shown in Figure 7-19.

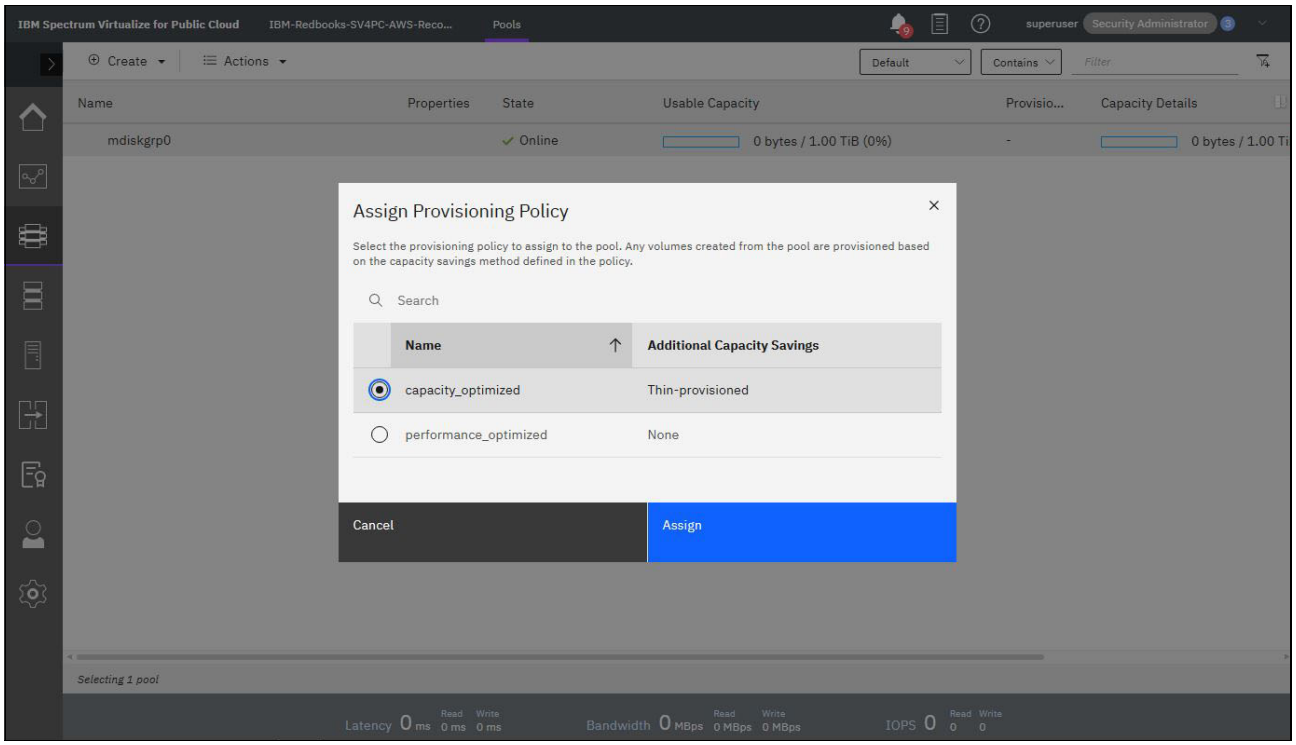

Figure 7-19 Assigning a provisioning policy

2. For our example, we link the pools by using the production cluster GUI. You can either link the pools by right-clicking the pool and selecting **Add Pool Link for Replication**, or by clicking **Link Pools** on the replication setup checklist.
3. Select the local pool and local provisioning policy, and the remote pool. Click **Apply**, as shown in Figure 7-20 on page 211.

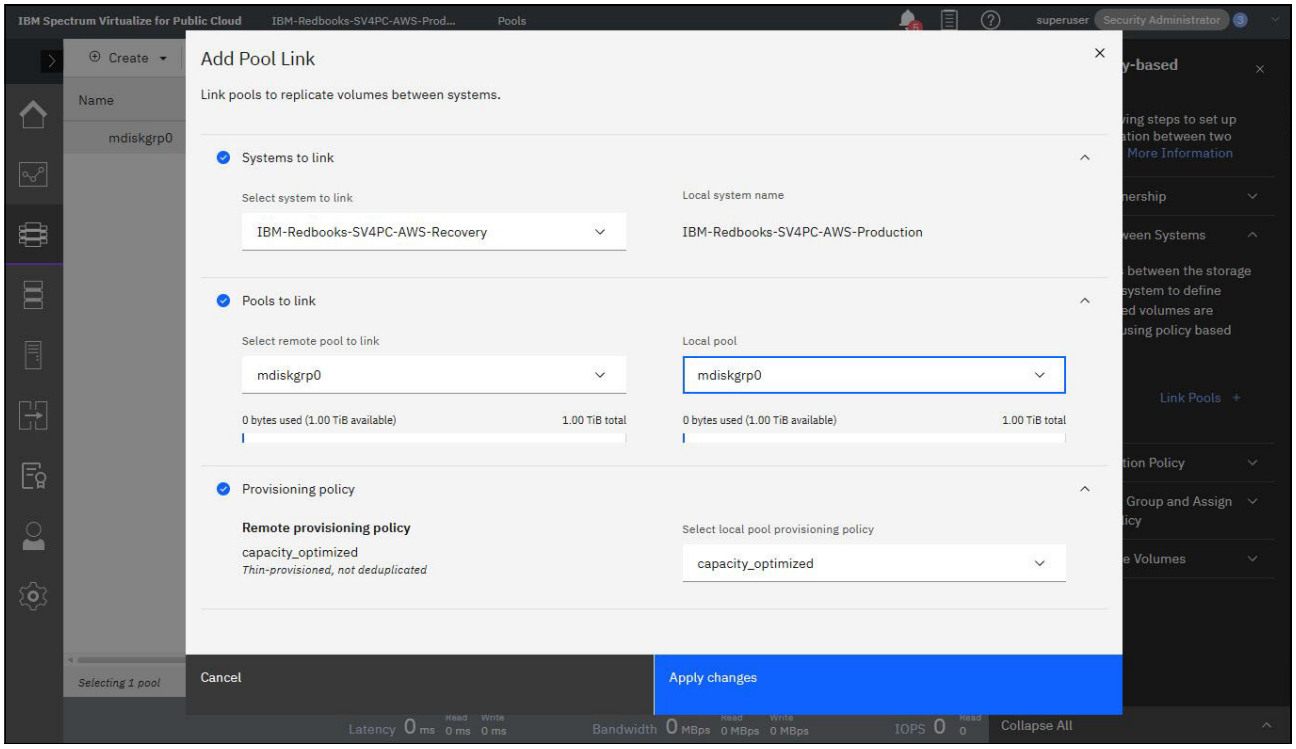

Figure 7-20 Creating a pool link

The pool linking step does *not* need to be repeated on the remote cluster.

The next step on the checklist is to create a replication policy. Complete the following steps:

1. Select **Yes** to use the local system as the production system, as shown in Figure 7-21 on page 212. Click **Create Replication Policy**.

2. Enter the name of the replication policy. Select location1 and location2, and specify the recovery point objective (RPO) for the policy in minutes, as shown in Figure 7-21. Click **Create**. This action automatically creates the policy both in the local and remote systems.

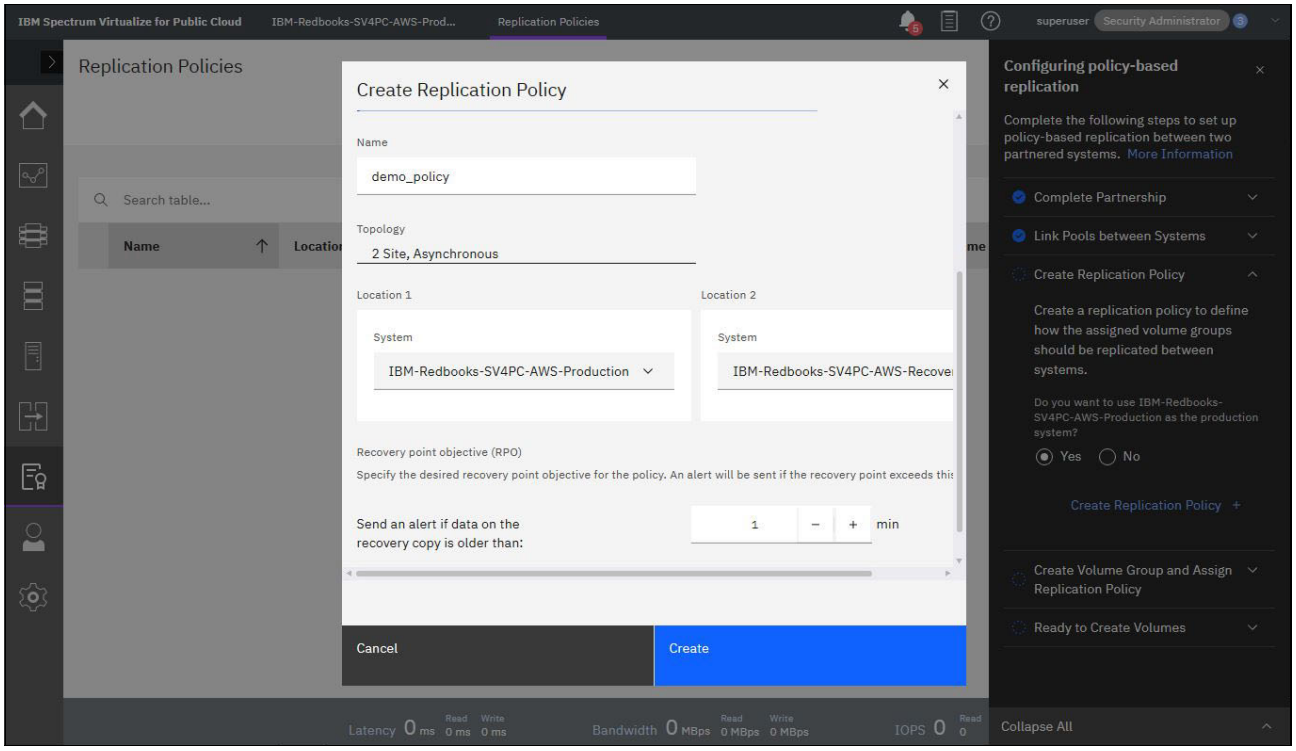

Figure 7-21 Creating a replication policy

3. Click **Create Volume Group and Assign Replication Policy** in the replication checklist. Enter the name of the VG to create, as shown in Figure 7-22 on page 213. Click **Create Empty Group**.

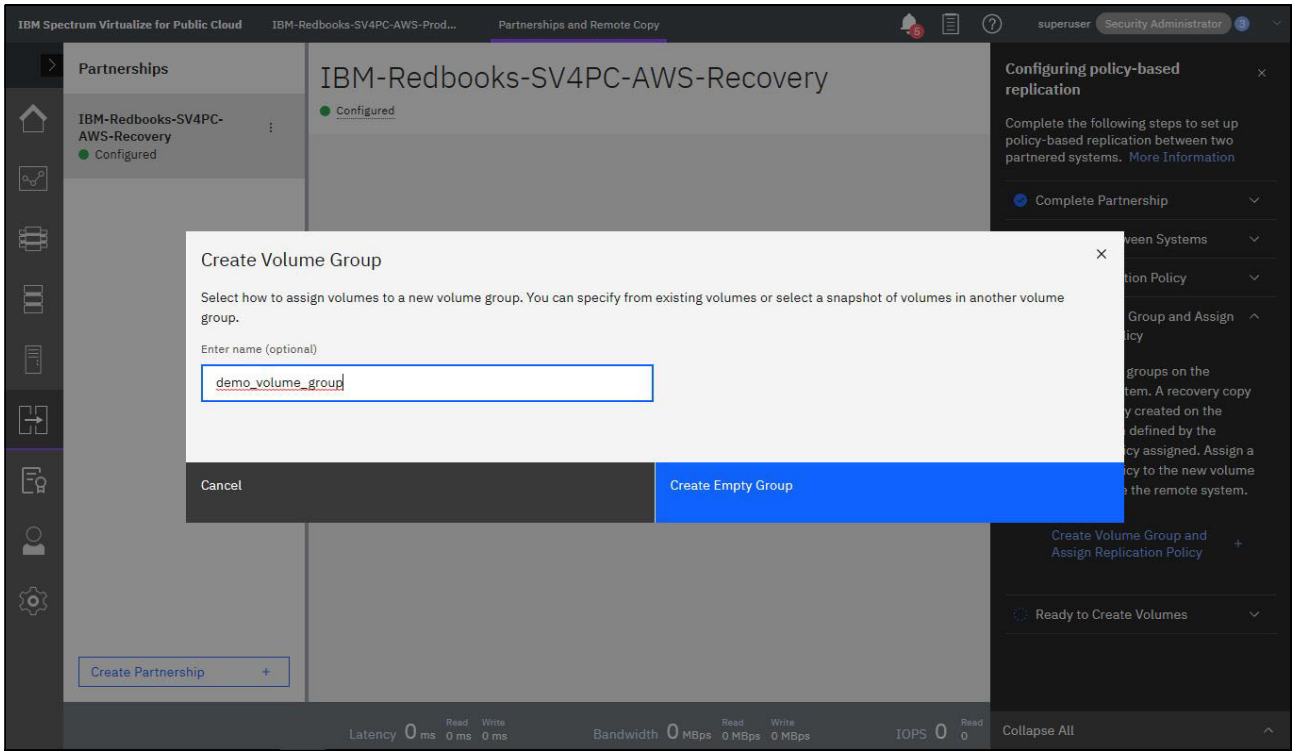

Figure 7-22 Creating a volume group

4. Select the replication policy to assign to the newly created VG, as shown in Figure 7-23. Click **Assign**.

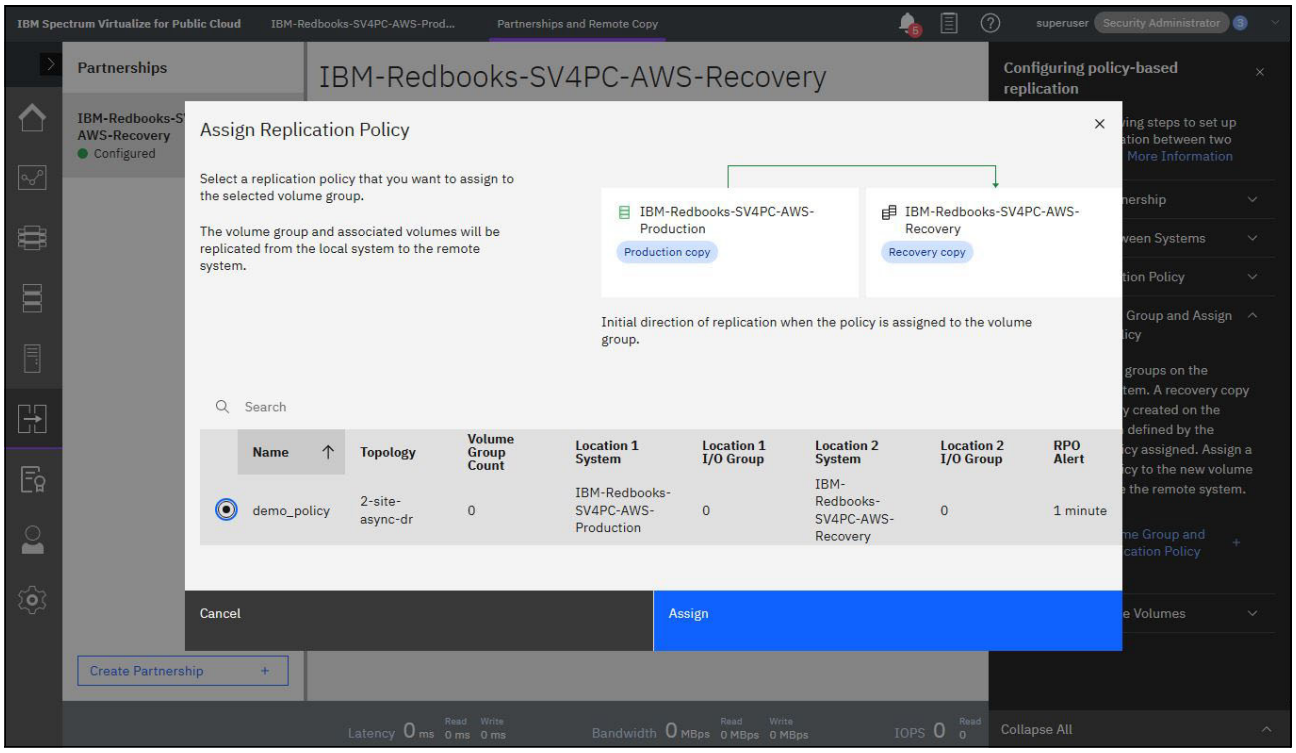

Figure 7-23 Assigning a replication policy

Figure 7-24 shows the new VG with the replication policy that is assigned to it. A replication status arrow shows that replication is running from the production system to the recovery cluster system.

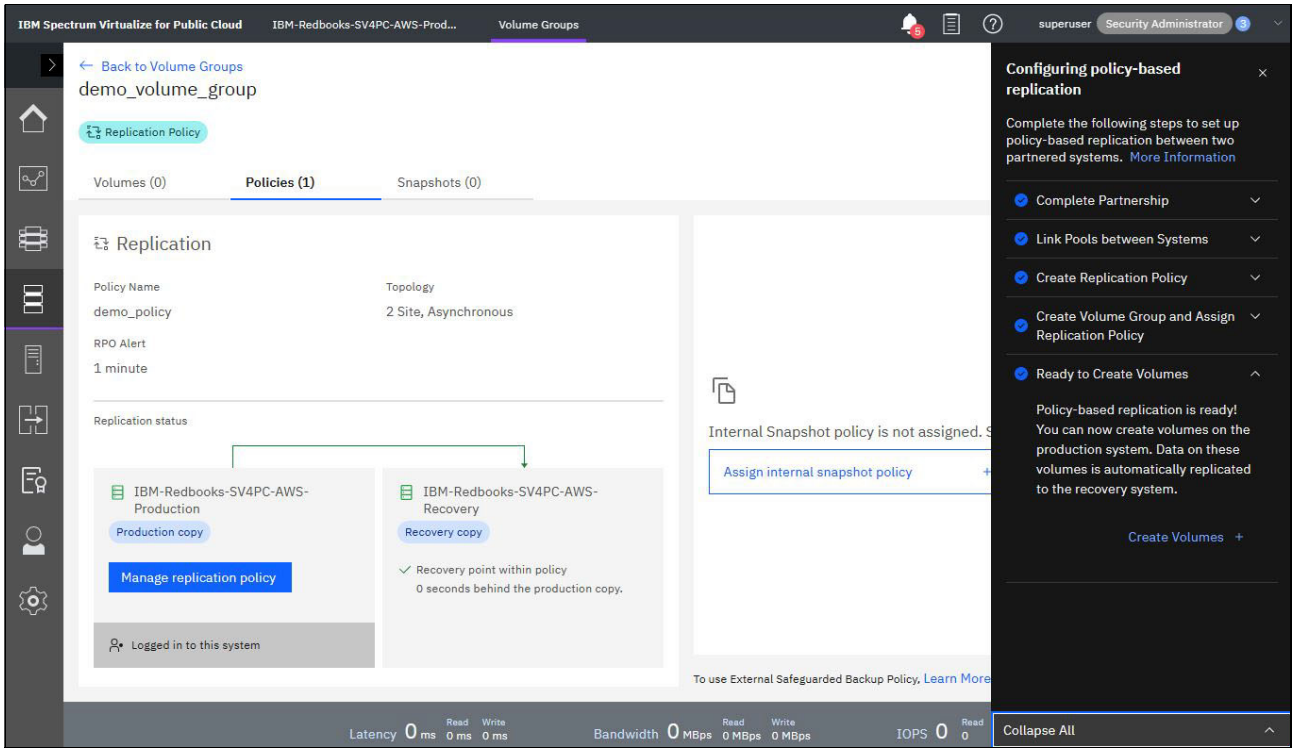

Figure 7-24 Volume group with an assigned replication policy

The checklist in Figure 7-24 now shows that PBR is ready for use. Volumes can be created in the VG, and data on these volumes is automatically replicated to the recovery system.

Complete the following steps:

1. Select **Create Volumes** → **Create Volumes**. Select the VG and pool and click **Define Volume Properties**. Enter a volume name and a value for capacity and click **Save**. Click **Create volumes**, as shown in Figure 7-25 on page 215.

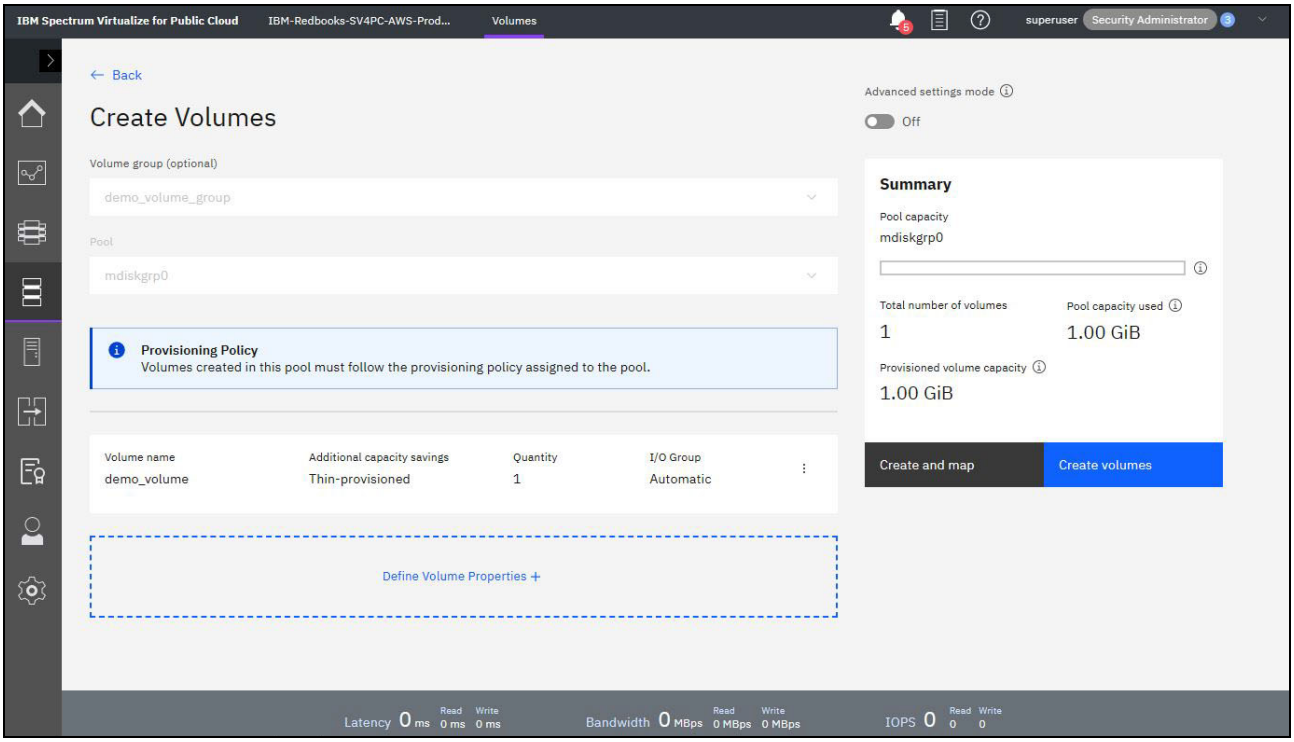

Figure 7-25 Create Volumes

Creating volumes directly in a VG optimizes the initial synchronization, even if the systems are disconnected. When you create volumes in the production copy of a VG, the system automatically creates the recovery copies of those volumes and configures replication according to the assigned replication policy.

2. To view the replication status of a VG, select **Volumes** → **Volumes Groups**. Click the VG name. The replication status pane displays the replication mode of each system and information about the recovery point, as shown in Figure 7-26.

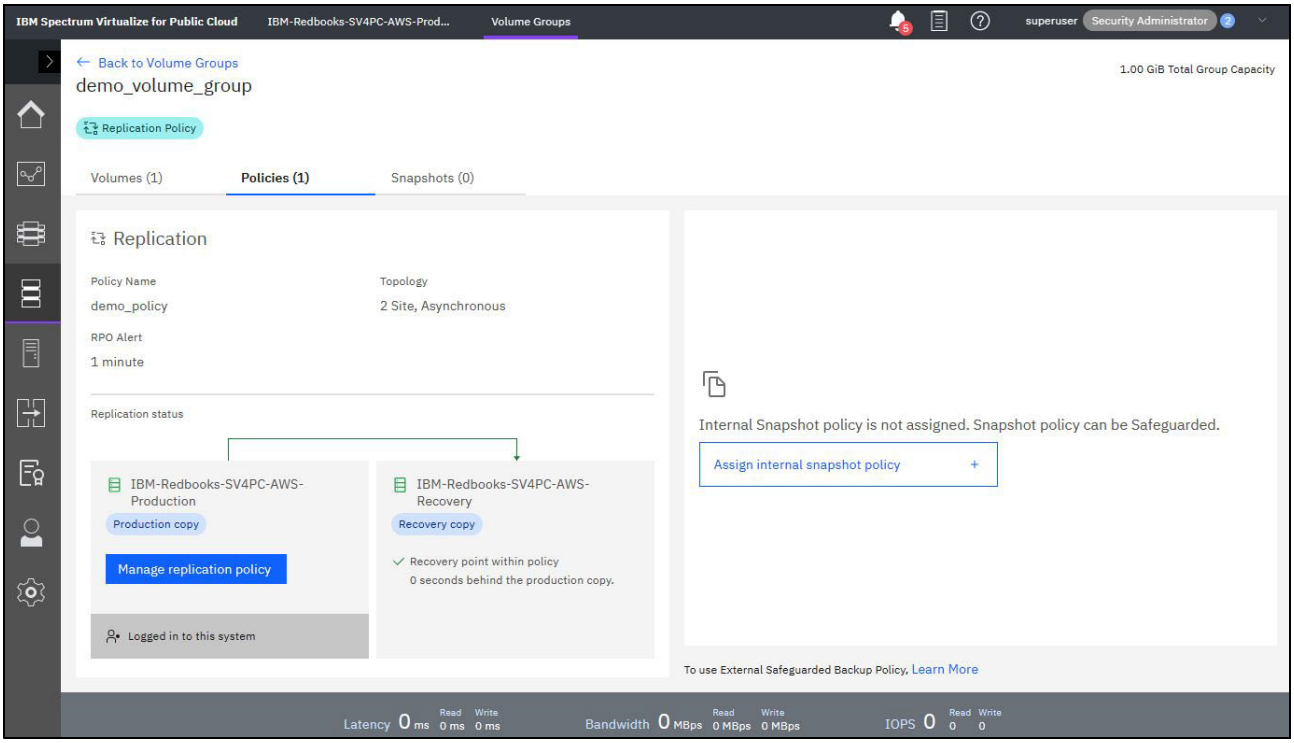

Figure 7-26 Replication status

For more information about PBR, see *Policy-Based Replication with IBM Storage FlashSystem, IBM SAN Volume Controller and IBM Storage Virtualize*, REDP-5704.

8

Monitoring and supporting the solution

This chapter provides guidance about supporting and monitoring IBM Spectrum Virtualize for Public Cloud on Azure and Amazon Web Services (AWS).

This chapter includes the following topics:

- Monitoring IBM Spectrum Virtualize for Public Cloud on Microsoft Azure or AWS through a GUI, IBM Spectrum Control, or IBM Storage Insights
- Call Home function and email notification
- Monitoring capacity reporting in IBM Spectrum Virtualize for Public Cloud on Microsoft Azure or AWS

8.1 Monitoring IBM Spectrum Virtualize for Public Cloud on Microsoft Azure or AWS through a GUI, IBM Spectrum Control, or IBM Storage Insights

This section describes the following procedures:

- Setting up Call Home and email notification
- Viewing capacity and performance in IBM Spectrum Virtualize GUI
- Viewing capacity and performance in Public Cloud GUI (Microsoft Azure and AWS)
- Monitoring IBM Spectrum Virtualize for Public Cloud on Azure or AWS in IBM Spectrum Control and IBM Storage Insights

8.2 Call Home function and email notification

The Call Home function of IBM Spectrum Virtualize uses the Cloud services and email services to the specific IBM Support center. Table 8-1 lists the supported configurations for Call Home.

Table 8-1 Supported network configurations for Call Home with Cloud services

Supported configuration	DNS configuration	Firewall requirements
Call Home with Cloud services with an internal proxy server	Required	Configure the firewall to allow outbound traffic on port 443 to `esupport.ibm.com`.
Call Home with Cloud services with a DNS server	Defined, but not required	Configure the firewall to allow outbound traffic on port 443 to `essuport.ibm.com`. Optionally, allow outbound traffic on port 443 to the following IP addresses: - 129.42.56.189 - 129.42.54.189 - 129.42.60.189
Call Home with Cloud services	None	Configure the firewall to allow outbound traffic on port 443 to the following IP addresses: - 129.42.56.189 - 129.42.54.189 - 129.42.60.189

Note: Call Home with Cloud services is the optimal transmission type for IBM Spectrum Virtualize for Public Cloud on Microsoft Azure or AWS.

To configure Call Home, complete the following steps. (The screen captures show all possible configuration options.)

