

Implementation Guide for IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Version 8.4.3









IBM Redbooks

Implementation Guide for IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Version 8.4.3

March 2022

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

First Edition (March 2022)

This edition applies to IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Version 8.4.3

© Copyright International Business Machines Corporation 2022. All rights reserved.

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

Contents

Notices	vii . viii
Preface Authors Now you can become a published author, too! Comments welcome Stay connected to IBM Redbooks	ix ix xii xii xii
Chapter 1. Introduction. 1.1 Introduction to IBM Spectrum Virtualize for Public Cloud 1.2 IBM Spectrum Virtualize for Public Cloud 1.2.1 Primers of storage virtualization and software-defined storage 1.2.2 IBM Spectrum Virtualize for Public Cloud benefits 1.2.3 IBM Spectrum Virtualize for Public Cloud features 1.3 IBM Spectrum Virtualize for Public Cloud on Azure 1.3.1 Application examples on Azure.	1 2 3 4 5 7 8
Chapter 2. Typical use cases for IBM Spectrum Virtualize for Public Cloud 2.1 Deploying whole IT services in the public cloud 2.1.1 Business justification. 2.1.2 Highly available deployment models. 2.2 Disaster Recovery. 2.2.1 Business justification. 2.2.2 Two common DR scenarios with IBM Spectrum Virtualize for Public Cloud. 2.3 IBM FlashCopy in the public cloud 2.3.1 Business justification. 2.3.2 FlashCopy mapping 2.3.3 Consistency groups 2.3.4 Crash-consistent copy and host considerations 2.4.1 Business justification. 2.4.2 Solution design 2.4.3 Component summary 2.5 Workload relocation into the public cloud 2.5.1 Business justification. 2.5.2 Data migration. 2.5.3 Host provisioning 2.5.4 Implementation considerations	9 10 11 12 15 15 15 16 17 18 19 21 21 22 23 23 23 23 24 24
Chapter 3. Solution architecture 3.1 IBM Spectrum Virtualize 3.1.1 Nodes 3.1.2 I/O groups 3.1.3 Systems 3.1.4 MDisks 3.1.5 Storage pools 3.1.6 Child pools 3.1.7 Volumes	. 25 . 26 . 26 . 26 . 26 . 27 . 27 . 27 . 28 . 28 . 28

3.1.8 Hosts	29
3.1.9 Host clusters	29
3.1.10 iSCSI	29
3.1.11 Cache	30
3.1.12 IBM Easy Tier	31
3.1.13 IP replication	31
3.1.14 IBM FlashCopy	32
32 Azure terminology	33
3.3 Key components of the Azure solution	34
3.3.1 Azure Besource Manager template-based publishing	34
3.3.2 Resource group and white	25
2.4 Highly available infractructure	26
	30
3.4.1 Controller node proximity	37
	37
3.4.3 Load balancer for management traffic failover	38
3.4.4 Multipath for data traffic failover and better performance	39
3.5 Security design considerations	40
3.5.1 Restrictive network security rules	41
3.5.2 Bastion access	41
3.5.3 Operating system patch update interface and Azure Update Manager integration	า 41
3.5.4 Manual patching using IBM Spectrum Virtualize CLIs	42
3.5.5 Automated patching by using Azure Update Manager	42
3.6 Solution architecture: IBM Spectrum Virtualize as a storage for all-in-cloud model	43
Chapter 4. Planning an installation on Microsoft Azure	45
4.1 Introduction	46
4.2 General planning on Microsoft Azure	46
4.2.1 IBM Spectrum Virtualize for Public Cloud prerequisites	46
4.2.2 Microsoft Azure prerequisites	47
4.3 Planning security access control on Microsoft Azure	48
4.3.1 Installer user role permissions	49
4.3.2 Management user role permissions	51
4.3.3 Bastion user role permissions	52
4.4 Planning networking for Microsoft Azure	52
4.4.1 Network considerations for basic deployment	52
4.4.2 Network considerations for cross-nublic network deployments	54
4.5. Planning an Azuro virtual maching	56
4.5 Flamming all Azure villual machine	50
4.0 Flamming Azure managed ulsks	50
	50
4.6.2 MDISK Support	58
4.7 Planning deployment access	59
4.8 Storage performance optimization	60
4.9 Planning for data reduction pools	61
Chanter F. Juniomenting on IDM Creatury Vistualize for Dublic Cloud on Microsoft	
Chapter 5. Implementing an IBM Spectrum virtualize for Public Cloud on Microsoft	~~
Azure environment	63
5.1 Installing IBM Spectrum Virtualize for Public Cloud on Azure.	64
5.2 Logging in to IBM Spectrum Virtualize for Public Cloud on Azure	73
5.2.1 Contiguring the Azure Bastion Service	74
5.2.2 Connecting to the Azure virtual machine by using Azure Bastion Service	76
5.2.3 Configuring GUI for Spectrum Virtualize for Public Cloud cluster on Azure and	
completing the installation	79
5.3 Configuring the cloud quorum	88
5.4 Configuring the back-end storage	90

5.4.2 Configuring host and volume mapping 98 5.5 Configuring a site to site virtual private network gateway for hybrid cloud connectivity in Azure Cloud 100 5.5.1 Azure configuration for VPNGW IPsec tunnel. 100 5.6 Configuring replication from on-premises IBM Spectrum Virtualize to IBM Spectrum Virtualize for Public Cloud on Azure. 101 Chapter 6. Monitoring and supporting the solution. 111 6.1 Monitoring and supporting the solution. 111 6.1 Monitoring IBM Spectrum Virtualize for Public Cloud on Microsoft Azure through GUI, Spectrum Control, or Storage Insights 112 6.1.1 Call Home function and email notification 112 6.1.2 Monitoring SV4PC on Azure in IBM Spectrum Control or IBM Storage Insights 124 6.2 Troubleshooting Spectrum Virtualize for Public Cloud on Microsoft Azure 137 6.2.1 Collecting diagnostic data for IBM Spectrum Virtualize 137 6.2.2 Uploading files to the Support Center 141 6.2.3 Service Assistant Tool 142 6.2.4 Remote Support Assistance 145 6.2.5 Troubleshooting in Microsoft Azure. 149 6.2.6 Hints and tips 152 6.3 IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Support contact 154 6.3.1 Who to ca
5.5 Configuring a site to site virtual private network gateway for hybrid cloud connectivity in Azure Cloud 100 5.5.1 Azure configuration for VPNGW IPsec tunnel. 100 5.6 Configuring replication from on-premises IBM Spectrum Virtualize to IBM Spectrum 100 5.6 Configuring replication from on-premises IBM Spectrum Virtualize to IBM Spectrum 101 Chapter 6. Monitoring and supporting the solution. 111 6.1 Monitoring IBM Spectrum Virtualize for Public Cloud on Microsoft Azure through GUI, Spectrum Control, or Storage Insights 112 6.1.1 Call Home function and email notification 112 6.1.2 Monitoring SV4PC on Azure in IBM Spectrum Control or IBM Storage Insights 124 6.2 Troubleshooting Spectrum Virtualize for Public Cloud on Microsoft Azure 137 6.2.1 Collecting diagnostic data for IBM Spectrum Virtualize 137 6.2.2 Uploading files to the Support Center 141 6.2.3 Service Assistant Tool 142 6.2.4 Remote Support Assistance 145 6.2.5 Troubleshooting in Microsoft Azure 149 6.2.6 Hints and tips 152 6.3 IBM
Azure Cloud1005.5.1 Azure configuration for VPNGW IPsec tunnel.1005.6 Configuring replication from on-premises IBM Spectrum Virtualize to IBM Spectrum Virtualize for Public Cloud on Azure.101Chapter 6. Monitoring and supporting the solution.1116.1 Monitoring IBM Spectrum Virtualize for Public Cloud on Microsoft Azure through GUI, Spectrum Control, or Storage Insights1126.1.1 Call Home function and email notification.1126.1.2 Monitoring SV4PC on Azure in IBM Spectrum Control or IBM Storage Insights1246.2 Troubleshooting Spectrum Virtualize for Public Cloud on Microsoft Azure.1376.2.1 Collecting diagnostic data for IBM Spectrum Virtualize1376.2.2 Uploading files to the Support Center.1416.2.3 Service Assistant Tool1426.2.4 Remote Support Assistance1456.2.5 Troubleshooting in Microsoft Azure.1496.2.6 Hints and tips1526.3 IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Support contact1546.3.1 Who to call for support1546.3.2 Working with IBM Support.1556.3.3 Working with Microsoft Azure Support156
5.5.1 Azure configuration for VPNGW IPsec tunnel. 100 5.6 Configuring replication from on-premises IBM Spectrum Virtualize to IBM Spectrum Virtualize for Public Cloud on Azure. 101 Chapter 6. Monitoring and supporting the solution. 111 6.1 Monitoring IBM Spectrum Virtualize for Public Cloud on Microsoft Azure through GUI, Spectrum Control, or Storage Insights 112 6.1.1 Call Home function and email notification. 112 6.1.2 Monitoring SV4PC on Azure in IBM Spectrum Control or IBM Storage Insights 124 6.2 Troubleshooting Spectrum Virtualize for Public Cloud on Microsoft Azure. 137 6.2.1 Collecting diagnostic data for IBM Spectrum Virtualize 137 6.2.2 Uploading files to the Support Center. 141 6.2.3 Service Assistant Tool 142 6.2.4 Remote Support Assistance 145 6.2.5 Troubleshooting in Microsoft Azure. 149 6.2.6 Hints and tips 152 6.3 IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Support contact 154 6.3.2 Working with IBM Support. 155 6.3.3 Working with Microsoft Azure Support 155
5.6 Configuring replication from on-premises IBM Spectrum Virtualize to IBM Spectrum Virtualize for Public Cloud on Azure. 101 Chapter 6. Monitoring and supporting the solution 111 6.1 Monitoring IBM Spectrum Virtualize for Public Cloud on Microsoft Azure through GUI, Spectrum Control, or Storage Insights 112 6.1.1 Call Home function and email notification 112 6.1.2 Monitoring SV4PC on Azure in IBM Spectrum Control or IBM Storage Insights 124 6.2 Troubleshooting Spectrum Virtualize for Public Cloud on Microsoft Azure 137 6.2.1 Collecting diagnostic data for IBM Spectrum Virtualize 137 6.2.2 Uploading files to the Support Center 141 6.2.3 Service Assistant Tool 142 6.2.4 Remote Support Assistance 145 6.2.5 Troubleshooting in Microsoft Azure. 149 6.2.6 Hints and tips 152 6.3 IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Support contact 154 6.3.2 Working with IBM Support. 155 6.3.3 Working with Microsoft Azure Support 155
Virtualize for Public Cloud on Azure101Chapter 6. Monitoring and supporting the solution1116.1 Monitoring IBM Spectrum Virtualize for Public Cloud on Microsoft Azure through GUI, Spectrum Control, or Storage Insights1126.1.1 Call Home function and email notification1126.1.2 Monitoring SV4PC on Azure in IBM Spectrum Control or IBM Storage Insights1246.2 Troubleshooting Spectrum Virtualize for Public Cloud on Microsoft Azure1376.2.1 Collecting diagnostic data for IBM Spectrum Virtualize1376.2.2 Uploading files to the Support Center1416.2.3 Service Assistant Tool1426.2.4 Remote Support Assistance1456.2.5 Troubleshooting in Microsoft Azure1496.2.6 Hints and tips1526.3 IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Support contact1546.3.1 Who to call for support1546.3.2 Working with IBM Support1556.3.3 Working with Microsoft Azure Support156
Chapter 6. Monitoring and supporting the solution. 111 6.1 Monitoring IBM Spectrum Virtualize for Public Cloud on Microsoft Azure through GUI, Spectrum Control, or Storage Insights 112 6.1.1 Call Home function and email notification. 112 6.1.2 Monitoring SV4PC on Azure in IBM Spectrum Control or IBM Storage Insights 124 6.2 Troubleshooting Spectrum Virtualize for Public Cloud on Microsoft Azure. 137 6.2.1 Collecting diagnostic data for IBM Spectrum Virtualize 137 6.2.2 Uploading files to the Support Center 141 6.2.3 Service Assistant Tool 142 6.2.5 Troubleshooting in Microsoft Azure. 149 6.2.6 Hints and tips 152 6.3 IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Support contact 154 6.3.2 Working with IBM Support. 155 6.3.3 Working with Microsoft Azure Support 156
6.1 Monitoring IBM Spectrum Virtualize for Public Cloud on Microsoft Azure through GUI, Spectrum Control, or Storage Insights 112 6.1.1 Call Home function and email notification 112 6.1.2 Monitoring SV4PC on Azure in IBM Spectrum Control or IBM Storage Insights 124 6.2 Troubleshooting Spectrum Virtualize for Public Cloud on Microsoft Azure 137 6.2.1 Collecting diagnostic data for IBM Spectrum Virtualize 137 6.2.2 Uploading files to the Support Center 141 6.2.3 Service Assistant Tool 142 6.2.5 Troubleshooting in Microsoft Azure 149 6.2.6 Hints and tips 152 6.3 IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Support contact 154 6.3.1 Who to call for support 154 6.3.2 Working with IBM Support. 155 6.3.3 Working with Microsoft Azure Support 156
6.1 Monitoring IBM Spectrum Virtualize for Public Cloud on Microsoft Azure through GUI, Spectrum Control, or Storage Insights 112 6.1.1 Call Home function and email notification 112 6.1.2 Monitoring SV4PC on Azure in IBM Spectrum Control or IBM Storage Insights 124 6.2 Troubleshooting Spectrum Virtualize for Public Cloud on Microsoft Azure 137 6.2.1 Collecting diagnostic data for IBM Spectrum Virtualize 137 6.2.2 Uploading files to the Support Center 141 6.2.3 Service Assistant Tool 142 6.2.4 Remote Support Assistance 145 6.2.5 Troubleshooting in Microsoft Azure 149 6.2.6 Hints and tips 152 6.3 IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Support contact 154 6.3.1 Who to call for support 154 6.3.2 Working with IBM Support. 155 6.3.3 Working with Microsoft Azure Support 156
Spectrum Control, or Storage Insights1126.1.1 Call Home function and email notification1126.1.2 Monitoring SV4PC on Azure in IBM Spectrum Control or IBM Storage Insights1246.2 Troubleshooting Spectrum Virtualize for Public Cloud on Microsoft Azure1376.2.1 Collecting diagnostic data for IBM Spectrum Virtualize1376.2.2 Uploading files to the Support Center1416.2.3 Service Assistant Tool1426.2.4 Remote Support Assistance1456.2.5 Troubleshooting in Microsoft Azure1496.2.6 Hints and tips1526.3 IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Support contact1546.3.1 Who to call for support1546.3.2 Working with IBM Support.1556.3.3 Working with Microsoft Azure Support156
6.1.1 Call Home function and email notification. 112 6.1.2 Monitoring SV4PC on Azure in IBM Spectrum Control or IBM Storage Insights 124 6.2 Troubleshooting Spectrum Virtualize for Public Cloud on Microsoft Azure. 137 6.2.1 Collecting diagnostic data for IBM Spectrum Virtualize 137 6.2.2 Uploading files to the Support Center. 141 6.2.3 Service Assistant Tool 142 6.2.4 Remote Support Assistance 145 6.2.5 Troubleshooting in Microsoft Azure. 149 6.2.6 Hints and tips 152 6.3 IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Support contact 154 6.3.1 Who to call for support 154 6.3.2 Working with IBM Support. 155 6.3.3 Working with Microsoft Azure Support 156
6.1.2 Monitoring SV4PC on Azure in IBM Spectrum Control of IBM Storage Insights 124 6.2 Troubleshooting Spectrum Virtualize for Public Cloud on Microsoft Azure 137 6.2.1 Collecting diagnostic data for IBM Spectrum Virtualize 137 6.2.2 Uploading files to the Support Center 141 6.2.3 Service Assistant Tool 142 6.2.4 Remote Support Assistance 145 6.2.5 Troubleshooting in Microsoft Azure 149 6.2.6 Hints and tips 152 6.3 IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Support contact 154 6.3.1 Who to call for support 154 6.3.2 Working with IBM Support. 155 6.3.3 Working with Microsoft Azure Support 156
6.2 Troubleshooting Spectrum Virtualize for Public Cloud on Microsoft Azure
6.2.1 Collecting diagnostic data for IBM Spectrum Virtualize1376.2.2 Uploading files to the Support Center1416.2.3 Service Assistant Tool1426.2.4 Remote Support Assistance1456.2.5 Troubleshooting in Microsoft Azure1496.2.6 Hints and tips1526.3 IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Support contact1546.3.1 Who to call for support1546.3.2 Working with IBM Support1556.3.3 Working with Microsoft Azure Support156
6.2.2 Oploading files to the Support Center1416.2.3 Service Assistant Tool1426.2.4 Remote Support Assistance1456.2.5 Troubleshooting in Microsoft Azure1496.2.6 Hints and tips1526.3 IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Support contact1546.3.1 Who to call for support1546.3.2 Working with IBM Support1556.3.3 Working with Microsoft Azure Support156
6.2.3 Service Assistant roof 142 6.2.4 Remote Support Assistance 145 6.2.5 Troubleshooting in Microsoft Azure. 149 6.2.6 Hints and tips 152 6.3 IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Support contact 154 6.3.1 Who to call for support 154 6.3.2 Working with IBM Support. 155 6.3.3 Working with Microsoft Azure Support 156
6.2.4 Hemote Support Assistance 145 6.2.5 Troubleshooting in Microsoft Azure. 149 6.2.6 Hints and tips 152 6.3 IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Support contact 154 6.3.1 Who to call for support 154 6.3.2 Working with IBM Support. 155 6.3.3 Working with Microsoft Azure Support 156
6.2.5 Troubleshooting in Microsoft Azure
6.2.6 Hints and tips 152 6.3 IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Support contact 154 6.3.1 Who to call for support 154 6.3.2 Working with IBM Support 155 6.3.3 Working with Microsoft Azure Support 156
6.3 IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Support contact 154 6.3.1 Who to call for support 154 6.3.2 Working with IBM Support 155 6.3.3 Working with Microsoft Azure Support 156
6.3.1 Who to call for support 154 6.3.2 Working with IBM Support 155 6.3.3 Working with Microsoft Azure Support 156
6.3.2 Working with IBM Support. 155 6.3.3 Working with Microsoft Azure Support 156
6.3.3 Working with Microsoft Azure Support
Related publications
IBM Redbooks
Online resources
Help from IBM

vi Implementation Guide for IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Version 8.4.3

Notices

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

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

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

IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

Trademarks

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

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

HyperSwap®
IBM®
IBM Cloud®
IBM FlashSystem®
IBM Spectrum®
Passport Advantage®

Redbooks® Redbooks (logo) @ ® Storwize® XIV®

The following terms are trademarks of other companies:

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

ITIL is a Registered Trade Mark of AXELOS Limited.

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.

OpenShift, Red Hat, 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 is now available on Microsoft Azure. With IBM Spectrum® Virtualize for Public Cloud 8.4.3, users can deploy a highly available two-node cluster running IBM Spectrum Virtualize for Public Cloud on supported Microsoft Azure virtual machines (VMs). This all-inclusive, bring your own license (BYOL) software offering provides the ability to virtualize, optimize, and provision supported Azure Managed Disk to applications that require the performance of block storage in the cloud with the added efficiencies that IBM Spectrum Virtualize for Public Cloud brings to native infrastructure as a service (IaaS) that is provided by Microsoft Azure.

This IBM Redbooks publication gives a broad understanding of the IBM Spectrum Virtualize for Public Cloud on Azure architecture. It also provides planning and implementation details of the common use cases for this new product.

This book helps storage and networking administrators plan, implement, install, and configure the IBM Spectrum Virtualize for Public Cloud on Azure offering. It also provides valuable troubleshooting tips.

Authors

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







Jackson Shea is a Level 2 certified IBM Information Technology Specialist/Architect performing design and implementation engagements through 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 had over 12 years of experience with IBM Spectrum Virtualize (formerly known as the IBM SAN Volume Controller) and related technology. Jackson is based out of 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.



Katja Kratt is a Product Field Engineer for IBM Spectrum Virtualize products and IBM V9000 in the EMEA Storage Competence Center (ESCC) in Kelsterbach, Germany. She has 34 years with IBM, mostly with technical support and technical education. She has co-authored several IBM Redbooks publications.



Lars Dill started with IBM in 1998 as a hardware technician for IBM high-end printers. After the spin off of IBM Printing Systems Division in 2006, he worked as a hardware technician for IBM DASD (storage) and IBM Power servers. In March 2009, he changed his role from hardware technician to a support center representative working for DASD midrange, SAN, and IBM SAN Volume Controller products. Since July 2012, he has been working as a SME for Spectrum Virtualize Products, working for the EMEA Storage Competence Center (ESCC), in Zurich. He provides remote technical support for IBM SAN Volume Controller, IBM Storwize®, and IBM FlashSystem® in Europe and worldwide. This work includes analyzing complex field issues, working on critsits, and collaborating with Development.



Leandro Torolho works as a Client Technical Specialist (CTS) - IBM Public Cloud Platform and is responsible for supporting technical sales of IBM's Public Cloud portfolio. His previous roles range from working as a SAN/Storage SME for several outsourced clients to be a CTS for IBM Storage. Leandro is a Certified Distinguished IT Specialist by The Open Group and holds a Bachelor's degree in Computer Science.



Nils Olsson is a SME (previously, Product Engineering Professional - Level 2 Support) working in the EMEA Storage Competence Center (ESCC) in Kelsterbach, IBM Germany. He provides remote technical support for Spectrum Virtualize Products (IBM SAN Volume Controller, IBM Storwize, and IBM FlashSystem) in Europe and worldwide. Nils joined IBM in 2008 and is skilled in SAN, storage, and storage virtualization. He has over 10 years of experience with troubleshooting IBM Spectrum® Virtualize and is certified on the IBM Storwize portfolio. In his current role, he delivers analysis of complex field issues and provides technical expertise during critical situations, collaborating closely with development.



Pankaj Deshpande is the IBM Spectrum Virtualize for Public Cloud Architect, working in IBM Pune, India. A key part of his role in IBM is in architecture, designing and developing products, and evolving technology and product strategy for IBM. Key technical contributions consist of storage technologies, storage security, cybersecurity, public and private cloud integration, storage for VMs and Kubernetes. Pankaj also represents IBM some of the 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 spread over more than 20 years, he has worked on various designs involving scale-out, distributed, high performance, and highly available architectures.



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 Project since 2015. He has worked and delivered features, including 3-Site Replication (Metro-Mirror, Hyperswap, and GUI), iSCSI/iSER Support on IBM SAN Volume Controller, and IBM Spectrum Virtualize for Public Cloud. He has significantly 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 degree in Computer Application from Mumbai University, India.



Vasfi Gucer is a project leader with the IBM Systems Worldwide Client Experience Center. He has more than 20 years of experience in the areas of systems management, networking hardware, and software. He writes extensively and teaches IBM classes worldwide about IBM products. His focus has been primarily on storage and cloud computing for the last 8 years. Vasfi is also an IBM Certified Senior IT Specialist, Project Management Professional (PMP), IT Infrastructure Library (ITIL) V2 Manager, and ITIL V3 Expert.

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

Mary Connell Michelle Tidwell Erica Wazewski IBM US

John Clifton IBM UK

Saurabh Singh IBM India

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 books to be as helpful as possible. Send us your comments about this book 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

Look for us on LinkedIn:

http://www.linkedin.com/groups?home=&gid=2130806

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

https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

Stay current on recent Redbooks publications with RSS Feeds:

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

1

Introduction

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

A brief overview of the technology that is behind the product introduces the drivers and business value of the use of IBM Spectrum Virtualize in the context of public cloud services. It also describes how the solution works from a high-level perspective.

This publication describes IBM Spectrum Virtualize for Public Cloud V8.4.3.

This chapter includes the following topics:

- ▶ 1.1, "Introduction to IBM Spectrum Virtualize for Public Cloud" on page 2
- ► 1.2, "IBM Spectrum Virtualize for Public Cloud" on page 2
- ▶ 1.3, "IBM Spectrum Virtualize for Public Cloud on Azure" on page 7

1.1 Introduction to IBM Spectrum Virtualize for Public Cloud

Companies are currently undergoing a digital transformation and making architectural decisions that determine how their businesses are going to operate in the next few years. They recognize the value of delivering services by using the cloud, and many use public clouds to some degree.

The role of the cloud is maturing and it is more often being considered as a platform for innovation and business value. The cloud is a key enabler to drive transformation and innovation for IT agility and new capabilities.

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

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

In this sense, coming from the IBM Spectrum Storage family, IBM Spectrum Virtualize for Public Cloud supports clients in their IT architectural transformation and migration toward the cloud service model. It enables hybrid cloud strategies or for a cloud-native workload, provides the benefits of familiar and sophisticated storage functions on public cloud data centers by enhancing the cloud offering.

Running on-premises, IBM Spectrum Virtualize software can virtualize over 400 different storage systems from IBM and other vendors. This wide range of storage support means that the solution can be used with almost any storage in a data center today and is integrated with its counterpart IBM Spectrum Virtualize for Public Cloud, which supports IBM Cloud®, Amazon Web Services (AWS), and now Microsoft Azure.

For more information about Azure, see Chapter 3, "Solution architecture" on page 25, and Chapter 4, "Planning an installation on Microsoft Azure" on page 45.

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. It includes technology that complements and enhances public cloud offering capabilities.

For example, traditional practices that provide data replication by copying storage at one facility to largely identical storage at another facility are not an option for public cloud. Also, the use of conventional software to replicate data imposes unnecessary loads on application servers.

