

IBM Cloud Object Storage System Product Guide

For on-premises IBM Cloud Object Storage installations

Vasfi Gucer

Chris de Almeida

Joe Dorio

Israel Feygelman

Max Huber

Michael Knieriemen

Lars Lauber

Jussi Lehtinen

Jaswinder Singh Saini



Storage



IBM Redbooks

IBM Cloud Object Storage System Product Guide

June 2023

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

Fourth Edition (June 2023)

This edition applies to IBM Cloud Object Storage System Version 3.17.2.

© Copyright International Business Machines Corporation 2019, 2023. All rights reserved.

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

Contents

Notices	vii
Trademarks	viii
Preface	ix
Authors	ix
Now you can become a published author, too	x
Comments welcome	xi
Stay connected to IBM Redbooks	xi
Summary of changes	xiii
June 2023, Fourth Edition	xiii
Chapter 1. IBM Cloud Object Storage System overview	1
1.1 Introduction	2
1.1.1 Key concepts and terminology	2
1.2 How IBM Cloud Object Storage System works	4
1.2.1 Dispersed storage-defined solution	4
1.2.2 Information Dispersal Algorithm	4
1.2.3 Security	5
1.2.4 Access method	5
1.3 Software-defined storage	6
1.3.1 Software-defined storage in storage-defined architecture framework	6
1.3.2 IBM Cloud Object Storage for software-defined storage	8
1.4 Typical workloads and use cases	9
1.4.1 Use case 1: Backup repository	10
1.4.2 Use case 2: Internet of Things	13
1.4.3 Use case 3: Analytics and cognitive systems	15
1.4.4 Use case 4: Active archive	17
1.4.5 Use case 5: Enterprise file services	19
1.4.6 Use case 6: Content repository	21
Chapter 2. Planning and sizing an IBM Cloud Object Storage System	23
2.1 Planning for capacity	24
2.1.1 Initial capacity requirement	24
2.1.2 Alternative method to plan for capacity for Standard Dispersal mode	30
2.1.3 Incremental capacity requirements	33
2.1.4 Summary	36
2.2 Performance planning	37
2.2.1 Accesser node layer performance for a single-site system	38
2.2.2 Accesser node layer performance for a multisite system	39
2.2.3 Accesser node layer performance for a two-site mirrored system	39
2.2.4 Slicestor node layer performance	40
2.2.5 IDA effect on performance	41
2.2.6 Storage engine choice	41
2.2.7 Network performance	42
2.2.8 Measuring performance	42
2.3 Planning for high reliability and availability	44
2.3.1 IDA selection	44
2.3.2 Best practices for data center planning	45

2.3.3 Multiple Manager devices	46
2.4 Network planning	46
2.4.1 Multi-networks with IBM Cloud Object Storage	46
2.4.2 Network Time Protocol	48
2.4.3 Load balancers	49
2.4.4 Firewalls	50
2.4.5 Differences between S3 and IBM Cloud Object Storage APIs	51
Chapter 3. IBM Cloud Object Storage Gen2 hardware appliances	53
3.1 Gen2 hardware appliance overview	54
3.1.1 Highlights	54
3.2 Appliance overview	55
3.2.1 Manager appliance	57
3.2.2 Accesser appliance	58
3.2.3 Slicestor appliances	59
3.3 Appliance specifications	71
3.4 Hardware options	73
3.4.1 Processor and memory upgrade	73
3.4.2 Network interface upgrade	74
3.5 Performance	74
3.5.1 Accesser performance	74
3.5.2 Slicestor performance	74
3.6 Rack guidance	75
3.6.1 Appliance weight	75
3.6.2 Internal dimensions	76
3.6.3 Power and PDU placement	77
Chapter 4. Deployment options	79
4.1 Introduction	80
4.2 IBM hardware appliances	80
4.3 Third-party appliances	81
4.4 Embedded Accesser	81
4.4.1 Enabling Embedded Accesser functions	81
4.5 IBM Cloud Object Storage System virtual appliances	82
4.5.1 Configuring the appliance environment	82
4.6 Appliance Docker Containers	83
4.6.1 Benefits	84
4.6.2 Workflow, use cases, and feature impact	84
4.6.3 Accesser container	85
4.6.4 Manager container	85
4.6.5 System and network configuration	86
4.6.6 Configuring the appliance container	86
4.6.7 Deployment	87
Chapter 5. Initial setup and configuration	93
5.1 Needed installation information	94
5.1.1 Required information	94
5.1.2 Optional information	94
5.2 Example information	95
5.3 Step 1: Installing the solution	96
5.3.1 Physical appliance	96
5.3.2 Virtual appliance	96
5.3.3 Container appliance	97
5.4 Step 2: Installing IBM Cloud Object Storage Appliance software	97

5.5 Step 3: Configuring the appliance	101
5.5.1 Configuring the Manager	101
5.5.2 Configuring the Accesser appliance	102
5.5.3 Configuring the Slicestor appliance.	104
5.6 Step 4: Configuring the Manager GUI	106
5.6.1 Initial login	106
5.6.2 Creating a system	109
5.6.3 Creating a site	110
5.6.4 Accept pending devices	110
5.6.5 Creating a storage pool.	114
5.6.6 Creating a vault.	116
5.6.7 Creating an access pool	118
5.6.8 Configuring HTTPS certificates for access pools	119
5.6.9 Enabling Access Key Authentication	120
5.6.10 Creating a user	121
5.6.11 Generating Access Key ID	123
5.6.12 Granting CLI Access	123
5.6.13 Manager configuration backup	125
5.6.14 Organizations	127
5.6.15 Notification Service	130
5.6.16 Vault Index Version.	132
5.6.17 Vault Deletion Authorization	134
5.6.18 Storage Account Portal.	136
5.7 Step 5: Verifying the solution	139
5.7.1 Programs to verify and test IBM Cloud Object Storage	139
5.7.2 Configuring AWS CLI	139
5.7.3 Uploading an object	140
5.7.4 Listing objects	140
5.7.5 Downloading an object	140
5.7.6 Deleting an object	140
5.7.7 Differences between S3 and the IBM Cloud Object Storage System APIs.	141
5.7.8 For more information.	142
5.8 Basic installation troubleshooting	142
5.8.1 Networking issues.	142
5.8.2 Installing IBM Cloud Object Storage	143
5.8.3 Pending appliances.	144
5.8.4 S3 API issues	144
5.9 IBM Call Home and log collection	144
5.9.1 Configuring IBM Call Home.	144
5.9.2 Log collection	146
5.10 Upgrading your IBM Cloud Object Storage instance	150
5.10.1 Upgrade procedure considerations.	151
5.10.2 Upgrading the Manager	152
5.10.3 Upgrading IBM Cloud Object Storage devices	153
5.11 Configuring multiple Manager devices	154
5.11.1 Adding a second Manager device.	154
5.11.2 Multiple Manager teardown.	156
5.12 Multi-factor authentication	158
5.13 Enabling Write Once Read Many capabilities	162
5.13.1 IBM retention vaults	162
5.13.2 S3 Object Lock	164
Chapter 6. Scalability	167

6.1	Scaling an IBM Cloud Object Storage System	168
6.1.1	Non-disruptive upgrade	168
6.2	Scaling for performance	168
6.2.1	Adding Accesser nodes	169
6.2.2	Removing Accesser nodes	170
6.2.3	Automating performance with Docker and Kubernetes	171
6.3	Scaling for capacity	171
6.3.1	Adding a device set to a storage pool	172
6.3.2	Replacing a device set	178
6.3.3	Removing a device set	180
6.3.4	Adding a storage pool	183
6.3.5	Planning for scalability	184
Chapter 7. IBM Cloud Object Storage System File Access		185
7.1	Introduction	186
7.2	Features	186
7.3	IBM Cloud Object Storage example use case	187
7.4	IBM Cloud Object Storage File Access deployment architecture	188
7.4.1	IBM Cloud Object Storage File Access Portal Application Server	188
7.4.2	Database Server	188
7.4.3	IBM Cloud Object Storage File Access Gateway	188
7.5	Conclusion	189
Abbreviations and acronyms		191
Related publications		193
	IBM Redbooks	193
	Online resources	193
	Help from IBM	193

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](http://www.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.

Accesser®
Cleversafe®
IBM®
IBM Cloud®

IBM Spectrum®
Insight®
PartnerWorld®
Redbooks®

Redbooks (logo) ®
Slicestor®

The following terms are trademarks of other companies:

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

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.

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

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

VMware, VMware vSphere, 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

Object storage is the primary storage solution that is used in the cloud and on-premises solutions as a central storage platform for unstructured data. IBM Cloud Object Storage is a software-defined storage (SDS) platform that breaks down barriers for storing massive amounts of data by optimizing the placement of data on commodity x86 servers across the enterprise.

This IBM Redbooks® publication describes the major features, use case scenarios, deployment options, configuration details, initial customization, performance, and scalability considerations of IBM Cloud Object Storage on-premises offering. For more information about the IBM Cloud Object Storage architecture and technology that is behind the product, see *IBM Cloud Object Storage Concepts and Architecture*, REDP-5537.

The target audience for this publication is IBM Cloud Object Storage IT specialists and storage administrators.

Authors

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

Vasfi Gucer is a project leader with the IBM Systems WW 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 an IBM Certified Senior IT Specialist, Project Management Professional (PMP), IT Infrastructure Library (ITIL) V2 Manager, and ITIL V3 Expert.

Chris de Almeida is a Software Architect at IBM Cloud Object Storage, where he works primarily on the Manager software, which is used for configuring and monitoring dispersed storage networks. He is involved in web standards at TC39, the standards body for JavaScript, W3C, the standards body for CSS and the web, and WHATWG, the standards body for HTML. He is an open source contributor and maintainer, most notably of Afterburner.js, a testing framework.

Joe Dorio is an IBM Sr. Storage Brand Technical Specialist, with a focus in IBM Cloud Object Storage, and several other Storage Software solutions. Joe came to IBM through the IBM Cleversafe® acquisition and has over 20 years of experience in the Storage and Data Protection environments. Joe supports IBM account teams selling Software systems in the Northeast National Market segment.

Israel Feygelman is a Senior Storage Brand Technical Specialist focusing on storage software and object storage solutions. Israel came to IBM after working for over 8 years as a systems engineer for data protection and storage systems for several clients. He supports the Northeast National Markets.

Max Huber is a Storage Technical Specialist with a focus on the Data and AI SDS portfolio at IBM Germany. He has 3 years of experience working with IBM Cloud Object Storage, and before that was working for 5 years in the web and Internet of Things (IoT) Application development area. He holds a master's degree in Information Systems and Management from Munich University of Applied Sciences.

Michael Knieriemen is a Program Manager with IBM Cloud® in Chicago, Illinois, US. He has more than 25 years of experience in enterprise systems management, software sales, training, and cloud computing. He leads a team of technical Cloud Object Storage consultants and architects supporting the sale of IBM Cloud Object Storage solutions and training for IBM Cloud Object Storage clients. As an IBM Redbooks author, certification and digital learner developer, and an IBM-recognized teacher and educator, he enjoys sharing knowledge with others. He holds a Bachelor of Science degree in Management from Purdue University, Indiana, US.

Lars Lauber is a Storage Technical Specialist who is focused on unstructured data in SDS environments. Lars has planned and deployed many object storage systems, mainly in Central Europe. He holds a master's degree in Technology and Innovation Management from FOM University of Applied Sciences for Economics and Management in Stuttgart.

Jussi Lehtinen is a Solutions Architect for IBM File and Object Storage working for IBM Systems in Europe. He has over 30 years of experience working in IT, with the last 20 years with Storage. He holds a bachelor's degree in Management and Computer Studies from Webster University in Geneva, Switzerland.

Jaswinder Singh Saini is a Cloud Architect for IBM Cloud for VMware Solutions. In his current role, he is involved in designing, implementing and documenting new cloud solutions. He has over 18 years of experience working in IT. In his previous role, he has been involved in designing, implementation, POC, of IBM Systems and Storage.

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

- ▶ Erica Wazewski
IBM Redbooks, Poughkeepsie Center
- ▶ Bill Mckenna, Doug Bohrer, John K Butler, John Shubeck, Matt Houghton, Steven Keller, Timothy Ranttila
IBM US

Thanks to the authors of the previous editions of this paper.

- ▶ Authors of the second edition, *IBM Cloud Object Storage Concepts and Architecture*, which was published on October 2021, were: Alexander Gavrin, Bradley Leonard, Hao Jia Johan Verstrepen, Jussi Lehtinen, Lars Lauber, Patrik Jelinko, Raj Shah, Steven Pratt.

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:

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:

<http://www.redbooks.ibm.com>

- ▶ Send your comments in an email:

redbooks@us.ibm.com

- ▶ Mail your comments:

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>

Summary of changes

This section describes the technical changes that are made in this edition of this IBM Redbooks publication and in previous editions. This edition might also include minor corrections and editorial changes that are not identified.

This book is the third edition of the book *IBM Cloud Object Storage System Product Guide*, SG24-8439, which was originally published on July 15, 2019. The new information that is included in this revision is described next.

June 2023, Fourth Edition

This revision includes the following new and changed information.

New information

- ▶ Multiple manager support
- ▶ S3 Object Lock
- ▶ Expanded section on the backup use case
- ▶ Multi-factor authentication (MFA)
- ▶ Gen2 hardware enhancements
- ▶ IBM Cloud Object Storage business continuity and disaster recovery (BCDR) manager database backup to cloud bucket

This edition applies to IBM Cloud Object Storage System 3.17.2.



IBM Cloud Object Storage System overview

This chapter provides an overview of IBM Cloud Object Storage System.

This chapter includes the following topics:

- ▶ 1.1, “Introduction” on page 2
- ▶ 1.2, “How IBM Cloud Object Storage System works” on page 4
- ▶ 1.3, “Software-defined storage” on page 6
- ▶ 1.4, “Typical workloads and use cases” on page 9

1.1 Introduction

This section provides an introduction to the IBM Cloud Object Storage System.

The IBM Cloud Object Storage System is a breakthrough cloud platform that helps solve petabyte and beyond storage challenges for enterprises worldwide. It uses an innovative and cost-effective approach for storing large volumes of unstructured data while still ensuring scalability, security, availability, reliability, manageability, and flexibility.

Consider the following points:

- ▶ Scalability offers a single storage system and namespace versus an ever-increasing number of limited-capacity storage silos.
- ▶ Security features include a wide range of capabilities that help meet security requirements.
- ▶ Availability and reliability characteristics of the system are configurable to best suit different use cases and requirements.
- ▶ Manageability helps storage administrators to handle large storage capacity.
- ▶ The flexibility of a software-defined storage (SDS) solution does not require specific or proprietary hardware.

Exponential growth of unstructured data: To understand the value of IBM Cloud Object Storage, you must understand the growth of unstructured data. Today, 80% of all data is unstructured, with 90% or more of that data in object storage in the cloud. On-premises, the IDC estimates 87% of file and object data will be in object storage by 2020. Object storage is also growing at a 35% compound annual growth rate (CAGR) from 2017 to 2022.^a

The increasing volume of data and presents a high demand to store securely more than 140 zettabytes (ZB) of unstructured data in a cost-effective and secure manner. IBM Cloud Object Storage is designed to meet this need with its core design tenets of scalability, fault-tolerance, and data protection. The system can manage over 1000 devices and store over 1 exabyte (EB) of data, with petabytes of storage capacity in a single namespace for big data analysis. With the system's configuration options, you can adjust the data protection level based on the data's criticality.

IBM Cloud Object Storage offers high data availability with the potential for more than fourteen 9s of data availability, and it can be configured, managed, and adjusted through a GUI or API. This approach minimizes day-to-day administration, freeing up time for other tasks.

a. <https://www.ibm.com/downloads/cas/NDLDMXKP>

1.1.1 Key concepts and terminology

This section describes the following the key concepts and terminology that are used in IBM Cloud Object Storage. Figure 1-1 on page 3 shows a logical view of an IBM Cloud Object Storage System.

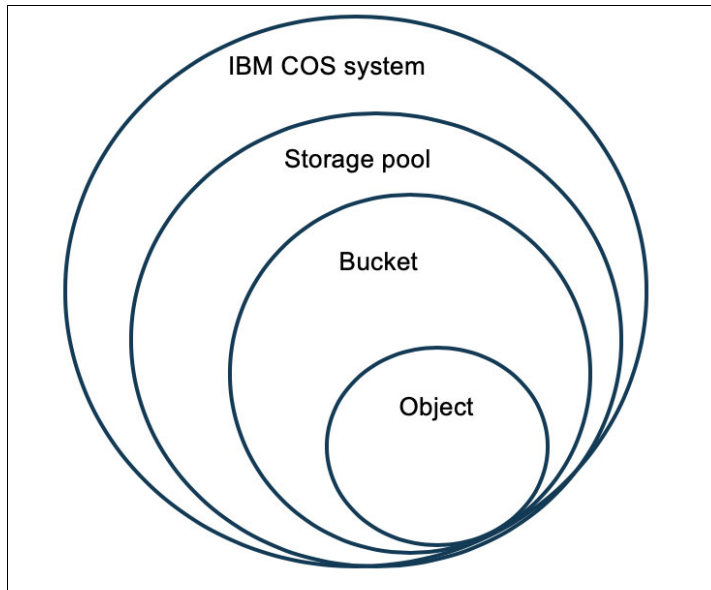


Figure 1-1 Logical concepts in IBM Cloud Object Storage

- ▶ **Object**
An *object* refers to user data that is uploaded to an IBM Cloud Object Storage System. Typically, it is a file and the object metadata that is stored together with the file.
- ▶ **Bucket**
A *bucket* refers to vault (also known as standard vault) in vault mode or a container in Container Mode in IBM Cloud Object Storage. Bucket is a logical abstraction that is used to store the data.
- ▶ **IBM Cloud Object Storage Manager node**
A system component that provides a management interface that is used for administrative tasks, such as system configuration, storage provisioning, and monitoring the health and performance of the system. A Manager node (also referred to as *Manager*) can be deployed as a physical appliance, VMware virtual machine, or Docker container.
- ▶ **IBM Cloud Object Storage Accesser node**
A system component that encrypts and encodes data on write, or decodes and decrypts data on read. It is a stateless component that manages the transformation of the data and presents the storage interfaces to the client applications. An IBM Accesser® node can be deployed as a physical appliance, VMware virtual machine, Docker container, or an embedded Accesser on an IBM Slicestor® appliance.
- ▶ **IBM Cloud Object Storage Slicestor node**
A system component that is responsible for storing the data. It receives data from the Accesser node on write and returns data to the Accesser node as required on read. Slicestor nodes are deployed as physical appliances.
- ▶ **Device set**
A *device set* is defined by a group of Slicestor devices. Device sets can be spread across one or multiple data centers.

- ▶ Storage pool

A *storage pool* is a logical grouping of one or more device sets that together provide the physical storage resources for one or more buckets.

- ▶ Access pool

An *access pool* is a logical grouping of one or more Accesser nodes that are used to access the data.

More information: For more information about IBM Cloud Object Storage concepts and terminology, see *IBM Cloud Object Storage Concepts and Architecture*, REDP-5537.

1.2 How IBM Cloud Object Storage System works

This section describes how the IBM Cloud Object Storage System uses a dispersed storage-defined solution to ensure availability, reliability, scalability, and security.

1.2.1 Dispersed storage-defined solution

Dispersed storage is a commercial-grade implementation of a data storage technology called *information dispersal*. Information dispersal uses *erasure code* as a means to create redundancy for transferring and storing data. An erasure code is a forward error correction (FEC) code that transforms a message of k symbols into a longer message with n symbols, such that the original message can be recovered from a subset of the n symbols.

A dispersed storage network consists of a collection of dispersed storage nodes. With the dispersed storage network, transmission and storage of data are inherently private and secure. No complete set of data for an object exists in any single storage node. Only a subset of nodes must be available to fully retrieve the data on the network.

A dispersed storage-defined solution provides the following benefits:

- ▶ Dispersed storage provides massive scalability with minimal administrative tasks. Systems can grow easily from terabytes to petabytes to exabytes and beyond.
- ▶ Dispersed storage maintains 100% data integrity. Data is accessible from anywhere, anytime. Data is always available with an architecture that can tolerate simultaneous failures.
- ▶ Dispersed storage ensures data confidentiality, even when multiple drives or servers are compromised. Data in motion and data at rest is encrypted to make it unrecognizable and inherently secure to minimize opportunities for security breaches.

1.2.2 Information Dispersal Algorithm

IBM Cloud Object Storage System uses the Information Dispersal Algorithm (IDA) to break objects into slices that are distributed by way of network connections to Slicestors. The original data follows a series of transformations of encryption, slicing, and erasure coding into slices. These slices are stored across multiple Slicestors by using a dispersal algorithm to attain a high degree of failure independence.

With dispersed storage, only a subset of slices is needed to retrieve the data. A dispersed storage system can tolerate appliance failures within a single site and across multiple sites.

In IBM Cloud Object Storage, IDA consists of three parameters at the bucket level that define the overall availability and reliability of the storage. *Width* is the total number of slices that is generated by erasure coding of a data segment. *Read threshold (RT)* is the number of slices that is required to read a segment. *Write threshold (WT)* is the number of slices that must be written to the Slicestor nodes for a successful write operation.

IBM Cloud Object Storage System also uses SmartWrite and SmartRead technology to optimize writes and reads of slices, which results in improved throughput and efficiency.

1.2.3 Security

The following security features are built in and available in IBM Cloud Object Storage System:

- ▶ Data security is achieved by SecureSlice technology that is used to guarantee confidentiality, integrity, and availability of data that is stored on the IBM Cloud Object Storage System. SecureSlice uses All-or-Nothing Transform (AONT) as an encryption method where information can be deciphered only if all the information is known.
- ▶ Network security is assured by the fact that all network traffic that is flowing within IBM Cloud Object Storage System, and the applications are encrypted by using Transport Level Security (TLS) with AES. Slicestors can be placed anywhere without requiring a complex firewall or VPN setup.
- ▶ Retention enabled bucket is supported to meet compliance requirements.

More information: For more information about the IBM Cloud Object Storage security features and standards compliance information, see *IBM Cloud Object Storage Concepts and Architecture*, REDP-5537.

1.2.4 Access method

With the underlying dispersed storage, IBM Cloud Object Storage System can be shared and is jointly accessible by the following access protocols:

- ▶ Object-based access method

An S3-compatible interface is accessed by way of an HTTP/REST API. Simple **PUT**, **GET**, **DELETE**, and **LIST** commands enable applications to access the data.

REST API access to storage offers the following advantages:

- Tolerates internet latency
- Provides for programmable storage
- Provides efficient global access to large amounts of data

- ▶ File-based access method

IBM Cloud Object Storage supports the traditional network-attached storage (NAS) protocols Server Message Block/Common Internet File System (SMB/CIFS) and Network File System (NFS) through integration with gateway appliances or virtual file systems. Users and storage administrators can easily transfer, access, and preserve data assets over standard file protocols.

More information: For more information about the IBM Cloud Object Storage System architecture, see *IBM Cloud Object Storage Concepts and Architecture*, REDP-5537.

1.3 Software-defined storage

This section describes the business context for SDS within a software-defined architecture framework, and how IBM Cloud Object Storage continues to transform SDS.

1.3.1 Software-defined storage in storage-defined architecture framework

A software-defined architecture is a framework that encompasses orchestrators and software-defined infrastructure, including SDS. Such a framework creates and implements optimized IT infrastructures that can help enterprises attain competitive advantage by delivering higher value and profitability through speed and efficiency in provisioning IT services (see Figure 1-2).

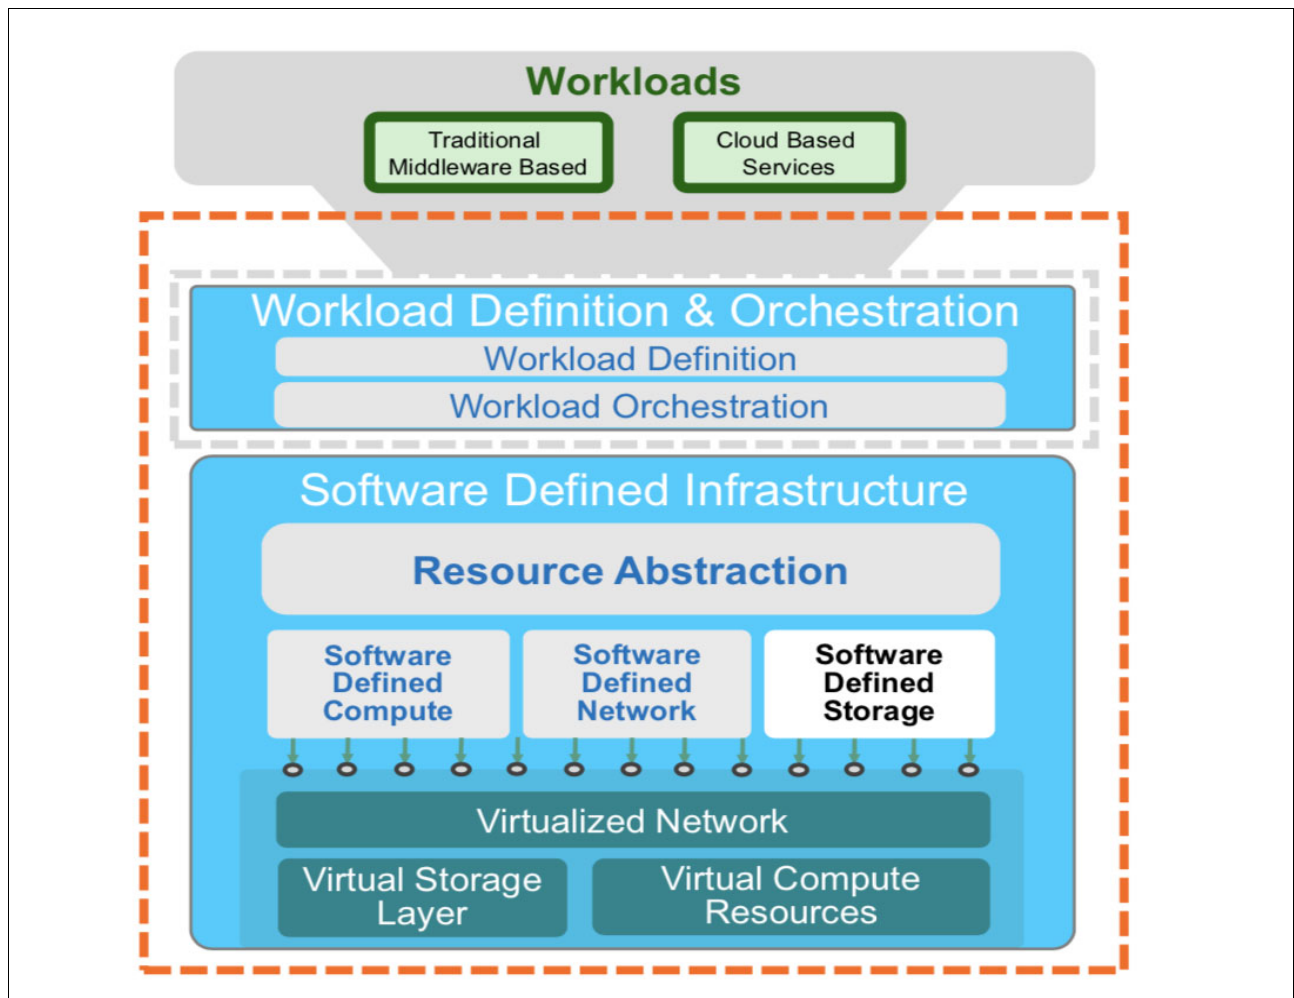


Figure 1-2 Software-defined architecture framework

SDS provides the following interrelated characteristics and benefits:

- ▶ Facilitates IT automation to improve business and IT agility at lower cost
- ▶ Optimizes systems administration and control to allow effective and efficient resource usage that lowers cost and supports business requirements
- ▶ Ease of deployment and redeployment of infrastructure resources

- ▶ Performance tuning with optimal alignment of available resources to application requirements
- ▶ Capacity planning simplification with larger storage pools that support multiple service levels
- ▶ Enables advanced application deployment of Systems of IBM Insight® by using Systems of Engagement (SoE) and Systems of Record (SoR):
 - Cognitive: Cloud
 - Analytics: Mobile
 - Social
 - Security
- ▶ Simplifies architecture to reduce specialized components and skills requirements
- ▶ Limitless elastic data scaling
- ▶ Supports block, file, and object data types

These business reasons for implementing SDS result from the combination of business pressures, new applications, data growth, and hybrid cloud that are challenging traditional storage approaches. A need exists to free data from hardware constraints and realign with new business processes and applications.

Figure 1-3 shows SDS in a software-defined infrastructure.

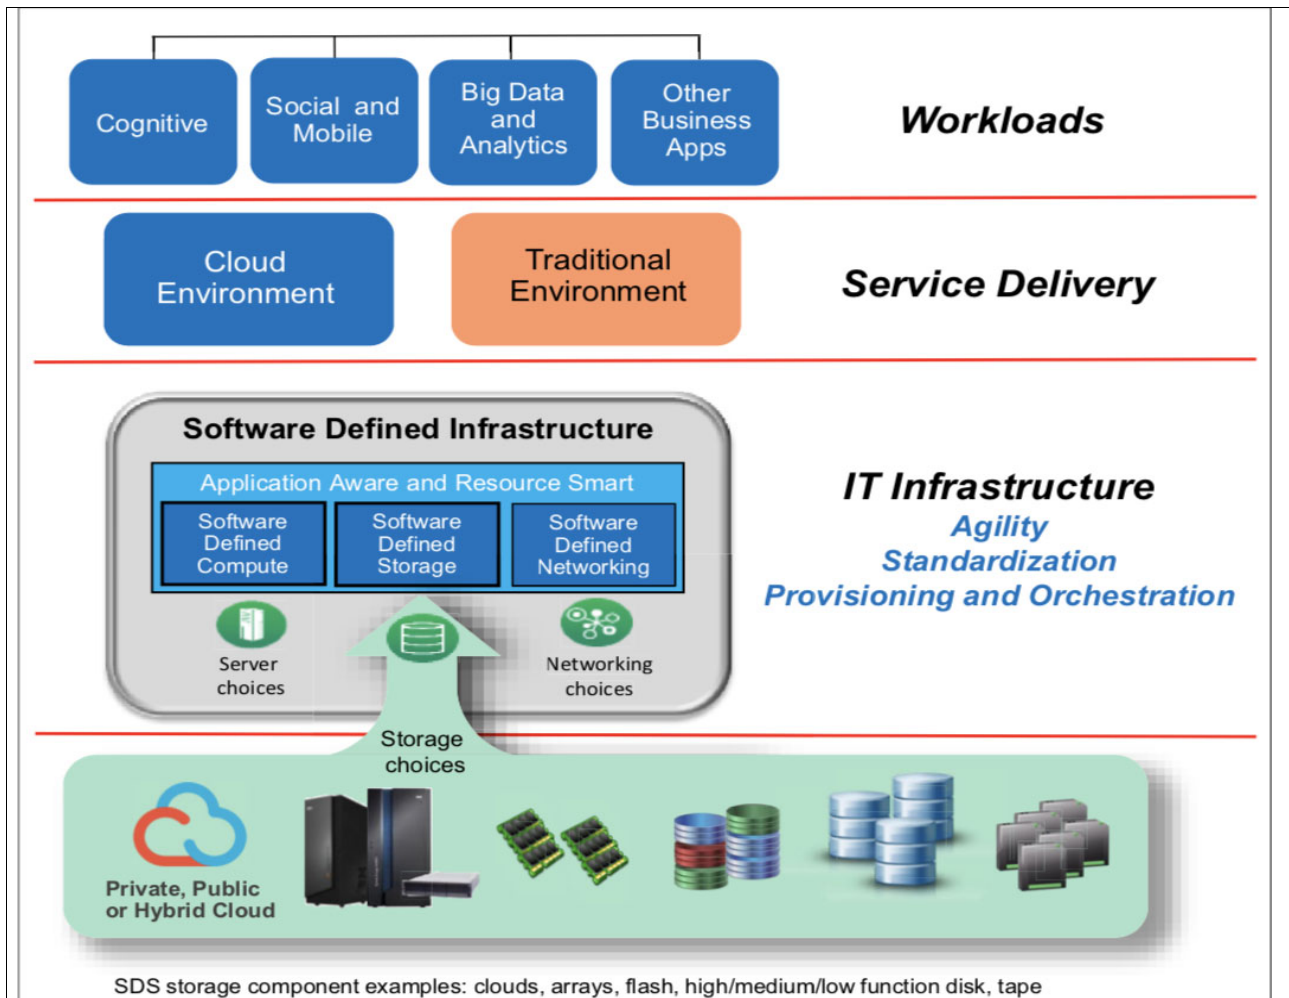


Figure 1-3 Software-defined storage in a software-defined infrastructure

1.3.2 IBM Cloud Object Storage for software-defined storage

IBM Cloud Object Storage continues to transform SDS by new software updates and a new hardware verification process.

New software update for software-defined modernization

IBM Cloud Object Storage System is a software-defined system that is also *hardware aware*. IBM Cloud Object Storage software supports newer models and hardware capabilities. This support enables the most advanced servers from various hardware vendors to run IBM Cloud Object Storage software. This support also prepares the way for supporting future technologies, such as NVMe (non-volatile memory express) and larger or more efficient storage servers and disk drives.

Figure 1-4 shows IBM Cloud Object Storage System components and appliances.

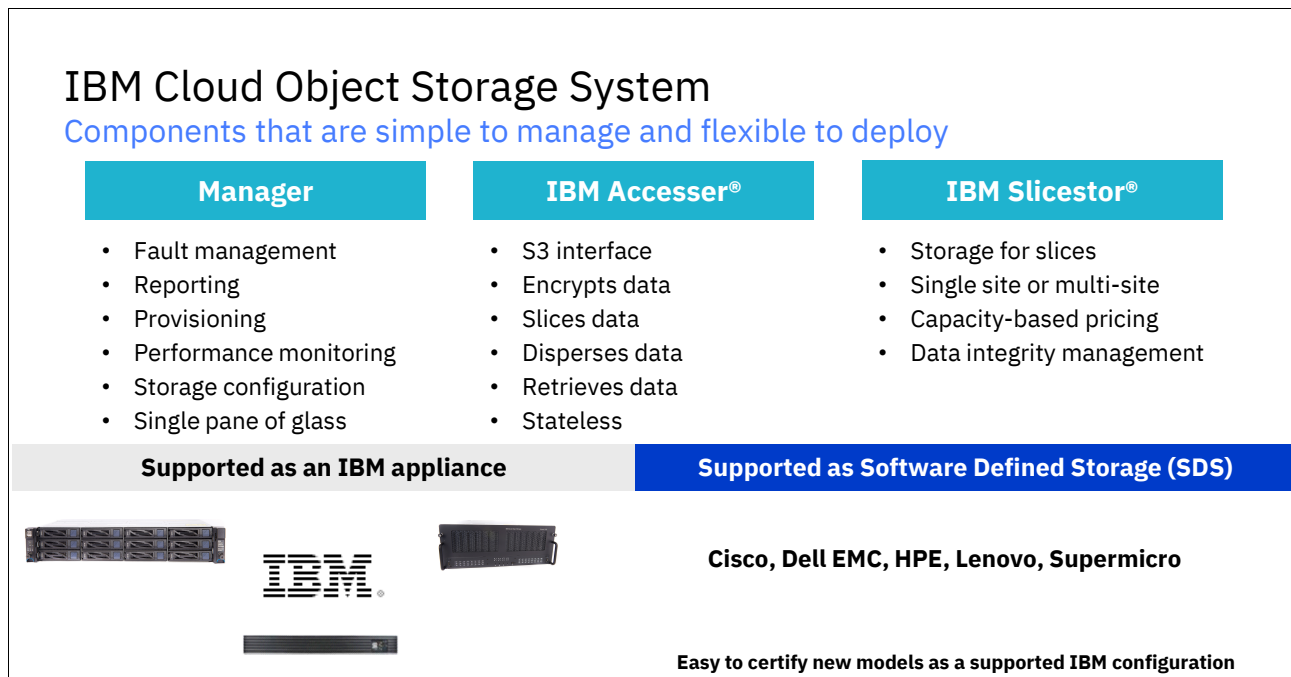


Figure 1-4 IBM Cloud Object Storage System components and appliances

New hardware verification process

IBM Cloud Object Storage supports over 80 different hardware configurations by using a selection of various models and vendors, including IBM, HPE, Seagate, Cisco, Dell, and Lenovo. With the new hardware verification process that is available now, an x86 server can be verified as a configuration within weeks by working directly with IBM and running a set of validation checks and sending the results to IBM. When verified, IBM adds the new server into the supported configuration and then the hardware can be used as a valid platform to use IBM Cloud Object Storage software.

1.4 Typical workloads and use cases

In this section, typical workloads and use cases for an IBM Cloud Object Storage based solution are described.

IBM Cloud Object Storage workloads and use cases can be grouped with the following broad categories, as shown in Figure 1-5:

- ▶ The first category targets known and mature workloads that were traditionally serviced by file, block, or tape storage systems.
- ▶ The second category targets the “new” workloads, which are newer applications and solutions that are more programmatically flexible and are more likely to directly access and derive insights from data and then expose them to users.

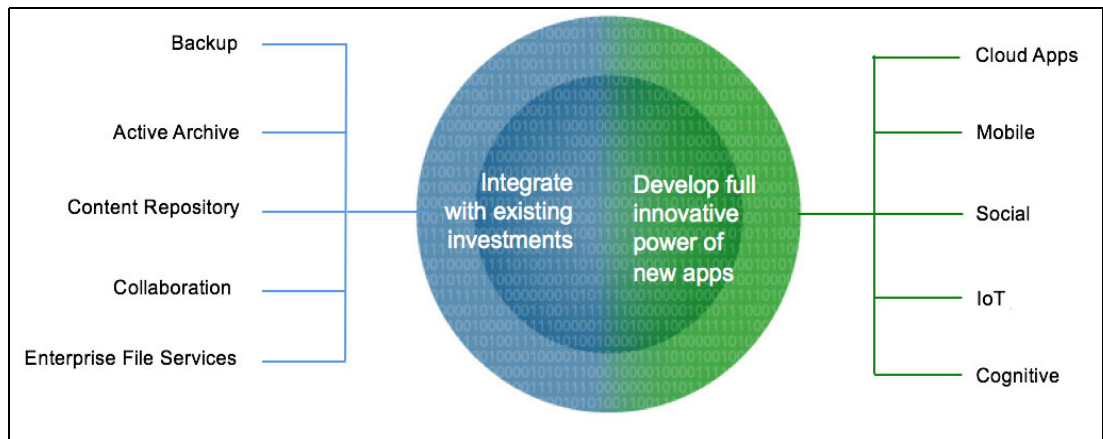


Figure 1-5 IBM Cloud Object Storage workloads

Key point: IBM Cloud Object Storage can service traditional workloads and new applications by using the same back-end service.

Figure 1-6 shows a high level of multiple use cases with multi-tenant access with IBM Cloud Object Storage.

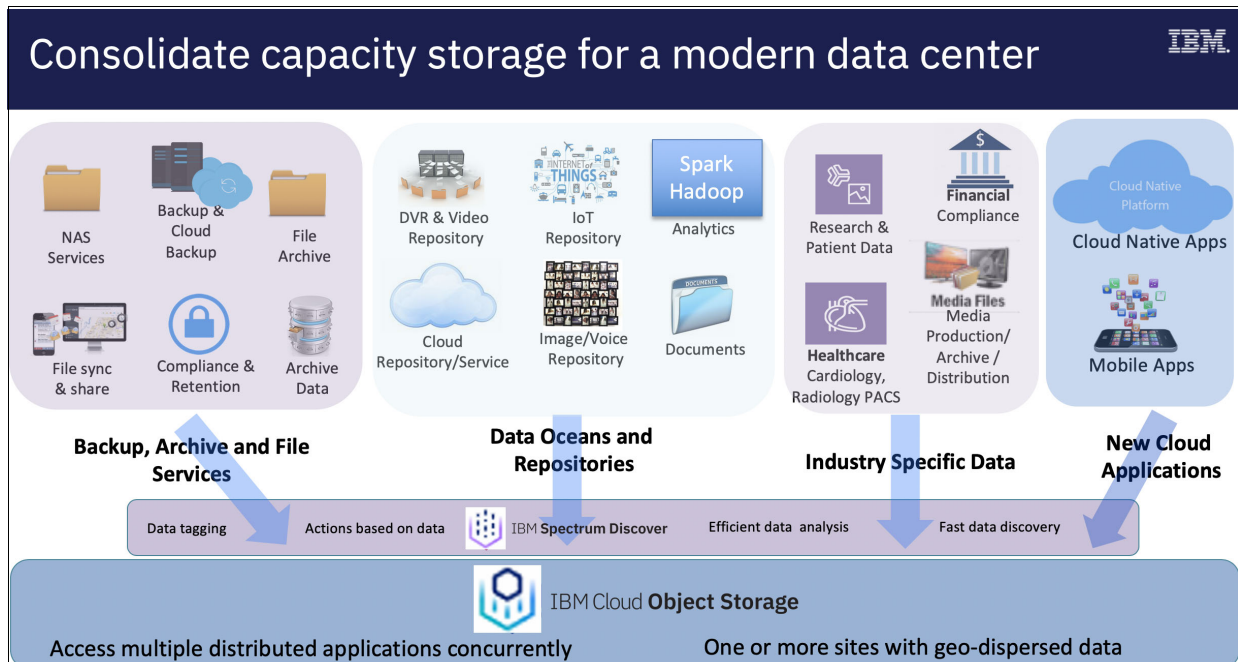


Figure 1-6 Multiple use cases with multi-tenant access with IBM Cloud Object Storage

Key point: IBM Cloud Object Storage is flexible in deployment with multiple use cases with multi-tenant access.

1.4.1 Use case 1: Backup repository

A *backup repository* is a storage system that is configured to store backup data. Traditionally, block storage, file storage, and tape systems are used for this purpose.

IBM Cloud Object Storage acts as a primary or a secondary backup repository tier. Some backup applications offer both possibilities for object storage, and some applications support only object storage as a secondary backup repository tier.

IBM Cloud Object Storage as a primary backup tier

IBM Cloud Object Storage as a primary backup repository tier means that backup data is written directly to IBM Cloud Object Storage and all restores happen directly from IBM Cloud Object Storage, as shown in Figure 1-7 on page 11.