1. From the left window of the GUI, select **Settings** → **Support** → **Call Home** (see Figure 8-1).

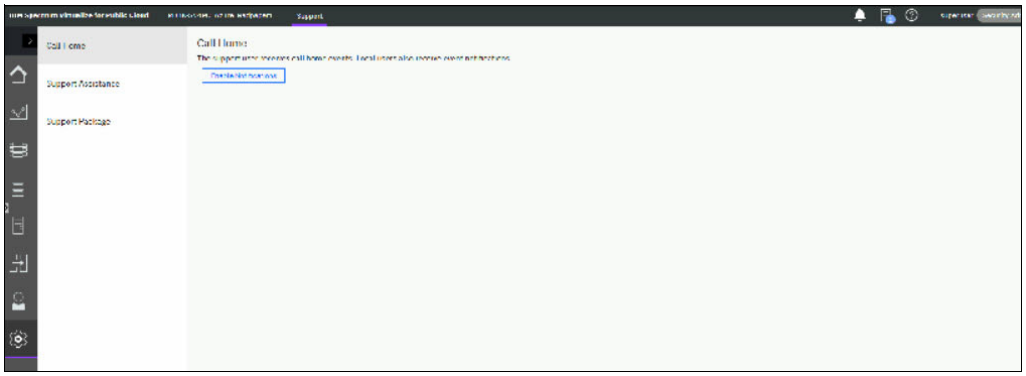

Figure 8-1 Call Home

2. Select **Enable Notifications** to start the Call Home configuration wizard (see Figure 8-2).

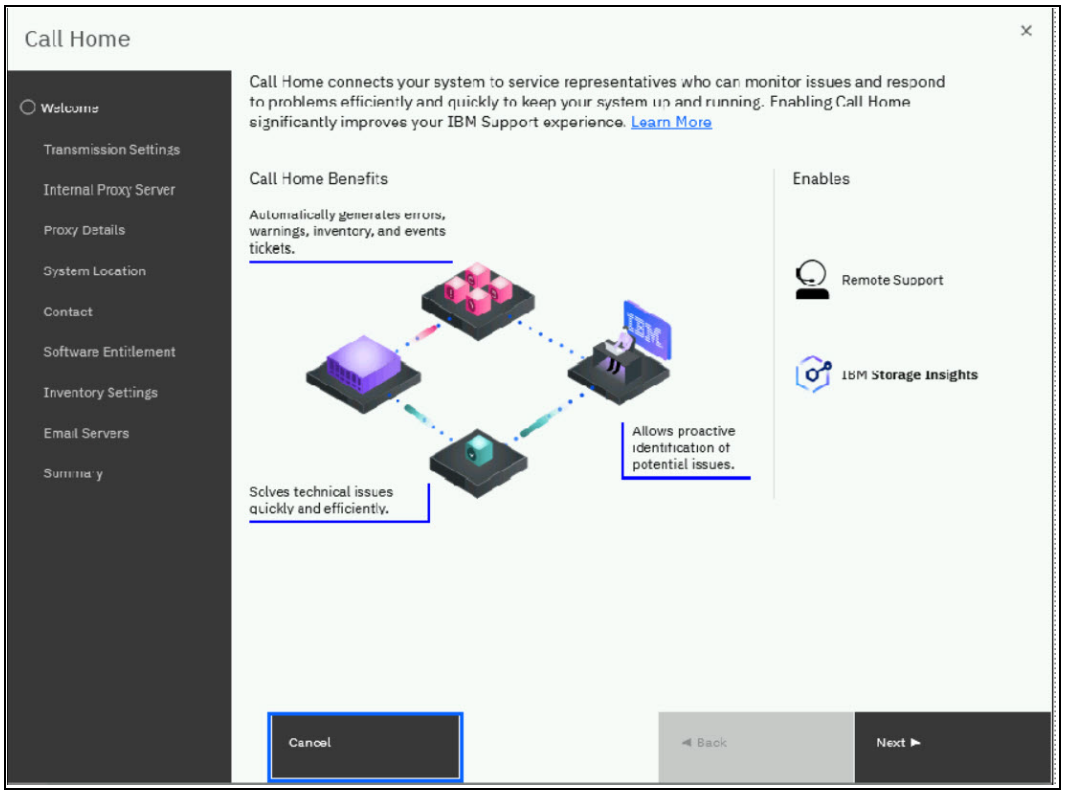

Figure 8-2 Call Home configuration wizard

3. After clicking **Next** in the welcome window, select the transmission type for Call Home. Select **Send using Cloud Services**, **Send using Email Services**, or both.

   Cloud services for Call Home are the optimal transmission method for error data because it ensures that notifications are delivered directly to the support center. Filters on email servers can prevent error notifications from arriving at the support center and delay error resolution (see Figure 8-3).

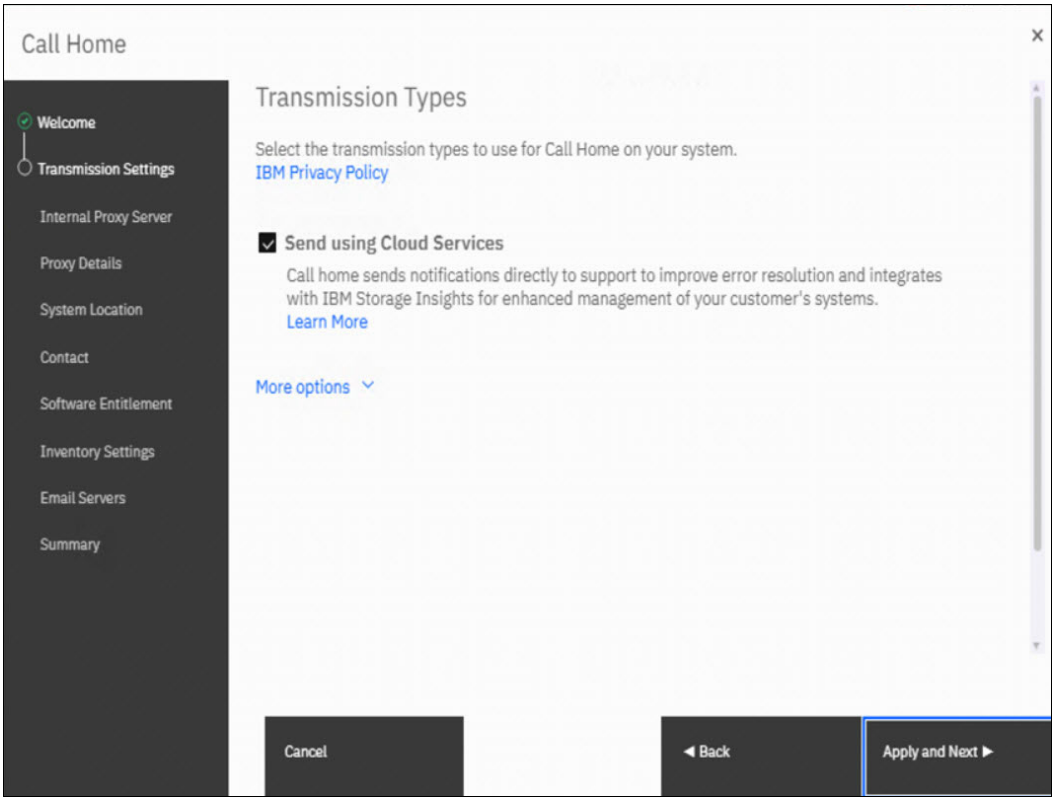

Figure 8-3 Selecting the Call Home transmission type

4. Configure an internal proxy server. You can choose between Open Proxy, Certificate, or Basic authentication. The proxy server also can be added later.

   A DNS server is required to use a proxy server for Call Home with Cloud services (see Figure 8-4).

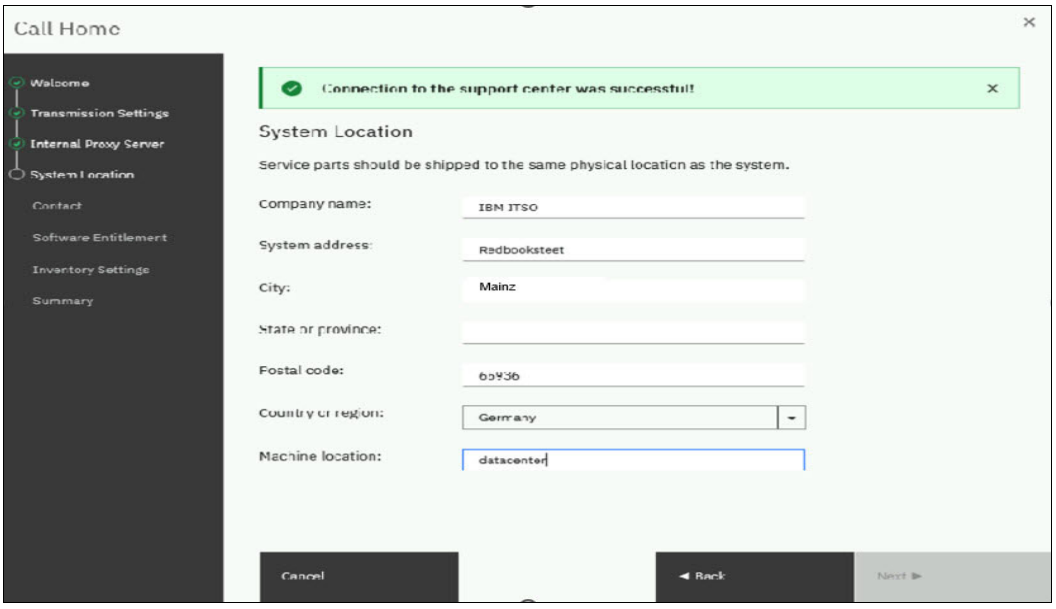

Figure 8-4 Configuring the proxy server

5. After clicking **Next** in the welcome window, enter the information about the location of the system, as shown in Figure 8-5.

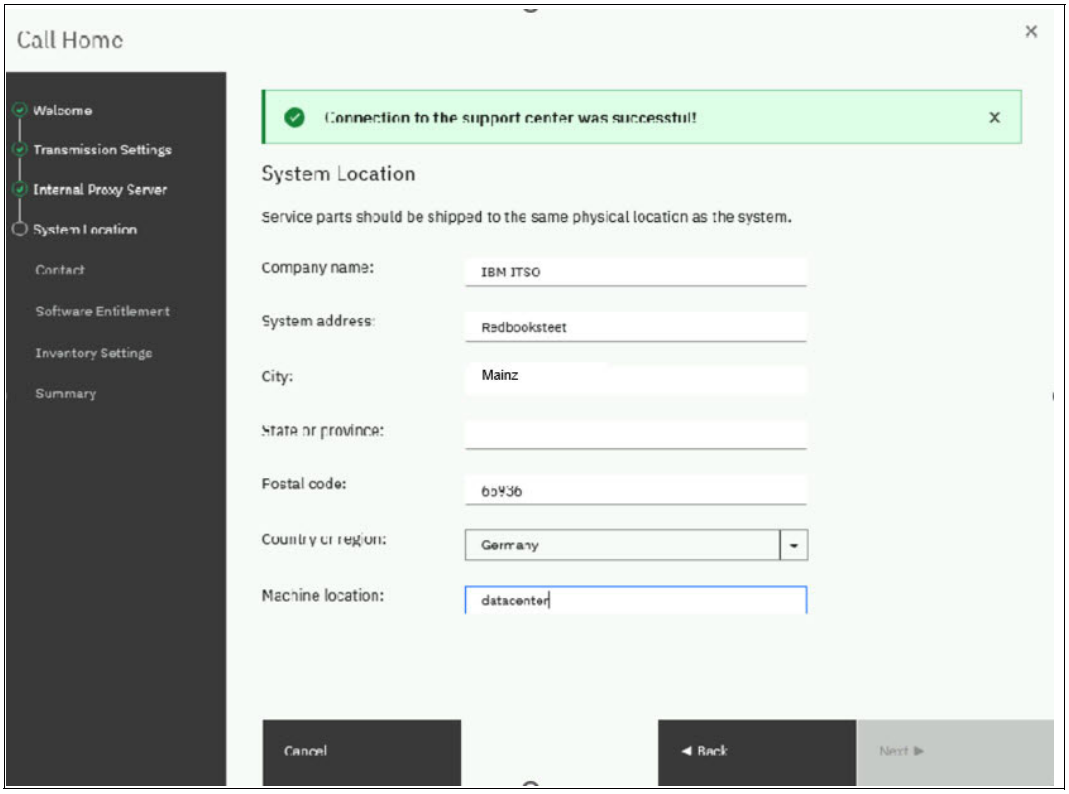

Figure 8-5 Location of the device

Figure 8-6 shows the contact information of the owner.

Tip: Add a specific contact name and number. Entering generic contact information might delay direct contact with the customer by the IBM Support Teams.

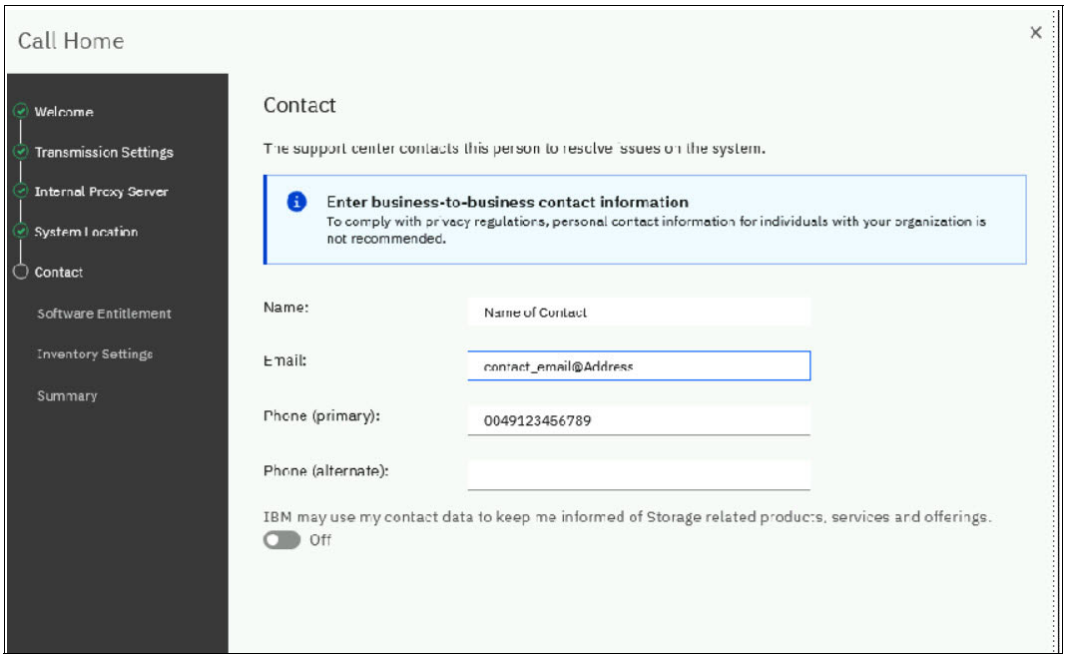

Figure 8-6 Contact information

6. In the next window of the wizard, enter the software entitlements details (see Figure 8-7).

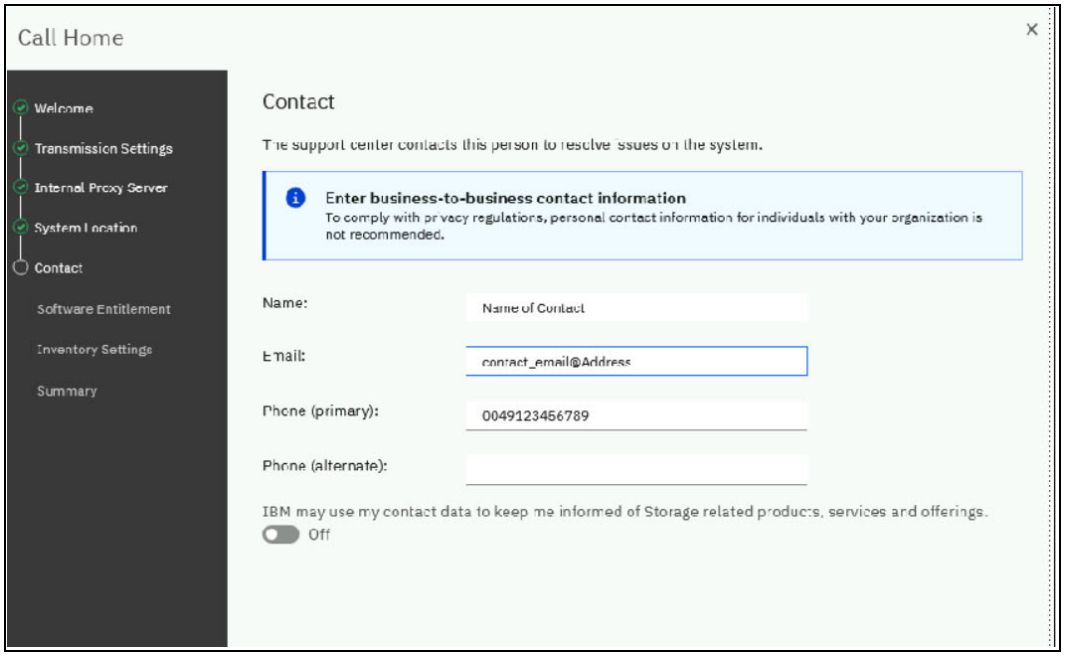

Figure 8-7 Software entitlement

7. If you specified only cloud service as the transmission type, the next window in the wizard contains only the configuration for Configuration Reporting. Set **Configuration Reporting** to `On` to generate enhanced reports.

   IBM Support uses this detailed information to automatically generate best practices that are specific to your configuration. You also can configure the inventory if the transmission type is `Email services`. Emails that are sent to the service center contain information about inventory.

   If both transmission types are used, inventory and configuration reporting can be enabled. (see Figure 8-8).

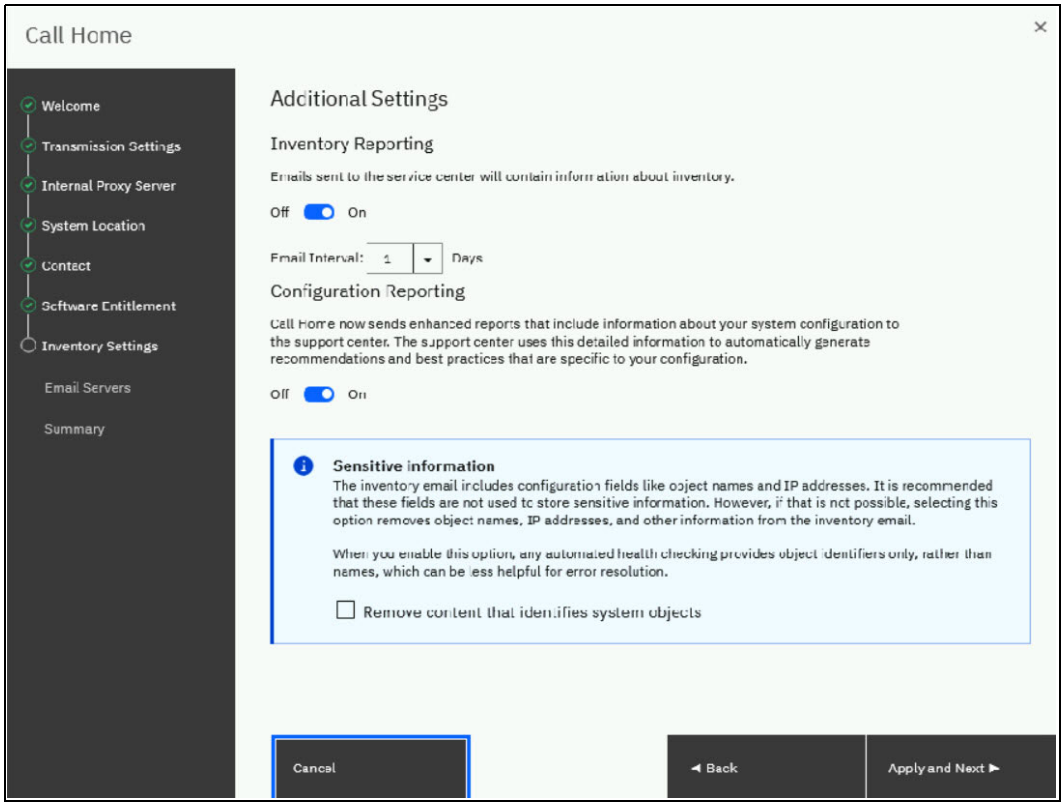

Figure 8-8 Additional Settings window

8. In the next window of the wizard, configure the email server. This window is part of the wizard only if the transmission type `Email services` was selected (see Figure 8-9).

> **Note:** The `Email services` transmission method is not recommended as the only way to send notifications to the support center. Use Call Home with email notifications as a backup method when Call Home with Cloud services is configured.
>
> Depending on your IBM Spectrum Virtualize for Public Cloud configuration, you might need to set up an SMTP service in your cloud portal (Azure or AWS).

Figure 8-9 Configuring email servers

9. The last window of the wizard shows a summary of the configuration. Verify the information, and then click **Finish**.
10. To test Call Home, left window of the GUI, select **Settings** → **Call Home** → **Test Support Notification** (see Figure 8-10 on page 225).

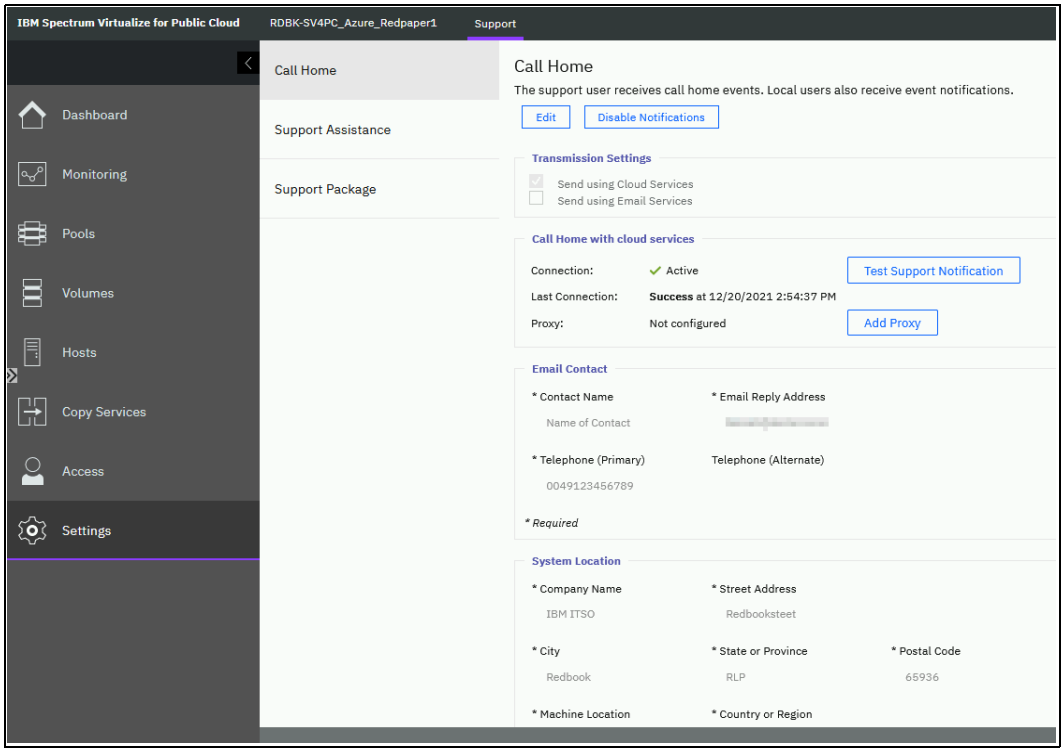

Figure 8-10 Test Support notification

8.2.1 Disabling and enabling notifications

At any time, you can temporarily or permanently disable notifications, as shown in Figure 8-11.

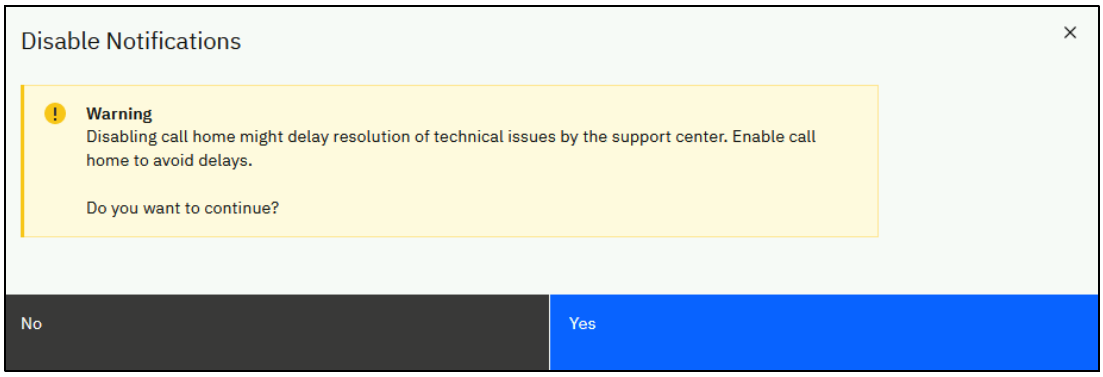

Figure 8-11 Disable Notifications

8.3 Monitoring capacity reporting in IBM Spectrum Virtualize for Public Cloud on Microsoft Azure or AWS

The Capacity section on the Dashboard provides an overall view of system capacity. This section displays usable capacity, provisioned capacity, and capacity savings (see Figure 8-12).

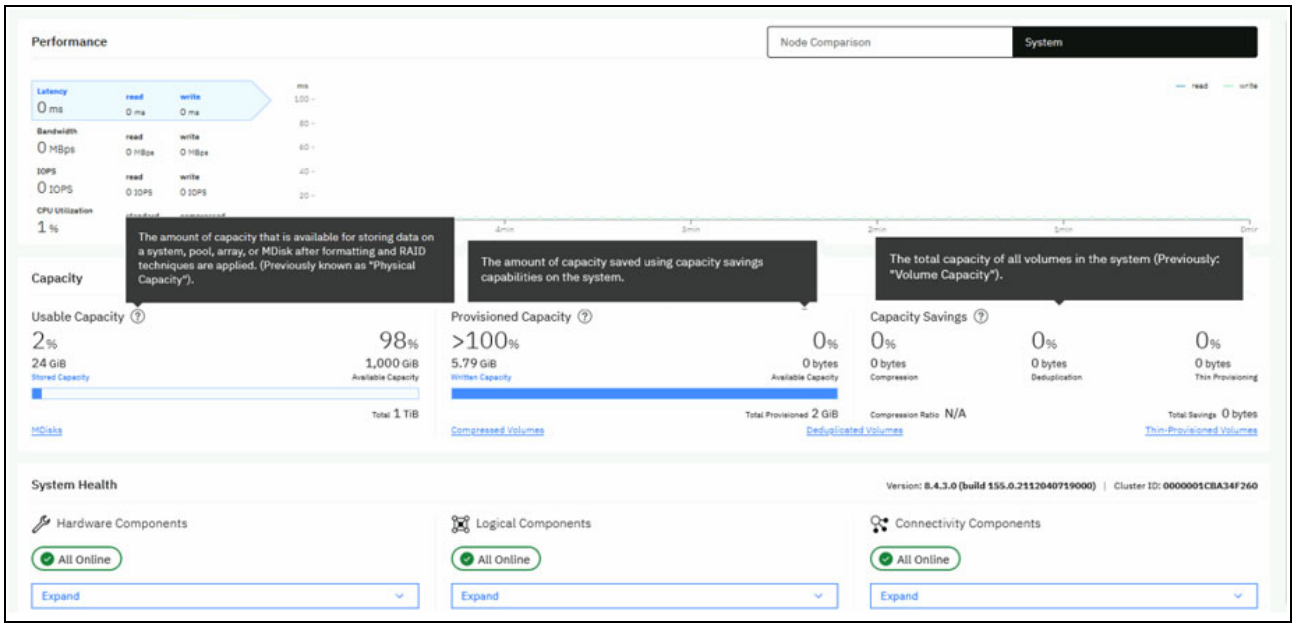

Figure 8-12 Dashboard capacity view

8.3.1 Usable Capacity

Usable Capacity indicates the total capacity in all storage on the system. It includes all the storage that can be virtualized and assigned to pools. Usable capacity is displayed in a bar graph and is divided into the following categories:

- Stored Capacity

  Stored Capacity indicates the amount of capacity that is used on the system after capacity savings. The system calculates the stored capacity by subtracting the available capacity and any reclaimable capacity from the total capacity that is allocated to managed disks (MDisks). To calculate the percentage, the stored capacity is divided by the total capacity that is allocated to MDisks. On the left side of the bar graph, the stored capacity is displayed in the total capacity and as a percentage.

- Available Capacity

  The total *Available Capacity* appears on the right side of the bar graph. Available capacity is calculated by adding the available capacity and the total reclaimable capacity. To calculate the percentage of available capacity on the system, the available capacity is divided by the total amount of capacity that is allocated to MDisks.

- Total

  The *Total Capacity* appears on the right under the bar graph and shows all the capacity that is available on the system. The bar graph is a visual representation of capacity usage and availability, and can be used to determine whether more storage must be added to the system.

Select MDisks to view more information about the usable capacity of the system on the MDisks by Pools page. You can also select Compressed Volumes, Deduplicated Volumes, or Thin-Provisioned Volumes.

8.3.2 Provisioned Capacity

Provisioned Capacity is the total capacity of all virtualized storage on the system. Provisioned capacity is displayed in a bar graph and is divided into two categories: `Written Capacity` and `Available Capacity`.

Written Capacity appears on the left side of the bar graph and indicates the amount of capacity that has data that is written to all the configured volumes on the system. The system calculates the written capacity for volumes by adding the stored capacity to the capacity savings. The percentage of written capacity for volumes is calculated by dividing the written capacity by the total provisioned capacity for volumes on the system.

The *Available Capacity* appears on the right side of the bar graph and indicates the capacity on all configured volumes that is available to write new data. The Available Capacity is calculated by subtracting the written capacity for volumes from the total amount of capacity that is provisioned for volumes.

The percentage of Available Capacity is calculated by dividing the Available Capacity for volumes by the total amount of capacity that is provisioned to volumes on the system. The Total Provisioned capacity displays under the Available Capacity and indicates the total amount of capacity that is allocated to volumes.

The Provisioned Capacity also shows the percentage for over-provisioned volumes. The Overprovisioned value indicates the percentage of provisioned capacity that is increased because of capacity savings.

8.3.3 Capacity Savings

Capacity Savings indicates the amount of capacity that is saved on the system by using compression, deduplication, and thin-provisioning. The percentage value for each of these capacity savings methods compares the stored capacity before and after capacity savings is applied. Compression shows the total capacity savings that are gained from the use of compression on the system. Deduplication indicates the total capacity savings that the system is saved from all deduplicated volumes. Thin-Provisioning displays the total capacity savings for all thin-provisioned volumes on the system. You can view all the volumes that use each of these technologies.