For more information about use cases, see Chapter 2, "Typical use cases for IBM Spectrum Virtualize for Public Cloud" on page 9.

IBM Spectrum Virtualize for Public Cloud delivers a powerful solution for the deployment of IBM Spectrum Virtualize software in public clouds. This capability allows the use of IBM Spectrum Virtualize for Public Cloud on Azure to enable hybrid cloud solutions, which offer the ability to transfer data between on-premises data centers by using any IBM Spectrum Virtualize -based storage system and multiple cloud environments.

With a deployment that is designed for the cloud, IBM Spectrum Virtualize for Public Cloud can be deployed in cloud data centers around the world.

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 use 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 more wanted characteristics.

It is a logical representation of resources that is not constrained by physical limitations and hides part of the complexity. It also adds or integrates new functions with existing services and can be nested or applied to multiple layers of a system.

The aggregation of volumes into storage pools enables you to better manage capacity, performance, and 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. For this reason, it is referred to as *block-level virtualization*, or the block aggregation layer. For the sake of clarity, the block-level volumes that are provided by the cloud are exposed as target volumes, and are seen by IBM Spectrum Virtualize as a managed disk (MDisk).

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

The virtualization terminology is included in 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 more generally, installed on any commodity x86 hardware and hypervisor.

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

In general, SDS technology uses the following concepts:

- A shared-nothing architecture (or in some cases, a partial or fully shared architecture) with no single point of failure and nondisruptive upgrades.
- Scale-up or scale-out mode: Add building blocks for a predictable increase in capacity, performance, and resiliency.
- Multiple classes of service: File-based, object-based, block-based, and auxiliary and storage support service. SDS solutions might also be integrated together into a hybrid or composite SDS solution.

- High availability (HA) and Disaster Recovery (DR): Tolerates levels of availability and durability as self-healing and adjusting.
- Lower total cost of ownership (TCO): Lower the TCO for those workloads that can use 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 cloud.

By using standard x86 servers, IBM Spectrum Virtualize for Public Cloud can be easily added to cloud infrastructures to deliver more features and functions, which 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 catalog by increasing standard storage, and offering the following capabilities and features that decrease 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, the administrator has more granular control, which enables a production volume to have a snapshot that is stored on lower-end 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 this writing, IBM Spectrum Virtualize allows for up to 256 TB and up to 20,000 host connections.

- Native storage-based replication

Replication features are natively supported, but are typically limited to specific data center pairs and a predefined minimum recovery point objective (RPO). They are accessible only when the primary volume is down.

IBM Spectrum Virtualize provides greater flexibility in storage replication to allow for user-defined RPO and replication between any other system running IBM Spectrum Virtualize.

New features for public cloud storage offering

IBM Spectrum Virtualize for Public Cloud introduces to the public cloud catalog new storage capabilities. These features are available on IBM SAN Volume Controller and IBM Spectrum Virtualize, but are not available by default. These features that are provided on public cloud are related to hybrid cloud scenarios and its support to foster 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 FlashSystem family, IBM SAN Volume Controller, or VersaStack and Public Cloud and extends replication to all types of supported virtualized storage on-premises.

Working together, IBM Spectrum Virtualize and IBM Spectrum Virtualize for Public Cloud support synchronous and asynchronous mirroring between the cloud and on-premises for more than 400 different storage systems from various vendors. In addition, they support other services, such as IBM FlashCopy® and IBM Easy Tier.

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

One of the reasons to replicate is to have a copy of the data from which to restart operations if an emergency occurs. IBM Spectrum Virtualize for Public Cloud enables DR for virtual and physical environments, which adds possibilities compared to the software replicators that are in use today that 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 from a cloud provider back to on-premises or to another cloud provider.

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

Capabilities, such as rapid, flexible provisioning; simplified configuration changes; nondisruptive movement of data among tiers of storage; and a single user interface helps make the storage infrastructure (and the hybrid cloud) simpler, more cost-effective, and easier 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 including functions that are not natively available on the public cloud catalogs and that are traditionally deployed within disk array systems in the on-premises environment. For this reason, IBM Spectrum Virtualize for Public Cloud improves and expands the capabilities of the cloud offering.

Table 1-1 on page 6 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. Designed to help support application availability. 		
Pools the capacity of multiple storage volumes	 Helps overcome the volume size limitations. Helps manage storage as a resource to meet business requirements, and not just 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 servers that are configured as IBM Spectrum Virtualize for Public Cloud engines	 Use of cloud-catalog Intel servers foundation. Designed to avoid single point of hardware failures. 		
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 or classes. 		
Easy-to-use IBM FlashSystem family management interface	 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. IBM Spectrum Control Insights and IBM Spectrum Protect provide more capabilities to manage capacity and performance. 		
Dynamic data migration	 Migrate data among volumes or LUNs without taking applications that use that data offline. Manage and scale storage capacity without disrupting applications. 		
Advanced network-based copy services	 Copy data across multiple storage systems with IBM FlashCopy. This version supports consistency groups, incremental copies, and cascaded copies. Copy data across metropolitan and global distances as needed to create high-availability storage solutions between multiple data centers. 		
Thin provisioning and snapshot replication	 Reduce volume requirements by using storage only when data changes. Improve storage administrator productivity through automated on-demand storage provisioning. Snapshots are available on lower-tier storage volumes. Safe Guarded Copy (SCG) to cloud 		
IBM Spectrum Protect Snapshot appli- cation-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	 Embedded compress replication traffic for WAN optimization. Reduces network costs or speed replication cycles, which improves remote data accuracy. 		
IBM Spectrum Connect Cloud Storage Management	 Manages container storage in Kubernetes. 		

Important: The following features are not supported in the first release of IBM Spectrum Virtualize for Pubic Cloud on Azure:

- Stretched cluster
- ► IBM HyperSwap®
- Encryption
- Hot spare node
- Multiple I/O groups

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

1.3 IBM Spectrum Virtualize for Public Cloud on Azure

Version 8.4.3 of IBM Spectrum Virtualize for Public Cloud is now available on 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 product and between public clouds.
- 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 thin provisioning volumes and paying for Azure storage as you grow.

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



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

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

In Azure, *Locally Redundant Storage* (LRS) is attached to the IBM Spectrum Virtualize node instances that compose a single node pair (or I/O group) that provides a shared storage pool that is used by IBM Spectrum Virtualize. IBM Spectrum Virtualize supports the following Azure LRS types:

- Standard solid-state drive (SSD)
- Premium SSD

1.3.1 Application examples on Azure

Depending on your environment, the following applications and workloads might be a good fit for IBM Spectrum Virtualize for Public Cloud on Azure:

- ► Traditional database applications, such as Microsoft SQL and Oracle databases.
- Cloud Analytics for read intensive workloads. Customers can benefit from the caching capabilities of IBM Spectrum Virtualize because Azure customers are charged on the storage capacity and how many times they are accessing data.
- Red Hat OpenShift and Kubernetes Services in Cloud, which provides persistent storage for containers (for example, by way of the CSI driver).
- Customer applications that are ported to Azure require low latency and high performance.

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:

- ▶ 2.1, "Deploying whole IT services in the public cloud" on page 10
- 2.2, "Disaster Recovery" on page 15
- ► 2.3, "IBM FlashCopy in the public cloud" on page 17
- ▶ 2.4, "Safeguarded Copy" on page 21
- ► 2.5, "Workload relocation into the public cloud" on page 23

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 that 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: 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 VPN gateways.
- Platform as a Service: PaaS is the intermediate and most typically the most common cloud environment. Microsoft Azure and Amazon Web Services, and 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 between the off-premises public cloud services with an on-premises IT environment.
- Cloud-native: The full application's stack is moved to cloud as SaaS, PaaS, IaaS, or as a combination of the three delivery models.



Figure 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 laaS 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 laaS 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 generally 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 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 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 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 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 guarantees service availability and uptime. Companies that use a public cloud IaaS can meet required SLAs by implementing highly available solutions and duplicating the infrastructure in the same data center or in two data centers to maintain business continuity in case of 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 highly available deployment models for an application that is fully deployed on public cloud are summarized as follows:

Highly available 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 (SPOF), but it does not function if the data center is unavailable.

Highly available cloud deployment on multi-site

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 laaS that is the primary site for service delivery, a reasonable approach is implementing at least a highly available 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 take advantage of the intrinsic availability (and SLAs) of the storage service, which is this case, is Microsoft Azure Managed Disk.

When IBM Spectrum Virtualize for Public Cloud is deployed as clustered pair of Azure VM 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 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 Spectrum Virtualize instances.

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

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. Copy data across metropolitan and global distances as needed to create high-availability storage solutions between multiple data centers.
Thin provisioning and snapshot replication	 Reduces volume 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.
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.
Third-party native integration	Integration with VMware vRealize.

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

Feature	Benefits
Safeguarded Copy	The new Spectrum Virtualize functions provides a valuable ransomware mitigation solution, especially when combined with implementation of IBM Spectrum Virtualize for Public Cloud.

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 15. 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.

¹ Spectrum Virtualize Highly Available multi-site topologies, such as HyperSwap and Enhanced Stretch Cluster, are not supported by Spectrum Virtualize Public Cloud as of this writing.

2.2 Disaster Recovery

Customers have long been adopting DR strategies to harness and secure proliferating data in their environment and infrastructure workloads in a cost-effective manner when a highly available (HA) level of recovery point objective (RPO) 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 DR approach and benefits of IBM Spectrum Virtualize for Public Cloud on Azure.

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.

Adoption drivers	Challenges	IBM Spectrum Virtualize for IBM public cloud capabilities	
The promise of reduced operational expenditures and capital expenditures	 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 	
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) GMCVs for Asynchronous replication with a tunable RPO (recommended for Public Cloud deployments) 	

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

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

- Can be implemented at several locations in Microsoft Azure and installed by using Azure Marketplace.
- ► Is deployed on an Azure VM instance.
- Offers data replication with the FlashSystem family, V9000, IBM SAN Volume Controller, or VersaStack and public cloud.
- Supports two node clusters in Microsoft Azure.
- Offers data services for Azure Managed Disks.
- 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 locations.
- Replicates between on-premises and Microsoft Azure running IBM Spectrum Virtualize on-premises and IBM Spectrum Virtualize for Public Cloud on Azure.

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.



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

As shown in Figure 2-2 on page 16, 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 an Amazon EBS volume.

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

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

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 section 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, it is easiest to think of these objects as "layers" that sit on top of one another. The application is the topmost layer, and beneath it is the operating system layer.

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

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

The resulting copy requires the same type of recovery procedure, such as log replay and file system checks, that is required following a host crash. 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, it is important that any data that is held on the host operating system (or application) caches for the target volumes is discarded. The easiest way to ensure that no data is held in these caches is to unmount the target volumes before the FlashCopy operation starts.

Best practice: 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 in most cases to use the write-suspend method that is available in most modern databases because the database maintains strict control over I/O.

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

2.4 Safeguarded Copy

Combining these use cases of remote replication and FlashCopy with the new safeguarded child pool function that was introduced in Spectrum Virtualize 8.4.2.0, we now have a powerful cyber-resilience use case for Spectrum Virtualize in Public Cloud on Azure. On-premises workloads can be protected from ransomware and other data corruption attacks with a truly air-gapped solution in the public cloud.

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 in some cases 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 by way of IP to a Spectrum Virtualize in Public Cloud instance on Azure.



Figure 2-4 Safeguarded Copy

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

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

Because Copy Services Manager is primarily a replication orchestration tool, it is perfectly positioned to also 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

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.
- Volume Group: 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 volume group, 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 CLI command. The three default policies are listed in Table 2-3.

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

Table 2-3 Default policies

Copy Services Manager (CSM): Application that has long existed as a replication and point-in-time copy orchestration tool for IBM storage (Spectrum Virtualize, DS8K, XIV). With version 6.2, CSM integrates with Spectrum Virtualize 8.4.2.0 to periodically scan for volume groups 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).
2.5 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 with regard to 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 premigration.

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 17, 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 it 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.

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		

Table 2-4 Migration methods

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. Currently, in Azure 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 managed disks (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 the Microsoft Azure environment regarding architectural aspects of IBM Spectrum Virtualize for Public Cloud in Azure. It also explains various components of the Azure solution and how they interact and interrelate with each other.

This chapter includes the following topics:

- ► 3.1, "IBM Spectrum Virtualize" on page 26
- ► 3.2, "Azure terminology" on page 33
- ► 3.3, "Key components of the Azure solution" on page 34
- ► 3.4, "Highly available infrastructure" on page 36
- ► 3.5, "Security design considerations" on page 40
- 3.6, "Solution architecture: IBM Spectrum Virtualize as a storage for all-in-cloud model" on page 43

3.1 IBM Spectrum Virtualize

IBM Spectrum Virtualize is a software-enabled storage virtualization engine that provides a single point of control for storage resources within the data centers. IBM Spectrum Virtualize is a core software engine of 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 Azure, providing increased flexibility in data center infrastructure and cloud systems. This section describes the components of IBM Spectrum Virtualize as they are deployed in the cloud.

3.1.1 Nodes

IBM Spectrum Virtualize software is installed on Azure 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 this writing, IBM Spectrum Virtualize for Public Cloud on Azure is limited to a single I/O group. However, plans are in place to expand to four I/O groups in the future.

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 active nodes are installed in pairs, each node maintains cache coherence with its partner to provide seamless failover functions and fault tolerance.

3.1.2 I/O groups

Each pair of IBM Spectrum Virtualize nodes is referred to as an I/O group. A specific volume is always presented to a host server by a single I/O group in the system.

When a host server performs I/O to one of its volumes, all the I/Os for a specific volume are directed to one specific I/O group in the system. Under normal conditions, the I/Os for that specific volume are always processed by the same node within the I/O group. This node is referred to as the *preferred node* for this specific volume. 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

The current IBM Spectrum Virtualize for Public Cloud in Azure supports a clustered system that consists of one I/O group (it is planned to support more soon). Specific configuration limitations are then set for the individual system. For example, the maximum managed disks (MDisks) that are supported is 992 TB per system. The current 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 IP* or *cluster IP* and is active on the configuration node. Each node in the system also is 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, see this IBM Documentation web page.

3.1.4 MDisks

IBM Spectrum Virtualize for Public Cloud on Azure uses Azure MDisks as a 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 does LRS. 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 (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 suitably 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*. These virtual disks 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.

MDisks that are presented to IBM Spectrum Virtualize can have the following operating modes:

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 by using volume mirroring). 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 subset of a pool can be 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 use of child pools was introduced in IBM Spectrum Virtualize Version 8.4.2: the *safeguarded child pool*. This object allows the creation of unalterable snapshots of important volumes to guard against accidental or malicious corruption.

For more information, see 2.4, "Safeguarded Copy" on page 21.

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 allow 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

The *iSCSI function* is a software function that is provided by the IBM Spectrum Virtualize code. IBM introduced software capabilities to allow the underlying virtualized storage to attach to IBM Spectrum Virtualize 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 de-staging 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 (see Figure 3-1).



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, and then, dynamically moves high-activity or hot extents to a higher disk tier within the storage pool. It also moves extents whose activity dropped off or cooled down from the high-tier 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 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 on the web at no initial cost at IBM Storage Tier Advisor.

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 and 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 worsens 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 take advantage of faster data transfer to speed replication cycles, improve remote data currency, and enjoy faster recovery.

3.1.14 IBM FlashCopy

FlashCopy is sometimes described as an instance of a time-zero (T0) copy or a 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 to this is provided by IBM Spectrum Protect, which is described in What can IBM Spectrum Protect do for your business?.

3.2 Azure terminology

Azure terminology is listed in Table 3-1.

Table 3-1	Azure terminology
-----------	-------------------

Azure term	Explanation
Azure resource group	A container that holds all of 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 and are deleted after the cluster configuration is complete. For more information, see this web page.
Azure network security groups	This group contains security rules that define rules for inbound network traffic, outbound network traffic between Azure resources. For more information, see this web page.
VNet	The Azure virtual network (VNet) facilitates secured communication between different types of Azure resources, such as Azure virtual machines (VM), external network, and the internet. For more information, see this web page.
Subnet	A range of IP addresses in the virtual network. You can divide a virtual network into multiple subnets for organization and security. Each NIC in a VM is connected to one subnet in one virtual network. 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 located 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 single point of failure.
Update domain	This group of resources can be updated and brought down for maintenance and patching (for example, system patches and software updates) possibly at the same time. Typically, only one update domain is brought down at a time for patching.

3.3 Key components of the Azure solution



Figure 3-2 shows the high-level architecture of Spectrum Virtualize for Public Cloud on Microsoft Azure.

Figure 3-2 Components of the Azure solution

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

- Azure virtual machines
- Azure managed disks
- Azure networking elements, such as vNET, subnets, loadbalancer, and vNICs
- Azure constructs, such as resource group, availability set, and proximity
- Azure policies for network security groups

3.3.1 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 nested Azure Resource Manager (ARM) template, which is an in*frastructure as a code primitive* in Azure. Customer needs to manage the solution and the transactions are billed through a separate plan. ARM template-based deployment automates deployment of several VMs along with customized network infrastructure, security groups, key vault management, managed disk in customer's Azure subscription.

3.3.2 Resource group and vNets

A *resource group* in Azure is a unit of management; therefore, each Spectrum Virtualize cluster is deployed in a fresh resource group. All resources that are part of the cluster are held in their respective resource group.

Although multiple clusters cannot be deployed in the same resource group, multiple resource groups can share virtual network across different clusters (see Figure 3-3).

SG E FOR CONTRACTOR OF CONTRAC	NSG Quorum Node Quorum VM Subnet	
	Virtual Network	

Figure 3-3 Resource group and vNets

VNet is the building block for a customer's private network in Azure. VNet facilitates various Azure resources to securely communicate with each other, with external resources over internet, and on-premises networks. Azure VNet is a virtual network 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 resource groups can share a vNet across different clusters.

3.4 Highly available infrastructure

IBM Spectrum Virtualize nodes from an I/O group are required to be isolated in terms of different fault domain and different update domain so that cloud does not bring them down for any maintenance activity simultaneously as far as possible (see Figure 3-4).



Figure 3-4 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. Azure's Availability Set construct allows us to express logical groupings of IBM Spectrum Virtualize VMs and define the high availability requirement.

Note: Consider the following points:

- ► Update domains indicate groups of VMs and underlying physical hardware that are brought down at the same time for Azure maintenance. With each node VM from an I/O group in a different fault domain, Azure ensures high availability 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.

3.4.1 Controller node proximity

A *proximity placement group* defined logical grouping makes sure that Azure places the node compute resources close to each other.

IBM Spectrum Virtualize Node VMs must be separated by a 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 a network latency that is not acceptable for cluster functioning. Therefore, to get node VMs as close as possible with lowest possible latency, we make use of construct of proximity placement group.

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

3.4.2 Disk types and multi-attach 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 3-5) allows two or more node VMs to simultaneously connect to the same disk and 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 3-5 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-managed disks to two VMs from same I/O group (see Figure 3-5 on page 37).

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) to be 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 3-2 lists the Azure disk types that are required for different use cases and workloads that are best served by it.

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

Table 3-2 Disk types required for different use cases and workloads

3.4.3 Load balancer for management traffic failover

In Azure, IP reassignment through IP failover from one VM to another VM can take much time. To solve this issue, Azure suggests the use of an Azure Load Balancer based design.

IBM Spectrum Virtualize design uses the services of the Azure Load Balancer for management IP failover. 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, load balancer service switches the traffic to another back-end IP from the surviving node (see Figure 3-6).

Figure 3-6 Loadbalancer for management traffic failover

3.4.4 Multipath for data traffic failover and better performance

For I/O traffic failover, IBM Spectrum Virtualize design relies on proven 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 and consequently, that path is stopped until the failed node is repaired. For this reason, it is required to 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:

- ► 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, it is suggested to increase the iSCSI sessions from the host to realize the high throughput performance.

3.5 Security design considerations

Cloud features shared the responsibility model of security between the cloud provider, IBM Spectrum Virtualize, and customers. It is important to understand the shared responsibility model.

Azure cloud provides security of compute, network, physical resources in cloud data centers and various virtual resources, such as VM, virtual network, 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 in 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 3-7) that factor in Azure-specific 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 3-7 Security design aspects

3.5.1 Restrictive network security rules

When IBM Spectrum Virtualize network is deployed by using the Azure Resource Management (ARM) template, network security groups are created and applied at the subnet and network interface level. Network security groups also are automatically configured for each subnet or interface to protect the virtual network. Default network security group (NSG) rules tighten security to achieve the goal of network security.

For more information, see this Microsoft Docs web page.

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

- Default network security group 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 to IBM Spectrum Virtualize configuration IPs or data traffic IPs.
- To reach out to 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.

3.5.2 Bastion access

Azure Bastion is a fully managed service that provides secure and seamless Remote Desktop Protocol (RDP) and Secure Shell Protocol (SSH) 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.

3.5.3 Operating system patch update interface and Azure Update Manager integration

Figure 3-8 on page 42 shows the construction of a node. In Azure, IBM Spectrum Virtualize is packaged in a container and run 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 public cloud. The following methods are available that can be used to meet this security requirement:

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



Figure 3-8 Operating system patch update interface and Azure Update Manager integration

3.5.4 Manual patching 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 RHELOS operating system 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. Azure Red Hat repository is used to retrieve and apply security patches to the node VMs.

3.5.5 Automated patching by using Azure Update Manager

Azure Update Management 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 the Spectrum Virtualize nodes:

- 1. Create an Azure automation account (for more information, see this web page).
- 2. Configure this automation account with 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, Update Manager automatically evaluates Azure VMs to maintain security compliance regarding released Critical and Security updates.

3.6 Solution architecture: IBM Spectrum Virtualize as a storage for all-in-cloud model

In the all-in-cloud model, as shown in Figure 3-9 on page 43, 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.



Figure 3-9 IBM Spectrum Virtualize as a storage for all-in-cloud model

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 Spectrum Virtualize are defined on volumes. These volumes are created from managed disks (Azure MDisks). All data transfer occurs through the IBM Spectrum Virtualize node, which is described as *symmetric virtualization*.

44 Implementation Guide for IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Version 8.4.3

4

Planning an installation on Microsoft Azure

This chapter describes the planning considerations for implementing IBM Spectrum Virtualize for Public Cloud on Microsoft Azure.

This chapter includes the following topics:

- ▶ 4.1, "Introduction" on page 46
- ▶ 4.2, "General planning on Microsoft Azure" on page 46
- ▶ 4.3, "Planning security access control on Microsoft Azure" on page 48
- ▶ 4.4, "Planning networking for Microsoft Azure" on page 52
- ▶ 4.5, "Planning an Azure virtual machine" on page 56
- ▶ 4.6, "Planning Azure managed disks" on page 58
- ▶ 4.7, "Planning deployment access" on page 59
- 4.8, "Storage performance optimization" on page 60
- 4.9, "Planning for data reduction pools" on page 61

4.1 Introduction

Before installing IBM Spectrum Virtualize for Public Cloud software in your Microsoft Azure cloud environment, ensure that all necessary planning steps are complete. This chapter describes the planning and preparation steps to provision the cloud infrastructure (network, server, and storage) components on Microsoft Azure, which are required for installing IBM Spectrum Virtualize for Public Cloud.

Background information about the Microsoft Azure cloud infrastructure requirements also is provided to help the reader who is unfamiliar with the Microsoft Azure plan for the IBM Spectrum Virtualize for Public Cloud placement into the larger context of an application environment.

To correspond with the pricing flexibility of supported cloud service providers, IBM Spectrum Virtualize for Public Cloud provides monthly and perpetual licensing options.

For more information about obtaining the IBM Spectrum Virtualize for Public Cloud license, contact your IBM Storage sales representative.

4.2 General planning on Microsoft Azure

Microsoft Azure offers 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 on Azure Marketplace.

The IBM Spectrum Virtualize for Public Cloud installation uses the Azure Resource Manager template to simplify provisioning of resources. This service uses a text file template to configure these resources.

When the installation template is started from Azure Marketplace, the user is prompted to provide information (such as customer ID) for entitlement check. The process provisions resources that are based on the settings that are defined in the installation template for the IBM Spectrum Virtualize for Public Cloud software.

4.2.1 IBM Spectrum Virtualize for Public Cloud prerequisites

The IBM Spectrum Virtualize for Public Cloud software is a "bring your own license" (BYOL) offering in Azure Marketplace. During deployment of the IBM Spectrum Virtualize for Public Cloud installation, the installation template verifies the proof of entitlement that indicates that a valid license is purchased from IBM.

To obtain the license and proof of entitlement for the software, complete the following steps:

- 1. See 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 your IBM Customer Number also are provided by email to the person who acquired the license.

4.2.2 Microsoft Azure prerequisites

Ensure that you complete the following tasks before installing IBM Spectrum Virtualize for Public Cloud software from the Microsoft Azure Marketplace:

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 the 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 recommended for users that are completing daily work on the system. These users have limited permissions, which restricts them from specific actions that are based on Azure policies.

For information more information about these recommended user profiles and security group settings, see 4.3, "Planning security access control on Microsoft Azure" on page 48.

3. Create an SSH key pair

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

For more information about creating SSH keys, see 5.1, "Installing IBM Spectrum Virtualize for Public Cloud on Azure" on page 64.

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 Azure Resource Manager (ARM) template services is used in IBM Spectrum Virtualize for Public Cloud installation. These ARM templates to 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 the 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 (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.