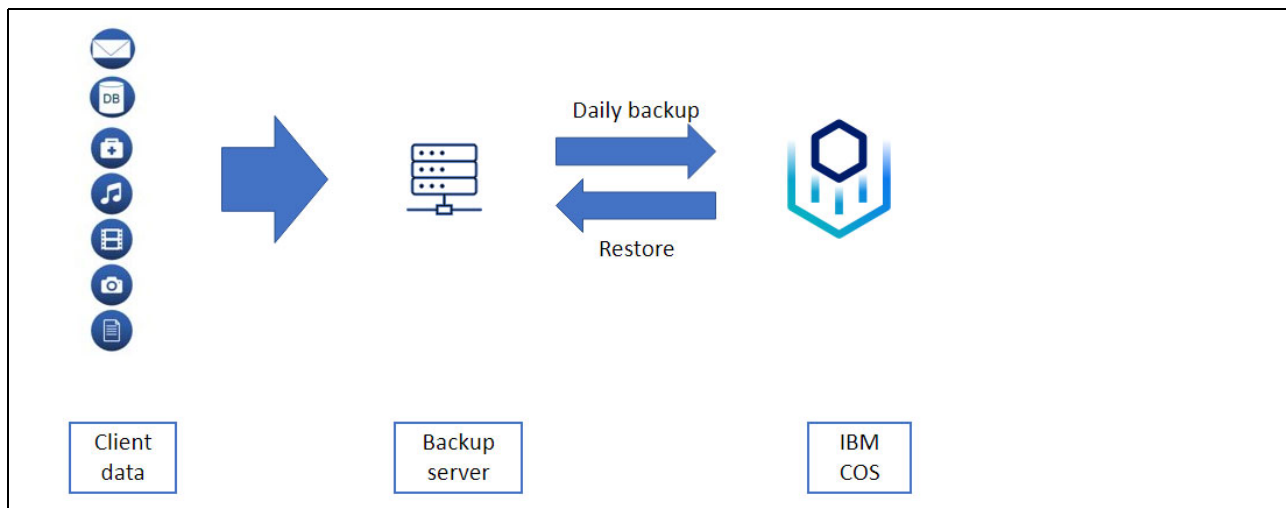


Figure 1-7 Showing IBM Cloud Object Storage deployed as a primary backup tier

There are pros and cons to this architecture, as shown in Figure 1-8.

- Pros
 - Simple design, only one backup tier to consider.
 - Lower cost than 2-tier backup repositories.
- Cons
 - Backup performance maybe lower than with traditional block or file storage system.
 - Restore performance can be significantly lower from object storage than what it would be from block or file storage.

Figure 1-8 Pros and cons of using IBM Cloud Object Storage as a primary backup tier

The main reason for lower restore performance is data deduplication. Many backup applications deduplicate backup data before sending it to object storage. The deduplication process creates small chunks of data that backup applications typically combine in larger objects to achieve high backup performance with object storage. The drawback of this method is that restoring often requires reading only a subset of deduplicated chunks from within an object, which lowers the restore performance. Deduplication also might require more object storage capacity if individual data chunks within single objects have different retention periods and the object storage system cannot free up the object's capacity until all its chunks expire.

IBM Cloud Object Storage as a secondary backup tier

IBM Cloud Object Storage as a secondary backup repository tier means that backup data is first written to block or file storage and then sent to IBM Cloud Object Storage for long-term retention. The primary tier provides fast restore capabilities from the latest backup copy. Data must be retrieved from IBM Cloud Object Storage only if it is not in the primary backup tier.

Figure 1-9 shows IBM Cloud Object Storage that is deployed as a secondary backup tier.

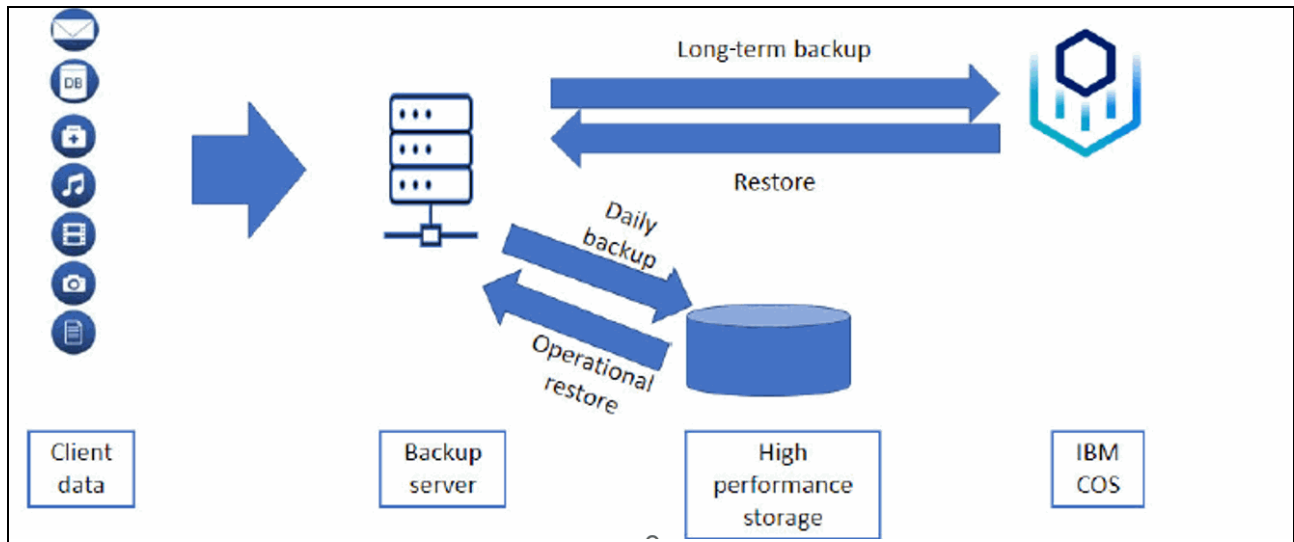


Figure 1-9 Showing IBM Cloud Object Storage deployed as a secondary backup tier

There are pros and cons to this architecture, as shown in Figure 1-10.

- Pros
 - Fast backup to primary tier
 - Fast restore from primary tier
- Cons
 - Higher cost
 - Two backup tiers to manage

Figure 1-10 Pros and cons of using IBM Cloud Object Storage as a secondary backup tier

Direct backup from an application server to IBM Cloud Object Storage

Some applications (such as databases) have a built-in backup capability that can do a direct backup from an application server to IBM Cloud Object Storage without a separate backup server, as shown in Figure 1-11 on page 13.

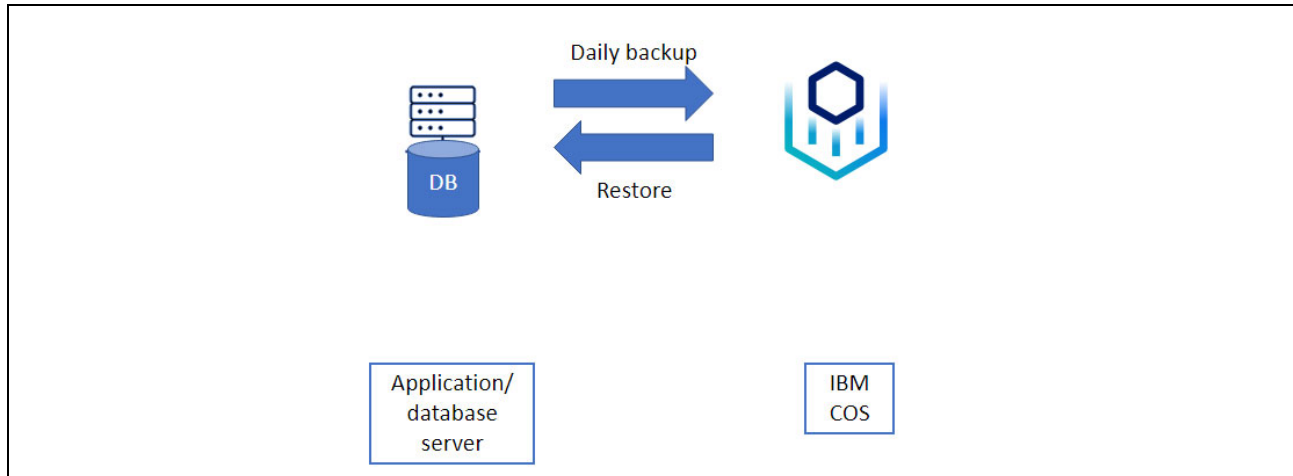


Figure 1-11 Showing IBM Cloud Object Storage deployed as a backup target directly from an application server

Ransomware protection

Ransomware is one of the biggest cybersecurity threats today. IBM Cloud Object Storage provides excellent protection against ransomware through its retention-enabled buckets and compatibility with S3 Object Lock, which makes data that is stored in protected buckets immutable for a period during which it cannot be overwritten. Most ransomware attacks encrypted backups first and then primary storage. By using retention-enabled buckets or S3 Object Lock in IBM Cloud Object Storage, you can prevent this attack, and data is kept safe from ransomware encryption efforts.

Why IBM Cloud Object Storage

IBM Cloud Object Storage is a perfect solution for a backup repository. Consider the following points:

- ▶ IBM Cloud Object Storage is a lower-cost storage solution compared to block or file storage.
- ▶ IBM Cloud Object Storage provides ransomware protection for your backup data through its retention-enabled buckets and S3 Object Lock compatibility.
- ▶ Most commercial backup applications work seamlessly with IBM Cloud Object Storage because of its compatibility with native S3 API.
- ▶ Adding capacity to IBM Cloud Object Storage is a simple online procedure that increases the backup storage repository size to whatever capacity is required.

1.4.2 Use case 2: Internet of Things

The *Internet of Things (IoT)* is a system of interrelated computing devices, mechanical and digital machines, objects, and people that are each provided with a unique identifier and the ability to transfer data over a network without requiring human-to-human or human-to-computer interaction. Different types of IoT device data are transformed through IoT platforms into long-term data storage, live dashboards, and real-time presentations

Figure 1-12 on page 14 shows a general workflow for an IoT system. Different devices, sensors, mobile, applications, web, and social networks produce different events that pass to the IoT platform. Data connection and secure transformation are processed through a connection and pipeline framework, such as REST API, MQTT, and Kafka, which transfers data through Spark to IBM Cloud Object Storage.

These customer-interaction-oriented models are often referred to as *SoE*, which are connected to back-end SoR where real-time data is stored, retrieved, and archived through IBM Cloud Object Storage. These solutions help users in various industries, homes, and cities to monitor activities and control devices by using the IoT platform.

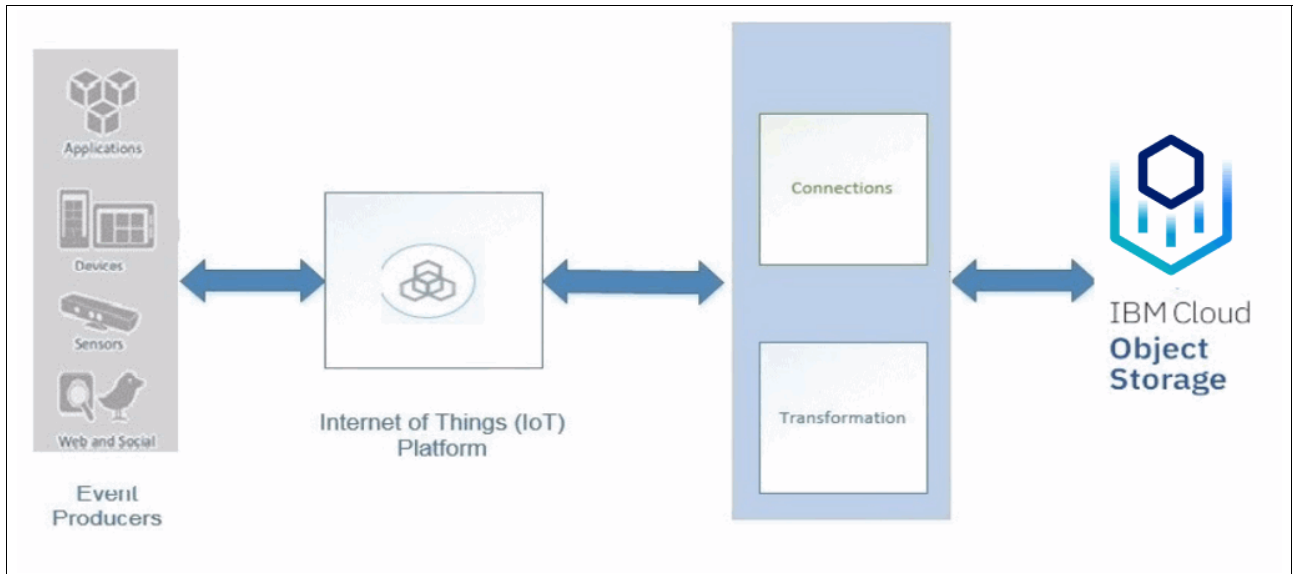


Figure 1-12 General workflow for an IoT system

Figure 1-13 shows an IoT use case that uses Spark as the front-end and IBM Cloud Object Storage as the back-end storage solution.

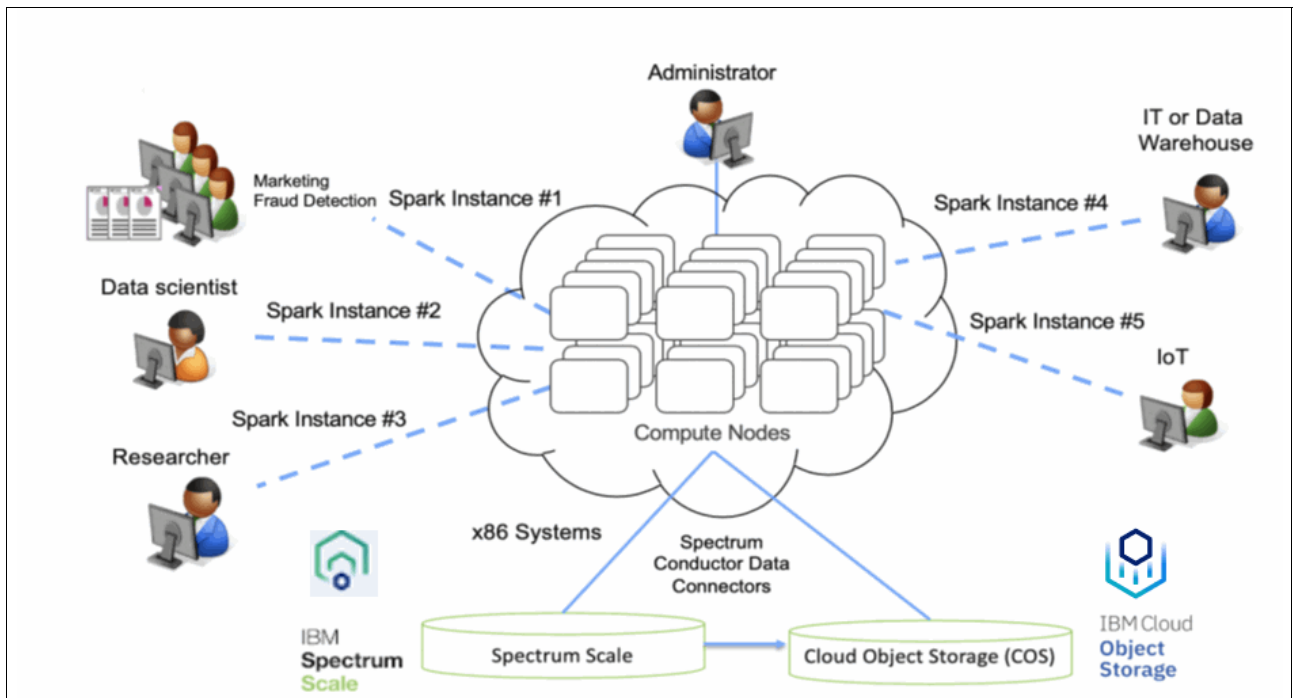


Figure 1-13 IoT use case with Spark and IBM Cloud Object Storage

Why IBM Cloud Object Storage

IBM Cloud Object Storage is well suited for IoT solutions for the following reasons:

- ▶ IBM Cloud Object Storage can store IoT content repository, real-time accessible data, large images, and video files.
- ▶ Data that is stored in IBM Cloud Object Storage is accessible anywhere, anytime for IoT devices, applications, and front-end consumers.
- ▶ IBM Cloud Object Storage works with on-premises and cloud data, visualizes data lineage, and feeds applications with data while integrating with IoT platforms.

1.4.3 Use case 3: Analytics and cognitive systems

Cognitive systems learn and “reason” based on interactions with humans and experiences with the environment. They do not give predetermined responses but make probabilistic prediction based on these experiences. Cognitive systems generate hypotheses, reasoned arguments, and recommendations, and can scale to keep pace with the complexity and unpredictability of information in the modern world.

Cognitive systems can “understand” unstructured information, such as the imagery, natural language, and sounds in books, email, tweets, journals, blogs, images, sound, and videos. They unlock meaning because they can reason through it and give us new contexts to weigh and consider. Cognitive systems also learn continually, honing our own expertise so that we can immediately take more informed actions.

As organizations embed data and analytics into every business process and every customer experience, they are finding that their IT infrastructure was not designed to drive performance for systems. Cognitive workloads require a reimagined IT infrastructure, one that can synthesize massive amounts of data quickly, accelerate analytics, and be available anytime, anywhere with trust.

A cognitive solution explains cognitive analytics, operations, and engagement, and how it interacts with customers, users, and ecosystem partners. These systems must store large quantities of unstructured data for the ongoing analysis that is required to build their understanding and cognitive capabilities.

A cognitive solution helps to build industry-leading solutions that address a faster time to market, a simple approach, and flexibility, while delivering accurate, evidence-based information. With analytics, you get key insights from data; with cognitive systems, you can turn those key insights into knowledge.

Figure 1-14 shows an analytics and cognitive system that uses IBM Cloud Object Storage as back-end storage being connected through Stocator.

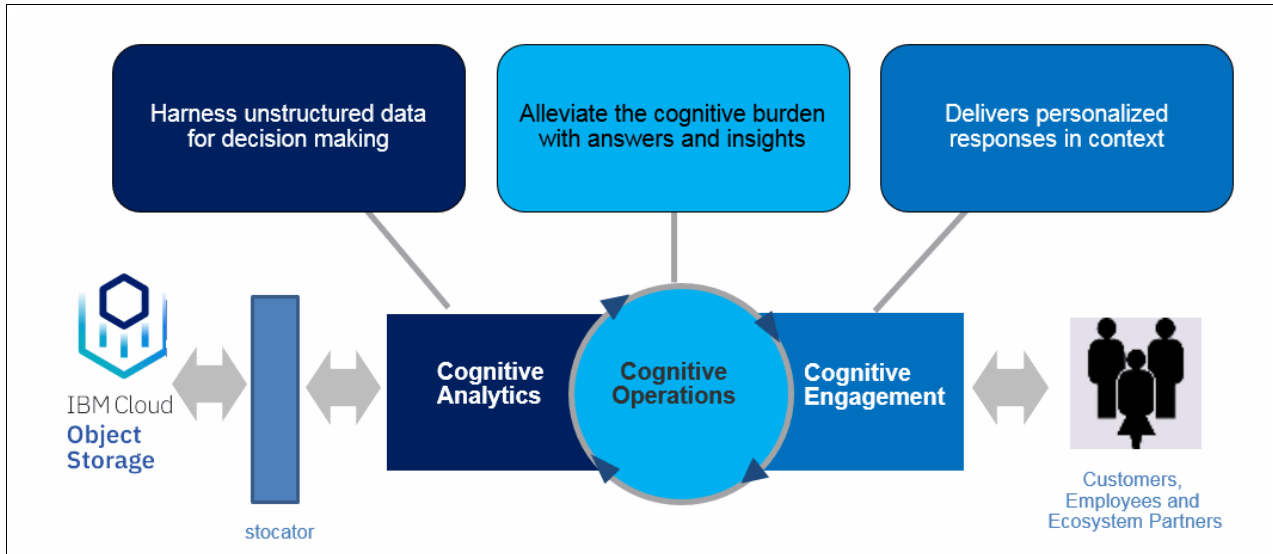


Figure 1-14 Analytics and cognitive system

Different applications evolved to store large volumes of unstructured data. More machine data is being created and must be stored and protected for further analysis and for potential cognitive uses. Software applications that create this type of data work with IBM Cloud Object Storage to deliver carrier-grade system reliability and data security, with the ability to scale to petabytes and larger in a cost-effective way. Cognitive systems store and process large volumes of data by using IBM Cloud Object Storage, and integrate with other analytics services.

More information: For more information about the integration between IBM Cloud Object Storage and other IBM storage products, see *Cloud Object Storage as a Service: IBM Cloud Object Storage from Theory to Practice - For developers, IT architects and IT specialists*, SG24-8385.

Why IBM Cloud Object Storage

IBM Cloud Object Storage is well suited for cognitive infrastructure solutions for the following reasons:

- ▶ IBM Cloud Object Storage works with on-premises and cloud data, with a cognitive solution that is deployed on a cognitive infrastructure.
- ▶ Data that is stored in IBM Cloud Object Storage is accessible anywhere and anytime for IoT devices, applications, and front-end consumers.
- ▶ IBM Cloud Object Storage features carrier-grade system reliability, availability, and scalability.

1.4.4 Use case 4: Active archive

Active archive is a data store that transparently stores important, but infrequently accessed, data, such as legal documentation, and finance documents. Because of these two seemingly conflicting requirements, this data store must be cost effective but resilient because data might be required at any time (often months or years) after intake. However, the data must be quickly accessible, unlike deep archive storage systems.

Figure 1-15 shows an older active archive system that is constructed of a mix of inexpensive disk systems and tape, with replication of data between sites being done by the active archive application or the underlying storage layer.

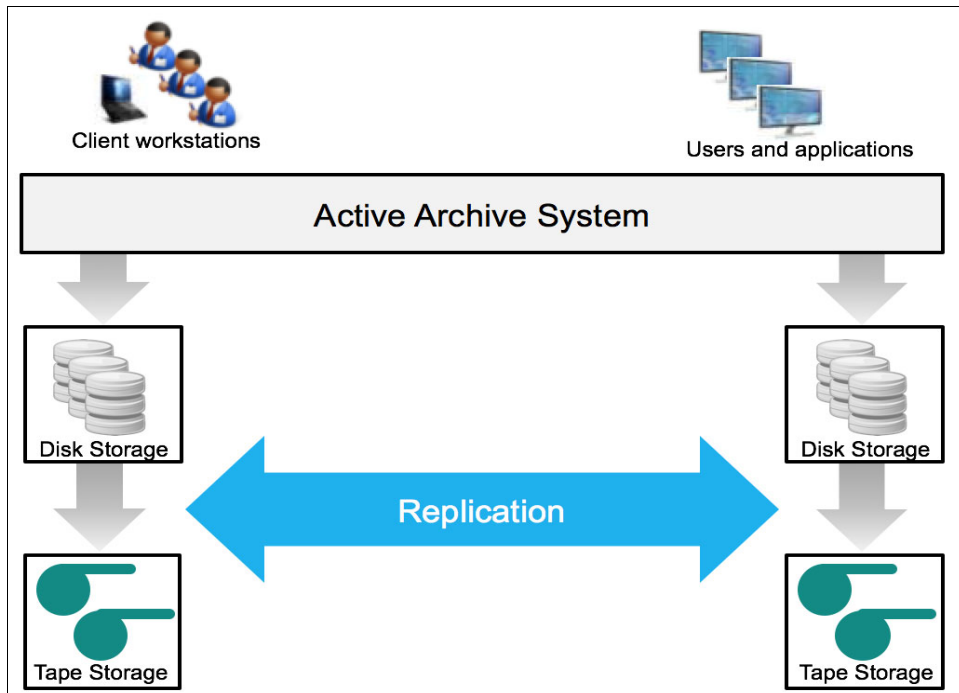


Figure 1-15 Older active archive system

The requirements for resilience, availability, and low cost are key to this use case, and typically mandate this replicated approach. The need for replication results in at least two data copies, even before backup copies are counted. Data replication must be carefully monitored because failure to do so reduces system resiliency.

Where tape is a part of the solution, architects must ensure that adequate disk staging space for initial intake and recall is available. Sufficient tape drives to service these components also must be included. In addition, the application must deal with extended recall times that a busy tape system might experience.

Figure 1-16 shows a solution that uses IBM Cloud Object Storage as the active archive tier where the following conditions exist:

- ▶ The data tiers are reduced or eliminated so that data can be directly pushed or retrieved from the IBM Cloud Object Storage repository.
- ▶ Site loss resilience is achieved through geographically dispersed erasure coding rather than replication.

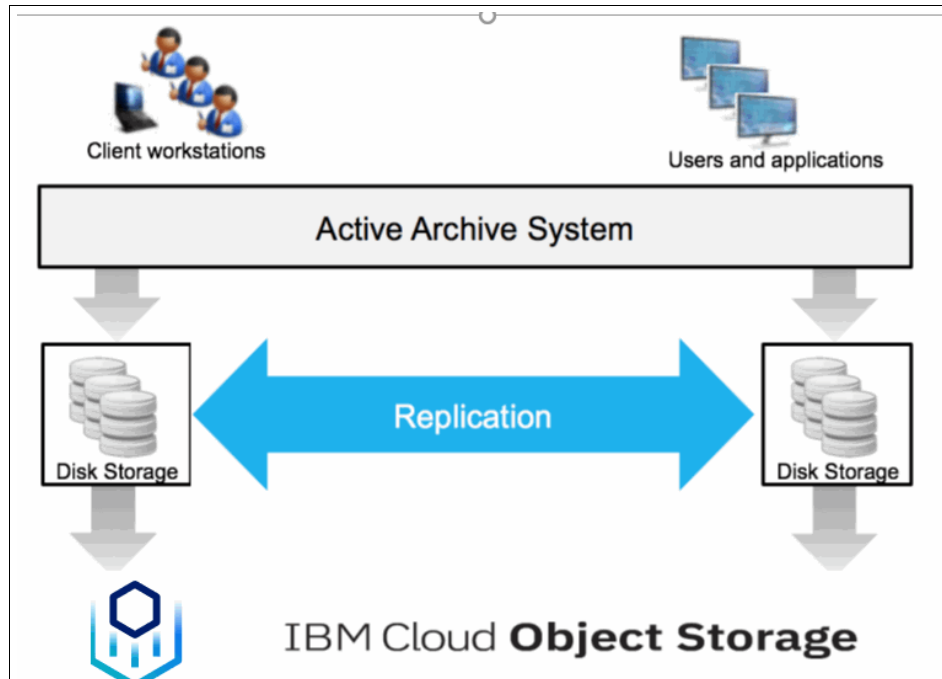


Figure 1-16 IBM Cloud Object Storage enabled active archive system

Why IBM Cloud Object Storage

Serving as an archive repository back-end, IBM Cloud Object Storage solves the following key issues that are present in traditional archive architecture:

- ▶ The geographically dispersed nature of IBM Cloud Object Storage, which provides the following key benefits:
 - Site fault loss resilience without multiple copies greatly reduces the capacity requirements, and removes the need to monitor replication.
 - Removal of the primary and secondary (read only) copy concept means a disaster involving a loss of a site, or loss of access to a site, does not require any traditional high availability (HA) processes (for example, failover to secondary, promote secondary to read/write access, reverse replication, demote primary to read only, or the reverse on fail back).
- ▶ The flattening of the data tiers enables simplified administration and uniform file access, which are expectations of users. The data access speed is the same, regardless of file age or type.
- ▶ The need to manage a large landing space is typically removed or drastically reduced because the data can be read or streamed directly from IBM Cloud Object Storage to the user, who can then view or edit the file locally and then upload through the active archive system.

- ▶ Capacity limitations that are present in traditional storage systems are largely removed because applications can access a “limitless” pool of storage. Therefore, although capacity usage must be measured for fiscal planning, management and design of the storage infrastructure is no longer a burden for the application, system administrators, or architects.
- ▶ The improved operational simplicity, reduced capacity, and increased resilience typically result in a lower total cost of ownership (TCO) for an active archive system when compared with traditional on-premises or replication-based cloud-based deployments.

1.4.5 Use case 5: Enterprise file services

Enterprise file services have existed for decades, with the initial implementations being constrained to local machine access. As enterprise dependence on IT grew, the need for centralized file services was apparent, and several vendors released various devices to satisfy these requirements.

Figure 1-17 shows a silo approach for enterprise file services.

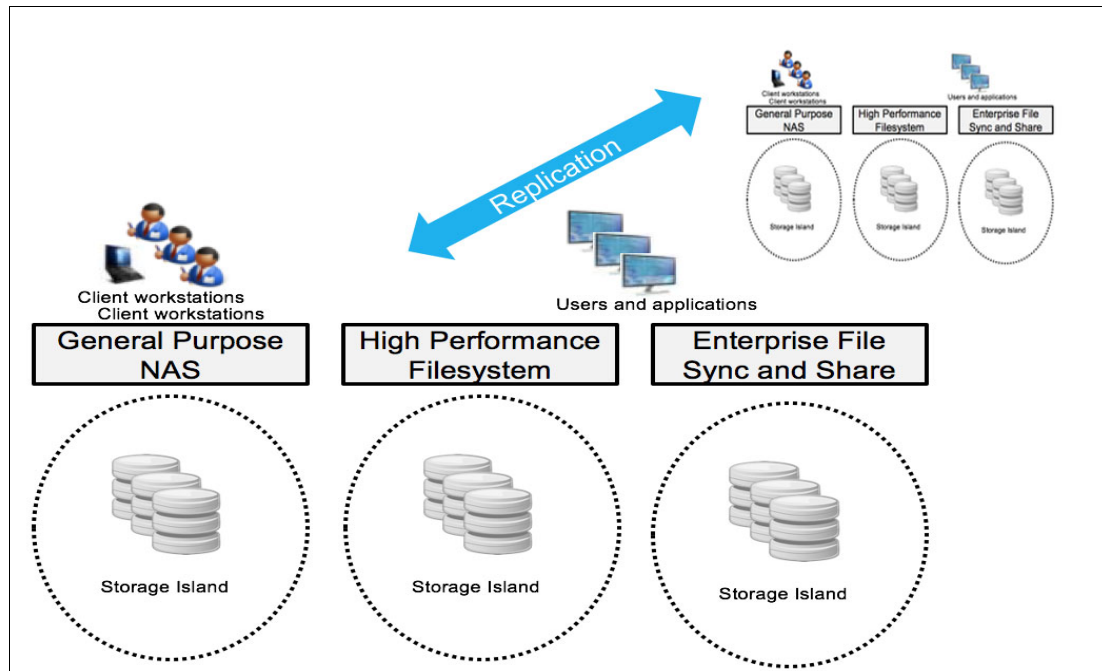


Figure 1-17 Enterprise file services: A silo approach

Consider the following points:

- ▶ Array resiliency is achieved through RAID schemes on the underlying block storage, and multisite resiliency is achieved through replication. Depending on RAID schemes, large-scale data protection has become cumbersome, especially on large capacity spinning disk drives. More disks are required for data protection to ensure resiliency, which requires significant administrator resources to manage because loss of one or two disks in the same RAID stripe can result in data loss.
- ▶ Multisite data resilience requires data replication by using significant network bandwidth and needing at least twice the capacity of the data that is being stored.
- ▶ Also, 3 - 5 times the capacity is required for the backup storage.

- ▶ At a high level, file access systems structure data into directories, subdirectories, and files that are shared among many users, workstations, and servers. Security is based on user and group concepts. Typically, a metadata repository stores the relationship between these components.
- ▶ Although NAS systems in particular are mature and understood, the underlying architecture of these systems poses several challenges because of the immense amounts of data being generated by users, applications, and sensors.
- ▶ The controller-enclosure scaling approach means that rigid limits exist regarding the number of NAS heads that a cluster can contain, which limits the capacity that is available in a single namespace. This approach creates several discrete storage islands within enterprises. This problem is what modern storage systems were designed to solve.
- ▶ The dependency on metadata access for file operations can place an enormous load on this component. It can often be a choke point on large file-based systems, especially if they contain millions to billions of objects.

An IBM Cloud Object Storage based architecture overcomes these challenges and makes many use cases for file services available, including the following functions:

- ▶ **General-purpose file service**
These solutions provide user file and application access through standard protocols, such as SMB/CIFS and NFS, and general-purpose file gateways, such as Nasuni, Panzura, and CTERA.
- ▶ **IBM Cloud Object Storage File Access (IBM Cloud Object Storage FA)**
IBM Cloud Object Storage FA is gateway software that provides SMB and NFS protocol interface to legacy applications to store and retrieve files on IBM Cloud Object Storage for Active Archiving use cases. It is deployed on-premises as a virtual machine, where back-end IBM Cloud Object Storage bucket endpoint can be on-premises or in IBM public Cloud. Cloud Object Storage FA is discussed in more details in Chapter 7, “IBM Cloud Object Storage System File Access” on page 185.
- ▶ **High-performance computing (HPC)**
These systems provide extreme performance and scalability by using parallelism and high-speed network interconnect over protocols that are implicitly designed for high performance, such as IBM Spectrum® Scale.
- ▶ **File sync and share**
These systems provide user file sharing capabilities by using web interfaces, agent software on workstations and mobile devices, and sometimes tie in with NAS services.

IBM Cloud Object Storage based solution for file services can include one or more of the gateway offerings because any enterprise can have a mixture of requirements, all of which can use a common IBM Cloud Object Storage back end.

Note: There is also the IBM Cloud Object Storage gateway software that provides SMB and NFS protocol interface to legacy applications to store and retrieve files on IBM Cloud Object Storage for Active Archiving use cases. It is deployed on-premises as a virtual machine, where back-end IBM Cloud Object Storage bucket endpoint can be on-premises or in IBM public Cloud.

For more information about IBM Cloud Object Storage gateway offerings, see *Cloud Object Storage as a Service: IBM Cloud Object Storage from Theory to Practice - For developers, IT architects and IT specialists*, SG24-8385.

Why IBM Cloud Object Storage

IBM Cloud Object Storage works well with enterprise file services because of the following functions:

- ▶ IBM Cloud Object Storage acts as a single unified back-end for file-based access. It offloads the enterprise from management and operational costs by removing the older storage silo approach.
- ▶ The geographically dispersed erasure coding data protection, which is unique to IBM Cloud Object Storage, means that data is instantly available at multiple sites. This availability eliminates RAID and replication management, and the extra capacity that this feature requires. Storage overhead is less than 2x, including backups.
- ▶ Depending on the solution, the metadata for the file access can be stored in the cloud, which is replicated between gateway devices, or stored on the user's local host in agent-based solutions. This approach preserves the security constructs that are necessary for file-based access, which means that no rewriting of applications is necessary.
- ▶ The best fit solution can be used for each workload so that enterprises are not locked into one NAS solution. The economies of scale are achieved at the underlying IBM Cloud Object Storage layer as opposed to the NAS heads, which are tied to deploying large on-premises NAS systems with hundreds to thousands of disk spindles that are dedicated to one storage silo.

1.4.6 Use case 6: Content repository

Content repository systems manage large catalogs of digital content, such as documents, images, and videos. These systems differ from active archive in that they provide comprehensive search and file management functions that include, but are not limited to, access control, checking in and checking out documents, file locking, and file tiering.

These systems typically store a mix of data types and must provide resilient services to their users who can use the data in these systems for their daily workflow. The locking features ensure that duplicate or forked versions are not unintentionally created, or that data is not lost when users accidentally overwrite by saving old files over an updated version.

In an older content repository system, the following conditions exist:

- ▶ Resiliency is typically provided through replication. Data tiering can be managed by the content repository system, or it might sit on top of an active archive system that is performing the underlying data protection.
- ▶ Content repository systems are typically constructed from a mix of disks and tapes, and are not performance sensitive. However, user experience and productivity can be compromised by extended wait times if data must be recalled from cold storage, such as tape.
- ▶ Disaster recovery often includes traditional HA procedures of promoting and demoting copies, and failing back is not possible until new data is replicated to original sites.

Figure 1-18 shows a content repository solution that is constructed around IBM Cloud Object Storage. In this solution, the application sits on traditional storage that hosts the catalog of data that is browsed by users while the data is held on IBM Cloud Object Storage.

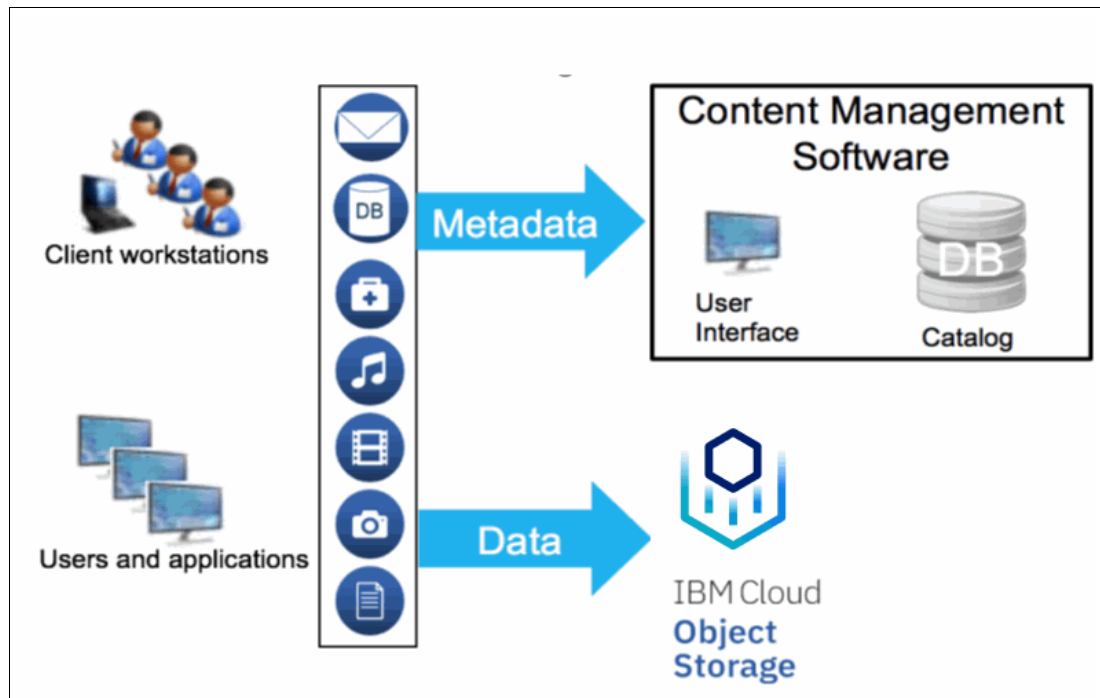


Figure 1-18 IBM Cloud Object Storage enabled content repository systems

Why IBM Cloud Object Storage

When used as the back-end storage for a content repository and management system, IBM Cloud Object Storage provides the following benefits:

- ▶ Offers a cost-effective solution that can solve several issues in traditional architectures.
- ▶ Removal of offline cold storage ensures that users can access data instantly, as opposed to waiting for extended recall times for old data.
- ▶ System administrators benefit from the geographically dispersed nature of IBM Cloud Object Storage because it removes the resources that are required to manage data replication. It also eliminates complex failover of the data store that is used by the repository.



Planning and sizing an IBM Cloud Object Storage System

This chapter provides information about designing an IBM Cloud Object Storage System to suit customer requirements.

This chapter includes the following topics:

- ▶ 2.1, “Planning for capacity” on page 24
- ▶ 2.2, “Performance planning” on page 37
- ▶ 2.3, “Planning for high reliability and availability” on page 44
- ▶ 2.4, “Network planning” on page 46

2.1 Planning for capacity

The following aspects of capacity requirements must be considered when designing an IBM Cloud Object Storage solution:

- ▶ Initial capacity requirement
- ▶ Incremental capacity requirement
- ▶ Expected growth rate over time

2.1.1 Initial capacity requirement

To satisfy a usable capacity requirement, the following items must be considered:

- ▶ Slicestor model
- ▶ Number of Slicestors
- ▶ Drive size
- ▶ One-site, two-site, or three-site (or more) deployment
- ▶ Possible Information Dispersal Algorithm (IDA) alternatives

Verify whether the initial capacity requirement is for TB/PB (decimal) or TiB/PiB (binary):

- ▶ 1 TiB = 1.09951 TB
- ▶ 1 PiB = 1.1259 PB

For systems smaller than 2 PB, a Concentrated Dispersal (CD) Mode system is typically more cost efficient than a Standard Dispersal (SD) Mode system. CD Mode allows for smaller initial capacity (starting as low as 72 TB usable) because it requires less Slicestors than an SD Mode system. For more information about IBM Cloud Object Storage dispersal modes, see *IBM Cloud Object Storage Concepts and Architecture*, REDP-5537.

Note: A CD Mode system has lower performance than an SD Mode system, especially for reads.

Raw capacities of various Slicestor models with different drive sizes

Figure 2-1 shows the raw TB capacity of various Slicestor nodes.

Drive Size	Slicestor 12	Slicestor 53	Slicestor 92	Slicestor 106
4 TB	48	212	N/A	424
6 TB	N/A	N/A	552	N/A
8 TB	96	424	736	848
10 TB	N/A	N/A	920	N/A
12 TB	144	636	1104	1272
16 TB	192	848	1472	1696
18 TB	216	954	1656	1908
20 TB	240	1060	1840	2120

Figure 2-1 Raw TB capacity of various Slicestor models

Important: All Slicestor nodes in a single device set must be of the same model, and they are to be populated with same-sized drives.

Typical overall raw starting capacity range for various Slicestor nodes

Figure 2-2 shows the typical overall raw starting capacity for various Slicestor nodes. In SD Mode, Slicestors can be partially populated and hence starting point for larger Slicestors can be smaller.

	CD Mode	SD Mode	SD Mode Partially populated
Slicestor 12	<1PB	500TB - 2PB	500TB - 2PB
Slicestor 53	1PB - 4PB	5PB - 10PB	1PB - 2PB
Slicestor 92	2PB - 6PB	8PB - 16PB	2PB - 4PB
Slicestor 106	2PB - 8PB	10PB - 20PB	6PB - 10PB

Figure 2-2 Typical starting capacities for CD Mode and SD Mode

Typical expansion factors for one, two, three, and more sites

The typical expansion factors for one, two, three, and more than three sites are shown in Figure 2-3.

Number of sites	Typical expansion factor	
	CD Mode	SD Mode
One	1.5 - 2.0	1.3 - 1.6
Two	2.4 - 3.0	2.8 - 3.5
Three	1.8 - 2.0	1.8 - 2.0
More than three	1.5 - 1.8	1.7 - 1.8

Figure 2-3 Typical expansion factors for one and multiple sites for CD Mode and SD Mode systems

Note: Consider the following points:

- ▶ Expansion factor = (raw capacity)/(usable capacity)
- ▶ Expansion factor = (IDA width)/(IDA read threshold (RT))

Supported IDAs for one, two, three, or more sites

To decide between possible Slicestor node models, start with required usable capacity and multiply that value by a typical expansion factor from the list that is shown in Figure 2-3 on page 25 to determine the raw capacity requirement. Compare that value to typical starting capacity ranges that are shown in “Typical expansion factors for one, two, three, and more sites” on page 25.