8.3.4 Monitoring performance in IBM Spectrum Virtualize for Public Cloud on Microsoft Azure or AWS

From left window of the GUI, select **Monitoring** → **Performance** to monitor real-time statistics of CPU utilization, volume, interface, and MDisk bandwidth of your system and nodes. Each graph represents 5 minutes of collected statistics and provides a means of assessing the overall performance of your system (see Figure 8-13).

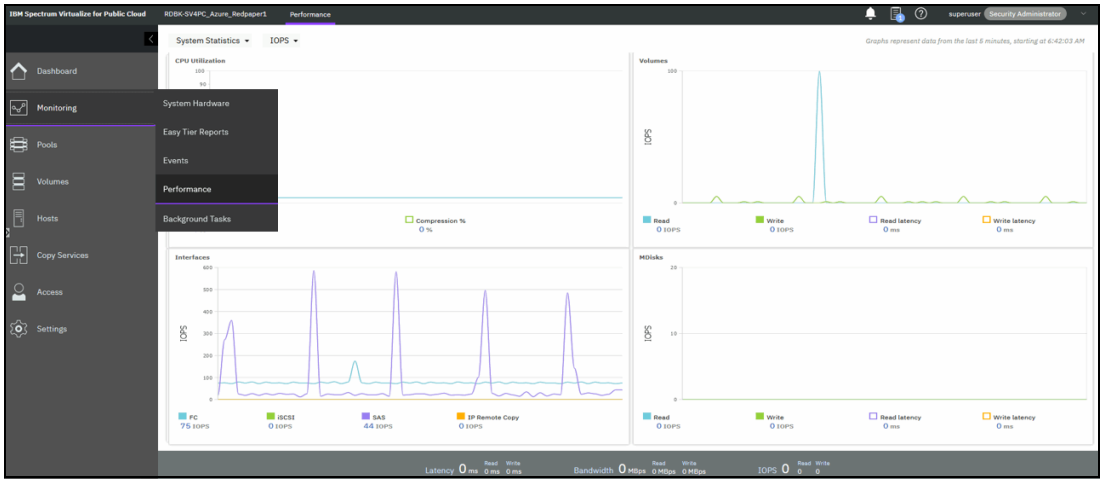

Figure 8-13 Monitoring performance

You can use system statistics to monitor the bandwidth of all the volumes, interfaces, and MDisks that are being used on your system. You can also monitor the overall CPU utilization for the system. These statistics summarize the overall performance health of the system and can be used to monitor trends in bandwidth and CPU utilization.

You can monitor changes to stable values or differences between related statistics, such as the latency between volumes and MDisks. These differences then can be evaluated further by using performance diagnostic tools.

Monitoring performance in Microsoft Azure

To monitor performance in Microsoft Azure, complete the following steps:

1. Log in to the Microsoft Azure Portal.
2. Select **Monitor**, as shown in Figure 8-14 on page 229.

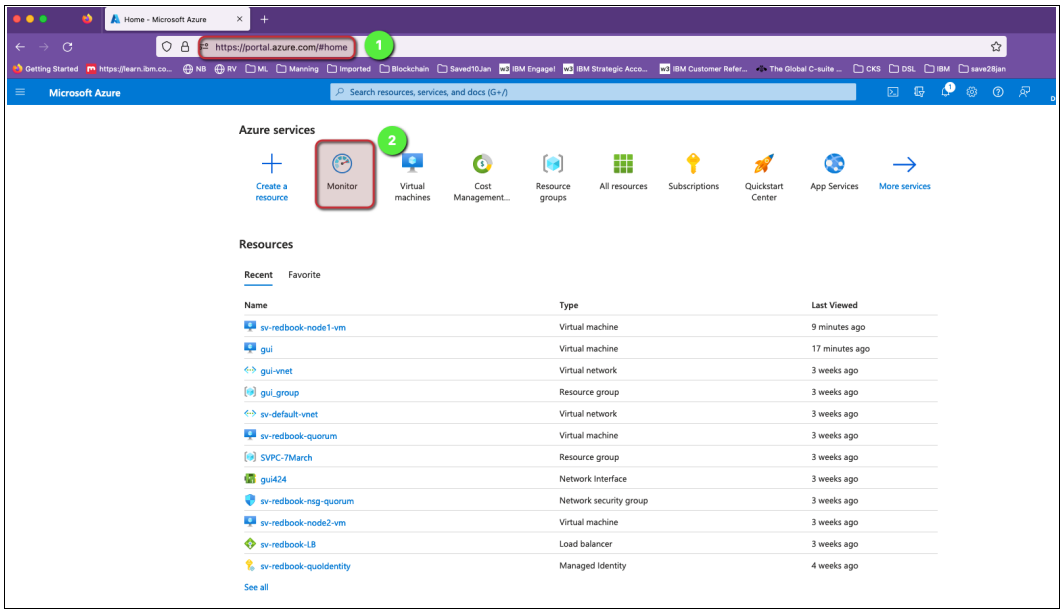

Figure 8-14 Microsoft Azure Monitor service

3. Select the resource and the metric to be monitored.

> **Note:** Azure provides different metrics for different resources, which can be used for monitoring performance and stats.

4. Select the metric that you want to view (see Figure 8-15). Our example shows the metric of the disk that is provided by Azure.

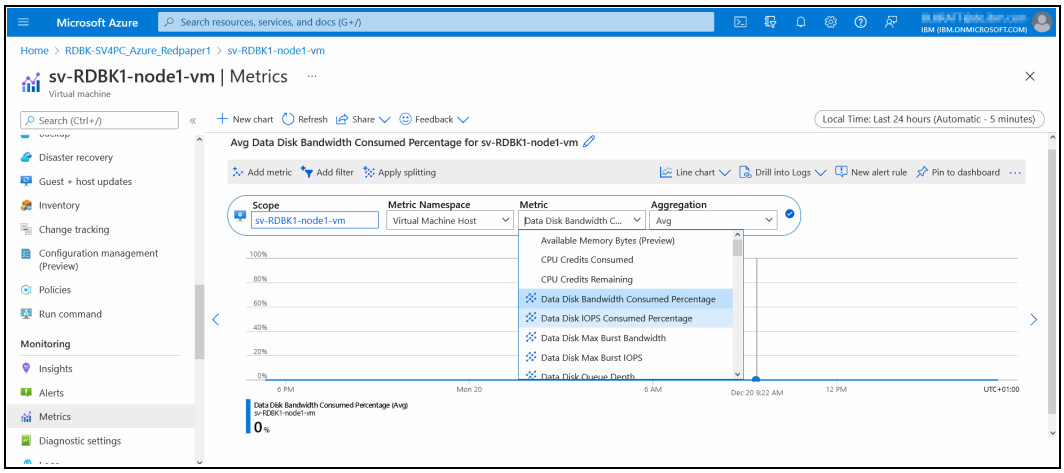

Figure 8-15 Microsoft Azure Metrics

> **Note:** For more information about metrics for Azure disks, see Supported metrics with Azure Monitor.

Monitoring performance in AWS

To monitor performance in AWS, complete the following steps:

1. Log in to the AWS Portal.
2. From the Elastic Compute Cloud (EC2) section, go to the disk or volume that is provisioned in your stack.
3. Select the **Monitoring** tab for the selected resource to view the performance stats that are collected by AWS (Figure 8-16).

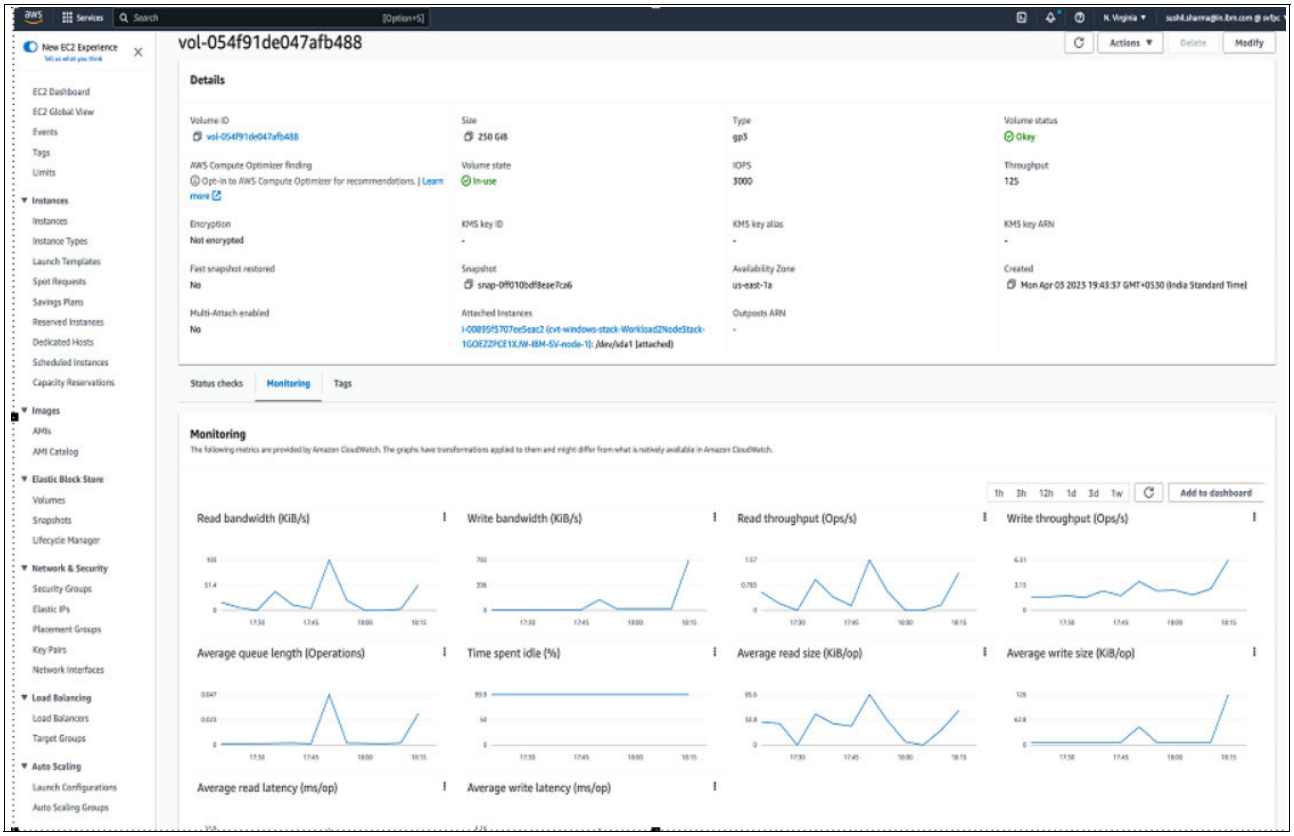

Figure 8-16 Performance stats that are collected by AWS

Note: For more information about gathering metrics for AWS resources, see What is Amazon EC2?

8.3.5 Monitoring IBM Spectrum Control for Public Cloud (Azure or AWS) in IBM Spectrum Control or IBM Storage Insights

IBM Spectrum Control is an on-premises storage management, monitoring, and reporting solution for storage systems, hypervisors, servers, fabrics, and switches. It uses the metadata that it collects about vendors' storage devices to provide services, such as custom alerting, analytics, and replication management.

For more information about the capabilities of IBM Spectrum Control, see IBM Spectrum Control documentation.

For a complete list of the storage systems that are supported by IBM Spectrum Control, see this IBM Support web page.

Because IBM Spectrum Control is an on-premises tool, it does not send the metadata about monitored devices offsite, which is ideal for dark shops and sites that do not want to open network ports to the internet.

IBM Storage Insights is an off-premises IBM Cloud service that provides cognitive support capabilities, monitoring, and reporting for storage systems. Because it is an IBM Cloud service, getting started is simple and upgrades are handled automatically.

By using the IBM Cloud infrastructure, IBM Support can monitor your storage environment to help minimize the time to resolution of problems and collect diagnostic packages without requiring you to manually upload them. This wraparound support experience, from environment to instance, is unique to IBM Storage Insights and transforms how and when you get help.

Both IBM Spectrum Control and IBM Storage Insights monitor storage systems, fabrics and switches, but IBM Spectrum Control also monitors hypervisors to provide you with unique analytics and insights into the topology of your storage network. It also provides more granular collection of performance data, with 1-minute interval rather than the 5-minute interval in IBM Storage Insights or IBM Storage Insights Pro.

For more information about the capabilities of IBM Storage Insights, see IBM Storage Insights documentation.

For more information about the storage systems that are supported by IBM Storage Insights, see IBM Documentation web page.

With IBM Spectrum Control or IBM Storage Insights, you can view the capacity, space usage, and performance of your IBM Spectrum Virtualize for Public Cloud storage systems. Other monitoring features, such as alerting, health checking, advanced analytics, and reporting are also supported.

Before you can add an IBM Spectrum Virtualize for Public Cloud storage system for monitoring, you must ensure that IBM Spectrum Control or IBM Storage Insights can connect to it.

To enable a connection, use one of the following methods:

- Method 1
  - Azure configuration

    Use a site-to-site virtual private network (VPN) IPsec tunnel that exists between the on-premises environment and the IBM Spectrum Virtualize on Azure storage system. Use this method if security or operations constraints exits that are related to controlling outbound internet on Azure connections in your cloud environment.

    Azure Virtual Private Network Gateways (VPNGW) provide cross-premises connectivity between customer premises and Azure. For more information about how to use the Azure portal to create a site-to-site VPNGW connection from your on-premises network to a vNet, see this Microsoft tutorial.
  - AWS configuration

    Configure an AWS VPNGW between the AWS Cloud VPN and on-premises setup to enable data transmission between AWS Cloud and an on-premises network.

To create a VPNGW, complete the following steps:

i. Create a virtual private gateway in your AWS account.

ii. Create a Customer Gateway (CGW) to represent your on-premises VPN device.

iii. Create a VPN connection to connect your virtual private gateway to your CGW. During this step, specify the public IP address of your on-premises VPN device and configure the encryption and authentication settings.

iv. Update the routing tables in your virtual private cloud (VPC) and on-premises network to allow traffic to flow between them. You might need to modify your network security groups (NSGs) to allow inbound and outbound traffic through the VPN connection.

v. Test the VPN connection to ensure that it is working correctly.

For more information about these steps, see the following reference materials:

- Virtual Private Gateway
- VPN Connection
- Routing and Security

- Method 2

  Install the Storage Insights data collector or deploy the IBM Spectrum Control software directly on a VM on Azure or AWS.

Storage Insights data collector implementation on a VM in a public cloud (Azure or AWS)

This section describes the implementation of a Storage Insights data collector installation off-premises on Windows Server 2019 Datacenter Gen2 operating system on Public Cloud (Azure or AWS).

As a baseline for the data collector application, a Windows Server 2019 Datacenter operating system was deployed as a VM on Azure or AWS by following the instructions that are available at this Microsoft Docs web page.

Note: The VM that contains the Storage Insights data collector requires a network connection to the vNet of an IBM Spectrum Virtualize for Public Cloud Azure instance and a VPC of IBM Spectrum Virtualize for Public Cloud on AWS.

To deploy the lightweight data collector off-premises to stream performance, capacity, and configuration metadata to IBM Storage Insights, complete the following steps:

1. Log in to your IBM Storage Insights instance.
2. Go to the **Configuration** → **Data Collector** page, and click **Deploy Data Collector** to download the data collector for Windows, as shown in Figure 8-17 on page 233.

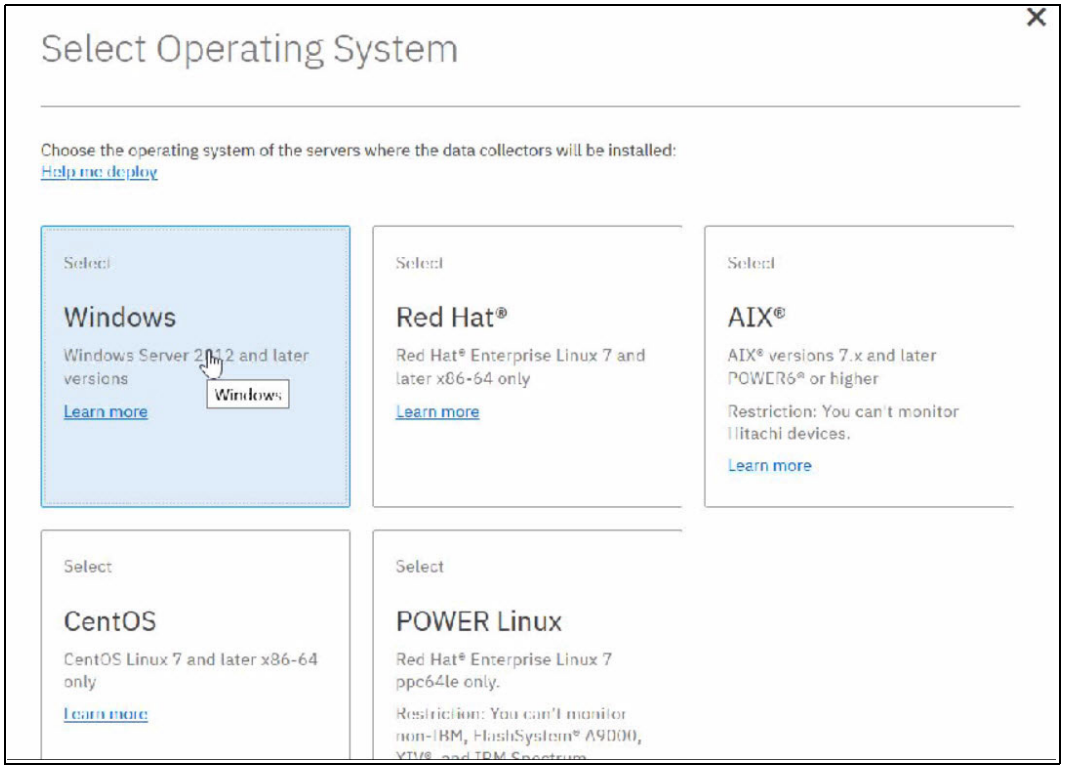

Figure 8-17 Select Operating System

3. Extract the contents of the `Cloud_DataCollector_windows.zip` file on Windows Server 2019 Datacenter.
4. Run **`installDataCollectorService.bat`**, as shown in Figure 8-18.

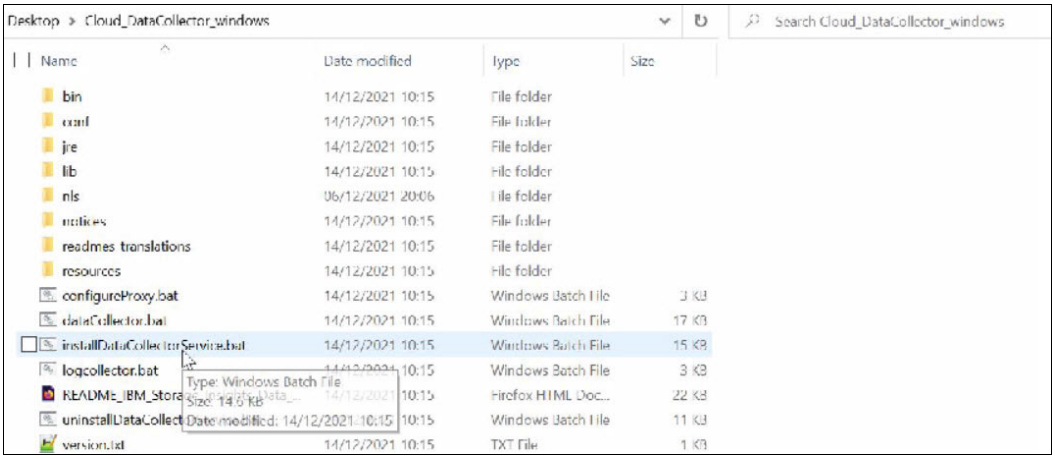

Figure 8-18 Windows Server 2019 Datacenter: installDataCollectorService.bat

For more information about downloading and installing data collectors, see this IBM Documentation web page.

After the data collector is deployed, it attempts to establish a connection to IBM Storage Insights. When the connection is complete, you are ready to start adding your storage systems.

To add the IBM Spectrum Virtualize system on Azure or AWS to Storage Insights, add the cluster with authentication type Secure Shell (SSH) and the private SSH key for the user who is used for authentication. Create a separate user for Storage Insights in IBM Spectrum Virtualize. The task of adding a system is shown in Figure 8-19.

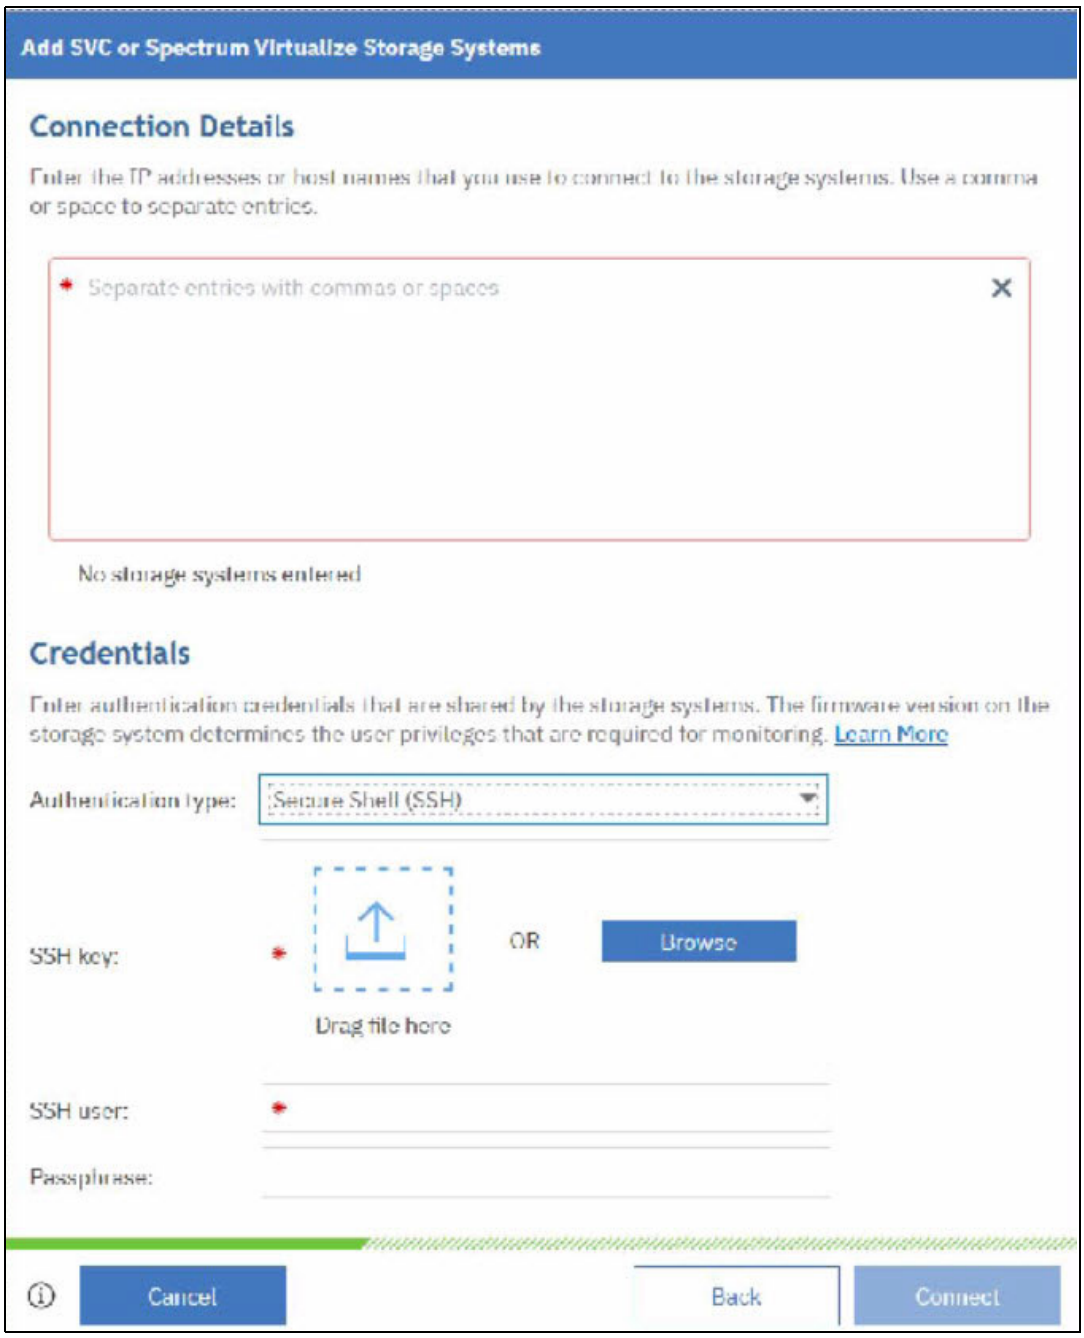

Figure 8-19 Add SAN Volume Controller or IBM Spectrum Virtualize Storage System window

The task to add the storage system completes automatically, as shown in Figure 8-20. The IBM Spectrum Virtualize system on Azure is visible in Storage Insights within a few minutes after the process completes.

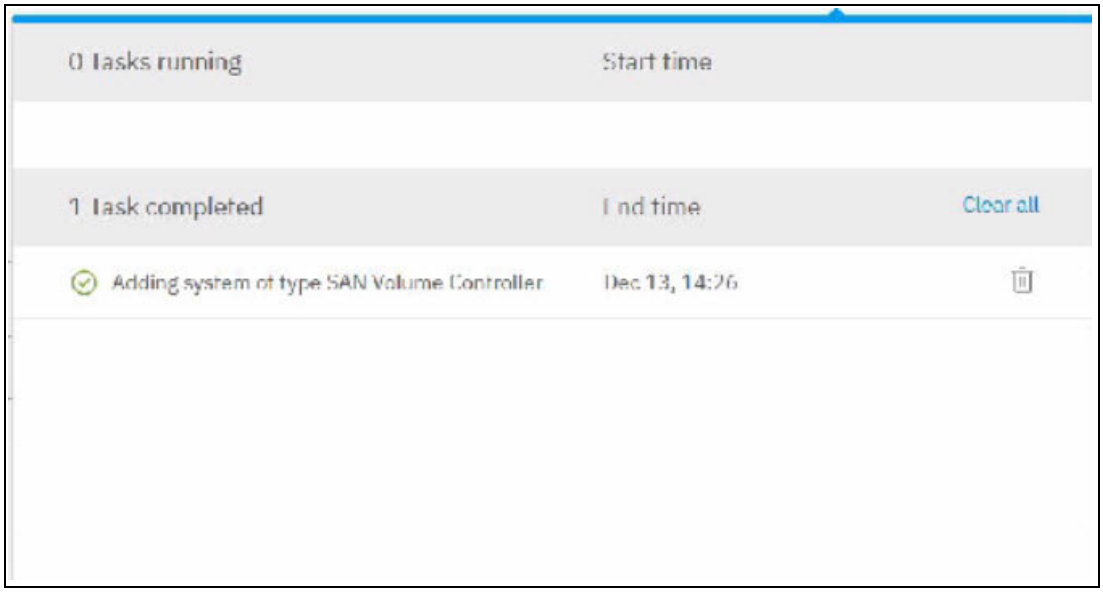

Figure 8-20 SAN Volume Controller or IBM Spectrum Virtualize Storage System successfully added

IBM Spectrum Control VM on Azure or AWS

This section shows a workflow about how the IBM Spectrum Control software can be implemented on a VM on Azure or AWS.

A VM with a supported operating system must be deployed on Azure or AWS as a baseline for IBM Spectrum Control. The IBM Spectrum Control software is available for IBM AIX, Linux, and Microsoft Windows operating systems.

For more information about part numbers, and required and optional part details, see this IBM Support web page.

A useful how-to guide about creating a VM on Azure is available at this Microsoft Docs web page.

The VM must establish a network connection to the IBM Spectrum Virtualize cluster on Azure. For more information about virtual network (vNet) settings on Azure, see this Microsoft Docs web page.

For more information about ab installation guide for the IBM Spectrum Control for all operating systems, see this IBM Documentation web page.

Performance monitoring in IBM Spectrum Control and Storage Insights

The Storage Insights and the IBM Spectrum Control dashboard are a quick way and the first instance to monitor the performance of your storage at a glance.

One of these key aspects is the Top Block Storage Performance section within the IBM Spectrum Control dashboard, which displays the I/O Rate of all added devices and gives you the first performance overview.

The *Performance* section within the dashboard of IBM Storage Insights is slightly different than the IBM Spectrum Control dashboard. Complete the following steps:

1. To get an overview of the I/O Rate (ops/s), Data Rate (MiBps) or Response Time (ms/op), select the IBM Spectrum Virtualize cluster on Azure. The Performance section shows these values, as shown in Figure 8-21.

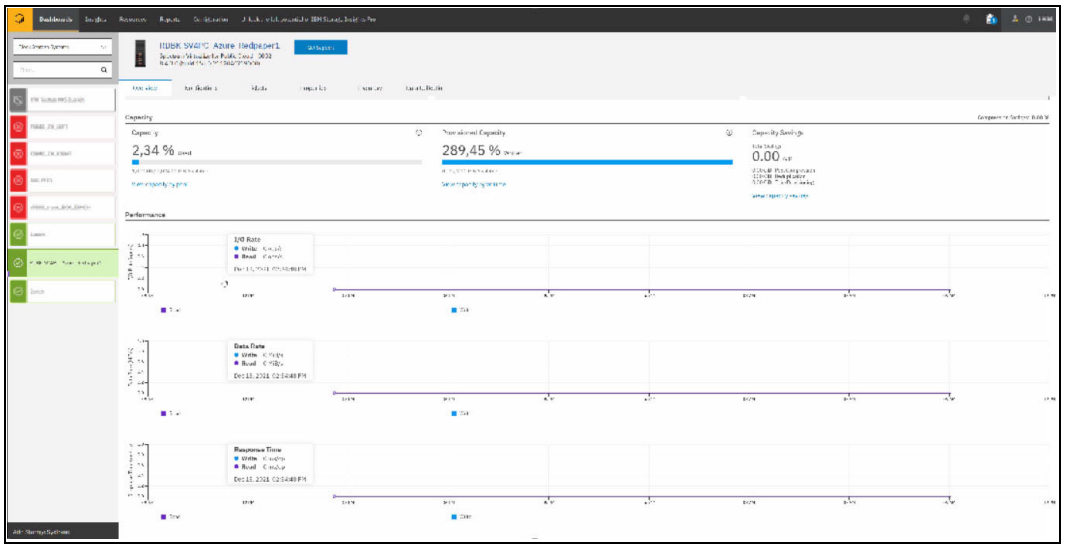

Figure 8-21 Dashboard of an IBM Spectrum Virtualize cluster on Azure

IBM Storage Insights Pro and IBM Spectrum Control also provide an enhanced performance view and more performance metrics.