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

Network Security Groups

Azure network security group is a static set of rules that protects each network. A network security group contains security rules that allow or deny network traffic to or from an Azure resource. The network security groups 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 failover and is provisioned as part of IBM Spectrum Virtualize for Public Cloud installation. The Load Balancer service is configured per resource group.

4.3 Planning security access control on Microsoft Azure

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

Microsoft Azure provides role-based access control (RBAC) to manage access to Microsoft Azure resources. It 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 resource groups 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 minimizes unauthorized access to resources within your environment.

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

Table 4-1 lists the roles that are required for the deployment and management of IBM Spectrum Virtualize for Public Cloud software in Microsoft Azure.

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, and attaching Azure disks.	Installation step: Deploy IBM Spectrum Virtualize for Public Cloud
Management	This 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 user can run system setup and any related configuration tasks in IBM Spectrum Virtualize for Public Cloud management interfaces.	 Postinstallation steps: Complete system setup Configuring pools and assigning storage Configure volumes Configure hosts and mappings
Bastion	This user can create a separate user role to manage all bastion connections between your public and private network or include these permissions as part of the user role or the installer user role. A bastion host allows public networks to access private virtual network.	Postinstallation step: Create a bastion host

Table 4-1 Required user roles for IBM Spectrum Virtualize for Public Cloud deployments

4.3.1 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/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"
```

4.3.2 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/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"
```