Important: Although a larger drive size is more economical in terms of cost per GB, a smaller drive size results in more spindles and more Slicestor nodes in the system, which leads to better performance.

Small capacity configuration for a 3-site example

In this example, a customer needs 200 TB of usable starting capacity. A typical IBM Cloud Object Storage expansion factor for three sites is 1.8 - 2.0.

Start by multiplying the required usable capacity by a typical expansion factor to determine the required raw capacity:

$$\text{Required raw capacity} = 200 \text{ TB} \times 2.0 = 400 \text{ TB}$$

Compare the initial raw capacity to the ranges that are shown in “Typical expansion factors for one, two, three, and more sites” on page 25.

For 400 TB raw capacity across three sites, we choose three Slicestor 12 nodes or six Slicestor 12 nodes in CD Mode. If performance is not a concern and we want to design the most cost-effective solution, we can proceed with three Slicestor 12 nodes if they can provide the required capacity.

Next, we calculate the total raw capacity for three Slicestor 12 nodes with various drive sizes:

$$\text{Raw capacity} = (\text{total number of disk drives}) \times (\text{drive size})$$

Consider the following points:

- ▶ 4 TB drives: $36 \times 4 \text{ TB} = 144 \text{ TB}$
- ▶ 8 TB drives: $36 \times 8 \text{ TB} = 288 \text{ TB}$
- ▶ 12 TB drives: $36 \times 12 \text{ TB} = 432 \text{ TB}$
- ▶ 16 TB drives: $36 \times 16 \text{ TB} = 576 \text{ TB}$
- ▶ 18 TB drives: $36 \times 18 \text{ TB} = 648 \text{ TB}$
- ▶ 20 TB drives: $36 \times 20 \text{ TB} = 720 \text{ TB}$

From the previous calculations, you see that 12 TB drives match the required initial raw capacity.

With three Slicestor 12 nodes, the only possible IDA for three sites is 18/9/11 (expansion factor 2.0).

Figure 2-4 on page 27 shows valid IDAs for three sites and CD Mode.

Geo-Dispersed 3 Site IDAs Using Slicestor 12 & Other Certified Slicestors with fewer than 24 disks - Concentrated Dispersal Mode

Total # of Slicestors	Optimization (S/P)	Width	RT	WT	System Exp Factor	Disk Size	Raw TB	Usable TB
3	S	18	9	11	2.000	4TB	144	72
6	S	36	20	23	1.800	4TB	288	160
6	P	18	9	11	2.000	4TB	288	144

Figure 2-4 Valid IDAs for three sites and Slicestor 12 by using CD Mode

Calculate the usable capacity of the planned system by using the following equation:

Usable capacity = (raw capacity) / (expansion factor)

Usable capacity = 432 TB / 2.0 = 216 TB

Medium capacity configuration for a three-site example

In this example, a customer needs 3 PB of usable starting capacity. A typical IBM Cloud Object Storage expansion factor for three sites is 1.8 - 2.0.

Start by multiplying the required usable capacity by a typical expansion factor to get the required raw capacity:

Required raw capacity = 3,000 TB × 2.0 = 6,000 TB

Compare the initial raw capacity to the ranges that are shown in Figure 2-5. A possible solution is to design an SD Mode IBM Cloud Object Storage System with Slicestor 53 nodes.

Geo-Dispersed 3 Site IDAs - Standard Dispersal Mode

Width	RT	WT	Exp Factor
12	6	8	2.000
15	8	10	1.875
18	9	12	2.000
21	11	13	1.909
24	13	15	1.846
27	14	17	1.929
30	17	19	1.765
33	18	21	1.833
36	20	23	1.800

Figure 2-5 Valid IDAs for three sites that use SD Mode

For 6,000 TB raw capacity across three sites, we can choose 12, 15, or 18 x Slicestor 53 nodes in SD mode, fully or partially populated with drives.

Build a table to compare raw and usable capacities with various drive sizes and IDAs.

Figure 2-6 shows raw capacity calculations for partially and fully populated Slicestor 53s with different drive sizes and different IDAs.

Drive size	IDA 12/6/8	IDA 15/8/10	IDA 18/9/12
SS53/14x4TB	672	840	1008
SS53/14x8TB	1344	1680	2016
SS53/14x12TB	2016	2520	3024
SS53/14x16TB	2688	3360	4032
SS53/14x18TB	3024	3780	4536
SS53/14x20TB	3360	4200	5040
SS53/28x4TB	1344	1680	2016
SS53/28x8TB	2688	3360	4032
SS53/28x12TB	4032	5040	6048
SS53/28x16TB	5376	6720	8064
SS53/28x18TB	6048	7560	9072
SS53/28x20TB	6720	8400	10080
SS53/41x4TB	1968	2460	2952
SS53/41x8TB	3936	4920	5904
SS53/41x12TB	5904	7380	8856
SS53/41x16TB	7872	9840	11808
SS53/41x18TB	8856	11070	13284
SS53/41x20TB	9840	12300	14760
SS53/53x4TB	2544	3180	3816
SS53/53x8TB	5088	6360	7632
SS53/53x12TB	7632	9540	11448
SS53/53x16TB	10176	12720	15264
SS53/53x18TB	11448	14310	17172
SS53/53x20TB	12720	15900	19080

Figure 2-6 Overall raw capacity of 12, 15, or 18 fully or partially populated Slicestor 53 nodes with different drive sizes

Table highlights raw capacities of around 6,000 TB - 7,000 TB. Partially populated options are highlighted in gold and fully populated options are highlighted in green.

A cost-efficient system that provides an upgrade just by adding new drives is 12x Slicestor 53 nodes each populated with 28 x 18 TB drives. Total raw capacity from table shows 6048 TB.

Calculate the usable capacity of the planned system by using the following equation:

$$\text{Usable capacity} = (\text{raw capacity}) / (\text{expansion factor})$$

$$\text{Usable capacity} = 6,048 \text{ TB} / 2 = 3,024 \text{ TB}$$

Large capacity configuration for a three-site example

In this example, a customer needs 20 PB of usable starting capacity. A typical IBM Cloud Object Storage expansion factor for three sites is 1.8 - 2.0. Customer is also informed that they are not able to deploy extra deep racks in their data center.

Start by multiplying the required usable capacity by a typical expansion factor to get the required raw capacity:

$$\text{Required raw capacity} = 20,000 \text{ TB} \times 2.0 = 40,000 \text{ TB}$$

Compare the initial raw capacity to the ranges that are shown in Figure 2-7.

Since customer is not able to deploy deep racks but requires a large capacity, the best option is to design the system based on Slicestor 92 nodes.

For a system of this size, we tend to look at the wider IDAs because they provide the lowest expansion factors with high reliability and availability.

Figure 2-5 on page 27 shows valid IDAs for three sites and SD Mode.

Next, we compare the raw capacities of fully populated Slicestor 92 nodes with wide IDAs, starting at IDA 24/13/15, as shown in Figure 2-7.

Drive size	IDA 24/13/15	IDA 27/14/17	IDA 30/17/19	IDA 33/18/21	IDA 36/20/23
6TB	13248	14904	16560	18216	19872
8 TB	17664	19872	22080	24288	26496
10 TB	22080	24840	27600	30360	33120
12 TB	26496	29808	33120	36432	39744
16 TB	35328	39744	44160	48576	52992
18 TB	39744	44712	49680	54648	59616
20 TB	44160	49680	55200	60720	66240

Figure 2-7 Raw capacity of fully populated Slicestor 92 nodes with wide IDAs that use SD Mode

From Figure 2-7, we can see several possible solutions with roughly 40 PB of raw capacity.

Most cost-efficient system would be the 24-wide system with 18 TB drives.

$$\text{Usable capacity} = (\text{raw capacity}) / (\text{expansion factor})$$

- ▶ Usable capacity = 39,744 TB / 1.846 = 21,530 TB
- ▶ Highest performing system would be the 36-wide system with 12 TB drives.

$$\text{Usable capacity} = (\text{raw capacity}) / (\text{expansion factor})$$

$$\text{Usable capacity} = 39,744 \text{ TB} / 1.8 = 22,080 \text{ TB}$$

2.1.2 Alternative method to plan for capacity for Standard Dispersal mode

A quick alternative method to evaluate potential Slicestor models and drives is to build a table that lists the raw capacity of each Slicestor node model with every drive option.

Figure 2-8 shows the raw capacities of each IBM Slicestor node with various drive sizes.

Drive Size	Slicestor 12	Slicestor 53	Slicestor 92	Slicestor 106
4 TB	48	212	N/A	424
6 TB	N/A	N/A	552	N/A
8 TB	96	424	736	848
10 TB	N/A	N/A	920	N/A
12 TB	144	636	1104	1272
16 TB	192	848	1472	1696
18 TB	216	954	1656	1908
20 TB	240	1060	1840	2120

Figure 2-8 Raw TB capacity of each fully populated Slicestor model with various drive sizes

The following equation is used:

Usable capacity required = (RT of IDA) × (raw capacity of a Slicestor)

For the previous 20 PB usable capacity requirement, we can now calculate how many Slicestor nodes are required to match 20 PB usable capacity. The result is the same as the RT of an IDA.

Figure 2-9 on page 31 shows a calculation of RTs to match the 20 PB usable capacity requirement with various drive sizes and fully populated Slicestor nodes. Only Slicestor 12 and Slicestor 92 are considered since the customer does not want deep drive enclosures.

Drive size	Raw Capacity	# required to match 20 PB
	Slicestor 12	
4 TB	48	416.7
8 TB	96	208.3
12 TB	144	138.9
16 TB	192	104.2
18 TB	216	92.6
20 TB	240	83.3
	Raw Capacity	# required to match 20 PB
	Slicestor 92	
6 TB	552	36.2
8 TB	736	27.2
10 Tb	920	21.7
12 Tb	1104	18.1
16 TB	1472	13.6
18 TB	1656	12.1
20 TB	1840	10.8

Figure 2-9 Example table to calculate the required number of Slicestor nodes with various drive sizes to match a usable capacity requirement

The decimal numbers must be rounded up to the nearest full number. Because the listed capacities in Figure 2-9 are raw and we want to know how many Slicestor nodes are required for 20 PB usable capacity, we can consider these Slicestor node amounts to equal the RT of an SD Mode IDA.

From the table, we can see that many Slicestor 12 nodes would be required to satisfy the 20 PB usable capacity requirement. We should consider Slicestor 12 nodes only if customer performance requirements can't be matched with Slicestor 92 nodes.

As an example, Figure 2-9 shows that to reach 20 PB of usable capacity we need 18.1 Slicestor 92 nodes that are fully populated with 12 TB drives. Rounding up the number of required Slicestor 92 nodes to 19 means that the capacity of 19 Slicestor 92 nodes with 12 TB drives is required for 20 PB of usable capacity. The number 19 would become the RT of an SD Mode IDA.

We compare those calculated IDA RTs to what is supported for three sites in SD Mode (see Figure 2-10) and can now consider the following options to choose from:

- ▶ 6 TB drives -> calculated RT is 37 -> no IDA with such high RT
- ▶ 8 TB drives -> calculated RT is 28 -> no IDA with such high RT
- ▶ 10 TB drives -> calculated RT is 22 -> no IDA with such high RT
- ▶ 12 TB drives -> calculated RT is 19 -> possible IDA from the following list is 36/20/23
- ▶ 16 TB drives -> calculated RT is 14 -> possible IDA from the following list is 27/14/17
- ▶ 18 TB drives -> calculated RT is 13 -> possible IDA from the following list is 24/13/15

Geo-Dispersed 3 Site IDAs - Standard Dispersal Mode

Width	RT	WT	Exp Factor
12	6	8	2.000
15	8	10	1.875
18	9	12	2.000
21	11	13	1.909
24	13	15	1.846
27	14	17	1.929
30	17	19	1.765
33	18	21	1.833
36	20	23	1.800

Figure 2-10 Valid IDAs for a 3-site solution that uses SD Mode

Now, we can calculate the usable capacity for those three alternatives by using the following formulas:

Usable capacity = (raw capacity of single Slicestor) × (RT of the IDA)

- ▶ Usable capacity with 12 TB drives = 1104 TB × 20 = 22,080 TB
- ▶ Usable capacity with 16 TB drives = 1472 TB × 14 = 20,608 TB
- ▶ Usable capacity with 18 TB drives = 1656 TB × 13 = 21,528 TB

As shown in Figure 2-10, we can see that we can consider the following alternatives:

- ▶ 12 TB drives -> IDA 36/20/23
- ▶ 16 TB drives -> IDA 27/14/17
- ▶ 18 TB drives -> IDA 24/13/15

If high performance is not required, the option with 18 TB drives should be considered.

2.1.3 Incremental capacity requirements

To use the extra capacity in new Slicestor nodes, IDAs must remain unaltered. If our initial system has data in buckets that use a CD Mode IDA 18/9/11, the new hardware addition (device set) must be able to use the same IDA.

Note: A new device set can have different Slicestor models (even different hardware generations) than the original device set. It can also have other drive sizes than the original device set.

Incremental capacity for small capacity configuration for a 3-site example

In this example, the initial small capacity configuration consisted of a single device set with three Slicestor 12 nodes, each with twelve 12 TB drives and a CD Mode IDA 18/9/11.

A new device set can have the same or different Slicestor models, and the same or different drive size.

For example, assume that a 50 TB incremental capacity is required. The smallest possible device set that can be added to the system is three Slicestor 12 nodes, each with twelve 4-TB drives.

The total raw capacity of the new device set is calculated as:

$$3 \times 12 \times 4 \text{ TB} = 144 \text{ TB}$$

The CD Mode IDA is 18/9/11, so the expansion factor is calculated as follows:

$$18 / 9 = 2.0$$

To calculate the net capacity increment in this example, use the following formula:

$$\text{Net capacity} = (\text{raw capacity}) / (\text{expansion factor})$$

By using this formula, you can now calculate the net capacity of the new device set:

$$\text{Net capacity} = 144 \text{ TB} / 2.0 = 72 \text{ TB}$$

Incremental capacity for existing medium capacity configuration for a 3-site example

In this example, the initial medium capacity configuration consists of a single device set with 12 × Slicestor 53 appliances, each partially populated with 28 × 18 TB drives and SD Mode IDA 12/6/8.

Because we have empty drive slots available in the existing device set, it makes sense to use them to increase the system capacity.

For example, assume that 1-1.5 PB usable incremental capacity is required. Figure 2-11 shows possible partial population options for Slicestor 53.

Partial Population of Slicestor 53	
Starting Options	Growth Options
14 Drives	14 to 28 Drives
28 Drives	14 to 41 Drives
41 Drives	14 to 53 Drives
53 Drives	28 to 41 Drives
	28 to 53 Drives
	41 to 53 Drives

Figure 2-11 Partial population options for Slicestor 53

We must add drives that are the same size as existing drives. The options are:

- ▶ Increase the number of 18 TB drives from 28 to 41. 13 new drives per Slicestor are required.
- ▶ Increase the number of 18 TB drives from 28 to 53. 25 new drives per Slicestor are required.

The total raw capacity of drive add-ons can be calculated by using formula:

(Number of add-on drives per Slicestor)*(number of existing Slicestors)*(drive size)

- ▶ Add 13x drives per Slicestor -> 13x 12x 18 TB = 2808 TB raw
- ▶ Add 25x drives per Slicestor -> 25x 12x 18 TB = 5400 TB raw

We can see that adding 13 drives per Slicestor is a better match for capacity requirement.

To calculate the net capacity increment in this example, use the following formula:

Net capacity = (raw capacity)/(expansion factor)

By using this formula, we can now calculate the net capacity of the drive add-on:

Net capacity = 2,808 TB / 2 = 1,404 TB

Incremental capacity for existing large capacity configuration for a 3-site example

In this example, the initial configuration consisted of a single device set with 24x Slicestor 92 nodes, each with 92 x 18 TB drives and SD Mode IDA 24/13/15.

The new device set can have the same or different Slicestor models, and the same or different drive size.

For example, assume that an 8-10 PB usable incremental capacity is required. The smallest possible device set that can be added to the system is 24x Slicestor nodes to support existing vault IDA of 24/13/15.

SD Mode IDA is 24/13/15; therefore, the expansion factor is $24 / 13 = 1.846$. We are then investigating if 24x Slicestor 92 nodes can provide the required 8-10 PB of usable capacity, which equals 14,768-18,460 TB of raw capacity.

Figure 2-12 compares capacity upgrades for a large system that is based on Slicestor 92 nodes and 6 TB, 8 TB, 10 TB, 12 TB, 16 TB, 18 TB, and 20 TB drives, which are fully and partially populated.

Drive Size	IDA 24/13/15
SS92/28x6TB	4032
SS92/28x8TB	5376
SS92/28x10TB	6720
SS92/28x12TB	8064
SS92/28x16TB	10752
SS92/28x18TB	12096
SS92/28x20TB	13440
SS92/64x6TB	9216
SS92/64x8TB	12288
SS92/64x10TB	15360
SS92/64x12TB	18432
SS92/64x16TB	24576
SS92/64x18TB	27648
SS92/64x20TB	30720
SS92/92x6TB	13248
SS92/92x8TB	17664
SS92/92x10TB	22080
SS92/92x12TB	26496
SS92/92x16TB	35328
SS92/92x18TB	39744
SS92/92x20TB	44160

Figure 2-12 Raw capacity of partially populated and fully populated twenty-four Slicestor 92 nodes with various drive sizes

The options that are a match for the raw capacity requirement are highlighted in gold (partially populated) or green (fully populated).

As a best practice, choose the option with fully populated Slicestors because that provides linear performance scaling to the existing system.

To calculate the net capacity increment in this example, use the following formula:

Net capacity = (raw capacity)/(expansion factor)

By using this formula, we can now calculate the net capacity of the new device set:

Net capacity = 17,664 TB / 1.846 = 9,569 TB

2.1.4 Summary

Designing an IBM Cloud Object Storage solution that is based on capacity requirements is easy to get started with, but it must never be the only design criteria. Performance, reliability, and availability also must be considered carefully before deciding on the final solution.

Looking further into the future is meaningful because better performance and lower expansion factors can sometimes be reached if the initial system is not sized for only today's capacity requirement, but for what is expected in 1 - 2 years.

Tip: A capacity-based design often is more affected by future incremental capacity requirements than what is needed to get started. If it is known up front that the required capacity add-ons are less than half of the starting capacity, this information must be built into the design.

If anticipated growth rates of the system are difficult to estimate, IBM Storage Utility Offering can be a good option to implement a Capacity on Demand (CoD) type of solution.

Note: IBMers and Business Partners can use the IBM Storage Modeller tool for capacity planning of an IBM Cloud Object Storage System. Access to Storage Modeller requires you to complete the registration process by using your IBMid and create an IBM PartnerWorld® profile. Contact your IBM representative for more details.

2.2 Performance planning

Although it might seem that the only correct answer to performance questions is “it depends”, we cannot rely on that answer when an IBM Cloud Object Storage System is designed. A solution can be designed with a high likelihood of knowing its performance beforehand.

Performance metrics terms: The following terms are used regarding performance metrics:

▶ **Throughput**

The aggregate volume of data per second that is handled by the IBM Cloud Object Storage System across all Accesser nodes, as measured on the Client-Accesser links that are expressed in MBps and GBps.

▶ **Operations per second (OPS)**

The aggregate number of HTTP requests per second handled by the IBM Cloud Object Storage System across all Accesser nodes. This value is *not* disk IOPS.

▶ **Concurrency**

The aggregate number of concurrent requests that is handled by the IBM Cloud Object Storage System across all Accesser nodes.

The following important details affect the performance and sizing of the solution:

- ▶ Typical object sizes
- ▶ Read/Write/Delete ratio
- ▶ Storage engine that is used:
 - Zone Slice Storage (ZSS)
 - Packed Slice Storage (PSS)
- ▶ IBM Cloud Object Storage settings and options that are used:
 - SD Mode versus CD Mode
 - Vault Mode versus Container Mode
 - Chosen IDA
 - Vault mirroring
 - SecureSlice
 - The use of encrypted protocols (HTTPS and Transport Level Security (TLS))
 - Name Index on/off
- ▶ Available network bandwidth:
 - Inside the local data center
 - Between data centers (if two or more sites are used)
- ▶ Accesser models:
 - CPU and RAM
 - Network ports (speed and quantity)
 - Deployment model (physical, embedded, or virtual)
- ▶ Slicestor models:
 - CPU and RAM
 - Network ports (speed and quantity)
 - Number of disk drives

IBM Cloud Object Storage System can be deployed in one, two, or multiple data centers.

Deployments: The following deployments are available:

- ▶ Single data center

All clients, Accesser nodes, and Slicestor nodes are directly connected in a LAN environment with minimal latency and zero packet loss.

- ▶ Two-site mirrored

Clients, Accesser nodes, and a subset of Slicestor nodes that represent a “local” site are directly connected in a LAN environment with minimal latency and zero packet loss. Other Slicestor nodes might be in different sites, with latency and potential packet loss between them.

- ▶ Multisite geodispersed

Clients, Accesser nodes, and a subset of Slicestor nodes that represent a “local” site are directly connected in a LAN environment with minimal latency and zero packet loss. Other Slicestor nodes might be in different sites, with latency and potential packet loss between them.

This configuration is sometimes deployed in an asymmetric model, in which two sites are relatively close to each other and the third site is further away.

If the latency between Accesser nodes and Slicestor nodes exceeds 20 ms, IBM Cloud Object Storage automatically uses multiple connections to offset some of the effects of the wide area network (WAN) latency.

We examine performance characteristics of IBM Cloud Object Storage Accesser nodes and IBM Cloud Object Storage Slicestor nodes separately. The architecture of IBM Cloud Object Storage allows the sizing of the Accesser node layer and the Slicestor node layer individually.

2.2.1 Accesser node layer performance for a single-site system

Accesser nodes connect to a client data network and internal data network that is used to move erasure coded data between Accessers and Slicestors. If an Accesser node has two 10 GbE ports, they often are configured as one port for the client network and the other for the internal network.

Because the amount of data that is imported by the Accesser node is less than what it sends to Slicestor nodes because of erasure coding, the internal network is normally more likely to become a bottleneck at the Accesser node layer.

If the files or objects that are imported by Accesser nodes are small, the Accessers' CPUs might become a bottleneck.

For the purposes of this guide, we are assuming an average object size of 5 MB or larger, which should not make an IOPS bottleneck at the Accesser node layer.

Theoretical throughput on a single 10 GbE port is 9.4 Gbps, which considers TCP/IP resource requirements of 1.175 GBps.

With larger object sizes and with sufficient concurrency, we typically see that the Accesser A10 node can saturate the 10 GbE port to the internal network. Because the data that is sent to the internal network is expanded by a factor of 1.5 - 2.0 (erasure coded), we can estimate that the maximum import speed for a single Accesser A10 node with two 10 GbE ports is 600 MB - 800 MB per second.

Typical rule: A single external Accesser A10 node with two 10 GbE ports can import at 600 MB - 800 MB per second when average object size is 5 MB or larger (actual results can vary).

To scale out throughput performance, Accesser nodes are stateless devices and more can be added (or removed) without disruption.

Important: It is a best practice to have N+1 or N+2 Accesser nodes in the access pool to provide the required throughput, even if one (N+1) or even two (N+2) Accesser nodes are offline for any reason.

Because embedded Accesser nodes share Ethernet ports with the Slicestor nodes on which they run, do not expect the same throughput for embedded Accesser nodes.

2.2.2 Accesser node layer performance for a multisite system

If Slicestor nodes are in locations that are remote from Accesser nodes, bandwidth, latency, and possible packet loss of the WAN between the data centers can affect the performance.

SmartWrite allows Accesser nodes to achieve similar maximum throughput for three sites than for a single site, if two of the three sites are local and near each other and only one site is remote (asymmetric deployment).

2.2.3 Accesser node layer performance for a two-site mirrored system

An Accesser node performs erasure coding twice for mirrored vaults. It requires more CPU and RAM from Accesser nodes. In addition, the amount of data that is sent from a single Accesser node to an underlying Slicestor node layer is greater than in single site or multisite deployments.

The expansion factor for a two-site mirrored system is typically 2.4 - 2.8x. A single Accesser node with two 10 GbE ports writing large objects can saturate the port that sends data to Slicestor nodes. Therefore, the maximum import speed is the theoretical speed of the 10 Gbps port divided by the expansion factor.

Typical rule: A single external Accesser A10 node with two 10 GbE ports writing to a mirrored vault can import data at 400 MB - 500 MB per second when the average object size is 5 MB or larger (actual results can vary).

2.2.4 Slicestor node layer performance

This section provides information about Slicestor node layer performance.

Effect of I/O size and reads versus writes on performance

For objects smaller than 1 MB, the bottleneck is typically the HTTP OPS. For larger objects, the IBM Cloud Object Storage System becomes throughput-constrained, as shown in Figure 2-13.




Figure 2-13 Effect of I/O size on OPS and throughput

Important: IBM Cloud Object Storage typically performs sequential writes to disks and random reads from disks.

Any CPU, disk I/O, or network can become the bottleneck in Slicestor node performance. CPU or network are more likely to be bottlenecks for writes, and disk I/O the bottleneck for reads.

Example maximum throughputs of the Slicestor node layer (actual results can vary):

All Slicestor nodes fully populated with drives. Name index disabled:

▶ 100 KB I/O size example

Twelve Slicestor 12 nodes, SD Mode IDA 12/6/8, 100 KB object size:

- 100% reads, 500 MBps total
- 100% writes, 2,500 MBps total

▶ 1 MB I/O size example:

Twelve Slicestor 53 nodes, SD Mode IDA 12/6/8, 1 MB object size

- 100% reads, 14,000 MBps total
- 100% writes, 18,200 MBps total

▶ 5 MB I/O size example:

Twelve Slicestor 92 nodes, SD Mode IDA 12/6/8, 5 MB object size:

- 100% reads, 32,000 MBps total
- 100% writes, 19,000 MBps total

CD Mode systems feature lower read and write throughput.

2.2.5 IDA effect on performance

Implementing various IDAs in larger environments can be used to optimize performance.

Important: With the same amount of hardware, a narrow IDA typically performs better than a wider IDA, especially for IOPS constrained workloads and smaller object sizes.

For example, with 24 Slicestor nodes in a device set, a 12-wide IDA outperforms a 24-wide IDA because it uses two stripes instead of one.

2.2.6 Storage engine choice

IBM Cloud Object Storage System allows a choice between two different storage engines:

- ▶ PSS
- ▶ ZSS

PSS is the legacy storage engine for IBM Cloud Object Storage that packs small objects into 1-MB bin files that help improve performance with small objects. PSS writes objects to EXT4 file system.

ZSS is the new improved storage engine that no longer uses a file system but writes directly to formatted raw disk drives. ZSS outperforms PSS by a wide margin for larger objects (greater than 1 MB) and provides comparable performance to PSS with small objects.

ZSS should be used always unless there are billions and billions of small objects to be stored.

2.2.7 Network performance

Key points to consider when optimizing network performance with IBM Cloud Object Storage are jumbo frames and NIC bonding.

Jumbo frames

The main advantage of the use of jumbo frames is that it reduces the CPU resource use that is required for TCP processing, which provides optimum network usage and higher throughput. Because jumbo frames are larger than standard frames, fewer frames are needed and therefore CPU processing resource use is reduced.

However, a larger frame introduces some delay, approximately 6 - 7 times more than a standard 1500 maximum transmission unit (MTU) Ethernet frame. This issue can affect the application response time, and sensitive applications, such as high-performance computing (HPC), multimedia, video streaming, VoIP, and web services, might be affected because of this latency.

Jumbo frames also add a complexity to the design, and this issue should be factored into your decision. For example, if you use the public internet on the client channel for the Accessers, passing jumbo frames across the internet might be an issue.

If response time is your primary concern (and not throughput), use standard MTU; otherwise, you can consider the use of jumbo frames. However, be aware of the known limitations. If you need more guidance, contact your IBM representative.

For a specific channel, the MTU settings must be consistent across the entire system and all immediate switches. This consistency avoids packet fragmentation, which can affect performance. The MTU setting also must be consistent across the clients that are connecting to the IBM Cloud Object Storage and the client-facing interfaces of the IBM Cloud Object Storage. Accesser nodes and Slicestor nodes can be configured with bonded NICs.

Tip: The default MTU setting for IBM Cloud Object Storage is 1500, which can be changed to support jumbo frames (MTU=9000) if so required.

NIC bonding

IBM Cloud Object Storage NICs can be bonded for load-balancing (Link Aggregation Control Protocol (LACP)) and for redundancy. The network switches must be properly configured to match the settings on the IBM Cloud Object Storage ports if load-balancing is used.

Several hashing algorithms can be used to optimize traffic flow across the aggregated links. The suitable hashing algorithm can be different for traffic into the appliance and traffic out of the appliance.

2.2.8 Measuring performance

IBM advises that customers and third-party integrators always baseline the performance of an IBM Cloud Object Storage System before running any performance tests on the customer application.

A load generator, such as Object Generator (OG), can be used to create a baseline of IBM Cloud Object Storage performance.

The OG tool is an available open-source Java program that can be used as a traditional client for HTTP testing against an Accesser node.

The public OG repository can be found on [GitHub.com](https://github.com), with release artifacts available at [this web page](#).

OG can be run from a Windows or Linux host operating system, and it supports a wide range of functions, including the ability to perform the following functions:

- ▶ Run in thread limited or OPS modes.
- ▶ Support multiple file size ranges from a single invocation.
- ▶ Support different mixes of PUT, GET, DELETE, LIST, HEAD, and Multipart Upload traffic.
- ▶ OG creates random source data that the tool generates.

Best practice: It is considered a best practice to partially fill a vault *before* running any performance tests.

- ▶ The fill should match the planned workload, and use a thread-based load generator that responds to a change in system performance (rather than a fixed OPS load generator).
- ▶ The fill also should continue until stable performance is achieved.
- ▶ IBM Cloud Object Storage System baselining should use the same object sizes and workflow as the planned production workload.

The following common performance test types are available:

- ▶ OPS test
Concurrency is held constant while file size is varied. Primary metrics are OPS and throughput.
- ▶ Latency test
File size is held constant while OPS is varied. Primary metric is latency.
- ▶ Smoke test
Contains a mix of file sizes and operation types. It can be run in a concurrency or OPS mode.
- ▶ Concurrency test
File size is held constant while concurrency is varied. The primary metric is throughput; the secondary metric is latency.

Other considerations: It is also important to monitor a system that is being used for performance testing to ensure that the test results are valid. Consider the following points:

- ▶ The storage pool should always be at full width with no Slicestor nodes down and no pulled or quarantined disks, unless the test is specifically looking at performance on an impaired system.
- ▶ No rebuilding should be occurring during the test run.
- ▶ If disks become quarantined, they should be replaced without performing failing disk migration. Also, rebuilding should be allowed to complete before testing is resumed.
- ▶ Disk fill levels should be monitored to ensure that disk rebalancing is not a factor and that disks do not become too full.
- ▶ Vault deletion returns immediately to the client, but can take some time to run in the background. For this reason, tests that involve vault deletion must allow sufficient time for the vaults to be deleted before starting another test.

2.3 Planning for high reliability and availability

In this section, we discuss planning for high reliability and availability.

2.3.1 IDA selection

A suitable selection of IDA is the key to high reliability and availability. Both are typically measured in nines: the value for reliability estimates the Mean Time To Data Loss (MTTDL), and the value for availability shows the uptime of the system for reading and writing data.

Reminder: IDA of an IBM Cloud Object Storage is made of width (W), RT, and write threshold (WT).

Width describes how many slices in total are created by the Accesser node or, if the source file is greater than 4 MiB, the Accesser node creates 4 MiB segments from source data and applies erasure coding to those 4 MiB segments individually.

An example SD Mode IDA 15/8/10 means that the width is 15, RT is 8, and WT is 10.

IDA rules: Consider the following rules regarding IDAs:

- ▶ $W > WT > RT$
- ▶ $RT \geq W / 2$
- ▶ $RT + WT > W$

The difference between W and RT determines the reliability. The difference between W and WT determines the system availability, as shown in Figure 2-14.

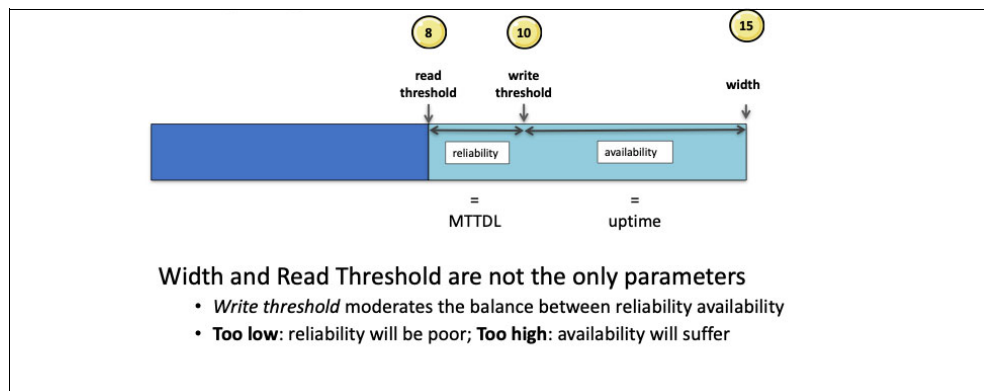


Figure 2-14 Write threshold that uses IDA 15/8/10 as an example

The setting of WT is critical to having a reliable and available system. Too small a difference between RT and WT results in poor reliability, and too large a difference degrades the system availability for writes.

Key point: A unique feature of IBM Cloud Object Storage is how the WT is being used to improve data reliability and to minimize the possibility of data loss.

IBM Cloud Object Storage allows new writes to the system if the WT number of slices can be written. However, if another drive failure takes the system below the WT, new writes are no longer accepted. Read requests continue to be serviced.

CD Mode systems include fixed IDAs that cannot be changed. SD Mode allows IDA parameters to be changed.

Important: A best practice is to have a difference of two or three between read and WTs.

2.3.2 Best practices for data center planning

IBM Cloud Object Storage ensures that an object is restored if any RT number of slices are preserved with the latest object state, even though a subset of IBM Cloud Object Storage Slicestor appliances may return no data, stale data, or even corrupted data. This property of the IBM Cloud Object Storage System provides a high level of confidence that write operations confirmed to a client before being fully committed to media will avoid being lost and data will be preserved for future read requests. However, this statement might not hold true in the case of correlated failures, when multiple devices experience an outage at approximately the same time.

Here are best practices for data center planning by area of consideration:

- ▶ Storage pool device set and stripe configuration:
 - If possible, for the highest availability and reliability, each storage pool device set should be distributed across multiple cabinets in such a way that an outage of any single cabinet will not result in the loss of (WT – RT) Slicestor appliances.
 - Within a storage pool, vaults may be configured to a width less than the width of the storage pool. In this case, the vault has multiple stripes spread across the cabinets.
- ▶ Site configuration:
 - If an uninterruptible power supply (UPS) backup is not available, single site deployments should be avoided and are not recommended until UPS backup is installed and functioning.
 - For multi-site deployments, the IBM Cloud Object Storage System design should allow for resiliency in the case of a single site outage by ensuring the recommendations that are stated above regarding the Storage pool set and stripes best practices are adhered to.
 - Distribute the Slicestor appliances to different locations within the data center if that is an option (mimic geo-dispersed configuration).
- ▶ Power source and circuit configuration:
 - Separate power circuits should be supplied to the networking equipment in a manner that ensures any loss of a single power source will not affect the operation of the IBM Cloud Object Storage System.
 - The IBM Cloud Object Storage cabinets should be connected to different power circuits to ensure that there are less than (WT – RT) Slicestor appliances under the same power source.
 - All IBM Cloud Object Storage appliances are equipped with redundant PSUs. All available PSUs within the IBM Cloud Object Storage appliance should be powered with separate power circuits.

- ▶ UPS configuration:
 - UPS should be available either directly at the cabinet level or through a data center wide capability. Electrical design of a data center following the “Tier 2” level (or higher) of the ANSI/TIA-942 standard provides sufficient protection against sudden power loss.
 - As a best practice, monitor the UPS to know when it takes over the distribution of power.
 - The UPS facilities should have a battery runtime that can power the IBM Cloud Object Storage System until grid power is restored.

If only a short-term UPS backup is available that is sufficient to flush data to disk but not to sustain operations, then a controlled shutdown of the IBM Cloud Object Storage System must be performed before the UPS power runs out. To do a controlled shutdown, run the **poweroff** command from the `localadmin` shell on each appliance. You can use the same command for a planned maintenance of the power distribution system.

2.3.3 Multiple Manager devices

You can add a second Manager device to the system to have continued visibility into system operation and the provisioning capabilities if one Manager device fails.

When both Manager devices are running, you can perform configuration changes concurrently, and observe events on either Manager device.

A maximum of two Manager devices may run simultaneously within the system.

2.4 Network planning

In this section, we discuss the different networks that can be used with IBM Cloud Object Storage, various network and switch settings, and the available load-balancing options.

2.4.1 Multi-networks with IBM Cloud Object Storage

The following network types and configurations are available:

- ▶ Data:
 - This network is the network for all data traffic.
 - Data is passed to and from Accesser nodes to the Slicestor nodes, and to and from other Slicestor nodes.
 - This network is for all the IBM Cloud Object Storage internal management traffic (internal monitoring, upgrades, and so on).
 - All Slicestor nodes and Accesser nodes communicate to the Manager node on this network.
 - This network is *required*.

- ▶ Management:
 - This network is the network to interface with your internal management infrastructure (Domain Name System (DNS), Network Time Protocol (NTP), SNMP, SMTP, AD, and so on).
 - Your administrators may access the IBM Cloud Object Storage solution through this network.
 - This network is *optional* (if not available, all management traffic is on the data network).
- ▶ Client:
 - This network is the network for client traffic (applications that use IBM Cloud Object Storage for storage).
 - All Accesser nodes communicate with user applications on this network.
 - Data comes into the Accesser node layer on this network, is processed, and then is distributed to Slicestor nodes on the data network.
 - This network is *optional* (if not available, all client traffic is on the data network).
- ▶ Intelligent Platform Management Interface (IPMI):
 - This network is a remote hardware health monitoring and management system that defines interfaces for use in monitoring the physical health of servers, such as temperature, voltage, fans, power supplies, and chassis.
 - Only used if you use IPMI in your environment (your networking or operations team are aware if IPMI is used).
 - This network is *optional*.

Attention: IBM Cloud Object Storage refers to these networks as *channels*.

An example of the different networks is shown in Figure 2-15.

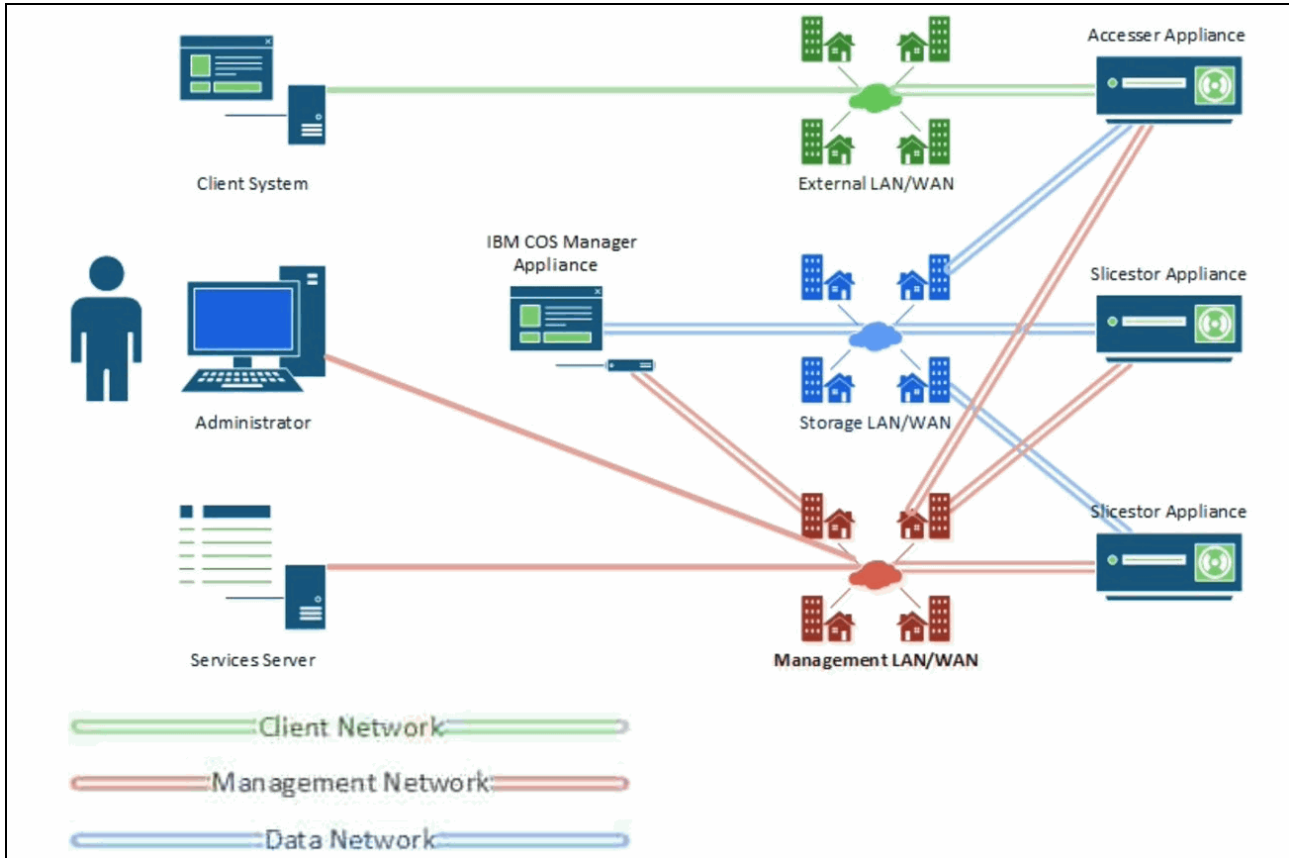


Figure 2-15 IBM Cloud Object Storage networks

For more information, see *IBM Multi-Network Configuration Guide*, which is available at [IBM Documentation](#).

Note: 802.1Q VLAN tagging is not supported by IBM Cloud Object Storage appliances.

2.4.2 Network Time Protocol

NTP configuration is required for suitable operations of an IBM Cloud Object Storage System. Time must be synchronized among the IBM Cloud Object Storage nodes and across all connecting clients. Typically, the IBM Cloud Object Storage Manager node synchronizes with an external time source, and all other nodes synchronize with the Manager node. Having the correct time is important because if the time difference between the client (application) and the server (IBM Cloud Object Storage) is too large, the S3 protocol will not process requests (*RequestTimeTooSkewed* error).