For more information about the differences between IBM Storage Insights and IBM Storage Insights Pro, see this IBM Documentation web page.

For more information about the difference between Storage Insights and IBM Spectrum Control, and about how the IBM Spectrum Control can be extended with IBM Storage Insights, see this IBM Documentation web page.

2. Select **Resources** → **Block Storage Systems** to open a detailed view of the performance metrics.
3. Double-click the system to open the Overview page. The Performance tab in the General section of the Overview page provides different key performance indicators, as shown in Figure 8-22.

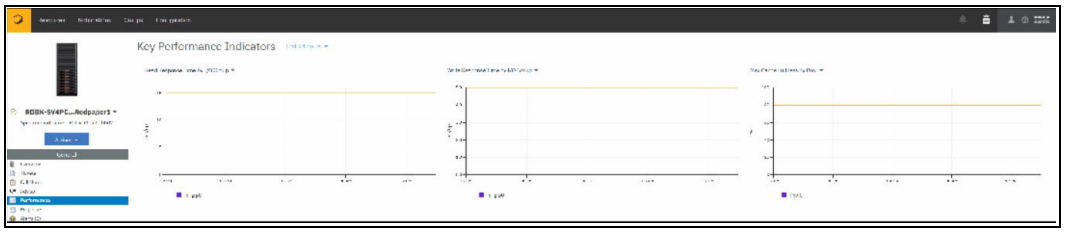

Figure 8-22 Key performance indicators of an IBM Spectrum Virtualize cluster on Azure

4. Choose **View performance** in the **Actions** menu to open the performance metrics for the cluster, as shown in Figure 8-23.

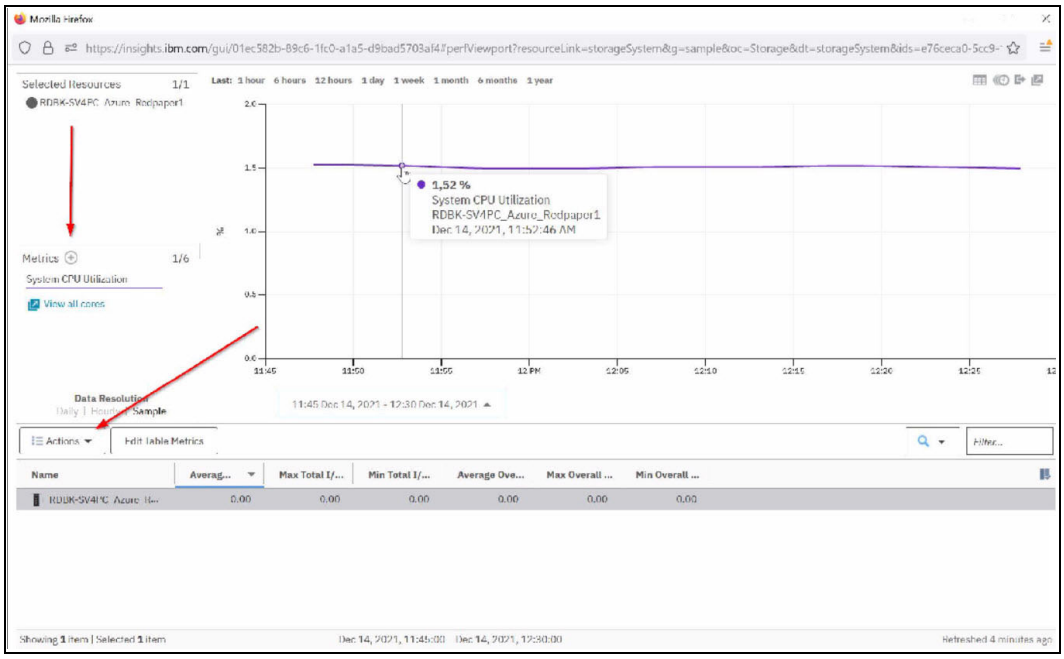

Figure 8-23 CPU utilization of an IBM Spectrum Virtualize cluster on Azure

5. Click **+** to change the chart metrics. To change the selected resources between Node, Pool, MDisk, and Volume, click **Actions** (see Figure 8-23).

   Alternatively, the performance view can be opened by selecting the system in the Block Storage Systems page and clicking **View Performance**.

   In IBM Spectrum Control, open View Performance and the Overview page by selecting **Storages** → **Block Storage Systems** in the top toolbar.

Note: The View Performance window of IBM Spectrum Control and IBM Storage Insights are the same.

Capacity monitoring in IBM Spectrum Control and IBM Storage Insights

The capacity section in IBM Spectrum Control and IBM Storage Insights provides an overall view of system capacity and illustrates the usable capacity, provisioned capacity, and capacity savings.

The capacity section of a system in IBM Spectrum Control is at the top of the **Overview** page each device. To open the Overview page of a device, the system must be selected by double-clicking it in **Storage** → **Block Storage Systems**.

In Storage Insights, the capacity segment is shown in the Overview section, which can be opened by selecting a device through the Dashboard, as shown in Figure 8-24.

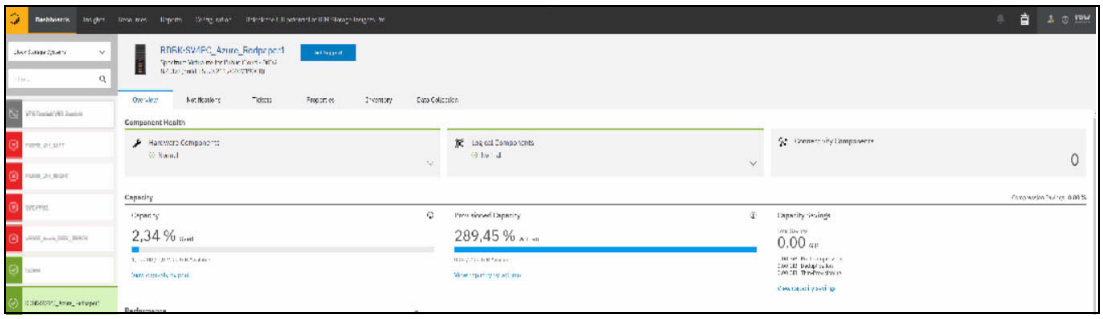

Figure 8-24 Capacity report of an IBM Spectrum Virtualize cluster on Azure

Note: The capacity overview of IBM Spectrum Control IBM Spectrum Control and IBM Storage Insights are the same.

The Capacity shows how much capacity is used and how much capacity is available for storing data. The *Provisioned Capacity* chart shows the written capacity values in relation to the total provisioned capacity values before data reduction techniques are applied.

The capacity of data that is written to the volumes is expressed as a percentage of the total provisioned capacity of the volumes. Available capacity is the difference between the provisioned capacity and the written capacity, which is the thin-provisioning savings.

A breakdown of the total *capacity savings* that are achieved when the written capacity is stored on the thin-provisioned volumes also is provided. In the capacity overview chart, a horizontal bar is shown when a capacity limit is set for the storage system. Hover over the chart to see the capacity limit and how much capacity is left before the capacity limit is reached.

Also, IBM Spectrum Control and IBM Storage Insights Pro provide to a detailed capacity view. This view can be opened in IBM Spectrum Control by right-clicking the device in **Storages** → **Block Storage Systems** or the **Actions** menu in the Overview page of the selected system.

In IBM Storage Insights, **View Capacity** can be selected in the Actions menu of the Overview page as shown in Figure 8-25. The Overview page can be opened by selecting **Resources** → **Block Storage Systems** in the top toolbar.

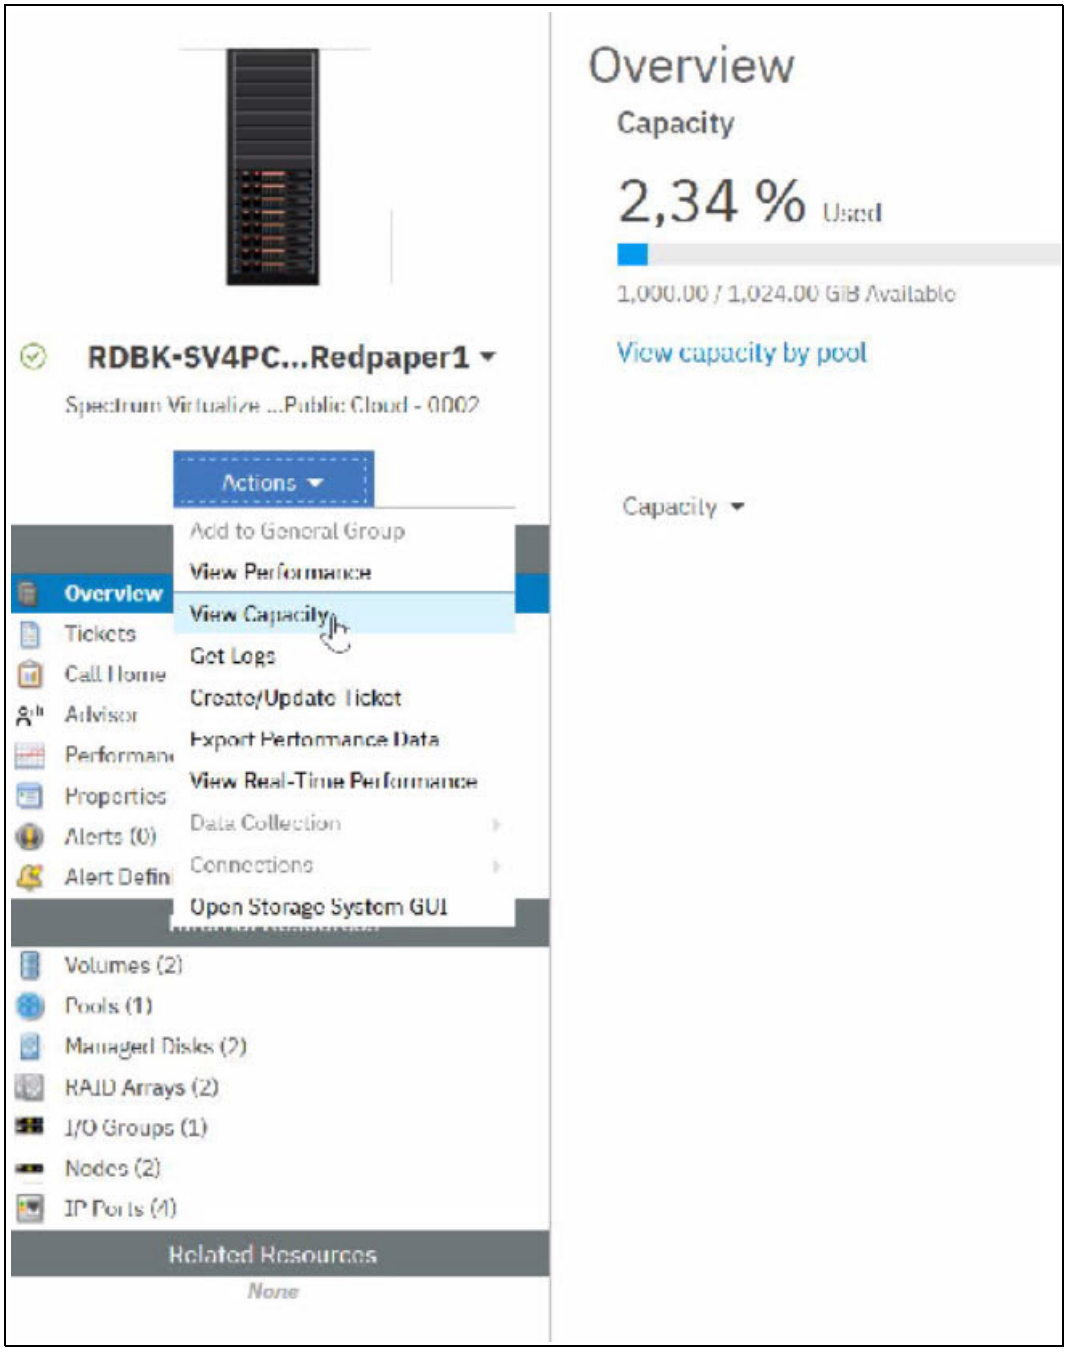

Figure 8-25 Viewing the capacity of an IBM Spectrum Virtualize cluster on Azure on the Overview page

Alternatively, the Capacity view can be opened by right-clicking a device in **Resources** → **Block Storage Systems**, as shown in Figure 8-26.

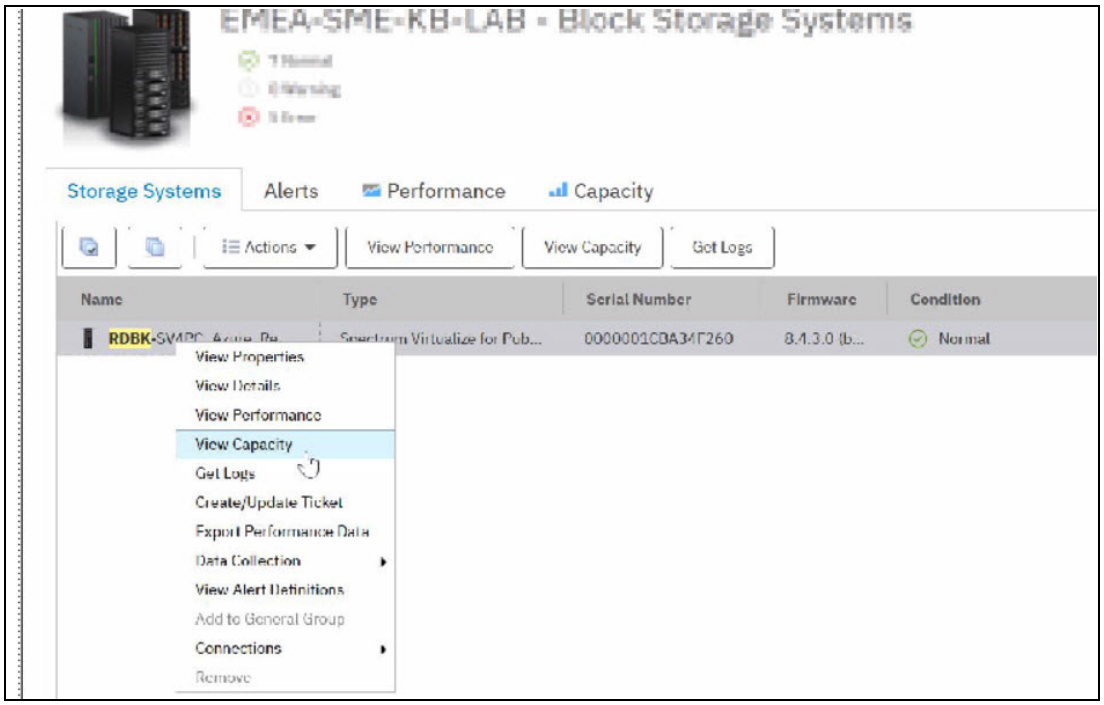

Figure 8-26 Viewing the capacity of an IBM Spectrum Virtualize cluster on Azure

For more information about the capacity metrics for Storage Insights block storage systems, see this IBM Documentation web page.

For more information about the capacity metrics for IBM Spectrum Control block storage systems, see this IBM Documentation web page.

8.3.6 Alerting in IBM Spectrum Control and IBM Storage Insights

In this section, the alerting function of IBM Spectrum Control and Storage Insights is described.

Alerting functions examine the attributes, capacity, and performance of resources. If the conditions that are defined for alerts are met, the actions that are specified for the alert are taken. Typically, the actions include sending a notification.

For example, if the status of an IBM Spectrum Virtualize storage system on Azure or AWS changes to Error, an alert is displayed in the Alerts page in the GUI, and an email might be sent to a storage administrator.

The conditions that trigger alert notifications depend on the type of monitored resource. In general, the following types of conditions can trigger alerts:

- A change of an attribute or configuration of a resource
- The capacity of a resource is outside of a specified range
- The performance of a resource is outside of a specified range
- Change of infrastructure, such as a new or removed resource
- Data is not being collected for a resource

For example, use performance thresholds to be notified when the total I/O rate for a storage system is outside a specified range. This information can help to identify areas in an infrastructure that is over-utilized.

The following example shows how the alerting functions can be implemented in IBM Storage Insights. It also describes the differences to IBM Spectrum Control. The policy section can be opened by using the toolbar in IBM Spectrum Control and IBM Storage Insights:

- IBM Spectrum Control: **Settings** → **Alert policies**
- IBM Storage Insights: **Configuration** → **Alert policies**

Complete the following steps in IBM Storage Insights:

1. Generate a policy by clicking **Create Policy**, as shown in Figure 8-27.

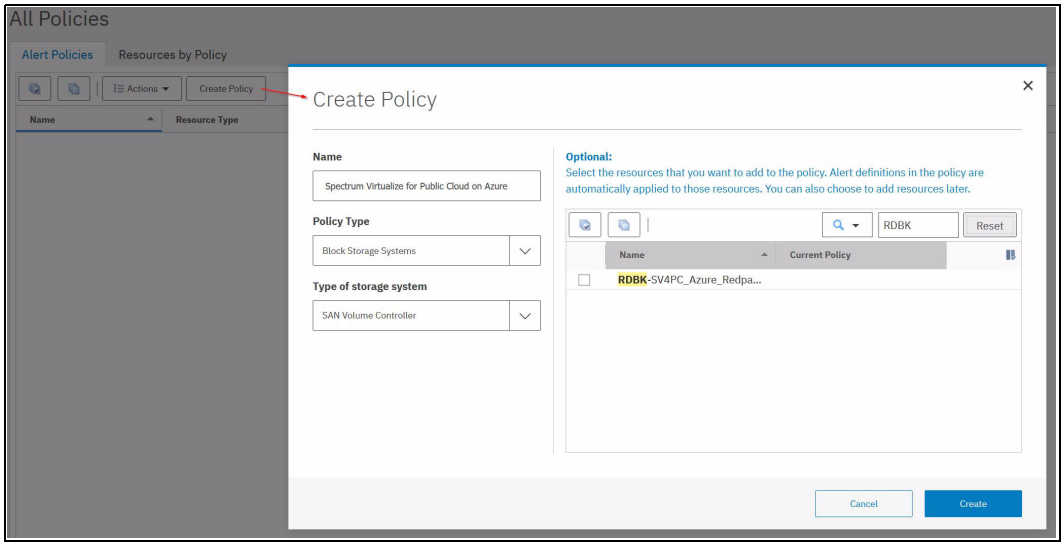

Figure 8-27 Creating a policy for an IBM Spectrum Virtualize cluster on Microsoft Azure

2. Define an alert in the policy, as shown in Figure 8-28.

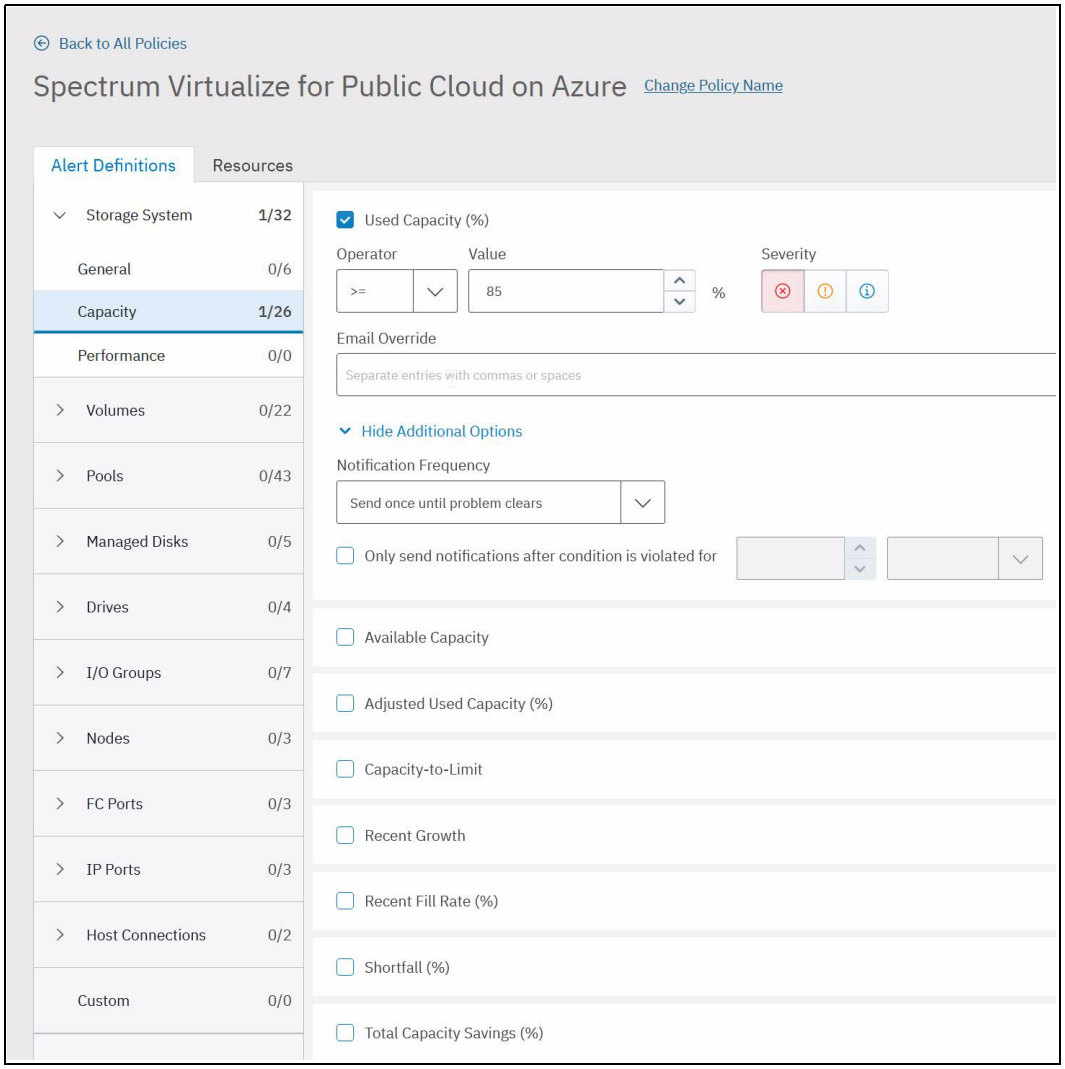

Figure 8-28 Alert definition in the policy for an IBM Spectrum Virtualize cluster on Microsoft Azure

3. Click **Save Changes**.
4. To change configuration settings, such as adding, removing, or changing alert definitions of the policy, click in IBM Storage Insights **Settings** → **Alert Policies in IBM Spectrum Control** or **Configuration** → **Alert Policies**.
5. Double-click the policy and open **Edit Alert Definitions** to make changes in the IBM Spectrum Control and IBM Storage Insights GUI interface.

Notification settings in IBM Spectrum Control and IBM Storage Insights

By using IBM Spectrum Control and IBM Storage Insights, global alert notifications, policy notifications, and alert definition notifications can be defined.

Complete the following steps:

1. Global Alert Notifications specifies the global notification settings for all alert definitions. To configure the global notification settings in IBM Storage Insights, click **Configuration** → **Settings**, and specify the email addresses that you want to notify when alerts are generated.
2. To configure Global email notification settings in IBM Spectrum Control, click **Settings** → **Notification Settings**.
3. Policy Notifications specifies the notification settings for an alert policy. To configure the notification settings for an alert policy in IBM Storage Insights, click **Configuration** → **Alert Policies**.
4. In IBM Spectrum Control, the policy notifications can be configured in **Settings** → **Alert Policies**. Double-click the policy whose notification settings you want to specify. Then, click **Edit Policy Notifications**.
5. Specify the email addresses that you want to notify when alerts are generated. The email addresses are applied to all the alert definitions for all resources in the policy, unless overridden.

Email Override specifies the notification settings for a specific alert definition in an alert policy. Complete the following steps:

1. In IBM Storage Insights, click **Configuration** → **Alert Policies**. In IBM Spectrum Control, click **Settings** → **Alert Policies**.
2. Double-click the policy whose notification settings you want to specify. For example, to change the notification settings for specific alerts in a custom alert policy, double-click the policy and click **Edit Alert Definitions**.
3. Edit the Email Override field or click **View Additional Options** for an alert definition to specify notification settings. The email addresses that you specify for the alert definition override any global notification settings, policy settings, and settings for the resource.

   Figure 8-29 shows the Email Override and Policy Notifications notification settings in IBM Storage Insights.

Figure 8-29 Notification settings for IBM Spectrum Virtualize for Public Cloud on Microsoft Azure

Also, by using IBM Spectrum Control and IBM Storage Insights, the notification settings for resources can be configured. Complete the following steps:

1. To specify the notification settings for an IBM Spectrum Virtualize cluster in Microsoft Azure, go to **Storage** in IBM Spectrum Control or **Resources** in **IBM Storage Insights** → **Block Storage Systems**.
2. Right-click the storage system, click **View Alert Definitions** and then, click **Edit Notifications**. The email addresses that you specify are applied to all the alert definitions that are specified for the selected storage system.

For more information about IBM Spectrum Control alert notifications, see this IBM Documentation web page.

For more information about alert notifications in IBM Storage Insights, see this IBM Documentation web page.

9

Troubleshooting the solution

This chapter provides guidance about supporting and monitoring this solution.

This chapter includes the following topics:

- Troubleshooting IBM Spectrum Virtualize for Public Cloud
- Collecting diagnostic data for IBM Spectrum Virtualize
- Uploading files to the Support Center
- Service Assistant Tool
- Remote Support Assistance
- Troubleshooting and fix procedures
- Managing the event log
- Troubleshooting in Microsoft Azure
- Troubleshooting in Amazon Web Services
- IBM Spectrum Virtualize for Public Cloud Support

9.1 Troubleshooting IBM Spectrum Virtualize for Public Cloud

The section describes how to collect support data in an IBM Spectrum Virtualize for Public Cloud.

9.2 Collecting diagnostic data for IBM Spectrum Virtualize

Occasionally, if a problem occurs and the IBM Support Center is contacted, you are prompted 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 then, the **Support Package** tab. Then, click **Upload Support Package**. The Upload Support Package button is available only if a DNS server is configured. When no DNS server is configured, use the Manual Upload instructions that are described in Step 4.

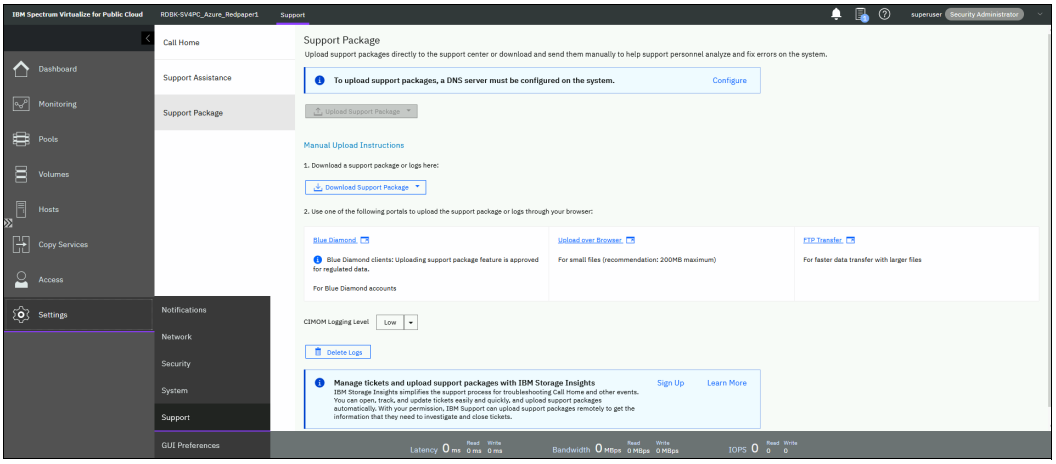

Figure 9-1 Support Package option

Assuming the problem that was encountered was an unexpected node restart that logged a 2030 error, collect the default logs and 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 IBM 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.

2. From the Upload Support Package window, four options of data collection are available. You are contacted by IBM Support when your system calls home. If you manually open a call with IBM Support, you receive a case number. Enter the case number into the field and select the snap type, often called snap option 1, 2, 3 or 4, as requested by IBM Support (see Figure 9-2). In our example, we enter a case number. Select snap type **3 (option 3)** because this choice automatically collects the statesave that was created at the time the node restarted. Click **Upload**, as shown in Figure 9-2.