4.3.3 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 Bastion User Role can be defined with the following permissions:

```
"Microsoft.Network/bastionHosts/read",
"Microsoft.Network/virtualNetworks/BastionHosts/action",
"Microsoft.Network/virtualNetworks/bastionHosts/default/action",
"Microsoft.Network/networkInterfaces/read",
"Microsoft.Network/virtualNetworks/read",
"Microsoft.Network/virtualNetworks/read",
"Microsoft.Network/virtualNetworks/subnets/virtualMachines/read",
"Microsoft.Network/virtualNetworks/virtualMachines/read",
"Microsoft.Network/virtualNetworks/virtualMachines/read",
"Microsoft.Network/publicIPAddresses/write",
"Microsoft.Network/bastionHosts/write",
"Microsoft.Network/publicIPAddresses/read",
"Microsoft.Network/publicIPAddresses/read",
"Microsoft.Network/publicIPAddresses/read",
"Microsoft.Network/publicIPAddresses/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 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 virtual network. If you are implementing your environment exclusively in Microsoft Azure as an all-in-cloud deployment, more networking objects are needed, such as a 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.

4.4.1 Network considerations for basic deployment

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

If you select an existing virtual network, ensure that the resource group 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 4.5, "Planning an Azure virtual machine" on page 56.

► 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. Load balancer is used for the management IP address only and the traffic is routed on all the ports.

Network security groups

In addition, network security groups are automatically configured for each subnet to protect the virtual network. A *network security group* filters network traffic to and from resources in the virtual network. It uses specific security rules that allow or deny inbound and outbound traffic to and from the virtual network.

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 network security group during a basic deployment are listed in Table 4-2.

Туре	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	ТСР	443	Secure HTTP (HTTPS) inbound traffic
Customer TCP rule	ТСР	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

Table 4-2 Default rules for Network security group

4.4.2 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 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/IKE parameters for VPN gateways 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/IKE parameters for Site-to-Site VPN Gateway connections
- Coordinate with your network administrator to ensure 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 virtual network 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 a network security group.

4.5 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 virtual central processing unit (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 random access memory that is allocated to the VM.

Data Disk

Indicates the number of virtual hard disk drives 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 4-3 compares the properties of each D-Series type VM that can be selected for your IBM Spectrum Virtualize for Public Cloud nodes.

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

Table 4-3 Properties of each D-Series type VM

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 4-4 lists the properties for the VM that is provisioned for the quorum node:

Table 4-4 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

4.6 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 resource group where the cluster was created by using Azure portal or 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.

4.6.1 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. You must 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 solid-state drive (SSD) or standard SSD MDisks, each extra mount of the MDisk is charged per month based on the disk size.

4.6.2 MDisk support

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

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
4.7 Planning deployment access

The basic deployment of IBM Spectrum Virtualize for Public Cloud creates a private virtual network 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 virtual network 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 4-5 lists the use case of a basic deployment in cloud model in which all setup and configuration occurs with the basic deployment of IBM Spectrum Virtualize for Public Cloud.

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.
Bastion service	Yes	For more information, see Permissions for Bastion service in the Azure documentation.	A bastion host can be created on the quorum node in your deployment or with the Azure Bastion Service.

Table 4-5 Access method for Basic deployment in cloud use case

Table 4-6 lists the use case of an all-in-cloud with replication model in which resources are in two private virtual networks that are separated by a public network. Data replication is configured between the two private virtual networks. You have the choice of two access methods: VNet peering and a VPN connection.

Table 4-6 Access method for 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 virtual networks 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 virtual networks. You might want to use a VPN if your security policy requires more protection.

Table 4-7 lists the use case of 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.

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.

Table 4-7 Access method for Hybrid cloud with replication use case

4.8 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.

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 4-8 is a best practice and must be adjusted manually by using the **chmdisk -tier -easytierload** command.

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

Table 4-8 Microsoft Azure MDisks assignment to Easy Tier level

After assigning an Azure MDisk to the suitable IBM Spectrum Virtualize for Public Cloud storage tier, run the **1smdisk** command to verify the change in the tier level.

4.9 Planning for data reduction pools

A deduplicated volume or volume copy can be created in a data reduction pool (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 memory to support deduplication.

Consider the following points:

- ► Avoid Global Mirror with Change Volumes 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 Comprestimator, 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 use of DRPs and deduplication on IBM Spectrum Virtualize for Public Cloud in Microsoft Azure is supported on all three VM type (see Table 4-3 on page 57).

62 Implementation Guide for IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Version 8.4.3

5

Implementing an IBM Spectrum Virtualize for Public Cloud on Microsoft Azure environment

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

- 5.1, "Installing IBM Spectrum Virtualize for Public Cloud on Azure" on page 64
- ▶ 5.2, "Logging in to IBM Spectrum Virtualize for Public Cloud on Azure" on page 73
- ► 5.3, "Configuring the cloud quorum" on page 88
- ► 5.4, "Configuring the back-end storage" on page 90
- 5.5, "Configuring a site to site virtual private network gateway for hybrid cloud connectivity in Azure Cloud" on page 100
- 5.6, "Configuring replication from on-premises IBM Spectrum Virtualize to IBM Spectrum Virtualize for Public Cloud on Azure" on page 101

5.1 Installing IBM Spectrum Virtualize for Public Cloud on Azure

Deployment video: As part of this IBM Redbooks publication, the authors also 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 the 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 Chapter 4, "Planning an installation on Microsoft Azure" on page 45 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 (ICN) 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 hAzure Marketplace.

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

- A valid Azure Account was created if such an account does not exist.
- The wanted Azure profiles for users were created.
- An SSH key pair was created. This key is used to access the Azure VM instance that is created and another key-based authentication.
- 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/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 5-1 on page 65. Review the Plans and the Support sections and then, click **Create**.

■ Microsoft Azure	Search resources, services, and docs (G+/)
Home > IBM Spectrum Virtualize for Public Cloud * IBM-Alliance-Global-Microsoft Partner-USA-NV-Armonk-HQ-IBMStorage-6201192	
IBM Spectrum Virtualize for Public Clo	Jd 👳 Add to Favorites
IBM-Alliance-Global-Microsoft Partner-USA-NY-Armonk-HQ-IBMStorage-62	01192
Create	
Overview Plans Usage Information + Support Reviews	
IBM Spectrum Virtualize for Public Cloud on Microsoft Azure provides software defined block s Microsoft Azure, as well as create all-in cloud data management using Azure Managed Disks. A native cloud block storage infrastructure and provides a way to mirror data from on premises also reduce the cost of your cloud storage, through greater efficiency in managing data and re based on the same software that runs in IBM's award-winning FlashSystem on premises storag a Service, specifically - in this case on Microsoft Azure.	torage that allows you to create solutions that span from on premises to mong many capabilities, this offering provides an optimization layer for to cloud for Disaster Recovery solutions, Dev/Test, or Data Migration. It can ducing infrastructure costs. IBM Spectrum Virtualize for Public Cloud is e, but it has been optimized to run on public cloud provider Infrastructure as
Working together, IBM Spectrum Virtualize and IBM Spectrum Virtualize for Public Cloud suppo data centers or between cloud data centers. These functions can be used to:	rt synchronous and asynchronous mirroring between on-premises and cloud
 Migrate data between on-premises and Microsoft Azure regions or between two Micros Implement disaster-recovery strategies between on-premises and Microsoft Azure or be Enable cloud-based DevOps with easy replication of data from on-premises sources Enhance performance and functionality, and lower cost of Azure Managed Disk block str provisioning, data reduction (compression and deduplication), and IBM Easy Tier. Enjoy consistent data management between on-premises storage and Microsoft Azure I 	oft Azure regions tween two Microsoft Azure regions. vrage with advanced data services such as IBM FlashCopy (snapshots), thin vlock storage.
IBM Spectrum Virtualize for Public Cloud capabilities include:	
 IBM Spectrum Virtualize on-premises and IBM Spectrum Virtualize for Public Cloud togel layer between on-premises and the cloud across heterogeneous storage pools that may Storage pooling and automated allocation with thin provisioning Easy Tier automated tiering Deduplication and compression to reduce cloud storage costs FlashCopy and remote mirror for local snapshots and remote replication Suport for virtualized and containerized server environments including VMware, Micros Microsoft Azure 	her enable a hybrid multi-cloud deployment with a single data management · exist in the data center. :oft Hyper-V, Red Hat OpenShift, CRI-O, and Kubernetes services on
Figure 5-1 IBM Spectrum Virtualize for Public Cloud or	n 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 5-2):
 - Subscription to use for deployment
 - Resource Group to group all resources
 - Region that is to host the resources
 - Project Tag to uniquely identify the deployed resources
 - (Optional) Enable Rollback support if a deployment fail occurs

■ Microsoft Azure	${\cal P}$ Search resources, services, and docs (G+/)					
Home > IBM Spectrum Virtualize for Pu	Home > IBM Spectrum Virtualize for Public Cloud >					
Create IBM Spectrum	/irtualize for Public Cloud					
Basics VM Selection Credentials	Networking Storage Review + create					
Project details						
Select the subscription to manage deplo manage all your resources.	yed resources and costs. Use resource groups like folders to organize and					
Subscription * 🕡						
Resource group * (i)						
	Create new					
Instance details						
Region * 🛈	East US 🗸					
Project Name						
Tag to identify deployment in a resource	group					
Tag * 🛈	redbook 🗸					
Rollback						
Rollback on failure. 🛈						

Figure 5-2 Basic configuration window for deployment on Azure

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

- 4. Select the size of 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 Chapter 4, "Planning an installation on Microsoft Azure" on page 45.

These VMs are used to run the 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 5-3 shows the information that is available during the VM selection.

≡ Microsoft Azure	۶ م	Search resources, services, and docs	s (G+/)						
Home > IBM Spectrum Virtualize for Public Cloud > Create IBM Spectrum Virtualize fo	r Public Cloud		Select a V	'M size					×
Create IBM Spectrum Virtualize for Basics VM Selection Credentials Networking Mid Spectrum Virtualize for Public Cloud Is deployed in a 3 M def the diverse as a guoren odd for the dust- different Auror Vist hat are supported for running IM Spect- Learn more Spectrum Virtualize for Public Cloud Node * O Fixed Size Quorum Node * O Is Standard IM 1 vcpu, 2 G mer Change size	Public Cloud ···	ing of 2 Azore VMs to select from 3	P Search by* Showing 3 VM Size ↑↓ > D-Series D165,33 D32,39 D44, y3	Mi size Displa Subscription: Microoff Azure Enterprise, John Family 12 General purpos General purpos	v cost : Monthly v Co v cost : Monthly v Co Essuito : s v CPUs †s = 16 = 22 = 04	PUIS : All RAM	(GB) : All term Learn more term for a for our discrete term more term for a for our discrete term for a for your general purpose n 32 32 32 32 32	ner choosing a or VV-1 size 25600 51200 80000	Group by series ✓ Temp storage (Gill) ↑↓ 128 256 512 ↓
Review + create < Previous Next : Crede	entials >		Select	Prices presented are estima include any applicable softv reseller, contact your reselle	tes in your local currency tl vare costs. Final charges wil r for full pricing details.	that include only Azure ill appear in your local	infrastructure costs and any disco currency in cost analysis and billi	unts for the subscript g views. If you purch	tion and location. The prices don't ased Azure services through a

Figure 5-3 VM Selection window during deployment

5. The Credentials window (see Figure 5-4 on page 68) shows all of the credentials that are related to IBM Spectrum Virtualize deployment; namely, the password for Spectrum Virtualize cluster, Customer entitlement check, SSH public key for access to VM and the notification email address.

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

≡ Microsoft Azure		♀ Search resources, services, and docs (G+/)						
Home > IBM Spectrum Virtualize for Publ	Home > IBM Spectrum Virtualize for Public Cloud >							
Create IBM Spectrum Virtualize for Public Cloud								
Pasies VM Selection Cradentials	Notworking Storage Devices + create							
	Networking Storage Review + create							
Spectrum Virtualize Management Crede	ntials							
Set password for the Security Administrato Learn more	r user profile (superuser) for management GUI.							
Password * (i)		✓						
Confirm password * 🕕	•••••	~						
Customer Entitlement								
Provide IBM Passport Advantage Customer purchase of the software license. The instal number Learn more	Number of BYOL offering. The IBM customer num lation template verifies entitlement to the software	ber is associated with the with this customer						
IBM Customer Number * 🛈		 ✓ 						
Notification								
The email address receives notifications on Learn more	the status of the installation							
Notification Email * 🕕		~						
1								
VM Credential								
Provide SSH public key to configure Spectr Learn how to generate SSH keys	um Virtualize VM nodes for secured access.							
SSH public key * 🕕		¥						
	and the second sec							
	/HTegbVmqVz+FRNDK+JX/YKz1h2G7oTfds= imp Learn more about creating and using SSH keys	orted-openssh-key						
Review + create < Previous	Next : Networking >							

Figure 5-4 Credentials window on deployment

Note: Use OpenSSH format for SSH keys. It is recommended to generate the SSH key from Azure.

6. In the Networking tab (see Figure 5-5 on page 69), select the Azure Virtual Network for this deployment. Users can select a VNet or create a virtual network 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 5-5 on page 69 shows the Network selection 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.

	$\mathcal P$ Search resources, services, and docs (G+/)						
Home > Marketplace > IBM Spectrum Virtualize for Public Cloud >							
Create IBM Spectrum \	/irtualize for Public Cloud						
-							
Basics VM Selection Credentials	Networking Storage Review + create						
Spectrum Virtualize is deployed in an Azu deployed in cluster subnet and a quorum create new VNet and new subnets or cho Learn more	re VNet across two subnets. The Spectrum Virtualize cluster VMs are VM is deployed in quorum subnet. User may provide new CIDR block to ose existing VNet and subnets in the same region						
Configure virtual networks							
Virtual Network * 💿	(new) sv-default-vnet						
	Create new						
Cluster Subnet * 🕠	(new) sv-cluster-subnet (10.1.0.0/24)						
Quorum Subnet * 🗊	(new) sv-quorum-subnet (10.1.1.0/27)						
Review + create < Previous	Next : Storage >						

Figure 5-5 Network selection during deployment

7. In the Storage tab, users select the default back-end storage that is to be attached to the Spectrum Virtualize nodes as part of the deployment. By default, two storage disks with similar configurations as provided by the user are provisioned and attached to the Spectrum Virtualize nodes. The user can choose from two supported types of Azure disks. The backend storage for Azure deployment is shown in Figure 5-6.

≡ Mic	crosoft Azure		$\mathcal P$ Search resources, s	services, and docs (G+/)				
Home > IB	Home > IBM Spectrum Virtualize for Public Cloud >							
Create	BM Spectrum Vi	rtualize for Public Cloud						
Basics	VM Selection Credentials	Networking Storage Review + crea	te					
Select the t provisionin installation. Check prici	Select the type of Azure Managed Disk to be used with IBM Spectrum Virtualize for Public Cloud for storage provisioning. A minimum of two volumes is required for initial cluster creation, and more can be added after installation. Check pricing details of Azure managed disks							
Azure Dis	k							
Disk Type *	* 🛈	Standard SSD (LRS)		✓				
Disk Size *	· (i)	512 GB		▼				

Figure 5-6 Storage selection during deployment

8. After all selections are made by the user, a basic validation can be run to verify the information that as provided by the user. In the Review and Create tab, the 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 5-7 on page 70.

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

	Microsoft Azure					𝒫 Search		
Home	> IBM Spectrum Vi	rtualize for Publ	ic Cloud >					
Cre	ate IBM Spe	ectrum Vi	rtualize f	or Pub	lic Cloud			
1								
0	Validation Passed							
5								
Basic	VM Selection	Credentials	Networking	Storage	Poviow + crost			
Dasic	s vivi selection	Credentiais	Networking	Storage	Review + create	_		
PRO	DUCT DETAILS							
IBM	Spectrum Virtualize f	or Public						
Clou	ł							
by IB	M-Alliance-Global-Mi	crosoft						
Partn	Partner-USA-NY-Armonk-HQ-							
IBMS	orage-6201192							
Term	s of use Privacy poli	cy						

Figure 5-7 Validation Passed message

 The deployment process takes approximately 20 minutes to complete. Progress can be monitored by reviewing the Deployment window, which is displayed after the Create button is clicked (see Figure 5-8).

	P Search resources, serv	ces, and docs (G+/)	Đ	i 6 🔎 🛛 🤉 🖉 🖓 🖬	0
Home > ibm-alliance-usa-ny Deployment	-armonk-hq-ibmstorage-6201192202112	23120002 Overview 🖈 …			×
	🗊 Delete 🚫 Cancel 🖄 Redeploy 💍 Refresh				
👶 Overview	\bigcirc We'd love your feedback! \rightarrow				
≅ Outputs ▶ Template	Deployment is in progress Deployment name: ibm-alliance-usa-my-armonk-hq-ibmstorage-62 subscription: Resource group: ibmyredbook_spip_aure Opployment details (Download)	Start time: 12/23/2021, 12:02:37 PM Correlation ID: 87252701-70c5-4Ca9-a1ed-e189691664ae			
	Resource	Туре	Status	Operation details	
	 sv-redbook-quorum 	Microsoft.Resources/deployments	Created	Operation details	
	Sv-redbook-LB/sv-LoadBalancerFrontend2	Microsoft.Network/loadBalancers/frontendlpConfigurations	ок	Operation details	
	sv-redbook-L8/sv-Load8alancerFrontend1	Microsoft.Network/loadBalancers/frontendlpConfigurations	ок	Operation details	
	Sv-redbook-LB	Microsoft.Network/loadBalancers	ок	Operation details	
	sv-redbook-nodes-1	Microsoft.Resources/deployments	Created	Operation details	
	sv-redbook-nodes-2	Microsoft.Resources/deployments	Created	Operation details	
	Sv-redbook-common	Microsoft.Resources/deployments	ок	Operation details	
	📀 default	Microsoft.Resources/tags	ок	Operation details	
	pid-1ffb4185-a9f9-4255-9085-91609153c548-partnercenter	Microsoft.Resources/deployments	ОК	Operation details	

Figure 5-8 Deployment in progress

Alternatively, the user can also monitor the deployment progress in the Deployment sections that are under the resource group that is used for deployment, as shown in Figure 5-9.

Microsoft Azure	P Search	resources, services, and docs (G+/)		E 4 4 6	0 R 🥯
Home > ibm_redbook_svpc_azure					
ibm_redbook_svpc_	azure Deployments …				×
P Search (Ctrl+/) «	💍 Refresh 🛇 Cancel 🁎 Redeploy 🗊 Delete	↓ View template			
(e) Overview	Filter by deployment name or resources in the deployme	nt			
 Activity log 					
R Access control (IAM)	Deployment name	Status	Last modified	Duration	Related events
Tags	sv-redbook-quorum	Deploying	12/23/2021, 12:07:07 PM	5 minutes 7 seconds	Related events
🛧 Resource visualizer	sv-redbook-nodes-2	Opploying	12/23/2021, 12:07:42 PM	5 minutes 19 seconds	Related events
🗲 Events	sv-redbook-nodes-1	Succeeded	12/23/2021, 12:08:27 PM	5 minutes 4 seconds	Related events
Settings	sv-redbook-common	Succeeded	12/23/2021, 12:03:13 PM	30 seconds	Related events
Deployments	pid-1ffb4185-a9f9-4255-9085-91609153c548-pa	L 🥝 Succeeded	12/23/2021, 12:02:42 PM	1 second	Related events
Security	ibm-alliance-usa-ny-armonk-hq-ibmstorage-620	😨 Deploying	12/23/2021, 12:08:38 PM	6 minutes 5 seconds	Related events
Policies					
Properties					
A Locks					
Cost Management					
Cost analysis					
Cost alerts (preview)					
(§) Budgets					
Advisor recommendations					
Monitoring					
Insights (preview)					
Alerts					
mi Metrics					
Diagnostic settings					
🗭 Logs					
Advisor recommendations					



Figure 5-10 shows the successful deployment for IBM Spectrum Virtualize for Public Cloud on Azure.

Microsoft Azure	P Search resources, services, and docs (G+/)	e f 🖓 🕸 🧿 🖓 🧶
Home > ibm-alliance-usa-ny Deployment	-armonk-hq-ibmstorage-620119220211223120002 Overview 💉 …	×
P Search (Ctrl+/) «	💼 Delete 🛇 Cancel 🏦 Redeploy 🖒 Refresh	
👶 Overview 👰 Inputs	Wed low your feedbackt →	
E Outputs	Vour deployment is complete Peployment anne: bin-allance-vis.ny-annotk-hq-binstorage-62. Start time: 12/23/2021, 12/2023 PM Correlation ID: 8/7522701-70c5-4cae-sted+189691644ae Correlation ID: 8/7522701-70c5-4cae-sted+189691644ae Veployment details (Download) Vent sets Correlation ID: 8/7522701-70c5-4cae-sted+189691644ae Correlation ID: 8/7522701-70c5-4cae-sted+189691644ae Correlation ID: 8/7522701-70c5-4cae-sted+189691644ae	Wiccode Defender for Cloud Secure your apps and infrastructure Go to Microsoft Defender for Cloud > Free Microsoft Tutorials Sart Learning today > Work with an expert Azare experts are service provider partners who can high manage your asset on Azare and be your first line of support. Find an Azare expert >

Figure 5-10 Successful deployment

10. Click the **Outputs** tab to check the various IP addresses that are deployed, as shown in Figure 5-11.

Microsoft Azure	𝒫 Search resources, services, and docs (G+/)	E 57 🖉 🕸 🖉 🔤 🛛	0						
Home > ibm-alliance-usa-ny-armonk-hq-ibmstorage-620119220211223120	Home > ibm-alliance-usa-ny-armonic-hq-ibmstorage-6201192-20211223120002								
ibm-alliance-usa-ny-armonk-hq-ibmstorage-	ibm-alliance-usa-ny-armonk-hq-ibmstorage-6201192.20211223120002 Outputs ···· ×								
P Search (Ctrla)			^						
primaryClusterIP									
10.21.0.6			D						
			_						

Figure 5-11 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 set-up process. This email includes more information about the deployment (see Figure 5-12).

[EXTERNAL root to] IBM SVPC Clu	ster Deployment: Cluster ibm_redbook_svpc_azure is created successfully
From	"root" <root@eastus.cl< th=""><th>oudapp.azure.com></th></root@eastus.cl<>	oudapp.azure.com>
То		
Your SVPC IP address Item	cluster ibm_red es created in re	book_svpc azure is created successfully. source group: T
Cluster IF	,	10.21.0.6
Nodel Serv	rice IP	0.21.0.27
Nodel Node	IP 1	10.21.0.25
Nodel Node	IP 2	10.21.0.20
Nodel Port	: IP 1	10.21.0.26
Nodel Port	: IP 2	10.21.0.21
Node2 Serv	rice IP	10.21.0.11
Node2 Node	IP 1	10.21.0.9
Node2 Node	IP 2	10.21.0.14
Node2 Port	: IP 1	10.21.0.10
Node2 Port	IP 2	10.21.0.15
To access	IBM Spectrum Vin	rtualize for Public Cloud management GUI, please follow this link:
https://ww	w.ibm.com/docs/e	en/spectrumvirtualizecl/8.4.x?topic=azure-completing-post-installation-tasks

Figure 5-12 Sample email of successful deployment

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

IBM Spectrum Virtualize for Public Cloud on Azure supports deployments in an-all-in-private network on Azure Cloud. The deployment does not provide any public IP address to any resource. All traffic from the internet to the Azure VMs that are deployed as part of an IBM Spectrum Virtualize deployment is disabled and controlled by the network security group (NSG) rules in Azure.

Note: To allow traffic from specific port/IP, 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 virtual network with no public inbound access. Therefore, to access the resources that are provisioned by the deployment, the user must use the Azure Bastion Service or configure an Azure VPN gateway. For the purposes of this publication, we configured the Azure Bastion Service, as explained next.

5.2.1 Configuring the Azure Bastion Service

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

Complete the following steps to configure an Azure Bastion Service:

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

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 5-13 shows a subnet that was created under a virtual network to be used for the Azure Bastion Service.

		₽ Search resour	ces, services, and docs (G+/)			E & @ @ R	
Home > into a set a set a set	Home > in vnet						×
vnet Su Virtual network	vnet Subnets ····						
₽ Search (Ctrl+A ≪	+ Subnet + Gateway si	ibnet () Refresh 🛛 💫 Manaci	e users 📋 Delete			Name	Copy to clipboard
(A) Quantinu	,,					AzureBastionSubnet	
	₽ Search subnets					Subnet address range * ()	
Activity log	Name 1.	IPv4 ↑.	IPv6 ↑.	Available IPs 1.	Delegated to 1.	10.21.21.0 - 10.21.21.255 (2)	i1 + 5 Azure reserved addresses)
	cuidafaulticubnat	10.21.0.0/24		223		Add IPv6 address space ①	
 Tags Diagnore and roke problem: 	sv-default-nuorum	10.21.1.0/24		250		NAT gateway ①	
2 Diagnose and solve problems	AzureBastionSubnet	10.21.21.0/24		248		None	~
Settings						Network security group	
Address space						None	~
S Connected devices						Route table	
Subnets						None	~
DDoS protection						CERTIFICE ENDPOINTE	
🖶 Firewall						SERVICE ENDFOINTS	
Security						Create service endpoint policies to allow traffic to specific azure res over service endpoints. Learn more	ources from your virtual network
🏠 Network manager						Services ()	
DNS servers						0 selected	~
Peerings							
Service endpoints						SUBNET DELEGATION	
Private endpoints						Delegate subnet to a service ③	
Properties						None	~
🔒 Locks							
Monitoring							
💵 Alerts							
ma Metrics							
Diagnostic settings						Save Cancel	

Figure 5-13 Azure Bastion Subnet in an existing Virtual Network

2. Create the Azure Bastion Service resource by searching for "Bastion" in the **Create a resource** window on the Azure Portal (see Figure 5-14).

Microsoft Azure		℅ Search resources	services, and docs (G+/)		5 G
Home > Create a resource >					
Bastion 🖈 … Microsoft					
Bastion Microsoft	Add to Favorites				
* 2.0 (5 Azure ra	atings)				
Create					
Create					
Overview Plans Usage Infor	mation + Support Reviews				
Bastion enables seamless secure RD	P/SSH connectivity to Azure Virtual Ma	chines in your Azure Virtual Networks	directly in your web browser and with	out the need of	
public IP on your Virtual Machines.				Care All	
More products from Microsof				Secon	
	•	•	•		
Device Update for IoT Hub	Front Door Standard/Premium (Preview)	Azure VMware Solution	API App		
Microsoft	Microsoft	Microsoft	Microsoft		
Azure Service	Azure Service	Azure Service	Azure Service		
Securely and Reliably update your devices with Device Update for IoT Hub.	Azure Front Door Standard/Premium (Preview) is security led, modern cloud CDN	Azure VMware Solution (AVS) combines the VMware Software Defined Data Center (SDDC) with	scalable RESTful API with enterprise grade security, simple access control and auto SDK generation		
Create 🗸 🗢	Create V 🗢	Create V 🗢	Create V 🗢		

Figure 5-14 Azure Bastion Service on Azure Portal

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

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

5.2.2 Connecting to the Azure virtual machine by using Azure Bastion Service

To connect the Azure virtual machine (VM) in all-in-cloud deployment use the Azure Bastion Service as configured in 5.2.1, "Configuring the Azure Bastion Service" on page 74.

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 (http://portal.azure.com).
- 2. Find the VM that is to be connected under the newly deployed resource group, as shown in Figure 5-15.

	∠P Search resources, services, and docs (G+/)		E 🛱 🗘 🍩 🕐	
Home >				
(ibm_redbook_svpc_a	azure 🖉 …			×
P Search (Ctrl+/) « (*) Overview Activity log Access control (IAM) Tars	+ Create III Edit columns ① Delete resource group ① Refresh ↓ Export to CSV % Oper ∼ Essentials Subscription (More) : Subscription (D : Tags (Edit) : Cluster(PFailover: LB_Based_Cluster(P, Deployment PFailoverEnabled : Faise	n query [®] Assign tags → Move [®] Delete s Deployments : 6 Succeedee Location : East US SVPC_Default_MTU : 1500 deploymentTag : sv-redbook	Export template R Feedback Open in mobile Cluster_Deployment_Status : DEPLOYMENT_STATUS_SUCCESS	JSON View V More (5)
 Resource visualizer Events 	Resources Recommendations (13)			
Settings	Filter for any field Type == all X Location == all X ⁺ \bigtriangledown Add filter			
Deployments	Showing 1 to 23 of 23 records. Show hidden types ①		Group by type	✓ Elist view ✓
Security		Type 🔿	Location 14	
Policies	> Availability set	96		^
😂 Properties	> Disk			
🔒 Locks	> Key vault			
Cost Management	> Load balancer			
\$. Cost analysis	> Managed Identity			
Cost alerts (preview)	> Network interface			
Budgets	> Network security group			
Advisor recommendations	> Proximity placement group			
Monitoring	Virtual machine			
Insights (preview)	🔲 🖳 sv-redbook-node1-vm	Virtual machine	East US	
Alerts	sv-redbook-node2-vm	Virtual machine	East US	
i Metrics	U 📮 💀 sv-redbook-quorum	Virtual machine	East US	··· 🗸
Diagnostic settings	Broulour Bage 1 X of 1 Novt >			
P Logs	s rievious rage 1 v or i ivent 2			
Advisor recommendations				

Figure 5-15 VMs in resource group

3. Go to the Quorum VM to connect and select the **Connect** option, as shown in Figure 5-16.

Microsoft Azure		$\mathcal P$ Search resources, services, and docs (G+/)			N 6 0 0 7 10000	-
Home > ibm_redbook_svpc_azure >						
sv-redbook-quorum	1 🖈 …					×
Search (Ctrl+/) «	💋 Connect ▷ Start 🤇 Restar	t 🔲 Stop 🞉 Capture 📋 Delete 💍 Refresh 🏮 Op	oen in mobile 🛛 🐰 CLI / PS 🛛 🖗 Feedback			
📮 Overview	SSH	nould be installed on your machines. →				
Activity log	RDP					
Access control (IAM)	Bastion				JSO	N View
Tags	DBC	DOK_SVPC_AZURE	Operating syste	m : Linux (redha	it 8.3)	
Diagnose and solve problems	Location : Fast US		SiZe Dublic ID addres	: Standard Bi	ms (1 vcpu, 2 GB memory)	
Sattings	Subscription (Move)		Virtual network	subnet : Nor v	met/sv-default-guorum	
Setungs	Subscription ID : monthline		DNS name			
Networking	Tags (Edit) : Click here t	to add tags				
Disks	Properties Monitoring Capal	bilities (7) Recommendations (7) Tutorials				
📮 Size	Virtual machine		Setworki	ng		
Security	Computer name	sv-redbook-quorum	Public IP	iddress		
Advisor recommendations	Health state	-	Public IP a	iddress (IPv6)		
Extensions + applications	Operating system	Linux (redhat 8.3)	Private IP	address	10.21.1.4	
🐔 Continuous delivery	Publisher	ibm-alliance-usa-ny-armonk-hq-ibmstorage-6201192	Private IP	address (IPv6)		
Availability + scaling	Offer	ibm-svpc-azure-843-base-vm	Virtual ne	work/subnet	_vnet/sv-default-quorum	
Configuration	Plan	svpc-8430-211209-2206	DNS nam	2	•	
s Identity	VM generation	VI	Size			
Properties	Agent status	Ready	Size		Standard B1ms	
A Lasta	Agent version	None	VCPUs		1	
Louis	Host	-	RAM		2 GiB	
Operations	Proximity placement group					
✓ Bastion	Colocation status	N/A	S Disk			
③ Auto-shutdown	Capacity reservation group		OS disk Azura dis	encryption	sv-reubook-quorum_Osbisk_i_ucq3ac5c3be3402cbo8b0e85app0g993	
🔗 Backup	Availability + scaling		Ephemera	l OS disk	N/A	

Figure 5-16 Virtual Machine Connect option

4. Select the **Bastion** option to connect and then, select the configured 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/private key pairs that was used during deployment.

Figure 5-17 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.

	ア Search resources, services, and docs (G+/)	
Home > ibm_redbook_svpc_azure > sv-	redbook-quorum	
Sv-redbook-quorum	n Bastion …	
Search (Ctrl+/)	Azure Bastion Service enables you to securely and seamlessly RDP & SSH to your VMs in your Azure virtual network, without exposing a public IP on the VM, directly from the Azure portal, without the need of any additional client/agent or any piece of software. Learn more cf Bastion (2 found) * 	
Settings	✓ Connection Settings	
2 Networking	✓ Open in new window	
🖉 Connect	Username * 💿	
Bisks	centos 🗸	
📮 Size	Authentication Type * ①	
Security	Password SSH Private Key SSH Private Key Tom Local File SSH Private Key Start	
Advisor recommendations		
Extensions + applications	Local File *	
🐔 Continuous delivery		
Availability + scaling	✓ Advanced	
Configuration	Connect	
💲 Identity		
Properties		
🔒 Locks		
Operations		
≯ Bastion		
Auto-shutdown		
🔗 Backup	R Tell us what you think of the Bastion experience	

Figure 5-17 Connection page for virtual machine by using 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 5-18.

Note: All connections that use Azure Bastion are done by using a web browser.



Figure 5-18 Connection to Quorum Node that uses Azure Bastion Service by using a web browser

5.2.3 Configuring GUI for Spectrum Virtualize for Public Cloud cluster on Azure and completing the installation

For all in Cloud deployments, the user must configure a Windows Host Server VM in the cloud under the existing Virtual Network that was used during the deployment process.

Note: A new virtual network also can be used. However, it must be peered by using the Virtual Network Peering service on Azure.

For the purposes of this document, a Windows server is preconfigured in an existing virtual network.

When the installation is complete, you can log in to the IBM Spectrum Virtualize for Public Cloud on Azure cluster through the WebGUI, as shown in Figure 5-19 by using a preconfigured Windows Server and Azure Bastion Service.

Complete the following steps to complete the installation:

1. Use the Cluster IP that is provided in email to connect the WebGUI for Spectrum Virtualize (see Figure 5-19).

8	Spectrum Virtualize for Public Cloud Storage Management (ibm_redbook_svpc_azure) superuser	
	Fassword	
<u>IBM</u> .		

Figure 5-19 WebGUI interface for SVPC on Azure

2. You are redirected to the Welcome window. Click Next (see Figure 5-20).



Figure 5-20 Welcome window

3. You are redirected to License Agreement window. Read the license agreement and select **I agree** and then, **Next**, as shown in Figure 5-21.

System Setup							×
🕑 Welcome	Read the lice	ense agr	eement care	fully.			
C License Agreement	License Ad	ddendum	Java Notices	Non-IBM Licen	ses Additional Licent	ses and Notices	
Change Password	International	Program Li	cense Agreemer	it			^
System Name	PARCY - GENERAL FERMS BY DOWNLOADING, INSTALLING, COPYING, ACCESSING, CLICKING ON AN "ACCEPT" BUTION, OR						
Licensed Functions	ARE ACCEPTING : AUTHORITY TO B	G THE PROG THESE TERM IND LICENS	RAM, LICENSEE IS ON BEHALF OF EE TO THESE TE	AGREES TO THE LICENSEE, YO RMS.	TERMS OF THIS AGREE U REPRESENT THAT YOU	MENT, IF YOU HAVE FULL	
DNS	IF YOU DO NOT I INSTALL, COPY,	AGREE IO I ACCESS. C	HESE TERMS OR LICK ON AN "AG	DO NOT HAVE AN CEPT" BUTTON.	UTHORITY: 1) DO NOT OR USE THE PROGRAM:	DOWNLOAD, AND 11)	
Date and Time	FROMPTLY REIUR FROM WHOM IT W	N THE UNUS AS OBTAINE	ED MEDIA, DOCU	MENTATION, AND OF THE AMOUNT	D PROOF OF ENTITLEME I PAID. IF THE PROGR	NT TO THE PARTY AM WAS	
Call Home	DOWNLOADED, DE	SINOT ALL	COPIES OF THE	PROGRAM.			
Storage Insights	INIS Internation Documents (togo IBM regarding this ISI3 rep)	onal Frogr ether the the use of	"Agreement") a "Agreement") a a Program. Tr	eement (IFLA) ire the comple: ie country requ	and applicable iran te agreement between wired terms included	Licensee and in Part 2 of	
Support Assistance	Transaction Do	cuments (T	Del provide a	description.	information, and ter	ma recarding	~
Summary	-		,				
	I agree with t	he terms in t	he license agreem	ent.			
	 I do not agree 	e with the ter	ms in the license	agreement.			
				_			
	Cancel			ľ	6 Back	Next ►	

Figure 5-21 License Agreement window

4. You are redirected to the Change Password window, as shown in Figure 5-22. Change your password, and then, click **Apply** and **Next**.

System Setup				×
🕐 Welcome	Change Password			
 License Agreement 	The password mu	st be reset before proceed	ling with system configuration.	
Change Password	User name:	superuser		
System Name	New password:	•••••		
Licensed Functions	Confirm password:	•••••		
DNS Date and Time	Password Minim	d requirements um 8 characters long		
Call Home	 Mustin 	ot include problematic chara	cters (ex: control characters), or sta	rt or end with a space
Ctorada Insidhte				
Storage Insigna				
Support Assistance				
Summary				
		_		
	Cancel		- Back	Apply and Next 🕨

Figure 5-22 Change Password window

5. Change your cluster name (if needed), which defaults to the resource group name, as shown in Figure 5-23. Click **Apply** and then, **Next**.

System Setup			×
Welcome License Agreement Change Password System Name Licensed Functions DNS Date and Time Call Home Storage Insights Support Assistance Summary	iystem Name inter a name for the system: Iom_recibecols_svcc_azore		
	Cancel	⊲ Back	Apply and Next ►

Figure 5-23 Change Cluster Name window

6. Enter your capacity license in accordance with your IBM agreement, as shown in Figure 5-24. Click **Apply** and **Next**.

System Setup		×
Welcome Ucense Agreement Change Password System Name Ucensed Functions DNS Date and Time Call Home Storage Insights Support Assistance Summary	Licensed Functions Additional licenses are required to use certain system functions. For auditing purposes, retain the license agreement for proof of compliance. External Virtualization:	
	Need Help Cancel Apply and Next	

Figure 5-24 Capacity License window

7. Enter the DNS Setting to manage the resources that are on an external network, as shown in Figure 5-25. For our example, we skip this step.

System Setup				:
Welcome	DNS (Optional)			
License Agreement	You can create, delete, or located on external netwo	change domain name serv rks. <u>Read More.</u>	ers, which manage names (of resources that are
System Name	Define DNS serve	ers to use domain names w	hen configuring servers on	the system.
Licensed Functions	DNS Server			
DNS Date and Time	IP Address:	Name:		⊕ ⊡
Call Home				
Storage Insights				
Support Assistance				
Junning y				
	Cancel		⊲ Back	Next ►

Figure 5-25 DNS Setting window

 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*.

 Ensure that the time zone is set. For ease of troubleshooting across multiple time zones, it is a best practice to use Greenwich mean time or Coordinated Universal Time+0, as shown in Figure 5-26.

System Setup				×
e Welcome	Date and Time			
License Agreement	Specify the appropriate tir	ne zone for your system.		
Change Password	Time Zone:			
System Name	(GMT) Monrovia		-	
Licensed Functions				
ONS				
Date and Time				
Call Home				
Storage Insights				
Support Assistance				
Summary				
	Cancel		< Back	Apply and Next ►

Figure 5-26 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 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 5-27.

System Setup				×
Welcome	Internal Proxy Se	rver		
Change Password	Call home with cloud server to support. A DNS server is <u>Information</u>	ices can be used with an in s required to use a proxy s	ternal proxy server to send erver for call home with clo	d call home messages oud services. <u>More</u>
System Name	Use proxy to send data. Ves			
DNS	No			
Call Home Call Home Internal Proxy Server System Location Contact				
Storage Insights				
Support Assistance Summary				
	Cancel		⊲ Back	Next ►

Figure 5-27 Call Home Setting by using Cloud Services

This verification should succeed, as shown in Figure 5-28, which is the System Location window.

System Setup					×
Welcome	Connection to the set	upport center was succ	essful!	×	
C License Agreement C License Agreement C License Agreement System Name C Licensed Functions	System Location Service parts should be shippe Company name:	ed to the same physical loo	cation as the system.		
DNS Date and Time	System address: City:				
Cata Pointe Transmission Settings Internal Proxy Server System Location Contact	State or province:	ж			
Storage Insights Support Assistance Summary	Country or region:	Not Set	•		
	Machine location: _	eastus			
	Cancel		⊲ Back	Next 🍽	

Figure 5-28 Successful verification of Call Home

11.Complete the Call Home configuration by entering the contact information, as shown in Figure 5-29.

System Setup				×		
Welcome License Agreement	Contact The support center contacts	this person to resolve issue	es on the system.			
Change Password Change Name System Name	Enter business-to-business contact information To comply with privacy regulations, personal contact information for individuals with your organization is not recommended.					
Ucensed Functions ONS Ons	Name:	IBM Redbook				
Call Home Call Home Transmission Settings Call Home Transmission Settings Call Home Contact	Email: Phone (primary): Phone (alternate):	ibmredbook@ibm.com				
Storage Insights Support Assistance Summary	IBM may use my contact da	ta to keep me informed of :	Storage related products, s	ervices and offerings.		
	Cancel		■ Back	Apply and Next 🕨		

Figure 5-29 Call Home contact information

12.IBM Storage Insights often is configured next. However, this configuration is not done during EasySetup. It requires registering for a no-charge account and the installation of a DataCollector, which can be installed on the Bastion host. Figure 5-30 shows the IBM Storage Insights configuration window.

Skip this step for now.

Note: For more information about IBM Storage Insights, see "Capacity monitoring in IBM Spectrum Control and IBM Storage Insights " on page 130.

System Setup					×
 Welcome License Agreement Change Password System Name Licensed Functions DNS Date and Time Call Home Storage Insights Support Assistance Summary 	I use IBM Storage Insigh To add the storage system: 1. Log in to Storage Insig 2. Click Add Storage Sy When you add the storage s address: 10.21.0.6 €	tts: stem ystem, enter this IP	Why should I use Stora Let's face it. Storage perfot to maintain and troublesh for every minute you can't Insights monitors perform collaboration with consult resolve issues faster. Best you will get all the credit. system to start. Sign Up Learn Mo	age Insights? prmance can be tough oot. Costs skyrocket access data. Storage nance for easy ants and experts to of all, it's free and Just register your pre	
	Cancel		◄ Back	Next ►	

Figure 5-30 Storage Insights window

13. You can configure your Remote Support Proxy (RSP), as shown in Figure 5-31. For our example, this step is skipped.

System Setup				×
 ⊘ Welcome ↓ ⊘ License Agreement 	Support Centers Support centers respond to ma support centers are configured	nual and automatic serv	vice requests from the syste	em. The following
Change Password	Name default_support_center0 default_support_center1	IP Address 129.33.206.139 204.146.30.139	Port 22 22	
Uicensed Functions	Remote Support Proxy (Op	tional)		
Date and Time	 A proxy is required to direct connections to 	r network configuration the network.	is using a firewall, or for sy	stems without
Call Home Control Call Home C	Name	IP	Port	•
Summary				
	③ Need Help Cancel		⊲ Back	Apply and Next 🕨

Figure 5-31 Remote Support Assistance Config window

14. Figure 5-32 shows a summary of your cluster configuration. After careful review, click **Finish** to complete the setup.

System Setup						×
Welcome Subscription Subscription Subscription Subscription Subscription	Summary System Information System name: Code level;	ibm_redbook 8.4.3.0	Date: Time; Time zone:	Jan 4, 2022 12:13:12 PM GMT		
Change Password System Name Licensed Functions	Licensed Functions External Virtualization:	1 T	В			
	Call Home A Call home is not enable	d. Will remind in 3	0 days.			
Date and Time Order Call Home	Proxy Details Configured:	No				
 Storage Insights 	Support Assistance					
Support Assistance	Remote Support: Proxy Servers:	Yes 0				
	Cancel			⊲ Back	Finish	

Figure 5-32 Summary window

5.3 Configuring the 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 virtual networks.

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 \rightarrow System \rightarrow IP Quorum, as shown Figure 5-33.

IBM Spectrum Virtualize for Public Cloud	ibm_redbook_svpc_azure System	👍 🖡 🕐 superuser (Security Administrator) 🗸
<	Date and Time	IP Quorum Download the quorum application and install it on your network. This application serves as tie
Dashboard	Licensed Functions	breaker for the system if communication is disrupted. For instructions on how to install it, <u>click</u> https://docenter.click.com/c
എ ⁹ Monitoring	Update System	Detected IP quorum Applications
Pools Pools	VVOL	IP Address System Name State
Volumes	Volume Protection	10.21.1.4 sv-redbook-quorum.internal.cl v Online
Hosts	Resources	<>
Copy Services	IP Quorum	Disks containing configuration backup No Quorum MBiaks found.
Access	I/O Groups	You can configure an IP quorum application as a quorum device on your system. Learn more
(O) Settings	Transparent Cloud Tiering	Quorum Setting
	External Scheduling Applications	
		Latency 0 ms 0 ms 0 ms Bandwidth 0 MBps 0 MBps 10 MBps 10 0 0 0

Figure 5-33 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 FlashSystem systems. Therefore, they are automatically allocated on the VM instance boot device for each IBM Spectrum Virtualize node.

The IBM Spectrum Virtualize command **1 squorum** 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 guaranteed for node-to-quorum traffic.

5.4 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 MDisks. As part of the initial default installation, two Azure MDisks are allocated to the IBM Spectrum Virtualize cluster (see Figure 5-34).

IBM Spectrum Virtualize for Public Cloud	ibm_redbook_svpc_azure MDisks by Pools					👍 🛃 🕜 superuser Security Administrator) 🗸
K	⊕ Create Pool 🛛 🗮 Actions 👻					Default V Contains V Filter V4
A Dashboard	Name	Cloud Disk ID	Cloud Disk Type	State	Usable Capacity	Written Capacity Limit
	 Unassigned MDisks (2) 					
می ^و Monitoring	mdisk1	sv-redbook-Mdisk-2	standardSSD	✓ Online	512.00 GiB	512.00 GiB
	mdisk0	sv-redbook-Mdisk-1	standardSSD	✓ Online	512.00 GiB	512.00 GiB
Pools						
Volumes						
Hosts						
Copy Services						
Access						
(O) Settings	<					
	Cloud Disk (standard 1.00 TB 2 Molas Asage					
		Latency Oms	Read Write	Bandwidth O MBns 0 M		

Figure 5-34 Default storage added as part of 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 5-35.

IBM Spectrum Virtualize for Public Cloud	ibm_redbook_svpc_azure MDisks by Pools					👆 湯 🥐 superuser Security Administrator	ă
X	⊕ Create Pool 🛛 🗮 Actions 👻					Default Contains Filter	74
A Dashboard	Name	Cloud Disk ID	Cloud Disk Type	State	Usable Capacity	Written Capacity Limit	
	V Unassigned MDisks (2)						
we Monitoring	mdisk1	sv-redbook-Mdisk-2	standardSSD	🗸 Online	512.00 GiB	512.00 GiB	
	mdisk0	sv-redbook-Mdisk-1	standardSSD	🗸 Online	512.00 GiB	512.00 GiB	
Pools		Create Pool			×		
Volumes		Pool Name: Pool0					
Hosts		Extent Size: 4.00 GiB -	Maximum Capacity: 512.00 TiB				
Copy Services		Additional Options:					
Access		! Warning If the usable	capacity usage of a da	ata reduction pool exceed	Is more		
¿O ∫ Settings	Selecting 1 MDisk	than 85%, I/ of usable cap capacity recl	(O performance can be pacity available in data lamation can be perfor	affected. The system ne reduction pools to ensu med efficiently.	eds 15% re that		
	Cloud Disk (standard 1.00 TiB						
	2 MDIsks	⑦ Need Help		Cancel Gree	ite		
	Assign						

Figure 5-35 Create Pool

- 2. Complete the following step to add the Unmanaged Disk to the created pool (see Figure 5-38):
 - 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.

IBM S	pectrum Virtualize for Public Cloud	ibm_redbook_svpc_azure MDisks by Pools			and the second second		0.00	
		⑦ Create Pool						
~		Name	Cloud Disk ID		Cloud Disk Type	State		Usable Capacity
		 Unassigned MDisks (1) 						
		mdisk2				✓ Online		512.00 GiB
		V Pool0				🗸 Online		0 bytes / 1.
Ħ		mdisk1	sv-redbook-Mdis	sk-2	standardSSD	🗸 Online		512.00 GiB
		mdisk0	sv-redbook-Mdi	Assign I	MDisk mdisk2	to Pool	×	512.00 GiB
				Pool:	Pool0	•		
				Tier:	Enterprise Disk	•		
2				0		Cancel Assign		
		C Selecting 1 MDisk						
		Cloud Disk (standard 512.00 GB 21. MDak Assign						

Figure 5-36 Assign unmanaged disk to Pool

Complete the following steps to add the new Azure MDisk to IBM Spectrum Virtualize:

- 1. Open the Azure portal.
- 2. Select Create Resource and search for "Managed Disk", as shown in Figure 5-37.

Microsoft Azure		₽ Search resources	${\cal P}$ Search resources, services, and docs (G+/)				
ome > Resource groups > ibm_redbook_svpc_azure > Create a resource >							
Managed Disks 🔗 ·							
Managed Microsoft	Disks 🗢 Add to Favorites						
★ 5.0 (1 Azure ra	tings)						
Create							
Overview Plans Usage Infor	mation + Support Reviews						
Managed Disks is an abstraction of (urrent Standard and Premium storage	disk in Azure Storage. You only need	to specify the type (Standard or Prem	nium) and size of disk			
you need in your selected Azure regi	on, and Azure will create and manage	the Disk accordingly.	to specify the type (standard of fren				
More products from Microsoft				See All			
4	•	•					
Device Update for IoT Hub	Front Door Standard/Premium (Preview)	Azure VMware Solution	API App				
Microsoft	Microsoft	Microsoft	Microsoft				
Azure Service	Azure Service	Azure Service	Azure Service				
Securely and Reliably update your devices with Device Update for IoT Hub.	Azure Front Door Standard/Premium (Preview) is security led, modern cloud CDN	Azure VMware Solution (AVS) combines the VMware Software Defined Data Center (SDDC) with	scalable RESTful API with enterprise grade security, simple access control and auto SDK generation				
Create V	Create V 🛇	Create 🗸 🗢	Create V ♡				

Figure 5-37 MDisk Resource on Azure portal

3. Enter the basic information for the new MDisk, as shown in Figure 5-38.

Note: Use the correct resource group to add the MDisk.

Microsoft Azure	\mathcal{P} Search resources, services, and docs (G+/)			
Iome > Resource groups > ibm_redbook_svpc_zzure > Create a resource > Managed Disks >				
Create a managed d	isk …			
Basics Encryption Networki	ng Advanced Tags Review + create			
Select the disk type and size needed	I for your workload. Azure disks are designed for 99.999% availability. Azure managed			
disks encrypt your data at rest, by d	efault, using Storage Service Encryption. Learn more about disks.			
Project details				
Select the subscription to manage d manage all your resources.	aployed resources and costs. Use resource groups like folders to organize and			
Subscription * 🕕				
Resource group * 🕕	ibm_redbook_svpc_azure			
	Create new			
Disk details				
Disk name * 🔅	sv-redbook-Mdisk-3 🗸			
Region * ③	(US) East US V			
Availability zone	None			
Source type ①	None			
Size * 🕕	512 GIB			
	Standard SSD LRS			
	strangs ass			
Review + create <	Previous Next : Encryption >			

Figure 5-38 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 5-39. Click **Next**.

Microsoft Azure		$\mathcal P$ Search resources, services, and	docs (G+/)	N 67			
Home > Resource groups > lbm_redbook_supc_azure > Create a resource > Managed Disks > Create a managed disk >							
Select a disk size						×	
Browse available disk sizes and the	heir features.						
Disk SKU ①							
Standard SSD (locally-redundan	nt storage) 🗸 🗸						
Locally-redundant storage (dat	ta is replicated within a single	Provisioned IOPS	Provisioned throughput	Max Shares 🛈	Max burst IOPS 🛈	Max burst throughput 🛈	
datacenter)		500	60	3	600	150	
Premium SSD Best for production and perform	mance sensitive workloads	500	60	3	600	150	
Standard SSD		500	60	3	600	150	
Best for web servers, lightly use	ed enterprise applications and	500	60	3	600	150	
dev/test		500	60	3	600	150	
Standard HDD		500	60	3	600	150	
Best for backup, non-critical, an	id infrequent access	500	60	3	600	150	
512 GiB	E20	500	60	3	600	150	
1024 GiB	E30	500	60	5	1000	250	
2048 GiB	E40	500	60	5	-	-	
4096 GiB	E50	500	60	5	-		
8192 GiB	E60	2000	400	10	-	-	
16384 GiB	E70	4000	600	10	-		
32767 GiB	E80	6000	750	10			
Custom disk size (GiB) * ①							
lander the state of the state o							
ок							

Figure 5-39 Selecting size for Azure MDisk

5. Select the encryption policy that is provided by Azure Cloud Infrastructure, as shown in Figure 5-40. Click **Next**.

	P Search resources, services, and docs (G+/)					
Home > Resource groups > ibm_redboo	Home > Resource groups > libm_redbook_svpc_azure > Create a resource > Managed Disks >					
Create a managed disk						
Basics Encryption Networking	Advanced Tags Review + create					
Azure offers server-side encryption with pl to use a customer-managed key. Learn m	atform-managed keys by default for managed disks. You may optionally choose ore					
Encryption type *	(Default) Encryption at-rest with a platform-managed key					
Review + create < Previ	aus Next : Networking >					

Figure 5-40 Encryption setting for Azure MDisk

6. Select the networking access for the MDisk, as shown in Figure 5-41. Click Next.

Microsoft Azure	P Search resources, services, and docs (G+/)					
Home > Resource groups > ibm_redb	Home > Resource groups > lbm_redbook_svpc_azure > Create a resource > Managed Disks >					
Create a managed disk	Create a managed disk					
Basics Encryption Networking	Advanced Tags Review + create					
Network connectivity						
You can import or export your disk eithe access resource is created to manage th	r publicly or privately, using a private endpoint. To support disks at scale, a disk e private endpoints.					
Connectivity method	Public endpoint (all networks)					
	Private endpoint (through disk access)					
	O Deny all					
	All networks will be able to access this managed disk. Learn more about concerding methods of					
Review + create < Pre	vious Next : Advanced >					

Figure 5-41 Network Setting for Azure MDisk

7. Enable **Shared Disk** as Yes and select the **Max share** as 2, as shown in Figure 5-42. Click **Next**.

Home > Resource groups > ibm_red	Home > Resource groups > ibm_redbook_svpc_azure > Create a resource > Managed Disks >				
Create a managed dis	ж				
Basics Encryption Networking	Basics Encryption Networking Advanced Tags Review + create				
Add additional configurations for your	managed disk				
Shared disk					
Allow this disk to be attached to two o is enabled host caching is unavailable.	r more virtual machines, depending on storage type and disk size. When shared disk Learn more about shared disks				
Enable shared disk	9 Yes O No				
Max shares ④	2 ~				
Review + create < P	revious Next : Tags >				

Figure 5-42 Disk Share setting for Azure MDisk

8. Leave the tags blank and click **Next** to review and create the Azure MDisk, as shown in Figure 5-43.

		${\cal P}$ Search resources, services, and docs (G+/)				
Home > Resource groups > ibm_redbo	Home > Resource groups > ibm_redbook_svpc_azure > Create a resource > Managed Disks >					
Create a managed disk	Create a managed disk					
Basics Encryption Networking	Basics Encryption Networking Advanced Tags Review + create					
Tags are name/value pairs that enable yo	ou to categorize resources and view con	isolidated billing by applying the same tag				
to multiple resources and resource group	ps. Learn more about tags is inge resource settings on other tabs, yo	ur tags will be automatically updated.				
Name ①	Value 🕕	Resource				
	:	2 selected V				
Review + create < Pre	vious Next : Review + create >					

Figure 5-43 Tags setting for 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 5-44.

Microsoft Azure		𝒫 Search resources, services, and docs (G+/)
Home > Resource groups	> ibm_redbook_svpc_azure > Create a resource	> Managed Disks >
Create a manad	aed disk	
	,	
Validation passed		
- ·		
Basics Encountion	Networking Advanced Tans Review + c	reate
ouses encyption	Terrer ago Terrer e	
Basics		
Subscription		
Resource group	ibm_redbook_svpc_azure	
Region	East US	
Disk name	sv-redbook-Mdisk-3	
Availability zone	None	
Source type	None	
Size		
Size	512 GiB	
Storage type	Standard SSD LRS	
Encryption		
Encryption type	Platform-managed key	
Advanced		
Auvanceu		
Enable shared disk	Yes	
Max shares	2	
Networking		
Connectivity method	AllowAll	
Tags		
-		
Create	< Previous Next > Download a	template for automation

Figure 5-44 Azure MDisk validation and review

A successful deployment message is shown after the disk is provisioned, as shown in Figure 5-45.

		ces, and docs (G+/)			
Home > Microsoft.Managedl Deployment	Disk-20220105163016 Overview 🛷 …				×
Search (Ctrl+/) « Search (Ctrl+/) « Overview Inputs	Delete Cancel Redeploy Refresh We'd love your feedback! →				
S≡ Outputs interplate	Your deployment is complete Deployment name: Microsoft.Managedbisk-20220105163016 Subscription: Resource group: Ibmredbook.spc_azure Deployment details (Download)	Start time: 1/5/2022, 4-44:11 PM Correlation ID: d54477Ta-17bb-497-b37c-a669738140eb			
	Resource	Туре	Status	Operation details	
	Syntebook-Midis-3 Next steps Co to resource	Microsoft.Compute/disks	OK .	Operation details	

Figure 5-45 Successful deployment of Azure MDisk

The disk is now available in the Azure resource group, as shown in Figure 5-46.

	𝒫 Search resources, services, and docs (G+/)		E 5 0 8 -	-
Home >				
ibm_redbook_svpc_	azure 🖈 …			×
P Search (Ctrl+/) ≪	🕂 Create 📰 Edit columns 📋 Delete resource group 🖒 Refresh 🞍 Export to CSV 😤 Open query	${\textcircled{\sc 0}}$ Assign tags ${\sc \rightarrow}$ Move ${\sc \sim}$ ${\begin{tabular}{ll}}$ Delete ${\begin{tabular}{ll}} {\begin{tabular}{ll}} {\begin{tabular}{ll}} {\sc 0} \end{array}$	Export template 🛛 🔗 Feedback 🔋 Open in mobile	
() Overview	↑ Essentials			JSON View
Activity log	Subscription (Move) :	Deployments : 7 Succeeded		
Access control (IAM)	Subscription ID :	Location : East US		
🗳 Tags	Tags (Edit) : ClusterIPFailover : L8_Based_Cluster_IP_Deployment IPFailoverEnabled : False SVPC_De	fault_MTU : 1500 deploymentTag : sv-redbook	Cluster_Deployment_Status : DEPLOYMENT_STATUS_SUCCESS	✓ More (5)
👃 Resource visualizer				
🗲 Events	Resources Recommendations (12)			
Settings	Filter for any field			
1 Deployments			· · · · · · · · · · · · · · · · · · ·	
Security	Showing 1 to 24 of 24 records. Show hidden types U		Group by type	iew 🗸
Policies	Name 🛧	Type ↑↓	Location $\uparrow \downarrow$	
🐲 Properties	> Availability set			^
🔒 Locks	✓ Disk			
Cost Management	Sv-redbook-Mdisk-1	Disk	East US	
S. Cost analysis	Sv-redbook-Mdisk-2	Disk	East US	
Cost alerts (preview)	C Sv-redbook-Mdisk-3	Disk	East US	
Budgets	Sv-redbook-node1-vm_OsDisk_1_f0a340a6648a41dc85e288f3b0a081ad	Disk	East US	
Advisor recommendations	sv-redbook-node2-vm_OsDisk_1_b68fb2d8f8ab420c913a6db49be65d28	Disk	East US	
Manifester	Sv-redbook-quorum_OsDisk_1_0cd39c5c3be3402cb68b0e859b60a5a3	Disk	East US	
Monitoring	> Key vault			
 Insights (preview) 	> Load balancer			
Alerts	Managed Identity			
NI WETICS	> Network interface			×
Diagnostic settings	< Previous Page 1 V of 1 Next >			
Logs				

Figure 5-46 Azure MDisk in Azure resource group

10.In the IBM Spectrum Virtualize GUI, select Pools \rightarrow MdiskByPools \rightarrow Discover Storage, as shown in Figure 5-47.

IBM S	pectrum Virtualize for Public Cloud	ibm_redbook_svpc_azur	e MDisks by Pools					💊 🛃 Ø	superuser Security Administrator	I ~
	X	① Create Pool	Actions					Default \checkmark Con	tains ~ Filter	~ 74
	Dashboard	Name	View All Throttles	Cloud Disk ID	Cloud Disk Type	State	Usable Capacity	Written Capa	acity Limit	
		Unassigned MD	Customize Columns 🔸							
<u>م</u> رم		mdisk2				 Online 	512.00 GiB	512.00 GiB		-
		> Pool0				✓ Online	0 bytes / 1.00 TiB (0%)		1,016.00 GiB / 1.00 TiB (99%)	
	Pools									
00	Volumes									
Đ										

Figure 5-47 Discover Storage

11. The newly created disk is available under UnManagedDisk, as shown in Figure 5-48.

IBM Spectrum Virtualize for Public Cloud	ibm_redbook_svpc_azure MDisks by Pools					💫 🔒 🕐 superuser Security Administrator 🗸
<	⊕ Create Pool 🛛 🗎 Actions 👻					Default V Contains V Filter V
A Dashboard	Name	Cloud Disk ID	Cloud Disk Type	State	Usable Capacity	Written Capacity Limit
	Unassigned MDisks (1)					
ം സംസ്ത്ര Monitoring	mdisk2			✓ Online	512.00 GiB	512.00 GiB
	V Pool0			✓ Online	0 bytes / 1.00 TiB (0%)	1,016.00 GiB / 1.00 TiB (99%)
Pools	mdisk1	sv-redbook-Mdisk-2	standardSSD	🗸 Online	512.00 GiB	512.00 GIB
Г	mdisk0	sv-redbook-Mdisk-1	standardSSD	✓ Online	512.00 GiB	512.00 GiB
Volumes						
▤						
>>> Hosts						
Copy Services						
Access						
So Settings	<					3
~~~ ·····•	Selecting 1 MDisk					
	Cloud Disk (standard 512 00 GB S 5 MDak Assign					
		Latency 0 ms	Read Write	Bandwidth O MBos O N	Mar OMBor IOPS 0 0 0	