2.4.3 Load balancers

A load balancer (LB) enables client applications to seamlessly distribute traffic among multiple Accesser nodes in an access pool. Many types of LB configurations are available that can be implemented. The following LB configurations are most commonly used:

- ▶ Global server load-balancing (GSLB) and server load-balancing (SLB)

The client applications are directed to selected endpoints based on the origination information of the client application's traffic. The endpoint can be an Accesser node or another LB, which can further route the traffic.

The GSLB is the implementation where the underlying endpoints can be in separate geographic regions, where the SLB is the implementation where the underlying endpoints might be in the same geographic region.

- ▶ DNS round robin

The DNS lookup of the fully qualified domain name (FQDN) of the endpoint returns an IP from a predefined pool of IPs. The returned IP can be an Accesser node, or another GSLB or SLB.

- ▶ Direct Server Return (DSR)

This implementation is preferred for READ traffic. This implementation is where the LB forwards traffic from the client application to the underlying Accesser node during session establishment, and the Accesser node sets up communication directly with the client application.

Consider the following points when you configure GSLB or SLB:

- ▶ The client connections can end at the LB or pass through the LB and be redirected to the Accesser node:
 - Ending connections at the LB helps avoid any direct access to the IBM Cloud Object Storage System, but requires the operator to ensure that the correct TCP buffer size is configured at the LB.
 - Ending connections at the LB can lead to added latency if SSL connections are used.
 - Passing connections through the GSLB avoids the resource use of SSL processing, and eliminates latency that is added by processing at the GSLB.
 - LBs should be configured to pass through x-forwarded-for request to Accesser nodes to reflect the client endpoints in the Access logs. Although this configuration is not required for an IBM Cloud Object Storage System to function properly, it is important for troubleshooting.
- ▶ Ensure that the cipher suite that is required by the client application is available on the GSLB.
- ▶ Ensure that the cipher suite is configured in the correct order at the GSLB because most applications establish communication with the first common cipher that they find.

Example LBs that can be deployed with IBM Cloud Object Storage include (but are not limited to) the following products:

- ▶ F5
- ▶ Citrix NetScaler
- ▶ HAProxy
- ▶ NGINX

HAProxy

HAProxy is a no-charge, fast, and reliable solution that provides high availability (HA), load-balancing, and proxying for TCP and HTTP-based applications. It is suited for high-traffic websites, and powers many the world's most visited sites.

Over the years, it became the de facto standard open source LB and is now included with most mainstream Linux distributions. It also is often deployed by default in cloud platforms. Its mode of operation makes its integration into architectures riskless.

Although IBM does not recommend the usage of HAProxy over commercial LBs, the fact that it is no-cost makes it attractive for users who want to set up an LB in a non-production environment to test it with IBM Cloud Object Storage.

NGINX

NGINX is another widely used open-source-based LB that can be used with IBM Cloud Object Storage.

2.4.4 Firewalls

Firewalls can exist anywhere in the environment and can be placed between the client applications, manager node, Accesser nodes, Slicestor nodes, and all points in between. Figure 2-16 lists the most common firewall ports used.

Firewall ports			
Appliance	Protocol	Port	Description
Mgr/Acc/SS	ICMP	Any	ICMP (Ping)
Mgr/Acc/SS	TCP	7	deNet application detection
Mgr/Acc/SS	TCP	22	ssh admin for cli
Mgr/Acc	TCP	80/8080	HTTP & object interface to vaults
Mgr/Acc/SS	TCP	161	SNMP
Mgr/Acc/SS	TCP	443	HTTPS for dsNet auth/registry data
SS	TCP	5000	IBM COS data dispersal protocol
Mgr/Acc/SS	TCP	6624	dsNet upgrade control
Mgr/Acc/SS	TCP	8088	IBM COS appliance management
Mgr/Acc/SS	TCP	11986	dsNet upgrade control
Mgr	TCP	25/487	SMTP for alert email
Mgr/Acc/SS	TCP	389	LDAP for AD integration
Mgr/Acc/SS	UDP	53	DNS protocol
Mgr/Acc/SS	UDP	123	NTP protocol

Figure 2-16 Most common firewall ports used

For more information about possible IBM Cloud Object Storage firewall ports, see [IBM Documentation](#).

Figure 2-17 shows various IBM Cloud Object Storage network channels and the most common firewall ports that are used.

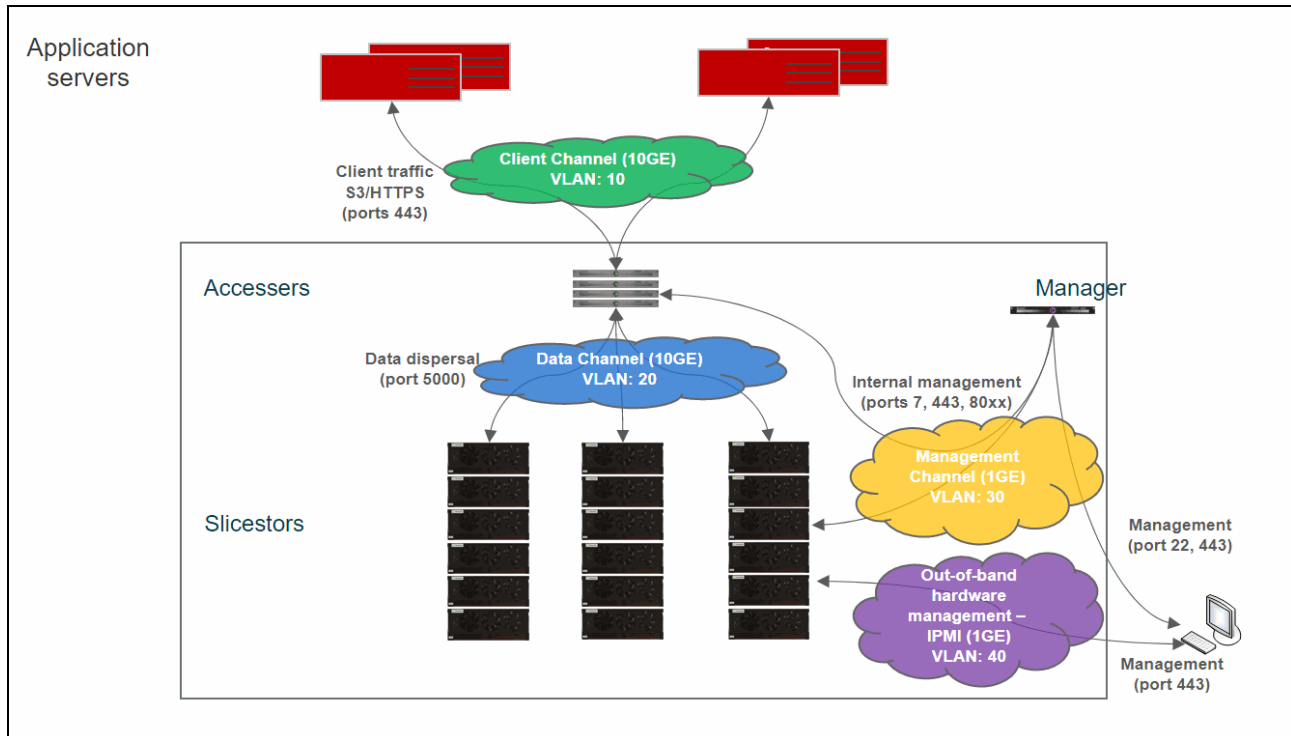


Figure 2-17 IBM Cloud Object Storage networks and common firewall ports

2.4.5 Differences between S3 and IBM Cloud Object Storage APIs

The following table highlights some functional differences between S3 API and IBM Cloud Object Storage API (see Table 2-1)

Table 2-1 Functional differences between S3 API and IBM Cloud Object Storage API

Feature	S3	IBM Cloud Object Storage
Object size limitations	5 TB.	Single objects up to 10 TB with streaming upload support or S3 Multipart Upload.
Retained Version Count Limitations	No explicit limit.	No limit for number of versions per object in Container Mode. A maximum of 1000 retained versions are allowed per object in vault mode.
Vault (Bucket) Granular ACL		Users who are configured in the Manager Web Interface can be granted read/write , read-only , or no-access permissions to any vault. These settings apply to the entire vault.
Vault (Bucket) Granular Data Reliability	Allows a storage class to be configured for each object. All objects that are stored in any vault share reliability characteristics.	Vault reliability characteristics are determined at vault creation time.

Feature	S3	IBM Cloud Object Storage
Traditional Authentication Mechanisms	Uses a custom HTTP scheme based on a keyed-HMAC.	In addition to Access Key authentication, these authentication methods are also supported: <ul style="list-style-type: none"> ▶ HTTP Basic over HTTP and HTTPS ▶ PKI over HTTPS ▶ Anonymous
Separated Audit and Logging Functions		Accesser node collects both access logs and audit trail information but does not expose it through the API.
Encryption and Cryptographic Integrity		<ul style="list-style-type: none"> ▶ An Object Vault can be configured to store information in an encrypted form. It must be configured at the vault/bucket level through the System Manager. These settings cannot be viewed or edited through the API. ▶ Request signing is supported. ▶ Non-cryptographic ▶ MD5 checksums are calculated and stored with objects.
Lifecycle Configuration		Does not support policy-based migration of data to alternative storage classes, or archiving of data. IBM Cloud Object Storage does not support expiration on vaults or containers with versioning enabled.
Vault (Bucket) Location Constraints	Allows buckets to be created with specific location constraints.	Can configure a system to allow data in one vault to be in a separate geographical location from data on another vault. It is configured when vaults are created in the Manager Web Interface.
Hard Quota Function	Does not support quotas for buckets.	A hard quota can be configured for an object vault. HTTP status code 507 (Insufficient Storage) is returned for a write request that would cause a hard quota to be exceeded.

For more information, see the *AWS CLI S3 Developer Guide*, which is available at [IBM Documentation](#).

Note: You can select the version of the IBM Cloud Object Storage software that you use by clicking **Change version or product** at the top of the page.



IBM Cloud Object Storage Gen2 hardware appliances

This chapter provides an overview of the IBM Cloud Object Storage second-generation (Gen2) hardware appliances.

This chapter includes the following topics:

- ▶ 3.1, “Gen2 hardware appliance overview” on page 54
- ▶ 3.3, “Appliance specifications” on page 71
- ▶ 3.4, “Hardware options” on page 73
- ▶ 3.5, “Performance” on page 74
- ▶ 3.6, “Rack guidance” on page 75

3.1 Gen2 hardware appliance overview

In May 2019, IBM introduced the second-generation hardware appliances for IBM Cloud Object Storage. The new appliances provide better performance, higher density, and more flexibility and cost-effective scaling than the first-generation models. The new appliances use the same server hardware components across all three functions: Manager, Accesser, and Slicestor nodes. This design simplifies installation, configuration, monitoring, management, and troubleshooting. In February 2021, a disk enclosure – the IBM Cloud Object Storage 5U92 disk enclosure – was added to the Gen2 range.

3.1.1 Highlights

IBM Cloud Object Storage Gen2 hardware delivers the following benefits:

- ▶ A total of 25% - 50% lower cost than first-generation hardware that is based on workload and use cases.
- ▶ Up to 15% more reads and writes can be completed in the same time frame, versus the present generation of hardware.
- ▶ Consolidated solution that supports up to 1.9 PB in a single node, and 15.2 PB in a single rack.
- ▶ Acquire Only the amount of performance or capacity that is initially required, and grow performance or capacity as needed, together or independently.

New servers in the IBM Cloud Object Storage Gen2 hardware product line include support for the following components:

- ▶ The latest Intel Xeon processors (scalable processor family)
- ▶ Higher performance serial-attached SCSI (SAS) controllers
- ▶ Higher speed memory registered dual inline memory modules (RDIMMs)

The Gen2 architecture provides the storage layer by using a controller node server and disk enclosure combination, including support for the following aspects:

- ▶ Higher capacity in fewer rack units (RUs)
- ▶ More capacity in a single node: 2.2 PB with 20 TB drives in 5 RU
- ▶ Adding capacity independently of performance

All improvements mean that clients can start with small systems in the areas of capacity and performance, and grow as needed in independent steps if that is what is appropriate for them.

All the capabilities, features, functions, and benefits of IBM Cloud Object Storage that clients are maintained in this generation of hardware, including the following examples:

- ▶ The ability to scale out to multi-petabyte or multi-exabyte capacity with faster performance
- ▶ Single-site, mirrored two-site, and geographically dispersed deployment options
- ▶ Configuration that provides up to 8 nines of availability and up to 15 nines of reliability

Clients also receive the following benefits:

- ▶ Lower cost
- ▶ Higher performance
- ▶ More options to tailor configurations for workload optimization
- ▶ Greater consolidation of data in a single RU or a rack configuration
- ▶ The ability to more easily and cost effectively grow their systems as needed
- ▶ The ability to intermix Gen2 hardware with previous generation or third-party hardware

Note: More information can be found on the hardware announcement at [IBM Offering Information](#).

3.2 Appliance overview

IBM Cloud Object Storage Gen2 hardware is an extension and enhancement of the existing hardware. It is an extension because it provides the same functions by using the same IBM Cloud Object Storage software. It also works with the existing generation seamlessly.

Clients can intermix past and present generations of hardware in the same system, even at the storage pool level. For example, a client that uses an IBM Cloud Object Storage System today that contains a Manager 3105 node, two Accesser 3105 nodes, and four Slicestor 2212A nodes can add a set of four Gen2 Slicestor nodes, as shown in Figure 3-1.

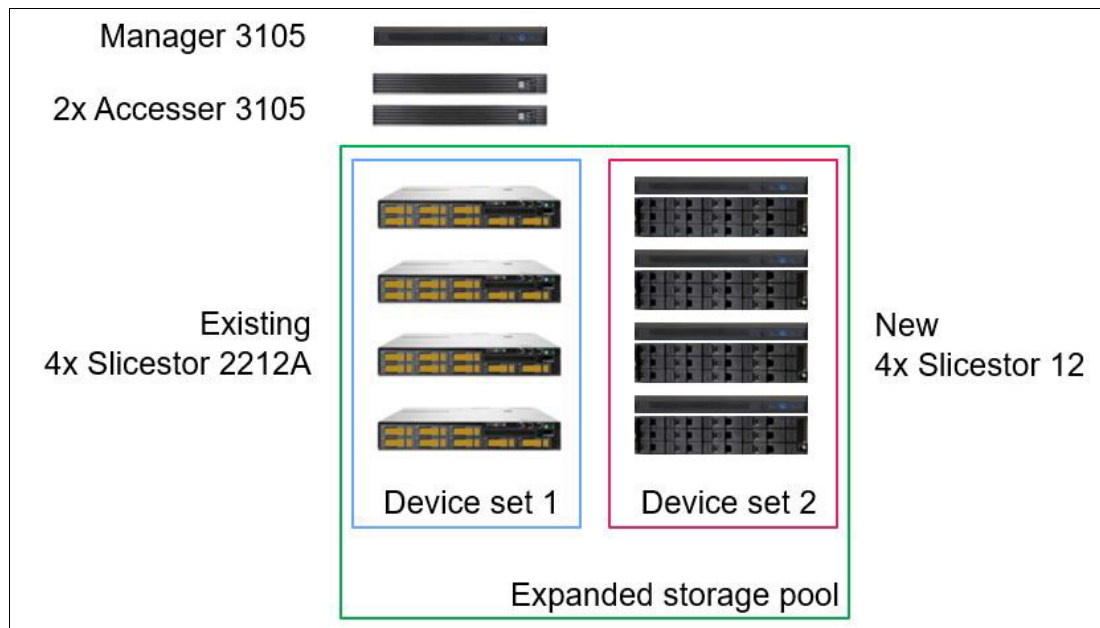


Figure 3-1 Gen1 system expansion with Gen2 Slicestor appliances

IBM Cloud Object Storage supports exabytes (EB) of data in a single namespace, and this generation supports that amount of data and more data per RU and per rack. With this new generation, systems feature increased capacity by up to 89% per RU. A single rack can now hold more than 16.96 PB when 20 TB drives are used if the rack can support the weight of eight fully populated J12 Slicestors. This approach is not recommended.

IBM Cloud Object Storage architecture does not change in the Gen2 hardware because the storage system is still provided by a Manager, Accesser, and a Slicestor function. The Manager and Accesser appliances are consistent with the previous Manager and Accesser appliances.

The key difference in Gen2 is in how the storage layer is designed. In the previous generation, the storage layer is provided by high-density server appliances. The performance components and the storage components in the storage layer are in the same physical box, and therefore must be installed, grown, and expanded together.

In Gen2, the storage layer is divided into two separate components, a *controller node* that contains the performance components and the *disk enclosures* that contain the storage components. In both the previous generation and in Gen2, the storage layer is called a *Slicestor node or appliance*, and the functional components, component names, and system management remains the same. This approach ensures that the same IBM Cloud Object Storage software is used across both the previous generation and Gen2.

The Gen2 server appliances are all based on the same 1U rack server with slightly different configuration for each function. The following appliances are available today:

- ▶ IBM Cloud Object Storage Manager M10
- ▶ IBM Cloud Object Storage Accesser A10
- ▶ IBM Cloud Object Storage Slicestor 12
- ▶ IBM Cloud Object Storage Slicestor 53
- ▶ IBM Cloud Object Storage Slicestor 92
- ▶ IBM Cloud Object Storage Slicestor 106

Note: The Gen2 Slicestor appliances include an IBM Cloud Object Storage Controller Node C10 and Small, Medium, Large, or 5U92 Disk Enclosure.

The Gen2 server appliances are based on the *Intel Xeon Scalable Processor* family (formerly Skylake) or the *2nd Generation Intel Xeon Scalable Processor family* (formerly known as Cascade Lake), giving the platform increased performance over the previous generation. The platform also enables deployment of more memory and multiple CPUs into the server appliances, which make scaling up performance within the boxes possible in the future.

The appliances include two optical 10 GbE ports that are integrated on the system board that come with short-wave SFP+ transceivers as default. The Intelligent Platform Management Interface (IPMI), for out-of-band hardware management, has a dedicated 1 GbE port with an RJ45 connector. The VGA and the two USB ports, which are on the rear of the nodes, can be used for console connection. The IPMI port also supports a virtual console.

The IPMI interface is set to DHCP by default. The setting can be modified in IBM Cloud Object Storage nut shell by the `ipmi` command or on the web-based IPMI GUI. The default username is ADMIN. The default password is ADMIN for Gen1 devices. For Gen2 devices, the password can be found on the pull-out tab above the first hard disk drive on the left side of the server (see Figure 3-2 on page 57) and is specific to each device.



Figure 3-2 Hard disk drive pull-out tab

All appliances include redundant, hot-swappable power supplies and cooling fans. The power supplies can be replaced from the rear, and the cooling fans can be accessed by removing the central cover from the top of the server appliance.

Tip: The IBM Cloud Object Storage Manager monitors the health of the underlying hardware components, and alerting can be set up in the Manager GUI. IBM Cloud Object Storage also allows SNMPv2c or SNMPv3 monitoring. The appliances can be connected to an SNMP server and send alerts in case of failures. SNMP polling is supported as well.

3.2.1 Manager appliance

The IBM Cloud Object Storage Manager M10 (Manager M10) is the successor of the existing Manager 3105. Table 3-1 lists the new machine type and model.

Table 3-1 Gen2 Manager machine type and model

Machine type	Model	Description
4958/4957	M10	IBM Cloud Object Storage Manager M10

The Manager M10 contains a RAID controller that protects the operating system drives from a single drive failure. The node can host up to 10 disks, but initially only two slots on the left side of the chassis are populated with the operating system drives.

The Gen2 Manager appliance supports up to 4,500 nodes in a single IBM Cloud Object Storage System.

Figure 3-3 shows the front of the Manager M10.



Figure 3-3 IBM Cloud Object Storage Manager M10 front view

Figure 3-4 shows the rear of the Manager M10 with the available ports.




Figure 3-4 IBM Cloud Object Storage Manager M10 rear view

3.2.2 Accesser appliance

The IBM Cloud Object Storage Accesser A10 (Accesser A10) is the successor of the existing Accesser appliance models. Table 3-2 lists the new machine type and model.

Table 3-2 Gen2 Accesser machine type

Machine type	Model	Description
4958/4957	A10	IBM Cloud Object Storage Accesser A10

The main difference between the Accesser A10 and the other Gen2 appliances is that it has by default twice the memory and a more powerful, Xeon Gold processor in it that allows high throughput for Accesser-related functions. Figure 3-5 shows the front of the Accesser A10.



Figure 3-5 IBM Cloud Object Storage Accesser A10 front view

Figure 3-6 on page 59 shows the rear of the Accesser A10 with the available ports.




Figure 3-6 IBM Cloud Object Storage Accesser A10 rear view

3.2.3 Slicestor appliances

The new IBM Cloud Object Storage Slicestor appliances are the successors of the existing Slicestor models with similar or slightly higher number of disk slots. Table 3-3 lists the new machine types and models.

Table 3-3 Gen2 Slicestor machine types

Machine type	Model	Description
4958/4957	C10	IBM Cloud Object Storage Controller Node C10 (Controller Node C10)
4958/4957	J10	IBM Cloud Object Storage Small JBOD Chassis (Small Disk Enclosure or Small J10 Disk Enclosure)
4958/4957	J11	IBM Cloud Object Storage Medium JBOD Chassis (Medium Disk Enclosure or Medium J11 Disk Enclosure)
4958/4957	J12	IBM Cloud Object Storage Large JBOD Chassis (Large Disk Enclosure or Large J12 Disk Enclosure)
4958/4957	J15	IBM Cloud Object Storage 5U92 Disk Enclosure (5U92 Disk Enclosure or J15 Disk Enclosure)

The Gen2 SliceStor appliances consist of a controller node and a disk enclosure. The upgrade options are similar to the previous generations, with disks that are added to the empty slots within the SliceStors or SliceStor nodes that are installed in a new device set, as shown in Figure 3-7.




Figure 3-7 Adding capacity and performance to a Gen2 system

All disk enclosures support the same type of 3.5-inch, hot-swappable NL-SAS drives. The Gen2 SliceStor appliances overview is shown in Figure 3-8.





IBM Cloud Object Storage Gen2 SliceStors				
				
	SliceStor™ 12	SliceStor™ 53	SliceStor™ 92	SliceStor™ 106
Rack depth	Normal	Extra deep (1,200 mm)	Normal	Extra deep (1,200 mm)
Total rack space	3U	5U	6U	5U
Disk configurations	12	14, 28, 41, 53	28, 64, 92	64, 78, 92, 106
Drive sizes	4, 8, 12, 14, 16, 18, 20 TB	4, 8, 12, 14, 16, 18, 20 TB	6, 8, 10, 12, 16, 18, 20 TB	4, 8, 12, 14, 16, 18, 20 TB
Node capacity raw (min / max)	48 TB / 240 TB	56 TB / 1.060 PB	168 TB / 1.840 PB	256 TB / 2.120 PB
Investment Protection	Lower Costs	Flexibility	Efficiency	
No migration required; Mix configurations	↓ Ultra dense, with high-capacity drives	More options to optimize configurations	Consolidated flexibility to scale-out	

Figure 3-8 Gen2 SliceStor appliance overview

IBM Cloud Object Storage Controller Node C10

The IBM Cloud Object Storage Controller Node C10 (Controller Node C10) is a 1U rack server that is based on the same hardware as the Gen2 Manager and Accesser appliances. The main difference is that the controller node has a 16-channel SAS HBA that is installed in it that enables connection of the disk enclosures, as shown in Figure 3-10 on page 61.

The controller node connects to the disk enclosure with redundant MiniSAS HD cables. The controller node has two operating system disks that are built in and mirrored by the operating system.

Tip: supplied MiniSAS HD cables for the J10, J11, and J12 are each 2 m. The supplied MiniSAS HD cables for the J15 are . As a best practice, install the controller node and the disk enclosure in the same rack.

Figure 3-9 shows the front of the Controller Node C10.



Figure 3-9 IBM Cloud Object Storage Controller Node C10 front view

Figure 3-10 shows the rear of the Controller Node C10.




Figure 3-10 IBM Cloud Object Storage Controller Node C10 rear view

Note: Controller Node C10 has sufficient memory to run the Embedded Accesser feature, which allows a low-cost option for small configurations when high performance is not a requirement.

Small Disk Enclosure

The Small Disk Enclosure (also known as *Small J10 Disk Enclosure* or *Small JBOD Chassis*) is the lowest capacity, 12-bay, 2U enclosure, which makes it possible to configure an entry-level IBM Cloud Object Storage System with 72 TB of usable capacity. The Small Disk Enclosure requires all drives to be populated. The SliceStor 12 that uses this enclosure is the successor of the SliceStor 2212 and 2212A models.

The minimum raw capacity of the 12-bay enclosure is 48 TB, but it can be as high as 2.12 PB when 20 TB disk drives are used. The drives are at the front of the enclosure, and they can be hot-swapped without moving or sliding out the chassis.

Note: IBM continuously certifies new types of and larger capacity drives. At the time of writing, the largest available drive size is 20 TB. For more information about certified drive options for each enclosure, see the IBM sales manuals at [IBM Offering Information](#).

The Small Disk Enclosure features a different drive carrier compared to the 53, 92, and 106-bay enclosures; therefore, drives are not interchangeable.

Figure 3-11 shows the front of the Small Disk Enclosure.



Figure 3-11 IBM Cloud Object Storage Small Disk Enclosure front view

The front left side of the enclosure features an operator's panel, which provides basic diagnostics functions. The overview of the panel is shown in Figure 3-12.

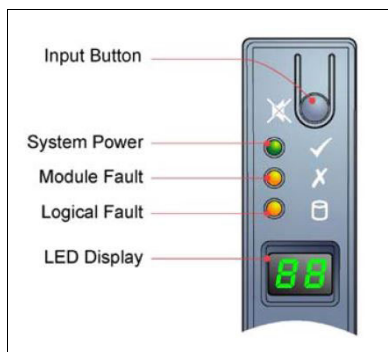


Figure 3-12 Operator's panel of the Small Disk Enclosure

Figure 3-13 shows the rear of the Small Disk Enclosure with a single I/O module.

Note: Three SAS ports are available on the enclosure; however, only the left two ports (A and B) are used to connect it with the controller node by using MiniSAS HD cables.

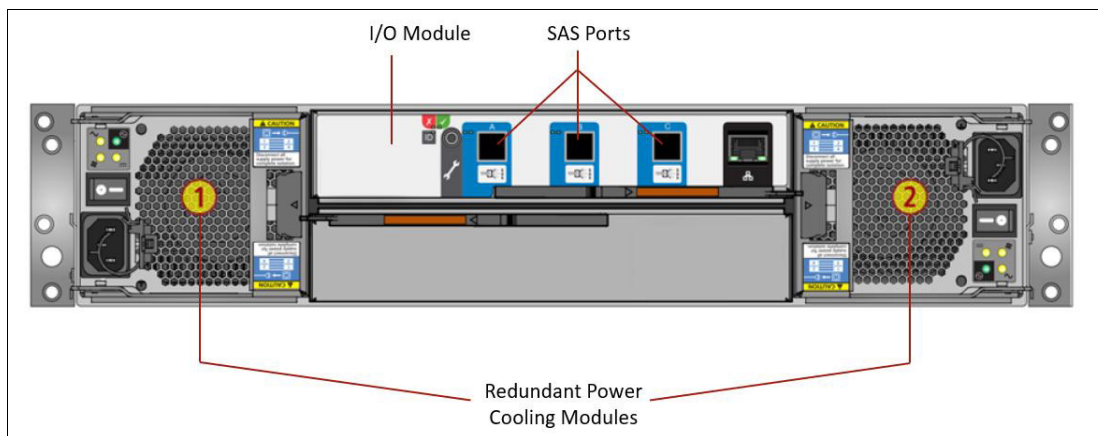


Figure 3-13 IBM Cloud Object Storage Small Disk Enclosure rear view

The power and cooling modules (PCMs) are at the rear, on the sides of the enclosure. A single PCM can supply power and cooling for the enclosure if the second module fails. The PCM has its own LEDs to provide fault information, as shown in Figure 3-14 on page 63. PCMs are hot-pluggable and replacement takes only a few seconds.




Figure 3-14 Power and cooling module

Attention: Never remove the failed PCM, unless the replacement is available. Replace the failed component within a few seconds.

Figure 3-15 shows the drive enumeration of the Small Disk Enclosure.




Figure 3-15 Small Disk Enclosure drive enumeration

Medium and Large Disk Enclosures

The Medium Disk Enclosure (also known as Medium J11 Disk Enclosure or Medium JBOD Chassis) and the Large Disk Enclosure (also known as Large J12 Disk Enclosure or Large JBOD Chassis) use the same enclosure chassis, drives, fan, and power modules and other components. In this section, we describe the similarities and the differences between the enclosures.

Medium Disk Enclosure capacity

The Medium Disk Enclosure is used for high-performance, higher density configurations. The enclosure is 4U and can host up to 53 hot-swappable NL-SAS drives. The Slicestor 53 that uses this enclosure (although it has slightly more drives) is the successor of the Slicestor 2448 and 3448 models.

The minimum raw capacity of the 53-bay enclosure is 212 TB when fully populated with 4 TB drives, and the maximum capacity is 1.06 PB when 20 TB disks are used. The drives can be added or removed from the top of the chassis. During installation, extra care is required for proper cabling because the chassis must slide out to swap failed drives.

Large Disk Enclosure capacity

The Large Disk Enclosure is used for high-density configurations. The enclosure can host up to 106 hot-swappable NL-SAS drives in only 4 RUs. The Slicestor 106 that uses this enclosure is the successor of Slicestor 2584 model.

The minimum raw capacity of the 106-bay enclosure is 424 TB when fully populated with 4 TB drives. The maximum capacity is 2.12 PB when 20 TB disks are used.

Important: Although the 53 and 106-bay 4U disk enclosures show physical similarities, the Medium Disk Enclosure cannot host more than 53 drives, and cannot be upgraded to a 106-bay enclosure. If more than 53 drives are required in a single enclosure, the Slicestor 106/Large Disk Enclosure must be deployed.

Physical view and modules

Figure 3-16 shows the front of the Medium or Large Disk Enclosure.




Figure 3-16 IBM Cloud Object Storage Medium or Large Disk Enclosure front view

The enclosure features a small LED panel at the lower left in the front, which displays basic diagnostics information, as shown on Figure 3-17. The fault LED's name indicates where the failed component is located, or which lid must be removed to access it.




Figure 3-17 Front LED panel of Medium or Large Disk Enclosure

Figure 3-18 shows the rear of the *Medium or Large Disk Enclosure*. Initially, the enclosure is supplied with a single I/O module that is installed on the right side.



Figure 3-18 IBM Cloud Object Storage Medium or Large Disk Enclosure rear view

Attention: During the planning phase, make sure that at least 1,200 mm (47.2 inch) deep racks are available for the enclosure, and that adequate space is available between the rack rows. You need a minimum of 5 inches for cabling. The enclosure is heavy, and it is advised to install it into lower rack positions. *Never* install the enclosure in the top positions of an empty rack because the rack might tip over if the enclosure slides out.

The enclosure includes a cable management arm, which enables the enclosure to slide out and the top lids to be removed if a disk must be replaced. Figure 3-19 shows the correct cabling. Both enclosure types are connected to the controller node by using four MiniSAS HD cables, regardless of the drive count. Therefore, the Medium Disk Enclosure has twice the bandwidth per drive when compared to the larger model.



Figure 3-19 Cable management arms and cabling

Drive population

The enclosure and the drives are packaged separately in the factory because of the weight. During installation, the enclosure must be racked first, and the drives are inserted later.

The Medium and the Large Disk enclosures support partial population of disk drives in Standard Dispersal (SD) Mode. The supported configurations and upgrade paths are listed in Table 3-4.

Table 3-4 Support configurations and upgrade paths

Disk enclosure	Supported number of drives in the enclosure	Upgrade options
Medium (J11)	14, 28, 41, 53	14 -> 28 14 -> 41 14 -> 53 28 -> 41 28 -> 53 41 -> 53
Large (J12)	64, 78, 92, 106	64 -> 78 64 -> 92 64 -> 106 78 -> 92 78 -> 106 92 -> 106

Note: Consider the following points:

- ▶ All drives must have the same capacity within a single enclosure.
For more information about the disk population rules, see the installation and maintenance manual that is available at [IBM Cloud Object Storage System](#).
- ▶ Partial population is not supported in Concentrated Dispersal (CD) Mode or for the Small Disk Enclosure.

Although the Medium Disk Enclosure features 106 disk slots, only 53 drives are inserted. Figure 3-20 on page 67 shows the correct drive population for the Medium Disk Enclosure when fully populated.

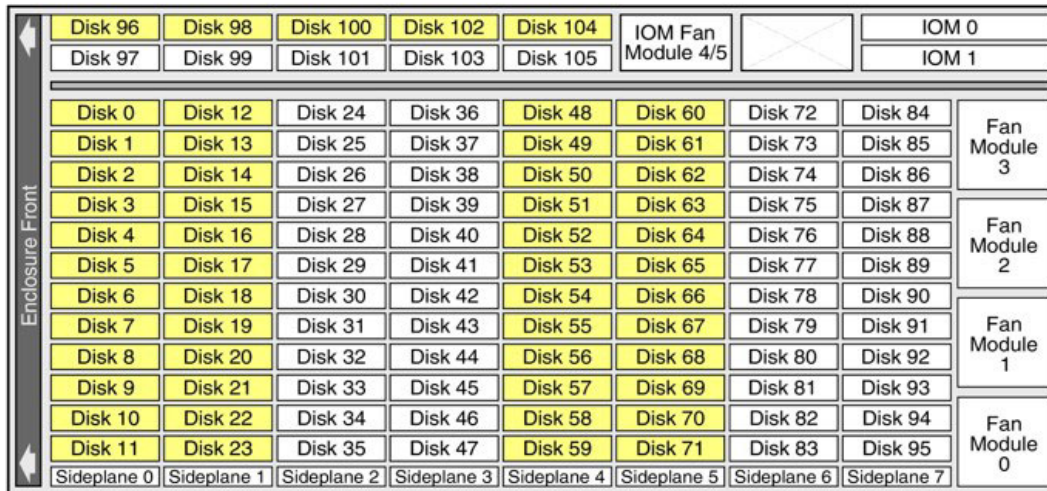


Figure 3-20 Full drive population for Medium Disk Enclosure (yellow highlighted slots to be populated)

Figure 3-21 shows the drive enumeration of the Large Disk Enclosure.

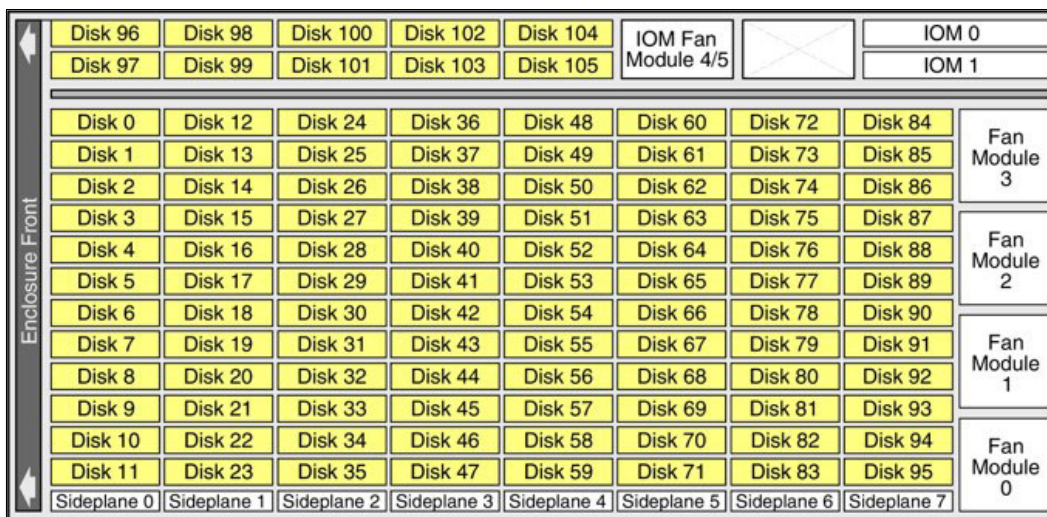


Figure 3-21 Large Disk Enclosure drive enumeration

5U92 Disk Enclosure

The 5U92 Disk Enclosure is the latest addition to the Gen2 hardware family. It is a high-density enclosure that sits between the Medium and the Large Disk Enclosures in terms of capacity. A main advantage is that it requires only a 1,000 mm deep rack, which allows more deployment scenarios when extra deep racks are not available at the client site.

The enclosure is 5U and can host up to 92 hot-swappable NL-SAS drives. The SliceStor 92 that uses this enclosure (although it has slightly more drives) is the successor of the SliceStor 2584 model.

The minimum raw capacity of the 92-bay enclosure is 552 TB when fully populated with 6 TB drives, and the maximum capacity is 1.84 PB when 20 TB disks are used. The drives can be added or removed from the top of the chassis. During installation, extra care is required for proper cabling because the chassis must slide out to swap failed drives.

Physical view and modules

Figure 3-22 shows the front of the 5U92 Disk Enclosure.



Figure 3-22 IBM Cloud Object Storage 5U92 Disk Enclosure front view

Figure 3-23 shows the rear of the 5U92 Enclosure.




Figure 3-23 IBM Cloud Object Storage 5U92 Disk Enclosure rear view

Attention: During the planning phase, make sure that at least 1,000 mm (39.4 inch) deep racks are available for the enclosure, and that adequate space is available between the rack rows. The enclosure is heavy, and it is advised to install it into lower rack positions. *Never* install the enclosure in the top positions of an empty rack because the rack might tip over if the enclosure slides out. Figure 3-24 shows the rack space requirements for the enclosure.




Figure 3-24 5U92 racking dimensions

The enclosure includes a cable management arm, which enables the enclosure to slide out and the top lids to be removed if a disk must be replaced. Figure 3-25 shows the correct cabling. The enclosure is connected to the controller node by using two MiniSAS HD cables, regardless of the drive count. The cables connect to the two ports of the expansion canister on the right side of the enclosure when viewed from the rear.




Figure 3-25 Cable management arms and cabling

Drive population

The enclosure and the drives are packaged separately in the factory because of the weight. During installation, the enclosure must be racked first, and the drives are inserted later.

The enclosure supports partial population of disk drives in SD Mode. The supported configurations and upgrade paths are listed in Table 3-5.

Table 3-5 Supported configurations and upgrade paths

Disk enclosure	Supported number of drives in the enclosure	Upgrade options
5U92 (J15)	28, 64, or 92	28 -> 64 28 -> 92 64 -> 92

Note: Consider the following points:

- ▶ All drives must have the same capacity within a single enclosure.
For more information about the disk population rules, see the documentation at [IBM Cloud Object Storage System](#).
- ▶ Partial population is not supported in CD Mode.

A label on the enclosure cover in Figure 3-26 shows the drive locations in the enclosure. The drive slots are numbered 1 - 14 from left to right and A - G from the back to the front of the enclosure. The drive slots must be populated sequentially, starting from the back-left position (slot 1, grid A1). Sequentially install the drive in the slots from left to right and back row to front. Always complete a full row before you install drives in the next row.




Figure 3-26 Drive locations in a 5U92 disk enclosure

3.3 Appliance specifications

Table 3-6 shows the Gen2 server appliances' specifications.

Table 3-6 Gen2 server appliances' specifications

Item	Manager M10	Accesser A10	Controller Node C10
CPU	Two Intel Xeon Silver 4410 or 4210R	Two Intel Xeon Gold 6126 or 6226	Two Intel Xeon Silver 4410 or 4210R
Number of cores/threads	8/16	12/24	8/16
Base frequency	2.1 GHz	2.6 GHz	2.1 GHz
Memory	192 GB	192 GB	192 GB
OS drives	Two 960 GB SSDs mirrored	Two 480 GB SSDs mirrored	Two 480 GB SSDs mirrored
RAID controller	Hardware RAID controller	Software RAID	Software RAID
RU	1	1	1
Width	43 mm (1.7 in.)	436 mm (17.2 in.)	436 mm (17.2 in.)
Height	43 mm (1.7 in.)	43 mm (1.7 in.)	43 mm (1.7 in.)
Length	762 mm (30 in.)	762 mm (30 in.)	762 mm (30 in.)
Weight (max)	21.7 kg (48 lbs)	21.7 kg (48 lbs)	21.7 kg (48 lbs)
Power supplies	Two 800 W	Two 800 W	Two 800 W
Power at startup	290 W	290 W	290 W
Power at max load	160 W	210 W	160 W
Input power	100 V - 240 V AC 50 Hz - 60 Hz	100 V - 240 V AC 50 Hz - 60 Hz	100 V - 240 V AC 50 Hz - 60 Hz
Supplied power connectors	Two C14 - C13	Two C14 - C13	Two C14 - C13
Operating temperature	10°C to 35°C (50°F to 95°F)	10°C to 35°C (50°F to 95°F)	10°C to 35°C (50°F to 95°F)
Nonoperating temperature	-40°C to 65°C (-40°F to 149°F)	-40°C to 65°C (-40°F to 149°F)	-40°C to 65°C (-40°F to 149°F)
Operating altitude	0 m - 3,000 m (0 ft - 10,000 ft)	0 m - 3,000 m (0 ft - 10,000 ft)	0 m - 3,000 m (0 ft - 10,000 ft)
Nonoperating altitude	0 m - 12,192 m (0 ft - 40,000 ft)	0 m - 12,192 m (0 ft - 40,000 ft)	0 m - 12,192 m (0 ft - 40,000 ft)
On-board networking	Two 10 GbE optical (SPFP+)	Two 10 GbE optical (SPFP+)	Two 10 GbE optical (SPFP+)

Item	Manager M10	Accesser A10	Controller Node C10
IPMI	Dedicated 1 GbE copper (RJ45)	Dedicated 1 GbE copper (RJ45)	Dedicated 1 GbE copper (RJ45)
Optional PCIe card	One 4-port Ethernet with 10 GbE SFP+ or 25 GbE SFP28 transceivers <i>or</i> one 4-port 10BaseT (RJ45)	One 4-port Ethernet with 10 GbE SFP+ or 25 GbE SFP28 transceivers <i>or</i> one 4-port 10BaseT (RJ45)	One 4-port Ethernet with 10 GbE SFP+ or 25 GbE SFP28 transceivers <i>OR</i> 1x 4-port 10BaseT (RJ45)

Table 3-7 shows the disk enclosures' specifications.