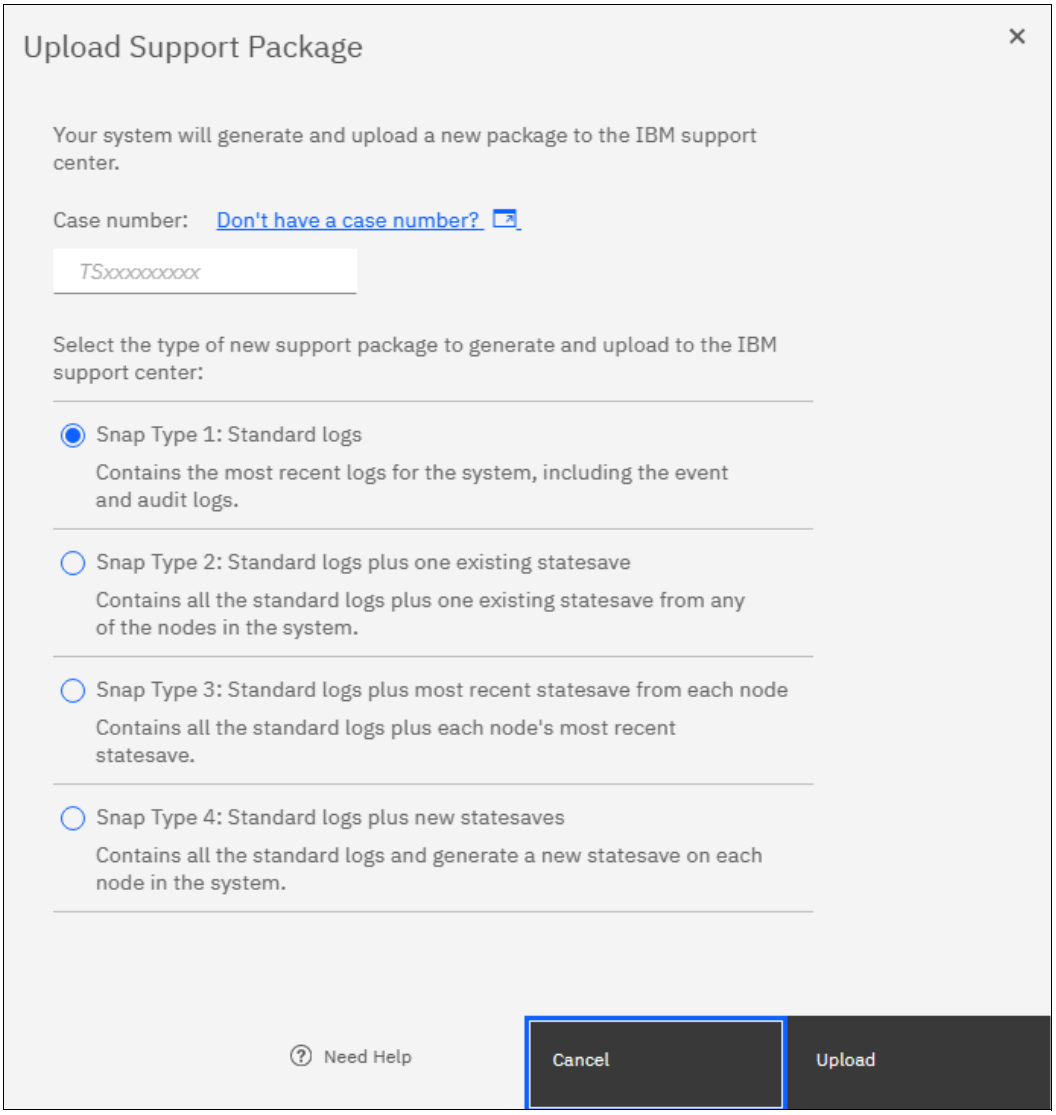

Figure 9-2 Upload Support Package window

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.

Collecting logs by using the command-line interface

Complete the following steps to use the command-line interface (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 that is requested by IBM Support:
   - Standard logs (type 1):

     ```
     svc_snap upload pmr=TSXXXXXXXXX gui1
     ```

   - Standard logs plus one existing statesave (type 2):

     ```
     svc_snap upload pmr=TSXXXXXXXXX gui2
     ```

   - Standard logs plus most recent statesave from each node (type 3):

     ```
     svc_snap upload pmr=TSXXXXXXXXX gui3
     ```

   - Standard logs plus new statesaves:

     ```
     svc_livedump -nodes all -yes
     svc_snap upload pmr=TSXXXXXXXXX gui3
     ```

2. We collect the type 3 (option 3) and have it automatically uploaded to the case number that is provided by IBM Support, as shown in Example 9-1.

Example 9-1 The svc_snap command

```
ssh superuser@IP_address
Password:
RDBK-SV4PC_Azure_Redpaper1:superuser>>svc_snap upload pmr=TSXXXXXXXXX gui3
```

3. If you do not want to automatically upload the snap to IBM, do not specify the `upload pmr=TSxxxxxxxxx` part of the commands. In this case, when the snap creation completes, it creates a file that is named by using 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** twistie → **Download Support Package**. Click **Download Existing Package**, as shown in Figure 9-3.

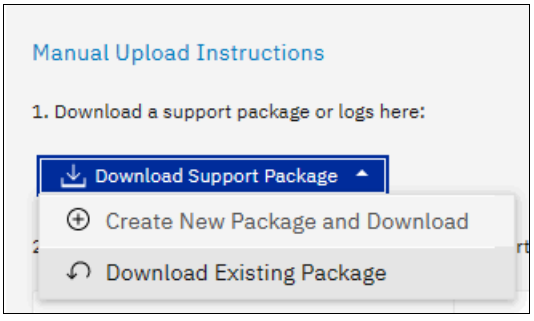

Figure 9-3 Download Existing Package

5. A new window opens. Enter `snap` into the Filter field, and then click **Enter**. A list of snap files is shown (see Figure 9-4). Find the name of the snap that was generated by using the **`svc_snap`** command that was issued earlier. Click to select that file, and then click **Download**.

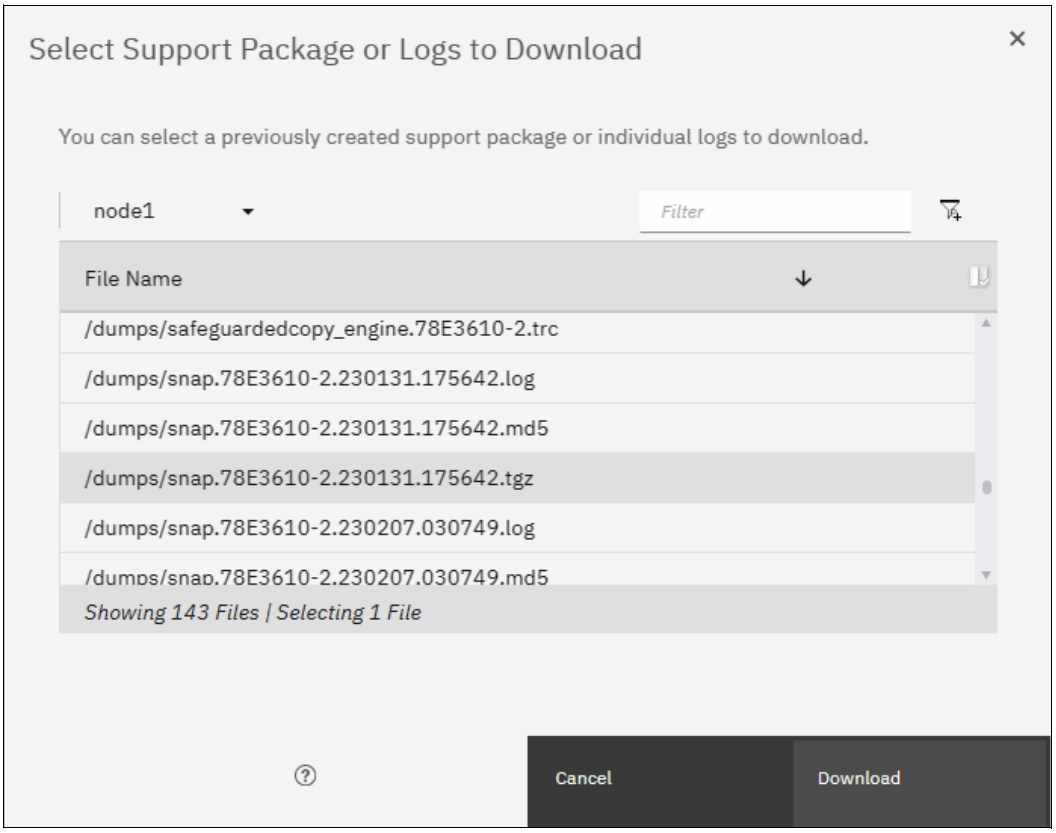

Figure 9-4 Filtering on snap to download

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

9.3 Uploading files to the Support Center

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

To upload information, complete the following steps:

1. Using a browser, navigate to the Enhanced Customer Data Repository webpage (see Figure 9-5).

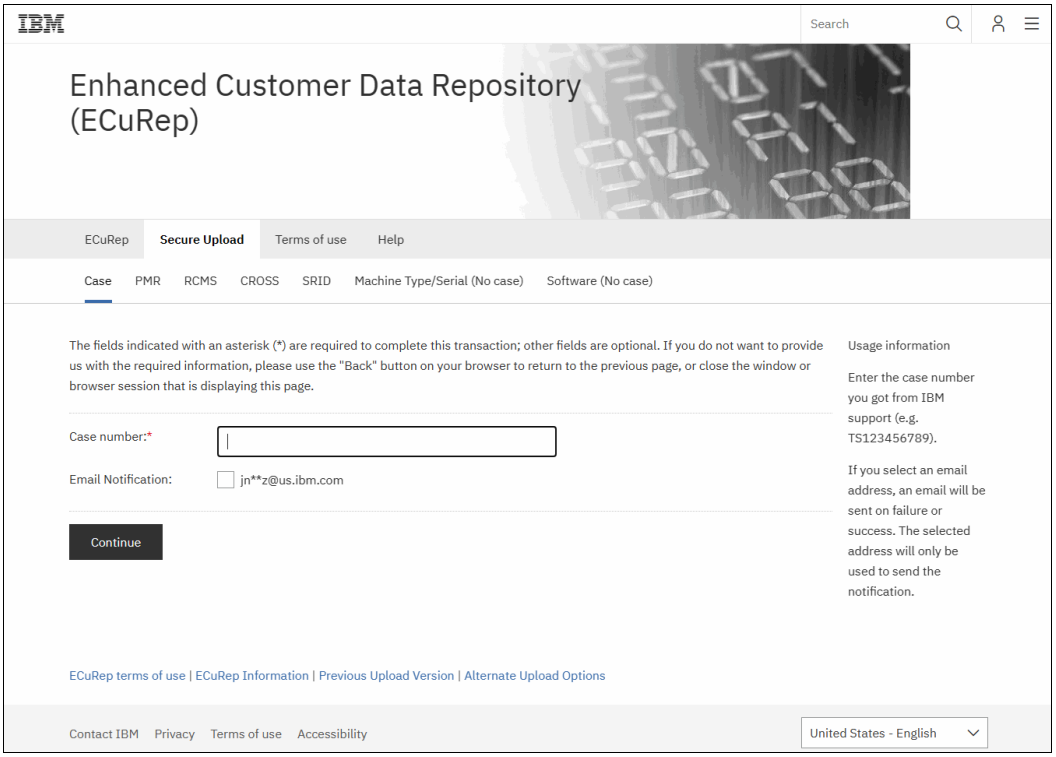

Figure 9-5 ECuRep website

2. Input the case number that was provided by IBM Support for your specific case. This number must be in the format of `TSxxxxxxxxx`; for example, `TS123456789`.

   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 9-6).

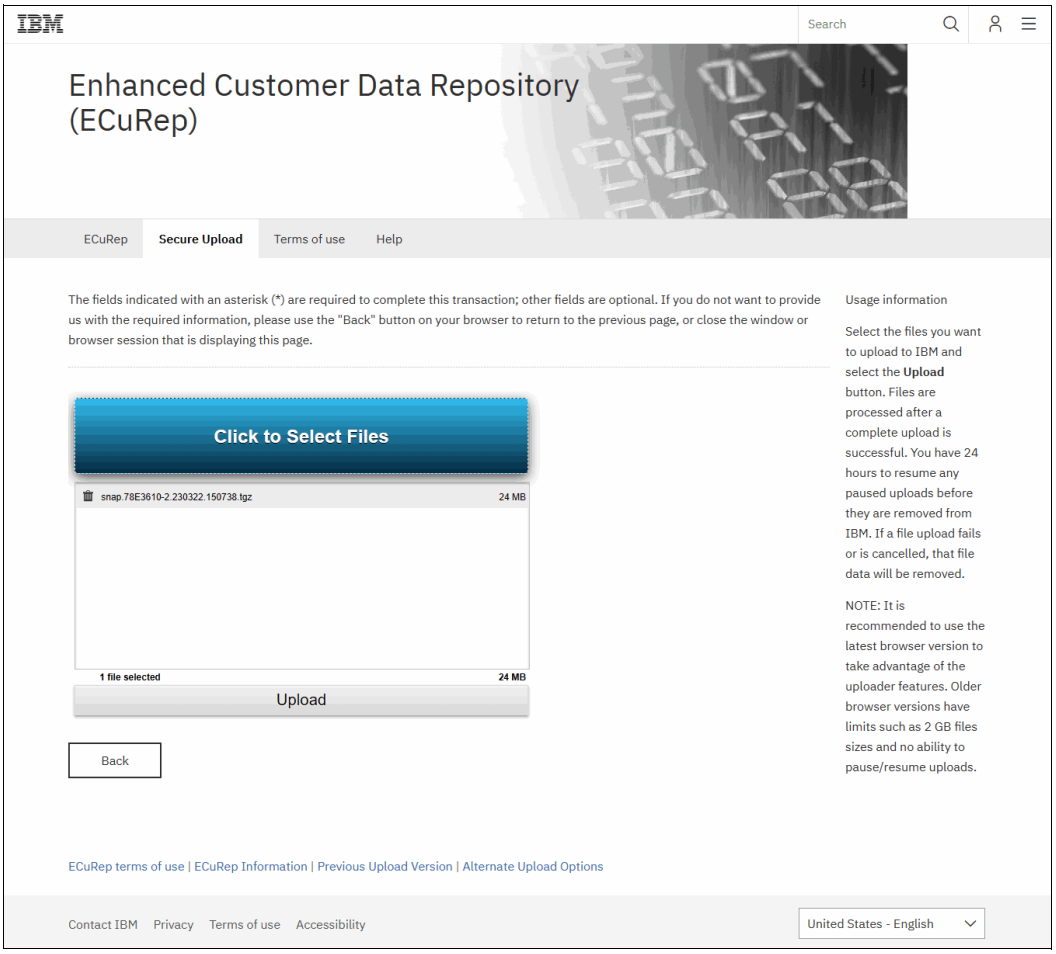

Figure 9-6 Support Package upload to ECuRep

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

9.4 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 in service state is not an active part of a clustered system.

IBM Spectrum Virtualize for Public Cloud on Microsoft Azure is initially configured with the following IP addresses:

- One service IP address for each IBM node.
- Two cluster management IP address.

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

- Status information about the connections and the nodes.
- Data collection for single nodes.
- Basic configuration information, such as configuring IP addresses.

- Service tasks, such as restarting the Common Information Model Object Manager (CIMOM) or web server (Tomcat).
- 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 also can be accessed through the cluster IP addresses by appending `/service` to the cluster management IP address.

If the cluster management IP is not accessible, the only method of communicating with the nodes is through the SAT IP address directly. Each node can have a single service IP address and must be configured for all nodes of the cluster.

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 9-7.

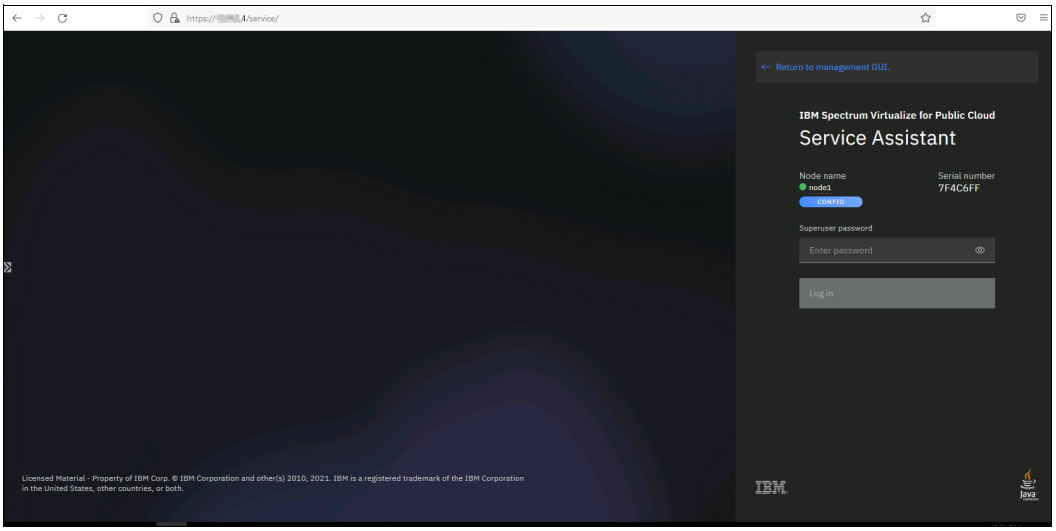

Figure 9-7 Service Assistant Tool Login GUI

2. After you are logged in, you see the SAT home page, as shown in Figure 9-8. The SAT can view the status and run service actions on other nodes, in addition to the node that the user is logged in to. The user is logged in to the node with relationship Local.

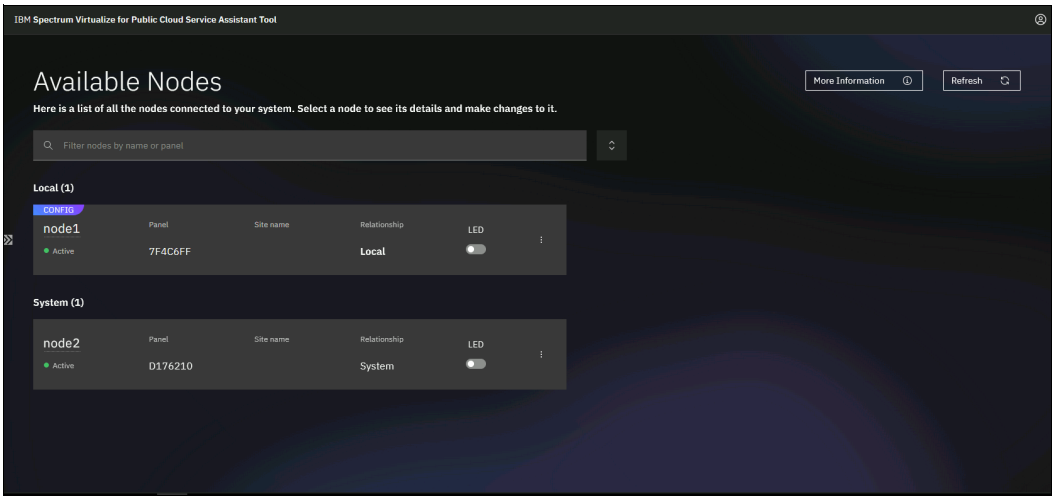

Figure 9-8 Service Assistant Tool

3. Click the node on which you want to run actions. The Node Details are displayed (see Figure 9-9).

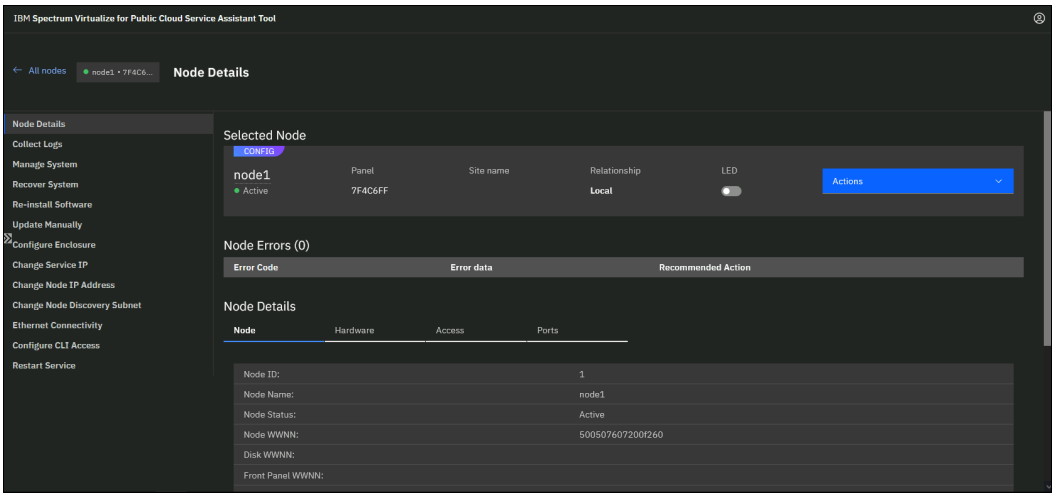

Figure 9-9 Service Assistant Tool: Node Details

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

For more information about how to use the SAT, see this IBM Documentation web page.

Collecting logs in the Service Assistant Tool

When a node is in service support, data can be collected from SAT. Select **Collect Logs** and choose between **Download with latest statesave** or **Download without latest Statesave** (see Figure 9-10).

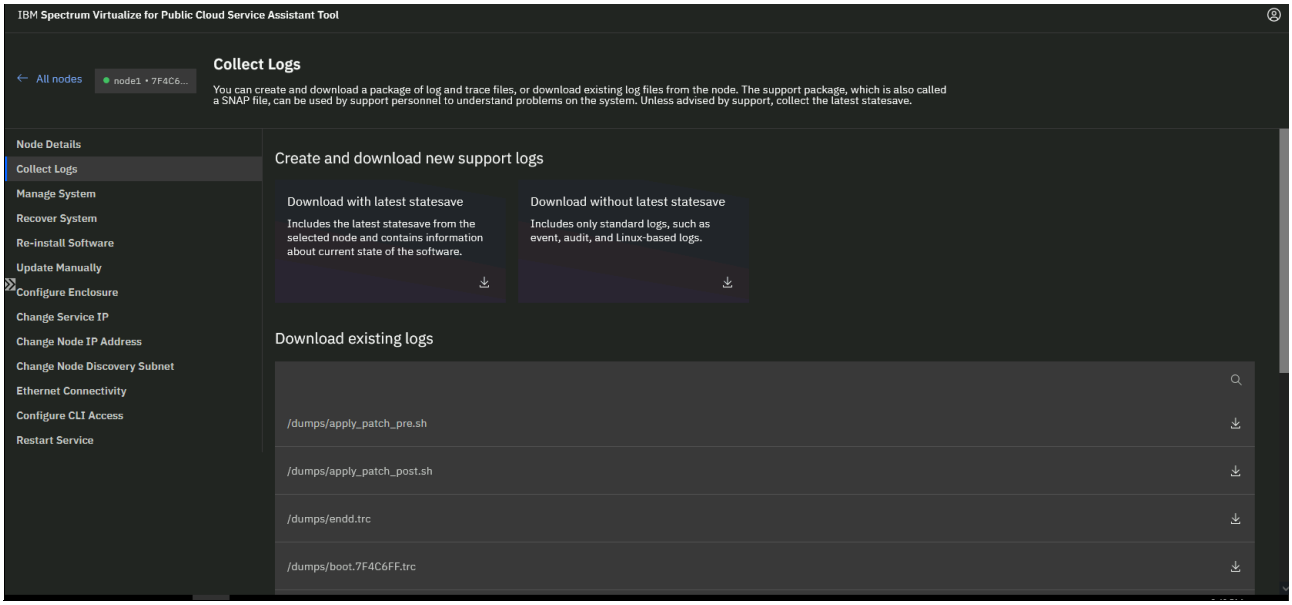

Figure 9-10 Collecting logs by using the Service Assistant Tool

Download the snap file to your workstation and upload file to IBM Support, as described in 9.3, “Uploading files to the Support Center” on page 250.

9.5 Remote Support Assistance

Remote Support Assistance allows IBM Support to remotely connect to the IBM Spectrum Virtualize by way of 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 wanted. If you are enabling Remote Support Assistance, ensure that prerequisites that are described at this IBM Documentation web page are met.

Complete the following steps:

1. Select **Settings** → **Support** → **Support Assistance** in the management GUI or in the system setup. Select **Set up Support Assistance** to start the configuration wizard (see Figure 9-11 on page 255).

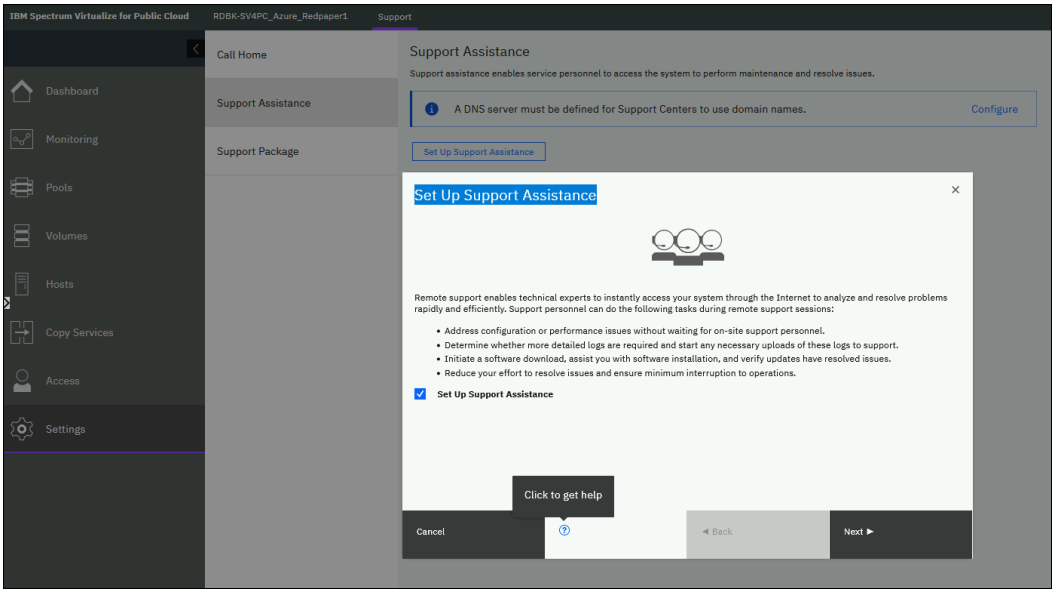

Figure 9-11 Remote Support Assistance menu

2. Click **?** to get more help. A window opens, in which the prerequisites for all configurations are listed (see Figure 9-12).

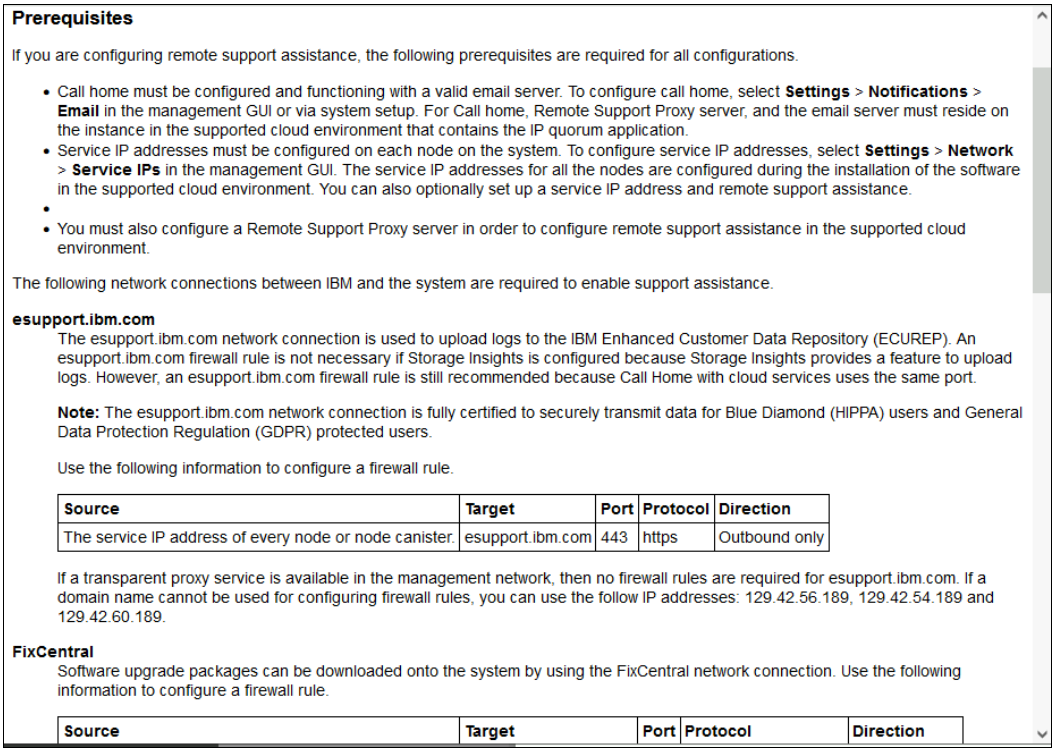

Prerequisites

If you are configuring remote support assistance, the following prerequisites are required for all configurations.

- Call home must be configured and functioning with a valid email server. To configure call home, select **Settings** > **Notifications** > **Email** in the management GUI or via system setup. For Call home, Remote Support Proxy server, and the email server must reside on the instance in the supported cloud environment that contains the IP quorum application.
- Service IP addresses must be configured on each node on the system. To configure service IP addresses, select **Settings** > **Network** > **Service IPs** in the management GUI. The service IP addresses for all the nodes are configured during the installation of the software in the supported cloud environment. You can also optionally set up a service IP address and remote support assistance.
-
- You must also configure a Remote Support Proxy server in order to configure remote support assistance in the supported cloud environment.

The following network connections between IBM and the system are required to enable support assistance.

esupport.ibm.com

The esupport.ibm.com network connection is used to upload logs to the IBM Enhanced Customer Data Repository (ECUREP). An esupport.ibm.com firewall rule is not necessary if Storage Insights is configured because Storage Insights provides a feature to upload logs. However, an esupport.ibm.com firewall rule is still recommended because Call Home with cloud services uses the same port.

Note: The esupport.ibm.com network connection is fully certified to securely transmit data for Blue Diamond (HIPPA) users and General Data Protection Regulation (GDPR) protected users.

Use the following information to configure a firewall rule.

Source	Target	Port	Protocol	Direction
The service IP address of every node or node canister.	esupport.ibm.com	443	https	Outbound only

If a transparent proxy service is available in the management network, then no firewall rules are required for esupport.ibm.com. If a domain name cannot be used for configuring firewall rules, you can use the follow IP addresses: 129.42.56.189, 129.42.54.189 and 129.42.60.189.