Figure 5-48 Newly created Azure MDisk seen as Unmanaged Disk

We can add the disk to the existing pool by following the steps as described in 5.4, "Configuring the back-end storage" on page 90.

Now, you can create the VDisk and assign the volume for host access by using ISCSI.

#### 5.4.1 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).

The default for thin-provisioned volumes that is indicated by the command-line interface as shown in Example 5-1, 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 5-1 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
```

				<b>10</b> Log U
	⑦ Create Volumes	Create Volumes	×	Default V Contains V Filter
Dashboard	Name	Basic Mirrored Custom		Capacity 📳
		Create a preset volume with all the basic features.	^	
		Pool: Capacity Details: Pool0  Total 1.00 TiB		
Volumes		Volume Details		
Hosts		Quantity: Capacity: Name:		
		1     -     30     GiB     •       Canacity sayings:     sypc-azure-thin-vol1		
		Thin-provisioned		
		(e) Define another volume		
		I/O group: Automatic ·		
			~	
	Showing 0 volumes   Selecting 0 volumes	Need Help     Cancel     Create and Hap     Create		
		Latency O ms 0 ms Bandwidth O MBps 0 MBps		

Figure 5-49 shows thinly provisioned (space-efficient) volume creation by using the GUI.

Figure 5-49 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.

#### 5.4.2 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 identifier (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** -11) to ensure that your VDisks are attached correctly through the multipath tool.

Typical output of a VDisk should resemble the output that is shown in Example 5-2.

Example 5-2 Multipath output

```
multipath -11
mpathah (36005076072a06dc4f0000000000000c) 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 (36005076072a06dc4f00000000000000) 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.

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

Connect to the discovered targets, as described at this IBM Documentation web page.

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.

## 5.5 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 the 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.

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.

#### 5.5.1 Azure configuration for VPNGW IPsec tunnel

This section describes the steps that are required at the virtual network level in Azure Cloud for establishing the IPsec tunnel. For more information, see this hAzure tutorial web page.

Complete the following steps:

- 1. Create a VPN gateway:
  - 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 virtual network 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.

## 5.6 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 a FlashSystem or IBM SAN Volume Controller system to an IBM Spectrum Virtualize for Public Cloud on Azure solution.

Our example uses a FlashSystem system in the on-premises data center and a two-node IBM Spectrum Virtualize for Public Cloud cluster on Azure as a DR storage solution.

This scenario uses IBM Spectrum Virtualize Global Mirror with Change Volume (GMCV) to replicate the data from the on-premises data center to Azure Cloud.

This implementation starts with the assumption that the IP connectivity between the on-premises data center and Azure Cloud is established through a Multiprotocol Label Switching (MPLS) or VPN connection.

**Note:** Because many methods are available to implement the IP connectivity, this section does *not* consider that specific configuration. For more information, contact your organization's network technical specialist.

To configure the GMCV, complete the following steps:

1. Configure your IBM Spectrum Virtualize Private IP ports so that they are enabled for remote copy. This configuration is required on both sites, as shown in Figure 5-50.

ibm_redbook_svpc_azure - Netv ×	GDLAZUDEV - Network × +											- o	x a
$\leftarrow \rightarrow C$	A https://10.21.0.6:8443/gui#config-net	twork-ethernet									☆	ſ	⊚ ≡
IBM Spectrum Virtualize for Public Cloud	ibm_redbook_svpc_azure								4	I () •	uperuser Secur	ity Administrator	•
	K Management IP Addresses	Ethernet Ports The Ethernet ports can be used	d for iSCSI conn	ections, host attachm	ent, and remote copy.								
Dashboard	Service IPs	$\equiv$ Actions $\checkmark$						[	Default	Contains	Filter		~ ~
Monitoring	Ethernet Connectivity	Name	Port 🛧	Link State	Speed	Host Attach	Storage	Replication					
<u> </u>	,	√io_grp0											
B Pools	Ethernet Ports	node2	1	Active	40Gb/s	Yes	Yes	Yes	1				
B		node1	1	Active	40Gb/s	Yes	Manage IP Addresses						
Volumes	iSCSI	node2	2	Active	40Gb/s	Yes	Modify Remote Copy						
Hosts	DNS	node1	2	Active	40Gb/s	Yes	Modify Storage Ports Modify Maximum Trans	mission Unit					
Copy Services	Internal Proxy Server												
Access	Portsets												
Settings													
				Read Write		Read Write		Read Write					
			Latency O m	s 0 ms 0 ms	Bandwidth 0 M	3ps 0 MBps 0 M	Bps IOPS O	0 0					

Figure 5-50 Enable Remote Copy

2. Enable IP Portset for remote copy With the new feature of multiple portset and Multiple IP, it is essential to enable the IP for portset. This enablement can be done only by using the CLI for IP address. as shown in Example 5-3. This process must be done at both clusters to be joined in IP partnership.

Note: The share IP option is not available on the IBM SAN Volume Controller GUI.

Example 5-3 Using same IP for 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. Create a cluster partnership between the on-premises data center and IBM Spectrum Virtualize for Public Cloud on Azure from the on-premises GUI, as shown in Figure 5-51.

IBM Spectrum Virtualize for Public Cloud	ibm_redbook_svpc_azure Remote Copy			🎭 🗉 🕐	superuser Security Administrator 🗸 🗸
<u> </u>	Partnerships	<pre> øibm_redbook_svpc_azure </pre>			Actions 👻
Dashboard	<pre> @ ibm_redbook_svpc_azure  ! </pre>	Consistancy Groups (0)			
പ്പ ⁹ Monitoring		Add Cor	×		
Pools		Select the type of partnership to create. If you have the other systems at each site.	ven't already, you need to create this partnership on		n q
Volumes		0			
Hosts					
Copy Services		2-site partnership	3-site partnership		
Access		2 sites.	susually located farther from the other sites, for disaster recovery purposes.		
Settings				Jependent	
		Cancel	Continue		
	Create Partnership +				
		Latency Oms Oms Oms	Bandwidth O MBps O MBps O MBps	IOPS 0 0 0	

Figure 5-51 Create Partnership from SVPC GUI to hybrid setup

4. Complete the IP Partnership from the on-premises setup, as shown in Figure 5-52.

IBM FlashSystem 7200	GDLAZUDEV	Remote Copy						ە 🛃 🍓	superuser Security Administrator 🗸 🗸
		Partnerships		Anuja	a T4test resgrp				Actions -
Dashboard		© GDLAZUDEV	:	Not Prese	(; ; ;		×		
୍ୱର Monitoring		Anula Tétant rondra		Consist	Create Partnership The system is in the replication layer, Partnership				
		Not Present			replication layer.				
B Pools		S Error Detected		Add Cor	Type Fibre Channel  IP				ti Q
Volumes				New	Partner IP Address				
Hosts					10.21.0.6		RELATIONSHIPS		
Conv Semicor				8	Link Bandwidth (Mbps)	Background Copy Rate (%)	0		
				8.	2048	50	RELATIONSHIPS 0		:
Access					Partner CHAP Secret	Compression Enabled			
Settings				8.	Enter Value Ø	Off Off	RELATIONSHIPS 0		:
					Portset Link 1	Portset Link 2 (Ontional)	RELATIONSHIPS		
				8.	Cancel	Create	0		:
				8 -			RELATIONSHIPS		1
		Create Partnership	+		Linky				~
			Î		Read Write Latency Oms Oms Oms	Read Write Bandwidth OmBps omBps omBps	IOPS 0 Read Write		

Figure 5-52 Complete Partnership from an on-premises setup

 After the partnership is configured on both clusters, you can see the configured partnership on the clusters (see Figure 5-53 for Azure Cloud GUI) and on-premises IBM SAN Volume Controller GUI (see Figure 5-53 on page 103).



Figure 5-53 Cloud Partnership GUI



Figure 5-54 On-premises IBM SAN Volume Controller GUI partnership

Now, you have a fully configured partnership. In our example, we are replicating 1 GB volume from on-premises data center to Azure Cloud. Figure 5-55 shows the volumes that were used for this example.

IBM FlashSystem 7200 GDLAZUDEV	Volumes by Pool		👍 🛃 🕐 superuser Security Administrator 🗸 🗸
Dashboard	C Filter Pools Pool Pool Pool Pool Pool Pool Pool	testpool          IMDisk, 2552 Volume copies         Easy Tier: Balanced         Creats Volumes         East Solumes	Pool Actions ↓ 18% Stored 5,73 Till (0,2%) Available 6,98 TiB Total Usable Default ↓ Contains ↓ Filter
Pools	Data Reduction	Name State Synchronized Protocol Type	UID Host Mappings Capacity
B Volumes	testpool 🗸	svpc-azure-redbook-cv 🗸 Online (formatting)	60050768108183F87000000000000 No 1.00 GiB
	2552 Volume copies	svpc-azure-redbook-vol 🗸 Online	60050768108183F8700000000000000000000000000000000000
E Hosts	1.25 TiB (18%) Stored	vola Online	60050768108183F8700000000004 No 20.00 MiB
	0.50 110 10121 052010	vol2 🗸 Online	60050768108183F8700000000004 No 20.00 MiB
		vol3 🗸 Online	60050768108183F8700000000004 No 20.00 MiB
		vol4 🗸 Online	60050768108183F8700000000004 No 20.00 MiB
Q Access		vol5 🗸 Online	60050768108183F8700000000004 No 20.00 MiB
		vol6 🗸 Online	60050768108183F8700000000004 No 20.00 MiB
Settings		vol7 🗸 Online	60050768108183F8700000000004 No 20.00 MiB
~ -		vol8 🗸 Online	60050768108183F8700000000004 No 20.00 MiB
		vol9 🗸 Online	60050768108183F8700000000004 No 20.00 MiB
		vol10 🗸 Online	60050768108183F8700000000004 No 20.00 MiB
		vol11 🗸 Online	60050768108183F8700000000004 No 20.00 MiB
	Create Pool	Showing 2552 volumes   Selecting 0 volumes	*
		Latency O ms O ms Bandwidth O MBps O MBps	IOPS 0 Read Write

Figure 5-55 Volumes for Remote Copy

6. Create a volume remote copy relationship for a GMCV from the on-premises data center, as shown in Figure 5-56.

IBM FlashSystem 7200 GDLAZUDEV	Remote Copy		👍 🛃 🕐 superuser Security Administrator 🗸
ŀ	Partnerships	ibm_redbook_svpc_azure	Actions 👻
Dashboard	⊗ GDLAZUDEV :	Configured	
orden Monitoring ∰ Pools	Anuja_T4test_resgrp : Not Present S Error Detected	Consist © Create Relationship × Select the type of relationships that you want to create. More Information	ent Relationships (0)           Default         Contains         filter           Auxiliary Volume         11
Volumes	ibm_redbook_svpc_azure : • Configured	Metro Mirror     Global Mirror     Global Mirror     Global Mirror with Change Volumes	
Copy Services		Migration O Non-disruptive System Migration	
Access			
		Cancel 4 Back Next >	1

Figure 5-56 Create Relationship for GMCV from on-premises data center

7. Select the remote system, as shown in Figure 5-57.

IBM FlashSystem 7200	Remote Copy						I 👆 🐁 🕐	superuser Security A	dministrator
	Partnerships	ibm_redb	ook_svpc_azure						Actions 🔹
Dashboard	◎ GDLAZUDEV	Configured							
୍ଦ୍ର Monitoring	Anuja_T4test_resgrp Not Present	Consiste ⊕ Create	Relationship		×	ent Relationships (0)	Default V	Contains 🗸 Filter	
Pools	S Error Detected	Name O	re the auxiliary volumes located? On this system			Auxiliary Volume			
Volumes	ibm_redbook_svpc_azure Configured	•	On another system ibm_redbook_svpc_azure						
Hosts									
Copy Services									
Access									
Settings									
		Cancel		≪ Back M	Next ►				
	Create Partnership	Showing 0 relationship:	;   Selected <b>0</b> relationships						
			Read Write Latency Oms Oms Oms	Bandwidth OMBps OM	d Write IBps OMBps	IOPS 0 0 0			

Figure 5-57 Remote system to be used for remote copy relationship

8. Select the volumes that must be in the relationship, as shown in Figure 5-58.

IBM FlashSyster	n 7200 GDLAZUDEV	Remote Copy		👍 🔒 🧿 superuser Security Administrator 🗸
		Partnerships	ibm_redbook_svpc_azure	Actions 👻
Dashbo		@ GDLAZUDEV :	Configured	
oي Monitol		Anuja_T4test_resgrp : Not Present  Error Detected	Consists     Create Relationships (0)       © Create Select the master and auxiliary volumes to use in the relationship.     Auxiliary Volume       Name     Image: Construction of the select the master and auxiliary volumes to use in the relationship.	Default V Contains V Filter
Volume		ibm_redbook_svpc_azure : Configured	Master     Auxiliary       supe-azure-rehin-rot     *       Capacity: 1.00 GiB     Add	
Copy Si				
Access				
			Cancel 4 Back Next >	
		Create Partnership +	Showing O relationships / Selected O relationships	
			Latency O ms oms Bandwidth O MBps o MBps IOPS O 8 0 0	

Figure 5-58 Volumes in remote copy relationship

9. In our example, we select **No, do not add a master change volume**, as shown in Figure 5-59 on page 106. They are added later.