Table 3-7 Disk enclosures' specifications

Item	Small (J10)	Medium (J11)	Med to Large (J15)	Large (J12)
Max number of drives	12	53	92	106
Partial population	-	14/28/41/53	28/64/92	64/78/92/106
Interface	12 Gb SAS	12 Gb SAS	12 Gb SAS	12 Gb SAS
Drives	4, 8, 12, 14, 16, or 18 TB NL-SAS	4, 8, 12, 14, 16, or 18 TB NL-SAS	6, 8, 10, 12, 16, or 18 TB NL-SAS	4, 8, 12, 14, 16, or 18 TB NL-SAS
RUs	2	4	5	4
Width	483 mm (19 in.)	441 mm (17.4 in.)	438 mm (17.24 in.)	441 mm (17.4 in.)
Height	88 mm (3.5 in.)	176 mm (6.95 in.)	220 mm (8.66 in.)	176 mm (6.95 in.)
Depth	630 mm (24.8 inch)	1139 mm (44.8 in.)	1013 mm (39.88 in.)	1139 mm (44.8 in.)
Weight (max)	26 kg (52.7 lbs)	82 kg (181 lbs)	139 kg (306 lbs)	141 kg (310 lbs)
Weight (chassis only)	-	62 kg (140 lbs)	43 kg (94 lbs)	62 kg (140 lbs)
Power supplies	Two 580 W	Two 2000 W	Power rating (Platinum): 2400 W	Two 2000 W
Power at startup	300 W	1300 W	-	1800 W
Power at max load	180 W	1440 W	-	2150 W
Input power	100 V - 240 V AC 50 Hz - 60 Hz	200 V - 240 V AC 50 Hz - 60 Hz	200 V - 240 V AC 50 Hz - 60 Hz	200 V - 240 V AC 50 Hz - 60 Hz
Supplied power connectors	Two C14 to C13	Two C20 to C19	Two C20 to C19	Two C20 to C19
Operating temperature	-5°C to 40°C (41°F to 104°F)	5°C to 35°C (41°F to 95°F) (de-rated by 1°C per 300 m above 900 m)	5°C to 35°C (41°F to 95°F) up to 950 m (3,117 ft) above sea level. Above 950 m, de-rate maximum air temperature 1 degree per 175 m.	5°C to 35°C (41°F to 95°F) (de-rated by 1°C per 300 m above 900 m)
Nonoperating temperature	-40°C to 70°C (-40°F to 158°F)	-40°C to 70°C (-40°F to 158°F)	1°C - 60°C (33.8°F - 140°F)	-40°C to 70°C (-40°F to 158°F)
Operating altitude	0 m to 3,000 m (0 ft to 10,000 ft)	-100 m to 3,000 m (-328 ft to 10,000 ft)	-	-100 m to 3,000 m (-328 ft to 10,000 ft)

Item	Small (J10)	Medium (J11)	Med to Large (J15)	Large (J12)
Nonoperating altitude	-305 m to 12,192 m (-1,000 ft to 40,000 ft)	-100 m to 12,192 m (-328 ft to 40,000 ft)	-	-100 m to 12,192 m (-328 ft to 40,000 ft)
Operating humidity	20% - 80% noncondensing	20% - 80% noncondensing	8% - 85%	20% - 80% noncondensing

Tip: SliceStor 53 and 106 appliances require 1,200 mm (47.2 inch) deep racks for the enclosure (not including cables). The IBM Enterprise Slim Rack (MTM 7965-S42) can be used with 2x rear rack extenders (feature code ECRK for a total of 10 extra inches) installed.

For more information about the rack and the rear extenders, see the following web pages of IBM Documentation:

- ▶ [Planning for the 7965-S42 rack](#)
- ▶ [Model 7965-S42 supported feature codes](#)

3.4 Hardware options

The upgrade possibilities for the Gen2 appliances after general availability are described in this section.

3.4.1 Processor and memory upgrade

IBM introduced the option to install another processor to the Gen2 nodes. This upgrade automatically includes a 96 GB memory upgrade for the Manager and Controller nodes for 192 GB of total memory. The upgrade can be added to Gen2 nodes that are in the field. The field upgrade can be installed by IBM only.

Table 3-8 lists the performance gains that are associated with the CPU and memory upgrades. Typically, the extra CPU helps to significantly increase the number of operations for small objects. For more information about performance implications for different object sizes, contact your IBM Cloud Object Storage specialist.

Table 3-8 Performance gains that are associated with the CPU and memory upgrades

Node	CPU	Memory	Rating
Accesser (A10)	One Xeon Gold 6126	192 GB	Baseline
Accesser (A10)	Two Xeon Gold 6126	192 GB	1.35x baseline
Controller (C10)	One Xeon Silver 4110	96 GB	Baseline
Controller (C10)	Two Xeon Silver 4110	192 GB	1.15x baseline

3.4.2 Network interface upgrade

By default, the Gen2 nodes include two 10 GbE network adapters with SFP+ transceivers on their system boards. There also is a dedicated 1 GbE port on the appliances for IPMI (hardware management).

The two new 4-port Ethernet adapters are introduced for the Gen2 nodes, which enable more flexible configurations for multi-networking requirements. With the additional network cards, up to six high-speed Ethernet ports are available per node, what allows separation of the Management and internal Data channels on all nodes and the Client channel on the Accesser nodes. All channels can be configured with redundant ports with bonding.

The optical adapter (AJ5W 4 x 10/25 GbE Network Interface Controller) supports 10 GbE SFP+ and 25 GbE SFP28 transceivers, the other adapter (AJ5J Quad Port 10BaseT Network Interface Controller) supports 1/10GbE with RJ45 ports.

Tip: For the 4-port Ethernet adapter to be functional, a second CPU and memory DIMMs that are associated with it must be present in the node. Existing Gen2 nodes can be upgraded, but the installation of the CPU and memory must be done by IBM. Customers can opt to add the network adapter themselves.

3.5 Performance

The second-generation appliances incorporate recent CPU, memory, SAS controller, and drive technologies that generally result in higher overall performance. In this section, we briefly compare the appliances' individual performance characteristics to the previous generation models. For more information comprehensive performance sizing, see 2.2, "Performance planning" on page 37, or contact an IBM Cloud Object Storage subject matter expert.

3.5.1 Accesser performance

The Gen2 Accesser A10 uses a newer generation CPU and 1.5x more memory compared to the Accesser 3105. When 10 GbE connections are used, this increase results in up to 15% higher performance for read and write operations.

3.5.2 Slicestor performance

The following are some guidelines for Slicestor performance:

- ▶ Slicestor 12 versus Slicestor 2212A: For a generic workload consisting of an equal mix of read, write, and delete operations with object sizes 10 K - 100 M, this configuration results in up to 50% better performance.
- ▶ Slicestor 53 versus Slicestor 2448/3448: For a generic workload consisting of an equal mix of read, write, and delete operations with object sizes 10 K - 100 M, this configuration results in up to 30% better performance.

- ▶ Slicestor 106 versus Slicestor 2584: For a generic workload consisting of an equal mix of read, write, and delete operations with object sizes 10 K - 100 M, this configuration results in up to 100% better performance. Using the 2-PU configuration is a best practice for the Slicestor 106 device, and it is standard.
- ▶ Slicestor 92: Performance is identical to the Slicestor 106 with marginal degradation for disk-bound workloads. Using the 2-PU configuration is a best practice for the Slicestor 92 device, and it is standard.

3.6 Rack guidance

This section covers the minimum specifications for racking all IBM Cloud Object Storage appliances.

Due to numerous manufacturer and models of racks, it is impossible to list every rack that supports IBM Cloud Object Storage appliances. Instead, what is provided is the minimum specifications that your rack needs. Also, IBM does sell a rack capable of supporting IBM Cloud Object Storage appliances.

All Generation 2 IBM Cloud Object Storage appliances can be installed in a rack that has either round holes or square holes.

The section covers what specifications that you need to consider when determining whether your racks can support IBM Cloud Object Storage appliances and the number you need.

3.6.1 Appliance weight

The maximum weight that your rack can accommodate determines the number of IBM Cloud Object Storage appliances you can put in your rack.

See 3.3, “Appliance specifications” on page 71 for the weight of each appliance.

In addition to appliance weight, you need to consider the following items:

- ▶ Rack location

Is your rack on cement or raised floor? If your rack is on a raised floor, there usually is a max weight that the floor can support per tile.

- ▶ Extensions

Rack extensions add weight to the rack, but also can support more weight. Also, your rack might support floor extensions, which increase the max weight that the rack can accommodate.

- ▶ Controller placement

For all IBM Cloud Object Storage Slice Storages, place both the controller node and JBOD into the same rack.

- ▶ Accessories

Calculate the weight of all the accessories that you plan on putting in the rack. These accessories include but are not limited to network switches, cable management tracks, and network cables.

- ▶ Power

Make sure that you can supply enough power and power receptacles for all the IBM Cloud Object Storage appliances (and all the other equipment) in the rack.

- ▶ Expansion needs

If you are running a partially populated solution, adding hard disk drives add weight to the appliance. Remember to take that into consideration.

- ▶ Aisle width

Appliances are racked from the front of a rack, so you must make sure that you have enough distance between racks (the width of your aisle) to install an appliance. IBM Cloud Object Storage appliance JBODs need be inserted parallel into their rails, so you generally need the max length of your JBOD appliance as the aisle width.

3.6.2 Internal dimensions

To determine whether your racks accommodate the IBM Generation 2 appliances, there are four key pieces of information that you need to know:

- ▶ The amount of space between the front door of your rack and the internal post, also known as the front rack space.
- ▶ The amount of space that is needed between from the front of the front pot to the rear of the rear post. All IBM Generation 2 appliance rail kits have a minimum and maximum adjustment and the internal posts need to fit within the range.
- ▶ The length needed to accommodate the appliance between the front rail to the back of the appliance (including the cable management arm).
- ▶ The amount of space between the back of the appliance (include cable management arms) and the rear door of the rack, also as the rear rack space. This space accommodates the power cords. As a best practice, provide extra space if it is needed for any power, SAS, or networks cables. This space is in addition to the rear rack space that is listed in Table 3-9, and it is 75 mm - 125 mm (3 inches - 5 inches). If you have sufficient space between the internal post and the side of the rack, you might need less space.

Table 3-9 lists the minimum dimensions for each IBM Generation 2 appliance that you need to check your rack against.

Table 3-9 Minimum dimensions for each IBM Generation 2 appliance

Appliance	Front rack space	Rail kit adjustment (min / max)	Appliance length (front post to CM arm)	Rear rack space
M10	51 mm (2.01 inches)	595 mm / 795 mm (23.4 inches / 31.3 inches)	762 mm (30 inches)	70 mm (2.76 inches)
A10	51 mm (2.01 inches)	595 mm / 795 mm (23.4 inches / 31.3 inches)	762 mm (30 inches)	70 mm (2.76 inches)
C10	51 mm (2.01 inches)	595 mm / 795 mm (23.4 inches / 31.3 inches)	762 mm (30 inches)	70 mm (2.76 inches)
J10	51 mm (2.01 inches)	595 mm / 795 mm (23.4 inches / 31.3 inches)	630 mm (24.79 inches)	70 mm (2.76 inches)
J11	51 mm (2.01 inches)	689 mm / 889 mm (27 inches / 35 inches)	1139 mm (44.84 inches)	70 mm (2.76 inches)

Appliance	Front rack space	Rail kit adjustment (min / max)	Appliance length (front post to CM arm)	Rear rack space
J15	51 mm (2.01 inches)	680 mm / 800 mm (26.77 inches / 34.65 inches)	968 mm (38.11 inches)	70 mm (2.76 inches)
J12	51 mm (2.01 inches)	689 mm / 889 mm (27 inches / 35 inches)	1139 mm (44.84 inches)	70 mm (2.76 inches)

3.6.3 Power and PDU placement

The amount of power you can place into a rack, the number of outlets that you can supply and the placement of the PDUs can affect how many IBM Cloud Object Storage appliances and which ones you can fit into a rack.

For the amount of power that you per rack, see the appliance specification section. Each IBM Cloud Object Storage appliance has dual (two) power supplies, and you must supply enough power for the startup and while running the appliances.



Deployment options

In this chapter, we describe the different deployment options for IBM Cloud Object Storage.

This chapter includes the following topics:

- ▶ 4.1, “Introduction” on page 80
- ▶ 4.2, “IBM hardware appliances” on page 80
- ▶ 4.3, “Third-party appliances” on page 81
- ▶ 4.4, “Embedded Accesser” on page 81
- ▶ 4.5, “IBM Cloud Object Storage System virtual appliances” on page 82
- ▶ 4.6, “Appliance Docker Containers” on page 83

4.1 Introduction

The IBM Cloud Object Storage appliances can be deployed on physical hardware as a Docker container, or as a VMware appliance.

The first two sections of this chapter describe the IBM Cloud Object Storage hardware appliances and the IBM certified third-party hardware appliances on which IBM Cloud Object Storage can be installed.

In 4.4, “Embedded Accesser” on page 81, we describe how to enable the Embedded Accesser service on a Slicestor appliance.

In 4.5, “IBM Cloud Object Storage System virtual appliances” on page 82 and 4.6, “Appliance Docker Containers” on page 83, we explain how to install and configure the Manager and Accesser appliances as Docker containers and VMware appliances. Docker containers are quick to deploy, require less resources than virtual machines, and do not need a hardware hypervisor. However, a VMware environment provides more security and isolation.

4.2 IBM hardware appliances

As of December 2022, the following IBM hardware appliances (Gen2) were released for IBM Cloud Object Storage:

- ▶ Manager node

Manager M10 supports up to 4,500 simultaneous appliances in a single Cloud Object Storage Network.

- ▶ Accesser node

Accesser A10 provides up to 15% more read and write performance (it can complete up to 15% more reads and writes in the same time frame) as the Accesser 3105 appliance.

- ▶ Slicestor node

The following options are available:

- Slicestor 12: 3U device that contains 12 drives of 4, 8, 12, 16,18 or 20 TB.
- Slicestor 53: 5U device that contains 53 drives of 4, 8, 12, 16,18 or 20 TB.
- Slicestor 92: 6U device that contains 92 drives of 6, 8, 10, 12, 16,18 or 20 TB.
- Slicestor 106: 5U device that contains 106 drives of 4, 8, 12, 16,18 or 20 TB.

For more information about the specifications of these devices, see Chapter 3, “IBM Cloud Object Storage Gen2 hardware appliances” on page 53.

In addition, the following Gen1 IBM hardware appliances are still supported for IBM Cloud Object Storage:

- ▶ 3401/3403 – M01 (Manager 3105)
- ▶ 3401/3403 – A00 (Accesser 3105)
- ▶ 3401/3403 – A01 (Accesser 4105)
- ▶ 3401/3403 – S10 (Slicestor 2212A)
- ▶ 3401/3403 – S01 (Slicestor 2448)
- ▶ 3401/3403 – S02 (Slicestor 3448)
- ▶ 3401/3403 – S03 (Slicestor 2584)

For more information about the specifications of these devices, see [IBM Documentation](#).

4.3 Third-party appliances

In addition to the IBM hardware appliances, IBM certifies several third-party servers, such as HPE, Dell, and Lenovo. For more information about the supported third-party appliances, contact your IBM Sales Representative.

Important: Hardware appliances from different hardware vendors can be deployed as part of an IBM Cloud Object Storage System. The only limitation is that all the appliances in a device set must be identical.

4.4 Embedded Accesser

The Embedded Accesser feature provides customers an opportunity to save on expenses by using one physical appliance for Accesser and Slicestor appliance functions.

The Embedded Accesser functions can be enabled on an existing storage pool or newly created storage pool. When enabled, all the Slicestor appliances in the storage pool have the Embedded Accesser functions activated. Consider the following points before you use this function:

- ▶ Hardware with a minimum of 10 GbE interconnect and a memory capacity of 96 GB is advised for a full-scale deployment of Slicestor devices with Embedded Accesser functions.
- ▶ Performance effects because not all workloads are suited for Embedded Accesser functions:
 - Spread the load on all the available Embedded Accesser appliances.
 - Some degree of performance degradation occurs on all workloads with Embedded Accesser appliances.
 - Some workloads, such as small file writes, are more severely affected than others.
- ▶ Service restart. The Slicestor appliance, which handles user I/O traffic, is restarted when this function is enabled.

4.4.1 Enabling Embedded Accesser functions

To enable this feature on storage pools, complete the following steps.

1. Go to the storage pool that is being targeted for Embedded Accesser functions.
2. Click **Configure** to reconfigure the storage pool.
3. Click **Change**.
4. Select **Enable the embedded Accesser service on all the Slicestor devices belonging to this storage pool**.
5. Click **Update**.
6. Upon the activation of the configuration change, the Slicestor appliance restarts. Wait for all the Slicestors to restart before you resume I/O operations.

4.5 IBM Cloud Object Storage System virtual appliances

IBM Cloud Object Storage virtual appliances (vAppliances) are bundled as Open Virtual Appliance (OVA) images to be deployed on VMware vSphere 5.5 or later. Table 4-1 lists the available vAppliance types.

Table 4-1 Available vAppliance types

vAppliance type	Purpose	High availability option
vAccesser	Provides access to buckets over S3 protocol.	Deploy multiple vAccesser appliances with load-balancing.
vManager	Configures, monitors, and administers an IBM Cloud Object Storage System.	Add a second Manager device to use the Multiple Manager feature, install two vManagers for manual failover, or use VMware high availability (HA) for automated failover.

Attention: Virtual Slicestor devices are *not* supported in an IBM Cloud Object Storage environment.

4.5.1 Configuring the appliance environment

In this section, we describe how to configure the appliance environment.

Deploying the OVA template

IBM Cloud Object Storage virtual appliances require VMware vSphere Hypervisor ESXi 5.5 or later.

Log in to the vSphere Client:

1. Select **File** → **Deploy OVA Template**.
2. Respond to the queries with information that is specific to your deployment.
3. Click **Power On After Deployment** and then, click **Finish**.

Setting virtual machine hardware properties

After deploying the virtual appliance on the host system, the hardware properties can be modified on the virtual Manager and virtual Accesser appliance.

These settings suffice for a demonstration or a lightly loaded production IBM Cloud Object Storage. For systems with higher performance expectations, you must provision more resources.

Note: Contact IBM Cloud Object Storage Customer Support for advised settings for a specific use case.

Table 4-2 lists the minimum settings for all virtual appliances.

Table 4-2 Minimum virtual appliance requirements

Component	vAccesser	vManager
Memory (GB)	16	16
vCPU	2 (see Note below)	4
SCSI Controller 0	Paravirtual	Paravirtual
Hard disk 1 (GB) OS	Virtual drive (128 GB)	Virtual drive (256 GB)
Network adapter 1	1 GbE	1 GbE

Note: Two vCPUs are not recommended for production use (only for functions testing). If you see 50x errors, you might need to add more resources to the vAccessers.

Using vAccesser Appliances on VMware in a production environment allocates more memory than the minimum system requirements. The usage of 2-CPU 16-GB vAccesser Appliances should be only for a non-production environment.

The actual IOPS capability of any vAccesser Appliance configuration varies depending on the underlying virtualized hypervisor capability. However, increasing the number of vAccesser Appliances scales the IOPS capability linearly until you reach a bottleneck in the underlying hypervisor infrastructure.

Depending on the workload, the vAccesser Appliance capabilities may be increased for network, CPU, and memory capacity to address resource constraints. Increasing NIC capacity, CPU, and memory can improve performance if those resources reach saturation. This situation is only valid up to the hypervisor capabilities.

A workload heavier in small object (less than 1 MiB) writes benefits from more CPUs and memory.

vAccesser Appliances on VMware leverage a VMware kernel module, which produces little advantage with more than eight CPUs.

Starting the virtual appliance for the first time

Navigate to the Console after powering on the virtual appliance. The initial IBM settings appear. For more information about starting the virtual appliance for the first time, see 5.5, “Step 3: Configuring the appliance” on page 101.

4.6 Appliance Docker Containers

The Docker Appliance Container provides the Manager Node and Accesser Node as Docker-compatible Container images.

The Appliance Container image can be to customer-managed hardware on a customer-managed operating system. The Appliance Container can be monitored, managed, and deployed by using the Manager Web Interface and Manager API.

4.6.1 Benefits

Support for Docker images was developed to provide flexible deployment of a system Manager and Accesser. Customers can use their familiar hardware and operating systems to run IBM Cloud Object Storage System software along existing applications.

General benefits

Appliance Containers allow customers to:

- ▶ Use existing infrastructure instead of purchasing dedicated appliances.
- ▶ Use existing customer workflows that use Docker to deploy and manage other containers in their existing infrastructure.
- ▶ Apply the benefits of virtualization with lower overhead than a traditional hypervisor.

Accesser Container benefits

The Accesser Container allows customers to:

- ▶ Add Accesser capabilities to a Manager application with less on-device configuration than either an Accesser Application or the Accesser Appliance.
- ▶ Combines the flexibility of the Accesser Application with the manageability and ease-of-use of an Accesser Appliance.

Manager Container benefits

The Manager Container allows customers to:

- ▶ Migrate from Manager Appliance to Manager Container.
- ▶ Backup up the Manager Appliance to a Manager Container, which then can be automated.
- ▶ Use the Manager Container on any compatible generic hardware.

The Slicestor node is *not available* as an Appliance Container.

4.6.2 Workflow, use cases, and feature impact

An Appliance Container is in a customer-supplied operating system. An administrator can load an Appliance Container image into the Docker daemon then deploy one or more instances of that image as a container subject to networking restrictions.

A running Appliance Container supports all the same client interfaces as its respective Appliance. The Appliance Container appears in the Manager Web Interface in a nearly indistinguishable manner to the corresponding appliance.

Note: Docker Containers cannot be run on a device running the ClevOS operating system.

Increasing deployment flexibility

An Appliance Container can be run on the same physical hardware as other applications by using an existing operating system.

Accesser container augments or replaces the Accesser Application

The Accesser Container requires less hand-configuration than the Accesser Application, and can be managed by using the Manager Web Interface or Manager API. If the customer is not using a Docker-compatible Linux distribution, the Accesser Application is still available.

Version upgrades

Appliance Containers cannot be upgraded through the Manager Web Interface.

4.6.3 Accesser container

This section includes the following topics:

- ▶ Memory requirements
- ▶ Storage requirements

Memory requirements

The Accesser Container tries to scale memory usage automatically based on the amount of system memory. The maximum amount of memory that is allocated to Accesser Container can be set through the `MAX_MEMORY` environment variable. Table 4-3 lists the memory requirements for the Accesser and Manager containers.

Table 4-3 Memory requirements for Accesser containers

Deployment scenario	Suggested system memory (in GB)	Suggested MAX_MEMORY setting (in MB)
Server Class Systems	16+	4000
Server Class Systems (Using Vault Mirrors)	32+	8000

Attention: For scenarios where the file size is large (> 10 GB) and the number of concurrent uploads and downloads is large (> 50), contact IBM customer support for more guidance.

Storage requirements

The Accesser container needs a modest amount of storage for logs and other state information. As a best practice, use 60 GB.

4.6.4 Manager container

The Manager container tries to scale memory usage automatically based on the amount of memory in the system.

The maximum amount of memory that is allocated to the Manager container can be set through the `DS_MYSQL_MAX_MEMORY` and `DS_MANAGER_MAX_MEMORY` environment variables.

Table 4-4 lists the memory requirements for the Manager container.

Note: Approximately 2 GB more base memory is required in addition to the listed settings.

Table 4-4 Memory requirements for Manager containers

Deployment Scenario	Suggested System Memory (GB)	Suggested DS_MYSQL_MAX_MEMORY setting (MB)	Suggested DS_MANAGER_MAX_MEMORY setting (MB)
Server Class Systems	16+	25% of system memory	25% of system memory

Storage requirements

As a best practice, use 1 TB per 1,000 vaults.

4.6.5 System and network configuration

Servers that are running IBM Cloud Object Storage containers must be configured for clock synchronization through Network Time Protocol (NTP). The host operating system should synchronize to the same NTP server as all the other nodes in the IBM Cloud Object Storage System.

Important: Unlike the Manager appliance, the Manager container cannot provide NTP synchronization services. Devices managed by a Manager container must be configured to use an external NTP server.

For more information about NTP synchronization, see [IBM Documentation](#).

You can select your current version of the IBM Cloud Object Storage software by clicking **Change version or product**.

Network ports

Appliance containers need connectivity to all system nodes. Table 4-5 lists the port usage for Docker containers.

Table 4-5 Port usage for appliance use

Destination	Port	Purpose
Slicestor nodes	TCP 5000	Data operations and registry lookup
Slicestor nodes	TCP 7 (OPEN or REJECT)	Round-trip time calculation
NTP	TCP or UDP 123	NTP messaging (configured in the host)
Manager node	TCP 443 (by default)	Manager API vault usage query by way of HTTPS
Manager CNC	TCP 8088	Management control port (non-Manager nodes)

Note: TCP port 7 can be closed, but any firewall rules should send REJECT messages and not drop packets.

4.6.6 Configuring the appliance container

IBM's implementation of Docker supports the following network operation modes:

- ▶ `--net="host"`

The container shares networking with the host operating system. Only one appliance container that uses `--net="host"` can run at one time on a single host. When `--net="host"` is used, some ports that are used by services inside the container can conflict with ports opened by services in your hosts network namespace. These ports can be remapped by using the `DS_CNC_*` variables that are noted in container environment variables or reconfigured through the Manager Web Interface for the Accesser node service.

Restriction: The Docker host must not use localhost as its hostname when an appliance container is used.

► --net="bridge"

The container uses a separate network namespace from the host operating system. Docker automatically allocates an internal NAT-like network. Some ports must be published from the container to the host.

Note: The `hostname` parameter must be set to a valid hostname other than localhost.

Container ports

Appliance containers run services on several ports. Depending on the use of `--net="host"` versus `--net="bridge"`, some of these ports might need to be published to host ports to access these services from outside the Docker host by using the `-p` flag. Port forwarding is operated by inserting `"-p"` in the `dock run` command, as shown in the following example:

```
"-p <CNC_PORT>:8088"
```

`CNC_PORT` is the port that is defined by the user and to which it is forwarded. The forwarded port can be defined as any port that is not used.

Table 4-6 lists the required ports for running appliance containers.

Table 4-6 Required ports for appliance containers

TCP Port	Accesser container purpose	Manager container purpose
80	Accesser software HTTP	N/A
443	Accesser software HTTPS	Manager software HTTPS
8080	Accesser software HTTP	N/A
8088	Manager CNC services	N/A
8192	Device API	N/A
8443	Accesser software HTTPS	N/A

4.6.7 Deployment

Note: Docker commands must be run with root privileges. In the following examples, this requirement is represented by using the `sudo` command before each Docker command.

Prerequisites

The following prerequisites must be met for deployment:

- Docker-compatible Linux operating system installation
- NTP synchronization configured on the host operating system
- Docker 1.3 or later

API compatibility

Consider the following points regarding API compatibility:

- ▶ All APIs that supported by the Accesser appliance also are supported by the Accesser container.
- ▶ All APIs supported by the Manager appliance also are supported by the Manager container.

Creating a container

Tip: For more information about Docker parameters, see [IBM Documentation](#).

You can select your current version of the IBM Cloud Object Storage software by clicking **Change version or product**.

To create a container, complete the following steps:

1. Load the container image into Docker running on your server:

```
# cat clevos-3.14.3.65-accesser-container.tar.gz | sudo docker load
```
2. List the container images to find the repository or tag pair or image ID to start a container (see Example 4-1).

Example 4-1 Listing the container images

```
# sudo docker images
```

REPOSITORY	TAG	IMAGE ID	CREATED	VIRTUAL SIZE
clevos-manager	3.14.3.65	bd60b4c172a3	2 weeks ago	2.64GB
clevos-accesser	3.14.3.65	4556463ecdd9	2 weeks ago	2.05GB


3. Start an appliance container, as shown in Example 4-2.

Example 4-2 Starting an appliance container

```
# sudo docker run -itd --net=host -v  
/var/lib/accesser-container01:/container-persistence:rw --hostname  
"docker-accesser9" --name="docker-accesser09"  
--env="DS_MANAGER_IP=192.168.99.11" --env="DS_MANAGER_AUTOACCEPT=true"  
clevos-accesser:3.14.3.65
```

```
bf6077050d639fa8e0c2ef48ab964fcf9d2682645670c3437247d6e2391e5a4c  
This is the container ID.
```

4. Approve the container instance in the Manager web interface (see Figure 4-1).



Devices Pending Approval: 1				Bulk Approve / Deny
Select devices to approve or deny from the system.				
<input type="checkbox"/>	Hostname	IP /address	Device type	Registered
<input type="checkbox"/>	docker-accesser9	192.168.99.7	accesser	2019-05-03 09:25:18 GMT

Figure 4-1 Device approval

Stopping a running container

Enter the **docker stop** command with the container ID in the CLI to stop the container:

```
# sudo docker stop bf6077050d63 -t 600
```


Resuming a stopped container

Enter the **docker start** command with the container ID in the CLI to resume the container:

```
# sudo docker start bf6077050d63
```

Running an interactive shell

To troubleshoot or debug a container, enter the **docker exec** command with the container ID and a shell file and path in the CLI:

```
# sudo docker exec -it bf6077050d63 /bin/bash
```

Note: If the **-i** and **-t** parameters are not specified in the original **docker run** statement when starting the container, terminal-related error messages might be displayed while trying to use commands inside the container.

Upgrading a container

Containers cannot be upgraded through the Manager device Web Interface. A container must be upgraded on the server on which it runs.

Note: The previous container must be run with a persistent volume that is mounted into container-persistence with **-v**, as described in “Creating a container” on page 88.

To upgrade a container, complete the following steps:

1. Load the new container image into Docker:

```
# gunzip -c clevos-3.14.3.81-accesser-container.tar.gz | sudo docker load
```

2. Stop the old container:

```
# sudo docker stop 27c60234bf89
```

3. Run the new container image by using the same persistent volume, environment variables, and hostname that is used for the previous container.

Note: If the hostname (**--hostname**) is not specified when **--net="host"** is used, the container inherits the hostname of the host OS.

```
# sudo docker run -itd --env="DS_MANAGER_IP=192.168.99.11"
--env="DS_MANAGER_AUTOACCEPT=true" -v
/home/data/container-data-1:/container-persistence:rw --net="host"
clevos-accesser:3.14.3.65
```

4. When the new container starts, remove the old container:

```
# sudo docker rm 27c60234bf89
```

Attention: This action removes the container instance from the Docker application, but does not remove the container image from the host operating system. That process must be done separately.

Converting from a physical or virtual Manager device to a Manager container

Complete the following steps:

1. Back up the Manager appliance.
2. Configure and start a Manager container instance that is running the same software version as the Manager appliance.

3. Restore the backup of the Manager appliance on the Manager container.
4. Reconfigure any appliances that are joined to the system that are served by the Manager appliance when an IP address or port changed regarding where Manager services originate.

Converting from Manager container to a physical or virtual Manager appliance

Complete the following steps:

1. Back up the Manager container.
2. Image and configure a Manager appliance that is running the same software version as the Manager container.
3. Restore the backup of the Manager container on the Manager appliance.
4. Reconfigure any appliances that are joined to the system that are served by the Manager container when an IP address or port changed regarding where Manager services originate.

Upgrading a system that is managed by a Manager container

A Manager container cannot be upgraded through the normal Manager UI orchestration. Upgrading of non-container devices in a Manager container system requires that the Manager container is upgraded externally before the Manager device Web Interface is used to upgrade the remaining supported devices.

Complete the following steps:

1. Upgrade the Manager container to the wanted version as described in “Upgrading a container” on page 89.
2. Upgrade any Accesser container instances to the wanted version as described in “Upgrading a container” on page 89.
3. Use the Manager Web Interface to upgrade the remaining hardware or virtual devices.

Migrating container devices to IPv6

The appliance containers support IPv6. Customers can migrate container devices to IPv6 to guard against the eventual depletion of IPv4 addresses and ensure that no outage in continuous connectivity occurs.

The solution can run several network implementations. For more information the different network implementations and the procedure to migrate to IPv6, [IBM Documentation](#).

You can select your current version of the IBM Cloud Object Storage software by clicking **Change version or product**.

Monitoring

Monitoring the appliance container in the Manager web interface is nearly identical to monitoring an appliance.

Because it is a software solution, the appliance container does not provide generalized hardware- and operating system-level monitoring.

Statistics that are not displayed (or provided in the Manager REST API or the Device API) include, but are not limited to, the following data:

- ▶ CPU temperatures
- ▶ Disk I/O
- ▶ Disk temperatures
- ▶ Fan speeds

Events and monitoring that are not performed include the following information:

- ▶ RAID monitoring
- ▶ CPU/Disk temperature alerts
- ▶ Device restart events
- ▶ Kernel dump events (does not apply to containers)

Some graphs and stats for appliance containers might behave differently than expected, including CPU usage and system load. These graphs and stats reflect the CPU usage and system load of the host machine as a whole, not an individual appliance container.

Network I/O graphs reflect the interfaces visible to the appliance container and are different depending on the containers network settings (`docker run --net={"host" | "bridge"}`).

Extra rules and restrictions

Consider the following extra rules and restrictions:

- ▶ The operator cannot restart a device from within a container.
- ▶ The Network Utility Tool (nut) commands do not work in a container.
- ▶ Unlike the Manager Appliance, the Manager container does not provide NTP synchronization services to other devices.
- ▶ NTP must be configured on each Docker appliance individually.
- ▶ With Docker containers, you cannot separate the data channel and client channel. Only the data channel can be used.

Troubleshooting

Container logs can be viewed in the following locations:

- ▶ On the host system of the persistent volume that is mounted into the container. If the container was started with `-v home/admin/container-data:/container-persistence:rw`, the logs are visible in `/home/admin/container-data/logs/`.
- ▶ In the container: `/container-persistence/logs` or `/var/log`.

In addition to logging, only non-hardware exceptions when various APIs are used show as events in the Manager device Web Interface.

For more information about troubleshooting a Docker environment, see [IBM Documentation](#).

You can select your current version of the IBM Cloud Object Storage software by clicking **Change version or product**.



Initial setup and configuration

In this chapter, we describe how to configure an IBM Cloud Object Storage System solution from the beginning to the point where you can use the system with your application.

This chapter includes the following topics:

- ▶ 5.1, “Needed installation information” on page 94
- ▶ 5.2, “Example information” on page 95
- ▶ 5.3, “Step 1: Installing the solution” on page 96
- ▶ 5.4, “Step 2: Installing IBM Cloud Object Storage Appliance software” on page 97
- ▶ 5.5, “Step 3: Configuring the appliance” on page 101
- ▶ 5.6, “Step 4: Configuring the Manager GUI” on page 106
- ▶ 5.7, “Step 5: Verifying the solution” on page 139
- ▶ 5.8, “Basic installation troubleshooting” on page 142
- ▶ 5.9, “IBM Call Home and log collection” on page 144
- ▶ 5.10, “Upgrading your IBM Cloud Object Storage instance” on page 150
- ▶ 5.11, “Configuring multiple Manager devices” on page 154
- ▶ 5.12, “Multi-factor authentication” on page 158
- ▶ 5.13, “Enabling Write Once Read Many capabilities” on page 162

5.1 Needed installation information

This section provides configuration steps for IBM Cloud Object Storage appliances: Manager, Accesser, and Slicestor. These steps apply to virtual Manager and Accesser appliances and physical Manager, Accesser, and Slicestor appliances (nodes).

5.1.1 Required information

For installation and configuration, the following information is required:

- ▶ Manager:
 - Hostname
 - Data channel information:
 - Data channel IP address
 - Data channel netmask address
 - Data channel gateway address
 - Network Time Protocol (NTP) servers
 - Domain Name System (DNS) servers
- ▶ Accesser appliance:
 - Hostname
 - Data channel information:
 - Data channel IP address
 - Data channel netmask address
 - Data channel gateway address
- ▶ Slicestor appliance:
 - Hostname
 - Data channel information:
 - Data channel IP address
 - Data channel netmask address
 - Data channel gateway address

Important: It is possible for the Manager, Accesser, and Slicestor appliances to have different values for any of this information. The specific information is based on your network location and physical locations.

5.1.2 Optional information

The following informational is optional and not required for all appliances:

- ▶ Maximum transmission unit (MTU) size for each network channel (defaults to 1500)
- ▶ Bonding requirements for all channels: requires Link Aggregation Control Protocol (LACP) enabled switch
 - Bonding (Yes/No)
 - Bonding type (Active-Active or Active-Passive)

- ▶ Management channel information:
 - Management channel IP address
 - Management channel netmask address
 - Management channel gateway address
- ▶ Client channel information:
 - Client channel IP address
 - Client channel netmask address
 - Client channel gateway address
- ▶ Organization name
- ▶ Appliance location:
 - City
 - State
 - County
- ▶ Intelligent Platform Management Interface (IPMI) information:
 - IPMI IP address
 - IPMI netmask address
 - IPMI gateway address

Important: It is possible for the Manager, Accesser, and Slicestor appliances to have different values for any of the information that is listed in this section. The specific information is based on your network location and physical locations.

5.2 Example information

For the installation and configuration in the rest of this chapter, the following information is used:

- ▶ Manager:
 - Hostname: icospdcm01
 - Data channel information:
 - Data channel IP address: 10.69.70.100
 - Data channel netmask address: 255.255.255.0
 - Data channel gateway address: 10.69.70.1
 - NTP Server: 10.69.70.10
 - DNS Server: 10.69.70.11
- ▶ Accesser appliance:
 - Hostname: icospdcacc01
 - Data channel information:
 - Data channel IP address: 10.69.70.110
 - Data channel netmask address: 255.255.255.0
 - Data channel gateway address: 10.69.70.1

- ▶ SliceStor appliance:
 - Hostname: icospdcsli01
 - Data channel information:
 - Data channel IP address: 10.69.70.130
 - Data channel netmask address: 255.255.255.0
 - Data channel gateway address: 10.69.70.1

This chapter does not cover the configuration of any of the optional information. More information about configuration can be found on [IBM Documentation](#).

Note: You can select your current version of the IBM Cloud Object Storage software by clicking **Change version or product**.

5.3 Step 1: Installing the solution

This section discusses the steps for installing the solution.

The IBM Cloud Object Storage appliance software (ClevOS) comes preinstalled on all Generation 2 IBM Cloud Object Storage appliances and some pre-Generation 2 IBM Cloud Object Storage appliances.

If you need to reinstall ClevOS or you plan on running ClevOS on IBM certified non IBM hardware, follow the installation instructions in 5.4, “Step 2: Installing IBM Cloud Object Storage Appliance software” on page 97. Otherwise, skip 5.4, “Step 2: Installing IBM Cloud Object Storage Appliance software” on page 97 and continue to 5.5, “Step 3: Configuring the appliance” on page 101.

Review the version installed to determine its release date. If the preinstalled version is within six months of the version that is recommended, upgrade to the recommended version after you configure the system *but before* you begin using the solution.

5.3.1 Physical appliance

This section *does not* cover how to rack a physical appliance. For more information about racking an appliance, see [IBM Documentation](#).

You can select your current version of the IBM Cloud Object Storage software by clicking **Change version or product**. Then, select the manual of the specific appliance that you must rack.

5.3.2 Virtual appliance

This section does *not* cover how to import a virtual appliance into VMware. For more information about how to import an Open Virtual Appliance (OVA), see “Deploying the OVA template” on page 82.

5.3.3 Container appliance

This section does *not* cover how to install a container infrastructure or the installation and configuration of the Manager or Accesser container. For more information about this process, see 4.6, “Appliance Docker Containers” on page 83.

5.4 Step 2: Installing IBM Cloud Object Storage Appliance software

Installation of the IBM Cloud Object Storage Appliance software requires console access to the appliance and the installation software.

Important: The term *IBM Cloud Object Storage appliance software* and the term *ClevOS* are interchangeable. The term ClevOS is the original name of the appliance software and is still used in the IBM Cloud Object Storage manuals that are available at [IBM Documentation](#).

Console access to the appliance is achieved by connecting a monitor and keyboard, KVM switch, or KVM device to the appropriate connections on rear of the appliance.

If IPMI is configured on the system, it is possible to connect to the appliance by way of the IPMI interface and open a console on the system.

The IBM Cloud Object Storage appliance installation software is provided by IBM upon purchase of the software. Contact your sales team for a copy of the software. Complete the following steps:

1. For installation purposes, the file you need is in the following format:

```
clevos-a.bb.cc.ddd-allinone-usbiso.iso
```

a.bb.cc.ddd references the specific version. Contact your sales team or IBM support to determine which version you should install. You want to make sure that you have the allinone-usbiso file. This file contains the code for the Manager, Accesser, and Slicestor appliances on the same installation media.

2. After obtaining the ISO installation media, copy the ISO file to a USB flash drive that is at least 8 GB by using any image burning software. An open source imaging software that is available at [balenaEtcher.com](https://balena.io/etcher/) is available for Mac, Windows, and Linux. The software also is known to work with the IBM Cloud Object Storage appliance software.
3. After connecting to the console and the USB flash drive, you must start the system from the USB flash drive. During the start process, whenever you see the white IBM splash window, press **F11**. Eventually, you are prompted with a blue startup window with options. At the bottom of the list, choose the USB flash drive to start from and press Enter.

You know you selected the correct drive when you are presented with window that is shown in Figure 5-1.



Figure 5-1 IBM Cloud Object Storage appliance installation boot window

You are prompted with the following screen, as shown in Example 5-1.

Example 5-1 IBM Cloud Object Storage appliance installation screen

```
*****  
IBM Cloud Object Storage System (r) Installer  
*****  
  
#1. Perform automatic installation  
#2. Perform manual installation  
#3. Reboot  
Choose Action: (1-3):
```

4. For an initial installation, select 1 and then, press Enter.

You are prompted for the type of installation you want, as shown in Example 5-2.

Example 5-2 IBM Cloud Object Storage appliance installation erase disk choice

```
*****  
IBM Cloud Object Storage System (r) Installer  
*****  
  
#1. OS Disk Only (Erase only OS disk and install)  
#2. Factory Install (Erase all disks and install)  
Choose Action: (1-2):
```

5. For an initial installation, select **2** and press Enter.

You are prompted to erase all disks. On a Slicestor, this erasure includes all data disks (see Example 5-3).

Example 5-3 Erasing all the disks

WARNING: This option will erase all disks attached to the system.
Enter 'erase' (no quotes) to confirm. Other input will cancel:

6. For an initial installation, enter erase and then, press Enter.

You are prompted about which type of appliance you are installing (see Example 5-4).