FixCentral

Software upgrade packages can be downloaded onto the system by using the FixCentral network connection. Use the following information to configure a firewall rule.

Source	Target	Port	Protocol	Direction

Figure 9-12 Setting up Support Assistance help

Figure 9-13 shows the first window in the wizard that is used to configure the optional Remote Support Proxy (RSP).

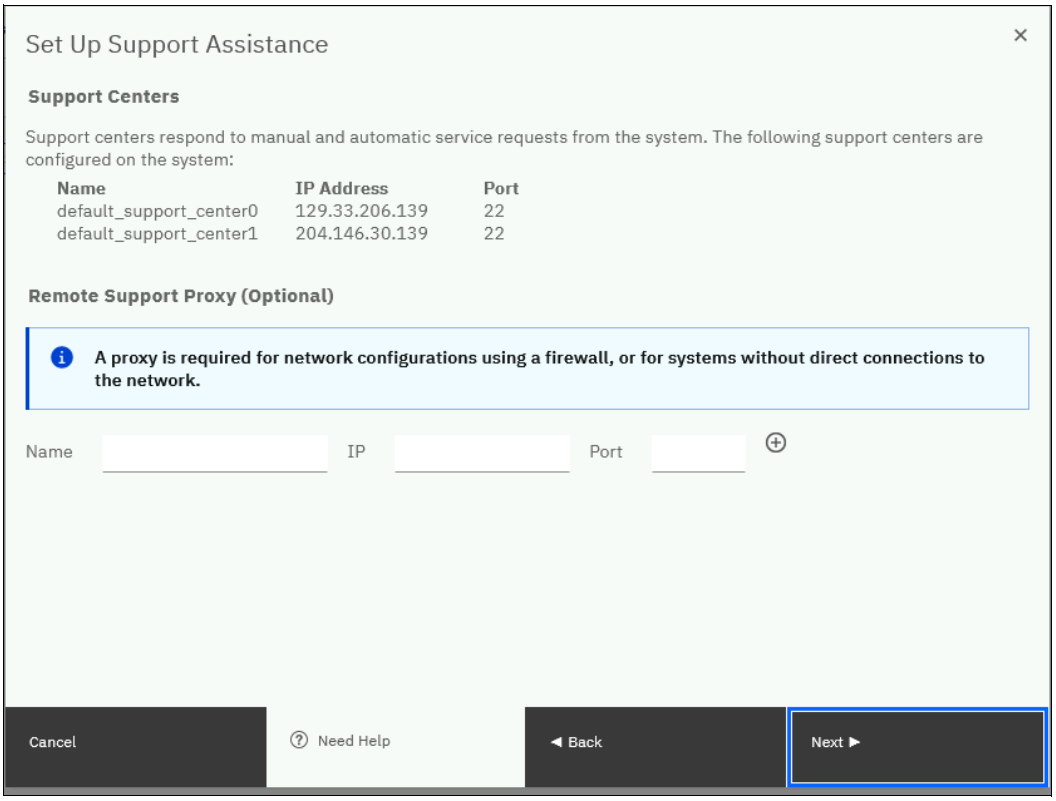

Figure 9-13 Configuring an optional Remote Support Proxy

3. In the next window (see Figure 9-14 on page 257), you are prompted to Make a choice: open a tunnel to IBM permanently, which allows IBM to connect to your IBM Spectrum Virtualize cluster at any time, or the **On Permission Only** option, which requires a storage administrator to log on to the GUI and enable the tunnel when required. Select one of the options and click **Finish**.

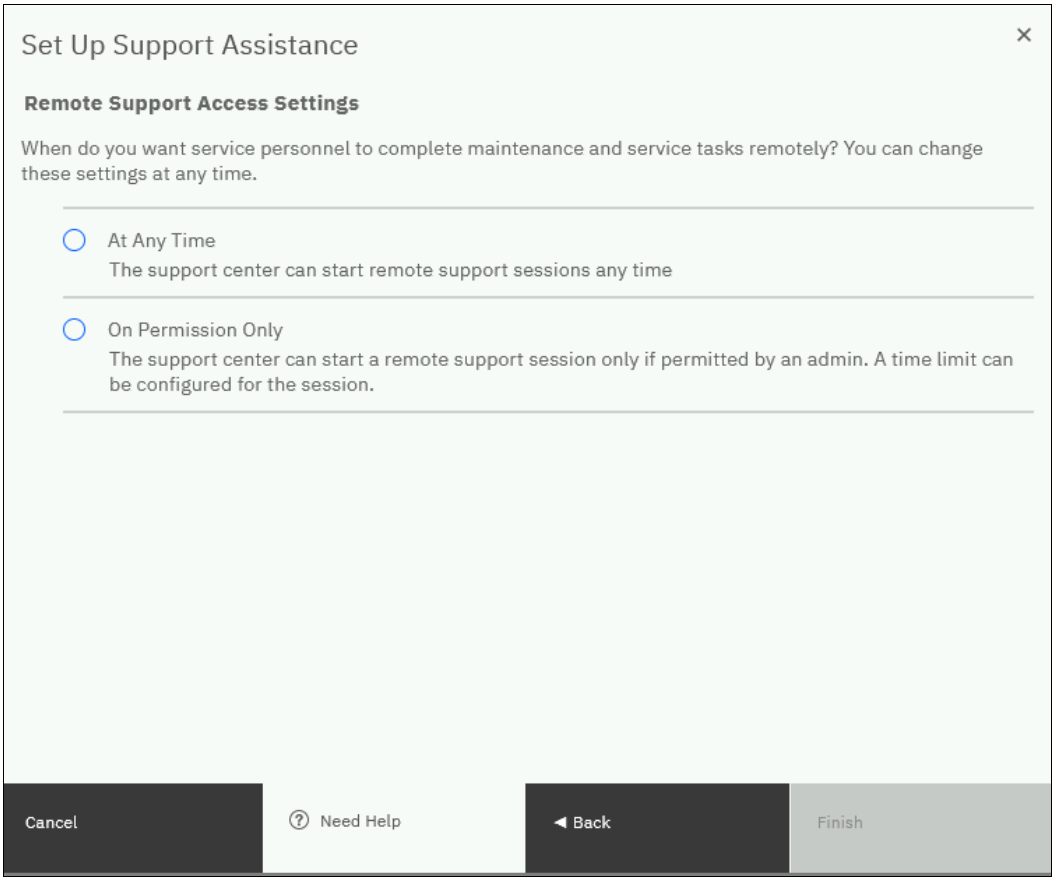

Figure 9-14 Remote Support Access settings

4. After the remote support setup is completed, you can view the status of any remote connection, start a new session, test the connection to IBM, and reconfigure the setup. Figure 9-15 shows a successfully tested connection. Now, click **Start New Session** to open a tunnel for IBM Support to connect.

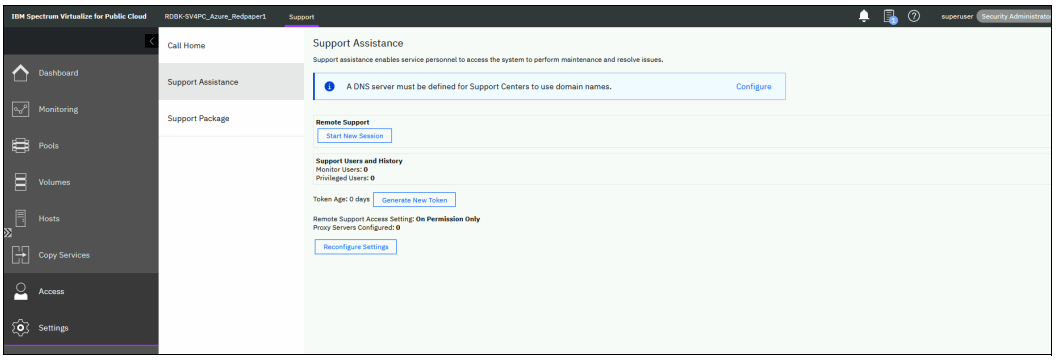

Figure 9-15 Remote Support Status and session management

A window opens, in which you are promoted to decide how long you want the tunnel to remain open if no activity occurs by setting a timeout value. Then, the connection establishes and is waiting for IBM Support to connect.

5. To disable Remote Support Assistance, restart the configuration wizard by reconfiguring the settings or by using the CLI command **`chsra -disable`**. To disable remote support assistance through the GUI, clear the checkmark for **Set Up support Assistance** in the first window of the configuration wizard. Then, click **Next** (see Figure 9-16).

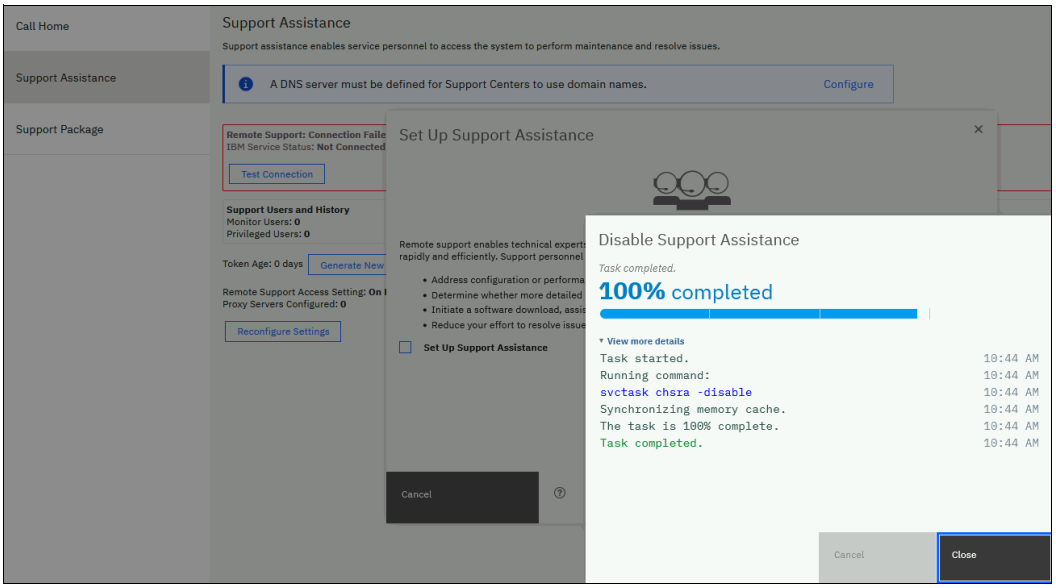

Figure 9-16 Disabling remote Support Assistance

9.6 Troubleshooting and fix procedures

The management GUI is a browser-based interface for configuring and managing all aspects of your system. It provides extensive facilities to help troubleshoot and correct problems. This section explains how to effectively use its features to avoid service disruption of your system.

Use the management GUI to manage and service your system. Select **Monitoring** → **Events** to list events that should be addressed and maintenance procedures that walk you through the process of correcting problems.

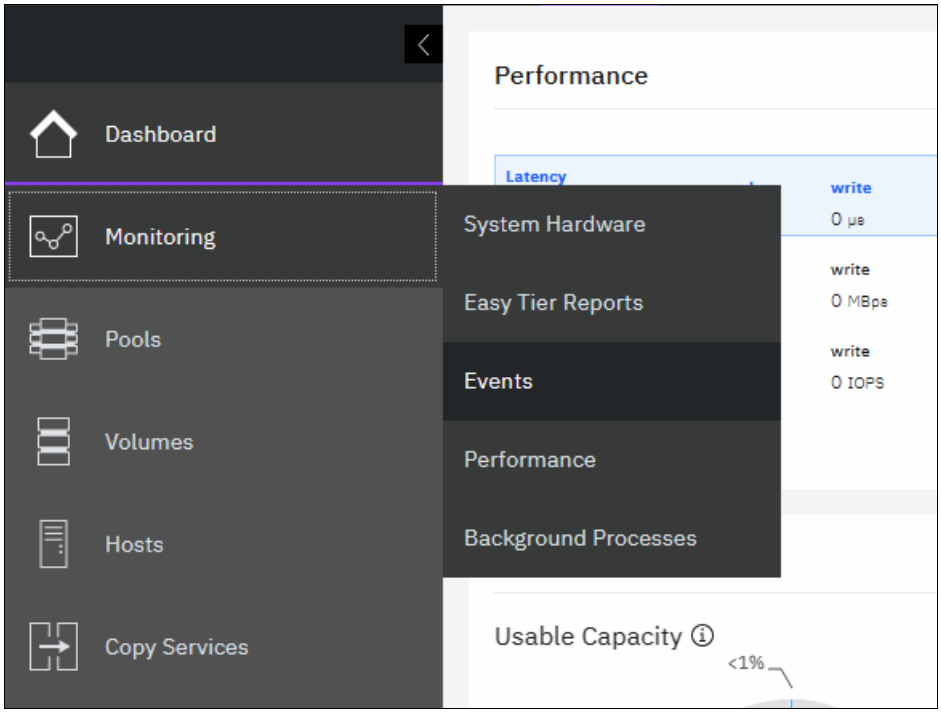

Figure 9-17 Events within Monitoring options

Information in the Events window can be filtered four ways:

- Recommended Actions

  Shows only the alerts that require attention. Alerts are listed in priority order, and should be resolved sequentially by using the available fix procedures. For each problem that is selected, you can perform the following tasks:

  - Run a fix procedure.
  - View the properties.

- Unfixed Alerts

  Displays only the alerts that are not fixed. For each entry that is selected, you can perform the following tasks:

  - Run a fix procedure.
  - Mark an event as fixed.
  - Filter the entries to show them by specific minutes, hours, or dates.
  - Reset the date filter.
  - View the properties.

- Unfixed Messages and Alerts

  Displays only the alerts and messages that are not fixed. For each entry that is selected, you can perform the following tasks:

  - Run a fix procedure.
  - Mark an event as fixed.
  - Filter the entries to show them by specific minutes, hours, or dates.
  - Reset the date filter.
  - View the properties.

- Show All

  Displays all event types whether they are fixed or unfixed. For each entry that is selected, you can perform the following tasks:

  - Run a fix procedure.
  - Mark an event as fixed.
  - Filter the entries to show them by specific minutes, hours, or dates.
  - Reset the date filter.
  - View the properties.

Some events require some occurrences in 25 hours before they are displayed as unfixed. If they do not reach this threshold in 25 hours, they are flagged as expired. Monitoring events are below the coalesce threshold, and are transient.

9.7 Managing the event log

Regularly check the status of the system by using the management GUI. If you suspect a problem, first use the management GUI to diagnose and resolve the problem. Use the views that are available in the management GUI to verify the status of the system, the hardware devices, the physical storage, and the available volumes by completing the following steps:

1. Select **Monitoring** → **Events** to see all problems that exist on the system

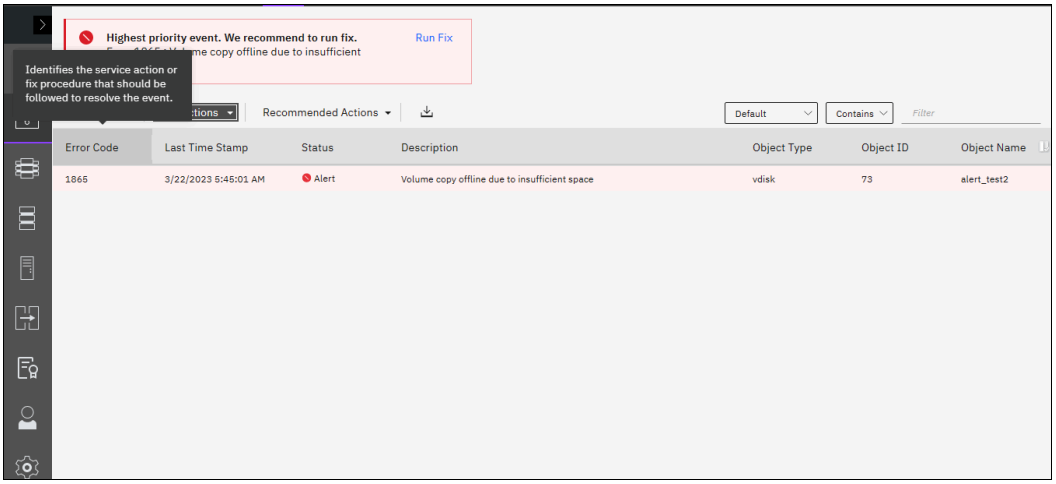

Figure 9-18 Messages in the event log

2. Select **Recommended Actions** from the drop-down list to display the most important events to be resolved. The **Recommended Actions** tab shows the highest priority maintenance procedure that must be run. Use the troubleshooting wizard so that the system can determine the proper order of maintenance procedures.

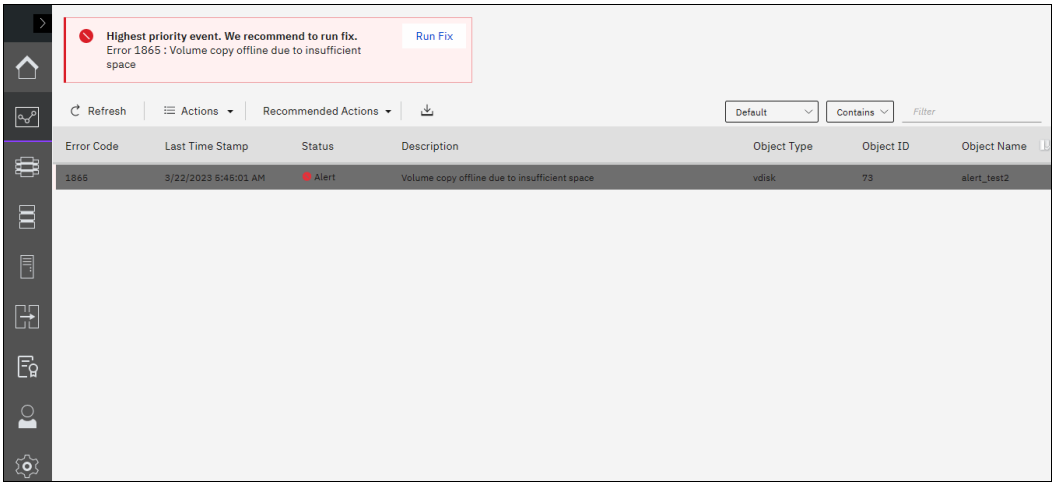

Figure 9-19 Recommended Actions

In this example, there is a canister that has a fault (service error code 1034). At any time and from any GUI window, you can directly go to this menu by clicking the **Status Alerts** icon at the top of the GUI (see Figure 9-20).

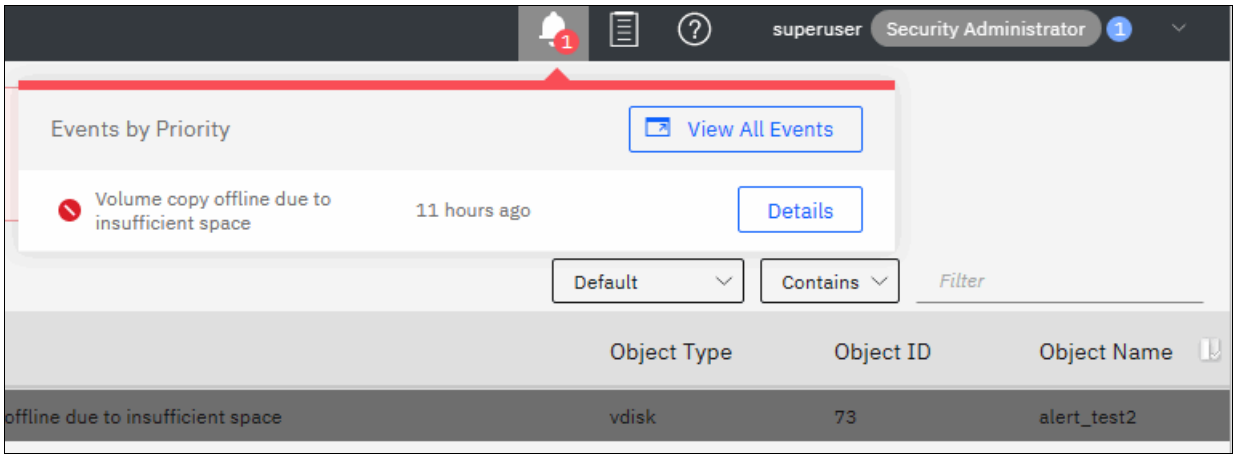

Figure 9-20 Status Alerts

9.7.1 Running a fix procedure

If an error code exists for the alert, run the fix procedure to help you resolve the problem. These fix procedures analyze the system and provide more information about the problem. They suggest actions to take and walk you through the actions that automatically manage the system where necessary while ensuring availability. Finally, they verify that the problem is resolved.

If an error is reported, always use the fix procedures from the management GUI to resolve the problem for both software configuration problems and hardware failures. The fix procedures analyze the system to ensure that the required changes do not cause volumes to become inaccessible to the hosts. The fix procedures automatically perform configuration changes that are required to return the system to its optimum state.

The fix procedure displays information that is relevant to the problem, and it provides various options to correct the problem. Where possible, the fix procedure runs the commands that are required to reconfigure the system.

The fix procedure also checks that any other existing problems do not result in volume access being lost. For example, if a PSU in a node enclosure must be replaced, the fix procedure checks and warns you whether the integrated battery in the other PSU is not sufficiently charged to protect the system.

Note: Always use **Run Fix**, which resolves the most serious issues first. Often, other alerts are corrected automatically because they were the result of a more serious issue.

Resolving alerts in a timely manner

To minimize any impact to your host systems, always perform the recommended actions as quickly as possible after a problem is reported. Your system is resilient to most single hardware failures.

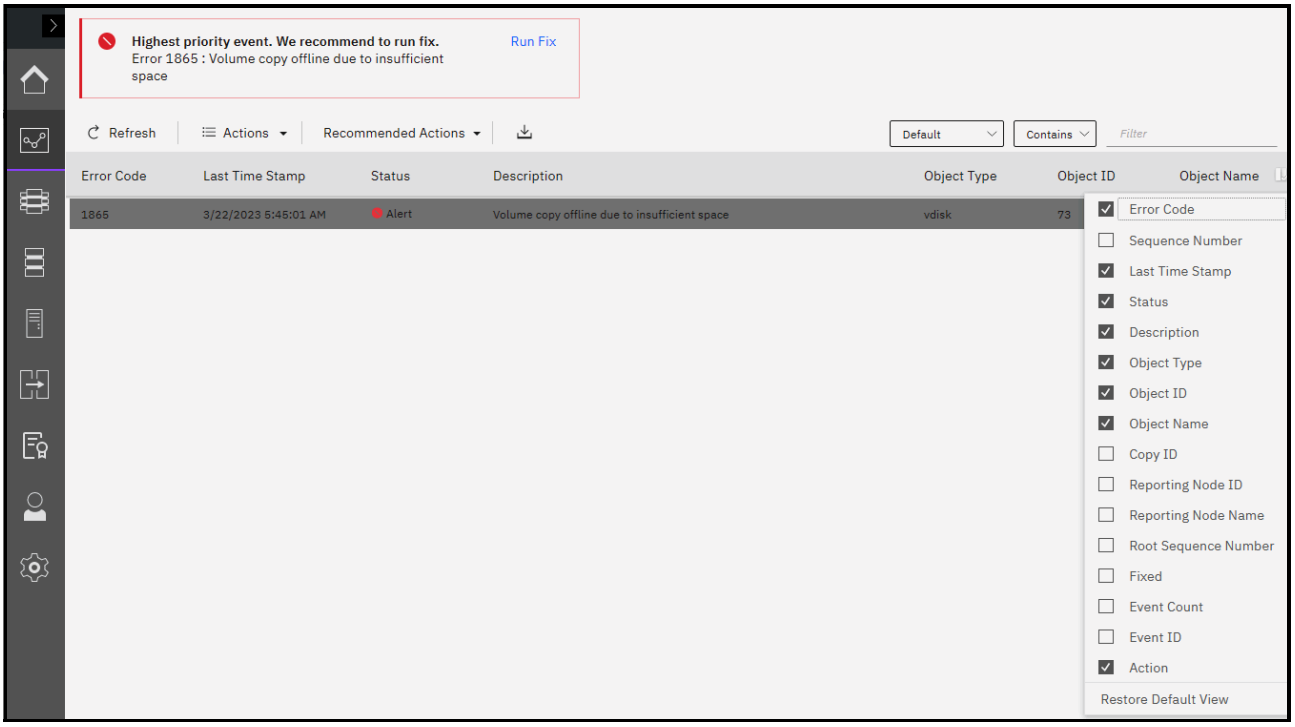

Figure 9-21 Grid options for the event log

However, if it operates for any period with a hardware failure, the possibility increases that a second hardware failure can result in some volume data unavailability. If several unfixed alerts exist, fixing any one alert might become more difficult because of the effects of the others.

9.7.2 Event log details

Multiple views of the events and recommended actions are available. When you click the column icon at the right end of the table heading, a menu for the column choices opens.

Select or remove columns as needed. You can also extend or shrink the width of columns to fit your window resolution and size. This method is relevant for most windows in the management GUI.

Every field of the event log is available as a column in the event log grid. Several fields are useful when you work with IBM Support. The preferred method in this case is to use the **Show All** filter, with events sorted by timestamp. All fields have the sequence number, event count, and the fixed state. Clicking **Restore Default View** sets the grid back to the defaults.

You might want to see more details about each critical event. Some details are not shown in the main grid. To access the properties and sense data of a specific event, double-click the specific event anywhere in its row.

The properties window opens with all the relevant sense data. This data includes the first and last time of an event occurrence, number of times the event occurred, worldwide port name (WWPN), worldwide node name (WWNN), enabled or disabled automatic fix, and other information.

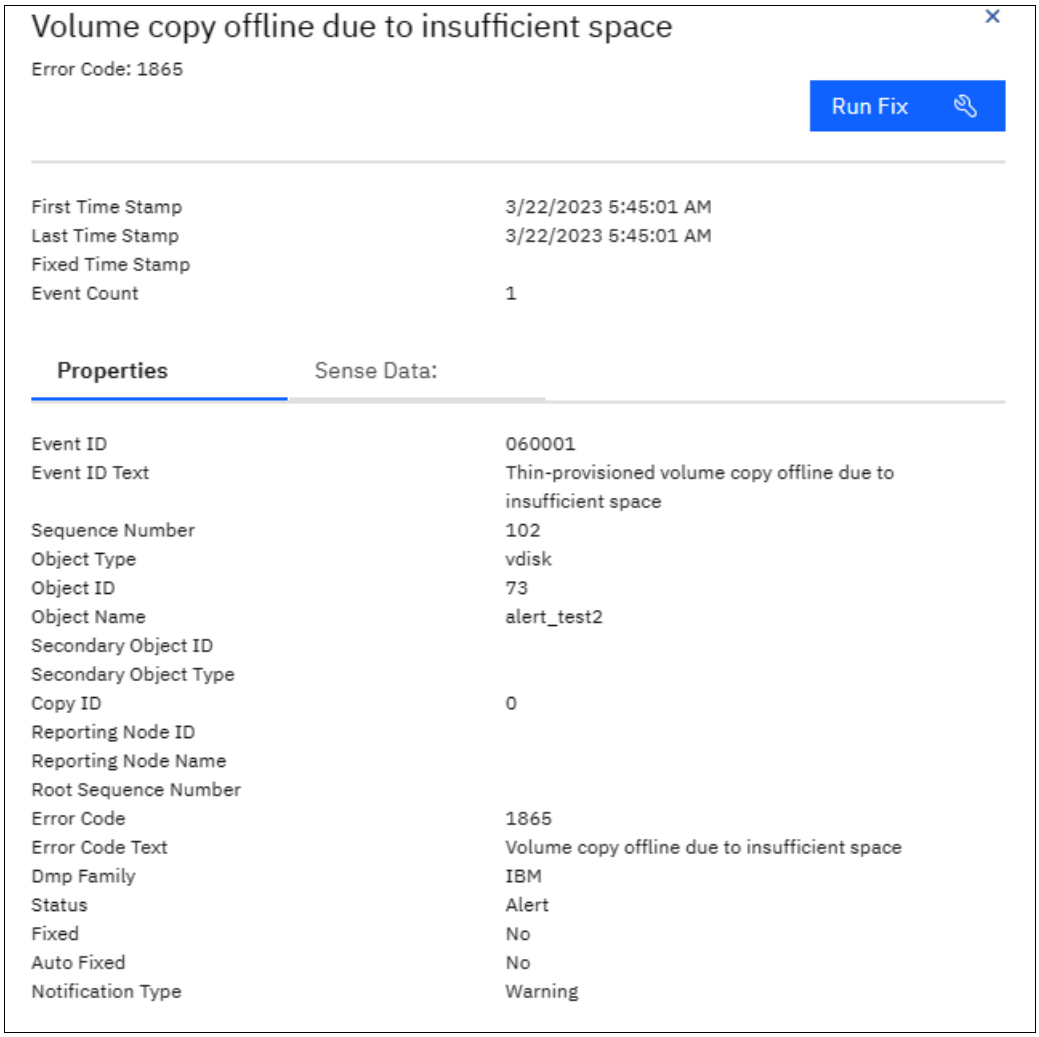

Figure 9-22 Event Sense Data and Properties

9.8 Troubleshooting in Microsoft Azure

In this section, we provide some examples of how to troubleshoot problems in Microsoft Azure.

9.8.1 Enabling boot diagnostics

Debug Microsoft Azure virtual machine (VM) boot problems by completing the following steps to enable boot diagnostics for the IBM Spectrum Virtualize node VMs:

1. In the Microsoft Azure portal, select the `node-vm` from your defined resource group (RG).
2. In the Help that is in left window, select **Boot diagnostic** → **settings**.
3. To connect to serial console, select **Enabling boot diagnostics with custom storage account**; otherwise, use the recommended option (see Figure 9-23).