IBM FlashSystem 7200 GDLAZUDEV	Remote Copy		👍 📳 🕐 superuser Security Administrator 🗸
C Dashboard	Partnerships	ibm_redbook_svpc_azure • configured	Actions 👻
Monitoring	Anuja_T4test_resgrp : Not Present	Consist Consist Consist Consist Add Change Volumes Do you want to add a master Global Mirror change	ps (0)           Default         V
Volumes	SError Detected	Name volume? Auxelian O Yes, add a master change volume. © No, do not add a master change volume.	vVolume 🗓
Hosts			
Access			
∑́⊙} Settings		Cancel < Back Finish	
	Create Partnership +	Showing 0 relationships   Selected 0 relationships	
		Latency 0 ms 0 ms 0 ms Bandwidth 0 MBps 0 MBps 0 MBps IDPS 0 0	

Figure 5-59 Add Volume without master change volume

IBM FlashSystem 7200 GDLAZUDEV	Remote Copy		a 🖡 🕐 superuser Security Administrator 🗸
	Partnerships	ibm_redbook_svpc_azure	Actions
Dashboard		Configured	
Monitoring Pools Volumes	Anuja_T4test_resgrp : Not Present : Error Detected :	Create Relationship × Create Relationship × Do you want to start copying now? Name Ves. start copying now. © Ves. start copying.	ent Relationships (0)           Default         Contains >>         //dor           Auxiliary Volume         III
	Configured		
Hosts			
Copy Services			
		Cancel 4 Back Pinish	
	Create Partnership +	Showing 0 relationships / Selected 0 relationships	
			IOPS 0 Read Write

#### 10.Select No, do not start copying, as shown in Figure 5-66.

Figure 5-60 Do not start copy for volumes in relationship

- 11. Edit the relationship and set the Cycling Mode and Cycling Period.
- 12.Add the CV volumes to your relationship on both sides, as shown in Figure 5-61 and Figure 5-62 on page 108.

IBM FlashSystem 7200 GDLAZUDEV	Remote Copy					👍 🗟	0	superuser Security	Administrator 🗸 🗸
C Dashboard	Partnerships            Ø GDLAZUDEV	ibm_redbook_svpc_azure • configured							Actions 🔹
مرم ² Monitoring	•1	Consistency Groups (0)		Independent	Relationships (1)				• 1
Pools	Anuja_T4test_resgrp : Not Present	③ Create Relationship IIII Actions ▼ ▲       Name     ↑     State	Master Volume	Replication Direction	Auxiliary Volume	Default	Conte	sins 🗸 Filter	
Volumes	S Error Detected	rcrel0 Inconsistent Stopped	svoc-azure-redbook-vol Rename	•>	spvc-azure-thin-vol				
Hosts	ibm_redbook_svpc_azure : Configured		Add to Consistency Group Convert to 3-site	Create New					
Copy Services	• 1		Change Volumes	Add Existing					
Access			Stop Switch	Properties (Master)					
Settings			Delete	Toperties (Auxiliary)					
			Edit Relationship	J					
	Create Partnership +	Showing <b>1</b> relationship   Selected <b>1</b> relationship							
		Latency <b>O</b> ms Oms Oms	Bandwidth <b>O</b> MBP	Read Write S OMBps OMBps IC	OPS 0 0 0				

Figure 5-61 Adding change volume to on-premises setup

IBM Spectrum Virtualize for Public Cloud	ibm_redbook_svpc_azure Remote Cop			🎭 🗐 🕐 superuser Security Administrator 🗸
<	Partnerships	øibm_redbook_svpc_azure		Actions 👻
Dashboard	⊗ ibm_redbook_svpc_azure :			
్లా ^D Monitoring		Consistency Groups (0) ⊕ Create Relationship ≡ Actions → 1 ³ / ₂	Independent Relationships (1)	
Pools	GDLAZUDEV : Configured	Name 🛧 State Master Volume Re	plication Direction Auxiliary Volume	
Volumes		Add Existing Change Volume	spvc-azure-thin-vol	
Hosts		Select the volume on the local system.		
Copy Services				
Access		Cancel Add		
کې Settings				
	Create Partnership +	Showing 1 relationship   Selected 1 relationship		
		Latency Oms Oms Oms Bandwidth OMBps OM	d Write Bps 0 MBps IOPS 0 0	

Figure 5-62 Add Change Volume on Azure Cloud

#### 13. Start your relationship from the on-premises site, as shown in Figure 5-63.

ibm_redbook_svpc_azure - Rem × GD	LAZUDEV - Remote Copy × +						– ø ×
$\leftarrow \rightarrow \circ \circ$	https://192.168.2.19/gui#copyservices-	remotecopy				☆	
IBM FlashSystem 7200 GDLAZUDEV	Remote Copy					👍 🔒 🧷 superuser (	Security Administrator
C Dashboard	Partnerships       Ø GDLAZUDEV	ibm_redbook_svpc_a • Configured	zure				Actions 🔹
a _o ₽ Monitoring	• 1	Consistency Groups (0)		Independer	nt Relationships (1)		1
Pools	Anuja_T4test_resgrp : Not Present		业. Master Volume	Replication Direction	Auxiliary Volume	Default $\checkmark$ Contains $\checkmark$	Filter
Volumes	Serror Detected	rcrel0 Inconsistent Stopped	Rename	•	spvc-azure-thin-vol		
Hosts	ibm_redbook_svpc_azure : Configured		Add to Consistency Group Convert to 3-site				
Copy Services	• 1		Change Volumes >				
Access			Stop Switch				
Settings			Delete Edit Relationship				
	Create Partnership +	Showing <b>1</b> relationship   Selected <b>1</b> relationship	and Write	Sast Write	Dand Write	_	
		Latency O ms 0	ms 0 ms Bandwidth Ом	Bps 0 MBps 0 MBps	IOPS 0 0 0		

Figure 5-63 Start relationship from on-premises setup

14. Create a GM consistency group and add your relationship to it, as shown in Figure 5-64 and Figure 5-65.

ibm_redbook_svpc_azure - Rem ×	DLAZUDEV - Remote Copy × +		_	- a ×
$\leftarrow \rightarrow $ C O	A https://192.168.2.19/gui#copyservices	remotecopy	☆	⊚ ≡
IBM FlashSystem 7200 GDLAZUDEV	Remote Copy	👍 🖡 🕐	superuser Security Admin	vistrator 🗸 🗸
	Partnerships	ibm_redbook_svpc_azure • continued	Ac	ctions 👻

Figure 5-64 Add relationship to consistency group

IBM F	lashSystem 7200 GDLAZUDEV	Remote Copy	👍 📳 🕐 superuser (Security Administrator) 🗠
		Partnerships	ibm_redbook_svpc_azure
			Configured
		Anuja_T4test_resgrp :	Consistency Groups (1) Independent Relationships (1)
æ		Not Present	🐵 Create Relationship 🗏 Actions 👻 🕹 👘 Err
		Error Detected	Name 🕈 State Master Volume Replication Direction Auxiliary Volume
		ibm_redbook_svpc_azure : Configured	Add Relationship to Consistency Group ×
<b>∑</b>			Select the consistency group to move the relationship rcret0 Consistency Group Tredbook_cg
H			
2			Cancel Add to Consistency Group
		Create Partnership +	Showing & relationship / Selected & relationship
			Jatency Ome new Bandwith Ower news 10PS O n n

Figure 5-65 Select the consistency group to add relationship

#### 👍 🗐 🕜 super ibm_redbook_svpc_azure Partnerships redbook_cg ← Back to GDLAZUDE ø ibm_redbook_svpc_azure Group info redbook_cg GDLAZUDEV of Global ication type GDLAZUDEV Primary site Cycle period ( 300 s 1/11/2022 5:39:04 PM Freeze time 🛈 Group ID Relationships (1) ≔ Actions -Default ✓ Contains ✓ ↑ State Master Volume Replication Direction Auxiliary Volume 💽 Set Create Partnership Showing 1 relationship | Selected 1 re

You can now see the status of your consistency group, as shown in Figure 5-66.

Figure 5-66 Status of the consistency group

In our example, we show the consistency group status from SVPC Azure Web GUI.

Our example is only one way to configure a GMCV relationship from an on-premises solution to an IBM Spectrum Virtualize for Public Cloud on Azure solution. It can be valuable to configure a snapshot (FlashCopy) of your GMCV auxiliary volume to be used for Disaster Recovery testing or other purposes.

For more information about how to manage FlashSystem, IBM Spectrum Virtualize, or IBM SAN Volume Controller copy functions, see the following publications:

- Implementing the IBM FlashSystem with IBM Spectrum Virtualize Version 8.4.2, SG24-8506
- Implementing the IBM SAN Volume Controller with IBM Spectrum Virtualize Version 8.4.2, SG24-8507
- IBM SAN Volume Controller Best Practices and Performance Guidelines for IBM Spectrum Virtualize Version 8.4.2, SG24-8509
- IBM FlashSystem Best Practices and Performance Guidelines for IBM Spectrum Virtualize Version 8.4.2, SG24-8508

## 6

# Monitoring and supporting the solution

This chapter provides guidance about supporting and monitoring this solution.

This chapter includes the following topics:

- 6.1, "Monitoring IBM Spectrum Virtualize for Public Cloud on Microsoft Azure through GUI, Spectrum Control, or Storage Insights" on page 112
- 6.2, "Troubleshooting Spectrum Virtualize for Public Cloud on Microsoft Azure" on page 137
- 6.3, "IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Support contact" on page 154

#### 6.1 Monitoring IBM Spectrum Virtualize for Public Cloud on Microsoft Azure through GUI, Spectrum Control, or Storage Insights

This section describes the following procedures:

- ► Setting up Call Home and email notification.
- ► Viewing capacity and performance in IBM Spectrum Virtualize GUI and Microsoft Azure.
- Monitoring IBM Spectrum Virtualize for Public Cloud (SV4PC) on Azure in IBM Spectrum Control and IBM Storage Insights.

#### 6.1.1 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 6-1 lists the supported configurations for Call Home.

Supported configuration	DNS configuration	Firewall requirements
Call Home with Cloud services with an internal proxy server	Required	Configure 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 firewall to allow outbound traffic on port 443 to essuport.ibm. Optionally, allow outbound traffic on port 443 to the following IP addresses: 129.42.56.189 129.42.60.189
Call Home with Cloud services	None	Configure 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

Table 6-1 Supported network configurations for Call Home with Cloud services

**Note:** Call Home with Cloud services is the optimal transmission type for Spectrum Virtualize for Public Cloud on Microsoft Azure.

Complete the following steps to configure Call Home (the screen captures show all possible configuration options):

1. From the left window of the GUI, select **Settings**  $\rightarrow$  **Support**  $\rightarrow$  **Call Home** see Figure 6-1.

Call Home	Call Home The support user receives call home events, Local users also receive event notifications.
Support Assistance	Enable Netfleations
Support Package	

Figure 6-1 Call Home

2. Select Enable Notifications to start the Call Home configuration wizard (see Figure 6-2).



Figure 6-2 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 is 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 6-3).

Call Home			×
Welcome Transmission Settings Transmission Settings	Select Transmission Types Select the transmission types to use for Call Hot policy info.	ne on your system. Click <u>here</u> for IBM priva	су
Internal Proxy Server Proxy Details System Location Contact Software Entitlement Inventory Settings Email Servers	Send using Cloud Services $\overleftrightarrow$ Select this option to send notifications using RESTFUL APIS, which are standard for transmitting data through web services. You can also specify an internal proxy server to manage outbound connections with the support center.	Send using Email Services Select this option to send notifications through an email server to support to fix errors or to international users or services that monitor activity on the system.	h
Summary	Cancel	■ Back Apply	and Next ►

Figure 6-3 Select 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.

 Call Home
 ×

 Welcome
 Internal Proxy Server

 Transmission Settings
 Call home with cloud services can be used with an internal proxy server to send call home messages to support. A DNS server is required to use a proxy server for call home with cloud services. More Information

 System Location
 Use proxy to send data.

 Contact
 Yes

 Summary
 No

A DNS server is required to use a proxy server for Call Home with Cloud services (see Figure 6-4).

Figure 6-4 Configure Proxy server

5. After clicking **Next** in the welcome window, enter the information about the location of the system, as shown in Figure 6-5.

Call Home				×
Welcome	Connection to the	support center was succe	ssful!	×
Internal Proxy Server     System Location	System Location Service parts should be shipp	ped to the same physical loca	ation as the system.	
Contact	Company name:	IBM ITSO		
Software Entitlement	System address:	Redbooksteet		
Inventory Settings Summary	City:	Redbook		
	State or province:			
	Postal code:	65936		
	Country or region:	Germany	•	
	Machine location:	datacenter		
		I		
	Cancel		◀ Back	Next 🕨

Figure 6-5 Location of the device

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

Call Home			×		
✓ Welcome	Contact				
Transmission Settings	The support center contacts	this person to resolve issues on the system.			
Internal Proxy Server     System Location	Enter business-to-business contact information     To comply with privacy regulations, personal contact information for individuals with your organization is     not recommended.				
Contact					
Software Entitlement	Name:	Name of Contact			
Inventory Settings	Email:	contact_email@Address			
Summary	Phone (primary):	0049123456789			
	Phone (alternate):				
	IBM may use my contact da	ta to keep me informed of Storage related products, services and offerings.			

Figure 6-6 Contact information

6. In the next window of the wizard, enter the software entitlements details (see Figure 6-7).

Call Home				×
Welcome     Section     Transmission Settings     Internal Proxy Server	Software Entitl Enter your IBM Spect on Passport Advanta	ement rum Virtualize software entitleme ge.	nt information that you used	
System Location Contact Software Entitlement Inventory Settings	IBM customer ID: IBM country code:	0123456 724		
Summary	IBM Remote Technica Contact the hardward Call home regularly se information. Object na	al Support responds only to Call Ho e manufacturer for support with ha nds emails to the support center that mes and other potentially sensitive i	me emails for software errors. Irdware errors. I describes your system hardware and nformation, such as IP addresses, are	critical configuration e not sent.
	Cancel		<b>⊲</b> Back	Apply and Next 🕨

Figure 6-7 Software entitlement

 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 0n to generate enhanced reports.

The Support center uses this detailed information to automatically generate recommendations and best practices that are specific to your configuration.

You also can configure 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 6-8).

Call Home		×
Welcome Transmission Settings Internal Proxy Server System Location	Additional Settings Inventory Reporting Emails sent to the service center will contain information about inventory. Off on Email Interval:	
Contact Software Entitlement Inventory Settings Email Servers Summary	Configuration Reporting Call Home now sends enhanced reports that include information about your system configuration to the support center. The support center uses this detailed information to automatically generate recommendations and best practices that are specific to your configuration. Off  On	
	<ul> <li>Sensitive information         The inventory email includes configuration fields like object names and IP addresses. It is recommended that these fields are not used to store sensitive information. However, if that is not possible, selecting this option removes object names, IP addresses, and other information from the inventory email.     </li> <li>When you enable this option, any automated health checking provides object identifiers only, rather than names, which can be less helpful for error resolution.</li> <li>Remove content that identifies system objects</li> </ul>	
	Cancel  Apply and Next	

Figure 6-8 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 6-9 on page 119).

**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 Spectrum Virtualize for Public Cloud configuration, you might need to set up an SMTP service in Microsoft Azure.

Call Home			×
Welcome	Email Servers	re routed through this email server.	
System Location	Server IP or Domain	Port:	
Contact		25 🕀 🖯	
Software Entitlement	Ping		
<ul> <li>Inventory Settings</li> </ul>			
🔆 Email Servers			
Summary			

Figure 6-9 Configure email servers

9. The last window of the wizard shows a summary of the configuration. Verify the information and then, click **Finish** (see Figure 6-10).

Call Home					×
💽 Welcome	IP address: Name:	1.2.3.4 dnsserver0			^
<ul> <li>Transmission Settings</li> <li>Internal Proxy Server</li> <li>System Location</li> </ul>	Contact Contact name: Email address: Telephone (primary): Telephone (alternate):	Name of Contact 0049123456789			
Contact	System Location Company name: Street address:	IBM ITSO Redbooksteet			
<ul> <li>Inventory Settings</li> <li>Email Servers</li> </ul>	City: State or province: Postal code: Country or region: Machine location:	Redbook RLP 65936 Germany datacenter			
   Summary	Email Servers Server IP 1.2.3.4	Port 25			
	Call Home Transmission setting: Support center: Alerts: Inventory Reporting: Email Interval: Configuration Reporting: Remove Sensitive Information:	Cloud, Email callhome1@de.ibm.com Errors, Inventory On Every day On Off			-
	Cancel		<b>⊲</b> Back	Finish	

Figure 6-10 Verify Summary

10. To test Call Home, left window of the GUI, select Settings  $\rightarrow$  Call Home  $\rightarrow$  Test Support Notification (see Figure 6-11).

1011.01					
	<	Call Home	Call Home The support user receives call hor	ne events. Local users also rec	eive event notifications.
	Dashboard	Support Assistance	Edit Disable Notifications		
Ŷ	Monitoring	Support Package	Send using Cloud Services Send using Email Services		
₿	Pools		Call Home with cloud services		
	Volumes		Connection: ✓ Active Last Connection: Success at Proxy: Not config	Tet 12/20/2021 2:54:37 PM ured Ad	st Support Notification
	Hosts		Email Contact		
[ <del>]</del>	Copy Services		* Contact Name Name of Contact	* Email Reply Address	
0	Access		* Telephone (Primary) 0049123456789	Telephone (Alternate)	
ক্ট	Settings		* Required		
			System Location		
			* Company Name	* Street Address	
			IBM ITSO	Redbooksteet	
			* City	* State or Province	* Postal Code
			Redbook	RLP	65936
			* Machine Location	* Country or Region	

Figure 6-11 Test Support notification

#### **Disabling and enabling notifications**

At any time, you can temporarily or permanently disable notifications, as shown in Figure 6-12.

Disa	ble Notifications		×
•	<b>Warning</b> Disabling call home might delay resolution of technical issue home to avoid delays. Do you want to continue?	s by the support center. Enable call	
No		Yes	

Figure 6-12 Disable notifications

### Monitoring capacity reporting in IBM Spectrum Virtualize for Public Cloud on Microsoft Azure

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 6-13).

Performance					Node Compari	son	System	
Catency O ms	o ma	write O ma	100-					- read - set
Bandwidth O MBps	read O MBpe	write O MBps	60 +					
IOPS O IOPS	read 0 10PS	write 0 sops	40 - 20 -					
CPU Utilization 1%	The a	mount of capaci	ty that is available for storing data on or MDisk after formatting and RAID	ž=in	T Imn	2min	1 Sein	Dirit Dirit
Capacity	techr Capa	iques are applie city").	d. (Previously known as "Physical	The amount of capacity sav capabilities on the system.	ed using capacity savings	The total capa "Volume Capa	city of all volumes in the s city").	ystem (Previously:
Usable Capac	ity 🕐			Provisioned Capacity ⑦		Capacity Savings	0	
2%			98%	>100%	0%	O%	0%	O _%
24 GiB Nored Capacity			1,000 GiB Available Capacity	5.79 GiB Written Capacity	O bytes Available Capacity	0 bytes Compression	O bytes Deduplication	O bytes Thin Provisionin
MDisks			Total 1 TiB	Compressed Volumes	Total Provisioned 2 GiB <u>Peduplicat</u>	Compression Ratio N/A ed Volumes		Total Savings Ö bytes Thin-Provisioned Volumes
System Healt	h					Version: 8.4.3.0 (b	uild 155.0.2112040719000)   (	Cluster ID: 0000001CBA34F260
& Hardware	e Compone	ents		🕱 Logical Components		Connectivity	Components	
All Online	$\mathbf{D}$			All Online		All Online		
						Concernment of		

Figure 6-13 Dashboard capacity view

#### 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 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 displays 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 displays 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.

#### **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 displays 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 displays 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 displays the percentage for over-provisioned volumes. The Overprovisioned value indicates the percentage of provisioned capacity that is increased because of capacity savings.

#### **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.

#### Monitoring performance in Spectrum Virtualize for Public Cloud on Microsoft Azure

From left window of the GUI, select **Monitoring**  $\rightarrow$  **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 6-14).



Figure 6-14 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

Complete the following steps to monitor performance in Microsoft Azure:

- 1. Login in to the Microsoft Azure Portal.
- 2. Select a node from your configured resource group; for example, a node-vm.
- 3. From the left, window select **Monitoring**  $\rightarrow$  **Metrics**.

4. Choose the metric that you want to display (see Figure 6-15).

≡ Microsoft Azure	∞ Search resources, services, and docs (G+/)	2	₽	Q	٢	0	8	
Home > RDBK-SV4PC_Azure	_Redpaper1 > sv-RDBK1-node1-vm							
SV-RDBK1-no	de1-vm   Metrics …							×
₽ Search (Ctrl+/)	$_{\ll}$ + New chart 🕐 Refresh 😰 Share $\checkmark$ 🙂 Feedback $\checkmark$				Loca	Time: I	ast 24 h	ours (Automatic - 5 minutes)
- sacrap	^ Avg Data Disk Bandwidth Consumed Percentage for sv-RDBK1-node1-vm 🖉							
Disaster recovery Guest + host updates	🏷 Add metric 🦙 Add filter 🐄 Apply splitting 🔝 Line chart N	- 🖪	Drill int	o Logs	~ Q	New a	lert rule	$\not\!\!\!\!\!\!\!\!\!\!\!\!\!\!\!\!\!\!\!\!\!\!\!\!\!\!\!\!\!\!\!\!\!\!\!\!$
😹 Inventory	Scope Metric Namespace Metric Aggregation							
Nange tracking	v-RDBK1-node1-vm Virtual Machine Host V þata Disk Bandwidth C V Avg		~	9				
Configuration managemen (Preview)	Available Memory Bytes (Preview) CPU Credits Consumed CPU Credits Consumed	<u> </u>						
Policies	CPO Credits Remaining	10						
🛓 Run command	∠ Data Disk IOPS Consumed Percentage							>
Monitorina	40% 🛠 Data Disk Max Burst Bandwidth							
Insights	20% 🛠 Data Disk Max Burst IOPS							
V Insignis	0% Set A Disk Oueue Denth	~			12 Ph	4		100+01-00
Alerts	Data Disk Bandwidth Consumed Percentage (Avg)	Dec 20	9:22 AM		1211			01010100
ni Metrics	9+H26K1+00001-Vm 0 %							
Diagnostic settings								

Figure 6-15 Microsoft Azure Metrics

#### 6.1.2 Monitoring SV4PC on Azure 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 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 the following methods:

Method 1

Use a site-to-site 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 VPN gateways 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 VPN gateway connection from your on-premises network to a VNet, see this Microsoft tutorial.

Configuration of site-to-site VPN IPsec tunnel is described in 5.5.1, "Azure configuration for VPNGW IPsec tunnel" on page 100.

Method 2

Install the Storage Insights data collector or deploy the Spectrum Control software directly on a virtual machine (VM) on Azure.

#### Storage Insights data collector implementation on a VM on Azure

This section describes the implementation of a Storage Insights data collector installation off-premise on Windows Server 2019 Datacenter Gen2 operating system on Azure.

As a baseline for the data collector application, a Windows Server 2019 Datacenter operating system was deployed as a VM on Azure 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 IBM Spectrum Virtualize for Public Cloud Azure instance. For more information, see 5.2, "Logging in to IBM Spectrum Virtualize for Public Cloud on Azure" on page 73.

Complete the following steps to deploy the lightweight data collector off-premises to stream performance, capacity, and configuration metadata to IBM Storage Insights:

- 1. Log in to your IBM Storage Insights instance.
- From the Configuration → Data Collector page, the data collector for Windows can be downloaded by using Deploy Data Collector, as shown in Figure 6-16 on page 126.



Figure 6-16 Select Operating System

- 3. Extract the contents of Cloud_DataCollector_windows.zip file on Windows Server 2019 Datacenter.
- 4. Run installDataCollectorService.bat, as shown in Figure 6-17.