Example 5-4 Choosing which appliance image to install

```
*****  
IBM Cloud Object Storage System (r) Installer  
*****
```

```
#1. CLEVOS-3.14.3.65-ACCESSER  
#2. CLEVOS-3.14.3.65-MANAGER  
#2. CLEVOS-3.14.3.65-SLICESTOR  
Choose Action: (1-3):
```

7. Select **1**, **2**, or **3** (depending on which appliance you are on) and then, press Enter.
8. The installer then performs the following steps:
 - a. Partitions and formats the drives.
 - b. Copies files and verifies the integrity of the installation.
 - c. Performs other checks.
 - d. Restarts the appliance.
9. After the appliance restarts, you see the BIOS window. Then, you can remove the USB flash drive.

You then see the boot window, as shown in Figure 5-2.



Figure 5-2 IBM Cloud Object Storage appliance boot window

After the appliance completely starts, you see the information that is shown in Example 5-5.

Example 5-5 IBM Cloud Object Storage appliance CLI login screen

```
IBM Cloud Object Storage
Appliance - ClevOS 3.14.3.65 tty1
Slicestor login:
```

The appliance is now installed.

5.5 Step 3: Configuring the appliance

This section describes how to complete the initial configuration of the appliance after the IBM Cloud Object Storage appliance software is installed. The default username and password is the same for all appliances after the initial installation. The default username is localadmin and the default password is password.

At the login prompt, enter the default username and password.

Important: SSH access is disabled by default on all IBM Cloud Object Storage appliances. To enable SSH access to the appliance, change the default password for the user localadmin. IBM advises that you change the password immediately.

5.5.1 Configuring the Manager

Complete the following steps to configure the Manager:

1. Log in to the Manager. You see the prompt that is shown in Example 5-6.

Example 5-6 Manager appliance default prompt

```
manager#
```

2. Run the **port list** command to determine which Ethernet ports are available to you (see Example 5-7).

Example 5-7 Listing the available ports on the Manager

```
manager# port list
PORT ADDRESS MAX SPEED STATUS
eth0 d2:8a:ba:26:9c:86 Unknown disconnected
```

3. On your appliance, you see more than one port. The names of the ports also might be different. Identify to which ports your Ethernet cable is connected.
4. Enter edit to begin configuring the appliance, and then enter all the configuration information (the output is shown in Example 5-8).

Example 5-8 Setting the required information for a Manager

```
manager# edit
manager (working)# channel data port eth0
manager (working)# channel data ip 10.69.70.100
manager (working)# channel data netmask 255.255.255.0
manager (working)# channel data gateway 10.69.70.1
manager (working)# system hostname icospdcmn01
manager (working)# system ntpservers 10.69.70.10
manager (working)# system dns 10.69.70.11
manager (working)#
```

5. Save and activate the appliance configuration, as shown in Example 5-9.

Example 5-9 Activating the Manager configuration

```
manager (working)# activate
Please wait, this may take a few minutes....
check OK
activate OK
icospdcmn01#
```

6. The prompt changes to the hostname. If you do not receive any errors, you should be fine. You can test your network connection to see whether you can ping the appliance's gateway, as shown in Example 5-10.

Example 5-10 Testing the Manager network configuration

```
icospdcmn01# ping -c 10 10.69.70.1
PING 10 10.69.70.1 (10 10.69.70.1) 56(84) bytes of data.
64 bytes from 10 10.69.70.1: icmp_seq=1 ttl=64 time=0.255 ms
64 bytes from 10 10.69.70.1: icmp_seq=2 ttl=64 time=0.394 ms
64 bytes from 10 10.69.70.1: icmp_seq=3 ttl=64 time=0.345 ms
64 bytes from 10 10.69.70.1: icmp_seq=4 ttl=64 time=0.417 ms
64 bytes from 10 10.69.70.1: icmp_seq=5 ttl=64 time=1.03 ms
64 bytes from 10 10.69.70.1: icmp_seq=6 ttl=64 time=11.7 ms
64 bytes from 10 10.69.70.1: icmp_seq=7 ttl=64 time=1.9 ms
64 bytes from 10 10.69.70.1: icmp_seq=8 ttl=64 time=0.615 ms
64 bytes from 10 10.69.70.1: icmp_seq=9 ttl=64 time=0.400 ms
64 bytes from 10 10.69.70.1: icmp_seq=10 ttl=64 time=0.398 ms

--- 10 10.69.70.1 ping statistics ---
10 packets transmitted, 10 received, 0% packet loss, time 9133ms
rtt min/avg/max/mdev - 0.255/1.752/11.764/3.370 ms
```

The Accesser and Slicestor appliances are configured as described next.

5.5.2 Configuring the Accesser appliance

Complete the following steps to configure the Accesser appliance:

1. After logging in to the Accesser appliance, you should see the prompt that is shown in Example 5-11.

Example 5-11 Accesser appliance default prompt

```
accesser#
```

2. Run the **port list** command to determine which Ethernet ports are available to you (see Example 5-12).

Example 5-12 Listing the available ports on an Accesser appliance

```
accesser# port list
PORT ADDRESS MAX SPEED STATUS
eth0 d2:8a:ba:26:9c:86 Unknown disconnected
```

3. On your appliance, you see more than one port. The names of the ports also might be different. Identify to which port your Ethernet cable is connected.

4. Enter edit to begin configuring the appliance. Also, enter all the configuration information (see Example 5-13).

Example 5-13 Setting the required information for an Accesser appliance

```
accesser# edit
accesser (working)# channel data port eth0
accesser (working)# channel data ip 10.69.70.110
accesser (working)# channel data netmask 255.255.255.0
accesser (working)# channel data gateway 10.69.70.1
accesser (working)# system hostname icospdcacc01
accesser (working)# system dns 10.69.70.11
accesser (working)# manager ip 10.69.70.100
ERROR: couldn't retrieve manager certificate: curl returned exit code 7
Automatically accept the manager certificate when it is available? [y/N]: y
Enter prefix of manager fingerprint to verify (press enter to skip)
>
accesser (working)#
```

5. You see that curl returned an error while attempting to retrieve the manager certificate. This error is normal because the network is not yet activated. Save and activate the appliance configuration, as shown in Example 5-14.

Example 5-14 Activating the Accesser configuration

```
accesser (working)# activate
Please wait, this may take a few minutes....
check OK
activate OK
icospdcacc01#
```

6. Notice that the prompt changed to the hostname. If you do not receive any errors, you should be fine. You can test your network connection to see whether you can ping the appliance's gateway, as shown in Example 5-15.

Example 5-15 Testing the Accesser appliance network configuration

```
icospdcacc01# ping -c 10 10.69.70.1
PING 10 10.69.70.1 (10 10.69.70.1) 56(84) bytes of data.
64 bytes from 10 10.69.70.1: icmp_seq=1 ttl=64 time=0.255 ms
64 bytes from 10 10.69.70.1: icmp_seq=2 ttl=64 time=0.394 ms
64 bytes from 10 10.69.70.1: icmp_seq=3 ttl=64 time=0.345 ms
64 bytes from 10 10.69.70.1: icmp_seq=4 ttl=64 time=0.417 ms
64 bytes from 10 10.69.70.1: icmp_seq=5 ttl=64 time=1.03 ms
64 bytes from 10 10.69.70.1: icmp_seq=6 ttl=64 time=11.7 ms
64 bytes from 10 10.69.70.1: icmp_seq=7 ttl=64 time=1.9 ms
64 bytes from 10 10.69.70.1: icmp_seq=8 ttl=64 time=0.615 ms
64 bytes from 10 10.69.70.1: icmp_seq=9 ttl=64 time=0.400 ms
64 bytes from 10 10.69.70.1: icmp_seq=10 ttl=64 time=0.398 ms

--- 10 10.69.70.1 ping statistics ---
10 packets transmitted, 10 received, 0% packet loss, time 9133ms
rtt min/avg/max/mdev - 0.255/1.752/11.764/3.370 ms
```

The rest of your Accesser and Slicestor appliances can now be configured, as described next.

5.5.3 Configuring the Slicestor appliance

Complete the following steps to configure the Slicestor appliance:

1. Log in to the Slicestor appliance. You should see the prompt that is shown in Example 5-16.

Example 5-16 Slicestor appliance default prompt

```
slicestor#
```

2. Run the **port list** command to determine which Ethernet ports are available to you (see Example 5-17).

Example 5-17 Listing the available ports on an Accesser appliance

```
slicestor# port list
PORT ADDRESS  MAX SPEED STATUS
eth0 d2:8a:ba:26:9c:86 Unknown disconnected
```

3. On your appliance, you see more than one port. The names of the ports might be different as well. Identify to which ports your Ethernet cable is connected.
4. Enter edit to begin configuring the appliance. Also, enter all the configuration information (see Example 5-18).

Example 5-18 Setting the required information for a Slicestor appliance

```
slicestor# edit
slicestor (working)# channel data port eth0
slicestor (working)# channel data ip 10.69.70.130
slicestor (working)# channel data netmask 255.255.255.0
slicestor (working)# channel data gateway 10.69.70.1
slicestor (working)# system hostname icospdcs1c01
slicestor (working)# system dns 10.69.70.11
slicestor (working)# manager ip 10.69.70.100
ERROR: couldn't retrieve manager certificate: curl returned exit code 7
Automatically accept the manager certificate when it is available? [y/N]: y
Enter prefix of manager fingerprint to verify (press enter to skip)
>
accesser (working)#
```

5. Notice that curl returned an error while it was attempting to retrieve the manager certificate. This error is normal because the network is not yet activated. Save and activate the appliance configuration, as shown in Example 5-19.

Example 5-19 Activating the Slicestor appliance configuration

```
slicestor (working)# activate
Please wait, this may take a few minutes....
check OK
activate OK
icospdcs1c01#
```

6. Notice that the prompt changed to the hostname. If you do not receive any errors, you should be fine. Test your network connection to see whether you can ping the appliance's gateway (see Example 5-20).

Example 5-20 Testing the Slicestor network configuration

```
icospdcs1c01# ping -c 10 10.69.70.1
PING 10 10.69.70.1 (10 10.69.70.1) 56(84) bytes of data.
64 bytes from 10 10.69.70.1: icmp_seq=1 ttl=64 time=0.255 ms
64 bytes from 10 10.69.70.1: icmp_seq=2 ttl=64 time=0.394 ms
64 bytes from 10 10.69.70.1: icmp_seq=3 ttl=64 time=0.345 ms
64 bytes from 10 10.69.70.1: icmp_seq=4 ttl=64 time=0.417 ms
64 bytes from 10 10.69.70.1: icmp_seq=5 ttl=64 time=1.03 ms
64 bytes from 10 10.69.70.1: icmp_seq=6 ttl=64 time=11.7 ms
64 bytes from 10 10.69.70.1: icmp_seq=7 ttl=64 time=1.9 ms
64 bytes from 10 10.69.70.1: icmp_seq=8 ttl=64 time=0.615 ms
64 bytes from 10 10.69.70.1: icmp_seq=9 ttl=64 time=0.400 ms
64 bytes from 10 10.69.70.1: icmp_seq=10 ttl=64 time=0.398 ms

--- 10 10.69.70.1 ping statistics ---
10 packets transmitted, 10 received, 0% packet loss, time 9133ms
rtt min/avg/max/mdev - 0.255/1.752/11.764/3.370 ms
```

The rest of your Accesser and Slicestor appliances can be configured next.

5.6 Step 4: Configuring the Manager GUI

After you install ClevOS on all of your appliances, you must configure the GUI. This process is done by using the Manager GUI.

For the rest of the chapter, it is assumed that the IBM Cloud Object Storage environment that is shown in Figure 5-3 is used.

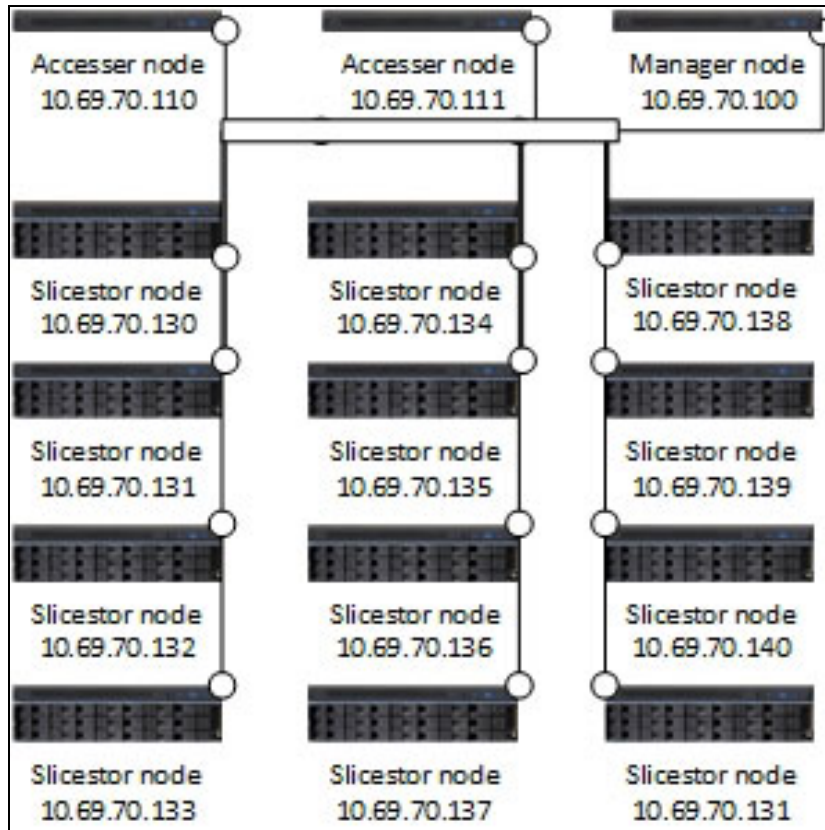


Figure 5-3 Example solution

5.6.1 Initial login

Complete the following steps for the initial login:

1. Open a web browser and go to the Manager by using the following URL:

`https://icospdcmn01`

Tip: You can use the IP address or the hostname to log in to the Manager GUI. The hostname is available only if you entered the hostname into your organization's DNS server.

You see the window that is shown in Figure 5-4 on page 107.

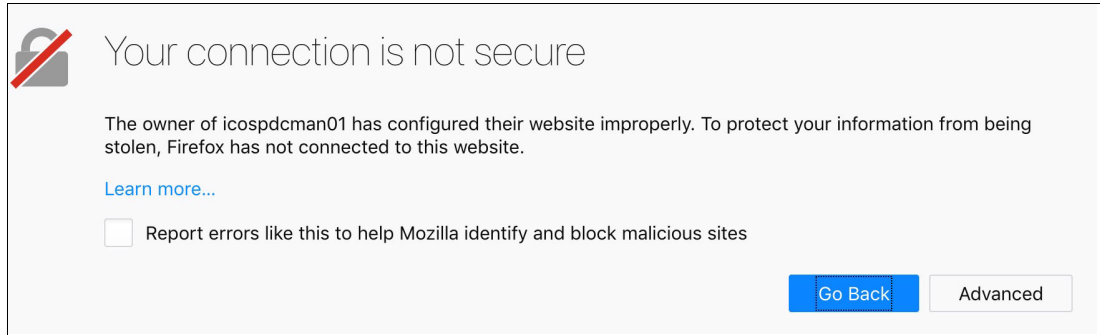


Figure 5-4 Default HTTPS certificate error

2. Accept the default HTTPS certificate that is generated by IBM Cloud Object Storage to continue.

Note: The steps regarding how to accept a user-generated certificate into a specific browser is out of the scope of this document. For more information, see your web browser's documentation.

After accepting the certificate, the login window opens (see Figure 5-5).

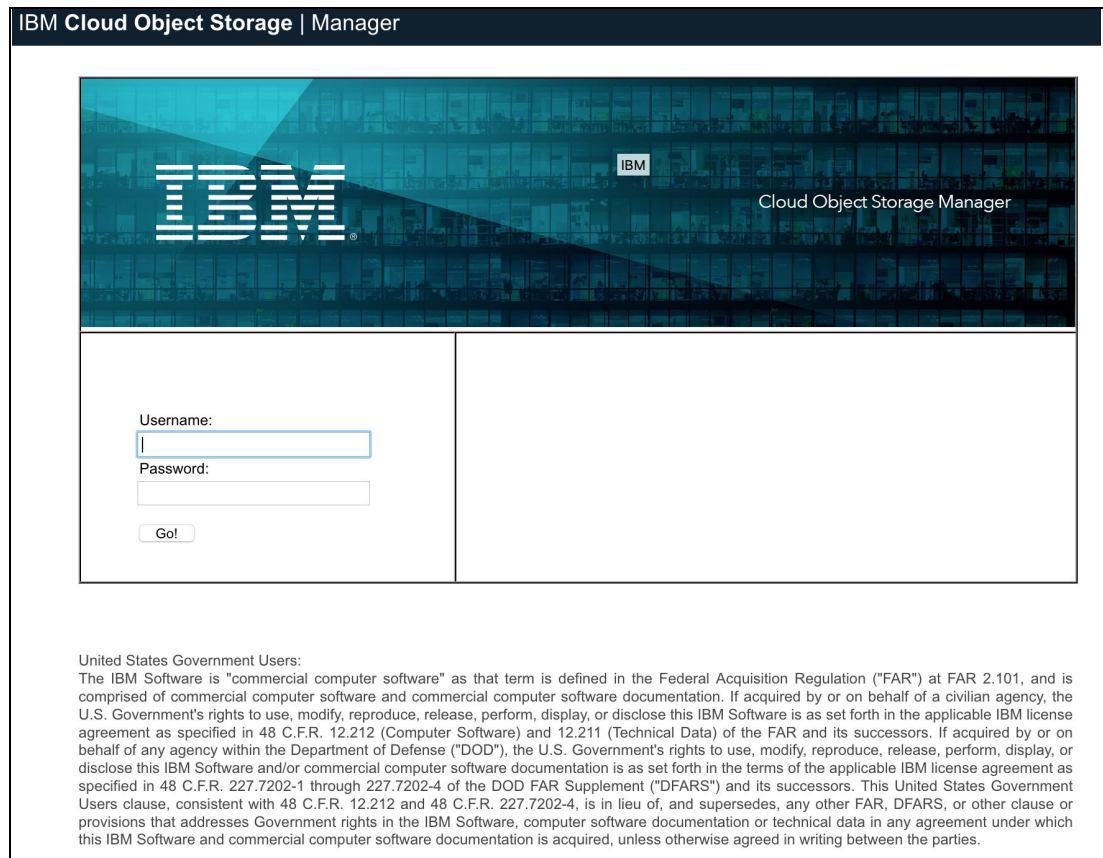


Figure 5-5 Manager GUI login window

3. Log in to IBM Cloud Object Storage by using the default username and password.

Note: The default username is admin; the default password is password. You are prompted as part of the initial setup to change the password for the admin user. Change the admin password.

The End-User License Agreement is shown next (see Figure 5-6).

Please read and accept the following IBM and non-IBM license agreements before proceeding. Language: English

LICENSE INFORMATION

The Programs listed below are licensed under the following License Information terms and conditions in addition to the Program license terms previously agreed to by Client and IBM. If Client does not have previously agreed to license terms in effect for the Program, the International Program License Agreement (Z125-3301-14) applies.

Program Name (Program Number):
IBM Cloud Object Storage System 3.14.3 (5725-281)
IBM Cloud Object Storage 1YR 3.14.3 (5641-C01)
IBM Cloud Object Storage 2YR 3.14.3 (5641-C02)
IBM Cloud Object Storage 3YR 3.14.3 (5641-C03)
IBM Cloud Object Storage 4YR 3.14.3 (5641-C04)
IBM Cloud Object Storage 5YR 3.14.3 (5641-C05)
IBM Cloud Object Storage System FIPS 3.14.3 (5725-281)
IBM Cloud Object Storage FIPS 1YR 3.14.3 (5641-C01)
IBM Cloud Object Storage FIPS 2YR 3.14.3 (5641-C02)
IBM Cloud Object Storage FIPS 3YR 3.14.3 (5641-C03)
IBM Cloud Object Storage FIPS 4YR 3.14.3 (5641-C04)
IBM Cloud Object Storage FIPS 5YR 3.14.3 (5641-C05)

The following standard terms apply to Licensee's use of the Program.

Limited use right

As described in the International Program License Agreement ("IPLA") and this License Information, IBM grants Licensee a limited right to use the Program. This right is limited to the level of Authorized Use, such as a Processor Value Unit ("PVU"), a Resource Value Unit ("RVU"), a Value Unit ("VU"), or other specified level of use, paid for by Licensee as evidenced in the Proof of Entitlement. Licensee's use may also be limited to a specified machine, or only as a Supporting Program, or subject to other restrictions. As Licensee has not paid for all of the economic value of the Program, no other use is permitted without the payment of additional fees. In addition, Licensee is not authorized to use the Program to provide commercial IT services to any third party, to provide commercial hosting or timesharing, or to sublicense, rent, or lease the Program unless expressly provided for in the applicable agreements under which Licensee obtains authorizations to use the Program. Additional rights may be available to Licensee subject to the payment of additional fees or under different or supplementary terms. IBM reserves the right to determine whether to make such additional rights available to Licensee.

Specifications

IBM License Agreement
 Non-IBM License Agreement

I have read and agreed to the terms provided in the IBM and non-IBM license agreements (required for acceptance).

Print Name (License Acceptor):

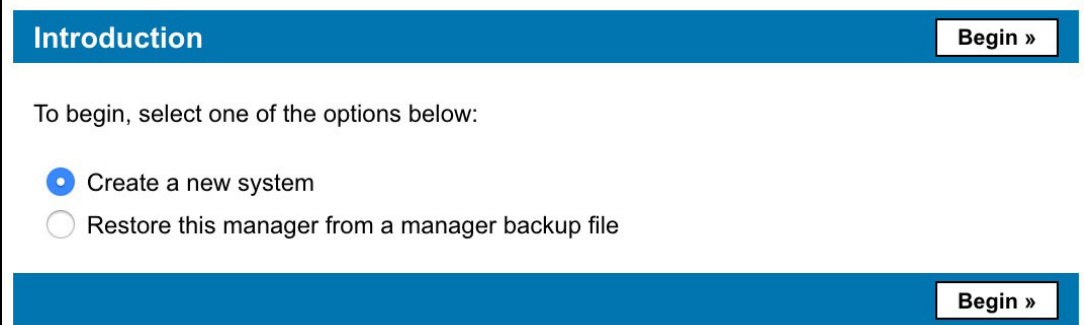
Figure 5-6 Accept License Agreement

4. Select **I have read and agreed to the terms provided in the IBM and non-IBM license agreements (required for acceptance)**.
5. Enter your initials to the right of **Print Name (License Acceptor)** field.
6. Click **Accept IBM & non-IBM Licenses**.

5.6.2 Creating a system

Complete the following steps to create a system:

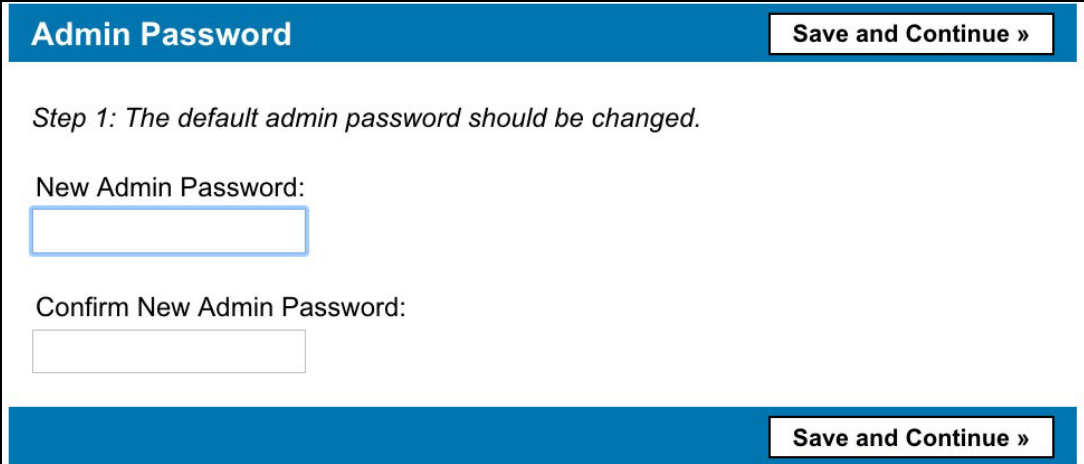
1. Select the **Create a new system** option and select **Begin** (see Figure 5-7).



The screenshot shows a web interface titled "Introduction". At the top right is a "Begin »" button. Below the title, the text reads "To begin, select one of the options below:". There are two radio button options: "Create a new system" (which is selected) and "Restore this manager from a manager backup file". At the bottom right, there is another "Begin »" button.

Figure 5-7 Creating a system

Next, you are prompted to change the password for the admin user (see Figure 5-8).



The screenshot shows a web interface titled "Admin Password". At the top right is a "Save and Continue »" button. Below the title, the text reads "Step 1: The default admin password should be changed.". There are two input fields: "New Admin Password:" and "Confirm New Admin Password:". At the bottom right, there is another "Save and Continue »" button.

Figure 5-8 Change administrator password

2. Choose a new password, confirm it, and select **Save and Continue**.

5.6.3 Creating a site

Next, you are prompted to create a site in which you can place appliances. A site in IBM Cloud Object Storage is a logical representation of a physical site. You can create more sites later and move appliances between sites (see Figure 5-9).

Create Sites Finish

Step 2: Configure a Site. At a minimum, the site name should be set appropriately (for example, "Chicago"). (Additional sites may be created for a multi-site network.) For more information, please refer to the [help](#).

* indicates required field

+ Add Additional Site

Site 1*:

Abbreviation:

Description:

Company:

Address:

Phone:

Latitude:

Longitude:

Finish

Figure 5-9 Creating a site

Enter the name of the site that you want to create and select **Finish**.

5.6.4 Accept pending devices

The next step is to accept all the appliances that you have configured into the system. Complete the following steps:

1. You see the prompt that is shown in Figure 5-10. Click **see below**.

There are 11 devices pending approval ([see below](#)).

Figure 5-10 Devices pending approval notification

A list of all devices pending is shown (see Figure 5-11 on page 111).

Devices Pending Approval: 11
Bulk Approve / Deny

Select devices to approve or deny from the system.

<input type="checkbox"/>	Hostname	IP Address	Device Type	Registered
<input type="checkbox"/>	<u>icospdcacc01</u>	10.69.70.110	accesser	2019-04-25 18:46:47 GMT
<input type="checkbox"/>	<u>icospdcacc02</u>	10.69.70.111	accesser	2019-04-25 18:52:20 GMT
<input type="checkbox"/>	<u>icospdcslc01</u>	10.69.70.130	slicestor	2019-04-25 19:09:57 GMT
<input type="checkbox"/>	<u>icospdcslc02</u>	10.69.70.131	slicestor	2019-04-25 19:41:26 GMT
<input type="checkbox"/>	<u>icospdcslc03</u>	10.69.70.132	slicestor	2019-04-25 19:41:58 GMT
<input type="checkbox"/>	<u>icospdcslc04</u>	10.69.70.133	slicestor	2019-04-25 19:20:46 GMT
<input type="checkbox"/>	<u>icospdcslc05</u>	10.69.70.134	slicestor	2019-04-25 19:24:27 GMT
<input type="checkbox"/>	<u>icospdcslc06</u>	10.69.70.135	slicestor	2019-04-25 19:27:55 GMT
<input type="checkbox"/>	<u>icospdcslc07</u>	10.69.70.136	slicestor	2019-04-25 19:33:24 GMT
<input type="checkbox"/>	<u>icospdcslc08</u>	10.69.70.137	slicestor	2019-04-25 19:37:39 GMT
<input type="checkbox"/>	<u>icospdcslc09</u>	10.69.70.138	slicestor	2019-04-25 19:43:17 GMT

Figure 5-11 Devices pending approval

2. Select all the devices, and then click **Bulk Approve / Deny**.

Note: The number of pending appliances varies and is based on your specific architecture. Ensure that the number of appliances that you installed and configured are included in this list.

3. Approve all the devices (see Figure 5-12).

Bulk Device Registration
Approve Deny Cancel

Verify whether the devices listed below should be approved for inclusion in the system.

Hostname	IP Address	Device Type	Key Fingerprint	Registered
icospdcacc01	10.69.70.110	accesser	30:05:14:8c:0b:f4:f1:80:93:0b:21:0d:6b:6f:41:84:90:b9:27:4f	2019-04-25 18:46:47 GMT
icospdcacc02	10.69.70.111	accesser	bb:2e:24:95:5c:1f:3a:a3:99:8e:07:c0:85:0c:23:75:8c:c7:6f:0d	2019-04-25 18:52:20 GMT
icospdcslc01	10.69.70.130	slicestor	39:48:79:a6:47:37:48:91:fa:3b:0a:c2:41:1b:60:6d:c1:4c:e2:b5	2019-04-25 19:09:57 GMT
icospdcslc02	10.69.70.131	slicestor	46:2b:c4:fe:59:ba:c4:58:90:84:5f:02:dd:6b:a8:e6:44:5d:f1:1e	2019-04-25 19:41:26 GMT
icospdcslc03	10.69.70.132	slicestor	bf:71:2d:a5:c7:11:f3:b5:4d:8b:3c:12:41:cd:f2:58:ca:e8:da:92	2019-04-25 19:41:58 GMT
icospdcslc04	10.69.70.133	slicestor	e5:cc:8e:de:52:70:6c:39:f9:3b:a1:52:a4:86:3a:7d:f3:cb:f7:7f	2019-04-25 19:20:46 GMT
icospdcslc05	10.69.70.134	slicestor	60:25:ea:47:fa:59:e5:4d:63:99:85:a1:2d:e8:70:8d:ce:a5:1c:2f	2019-04-25 19:24:27 GMT
icospdcslc06	10.69.70.135	slicestor	49:c5:53:d9:c8:f7:7b:60:4c:83:94:8c:60:fa:89:b0:6f:12:ad:40	2019-04-25 19:27:55 GMT
icospdcslc07	10.69.70.136	slicestor	9d:03:34:bb:c8:2d:6f:3f:55:d7:63:e2:5b:92:22:00:3c:92:61:b2	2019-04-25 19:33:24 GMT
icospdcslc08	10.69.70.137	slicestor	fa:ea:14:2c:7e:fd:7f:13:e3:25:af:6c:99:8b:53:a0:4e:54:86:15	2019-04-25 19:37:39 GMT
icospdcslc09	10.69.70.138	slicestor	fc:86:18:ee:73:cb:6b:e9:3a:61:17:37:5f:54:96:48:7a:cc:a5:91	2019-04-25 19:43:17 GMT

Approve Deny Cancel

Figure 5-12 Bulk device registration

4. Assign the devices to a site. Select all the appliances and assign them to the site you created (see Figure 5-13 on page 113).

Bulk Edit Device SiteSave

Select devices to assign them to a site:

<input type="checkbox"/>	Hostname
<input type="checkbox"/>	icospdcacc01 (10.69.70.110)
<input type="checkbox"/>	icospdcacc02 (10.69.70.111)
<input type="checkbox"/>	icospdcslc01 (10.69.70.130)
<input type="checkbox"/>	icospdcslc02 (10.69.70.131)
<input type="checkbox"/>	icospdcslc03 (10.69.70.132)
<input type="checkbox"/>	icospdcslc04 (10.69.70.133)
<input type="checkbox"/>	icospdcslc05 (10.69.70.134)
<input type="checkbox"/>	icospdcslc06 (10.69.70.135)
<input type="checkbox"/>	icospdcslc07 (10.69.70.136)
<input type="checkbox"/>	icospdcslc08 (10.69.70.137)
<input type="checkbox"/>	icospdcslc09 (10.69.70.138)

Assign the selected devices to site:

PDC

Or, create a new site for the selected devices

New Site Name:

Save

Figure 5-13 Assigning devices to a site

5. (Optional) Create an alias for each appliance (see Figure 5-14).

Bulk Edit Device Alias Save

Enter device aliases below and click the Save button when complete. (Alias field is optional)

icospdcmn01

icospdcacc01

icospdcacc02

icospdcslc01

icospdcslc02

icospdcslc03

icospdcslc04

icospdcslc05

icospdcslc06

icospdcslc07

icospdcslc08

icospdcslc09

Save

Figure 5-14 Bulk device alias

6. Click **Save**.

5.6.5 Creating a storage pool

Complete the following steps to create a storage pool:

1. From the main menu, click **Create Storage Pool**, as shown in Figure 5-15.

IBM Cloud Object Storage | Manager

[Monitor](#) [Configure](#) [Security](#) [Settings](#)

System	Summary
Vaults	22 Devices 1 Manager Device 6 Accesser® Devices 15 Slicestor® Devices
Storage Pools	2 Vaults
Access Pools	0 Mirror
Sites	2 Storage Pools
Devices	2 Access Pools 3 Sites 0 Cabinet

[Create Vault](#)
[Create Mirror](#)
[Create Storage Pool](#)
[Create Access Pool](#)
[Create Site](#)
[Create Cabinet](#)

Figure 5-15 Creating a storage pool

2. Complete the following steps, as shown in Figure 5-16:
 - a. Enter the name of the storage pool in the Name field.
 - b. Select the width of the storage pool based on the Information Dispersal Algorithm (IDA) of the solution.
 - c. Select **Packed Storage**.
 - d. Select all the devices that should be part of the storage pool.
 - e. Click **Save**.

Create New Storage Pool
Cancel Save

General

Name:

Width:

[Need to create a storage pool with some of the new devices temporarily missing?](#)

Storage Engine (cannot be changed later):

Packed Storage
Recommended for all S3 or OpenStack use cases as well as any use case involving mirrors or small (<32 kB) objects.

File Storage
Legacy engine for Simple Object use cases, particularly ones which include heavy delete activity.

Enable the embedded Accesser service on all Slicestor devices belonging to this storage pool

Suggest Devices Suggest Devices

Devices

[Select All](#) [Unselect All](#) Selected item count: 9

	Name	Appliance Name	Drive Count	Site	Total Size	Version
<input checked="" type="checkbox"/>	s icospdclsc01	Virtual Appliance - KVM	48	PDC	12.55 GB	3.14.3.65
<input checked="" type="checkbox"/>	s icospdclsc02	Virtual Appliance - KVM	48	PDC	12.55 GB	3.14.3.65
<input checked="" type="checkbox"/>	s icospdclsc03	Virtual Appliance - KVM	48	PDC	12.55 GB	3.14.3.65
<input checked="" type="checkbox"/>	s icospdclsc04	Virtual Appliance - KVM	48	PDC	12.55 GB	3.14.3.65
<input checked="" type="checkbox"/>	s icospdclsc05	Virtual Appliance - KVM	48	PDC	12.55 GB	3.14.3.65
<input checked="" type="checkbox"/>	s icospdclsc06	Virtual Appliance - KVM	48	PDC	12.55 GB	3.14.3.65
<input checked="" type="checkbox"/>	s icospdclsc07	Virtual Appliance - KVM	48	PDC	12.55 GB	3.14.3.65
<input checked="" type="checkbox"/>	s icospdclsc08	Virtual Appliance - KVM	48	PDC	12.55 GB	3.14.3.65
<input checked="" type="checkbox"/>	s icospdclsc09	Virtual Appliance - KVM	48	PDC	12.55 GB	3.14.3.65

Figure 5-16 Naming a storage pool

5.6.6 Creating a vault

Complete the following steps to create a vault:

1. From the main window, click **Create Vault**.

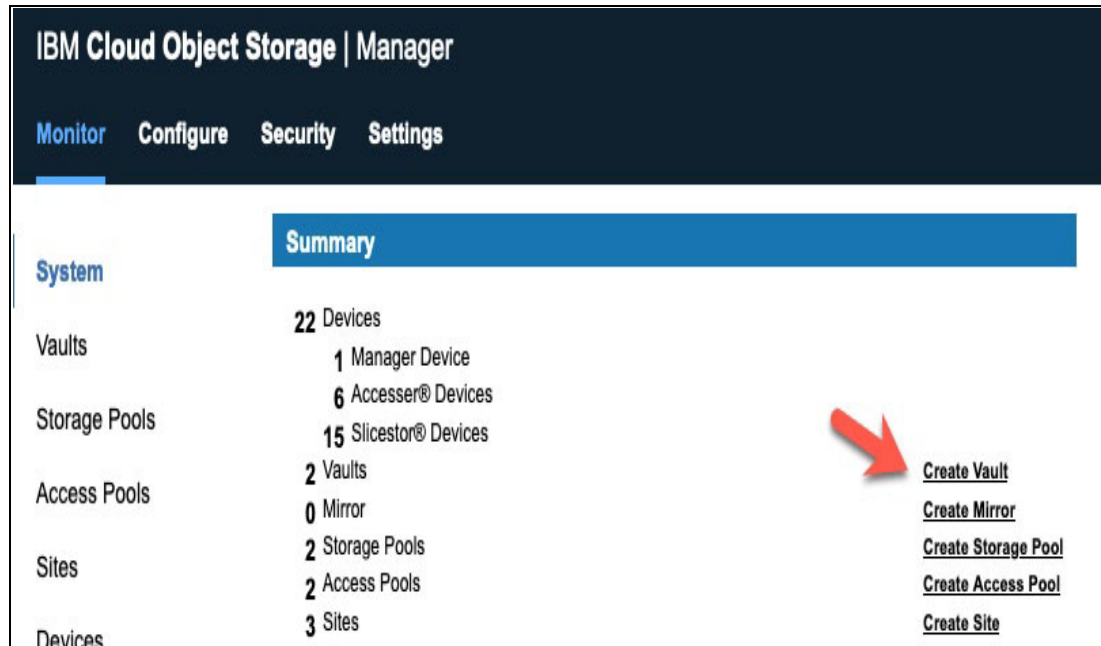


Figure 5-17 Creating a vault

2. Enter your vault name in the field. Then, set the Width, Threshold (read threshold (RT)), and Write Threshold (WT) based on the IDA for which you designed the system.
3. For Options, leave the default settings.
4. For Quotas, leave the default, settings.
5. For Advanced Index Settings, leave the default settings.
6. Click **Save** (see Figure 5-18 on page 117).

For more information about the IDA, see [IBM Documentation](#).

Attention: When running in Concentrated Dispersal (CD) mode, instead of setting a user-defined IDA, you are presented with two or more options. These options enable you to choose between a system with better performance or a system with more available storage.

Create New Standard Vault Cancel Save

General

Name: * pdc_vault_01

Description: (optional)

Tags: Select one or more tags...

Configuration

Width: * 9 Write Threshold: * 6

Threshold: * 4 Alert Level (optional): 8

Figure 5-18 Create New Standard Vault

Tip: When creating a vault, you must decide which vault index version is suitable for your environment (Name Index Format Version 4 (default) or Name Index Format Version 2). For more information about choosing between the two versions, see 5.6.16, “Vault Index Version” on page 132.

5.6.7 Creating an access pool

Complete the following steps to create an access pool:

1. From the window, click **Create Access Pool**, as shown in Figure 5-19.

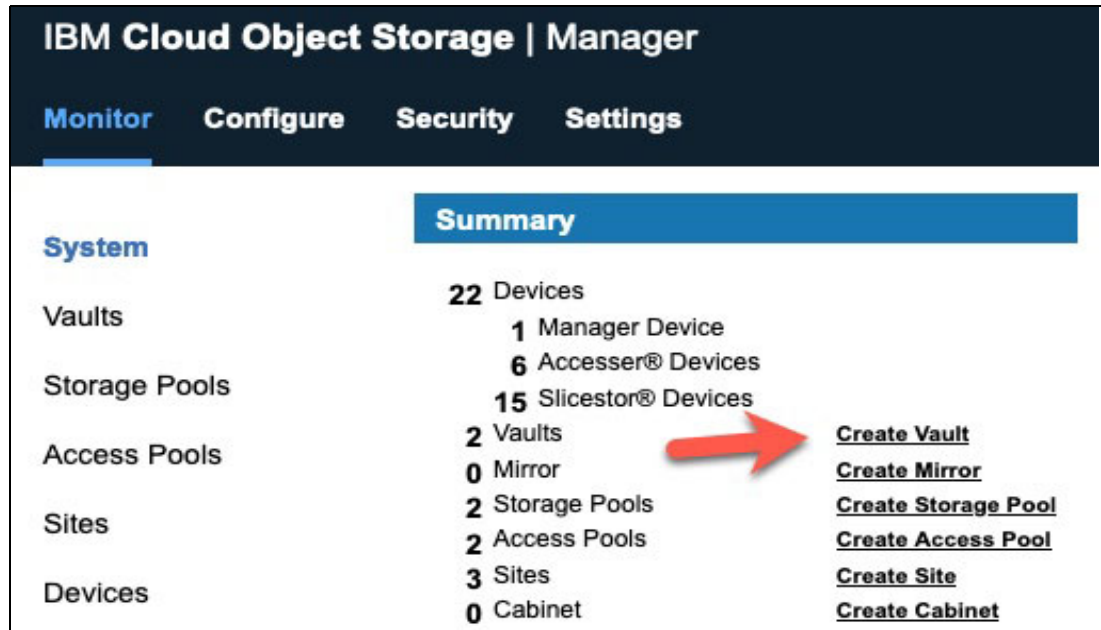


Figure 5-19 Create Access Pool link

2. Enter the name of the vault in the Name field. Leave the API type as Cloud Object Storage, which is the default API and is compatible with the S3 API (see Figure 5-20).

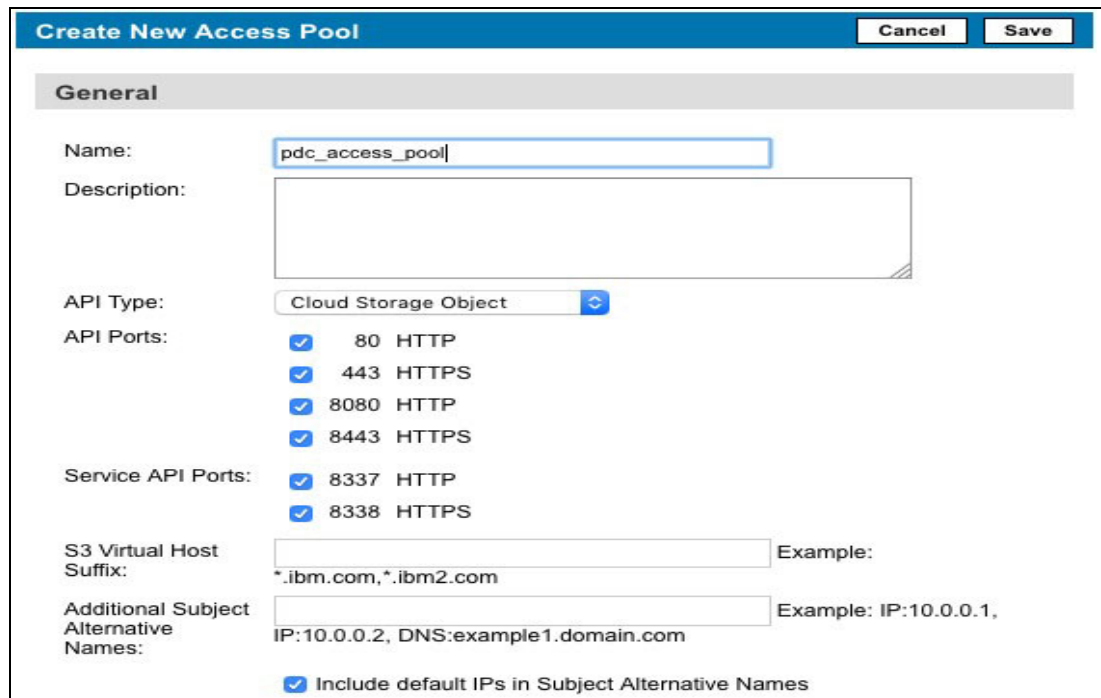


Figure 5-20 Create New Access Pool window

3. Select all the Accessers in your environment. This selection creates an Access pool for all Accessers (see Figure 5-21).