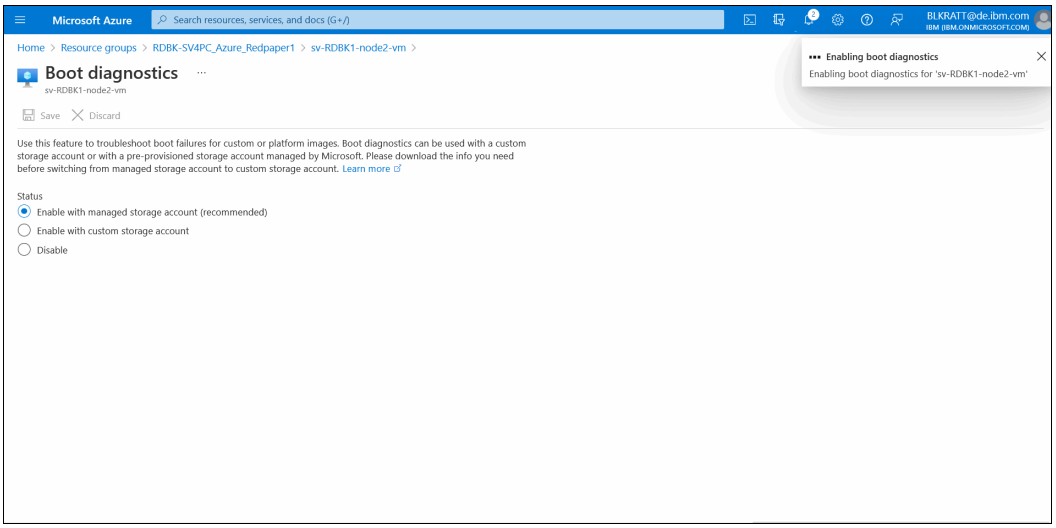

Figure 9-23 Microsoft Azure boot diagnostics

4. A node VM with boot diagnostics enabled collects serial log information and screen captures from start time (see Figure 9-24 on page 265).

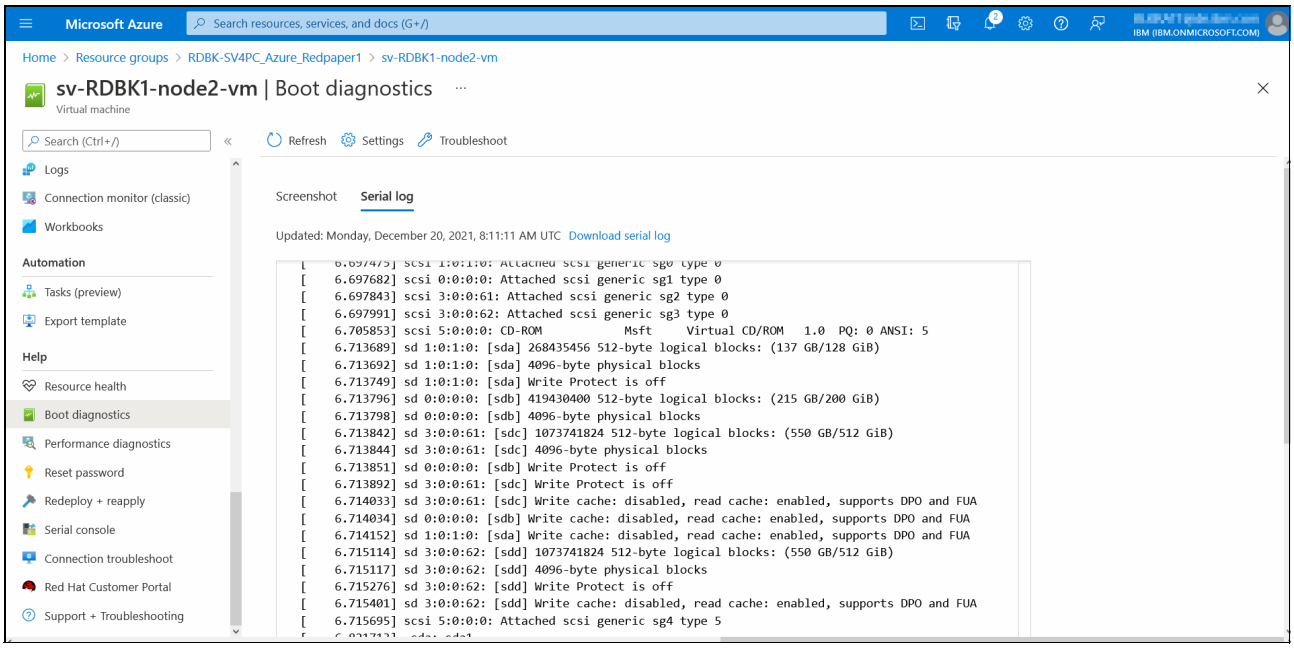

Figure 9-24 Microsoft Azure Boot diagnostics

For more information about Microsoft Azure Boot diagnostics, see this Microsoft Docs web page.

9.8.2 Connecting to a serial console

To monitor your node-vm during boot, connect the serial console. From Help in left window, select **Serial console**.

Before connecting to a serial console, you must enable **Boot diagnostic with custom storage account** (see Figure 9-25). For more information, see this Microsoft Docs web page.

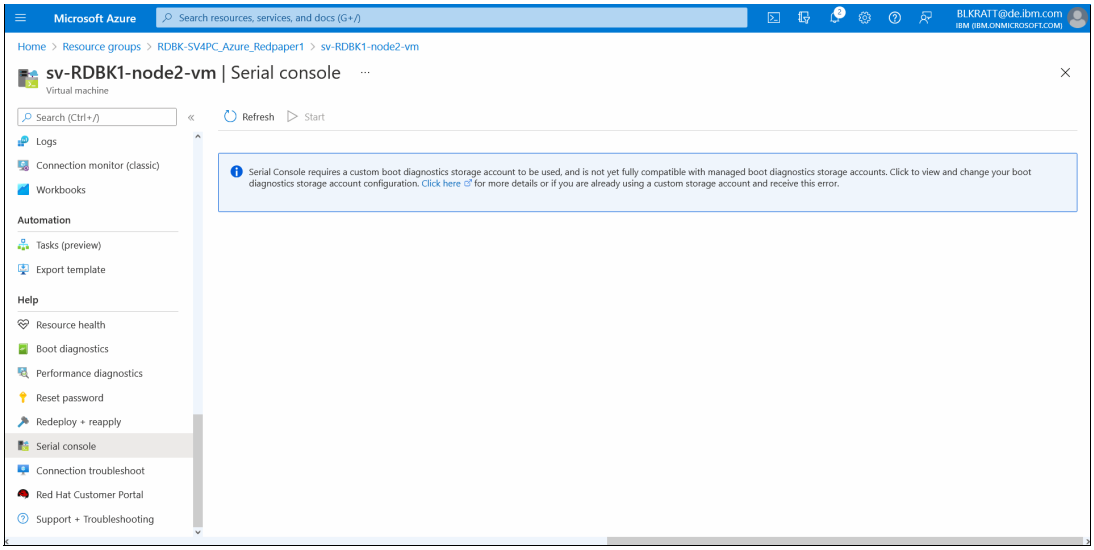

Figure 9-25 Serial console requirements

Figure 9-26 shows the serial console view after you enable a boot diagnostic with a custom storage account.

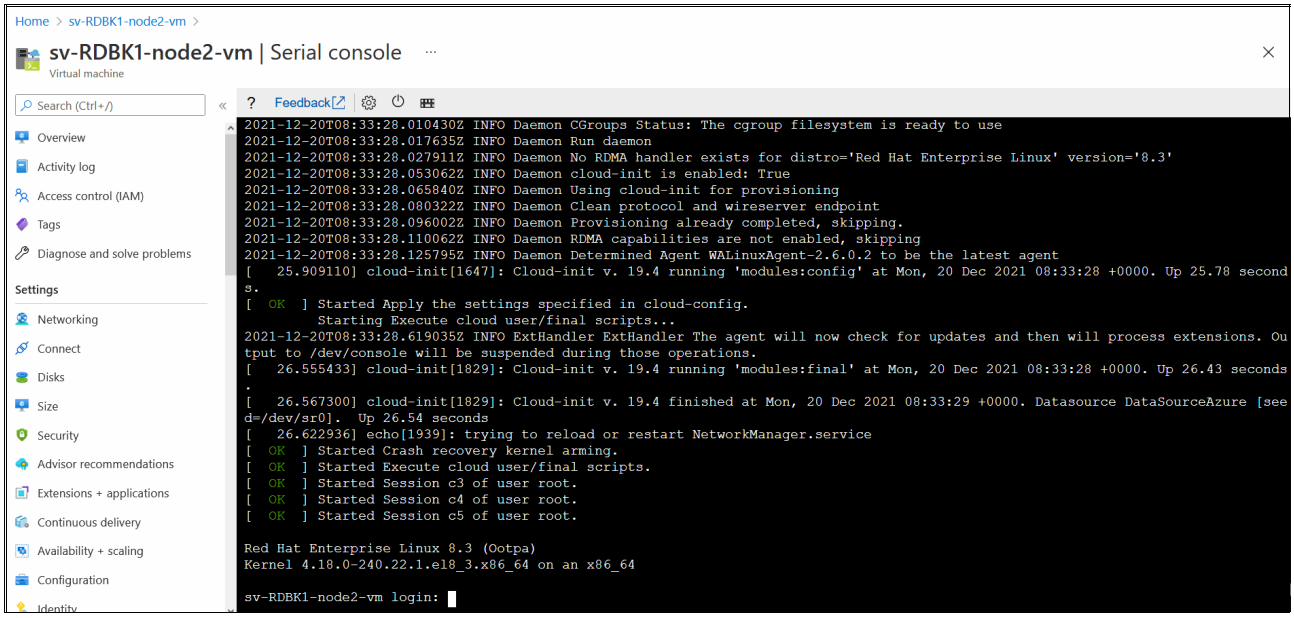

Figure 9-26 Serial console in Microsoft Azure

9.8.3 Deployment errors

This section provides insight into the potential causes for deployment failures. Depending on where the failure occurs in the deployment process, different troubleshooting methods are required.

Initial deployment errors

If the deployment fails in the initial phases due to a resource creation failure, an email notification is not sent and the final output logs are not generated.

Complete the following steps:

1. The first action to take is to access the Azure portal with the installer or administrator account.

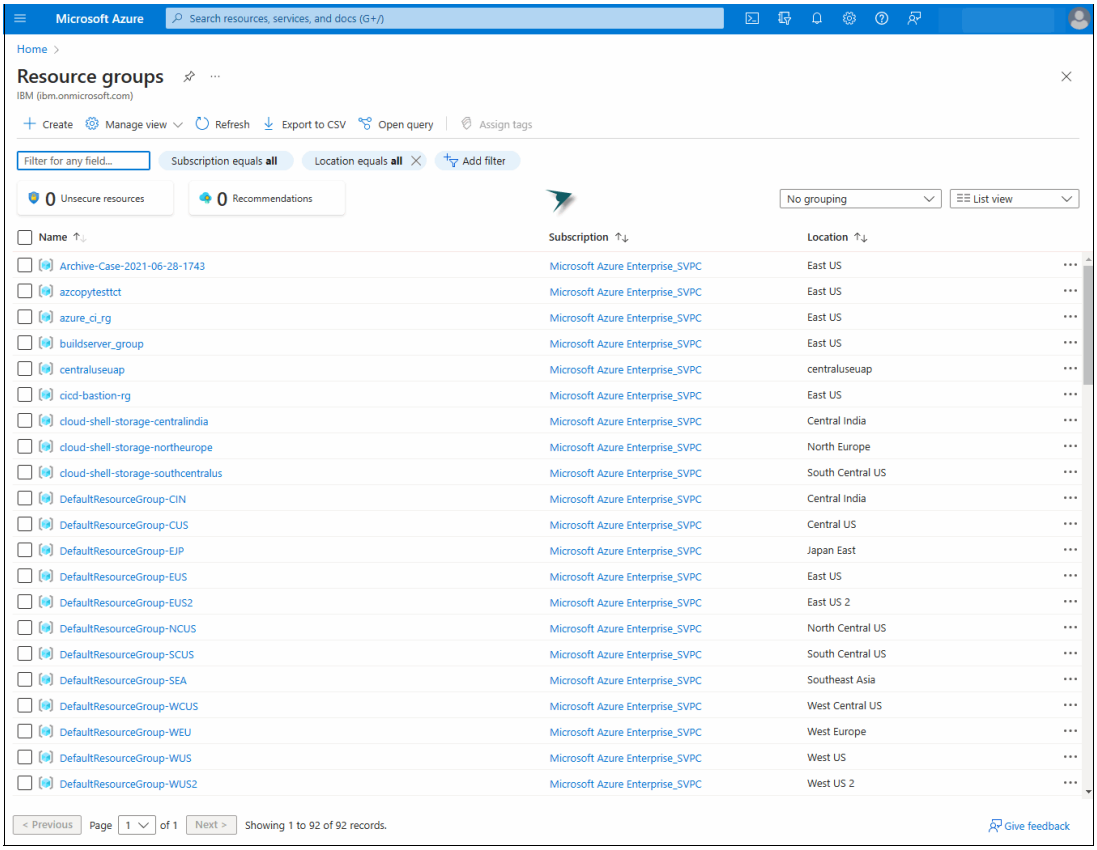

Figure 9-27 Resource groups view

2. Within the Resource groups view, you can filter the content of the view to limit the number of entries that are displayed (Figure 9-28).

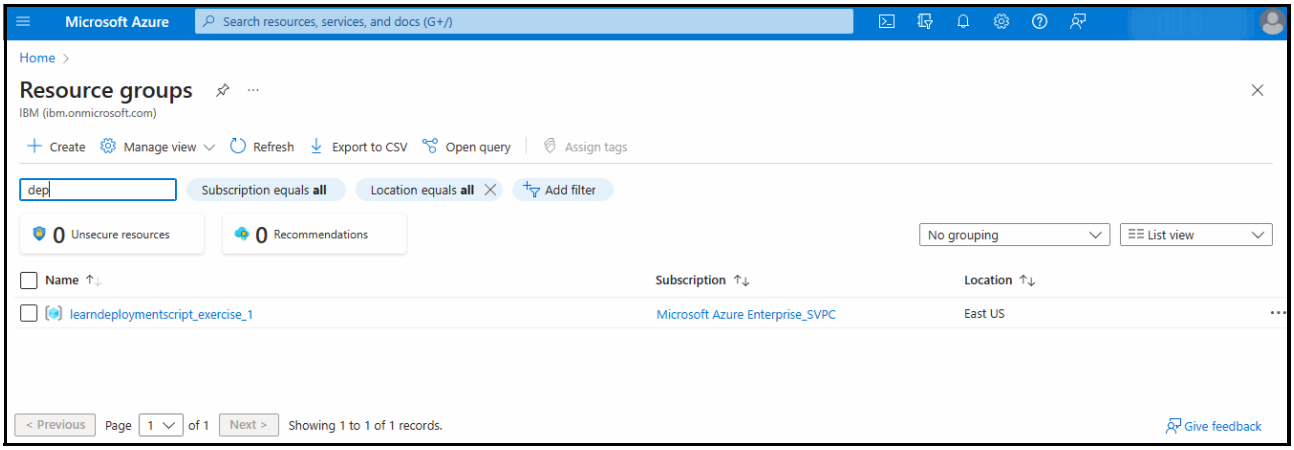

Figure 9-28 Selected resource group

3. Clicking the RG name shows the group details. To show the deployment details Within the navigation bar, select **Deployments** in the Settings pane, as shown in Figure 9-29.

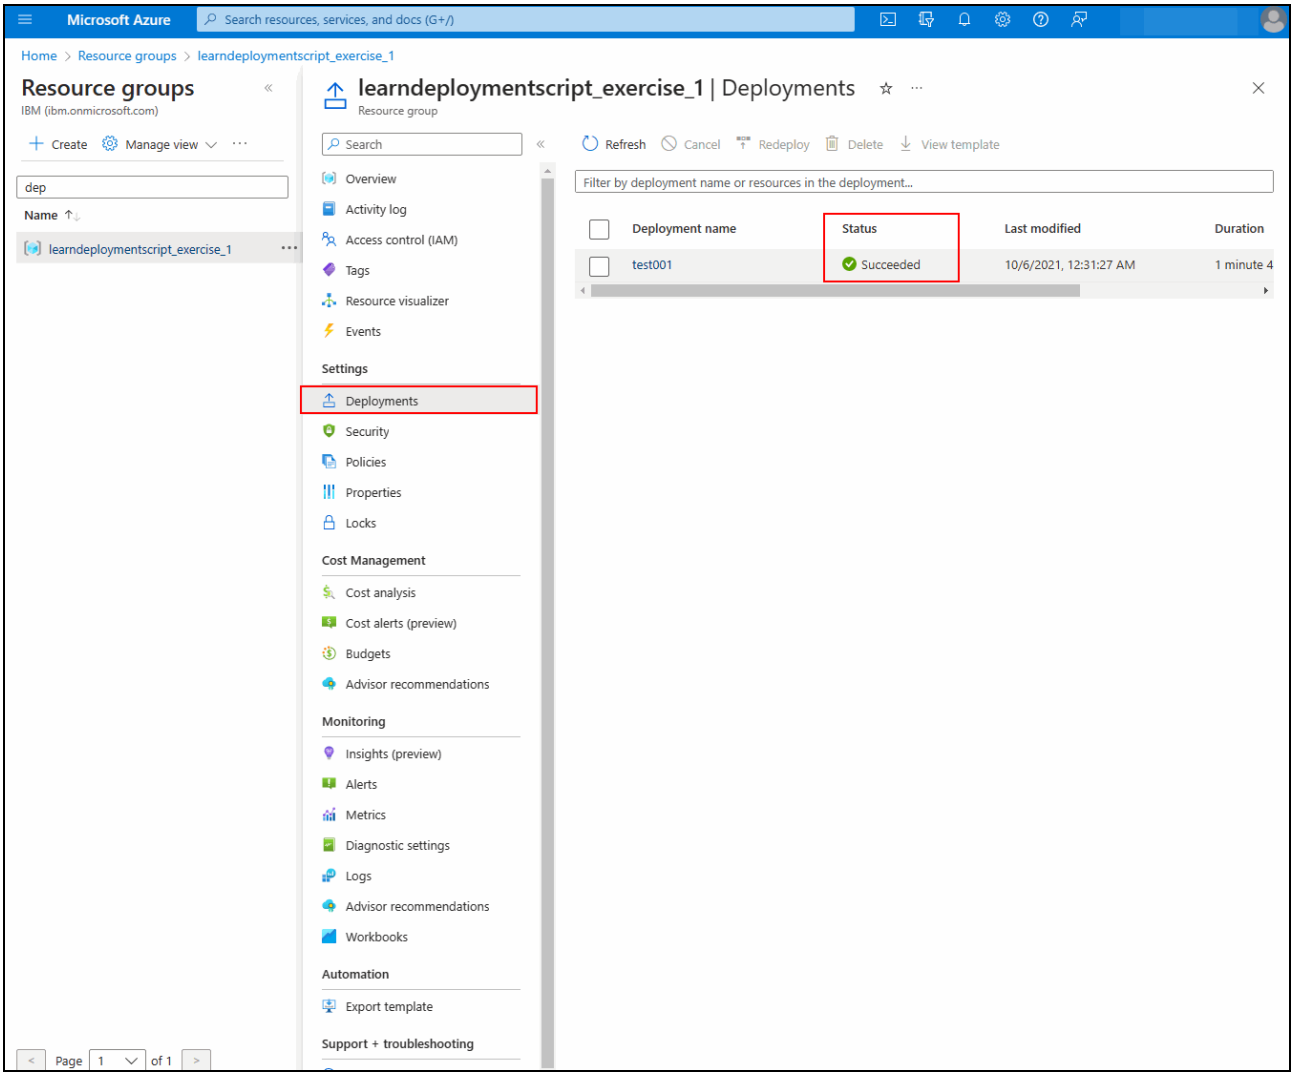

Figure 9-29 Deployment view

4. When an error is encountered during deployment the status field will be **Failed**. Selecting the name provides a summary of the actions that are taken during the deployment. Error codes provide an extensive set of information that can be used for understanding a failure.

Deployment configuration errors

If a deployment fails because of an IBM Spectrum Virtualize for Public Cloud configuration error, an email that contains details about the error is sent. The same information is available through the Azure portal.

Complete the following steps:

1. When access to the portal is available, find the deployment within the Resource groups view, and select **Overview** (see Figure 9-30 on page 269). Errors are listed in the Essentials pane.

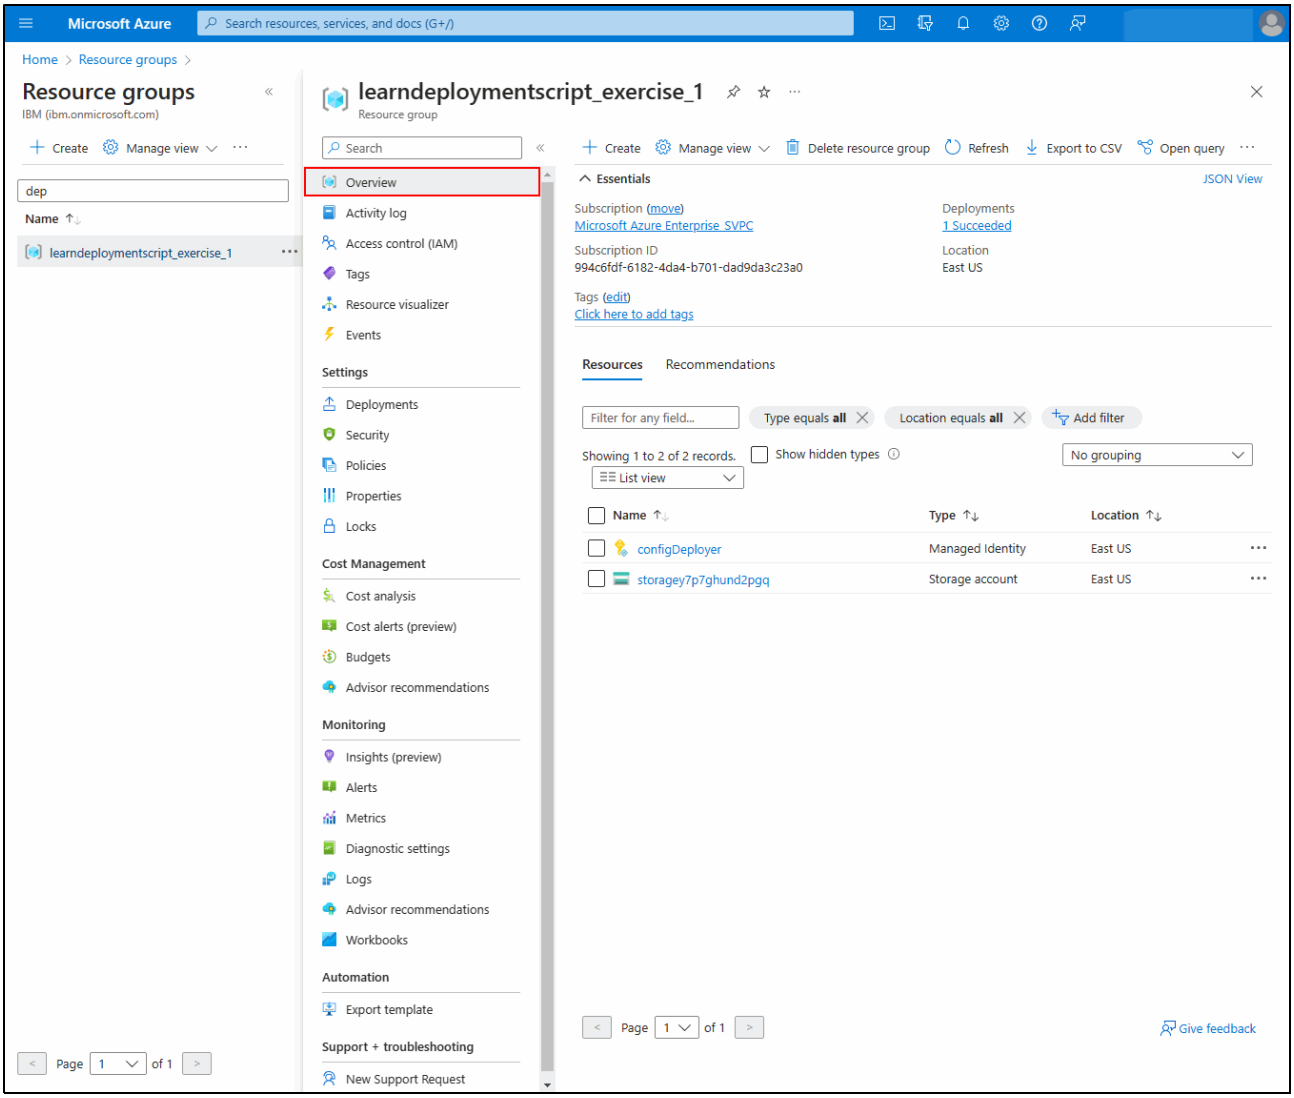

Figure 9-30 Overview

Table 9-1 contains information about common configuration errors.

Table 9-1 Common configuration errors

Error	Corrective action
NODE_DEPLOYMENT_ERROR_RPM_NOT_INSTALLED	Check your Azure network connectivity, and fix any network errors. Retry the deployment. If the problem persists, contact IBM Support.
NODE_DEPLOYMENT_ERROR_IMAGE_DOWNLOAD_FAILED_VERSION	Check your Azure network connectivity, and fix any network errors. Retry the deployment. If the problem persists, contact IBM Support.

Error	Corrective action
NODE_DEPLOYMENT_ERROR_ENTITLEMENT_CHECK_FAILED	Verify the customer ID that is used. Update and retry the deployment if the incorrect ID was used. If a correct ID was used, contact IBM Support.
CLUSTER_DEPLOYMENT_ERROR_PASSWORD_FAILED	Input the correct password and retry the deployment. If the problem persists, contact IBM Support.
CLUSTER_DEPLOYMENT_ERROR_KEYFILE_FAILED	Retry the deployment with the correct key file. If the problem persists, contact IBM Support.
ANY_OTHER_ERROR	Retry the deployment. If the problem persists, contact IBM Support.

Deleting failed deployments and resources

If deployment fails, all the resources from the failed deployment must be deleted before you attempt the deployment again. These resources can potentially conflict if you want to create a new deployment with the same settings. As part of the deployment template, you can optionally choose a rollback option that helps remove resources automatically in a failed deployment. If an error occurs in either the template or with the IBM Spectrum Virtualize for Public Cloud configuration settings, the rollback option automatically deletes resources to prevent a redeployment from failing. However, some objects still need to be deleted manually if the rollback option is selected.

If the rollback option is selected, the following resources need to be deleted manually:

- RG

  Manually delete the RG first. Deleting an RG also deletes the Virtual Network, Subnets, and Rollback script created by the deployment. For more information, see Delete Resource Groups in the Azure documentation.

- Key vault

  After RGs are deleted, you need to purge the key vault. If you plan to create a deployment by using the older deployment tag, it can cause conflicts with the existing key vault name. To avoid this, key vault must be purged; however, Azure supports a retention requirement that can prevent the deletion of the resource. For more information, see Azure Key Vault recovery management with soft delete and purge protection.

If the rollback option is *not* selected during the deployment, in addition to the *RG* and *Key Vault* that must be deleted manually, you must delete *Role assignments and role definitions*.

9.8.4 Azure hints and tips

This section provides some information about cluster management.

Identifying managed disks by using the Microsoft Azure view or the IBM Spectrum Virtualize CLI

To view all the attached cloud-managed disks (MDisks) in IBM Spectrum Virtualize, use the GUI or the CLI and complete the following steps:

1. In the GUI, select in **Pools** → **MDisk by pools**. Record the Cloud disk ID (Figure 9-31).

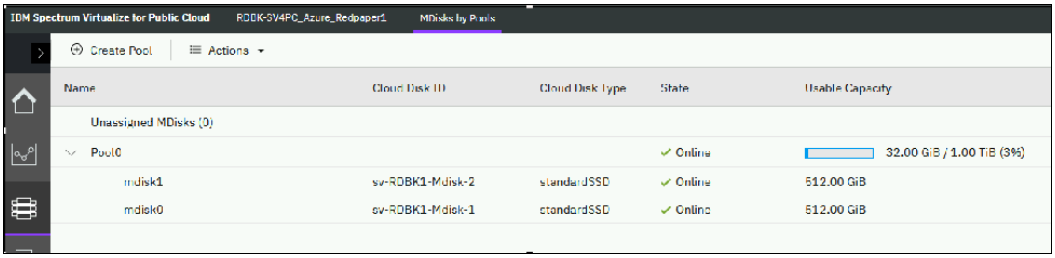

Figure 9-31 MDisks by pools

2. From the CLI, run the **`lslocaldisk`** command to view the same information. The Cloud disk ID is shown as `disk ID`. See Figure 9-32.