Desktop > Cloud_DataCollector_windows			~	U	Search Cloud_DataCollector_windows
Name ^	Date modified	Туре	Size		
📕 bin	14/12/2021 10:15	File folder			
🧵 conf	14/12/2021 10:15	File folder			
🧵 jre	14/12/2021 10:15	File folder			
📜 lib	14/12/2021 10:15	File folder			
🧵 nls	06/12/2021 20:06	File folder			
notices	14/12/2021 10:15	File folder			
readmes_translations	14/12/2021 10:15	File folder			
resources	14/12/2021 10:15	File folder			
ConfigureProxy.bat	14/12/2021 10:15	Windows Batch File		3 KB	
dataCollector.bat	14/12/2021 10:15	Windows Batch File	1	7 KB	
installDataCollectorService.bat	14/12/2021 10:15	Windows Batch File	1	5 KB	
S logcollector.bat	14/12/2021 10:15	Windows Batch File		3 KB	
README_IBM_Stora	14/12/2021 10:15	Firefox HTML Doc	2	2 KB	
S uninstallDataCollect Date-modified: 14/12	2/2021210:15 10:15	Windows Batch File	1	1 KB	
Version.txt	14/12/2021 10:15	TXT File		1 KB	

Figure 6-17 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 to Storage Insights, the cluster with authentication type Secure Shell (SSH) and the private SSH key for the user is used for authentication must be added. Create a separate user for Storage Insights in IBM Spectrum Virtualize. The task of adding a system is shown in Figure 6-18.

Add SVC or Spectrum	Virtualize Storage Systems
Connection Det	ails
Enter the IP addresse or space to separate e	s or host names that you use to connect to the storage systems. Use a comma entries.
* Separate entries	s with commas or spaces
No storage system	ns entered
Enter authentication of storage system determ	credentials that are shared by the storage systems. The firmware version on the nines the user privileges that are required for monitoring. <u>Learn More</u>
Authentication type:	Secure Shell (SSH)
SSH key:	* OR Browse
SSH user:	*
Passphrase:	
(i) Cancel	Back Connect

Figure 6-18 Add SVC or Spectrum Virtualize Storage System window

The task to add the storage system completes automatically, as shown in Figure 6-19. The IBM Spectrum Virtualize system on Azure is visible in Storage Insights within a few minutes after the process completes.

0 Tasks running	Start time	
1 Task completed	End time	Clear all
Adding system of type SAN Volume Controller	Dec 13, 14:26	ū

Figure 6-19 SVC or Spectrum Virtualize Storage System successfully added

#### **IBM Spectrum Control VM on Azure**

This section represents a workflow of how the Spectrum Control software can be implemented on a VM on Azure.

A VM with a supported operating system must be deployed on Azure 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 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 6-20.

G	Dashboards Insights	Resources Reports Confly	guration Unlock the full potential of IBN	l Storage Insights Pro							🄲 🏠 🕹 🔿 IBM.
•	lock Storage Systems V	RDBK-SV4PC Spectrum Virtualiz 8.4.3.0 (build 155.	Azure_Redpaper1 cers e for Public Cloud - 0002 0.2112040719000)	opport							
8	PR Tonia W5 Lovin	Overview Notificati	ions Tickats Propert	es Inventory Data O	offaction						
		Capacity									Compression Savings: 0.00 %
8	1988,20,391	Capacity		۵	Provisioned Capacity			① Capa	acity Savings		
8	Rest. DLRB-F	2,34 % tred			289,45 % Witten				ovigi DO gib		
8	96.995	1,000.00 / 1,024.00 GB Available View capacity by pool			0.00 / 2.00 GB Available View capacity by volume			0.00 0 0.00 0 0.00 0	38 Pool Compression 38 Deduptication 38 Thin-Provisioning		
8	VIIILINK. ROLINKK	Performance							Capacity savings		
e	Laws	2 11	I/O Rate								
0	ROBK-OVIPC_Azire_Redpiper1	0.5	Read 0 cps/s Dec 13, 2021, 02:34:48 PM								
e	244	61	6								
~		09 AM	63 PPI	03.991	06.PM	09.994	Tel 14	02.04	D6.MM	DR AM	12 PM
		Read			Wite						
		2	Data Rate								
			Write 0 MB/s								
		3 14-	<ul> <li>Nead 0 Millys</li> </ul>								
		3	Dec 13, 2021, 02:34:48 PM								
		09 AM	12 (99)	03.89	06.991	09.991	Tel: 3.6	02.04	ON AM	DR AM	12 PM
		Read			Wi2s						
		3 117	Pasananza Timo								
		1 ca-	• Write 0 ma/op								
		£ 0.4-	<ul> <li>Read 0 ms/op</li> </ul>								
		8 6.4-	Dec 13, 2021, 02:34:48 PM								
		\$ 02-									
		@ 63.	11 AN	0 	0. PF	(1) FM	2411	(11 MM	70 M	78 AM	1100
								60.504	01.04	68.04	117
		Read			Wiles						
Add											

Figure 6-20 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 Spectrum Control, and about how the IBM Spectrum Control can be extended with IBM Storage Insights, see this IBM Documentation web page.

- Click Resources → Block Storage Systems to open a detailed view of the performance metrics.
- 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 6-21.

<b>0</b> ×	essurces Notifications (	Groups Configuration									÷ 🖞	L @ IBM.
		Key Performance Indicators Land 2000 - March 1000 - March										
RDB     Spectrum     Decoires     Training	C-SV4PCRedpaper1 + Vrtusike _ Public Dised - 0002 Actives + Govern1				100 Hogers Hits of globalog				80- 60- 80- 80- 80- 80- 80- 80- 80- 8			
Call Horne & Advisor Performent Descention	9 800	12 bm oc/m io_gp0	Turi 14 ORÂNH	12 996	tijan orja	Ter S4	сылн	32.044	12 ⁰ H R	24 7.e54	oslar	12 04

Figure 6-21 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 6-22.



Figure 6-22 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 6-22).

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**  $\rightarrow$  **Block Storage Systems** in the top toolbar.

**Note:** The View Performance pop-up 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 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**  $\rightarrow$  **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 6-23.

0	Dashboards D											. 🕯	A 🔿 IBM.
Block 54	korago Systems	~	RDBK-SV4PC_Azure_Redpaper1 Out Support										
- Eliza		Q	1	8.4.3.0 (build 155.0.2112040719000)									
			Overview	Notifications	Tickets	Properties	Inventory De	ta Collecti	ion .				
6 v	K Teerlah VKI Zuerleh		Compone	nt Health				_					í
•	NUMBER		<b>₽</b> На ⊙	urdware Components Normal					Logical Components     O Normal		🛠 Connectivity Components		0
(i)	AND 24 MIGHT							~		~			0
-			Capacity									Compress	ion Savings: 0.00 %
0 **	CANEL		Capacity	y				6	Provisioned Capacity	0	Capacity Savings		
•	00,548,300,880	8	2,34	1 % Used					289,45 % Written		0.00 GB		
0	2010		1,000.00 /	1,024.00 GiB Available acity by pool					0.00 / 2.00 GB Avalable View capsolity by volume		0.00 GiB Pool Compression 0.00 GiB Deduplication 0.00 GiB Thin-Provisioning		
() R0	065WPC_Asre_Redg	paper1									View capacity savings		

Figure 6-23 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**  $\rightarrow$  **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 6-24. The Overview page can be opened by selecting **Resources**  $\rightarrow$  **Block Storage Systems** in the top toolbar.

$\odot$	RDBK	-SV4PCRedpaper1 -	Overview Capacity 2,34 % Used 1,000.00 / 1,024.00 GiB Available View capacity by pool
	Spectrum \	ni tualizePublic Cloud - 0002	
		Actions 🔻	
		Add to General Group	Capacity 🔻
	Overview	View Performance	
	Tickets	View Capacity	
	Call Home	Get Logs	
<u>انم</u>	Advisor	Create/Update Ticket	
	Performan	Export Performance Data	
	Properties	View Real-Time Performance	
	Alerts (0)	Data Collection	
æ	Alert Defin	Connections •	
	]	Open Storage System GUI	
8	Volumes (2	.)	
8	Pools (1)		
	Managed D	isks (2)	
19	RAID Array	rs (2)	
-	I/O Groups	(1)	
	Nodes (2)		
	IP Ports (4)	)	
	I	Related Resources	
		None	
1			

Figure 6-24 View Capacity an IBM Spectrum Virtualize cluster on Azure through overview page

Alternatively, the Capacity view can be opened by right-clicking a device in **Resources**  $\rightarrow$  **Block Storage Systems**, as shown in Figure 6-25.

Storage Systems Alerts	SME=KB=LAB =	Block Storage	e System	15	
Image: Actions ▼	View Performance Vi	iew Capacity Get Logs			
Name	Туре	Serial Number	Firmware	Condition	
RDBK-SV4PC Azure Re View Properties View Performance View Capacity Get Logs Create/Update Tic Export Performance Data Collection View Alert Definiti Add to General Gr Connections Remove	ket bons bup	0000001CBA34F260	8.4.3.0 (b	⊘ Normal	

Figure 6-25 View 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.

#### 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 a IBM Spectrum Virtualize storage system on Azure 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 6-26.

All Policies							
Alert Policies Resources by Policy							_
Create Policy	Create Policy						×
Name Resource Type							. [
	Name	Option	nal:				- 1
	Spectrum Virtualize for Public Cloud on Azure	autom	the resources that you want t atically applied to those resou	o add to the policy. Ale irces. You can also cho	ert definitions in the pol oose to add resources la	cy are ter.	- 1
	Policy Type			٩	RDBK	Reset	1
	Block Storage Systems	~	Name	▲ Current Poli	icy	ļļ	
	Type of storage system		RDBK-SV4PC_Azure_Rev	dpa			
	SAN Volume Controller	$\sim$					
							w.
					Cancel	Create	

Figure 6-26 Creating a policy for an IBM Spectrum Virtualize cluster on Microsoft Azure

2. Define an alert in the policy, as show in Figure 6-27.

Back to All Policies Spectrum Virtua	lize fo	or Public Cloud on Azure Change Policy Name
Alert Definitions Resc	ources	
✓ Storage System	1/32	✓ Used Capacity (%)
General	0/6	Operator Value Severity
Capacity	1/26	>= V 85 % () ()
Performance	0/0	Email Override Senarate entries with commas or spaces
> Volumes	0/22	<ul> <li>✓ Hide Additional Options</li> </ul>
> Pools	0/43	Notification Frequency     Send once until problem clears   V
> Managed Disks	0/5	<ul> <li>Only send notifications after condition is violated for</li> </ul>
> Drives	0/4	Available Capacity
> I/O Groups	0/7	Adjusted Load Conscity (%)
> Nodes	0/3	
> FC Ports	0/3	
> IP Ports	0/3	Recent Growth
> Host Connections	0/2	Recent Fill Rate (%)
Custom	0/0	Total Canacity Savings (%)

Figure 6-27 Alert definition in the policy for an IBM Spectrum Virtualize cluster on Microsoft Azure

- 3. Click Save Changes.
- 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:

- 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  $\rightarrow$  Notification Settings.
- 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.
- 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 of 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**.
- 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.
- 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 6-28 shows the Email Override and Policy Notifications notification settings in IBM Storage Insights.

Back to All Policies							
Spectrum Virtualize for	Public Cloud on	Azure Change Ba	olicy.Name				
Alert Definitions (2) Resources (	1)						
🔯 🔯 ill Actions 🕶					Q 🔻 Filter.	Policy Notifications	
Condition	Category	Subcategory	Severity	Notification Prequency	Email Override	B Frail Addresses (i)	
Storage System Used Capacity >= 95 %	Storage System	Copecity	<ul> <li>Oritical</li> </ul>	Send Once	email-advas@company.com	email-adress@company.com	
System CPU Utilization >= 90 %	Nado	Performance	Orisical	Seed Once	Nasa		

Figure 6-28 Notification settings for an IBM SV4PC 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:

- 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.
- 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.

# 6.2 Troubleshooting Spectrum Virtualize for Public Cloud on Microsoft Azure

The section describes how to collect support data in an Spectrum Virtualize for Public Cloud on Azure environment.

#### 6.2.1 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**  $\rightarrow$  **Support** menu.

#### Collecting information by using the GUI

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

 Click Settings → Support and then, the Support Package tab (see Figure 6-29). 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.

IBM Spectrum Virtualize for Public Cloud	RDBK-SV4PC_Azure_Redpaper1 Su	port 🌲 📳 🕐 superuser (Beounty Administrator) 🗸						
K	Call Home	Support Package function to the support center or download and send them manually to help support personnel analyze and fix errors on the system.						
Dashboard	Support Assistance	To upload support packages, a DNS server must be configured on the system.     Configure						
୍କ୍ତ ^e Monitoring	Support Package	1 Upload Support Package *						
Pools		Manual Upload Instructions						
Volumes		Download a vupport package or logs here:     Download Support Package *						
Hosts		2. Use one of the following portals to upload the support package or logs through your browner:						
Copy Services		Elise Siencod (3)         Valued one Economic (3)         P22 Transfer (3)           It is a Siencod climits: Uploading support package feature is approved or matcher official.         For small files (recommendation: 200H8 mailmun)         For faster data transfer with larger files						
Access		For the Diamond accounts						
کې Settings	Notifications Network Security	CBMMLaggig Law (						
	System Support	Manage tickets and uppload support packages with IDM Storage Indigits Sign Up Laurn More     to Sign Up Laurn More     to Sign Up Laurn More     down and the second storage storage storage storage storage storage     down and the second storage storage storage storage storage     down and the second storage storage storage storage storage						
	GUI Preferences	Laterncy O ms 0 ms 0 ms Bandwidth O Mose o Maps a Maga 10PS O 0 0						

Figure 6-29 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 states are 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 states aves 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 or PMR number. Enter the case or PMR number into PMR field and select the snap type, often called snap option 1, 2, 3 or 4, as requested by IBM Support (see Figure 6-30). In our example, we enter a PMR 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**.

Upload Support Package	×
Your system will generate and upload a new package to the IBM support center.	
PMR Number: Don't have PMR?	
ppppp,bbb,ccc	
Select the type of new support package to generate and upload to the IBM support center:	
Snap Type 1: Standard logs	
Contains the most recent logs for the system, including the event and audit logs.	
Snap Type 2: Standard logs plus one existing statesave	
Contains all the standard logs plus one existing statesave from any of the nodes in the system.	
Snap Type 3: Standard logs plus most recent statesave from each node	
Contains all the standard logs plus each node's most recent statesave.	
Snap Type 4: Standard logs plus new statesaves	
Contains all the standard logs and generate a new statesave on each node in the system.	
⑦ Need Help Cancel Upload	

Figure 6-30 Upload Support Package window

The procedure to create the snap on an IBM Spectrum Virtualize system, including the latest states ave from each node, starts. This process might take a few minutes.

#### Collecting logs by using the CLI

Complete the following steps to use the CLI to collect and upload a support package as requested by IBM Support. This example shows a case number, but you can use PMR number as well:

- 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=TSXXXXXXXX gui2

- Standard logs plus most recent states ave 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=TSXXXXXXXX gui3
```

2. We collect the type 3 (option 3) and have it automatically uploaded to the PMR number that is provided by IBM Support, as shown in Example 6-1.