The access pool that was created represents a many-to-many relationship between Accesser appliances and buckets. For more information about the access pool, see *IBM Cloud Object Storage Concepts and Architecture*, REDP-5537.


Access Devices			
	Name (Device IP)	Site	Version
<input checked="" type="checkbox"/>	A  .icospdcacc01 (10.69.70.110)	(PDC)	3.14.3.65
<input checked="" type="checkbox"/>	A  .icospdcacc02 (10.69.70.111)	(PDC)	3.14.3.65

Figure 5-21 Selecting Accessers for an access pool

4. In the deployment section, select your storage pool, **Standard Type** for the Item Type, and select the vault that you created earlier, as shown in Figure 5-22.

Deployment

Storage Pool:

Item Type:

Text:

To select multiple items at once click on the desired check boxes while holding shift key.

[Select All](#) [Unselect All](#) Visible items: 1 - Filtered from: 2
Selected item count: 1

pdc_vault_01

Figure 5-22 Selecting deployment options for an access pool

5. Click **Save**.

5.6.8 Configuring HTTPS certificates for access pools

Certificates can be used for HTTPS access that is trusted inside your organization instead of the default Manager signed certificates. You may obtain a certificate from a trusted certificate provider, such as Comodo SSL, DigiCert, GeoTrust, GlobalSign, Letsencrypt, or any other preferred provider in your organization. Follow the instructions from the certificate provider to generate a certificate request for your access pool, sign it with certificate authority and generate certificate in PEM format. Complete the following steps to upload the generated certificate to IBM Cloud Object Storage:

1. Click the **Configure** tab.
2. Click **Access Pools** in the navigation panel.
3. Click the link of an access pool to display the Access Pool: <access-pool-name> page.
4. In the **Access Pool HTTPS Certificate** section, click **Configure** to display the Editing Access Pool HTTPS Certificate page.

5. Paste PEM-formatted private key and certificate text into the corresponding **Private Key PEM** and **Certificate PEM** fields:
 - a. To add more certificates, paste one after another.
 - b. To remove a single certificate, delete text for that certificate.
 - c. To remove all certificates, delete all contents.
6. Click **Update** to update the certificates for the access pool.

5.6.9 Enabling Access Key Authentication

Enabling Access Key Authentication enables a user to use an Access Key and Secret Key to authenticate for vault access in IBM Cloud Object Storage. Complete the following steps:

1. From the main window, select the **Security** tab. Then, select **Configure** in the Vault Access Authentication bar (see Figure 5-23).

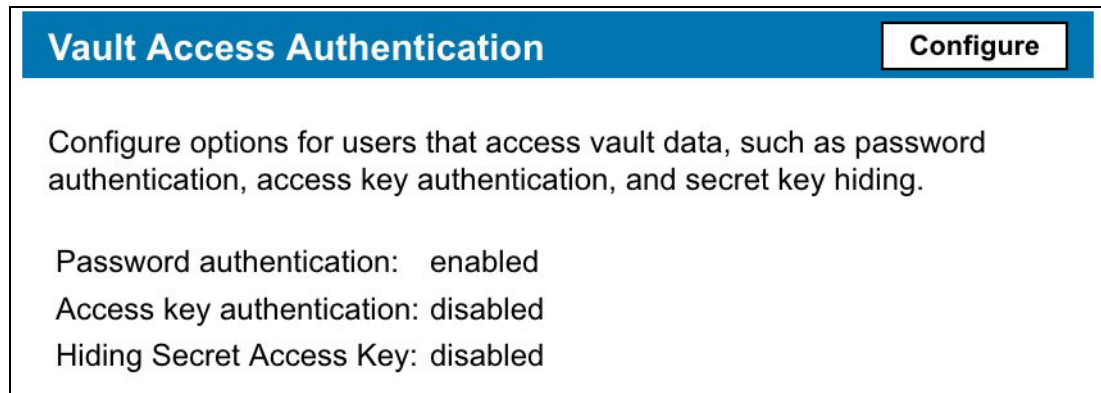


Figure 5-23 Configuring Vault Access Authentication

2. Select **Enable access key authentication**, and then select **Update** (see Figure 5-24).

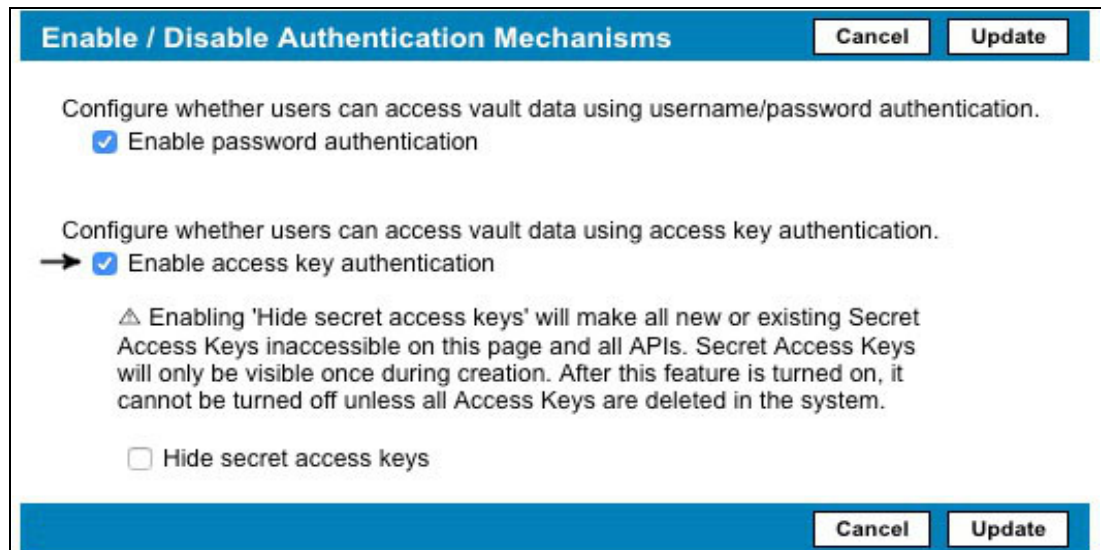


Figure 5-24 Enabling Authentication Mechanisms

5.6.10 Creating a user

Complete the following steps to create a user that can be used to store data in IBM Cloud Object Storage:

1. From the main window, select the **Security** tab.
2. From the Security tab, select **Create Account** (see Figure 5-25).

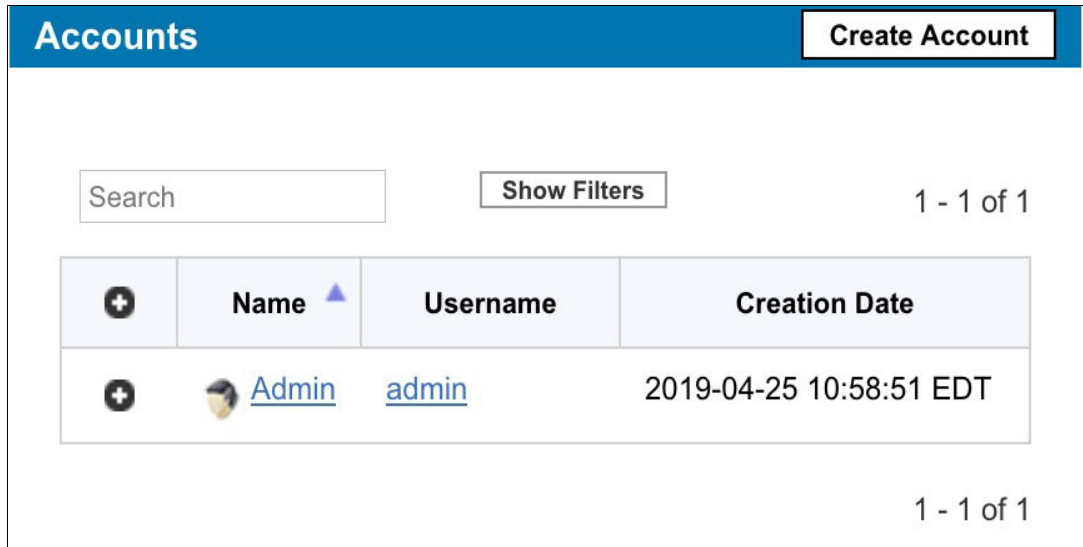


Figure 5-25 Creating an account

3. In the General section, enter the name of the user. This name is the username that is displayed in IBM Cloud Object Storage (see Figure 5-26).

General

Name:

Email:
 (optional)

Figure 5-26 Name of user

- In the Authentication section, clear the option that enables the user to authenticate with a username and password (see Figure 5-27).

Figure 5-27 Disabling username and password authentication

- In the Roles section, make sure that none of the options are selected. If you select any option, you give the corresponding IBM Cloud Object Storage roles to the user. For a user who needs only the ability to access the system for storage, none of these roles are required (see Figure 5-28).

Assign Role	Read Only	Role	Description
<input type="checkbox"/>		Super User	Perform any action within the Cloud Object Storage Manager except vault read/write.
<input type="checkbox"/>	<input type="checkbox"/>	System Administrator	Perform any action within the Cloud Object Storage Manager except security, account management and vault read/write.
<input type="checkbox"/>	<input type="checkbox"/>	Security Officer	Perform security and account management actions within the Cloud Object Storage Manager.
<input type="checkbox"/>		Operator	Perform monitoring actions within the Cloud Object Storage Manager.

Figure 5-28 Account roles

- In the Vault Access section, you should be in the No Access tab. Select the vault that you created earlier and then, select **Move to Read/Write**. The user now has read and write access to the vault (see Figure 5-29).

Figure 5-29 User Vault Access

- In the Device Access tab, ensure that **No Access** is selected.

8. In the Site Access tab, leave the default settings.
9. Select **Save**.

5.6.11 Generating Access Key ID

The next step is to generate an Access Key ID for the user to manage objects within IBM Cloud Object Storage.

Note: The default authentication method for S3 is Access Key ID and Secret Key authentication.

Complete the following steps:

1. From the main window, select the **Security** tab.
2. From the Security tab, select the username of the account that you created.
3. Select **Change Keys**, as shown in Figure 5-30.




Figure 5-30 Access Key Authentication

4. Click **Generate New Access Key** (see Figure 5-31).




Figure 5-31 Generate Access Keys

5. A new Access Key ID and Secret Access Key are generated. Record this information for use later on.
6. Exit this window by selecting **IBM Cloud Object Storage** in the upper left of your web browser.

5.6.12 Granting CLI Access

For IBM Cloud Object Storage Administrators, you might want to give them access to the CLI. Complete the following steps to grant a user CLI access to the individual IBM Cloud Object Storage appliances:

1. From the main window, select the **Security** tab.
2. From the Security tab, select the username of the account to which you want to grant CLI Access.

3. In the upper right of the window, select **Change**, as shown in Figure 5-32.

Figure 5-32 Changing access for a user

4. Scroll down to the bottom of the page to the Device Access section, as shown in Figure 5-33.

Devices	Root	Read/Write	Read Only	No Access
Manager Device	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Other Devices	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>

Figure 5-33 Device Access

5. Select an option to grant access to the Manager.

Note: The default device access is **No Access** for all users.

6. Select an option to grant access to all other appliances, except for the Manager.

7. For all non-manager devices, you can select more permissions based on the site of the appliances.

8. For each site, you can select the site and click **Root Access**, **Read Write Access**, or **Read Only Access**, as shown in Figure 5-34.

Figure 5-34 Site Level Access

Note: Site Level Access applies only to non-Manager appliances.

9. To set the configuration, select **Update**.

After waiting a few minutes, you should be able to SSH into the appliance by using the correct username and password.

Access to each device depends on the user's default Manager device permissions, default non-Manager device permissions, and site level permissions. For Manager devices, the most permissive default permission determines device access; *site level permissions do not apply*. For non-Manager devices, the most permissive of all default or site level permission determines device access.

Note: For a more information about a Device Role-Based Access configuration and its use cases and deployment options, see [IBM Documentation](#).

5.6.13 Manager configuration backup

Regularly back up the Manager configuration to have a recent configuration ready if the Manager becomes unusable for any reason. With a backup, you quickly can restore the manager functions by using the integrated recovery procedure when starting the Manager GUI for the first time.

You can configure IBM Cloud Object Storage to back up its configuration according to a daily or weekly schedule. Destinations for the backup file can be a local directory on the Manager appliance, a remote FTP or HTTP Server, or an S3 endpoint.

To activate a regular backup of the Manager configuration, complete the following steps:

1. From the main window, select the **Settings** tab.
2. Select **Operations** and then **Backup Configuration**, as shown in Figure 5-35.

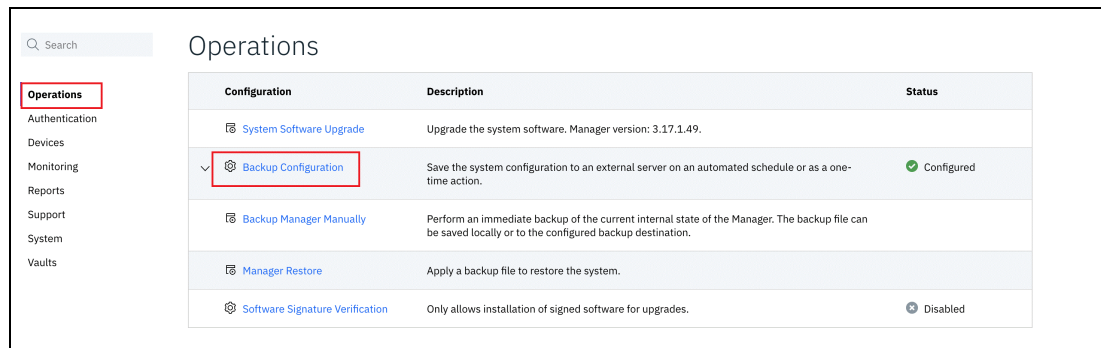
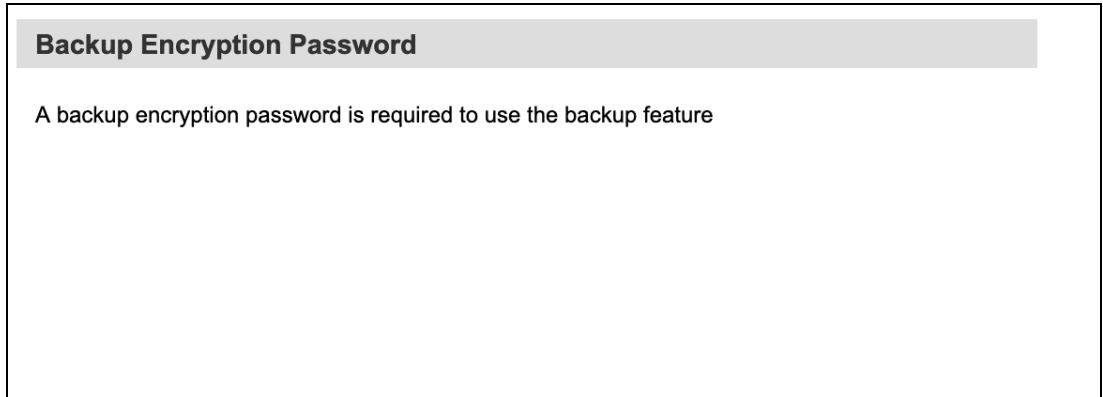


Figure 5-35 Backup configuration

3. Set a backup encryption password, as shown in Figure 5-36.



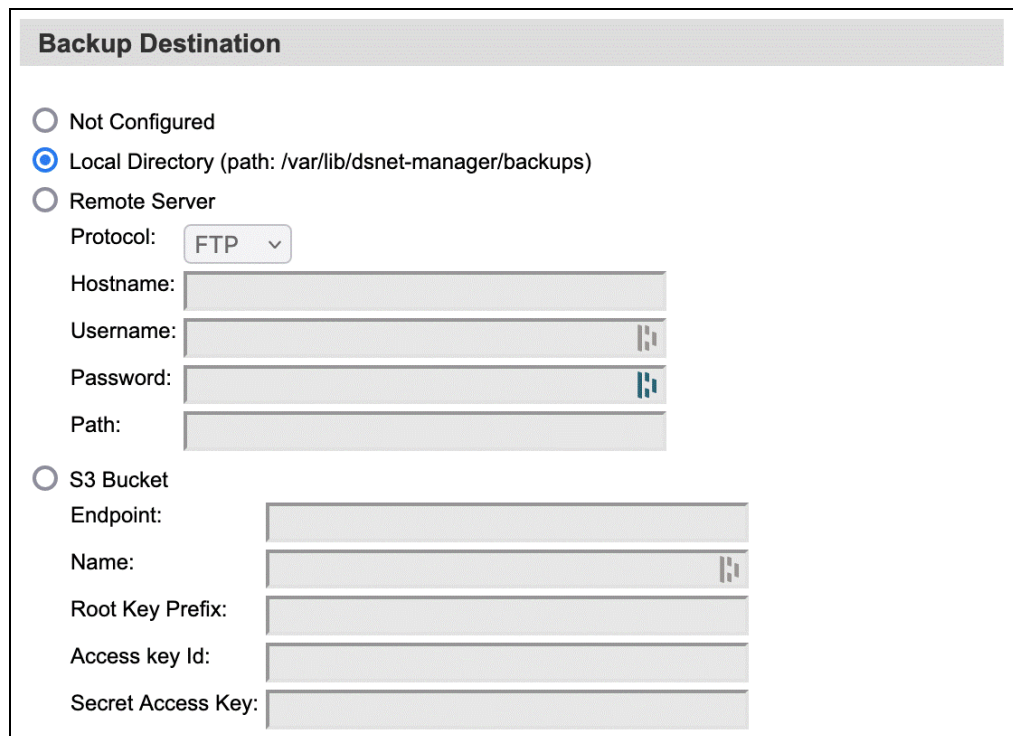
Backup Encryption Password

A backup encryption password is required to use the backup feature

Figure 5-36 Protecting the backup file with a password

4. Select backup targets:

- **Local Directory:** Save a backup file on the manager appliance.
- **Remote Server:** You can use an FTP or SFTP server. Enter the credentials and path.
- **S3 Bucket:** Uses an S3 endpoint with S3 credentials. You can set an object prefix.



Backup Destination

Not Configured

Local Directory (path: /var/lib/dsnet-manager/backups)

Remote Server

Protocol:

Hostname:

Username:

Password:

Path:

S3 Bucket

Endpoint:

Name:

Root Key Prefix:

Access key Id:

Secret Access Key:

Figure 5-37 Backup Destination

5. Set a schedule for automatic backups. You can set it to daily or weekly backups, as shown in Figure 5-38 on page 127.

Automatic Backup Schedule

The Backup Destination Information section is required to use the Automated Backup Feature.

Disable automatic backup

Enable automatic backup - times are in GMT

Daily at

Weekly on at

Figure 5-38 Backup scheduling

5.6.14 Organizations

Organizations is a feature that allows an IBM Cloud Object Storage administrator (with Super User or Security Officer role) to create entities called Organizations. They also can define the maximum number of vaults and the total storage capacity that can be allocated to each Organization.

IBM Cloud Object Storage administrator can create user accounts and vaults and map them to Organizations. However, the user also can provision new vaults that belong to an Organization by way of IBM Cloud Object Storage API if the user has the Vault Provisioner role.

Organizations feature is supported by IBM Cloud Object Storage System running in vault mode.

Note: This feature is not supported in Container Mode, which means standard vaults that are converted to container vaults lose the configurable grouping of vaults by Organization. The concept of multi-tenancy is implemented with by using *Storage Accounts* in Container Mode. For more information, see [IBM Container Mode Feature Description Document](#).

To create an Organization, complete the following steps:

1. From the main window, select the **Security** tab.
2. From the Security tab, select **Create Organization**, as shown in Figure 5-39.

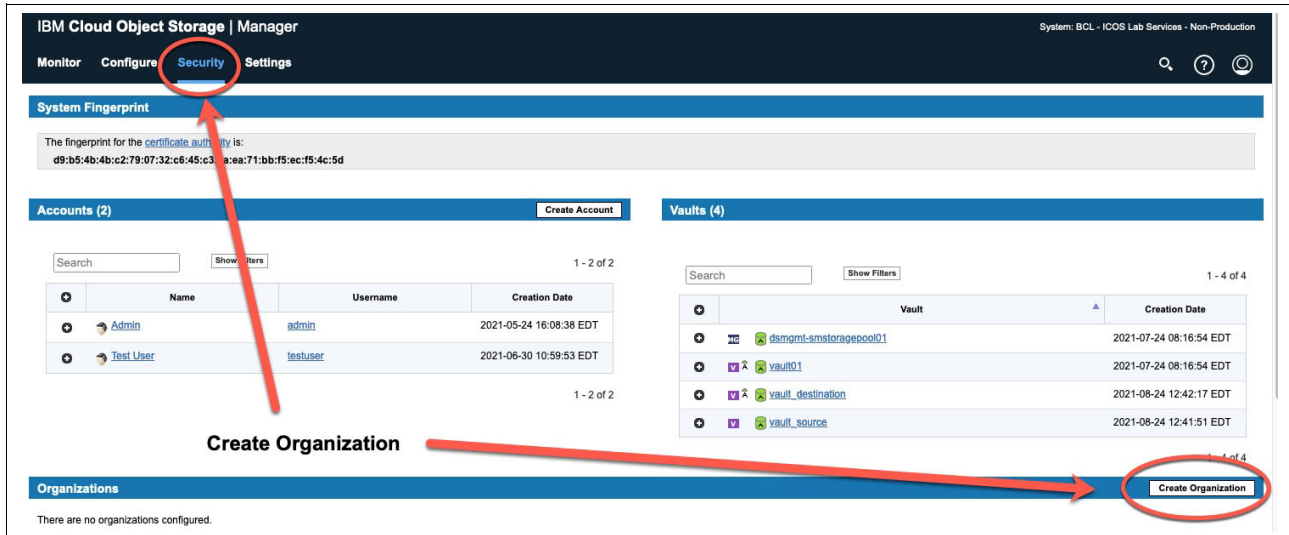


Figure 5-39 Selecting Create Organization option

3. Enter Name and Description for the new Organization.
4. Choose Unlimited access or maximum usable capacity and maximum number of vaults for the Organization.
5. Select **Save**, as shown in Figure 5-40.

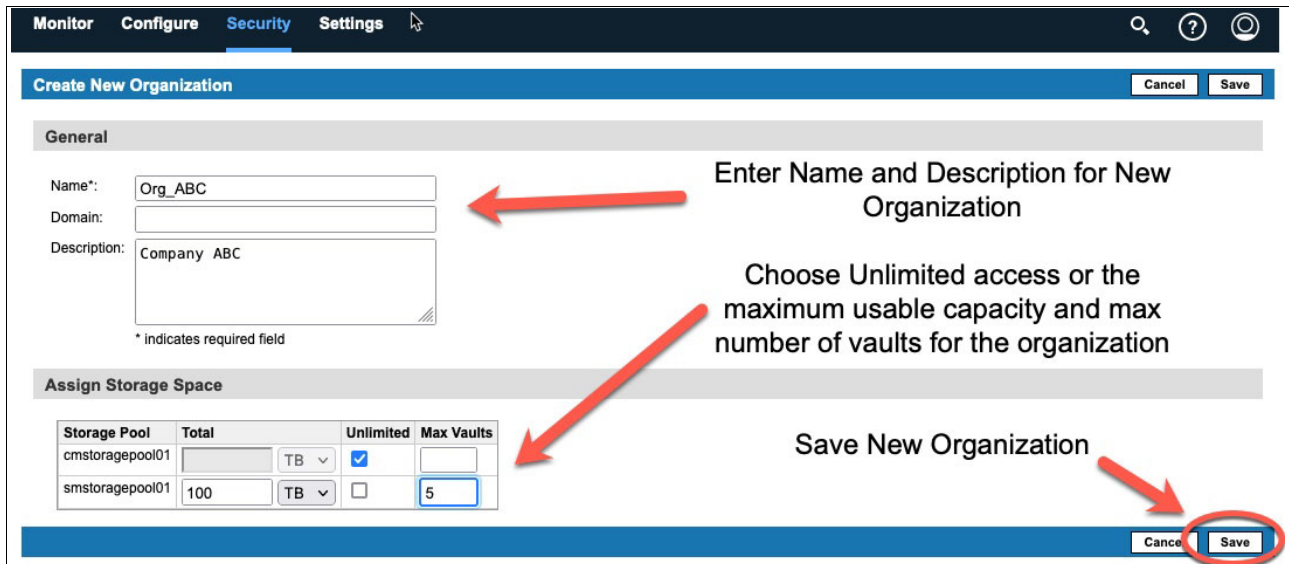


Figure 5-40 Selecting Save

6. Newly created Organizations are now shown in the Organizations section (see Figure 5-41 on page 129).

Name	Domain	Storage Pool	Used	Total	Vaults	Max Vaults	Action
BCL - IBM COS Lab Services - Non-Production - Stygian Research (System Owner)		smstoragepool01	479.48 MB	Unlimited	4		
Org_ABC		smstoragepool01	0 bytes	100 TB	0	5	Change

List of Organizations including default System Owner Organization

Capacity usage and Quotas per Organization

To change Existing Organization

Figure 5-41 New Organizations listed

- To create a user and associate that user with an Organization, complete steps 1 and 2 as described in 5.6.10, “Creating a user” on page 121.
- In the General section, the new user can now be associated with an Organization (see Figure 5-42).

Create New Account [Cancel] [Save]

General

Name:

Email: (optional)

Organization:

- Select Organization
- BCL - IBM COS Lab Services - Non-Production - Stygian Research (System Owner)
- Org_ABC

Allow authentication with a username and password maintained within the Cloud Object Storage Manager

Create new Account and associate is with an Organization

Figure 5-42 Create organization - 4/6

- A Vault Provisioner role should be assigned to the user if the user needs to create vaults by using API (see Figure 5-43).

Assign Role	Read Only	Role	Description
<input type="checkbox"/>		Super User	Perform any action within the Cloud Object Storage Manager except vault read/write.
<input type="checkbox"/>	<input type="checkbox"/>	System Administrator	Perform any action within the Cloud Object Storage Manager except security, account management and vault read/write.
<input type="checkbox"/>	<input type="checkbox"/>	Security Officer	Perform security and account management actions within the Cloud Object Storage Manager.
<input type="checkbox"/>		Operator	Perform monitoring actions within the Cloud Object Storage Manager.
<input checked="" type="checkbox"/>		Vault Provisioner	Create / delete vaults using the Provisioning API. This role alone does not grant access to the Cloud Object Storage Manager interface.

Figure 5-43 Create Organization - 5/6

10. A vault can also be created and associated to a specific Organization by the IBM Cloud Object Storage administrator (see Figure 5-44).

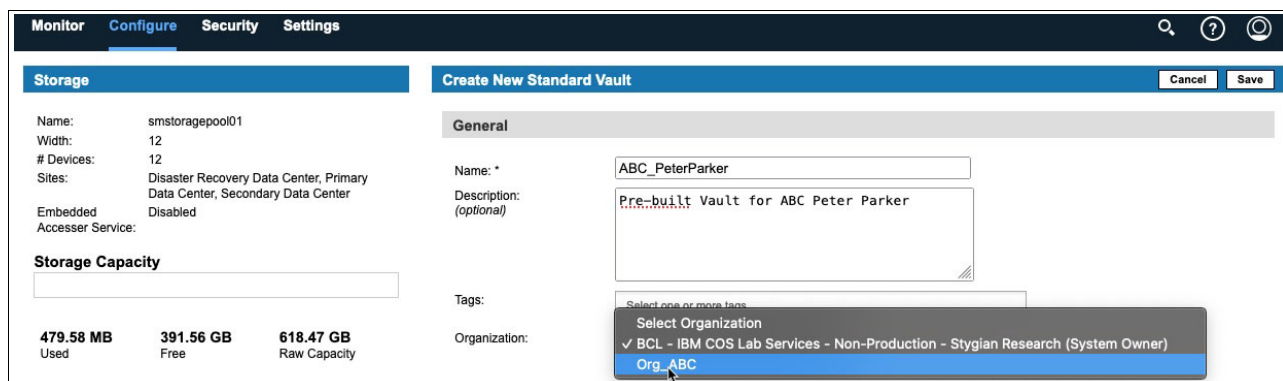


Figure 5-44 Create Organization - 6/6

11. If vault is created by IBM Cloud Object Storage administrator, follow step 6 in 5.6.10, “Creating a user” on page 121 to give access rights to the vault for the user.

12. Follow the steps as described 5.6.11, “Generating Access Key ID” on page 123 to generate access keys for the user.

5.6.15 Notification Service

The Notification Service provides the user flexibility to implement Apache Kafka notifications as needed, including notifications for all object writes and deletes, and can be used with IBM Spectrum Discover.

This section describes the IBM Cloud Object Storage notification service and how to enable it for vaults.

Notification concepts

Notifications enable IBM Cloud Object Storage to take part in a business process in a way that it can notify other applications of changes in a specific vault. When an object is written, overwritten, or deleted, IBM Cloud Object Storage sends out an Apache Kafka message into a Kafka topic. Other applications (for example, IBM Spectrum Discover for metadata management) can then trigger actions based on that new object.

Consider the following example:

- ▶ The user writes or deletes objects to an IBM Cloud Object Storage System.
 - The system allocates a notification intent that includes the notification data.
 - The system then writes the intent to Slicestor memory in parallel with object metadata and other intents.
 - The system replies to the user’s request.
 - If the user request succeeds, the system sends a notification to a Kafka topic. If the user request fails, the system does not send a notification, and the intent is deleted.
- ▶ The Kafka consumer on that topic can now act based on the JSON formatted information that is provided in the notification. The notification contains information about the object, vault, and the IBM Cloud Object Storage System.

Note: Vaults that are assigned to a notification configuration cannot be used in a mirror, for data migration, with a management vault, or with a vault proxy.

Notification Services are also available in Container Mode. In Container Mode, each Container can be given a separate notification setting.

Configure Notification Service

The notification service is enabled by completing the following steps:

1. In the Settings tab, click **Monitoring**, and then click **Notification Service**.
2. Click **Add** to create a Kafka connection.
3. Complete the required fields, as shown in Figure 5-45.

Add Notification Service [Cancel] [Save]

General

Name*: Discover_Connection

Topic: cos-le-connector-topic

Hosts*: 10.0.40.50:9092

Type*: General IBM Spectrum Discover

Authentication

Enable authentication

Username: cos

Password:

Encryption

Enable TLS for notification cluster network connections

Certificate PEM:

```
-----BEGIN CERTIFICATE-----
MIIEYzCCA7OgAwIBAgIJAPHO1GJ4/VaAMA0GCSqGSIb3DQEBCwUAMIGAMQswCQYD
VQQGEwJHQjeOMAwGA1UECAwFSEFOVFMxEDA0BgNVBACMB0h1cnNsZXkxDDAKBgNV
BAoMA0lCTTEaMBGGA1UECwwRc3B1Y3RydW0uZG1zY292ZXIxGjAYBgNVBAMMEXNw
ZWN0cnVtLmRpc2NvdmVyMSMwIQYJKoZIhvcNAQkBFhRtbGF3cmVudWY2VAdWsuaWJt
LmNvbTAeFw0yMDAyMTQxNjE2NDIaFw00MDAyMDkxNjE2NDIaMIGAMQswCQYDVQQG
EwJHQjeOMAwGA1UECAwFSEFOVFMxEDA0BgNVBACMB0h1cnNsZXkxDDAKBgNVBACM
A0lCTTEaMBGGA1UECwwRc3B1Y3RydW0uZG1zY292ZXIxGjAYBgNVBAMMEXNwZWN0
-----END CERTIFICATE-----
```

Figure 5-45 Create Notification Service

- Choose a name that reflects this connection in IBM Cloud Object Storage.
- Specify a Kafka topic to which to write. This specification also can be done later in the vault settings.
- Insert hosts where the Kafka cluster is running, including the port number.

- Specify whether notifications are sent to IBM Spectrum Discover or another Kafka environment.
- Enable authentication, if required by the Kafka cluster.
- Insert a PEM Certificate, if required by the Kafka cluster.

After setting up the Notification Service, complete the following steps to enable it on the vaults that use this service:

1. Go to the vault configuration.

A new section for the Notification Service is now available.
2. Choose the created service.

An individual Kafka topic for this vault can be specified here (see Figure 5-46).

The screenshot shows a configuration window titled "Notification Service". It contains a dropdown menu for "Notification Service" with "Discover_Connection" selected. Below it, there are two radio button options for "Topic": "Configured topic: cos-le-connector-topic" (which is selected) and "Custom:" followed by an empty text input field. At the bottom right of the window, there are two buttons: "Cancel" and "Update".

Figure 5-46 Enable Notification Service on vault

3. Click **Update** to enable the settings.

All writes, overwrites, and deletes are now pushed to the Kafka topic.

Note: In Container Mode, the notification setting on a bucket is done by specifying the topic at creation or metadata update, as shown in the following example:

```
PUT <accesser>:8338/container/{bucket.name}
{
  "storage_location":"us-south",
  "service_instance":"731fc6f265cd486d900f16e84c5cb594",
  "notifications":{
    "topic":"my-bucket_topic"
  }
}
```

5.6.16 Vault Index Version

When creating a vault, Name Indexing is enabled by default. When enabled, Name Index allows a user to list the contents of a vault in lexicographical order based on the object's name or key. The Name Index is updated whenever objects are added or removed from a vault. The Name Index must be enabled to provide prefix-based listing and sorted listing results for named object vaults. If you disable Name Index, it can be re-enabled only by contacting Customer Support.

Note: Name Index cannot be disabled for Protected Vaults.

Recovery Listing is another option that can be enabled at vault creation. This feature allows for limited listing capability, even when the contents of a vault are not indexed. When enabled, Recovery Listing lists the SourceNames of the metadata headers.

Recovery Listing is slower than the Name Index listing and the results are not sorted. It also can be used to list contents of a vault for which Name Index is corrupted or not enabled.

In IBM Cloud Object Storage, each standard vault includes an associated Vault Index. Vault Index is a distributed dispersed data structure that is in the vault and enables with S3 listing of objects for the vault. Before ClevOS 3.14.8, all configured standard vaults support Name Index Format Version 2 only, in which only the names of all S3 objects in the vault are maintained. All S3 listing requests for a standard vault are serviced by using the Vault Index Version 2 to get the object names, followed by the lookup of the object's metadata to retrieve information about the object's last modified time, e-tag, size, and owner information for inclusion in the listing response.

Starting with ClevOS 3.14.8, a feature function supports Name Index Format Version 4. Vaults that use the newer Vault Index Version 4 see significant latency improvements in servicing of S3 listing requests for the configured vault. However, depending upon the workflow for the standard vault, a performance degradation might occur with latencies while servicing S3 PUT requests for small objects (< 1 MB).

Note: The *Object Expiration* feature in this release, and beyond, can be enabled only on standard vaults that include Vault Index Version 4 configured. All planned data lifecycle management features will be supported on only standard vaults that have Vault Index Version 4 configured.

Benefits

A standard vault that uses the Name Index Format Version 4 features the following benefits over a standard vault that uses Name Index Format Version 2:

- ▶ Latencies improved while servicing all the S3 listing requests.
- ▶ Supports more S3 listing operations.
- ▶ Less CPU and disk usage on IBM Slicestor devices.

In addition, Vault Index Version 4 is a prerequisite for enabling data lifecycle management features, including the Object Expiration feature.

In most mixed work flows, the improvements in listing performance are more than enough to offset the increased cost of writing data by using Vault Index Version 4. However, you might want to conduct experiments before enabling Version 4 as the default on your systems. Considerations that can suggest such a course of action include the following examples:

- ▶ Vault Index is enabled, but the workflow uses minimal S3 Bucket Listing and includes a large percentage of small object writes.
- ▶ The client application is sensitive to changes in write latency.

Configuring Vault Index Version

Tip: If you are unsure or your application vendor did not specify the vault index requirement, use *Name Index Enabled* and *Vault Index Version 4* settings.

Complete the following steps to configure Vault Index Version:

1. At the global system level on the Configure page, set the **Vault Index Version** to be Version 2 or Version 4.

When any vaults are created with the Name Index Enabled feature, the configured global default for Index Version is used if the Index Version is not overridden by the setting at the storage pool or at the vault.

2. (Optional) Configure a different Vault Index Version at an individual storage pool. Also, override the global system setting for Default Vault Index Version the New Create Storage Pool page or the Edit: StoragePoolName page that is used.

The Name Index Format Default for the Storage Pool overrides the global default Name Index Format. All vaults that are created on this storage pool with Name Index enabled inherit this Name Index Format Setting from the storage pool if you do not set a specific Vault Index Version.

3. (Optional) Set the Vault Index Version at the time of standard vault creation. The specified Vault Index Version for the standard vault overrides the global default and storage pool settings for Vault Index Version. Specify the Vault Index Version that uses the Advanced Index Settings by using the Create New Standard **Vault** page:

- a. Go to **Create New Standard Vault Advanced Index Settings** and select **Name Index Enabled**.
- b. Set **Vault Index Version** to Version 2 or Version 4.

The Name Index Format for the vault overrides the Name Index Format configured for the storage pool and the global configuration.

Important: The Vault Index Version cannot be updated after standard vaults are created. This limitation applies to all preexisting vaults that are created on older releases and vaults that are created on the current release.

4. Repeat the previous step to configure the Vault Index Version for a mirror by using the **Create New Standard Mirror page**.

New mirrors that are configured by using a seed vault automatically set the Name Index Enabled and Name Index Format settings for the mirror based on the settings of the seed vault.

5.6.17 Vault Deletion Authorization

A vault often is deleted by one administrator. The *Vault Deletion Authorization* option can be enabled for 4-eyes authorization by two administrators. This option enhances security against misconfiguration and data loss.

To enable the Vault Deletion Authorization feature, go to the **Security** tab and select the section that is shown in Figure 5-47 on page 135. This system-wide setting is applied to all vaults.



Figure 5-47 Vault Deletion Authorization option

5. Select **Enable multi-user vault deletion** and click **Update**, as shown in Figure 5-48.

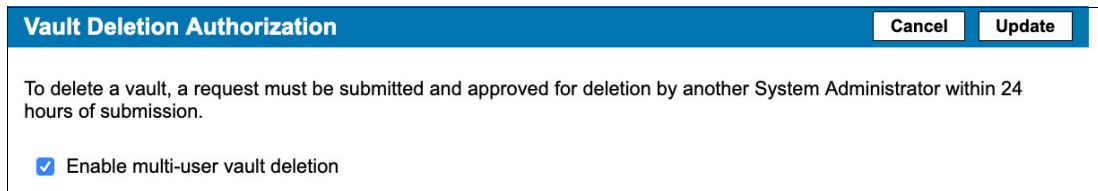


Figure 5-48 Enabling Vault Deletion Authorization

With that option enabled, a verification from a second user is needed to delete vaults. While deleting a vault, the message shown in Figure 5-49 is displayed.

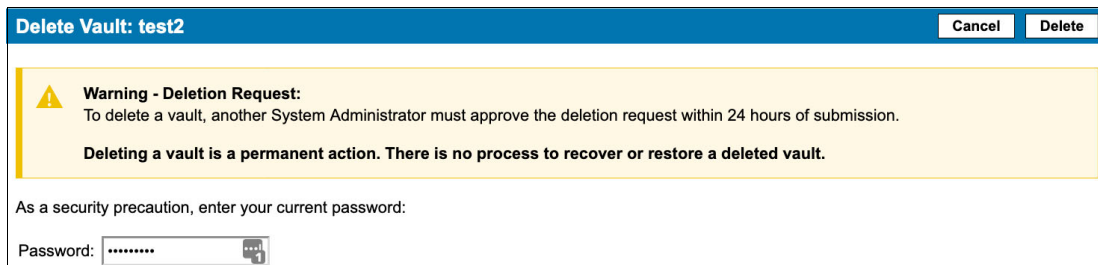


Figure 5-49 Deleting a vault shows this message

6. The second administrator sees the message that is shown in Figure 5-50 in the Configuration overview of the vault. Now, the administrator can permanently delete the vault by clicking **Delete Vault** and entering their password.




Figure 5-50 Deletion request review

7. The deletion request stays active for 24 hours. After that period, the entire process must be restarted.

5.6.18 Storage Account Portal

The *Storage Account Portal* enables GUI guided administration of accounts and containers while Container Mode in IBM Cloud Object Storage is enabled.

The Container Mode was enabled for different reasons (see *IBM Cloud Object Storage Concepts and Architecture*, REDP-5537), and was handled by using APIs only in previous versions. The Storage Account Portal is now being used to make the daily administration in Container Mode easier.

There are two prerequisites to use the Storage Account Portal:

- ▶ Access to Storage Account Portal is granted by the newly introduced role of *Storage Account Administrator* that can be given to Accounts in the Security section. Accounts can have this role in addition to other roles or as a single role. If this role is the only one for an account, users are redirected directly to the Storage Account Portal when logging in to the Manager GUI. Other Accounts will see a new section Storage Account Portal after login and can open the page by clicking **Configure**.
- ▶ There must be at least one Container Vault that is created in the system.

When accessing the Storage Account Portal, it presents a high-level overview of the different Storage Accounts as shown in Figure 5-51, the option to create a usage report (**Usage Summary**), or create a Storage Account.