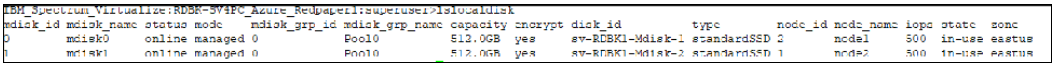

Figure 9-32 CLI command lslocaldisk

3. From the Microsoft Azure portal, select your RG, filter for `mdisks`. See Figure 9-33.

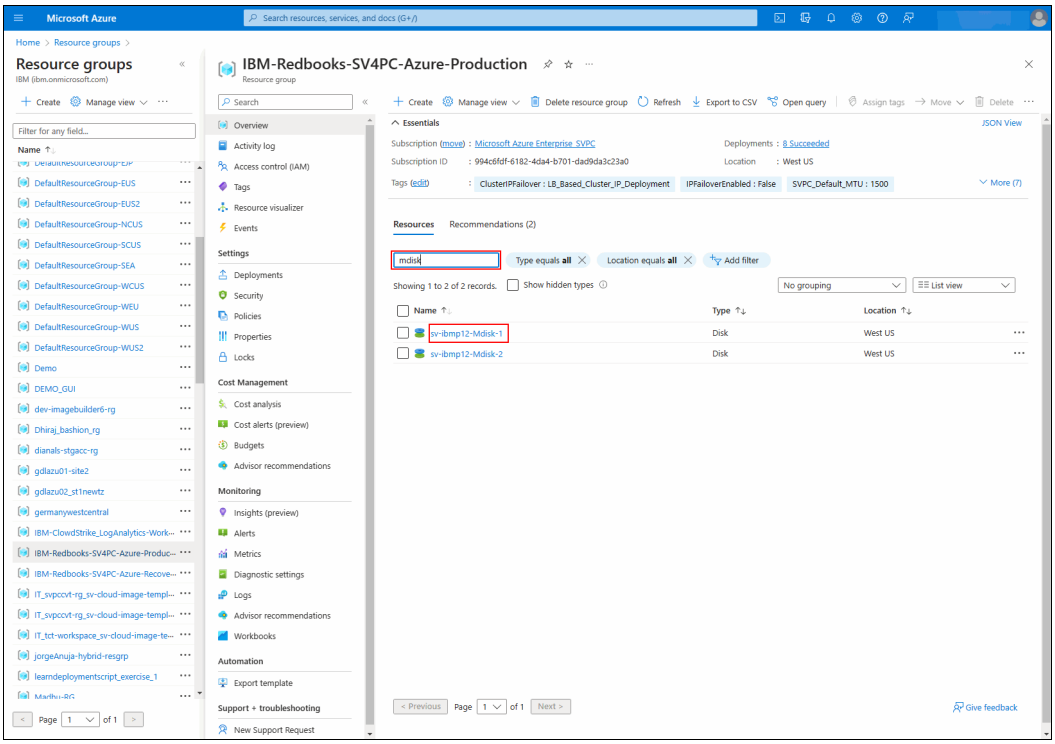

Figure 9-33 MDisk view

4. Select the MDisk to view or modify defined Tags, such as `Clustername` or `Cluster ID`. Make this selection only if you are directed to do so by the IBM Support team (see Figure 9-34).

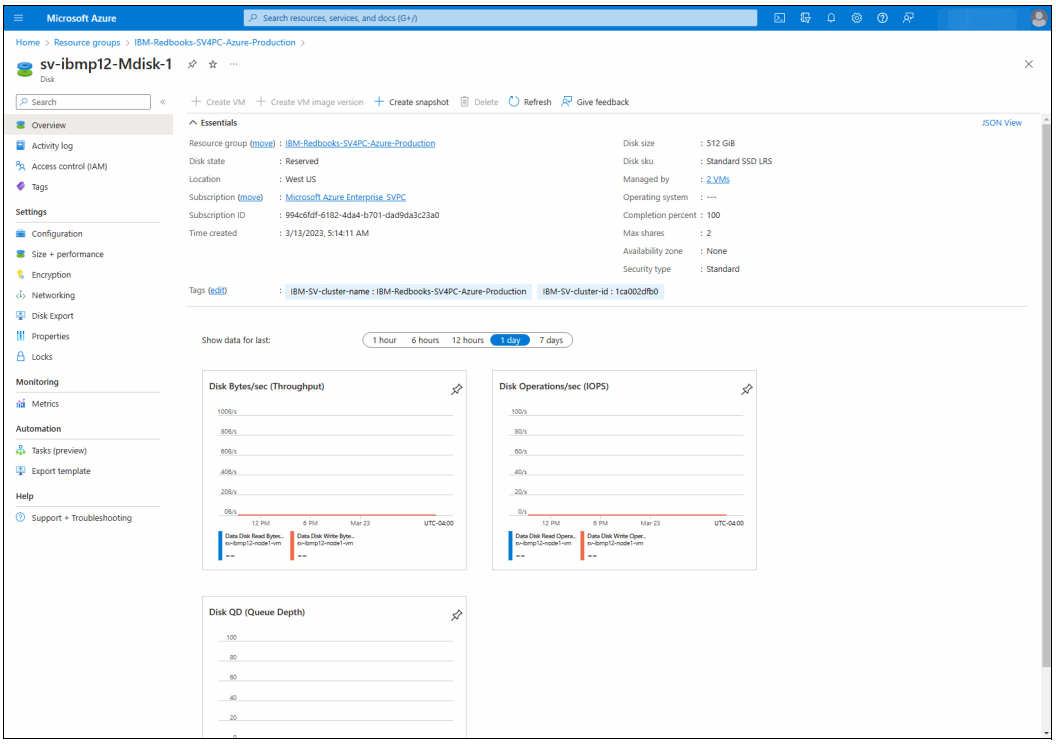

Figure 9-34 MDisk Tags view

9.9 Troubleshooting in Amazon Web Services

In this section, we provide some examples of how to troubleshoot problems in Microsoft Amazon Web Services (AWS).

9.9.1 Deployment errors

This section provides insight into the potential causes for deployment failures. Depending on where the failure occurs in the deployment process, different troubleshooting methods are required.

Initial deployment errors

If a deployment fails in the initial phases due to a resource creation issue an email nor the final output logs are not generated. To understand what occurred a user can use the AWS console. When a connection to the console is established, access the CloudFormation service and select the Region (Figure 9-35 on page 273).

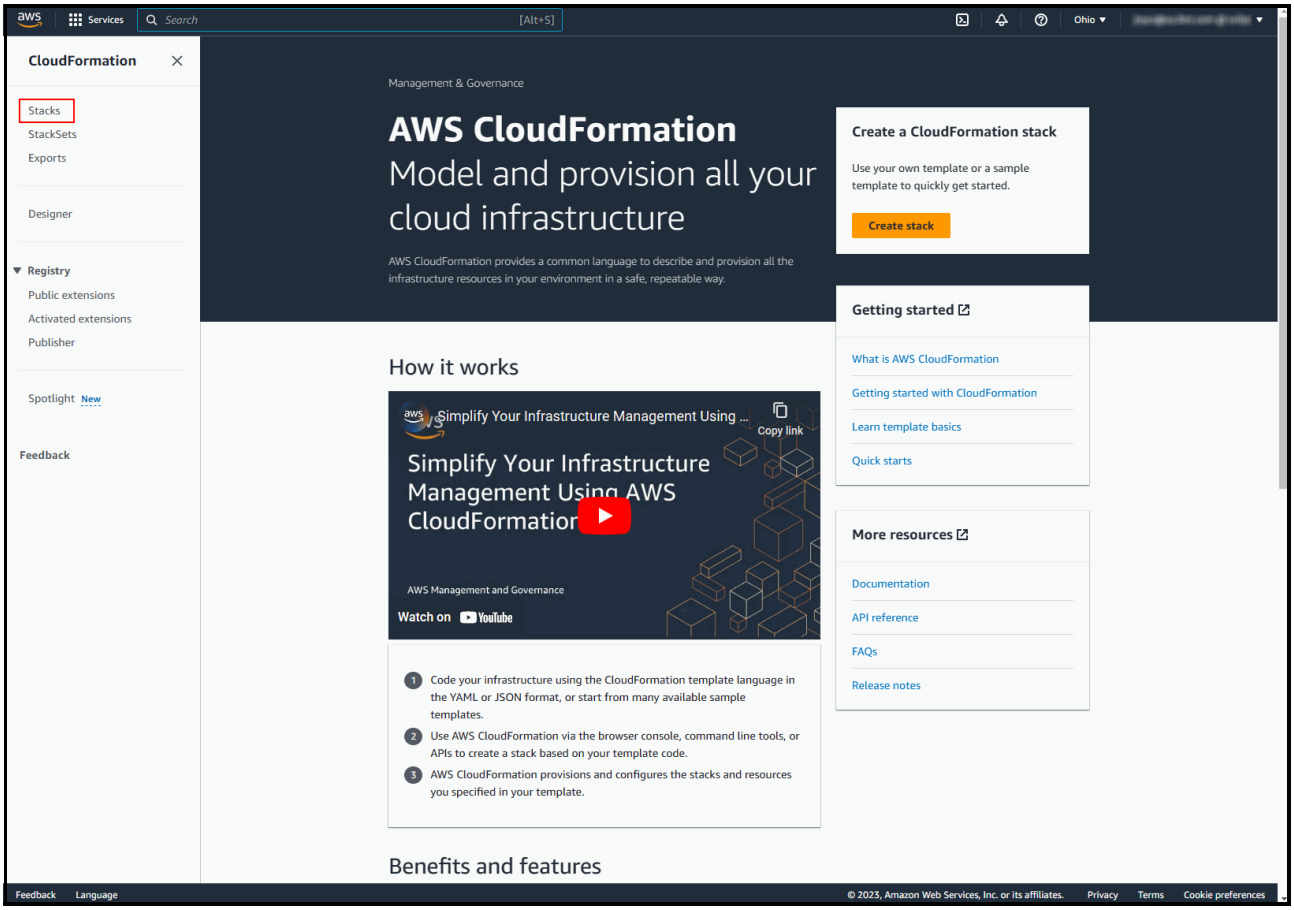

Figure 9-35 AWS CloudFormation service within the AWS portal

Select the Stack deployment with status `Create_Failed`. The Events section identifies which resource failed to create, delete, or update. For more information about errors, see Troubleshooting CloudFormation in the AWS documentation.

Deployment configuration errors

If the deployment fails because of a IBM Spectrum Virtualize for Public Cloud configuration error, you receive an email with details of errors. The errors are generated as part of the deployment, and they are available through the AWS console. The navigation options on the left enable access to Stacks (Figure 9-36).

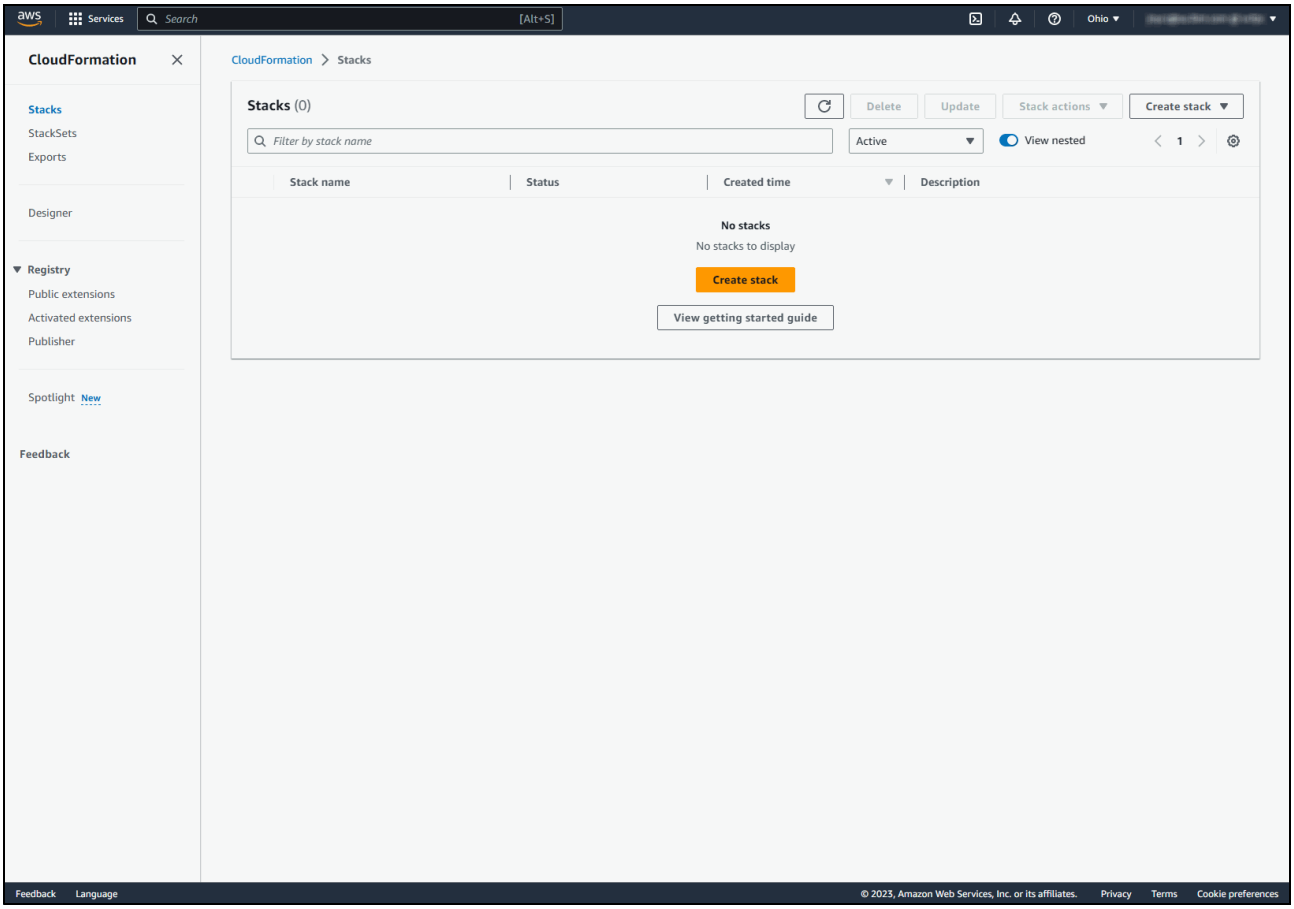

Figure 9-36 Stacks view

Selecting the name of a stack or nested stack deployment with the status `Create_Failed` provides deployment details. Check the Events section and identify which resource failed to create, delete, or update. For more information about errors, see Troubleshooting CloudFormation in the AWS documentation.

Note: Do not delete parent stack resources if they are being used by other stacks.

Table 9-2 contains information about common configuration errors.

Table 9-2 Common configuration errors

Error	Corrective action
NODE_DEPLOYMENT_ERROR_RPM_NOT_INSTALLED	Check your AWS network connectivity, and fix any network errors. Retry the deployment. If the problem persists, contact IBM Support.
NODE_DEPLOYMENT_ERROR_IMAGE_DOWNLOAD_FAILED_VERSION	Check your AWS network connectivity, and fix any network errors. Retry the deployment. If the problem persists, contact IBM Support.
NODE_DEPLOYMENT_ERROR_ENTITLEMENT_CHECK_FAILED	Verify the customer ID that is being used. Update and retry the deployment if an incorrect ID was used. If the correct ID was used, contact IBM Support.
CLUSTER_DEPLOYMENT_ERROR_PASSWORD_FAILED	Input the correct password and retry the deployment. If the problem persists, contact IBM Support.
CLUSTER_DEPLOYMENT_ERROR_KEYFILE_FAILED	Retry the deployment with the correct key file. If the problem persists, contact IBM Support.
ANY_OTHER_ERROR	Retry the deployment. If the problem persists, contact IBM Support.

Stack update errors

Errors can occur during stack updates. These errors are also accessible by using the Stacks view within the AWS console. A deployment that suffered a stack update failure has a status of `Update_Failed`. Clicking the deployment name and checking the Nested Stack Events section identifies which resource failed to create, delete, or update. For more information about errors, see Troubleshooting CloudFormation in the AWS documentation.

Deleting failed deployments and resources

As part of the deployment template, you can optionally choose a rollback option that helps remove resources automatically in a failed deployment. If an error occurs in either the template or with the IBM Spectrum Virtualize for Public Cloud configuration settings, the rollback option automatically deletes resources. However, the CloudFormation stack needs to be deleted manually. If the rollback option is not selected during the deployment, deletion of the CloudFormation stack deletes all the resources.

Note: CloudWatch logs are retained up to 7 days. Manually delete the CloudWatch logs if you do not need them.

9.10 IBM Spectrum Virtualize for Public Cloud Support

In this solution, the cloud provider is responsible for providing the infrastructure, network components and storage, and support and assistance. The cloud user or any involved third party is responsible for deploying and configuring the solution. IBM Systems support is responsible for providing support and assistance with the IBM Spectrum Virtualize application.

9.10.1 Who to call for support

The solution consists of multiple parties with different roles and responsibilities. For this reason, it is a 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 is specific for each solution deployment: how the cloud service is provided, who is the final user, who are the parties that are involved, and so on. To help with creating a workflow for handling problems when they arise, we created Table 9-3 as an example.

Table 9-3 Simplified workflow definition based on a RACI matrix

Situation	Client	Cloud provider	IBM Spectrum Virtualize
SV error 2030	Informed	Consulted	Responsible
MDisk is offline	Informed	Responsible	Accountable
Network port is down	Informed	Responsible	Consulted
Configuration question	Responsible	Consulted	Accountable

In the situations where the cloud provider is responsible or accountable, the customer must collect as much information about the problem as possible and open a ticket with the cloud provider.

In situations where IBM Spectrum Virtualize is responsible or accountable, the customer must collect as much information and diagnostic data surrounding the event about the problem as possible and open a support ticket with IBM.

In the situations where the customer is responsible, it is up to the customer to be as detailed as possible in any requests or questions that are submitted the cloud provider or IBM Spectrum Virtualize or any other third party that is involved in the support.

9.10.2 Working with IBM Support

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

- To open a case for IBM Spectrum Virtualize for Public Cloud, see the IBM Support web page.
- By phone (see the IBM Directory of worldwide contacts web page.
- Through IBM Storage Insights.
- IBM Call Home.

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

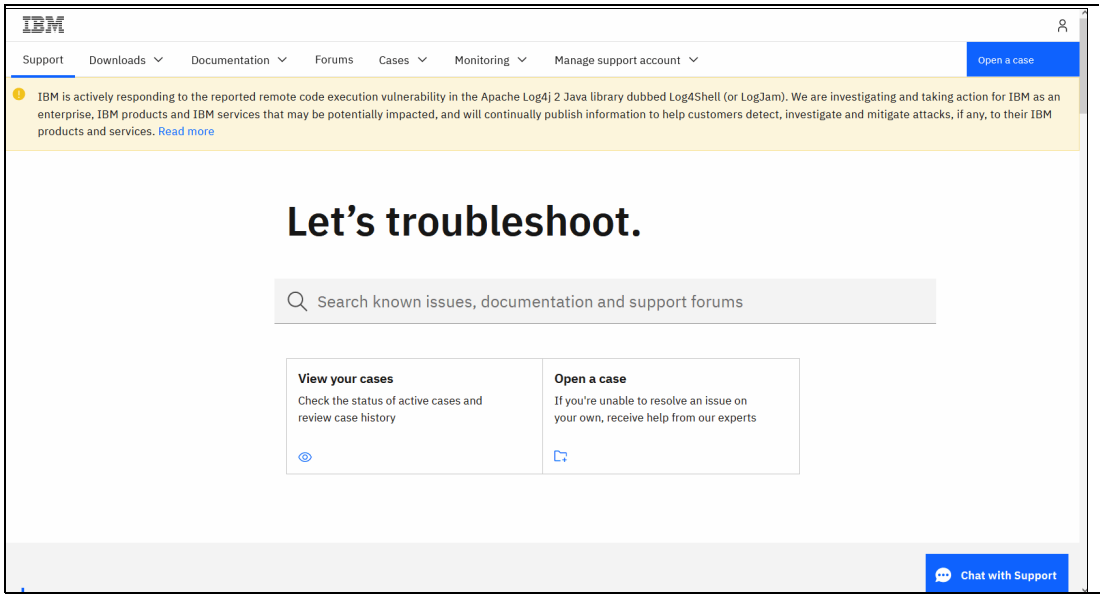

Figure 9-37 IBM Support web page

When opening a new case, select **IBM Spectrum Virtualize for Cloud** as the Product (see Figure 9-38).

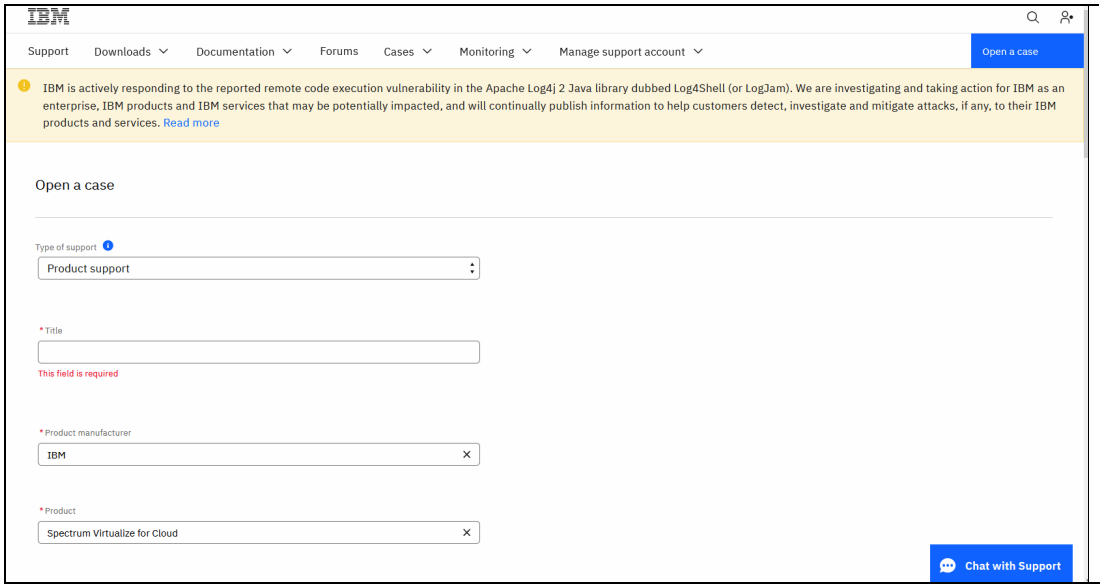

Figure 9-38 Opening a support case through the IBM Support site

Complete all fields, including severity and description.

9.10.3 Working with Microsoft Azure Support

Microsoft Azure is a service that provisions the infrastructure, network, operating systems, and back-end storage that is used in this solution. Microsoft Azure Support is responsible for helping resolve problems and answer questions about products and services that are acquired through the Microsoft Azure Marketplace portal.

Creating a support request

To engage the Microsoft Azure Support team, create a support request through the Microsoft Azure portal.

Complete the following steps:

1. In the Home view, select **Help + Support**.
2. In the Global Header, click **?** and then, select **Help + Support**.
3. For the Resource view of resource type VM, select **Support + Troubleshooting** from the left window.
4. For the Resource view of all other resource types, select **New support request**.

For more information, see this Microsoft Docs web page.

Figure 9-39 shows how to create a support request in Microsoft Azure portal by selecting **Help + support**.

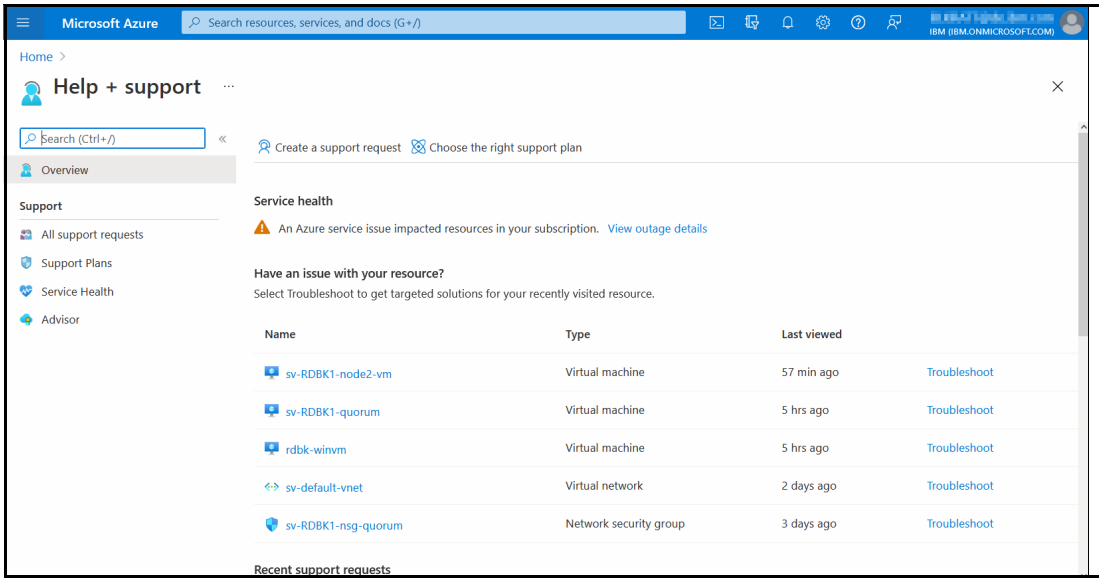

Figure 9-39 Creating a support request in Microsoft Azure from Help + support

A support request can be opened for the following issue types:

- Billing
- Service and Subscription limits
- Subscription management
- Technical

In this publication, we focus on a support request that was created from the **Resource** menu.

Figure 9-40 shows how to create a support request from the **Resource** menu.

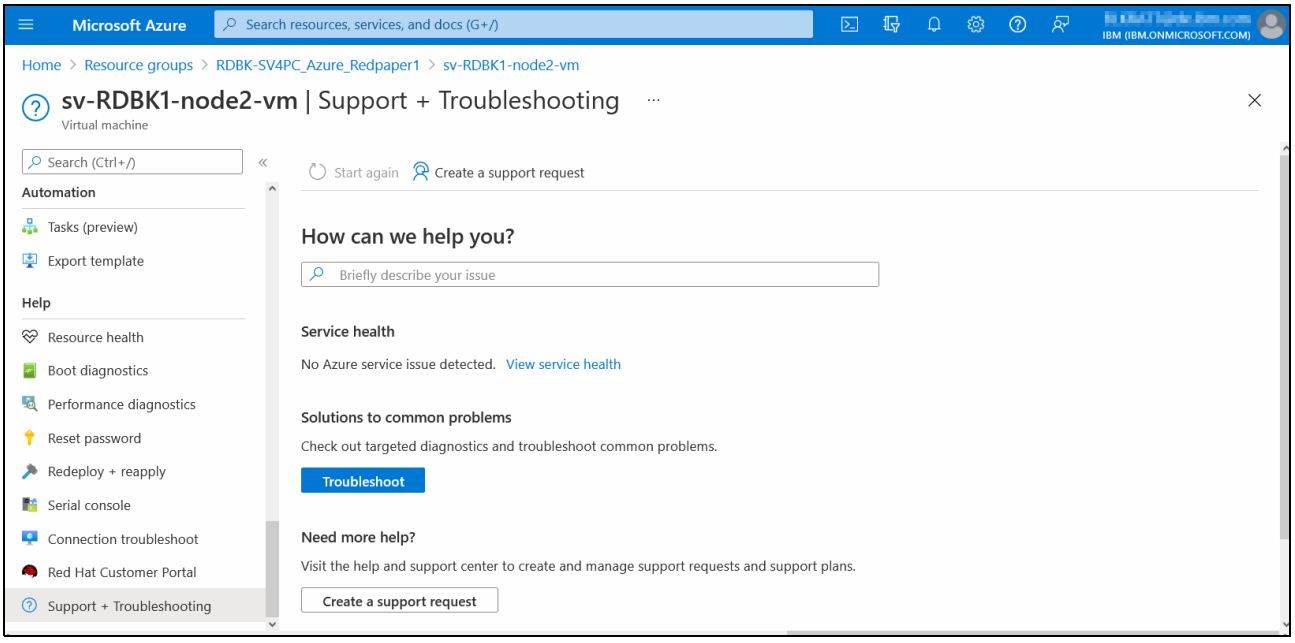

Figure 9-40 Creating a support request from the Resource menu for the resource type VM

Abbreviations and acronyms

ACL	access control list
AI	artificial intelligence
AMI	Amazon Machine Image
ARM	Azure Resource Manager
AWS	Amazon Web Services
BYOL	Bring Your Own License
CapEx	capital expenditure
CDB	Command Descriptor Block
CFT	Cloud Formation Template
CIMOM	Common Information Model Object Manager
CLI	command-line interface
CSM	Copy Services Manager
DHCP	Dynamic Host Configuration Protocol
DR	disaster recovery
DRET	Data Reduction Estimation Tool
DRP	data reduction pool
EBS	Elastic Block Store
EC2	Elastic Compute Cloud
ECuRep	Enhanced Customer Data Repository
GMCV	Global Mirror with Change Volumes
HA	high availability or highly available
HDD	hard disk drive
IaaS	infrastructure as a service
IAM	Identity and Access Management
IBM	International Business Machines Corporation
IQN	iSCSI-qualified name
ITIL	IT Infrastructure Library
LRS	locally redundant storage
MAC	Media Access Control
MDisk	managed disk
MPLS	Multiprotocol Label Switching
NDVM	Non-Disruptive Volume Move
NIC	network interface card
NPIV	N_Port ID virtualization
NSG	network security group
OpEx	operating expenditure
PaaS	platform as a service
PBR	Policy-Based Replication
PiT	point-in-time
PMP	Project Management Professional
RBAC	role-based access control
RDP	Remote Desktop Protocol
RG	resource group
RPO	recovery point objective
RSP	Remote Support Proxy
S3	Simple Storage Service
SAT	Service Assistant Tool
SDE	software-defined environment
SDS	software-defined storage
SLA	service level agreement
SPOF	single points of failure
SSD	solid-state drive
SSH	Secure Shell
SSO	single sign-on
STAT	Storage Advisor Tool
T0	time-zero
TCO	total cost of ownership
TCT	Transparent Cloud Tiering
vCPU	virtual central processing unit
VG	volume group
VHD	virtual hard disk
VM	virtual machine
vNet	virtual network
VPC	virtual private cloud
VPN	virtual private network
VPNGW	virtual private network gateway
WWNN	worldwide node name
WWPN	worldwide port name
ZRS	zone-redundant storage

Related publications

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

IBM Redbooks

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

- *Implementation Guide for IBM Spectrum Virtualize Version 8.5*, SG24-8520
- *Performance and Best Practices Guide for IBM Spectrum Virtualize 8.5*, SG24-8521
- *Policy-Based Replication with IBM Storage FlashSystem, IBM SAN Volume Controller and IBM Storage Virtualize*, REDP-5704

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

ibm.com/redbooks

Online resources

These websites are also relevant as further information sources:

- Amazon Web Services portal

  https://aws.amazon.com/
- Azure portal

  https://portal.azure.com
- IBM Spectrum Virtualize for Public Cloud on AWS - Redbooks Deployment Video

  https://mediacenter.ibm.com/media/IBM+IBMSpectrum+Virtualize+for+Public+Cloud+V8.5+installation+on+AWS/1_wlw60bds

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

REDP-5671-00

ISBN 0738461253

Printed in U.S.A.