```
Example 6-1 The svc_snap command
```

```
ssh superuser@IP_address
Password:
RDBK-SV4PC_Azure_Redpaper1:superuser>>svc_snap upload pmr=TSXXXXXXXX gui3
```

3. If you do not want to automatically upload the snap to IBM, do not specify the upload pmr=TSxxxxxxx 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 states aves are included).

 The generated file can then be retrieved from the GUI under the Settings → Support → Manual Upload Instructions twisty → Download Support Package. Click Download Existing Package, as shown in Figure 6-31.



Figure 6-31 Download Existing Package

5. A new window opens. Click in the **Filter** box and enter snap; then, click **Enter**. A list of snap files is shown (see Figure 6-32). Locate 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**.

Select Support Package or Logs to Download									
You can select a previously created support package or individual logs to download.									
nodel -		Filter Applied	×	A					
File Name		$\uparrow$		P					
/dumps/snap./F4C6FF.211214.090811.0	og			^					
/dumps/ <b>snap</b> .7F4C6FF.211214.090811.g	zip.out								
/dumps/ <b>snap</b> .7F4C6FF.211214.090811.g	zip.log								
/dumps/ <b>snap</b> .7F4C6FF.211213.144120.ld	og								
/dumps/snap.7F4C6FF.211213.144120									
/dumps/snap.7F4C6FF.211213.083441.t	gz			~					
Showing 53 Files   Selecting 1 File									
⑦ Need Help	Cancel	C	ownload						

Figure 6-32 Filtering on snap to download

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

## 6.2.2 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 or PMR 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 6-33).

IBM		Search	Q 8 ≡
	Enhanced Customer Data Repository (ECuRep)		
	ECuRep Secure Upload Terms of use Help		
	Case PMR RCMS CROSS SRID Machine Type/Serial (No case) Software (No case)		
	The fields indicated with an asterisk (*) are required to complete this transaction; other fields are optional. If you do not want to provide us with the required information, please use the "Back" button on your browser to return to the previous page, or close the window or browser session that is displaying this page. Case number:* Email Notification: BL****T@de.ibm.com Continue	Usage information Enter the case number you got from IBM support (e.g. TS123465789). If you select an email address, an email will be sent on failure or success. The selected address will noh be used to send the notification.	
	ECuRep terms of use   ECuRep Information   Previous Upload Version   Alternate Upload Options		
	Contact1BM Privacy Terms of use Accessibility	United States - English 🛛 🗸	

Figure 6-33 ECuRep details

- 2. Complete the following required fields:
  - Case number or PMR (mandatory) as provided by IBM Support for your specific case. This number must be in the format of TSxxxxxxxx; for example, TS123456789 or PMR 12345,789,021.

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

3. When completed, click Continue. The Input window opens (see Figure 6-34).

TDV			
Ibn	Enhanced Customer Data Re (ECuRep)	pository	
	ECuRep Secure Upload Terms of use Help		
	The fields indicated with an asterisk (*) are required to complete this tra nformation, please use the "Back" button on your browser to return to	ansaction; other fields are optional. If you do not want to provide us with the required the previous page, or close the window or browser session that is displaying this page.	Usage information Select the files you want to
\$	Click to Select Files		upload to IBM and select the Upload button. Files are processed after a complete upload is successful. You have 24 hours to resume any paused uploads before they
	map.7F4C8FF.211210.168089 tgz	2 M8	are removed from IBM. If a file upload fails or is cancelled, that file data will be removed.
			NOTE: It is recommended to use the latest browser version to take advantage of the uploader features. Older browser versions have limits such as 2 GB files sizes and po
	1 file selected	2 MB	ability to pause/resume
	Upload		uploads.
[	Back		

Figure 6-34 ECuRep File upload

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

#### 6.2.3 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 (CIM) 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.

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:

 When you are accessing SAT by using <cluster IP address>/service, the configuration node canister SAT GUI login window opens. Enter the Superuser Password, as shown in Figure 6-35.



Figure 6-35 Service Assistant Tool Login GUI

2. After you are logged in, you see the Service Assistant Home window, as shown in Figure 6-36. 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.

IBN	1 Spectrum Virtualize for Pu	ublic Cloud Service Assist	ant Tool								9
	Available Here is a list of all the Q Filter nodes by par	Nodes nodes connected to yo	our system. Select a no	de to see its details an	d make change	s to it.		More Infor	mation ()	Refresh	8
	Local (1)										
2	CONFIG node1 • Active	Panel 7F4C6FF		Relationship Local							
	System (1)										
	node2 • Active	Panel D176210		Relationship System							

Figure 6-36 Service Assistant Tool GUI

3. Click the node on which you want to run actions (see Figure 6-37).

**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.

IBM Spectrum Virtualize for Public Cloud Service	Assistant Tool						8
← All nodes • node1 • 7F4C6 Node De	etails						
Node Details	C-1						
Collect Logs	Selected Node						
Manage System	nodo1						
Recover System	Active	7F4C6FF		Local	•	Actions	×
Re-install Software							
Update Manually							
²² Configure Enclosure	Node Errors (0)						
Change Service IP	Error Code		Error data	Recommen	ded Action		
Change Node IP Address							
Change Node Discovery Subnet	Node Details						
Ethernet Connectivity	Node	Hardware					
Configure CLI Access							
Restart Service							
	Node WWNN:						

Figure 6-37 SAT node details

For more information about how to use the SA Tool, 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 6-38).

18M Spectrum Virtualize for Public Cloud Service Assistant Tool							
Collect Logs • node1 • 773C4C. Outcome te and download a package of log and trace files, or download existing log files from the node. The support package, which is also called a SMAP file, can be used by support personnel to understand problems on the system: Unless advised by support, collect the latest statesave.							
Node Details Collect Logs	Create and download new support logs						
Managis System Reciver System Re-install Software Update Manually Occonfigure Enclosure Chango Service IP Chango Node IP Address	Download with latest statesave     Download without latest statesave       Includes the latest statesave from the selected node and contains information about current state of the software.     Download without latest statesave $\pm$ $\pm$ Download existing logs $\pm$						
Change Node Discovery Subnet Ethernet Connectivity Configure CLI Access Distant Sandon		d A					

Figure 6-38 Collect Logs using SAT

Download the snap file to your workstation and upload file to IBM Support, as described in 6.2.2, "Uploading files to the Support Center" on page 141.

### 6.2.4 Remote Support Assistance

Remote Support Assistance allows IBM Support to remotely connect to the 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:

 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 6-39).



Figure 6-39 Remote Support Assistance menu

2. Click the **?** to get more help. A pop-up window appears, in which the prerequisites for all configurations are listed (see Figure 6-40).

#### 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	
 	<b>a</b>				

Figure 6-40 Set up Support Assistance help window

# Figure 6-41 shows the first window in the wizard that is used to configure the optional Remote Support Proxy.

Set Up Support Assis	tance		×
Support Centers			
Support centers respond to n configured on the system:	nanual and automatic se	rvice requests from the system	n. The following support centers are
Name default_support_center0 default_support_center1	IP Address 129.33.206.139 204.146.30.139	<b>Port</b> 22 22	
Remote Support Proxy (O	ptional)		
A proxy is required the network.	or network configuratio	ns using a firewall, or for syst	ems without direct connections to
Name	IP	Port	$\oplus$
Cancel	⑦ Need Help	<b>⊲</b> Back	Next ►

Figure 6-41 Configure optional Remote Support Proxy

3. In the next window (see Figure 6-42 on page 148), 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**.

et Up	Support Ass	sistance		
mote	Support Acces	s Settings		
en do se se	you want service ttings at any time.	personnel to complete r	naintenance and service ta	sks remotely? You can change
0	At Any Time			
	The support cent	er can start remote sup	port sessions any time	
0	On Permission O	nly		
	The support cent be configured for	er can start a remote su the session.	pport session only if permi	tted by an admin. A time limit can

Figure 6-42 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 6-43 shows a successfully tested connection. Now, click **Start New Session** to open a tunnel for IBM Support to connect.

IBM Spectrum Virtualize for Public Cloud	RDBK-SV4PC_Azure_Redpaper1	iupport	🌲 🖡 🛛	superuser Security Administrator
<	Call Home	Support Assistance Support assistance enables service personnel to access the system to perform maintenance and resolve issues.		
Dashboard	Support Assistance	A DNS server must be defined for Support Centers to use domain names.     Configure		
ംപ്പ് Monitoring	Support Package	Remote Support		
Pools		Start New Session		
Volumes		Sugart Users and History Monitor Users: 0 Piviliaged Users: 0		
Hosts		Token Age: Coxy Connential New Token Remote Support Access Berling: On Permission Only Provs Serves Configured: O		
Copy Services		Reconfigure Settings		
Access				
() Settings				

Figure 6-43 Remote Support Status and session management

A pop-up 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 6-44).



Figure 6-44 Disabling remote Support Assistance

#### 6.2.5 Troubleshooting in Microsoft Azure

In this section, we provide some examples of how to troubleshoot problems in Microsoft Azure.

#### **Enabling Boot diagnostics**

Debug Microsoft Azure VM boot problems by completing the following steps to enable boot diagnostics for the Spectrum Virtualize node VMs:

- 1. In the Microsoft Azure portal, select the node-vm from your defined resource group.
- 2. In the Help that is in left window, select **Boot diagnostic**  $\rightarrow$  **settings**.

3. To connect to serial console, select **Enabling boot diagnostics with custom storage account**; otherwise, use the recommended option (see Figure 6-45).



Figure 6-45 Microsoft Azure Boot diagnostics

A node VM with Boot diagnostic enabled collects serial log information and screen captures from boot time (see Figure 6-46).

≡	Microsoft Azure	, $\rho$ Search resources, services, and docs (G+/)	P	8	0	ନ୍ଦି	
Но	ome > Resource groups > R	RDBK-SV4PC_Azure_Redpaper1 > sv-RDBK1-node2-vm					
*	sv-RDBK1-nod	de2-vm   Boot diagnostics					×
5	Search (Ctrl+/)	« 🜔 Refresh 🔞 Settings 🤌 Troubleshoot					
P	Logs	^					
	Connection monitor (classic)	Screenshot Serial log					
×	Workbooks	Updated: Monday, December 20, 2021, 8:11:11 AM UTC Download serial log					
Au	itomation	[ 0.09/4/5] SCSI I:0:I:0: Attached SCSI Benelic SB0 thbe 0					
A	Tasks (preview)	<pre>[ 6.697682] scsi 0:0:0:0: Attached scsi generic sg1 type 0 [ 6.697843] scsi 3:0:0:61: Attached scsi generic sg2 type 0</pre>					
Ţ	Export template	[ 6.699991] scsi 3:00:0:62: Attached scsi generic sg3 type 0 [ 6.705853] scsi 5:0:0:0: CD-ROM Msft Virtual CD/ROM 1.0 PQ: 0 ANSI: 5					
He	elp.	[ 6.713689] sd 1:0:1:0: [sda] 268435456 512-byte logical blocks: (137 GB/128 GiB)					
$\otimes$	Resource health	[ 6.713632] Sa 1991101 [Saa] 4095-byte physical blocks [ 6.713749] Sa 1991101 [Saa] Write Protect is off [ 6.713749] Sa 199120 [Saa] Write Protect is off					
	Boot diagnostics	[ 6.713798] sd 0:0:0:0: [sdb] 419430400 512-byte logical blocks: (215 Gb/200 G1B) [ 6.713798] sd 0:0:0:0: [sdb] 4096-byte physical blocks					
1	Performance diagnostics	<pre>[ 6.713842] sd 3:0:0:61: [sdc] 1073741824 512-byte logical blocks: (550 GB/512 GiB) [ 6.713844] sd 3:0:0:61: [sdc] 4096-byte physical blocks</pre>					
+	Reset password	[ 6.713851] sd 0:0:0:0: [sdb] Write Protect is off					
*	Redeploy + reapply	<ul> <li>6.714033] sd 3:0:0:61: [sdc] Write Protect 13 011</li> <li>6.714033] sd 3:0:0:61: [sdc] Write cache: disabled, read cache: enabled, supports DPO and FUA</li> </ul>					
1	Serial console	[ 6.714034] sd 0:0:0:0: [sdb] Write cache: disabled, read cache: enabled, supports DPO and FUA [ 6.714152] sd 1:0:1:0: [sda] Write cache: disabled, read cache: enabled, supports DPO and FUA					
	Connection troubleshoot	[ 6.715114] sd 3:0:0:62: [sdd] 1073741824 512-byte logical blocks: (550 GB/512 GiB)					
	Red Hat Customer Portal	[ 6.715276] sd 3:0:0:62: [sdd] 4000-0yte physical blocks					
0	Support + Troubleshooting	<pre>[ 6.715401] sd 3:0:0:62: [sdd] Write cache: disabled, read cache: enabled, supports DPO and FUA [ 6.715695] scsi 5:0:0:0: Attached scsi generic sg4 type 5 </pre>					

Figure 6-46 Serial log in Microsoft Azure

For more information about Microsoft Azure Boot diagnostics, see this Microsoft Docs web page.

#### 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 6-47). For more information, see this Microsoft Docs web page.

= Microsoft Azure		resources, services, and docs (G+/)	Σ	R	Ç.	۲	0	8	BARAN Tigan dina ang 🧶
Home > RDBK-SV4PC_Azure_F	Redpaper	1 > sv-RDBK1-node2-vm							
sv-RDBK1-nod	e2-vr	n   Serial console 🦷							×
	«	🕐 Refresh 🕞 Start							
🧬 Logs	^								
Connection monitor (classic)		G Serial Console requires a custom boot diagnostics storage account to be used, and is not yet fully compatible with managed b	oot diag	nostics s	torage a	account	s. Click	to view :	and change your boot
🞽 Workbooks		diagnostics storage account configuration. Click here C ⁹ for more details or if you are already using a custom storage account	and rece	ve this e	rror.				
Automation									
🖧 Tasks (preview)									
Export template									
Help									
℅ Resource health									
Boot diagnostics									
Performance diagnostics									
📍 Reset password									
Redeploy + reapply									
Serial console									
Connection troubleshoot	- 1								
Red Hat Customer Portal									
O Support + Troubleshooting	~								

Figure 6-47 Serial console requirements

Figure 6-48 shows serial console view after enablement of boot diagnostic with custom storage account.

Home > sv-RDBK1-node2-vm >	
sv-RDBK1-node2-v	m   Serial console ···· ×
✓ Search (Ctrl+/) «	? Feedback[2] (3) (0) ##
📮 Overview	2021-12-20T08:33:28.0104302 INFO Daemon CGroups Status: The cgroup filesystem is ready to use 2021-12-20T08:33:28.0176352 INFO Daemon Run daemon
Activity log	2021-12-20T08:33:28.0279112 INFO Daemon No RDMA handler exists for distro='Red Hat Enterprise Linux' version='8.3' 2021-12-20T08:33:28.0530622 INFO Daemon cloud-init is enabled: True
Access control (IAM)	2021-12-20T08:33:28.065840Z INFO Daemon Using cloud-init for provisioning
🗳 Tags	2021-12-20708:33:28.0960022 INFO Daemon Frovisioning already completed, skipping.
Diagnose and solve problems	2021-12-20100-33-20.120022 Into Datamon Rever capacities are not enabled, skiping 2021-12-20108:33:28.1257952 INFO Daemon Determined Agent WALinuxAgent-2.6.0.2 to be the latest agent [ 25.901010 [c]nud-init[L647]. [c]nud-init x, 19.4 running 'moduleg:config' at Mon. 20 Dec 2021 08:33-28 40000. In 25.78 second
Settings	S. ( ov ) Charted Jupin, the action provided in sland and in
Networking	[ OK ] Started Apply the Settings Specified in cloud-config. Starting Execute cloud user/final scripts
🖉 Connect	2021-12-20T08:33:20.619035Z INFO ExtHandler ExtHandler The agent will now check for updates and then will process extensions. Ou tput to /dev/console will be suspended during those operations.
B Disks	[ 26.555433] cloud-init[1829]: Cloud-init v. 19.4 running 'modules:final' at Mon, 20 Dec 2021 08:33:28 +0000. Up 26.43 seconds.
🖳 Size	<pre>[ 26.567300] cloud-init[1829]: Cloud-init v. 19.4 finished at Mon, 20 Dec 2021 08:33:29 +0000. Datasource DataSourceAzure [see d=/dev/sr0]. Up 26.54 seconds</pre>
Security	[ 26.622936] echo[1939]: trying to reload or restart NetworkManager.service
Advisor recommendations	[ OK ] Started Execute cloud user/final scripts.
Extensions + applications	[ OK ] Started Session c4 of user root.
🐔 Continuous delivery	[ OK ] Started Session c5 of user root.
Availability + scaling	Red Hat Enterprise Linux 8.3 (Ootpa) Kernel 4.18.0-240.22.1.el8 3.x86 64 on an x86 64
Configuration	av-RDBK1-node2-vm login:
🔶 Idoptity	of their house in regime

Figure 6-48 Serial Console in Microsoft Azure

#### 6.2.6 Hints and tips

This section provides helpful hints and tips for your deployment and cluster management.

#### **Deployment hints and tips**

After a successful deployment, it is expected that you receive an email that includes the IP addresses of the cluster and link to next steps (see Figure 6-49).

	<b>(TERNAL) IBM</b> to: mailsacter, inf	SVPC Cluster Deployment: Cluster RDBK-SV4PC_Azure4 is created successfully
Your SVPC cl	luster RDBK-SV	'4PC_Azure4 is created successfully.
IP addresses	s created in 1	esource group:
Item		IP
Cluster IP		10.9.0.4
Nodel Servio	ce IP	10.9.0.14
Nodel Node I	IP 1	10.9.0.12
Nodel Node I	IP 2	10.9.0.6
Nodel Port 1	IP 1	10.9.0.13
Nodel Port 1	IP 2	10.9.0.7
Node2 Servio	ce IP	10.9.0.20
Node2 Node 1	IP 1	10.9.0.18
Node2 Node 1	IP 2	10.9.0.23
Node2 Port 1	IP 1	10.9.0.19
Node2 Port 1	IP 2	10.9.0.24
To access IH	3M Spectrum Vi	rtualize for Public Cloud management GUI, please follow this link:
https://www.	.ibm.com/docs/	en/spectrumvirtualizec1/8.4.x?topic=azure-completing-post-installation-tasks

Figure 6-49 Example of deployment email

**Note:** At time of this writing. the email was not delivered at times. In such a case, check the IP addresses through the Microsoft Azure portal and follow the postinstallation instructions that are available at this IBM Documentation web page.

In Microsoft Azure, select your virtual network of your resource group (see Figure 6-50).

	esources, services, and docs (G+/)		D G	R © © R	BLEWIT gebruikerunen 🥹						
Home > Resource groups > RDBK-SV4PC	lome > Resource groups > RDBK-SV4PC_Azue_Redpaper1 >										
← sv-default-vnet ☆ · Virtual network					×						
	$\bigodot$ Refresh $\rightarrow$ Move $\lor$ iii Delete	🖗 Give feedback									
<ul> <li>Overview</li> </ul>	Search connected devices				^						
Activity log	Device ↑↓	Туре ↑↓	IP Address ↑↓	Subnet ↑↓							
Access control (IAM)	sv-RDBK1-quorum-nic0	Network interface	10.9.1.4	sv-quorum-subne	t						
Tags	rdbk-winvm784	Network interface	10.9.1.5	sv-quorum-subne	t						
Diagnose and solve problems	NIC4quorumsubnet	Network interface	10.9.1.6	sv-quorum-subne	t						
Settings	rdbk-si-collector650	Network interface	10.9.1.7	sv-quorum-subne	t						
Address space	rdbk-si-collector-quorum-subnetNIC	Network interface	10.9.1.8	sv-quorum-subne	t						
6 Connected devices	sv-RDBK1-LB	Load balancer	10.9.0.4	sv-cluster-subnet							
<ul> <li>Subnets</li> </ul>	sv-RDBK1-LB	Load balancer	10.9.0.5	sv-cluster-subnet							
DDoS protection	sv-RDBK1-node1-nic2	Network interface	10.9.0.6	sv-cluster-subnet							
General Firewall	sv-RDBK1-node1-nic2	Network interface	10.9.0.7	sv-cluster-subnet							
Security	sv-RDBK1-node1-nic2	Network interface	10.9.0.8	sv-cluster-subnet							
Network manager	sv-RDBK1-node1-nic2	Network interface	10.9.0.9	sv-cluster-subnet							
DNS servers	sv-RDBK1-node1-nic2	Network interface	10.9.0.10	sv-cluster-subnet							
Peerings	sv-RDBK1-node1-nic0	Network interface	10.9.0.11	sv-cluster-subnet							
Service endpoints	sv-RDBK1-node1-nic1	Network interface	10.9.0.12	sv-cluster-subnet							
· · · · · · ·	sv-RDBK1-node1-nic1	Network interface	10.9.0.13	sv-cluster-subnet							

Figure 6-50 IP address of your deployment in Azure

#### Identifying MDisks in Microsoft Azure view and Spectrum Virtualize CLI

To view all attached cloud MDisks in Spectrum Virtualize, use the GUI or the CLI and complete the following steps:

 In the GUI, select in the left window Pools → MDisk by pools (see Figure 6-51). Record the Cloud disk ID.

IBM Spe	ctrum Virtualize for Public Cloud RDBK-SV4PC_	Azure_Redpaper1 MDisks by Pools			
>	$\oplus$ Create Pool $\begin{tabular}{lllllllllllllllllllllllllllllllllll$				
<u>^</u>	Name	Cloud Disk ID	Cloud Disk Type	State	Usable Capacity
	Unassigned MDisks (0)				
<del>مر</del> م	V Pool0			🗸 Online	32.00 GiB / 1.00 TiB (3%)
	mdisk1	sv-RDBK1-Mdisk-2	standardSSD	🗸 Online	512.00 GiB
ŧ	mdisk0	sv-RDBK1-Mdisk-1	standardSSD	🗸 Online	512.00 GiB

Figure 6-51 GUI MDisk view

2. From the CLI run the **lslocaldisk** command to view the same information. The Cloud disk ID is shown as disk ID (see Figure 6-52).

1	IBM Spect	rum Virtual	ize:RD	BK-SV4PC	Azure	Redp	aperl:sur	eruser>1	slocaldis	k									
	mdisk id	mdisk name	status	mode	mdisk	grp	id mdisk	grp name	capacity	encrypt	disk id		type	node	id nod	name	iops	state	zone
	o —	mdisk0	online	managed	0 -		Pool0		512.0GB	yes	sv-RDBK1-	Mdisk-1	standardSSD	2	nod	1	500	in-use	eastus
	1	mdiskl	online	managed	0		Pool0		512.0GB	yes	sv-RDBK1-	Mdisk-2	standardSSD	1	nod	2	500	in-use	eastus
1																			

Figure 6-52 CLI command Islocaldisk

3. From the Microsoft Azure portal, select your resource group, filter for mdisks (see Figure 6-53).

=	Microsoft Azure 🔑 Search resour	ces, services, and docs (G+/)		Σ	- G	Q	0	? ž	8	andre dens com 🔕
Hor	ne > Resource groups >									
»	RDBK-SV4PC_Azure_	_Redpaper1 🖈 …								×
		+ Create ≡≡ Edit columns 📋 Delete resource group 💍 Refresh 🞍 Export to C	SV 😽 Open q	query   🤇	Assign	tags	→ Mov	e 🗸 📋	Delete ···	
	(e) Overview	↑ Essentials								JSON View
	Activity log	Subscription (Move) : Microsoft Azure Enterprise_ikky	Deploymer	nts:1 Failed	l,23 Succe	eded				
	Access control (IAM)	Subscription ID ( ) All Shift of the share better share by the	Location	: East US						
	🗳 Tags	Tags (Edit) : ClusterIPFailover : LB_Based_Cluster_IP_Deployment IPFailoverEna	bled : False 9	SVPC_Defaul	t_MTU : 1	500	deployr	mentTag :	sv-RDBK1	✓ More (6)
	🛧 Resource visualizer									
	🗲 Events	Resources Recommendations (19)								
	Settings	mdisk Type all X Location all X the Add fill	tor							
	Deployments									
	Security	Showing 1 to 2 of 2 records.			No grou	ping		~	≔ List view	~
	Policies	Name ↑↓	Туре ↑↓				Locati	on ↑↓		
	Se Properties	Sv-RDBK1-Mdisk-1	Disk				East U	s		
	🔒 Locks	Sv-RDBK1-Mdisk-2	Disk				East U	S		
	Cost Management									
	💺 Cost analysis									
	Cost alerts (preview)	- Provinur Dage 1 Nd of 1 Mayts								
ttos//no	Budgets tal assure com/#8/hm pamicrosoft.com/resource/subscriptions/99	rege i      vii i rekt?								

Figure 6-53 MDisk in Azure

 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 6-54 on page 154).

≡ Microsoft Azure 🔎 Sea	rch resources, services, and c	locs (G+/)					₽ Q	© (	) R	BARANT pole deruces and an environmental services
Home > RDBK-SV4PC_Azure_Redpap sv-RDBK1-Mdisk-1 Disk	er1 > ☆ …									×
P Search (Ctrl+/) ≪	+ Create VM +	Create snapshot 📋 Delete 🏷 Refresh								
8 Overview	^ Essentials									JSON View
Activity log	Resource group (Move	: RDBK-SV4PC_Azure_Redpaper1 Attached			Disk size Disk sku	:512 GiB :Standard SSI	DIRS			
Access control (IAM)	Location Subscription (Move)	: East US : Microsoft Azure Enterprise ikky			Managed by Operating system	: 2 VMs	0 00			
Settings	Subscription ID	294/07/01/01/484-1011-04/06/02/9	(		Max shares	: 2				
Configuration	Time created	: 12/10/2021, 11:13:11 AM			Availability zone	: None				
Size + performance					Security type	: Standard				
Encryption	Tags (Edit)	BM-SV-cluster-name : RDBK-SV4PC_A	ure_Redpaper1	IBM-SV-cluste	-id : 1ca00e354e	2				
<li>Networking</li>										
Disk Export	Show data for last:	1 hour 6 hour	s 12 hours	1 day 7 days						
Properties										
🔒 Locks										
Monitoring	Disk Bytes/sec	(Throughput)	\$P	Disk Operation:	s/sec (IOPS)			Ś		
Metrics				100						
Automation				80						
🖧 Tasks (preview)	60			40						

Figure 6-54 Tag fields in Microsoft Azure

# 6.3 IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Support contact

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.

#### 6.3.1 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 6-2 as an example.

Situation	Client	Cloud provider	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

Table 6-2 Simplified workflow definition based on RACI matrix

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.

### 6.3.2 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 of the other solutions that are based on IBM Spectrum Virtualize. IBM Support can be engaged by using one of the following methods:

- Visit the IBM Support web page to open a case for IBM Spectrum Virtualize for Public Cloud
- ► By phone (see the IBM Directory of worldwide contacts web page
- Through IBM Storage Insights
- ► IBM Call Home

After you receive a case number or PMR, 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.

	IBM									ĉ
	Support	Downloads 🗸	Documentation $$	Forums	Cases 🗸	Monitoring 🗸	Manage support account 🗸		Open a case	
•	IBM is enterpr produc	actively responding t rise, IBM products ar ts and services. Read	to the reported remote nd IBM services that m d more	code executic ay be potentia	on vulnerability Illy impacted,	y in the Apache Log and will continually	(4) 2 Java library dubbed Log4Shell (or LogJam). V publish information to help customers detect, in	Ve are investigating and tak vestigate and mitigate attac	ting action for IBM as a cks, if any, to their IBM	an 1
			I	_et's	s tro	ubles	hoot.			
			C	λ Search	known iss	sues, docume	ntation and support forums			
				View your ca Check the stat review case hi ③	ases tus of active ca istory	ases and	Open a case If you're unable to resolve an issue on your own, receive help from our experts			
									💮 Chat with Suppo	rt

Figure 6-55 IBM Support web page

When opening a new case, select **Spectrum Virtualize for Cloud** as the Product (see Figure 6-56).

	Q 8.
Support Downloads V Documentation V Forums Cases V Monitoring V Manage support account V	Open a case
IBM is actively responding to the reported remote code execution vulnerability in the Apache Log4j 2 Java library dubbed Log4Shell (or LogJam). We are investigating an enterprise, IBM products and IBM services that may be potentially impacted, and will continually publish information to help customers detect, investigate and mitigate products and services. Read more	d taking action for IBM as an attacks, if any, to their IBM
Open a case	
Type of support <b>O</b>	
Product support ÷	
*Title	
This field is required	
* Product manufacturer	
твм х	
* Product	
Spectrum Virtualize for Cloud X	
	💬 Chat with Support

Figure 6-56 Open a support case in IBM support page

Complete all fields, including severity and description.

#### 6.3.3 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 6-57 shows how to create a support request in Microsoft Azure portal by selecting **Help + support**.

■ Microsoft Azure	h resources, services, and docs (G+/)		Ŗ	Q	© ©	ন্দ	
Home > Arrow Help + support							×
	lpha Create a support request 🛛 K Choose the right s	upport plan					
Support	Service health	ur subscription. View outage details					
<ul> <li>All support requests</li> <li>Support Plans</li> <li>Service Health</li> </ul>	Have an issue with your resource? Select Troubleshoot to get targeted solutions for you	r recently visited resource.					
🤤 Advisor	Name	Туре		Last vi	ewed		
	sv-RDBK1-node2-vm	Virtual machine		57 mir	ago		Troubleshoot
	sv-RDBK1-quorum	Virtual machine		5 hrs a	igo		Troubleshoot
	🖳 rdbk-winvm	Virtual machine		5 hrs a	igo		Troubleshoot
	<ul> <li>↔ sv-default-vnet</li> </ul>	Virtual network		2 days	ago		Troubleshoot
	💎 sv-RDBK1-nsg-quorum	Network security group		3 days	ago		Troubleshoot
	Recent support requests						,

Figure 6-57 Create 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 6-58 shows how to create a support request from the Resource Menu.

= Microsoft Azure	Search resources, services, and docs (G+/)	] [	3 (	<u>р</u> 8	ø	?	ন্দ	
Home > Resource groups > RI	DBK-SV4PC_Azure_Redpaper1 > sv-RDBK1-node2-vm							
? sv-RDBK1-node	<b>22-vm</b>   Support + Troubleshooting …							×
✓ Search (Ctrl+/)								Î
Automation								
🖧 Tasks (preview)	How can we help you?							
😫 Export template	Briefly describe your issue							
Help								
℅ Resource health	Service health							
Boot diagnostics	No Azure service issue detected. View service health							
Reformance diagnostics	Solutions to common problems							
📍 Reset password	Check out targeted diagnostics and troubleshoot common problems.							
🔎 Redeploy + reapply	Troubleshoot							
Nerial console								
📮 Connection troubleshoot	Need more help?							
🧠 Red Hat Customer Portal	Visit the help and support center to create and manage support requests and support plan	5.						
③ Support + Troubleshooting	Create a support request							

Figure 6-58 Create a support request from Resource Menu for resource type VM

Support requests that are created from the Resource Menu feature the issue type Technical preselected (see Figure 6-59).

≡ Microsoft Azure	$\mathcal{P}$ Search resources, services, and docs (G+/)	$\mathbf{\Sigma}$	Ģ	Q	٢	?	ନ୍ଦ	
Home > Resource groups >	RDBK-SV4PC_Azure_Redpaper1 > sv-RDBK1-node2-vm >							
New support rec	quest ····							×
								,
Basics Solutions Deta	ails Review + create							
Create a new support request issues. Complete the Basics tab by se help to solve your issues faste	to get assistance with billing, subscription, technical (including advisory) or quota management electing the options that best describe your problem. Providing detailed, accurate information can r.							
Summary *	Describe your issue							
Issue type *	Technical							
Subscription *	Microsoft Azure Enterprise_ikky (							
	Can't find your subscription? Show more $\mathbb O$							
Service	My services      All services							
Service type *	Virtual Machine running Linux							
Resource *	sv-RDBK1-node2-vm							
Next: Solutions >>								

Figure 6-59 New support request with preselected Issue type

Provide all requested information, such as the summary, problem type, and problem subtypes. The requested information depends on the specified service type.

Define the severity by using one of the following categories, the preferred contact method, and contact information (see Figure 6-60):

- Severity A Critical Impact: Business has significant loss or degradation
- Severity B Moderate Impact: Business can continue in an impaired manner
- Severity C Minimal impact: Business is functioning with minor impediments

Microsoft Azure	₽ Search resources, se	ervices, and docs (G+/)	
Home > sv-RDBK1-node2-vm >			
New support request	t		×
Support method			
Support plan	Azure Support Plan - Standard		
Severity *	C - Minimal impact	~	
Preferred contact method *	File Live chat     Get help new from a Support engineer via chat		
	🔿 🖾 Email		
	A Support engineer will contact you over email.		
	A Support engineer will contact you over the phone.		
Your availability	Business Hours		
Support language *	English	V	
Contact info			
First name *	Tela		
Last name *	104		
Email *	Instantige and the second		
Additional email for notification			
Phone			
Country/region *	Choose a country	Y	
Save contact changes for future	support requests.		
<< Previous: Solutions	Next: Review + create >>		

Figure 6-60 Enter details to Microsoft support request

#### Managing a support request

To change the severity, update the ticket, or upload files, select the support request from the Home view and then, select **Help + support** and **All support requests** (see Figure 6-61).

Home > sv-RDBK1-node2-vm > Help + support         Help + support   All support requests       **
Help + support   All support requests       ···       ×
Search (Ctrl+/)
Support     1 Subscriptions     Past 7 days     Open       all support requests     Can't find your subscription? Show more O
All support requests     Can't find your subscription? Show more      Can't find your subscription? Show more
Support Plans
P Filter results
Service Health Title ↑↓ ID ↑↓ Created ↑↓ Subscription ↑↓ Resource type ↑↓ Updated Status
🦂 Advisor TEST ticket 2112130050000591 Mon, Dec 13, 2021, 1 Microsoft Azure Enterprise_ikky Virtual Machine running Linux 1 day ago \cdots Open
<

Figure 6-61 Display support request in Microsoft Azure

Figure 6-62 shows details of an open support ticket.

■ Microsoft Azure	S & Q & Z
Home > Help + support >	~
A support request. Test licket	^
Status Severity	Updates
Open C _{Change}	Enter a new message
Title TEST ticket	
Support request 2112130050000591 ID	
Created on Mon, Dec 13, 2021, 10:25:57 AM	
Created by	
Support plan Azure Support Plan – Standard	
Advanced Permission not granted Change diagnostic information	Submit Cancel
File upload ① Select a file Upload	A Mon, Dec 13, 2021, 11:58:15 AM
	Hello 🖬 🖉
✓ Case details	Thank you for contacting Microsoft Azure laaS VM team. My name is Una and I'm reaching out to you regarding this ticket.
V Contact info	My understanding is that this is just a test ticket. So I will archive this ticket.
	However, if you have any other questions or concerns, please don't hesitate to let me

Figure 6-62 Update or change Microsoft Azure support requests

160 Implementation Guide for IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Version 8.4.3

# **Related publications**

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

## **IBM Redbooks**

The following IBM Redbooks publications provide more information about the topic in this document. Note that some publications that are referenced in this list might be available in softcopy only:

- ► IBM Spectrum Virtualize for Public Cloud on AWS Implementation Guide, REDP-5534
- IBM FlashSystem Best Practices and Performance Guidelines for IBM Spectrum Virtualize Version 8.4.2, SG24-8508
- Implementing the IBM FlashSystem with IBM Spectrum Virtualize Version 8.4.2, SG24-8506
- Implementing the IBM SAN Volume Controller with IBM Spectrum Virtualize Version 8.4.2, SG24-8507
- IBM SAN Volume Controller Best Practices and Performance Guidelines for IBM Spectrum Virtualize Version 8.4.2, SG24-8509
- Introduction and Implementation of Data Reduction Pools and Deduplication, SG24-8430

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

ibm.com/redbooks

## **Online resources**

The IBM Passport Advantage website also is relevant as a further information source:

http://www-01.ibm.com/software/passportadvantage/

# **Help from IBM**

IBM Support and downloads ibm.com/support

**IBM Global Services** 

ibm.com/services

162 Implementation Guide for IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Version 8.4.3

(0.2"spine) 0.17"<->0.473" 90<->249 pages Implementation Guide for IBM Spectrum Virtualize for Public Cloud on Microsoft Azure Version 8.4.3



SG24-8510-00

ISBN 0738460451

Printed in U.S.A.



Get connected