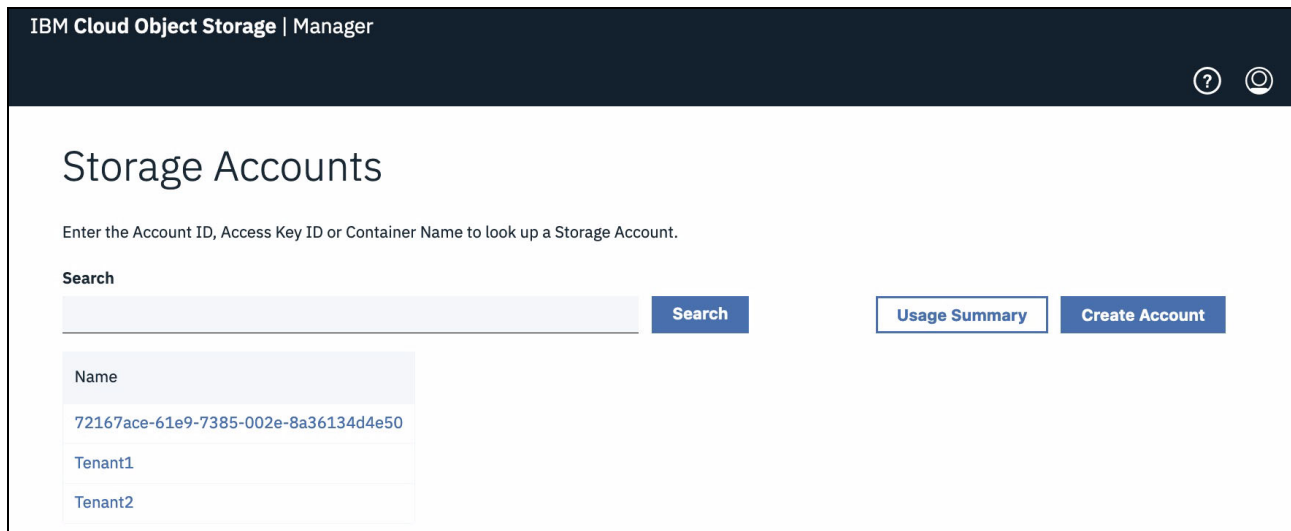


Figure 5-51 Storage Account overview

Selecting one of the Storage Accounts presents an overview to that account with credentials and buckets, as shown in Figure 5-52 on page 137. To change the settings, click **Edit**.

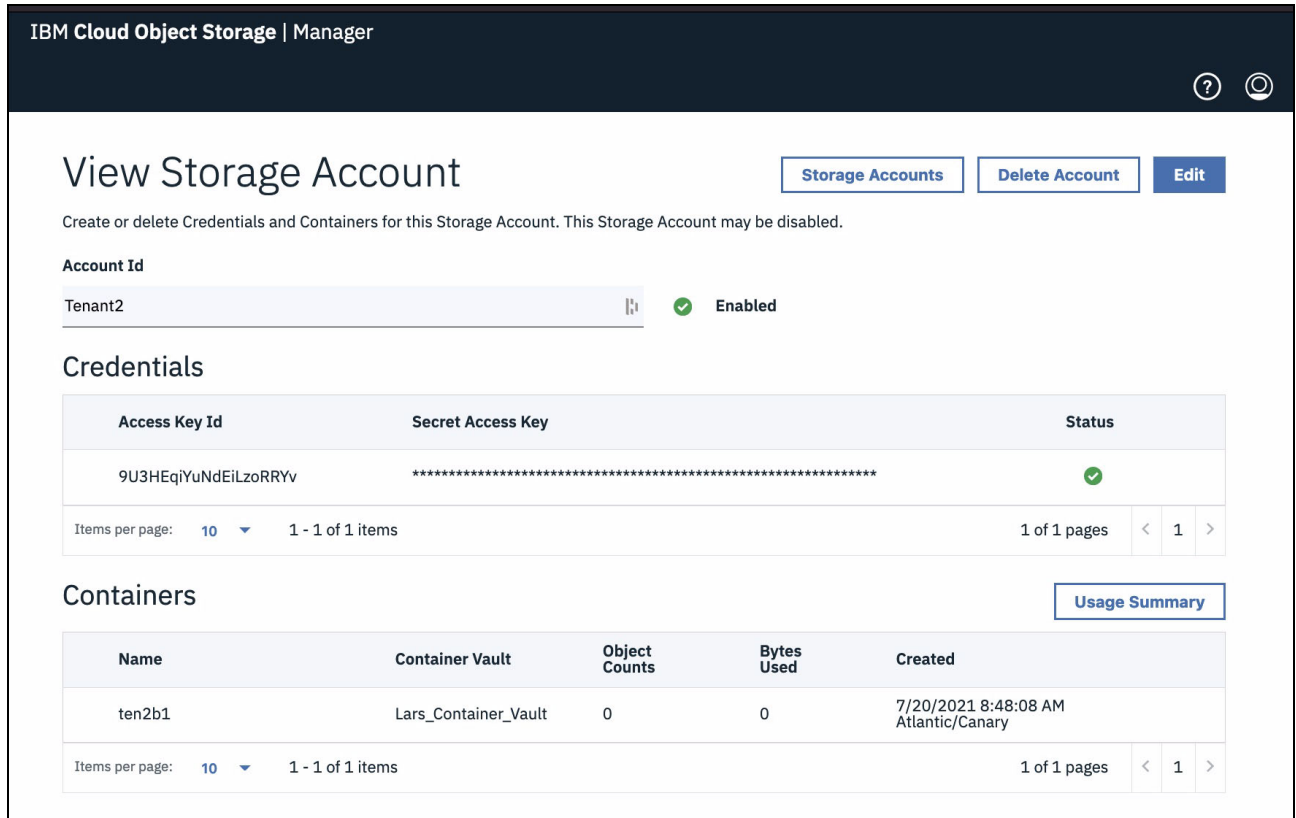


Figure 5-52 Storage Account Details

Creating credentials

Clicking **Create Credentials** generates a new S3 Access Key ID and Secret Access Key for that Account. Each Storage Account can have several key pairs. With these credentials, a user sees all buckets that belong to that Storage Account. More granular ACLs can be set by using the S3 API.

Note: Secret Access Keys are only shown right after creation. Afterward, they are masked and can be retrieved only through the Service API.

Creating a container

A container can be created by selecting a Container Vault (1), giving it a Name (2) and then clicking **Create Container** (3) as shown in Figure 5-53. This new bucket uses the IDA and setting of the Container Vault it is created in. A Container can be deleted by clicking the bin icon next to it (4).

The screenshot shows a web interface titled "Containers". At the top right is a "Usage Summary" button. Below the title, there are two input fields: "Container Vaults" with a dropdown menu showing "Lars_Container_Vault" and "Container Name" with a text input field containing "ten2b2". To the right of the "Container Name" field is a "Create Container" button. Below these fields is a table with the following columns: "Name", "Container Vault", "Object Counts", "Bytes Used", and "Created". The table contains one row for a container named "ten2b1" with 9 object counts, 1565170 bytes used, and a creation date of "7/20/2021 8:48:08 AM Atlantic/Canary". A bin icon is visible to the right of the "Created" column for the "ten2b1" row.

Name	Container Vault	Object Counts	Bytes Used	Created
ten2b1	Lars_Container_Vault	9	1565170	7/20/2021 8:48:08 AM Atlantic/Canary

Figure 5-53 Create new Container

Usage metrics

Usage Summaries can be exported on either the Storage Account overview level or the individual Storage Account. Exporting from the overview generates a report that shows usage in byte hours and objects hours per Storage account. This task can be used for charging different tenants.

Exporting the Usage Summary from a single Storage Account generates a report with individual buckets and the same metrics as above.

All reports can be generated for the current usage, a single calendar month or an individual date range, as shown in Figure 5-54.

The screenshot shows a dialog box titled "Export Usage Summary". The main text reads: "Select current usage or date range to get summary of the space used on account Tenant2." Below this is a section titled "Select Range" with a dropdown menu. The dropdown menu is open, showing three options: "current usage" (which is selected with a checkmark), "single month", and "custom date range". At the bottom right of the dialog box are two buttons: "Cancel" and "Export".

Figure 5-54 Export Usage Summary

5.7 Step 5: Verifying the solution

If you successfully completed all the steps that are described in 5.3, “Step 1: Installing the solution” on page 96 and 5.6, “Step 4: Configuring the Manager GUI” on page 106, you should have a fully functional IBM Cloud Object Storage System with a vault created, and a user with the correct permissions to manage data in IBM Cloud Object Storage.

5.7.1 Programs to verify and test IBM Cloud Object Storage

The following open source or no-charge programs can be used access IBM Cloud Object Storage:

- ▶ [Cyberduck](#) (available for Windows and Mac)
- ▶ [S3 Browser](#) (available for Windows)
- ▶ [AWS CLI](#) (available for Windows, Mac, and Linux)

Important: Most no-charge programs that are available to access IBM Cloud Object Storage are bandwidth-limited by design. Review the program to see whether it limits bandwidth usage. Do not use a no-charge program that limits bandwidth usage to test the performance of IBM Cloud Object Storage. In the previous list, Cyberduck and S3 Browser limit bandwidth usage in the no-charge version. The AWS CLI option does not limit bandwidth usage.

The rest of this section document uses the AWS CLI on Linux to show how to access IBM Cloud Object Storage from the CLI.

We also use HTTP instead of HTTPS.

The instructions for how to install AWS CLI on your instance of Linux are *not* covered in this publication. For more information about how to install the AWS CLI, see [this web page](#).

5.7.2 Configuring AWS CLI

To configure the AWS CLI, complete the following steps:

1. Run the command that is shown in Example 5-21.

Example 5-21 Configuring the AWS CLI

```
[root@linux ~]# aws configure
AWS Access Key ID [None]: rQpHeqhFPPUMAazJbe3T
AWS Secret Access Key [None]: MJ2SepQvnXGoYKmHtnqKrZQWz79P0dWSJH1JNTz0
Default region name [None]:
Default output format [None]:
```

2. Enter the Access Key ID and Secret Access Key for the user that you created.
3. Leave the Default region name and Default output format fields blank.

5.7.3 Uploading an object

To upload an object, run the command that is shown in Example 5-22 by using the IP address of an Accesser appliance and vault that you created.

Example 5-22 Uploading a file

```
[root@linux ~]# aws --endpoint-url http://10.69.70.110 s3 cp testfile
s3://pdc_vault_01
upload: ./testfile to s3://pdc_vault_01/testfile
[root@linux ~]#
```

The file `testfile` is now uploaded to IBM Cloud Object Storage as an object. The manager GUI should now show that Raw Space and Usable Space increased for your vault. The amount of increase depends on the size of `testfile`.

Note: This command assumes that the `testfile` file exists on your Linux server.

5.7.4 Listing objects

To list objects in your bucket, run the command that is shown in Example 5-23 by using the IP address of an Accesser appliance and vault that you created.

Example 5-23 Listing objects in a bucket

```
[root@linux ~]# aws --endpoint-url http://10.69.70.110 s3 ls s3://pdc_vault_01
2019-05-01 16:42:30 1073741824 testfile
[root@linux ~]#
```

The output should list the object `testfile`.

5.7.5 Downloading an object

To download an object, run the command that is shown in Example 5-24 by using the IP address of an Accesser appliance and vault that you created.

Example 5-24 Downloading an object

```
[root@linux ~]# aws --endpoint-url http://10.69.70.110 s3 cp
s3://pdc_vault_01/testfile testfile_new
download: s3://pdc_vault_01/testfile to ./testfile_new
[root@linux ~]#
```

You now downloaded the object `testfile` in IBM Cloud Object Storage to the file `testfile_new` on your Linux server.

Note: This command assumes that the `testfile_new` file did not exist on your Linux server.

5.7.6 Deleting an object

To delete an object in IBM Cloud Object Storage, run the command that is shown in Example 5-25 on page 141 by using the IP address of an Accesser appliance and vault that you created.

Example 5-25 Deleting an object

```
[root@linux ~]# aws --endpoint-url http://10.69.70.110 s3 rm  
s3://pdc_vault_01/testfile  
delete: s3://pdc_vault_01/testfile  
[root@linux ~]#
```

You now deleted the object `testfile` in IBM Cloud Object Storage. The manager GUI should now show that Raw Space and Usable Space decreased for your vault. The amount of decrease depends on the size of `testfile`.

Note: Deleting files in IBM Cloud Object Storage can take some time to complete. The decrease of space might also take some time to appear in the Manager GUI.

You also can list the objects in your vault by using the `ls` command to see whether it is removed.

5.7.7 Differences between S3 and the IBM Cloud Object Storage System APIs

Table 5-1 highlights some functional differences between S3 API and the IBM Cloud Object Storage API.

Table 5-1 Functional differences between S3 API and IBM Cloud Object Storage API

Feature	S3	IBM Cloud Object Storage
Object size limitations	5 TB	Single objects up to 10 TB with streaming upload support or S3 Multipart Upload.
Retained Version Count Limitations	No explicit limit.	No limit for number of versions per object in Container Mode. A maximum of 1000 retained versions are allowed per object in vault mode.
Vault (Bucket) Granular ACL	ACL specifies granular roles for vaults and objects.	In vault mode - users who are configured in the Manager Web Interface can be granted read/write , read-only , or no-access permissions to any vault. These settings apply to the entire vault.
Vault (Bucket) Granular Data Reliability	Allows a storage class to be configured for each object. All objects that are stored in any vault share reliability characteristics.	Vault reliability characteristics are determined at vault creation time.
Traditional Authentication Mechanisms	Uses a custom HTTP scheme based on a keyed-HMAC.	In addition to Access Key authentication, these authentication methods are also supported: <ul style="list-style-type: none">▶ HTTP Basic over HTTP and HTTPS▶ PKI over HTTPS▶ Anonymous
Separated Audit and Logging Functions	Access logs can be enabled, and the logs are stored in a separate bucket.	Accesser node collects both access logs and audit trail information but does not expose it through the API.

Feature	S3	IBM Cloud Object Storage
Encryption and Cryptographic Integrity	Server-side encryption is supported.	<ul style="list-style-type: none"> ▶ An Object Vault can be configured to store information in an encrypted form. It must be configured at the vault or bucket level through the System Manager. These settings cannot be viewed or edited through the API. ▶ Request signing is supported. ▶ Non-cryptographic ▶ MD5 checksums are calculated and stored with objects.
Lifecycle Configuration	Supported various transition and expiration rules configurations per object.	Does not support policy-based migration of data to alternative storage classes, or archiving of data. IBM Cloud Object Storage does not support expiration on vaults or containers with versioning enabled.
Vault (Bucket) Location Constraints	Allows buckets to be created with specific location constraints.	Can configure a system to allow data in one vault to be in a separate geographical location from data on another vault. It is configured when vaults are created in the Manager Web Interface.
Hard Quota Function	Does not support quotas for buckets.	A hard quota can be configured for an object vault. HTTP status code 507 (Insufficient Storage) is returned for a write request that would cause a hard quota to be exceeded.
Object retention	S3 Object Lock can be used to set retention period for objects, restricted to versioned buckets only, and provide means to lock a specific object version.	Retention period can be set for a bucket providing default retention rules for all objects in vault. Versioning is disabled for retention-enabled vaults. The S3 retention API extension can extend the retention period for objects and set or remove legal holds.

5.7.8 For more information

For more information, see the *AWS CLI S3 Developer Guide*, which is available at [IBM Documentation](#).

Note: You can select the version of the IBM Cloud Object Storage software that you use by clicking **Change version or product**.

5.8 Basic installation troubleshooting

The following sections describe some of the most common errors or issues that you might have during the installation and initial configuration of IBM Cloud Object Storage.

5.8.1 Networking issues

Networking-related issues are the most common cause of problems during an installation and configuration of IBM Cloud Object Storage. If one or more of your appliances cannot communicate with IBM Cloud Object Storage or any part of your network, the cause is almost always related to networking.

The basic networking rule of IBM Cloud Object Storage is that the Manager must communicate with every other appliance. The Slicestors also must communicate with every other appliance. Every Accesser must communicate with every Slicestor and the Manager appliance.

If the appliance cannot communicate (for example, ping) with its gateway, check one or more of the following items:

- ▶ Network settings: Review the network settings that you input in to the appliance. It is easy to make a mistake when entering in the settings.
- ▶ Firewall: Verify that your network allows pings.
- ▶ Network cable: Verify that the network cable between the appliance and your network switch is connected and valid.
- ▶ Network port: Verify that the network ports on the switch are on and active. Some switches require a manual configuration to become active.
- ▶ VLAN tagging: Ensure that VLAN tagging is turned off. IBM Cloud Object Storage does not support VLAN tagging.
- ▶ Access ports: Ensure that all ports on the switch are defined as access ports.
- ▶ Switch configuration: If you are bonding more than one interface, verify that your switch supports LACP and that it is configured correctly.

If the appliance cannot communicate with appliances in remote sites, verify one or more of the following items:

- ▶ The local site and remote site can communicate and that it is enabled.
- ▶ The local site and remote site can ping each other.
- ▶ All the correct firewall ports were opened.

If the appliance does not appear in the Manager GUI as a pending device, check one or more of the following items:

- ▶ Verify that you set the Manager with the correct IP address from the CLI.
- ▶ Ping the manager.

5.8.2 Installing IBM Cloud Object Storage

You might see the following issues when you are trying to install IBM Cloud Object Storage:

- ▶ Issue: You cannot boot from the USB key:
 - Answer: The USB key is not correctly installed with IBM Cloud Object Storage. Reimage the USB key and try in another appliance.
 - Answer: You need to press F11 when the IBM logo appears.
- ▶ Issue: The appliance starts with the wrong image (for example, Slicestor appliance boots with Accesser the appliance image).
Answer: Reinstall the appliance with the correct image.

5.8.3 Pending appliances

You might see the following issues when you are accepting appliances into IBM Cloud Object Storage:

- ▶ Issue: The appliance is rejected instead of being accepted into IBM Cloud Object Storage.
Answer: Reinstall IBM Cloud Object Storage on the appliance.
- ▶ Issue: The appliance is deleted instead of being accepted into IBM Cloud Object Storage.
Answer: Reinstall IBM Cloud Object Storage on the appliance.
- ▶ Issue: The appliances do not appear in the list of appliances to accept.
Answer: This problem often is caused by a network issue. For more information, see 5.8.1, “Networking issues” on page 142.

5.8.4 S3 API issues

You might see the following issues when you are trying to connect to IBM Cloud Object Storage externally:

- ▶ Issue: Cannot connect to device.
Answer: Validate that you are using the correct IP address and port of the Accesser appliance.
- ▶ Issue: Cannot authenticate.
Answer: Validate that you enabled Access Key ID authenticate in the Manager GUI.
- ▶ Issue: Cannot list directory:
 - Answer: Validate that you created the vault.
 - Answer: Validate that you deployed the vault to the Accesser through an access pool.
 - Answer: Validate that you were granted the correct rights to the vault.

5.9 IBM Call Home and log collection

You can use two product features to create a support case with IBM Support: IBM Call Home and log collection. These features are described in this section.

5.9.1 Configuring IBM Call Home

When IBM Call Home is enabled, the Manager automatically sends an email to open a support case with IBM Customer Support when an incident is created.

Note: For IBM-branded hardware, the IBM Call Home feature supports device-specific incidents such as disk-related problems and hardware incidents (for example, fan and power supply).

In addition, Call Home supports system-level (for example, storage pool) incidents on IBM-branded hardware and non-IBM branded hardware. SMTP must be configured to send outbound Call Home notifications.

Complete the following steps to configure Site Detail with accurate information to facilitate hardware part replacements:

1. Select **Sites** in the Configure section on the Manager UI (see Figure 5-55).

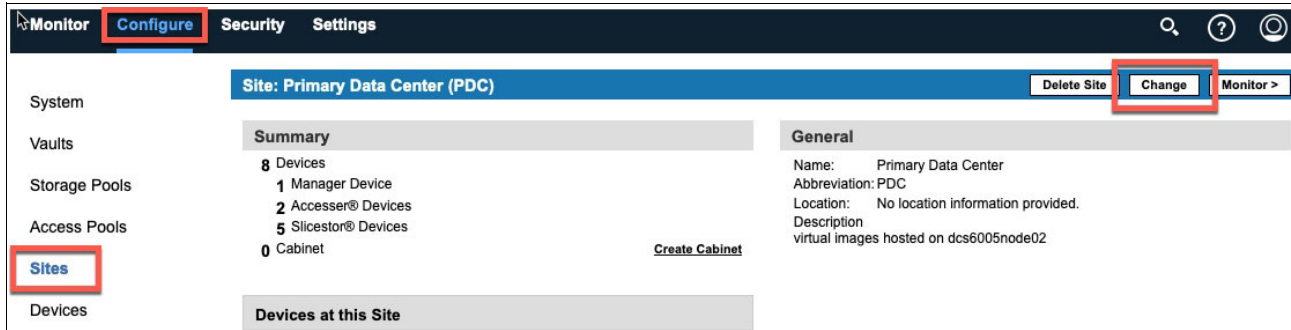


Figure 5-55 Configuring the Call Home feature (1/2)

2. Update the **Name** and **Address** of the sites in the Edit Site section and click **Update**. This section must be completed for all Sites (see Figure 5-56).

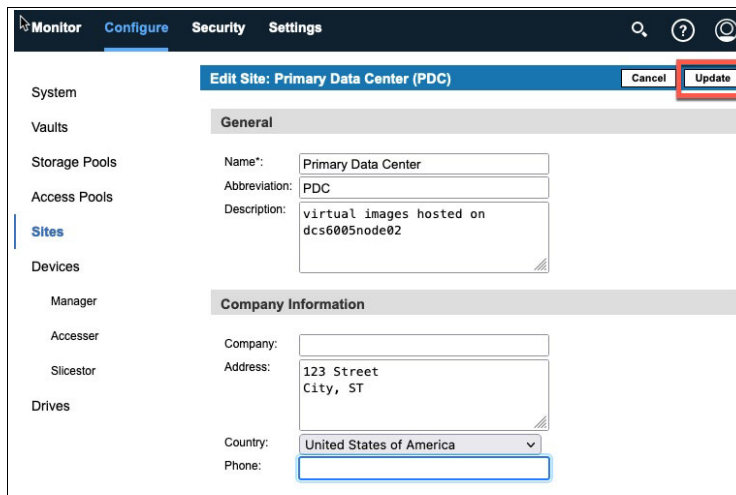


Figure 5-56 Configuring the Call Home feature (2/2)

Complete the following steps:

- a. Set up the **SMTP Server** by selecting the **Settings** tab.
- b. Click the **Monitoring** section, followed by the **SMTP** link.
- c. Complete the appropriate information for your company and save. You can alternatively test your settings by clicking **Test** in the upper right.
- d. Setup the **System Owner** by selecting the **Settings** tab.
- e. Select the **Systems** Section, followed by the **System Owner** link.
- f. Complete the appropriate information for your company and save.
- g. Setup the **System Display Name** by selecting the **Settings** tab.
- h. Select the **System** link on the left followed by the **Display Settings** on the right.
- i. Change the **System Display Name** followed and select **Update**.
- j. Click the **Settings** tab.

- k. Select the **Support** link on the left followed by the **Call Home** link to the right.
- l. Select **Enable Call Home** to notify IBM Customer Support of open incidents.
The configuration fields are available to be completed.
- m. Enter the IBM Customer Number.

Note: The IBM Customer Number is a 6- or 7-digit number that is given to the customer for use with IBM Customer Support. Support Cases that are generated by Call Home are opened under this customer number.

- n. Select the **Country** where the site is located.
- o. Optionally, select the Support Area where the customer's account is managed:
 - North or South America. This area includes Central America and the Caribbean Islands.
 - Africa, Asia, Australia, or Europe. This area includes the Middle East and the Pacific Islands.
 - Other: An email address required. Notifications are sent to only the email addresses that are entered in the Email Addresses to Copy field and are *not* sent to IBM Customer Support. At least one email address is required.

For systems that span across hemispheres, contact IBM Customer Support to determine which area to select.
- p. Enter any Email Addresses to Copy.
These email addresses receive the same notifications that are sent to IBM Customer Support. Enter email addresses in a comma-separated list or on separate lines.
- q. Click **Update**.

5.9.2 Log collection

The following approved secure methods are available for uploading IBM Cloud Object Storage log files to Enhanced Customer Data Repository (ECuRep) for support:

- ▶ Secure File Transfer Protocol (SFTP), which is FTP over SSH
- ▶ Hypertext Transfer Protocol Secure (HTTPS) web
- ▶ IBM Cloud Object Storage Manager User Interface

Note: For more information about these upload methods, see this [IBM Support web page](#).

ECuRep supports several methods for sending data to IBM. The file size of your data largely determines the methods that are available for use.

Sending data by using SFTP

The following prerequisites must be met for sending data by using SFTP:

- ▶ Transfer ID: An IBM Support File Transfer ID is required to authenticate to the SFTP server, and your password is displayed only at creation. To create an IBM Support File Transfer ID, see [this web page](#).

Log in with your IBMid (IBM intranet ID) and click **Create new transfer ID** (see Figure 5-57 on page 147).

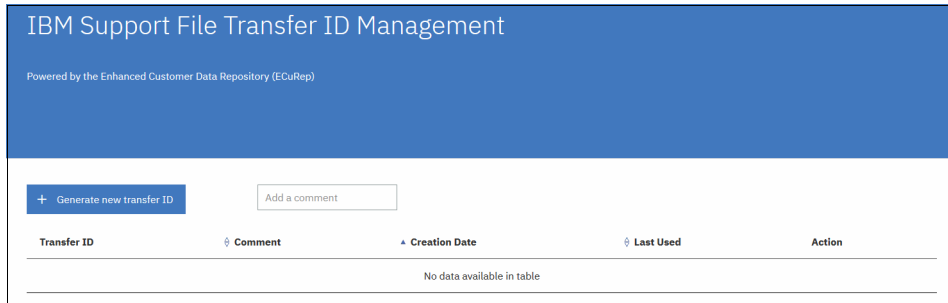


Figure 5-57 Creating a transfer ID

If you lose or forget the password, you must delete your Support File Transfer ID and create another transfer ID.

- ▶ Support case: An IBM Cloud Object Storage Support case must be created on the IBM Support Portal before uploading log files.
- ▶ File naming: When sending IBM Cloud Object Storage logs to IBM Support, the naming convention that is described next must be followed before uploading.

Example case

The following example support case is used in this section:

- ▶ File naming convention: TS<xxxxxxxx>.xxx.yyy
- ▶ Support case number: TS123456789.dump_logs.tar.gz (*ts<Case ID>. <filename>.<extension>*)

Table 5-2 lists the naming convention details for this example support case.

Table 5-2 Naming convention details for the example support case

Field	Description	Example
TSxxxxxx	My Support Case ID	TS123456789
xxx	Short description of the file	dump-logs
yyy	File extension	tar.gz

The ECuRep SFTP server is `sftp.ECuRep.ibm.com`.

Complete the following steps:

1. Log in to the EcuRep server by using the following command:

```
sftp <Transfer_ID>@sftp.ECuRep.ibm.com
```

2. Change to the `/toibm/COS` upload directory:

```
cd /toibm/cos
```

3. Upload the file by using the `put` command. For example:

```
put TS123456789.dump-logs.tar.gz
```

Note: Each file requires a unique name that cannot be overwritten after it is uploaded.

4. After the transfer is complete, enter `quit` and then, press Enter to exit.

Sending data by way of HTTPS Web

The IBM EcuRep - Secure Upload [web page](#) guides you through the upload process. The following steps are provided for general guidance.

The following prerequisites must be met for sending data by way of HTTPS Web:

- ▶ IBM HTTPS web upload site has access to the log file by network share or locally to a client machine.
- ▶ JavaScript is enabled on the web browser to upload log files that are larger than 2 GB.

Complete the following steps for sending data by using HTTPS Web:

1. Start the EcuRep [upload page](#) from a web browser.
2. On the Secure Upload page, ensure that you are on the Case page.
3. Enter the My Support case number that is provided by support (for example, TS123456789).
4. Click **Continue** to proceed to the file upload page.
5. Choose up to five files to upload. No file naming convention is required for HTTPS upload because the My Support Case number is automatically prefixed to each file that is uploaded.
6. Click **Submit** to upload the files.

Note: The browser window must remain open until a new web page is displayed that shows information about the upload that is performed. Leaving or closing the upload page ends the upload process and all data for failed uploads are removed from the web server.

Sending data by using the IBM Cloud Object Storage Manager UI

From the Maintenance tab within the IBM Cloud Object Storage Manager UI, you can collect internal logs across several different IBM Cloud Object Storage devices and manually send them to IBM Support.

Complete the following steps to send data by using the IBM Cloud Object Storage Manager UI:

1. Log in to the Manager UI.
2. Select the **Settings** tab, then **Support**, and then **Log Collection** (see Figure 5-58).

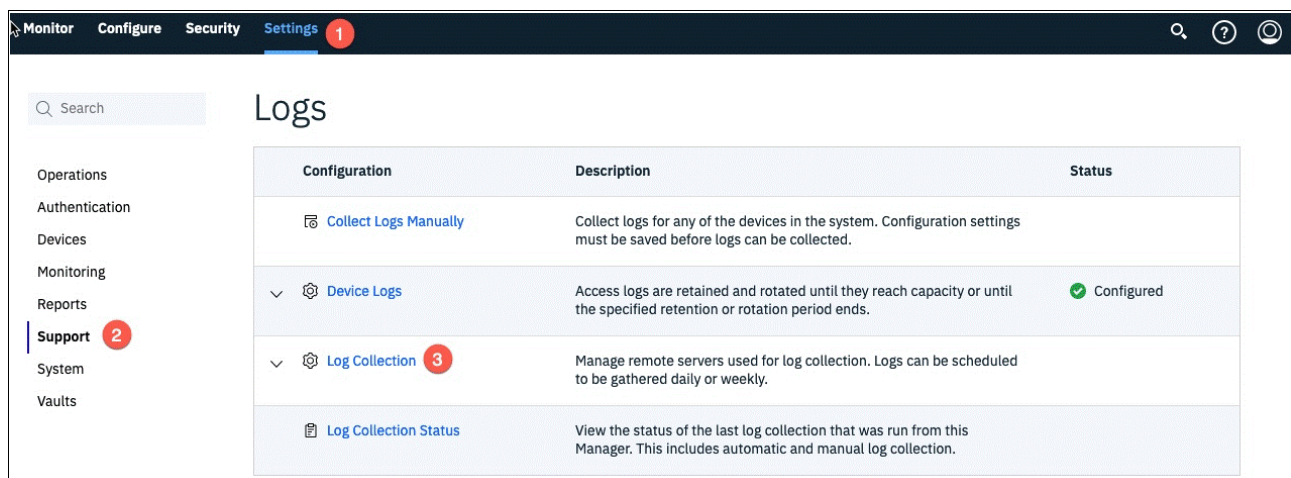


Figure 5-58 Log collection (1/3)

3. Complete the IBM Log Server section with the following information and click **Update** (see Example 5-59):
 - Hostname: sftp.ECuRep.ibm.com
 - Transfer ID: <IBM Transfer ID that was created for your account>
 - Password: <transfer ID password>

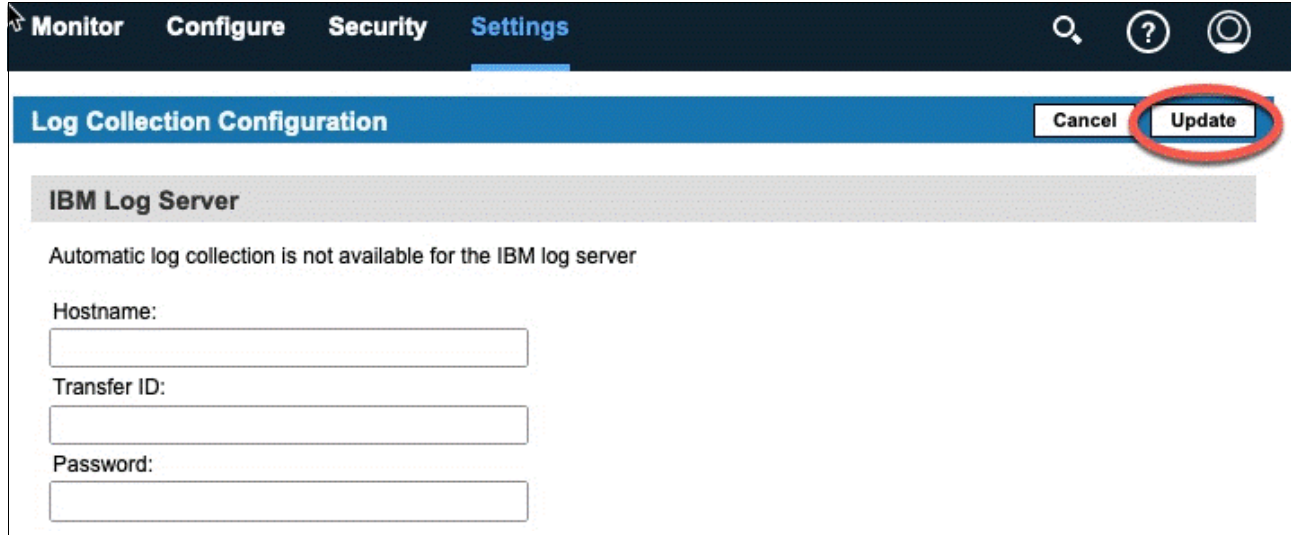


Figure 5-59 Log collection (2/3)

4. To collect logs manually, select the **Settings** tab, followed by the **Support** link on the left, followed by the **Collect Logs Manually** to the right.
5. In the Collect Log window in the Log Destination section, select **IBM Log Server**. Enter the My Support Case number in the Case Number field.
6. Complete the **Filter Criteria** field.
7. Select which devices to send log files from in the **Devices** section.
8. Click **Collect Logs** to start log transfer (see Figure 5-60).

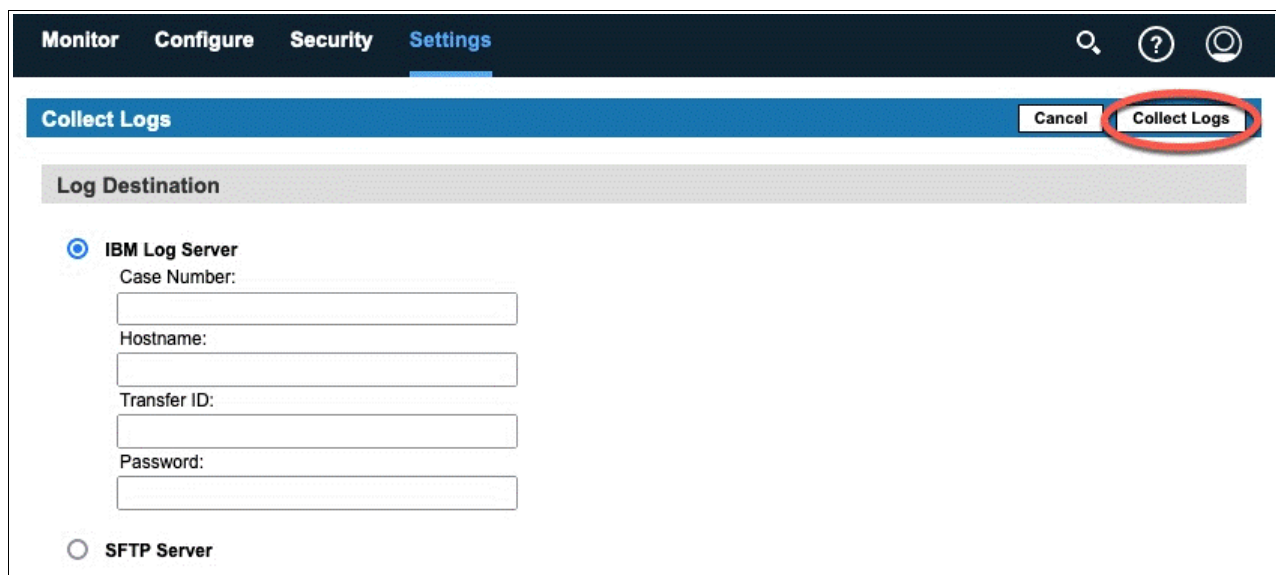


Figure 5-60 Log collection (3/3)

5.10 Upgrading your IBM Cloud Object Storage instance

You can upgrade the physical and virtual appliances to a new version of IBM Cloud Object Storage by using the Manager Web Interface.

Note: Do a manual backup just before performing an upgrade. If the Manager backup has not completed, contact customer support before doing a system upgrade. If an automatic backup is in progress, wait until it is complete before upgrading the Manager.

Upgrading is performed from the **Settings** tab in the Manage GUI. The upgrade process supports n-2 upgrades, which means the starting point of a direct upgrade is a release version whose *second digit* differs by less than or equal to two.

Upgrading is based on the concept of a single upgrade queue, containing the devices to be upgraded. A device that is selected for upgrade is placed in a queue and upgraded, while ensuring the health of the vaults across the system are not compromised. Before a device upgrade, the Manager application tests to see whether the health of any of the vaults that are associated with the device are negatively impacted if the device goes down.

Upgrade can be performed on:

- ▶ A single device
- ▶ Selected devices
- ▶ An entire access pool or storage pool

Checks will be performed to ensure that a device comes back online correctly after upgrade.

Here are best practices for upgrading IBM Cloud Object Storage devices:

1. Manager upgrade
2. Upgrade of storage pools and access pools

Order of upgrade path between Accessers and Slicestors appliances is not important. Both can be upgraded concurrently or separately.

Notes:

- ▶ Order of upgrade path between accessers and slicestors appliances is not important. Both can be upgraded concurrently or separately.
- ▶ As best practice, test the upgrade on a staging environment before upgrading the production. Many customers have staging environments where they test the new release for few days to make sure that there are no issues.
- ▶ Upgrade only a single Slicestor first. If it works with no issues, let the system upgrade the rest of the Slicestors in the storage pool automatically.
- ▶ In CD Mode, the system enforces a delay between individual Slicestor upgrades to avoid data unavailability and provide enough time to rebuild the missing slices that were created during the Slicestor upgrade procedure.
- ▶ Upgrade only a single Accesser to the new code. Test if applications are working sufficiently, and then upgrade the rest of the Accessers in the access pool.
- ▶ During the upgrade, individual Accessers are unavailable for a short amount of time. To minimize application impact, make sure that your applications are configured with automatic load-balancing and fail-over capability when accessing IBM Cloud Object Storage. If your load balancer (LB) or application does not detect and redirect traffic automatically from unavailable Accessers, manual steps may be required before and after each Accesser upgrade.

5.10.1 Upgrade procedure considerations

The high-level procedure consists of the following steps:

1. Transfer the upgrade compressed file to the machine running the browser.
See alternative procedure in the note below.
2. Browse and upload one upgrade compressed file from the Upgrade page.
See alternative procedure in the note below.
3. Initiate the Manager upgrade from the System Software Upgrade page.

For upgrade purposes, the file you need is in the following format:

`c1evos-a.bb.cc.ddd-upgrader.zip`

`a.bb.cc.ddd` references the target version.

The upgrade package is available at [IBM Fix Central](#).

Contact your sales team or IBM support to determine which version you should install. This file contains the upgrade version for the Manager, Accesser, and Slicestor appliances on the same installation media.

Note: In some circumstances, a desirable alternative to Steps 1 and 2 is to transfer the image file directly to the Manager device from an alternative source machine. It is beneficial when the operator is working over a low-bandwidth connection for which the upload time of the upgrade image is prohibitive. In such cases, the operator can log in to a jump host elsewhere on the network that already contains the upgrade image and send it directly to the Manager device. The following example curl command accomplishes this task:

```
curl --insecure -u admin -F upgradeFile=@clevos-3.15.7.60-upgrader.zip -F action=installUpgradeFile "https://<manager>:443/manager/api/json/1.0/installUpgradeRepo.adm"
```

5.10.2 Upgrading the Manager

Perform the following steps to upgrade the Manager:

1. Log in to Manager GUI.
2. Select **Maintenance** → **System Software Upgrade**. You see the current version number of your IBM Cloud Object Storage Manager (see Figure 5-61).



Figure 5-61 Showing your version of IBM Cloud Object Storage Manager

3. Click **Upload** and upload your latest upgrade file (or use the alternative method as described above to upload the upgrade file).
4. Click **Read and Accept License Agreements to Continue** (see Figure 5-62).

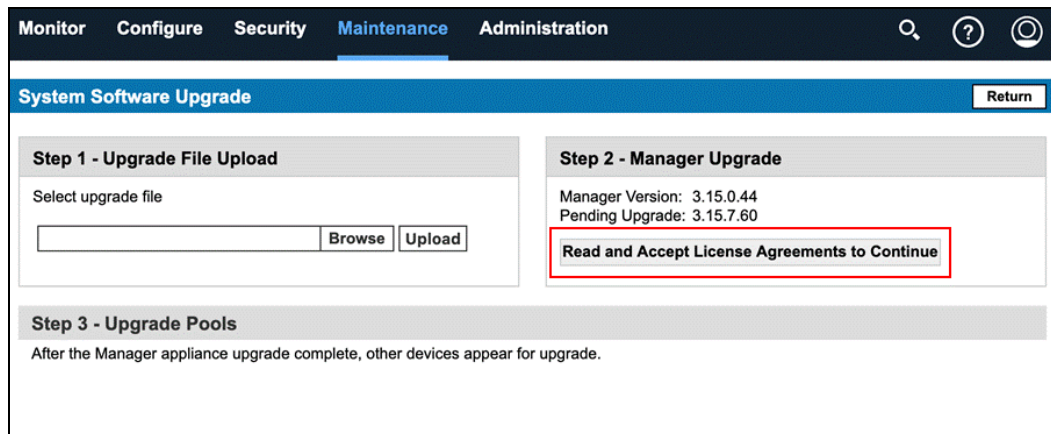


Figure 5-62 Clicking Read and Accept License Agreements to Continue

5. Click **Upgrade**. The process is started. You can monitor the progress in the status box (see Figure 5-63 on page 153).

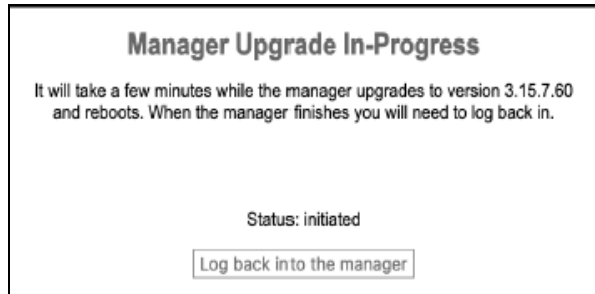


Figure 5-63 Monitoring the IBM Cloud Object Storage Manager upgrade progress

- After upgrade is finished, you must log in again to the Manager, and will see the message that your device's version is different from manager version (see Figure 5-64). It means your manager upgrade means that your manager upgrade completed, and you must upgrade your devices as a next step (see the next section).

5 of the devices in the system are running a different version of software than the Manager. [Please proceed to the Upgrade page »](#)

Figure 5-64 Message about a version mismatch

5.10.3 Upgrading IBM Cloud Object Storage devices

After your manager is upgraded, you should be able to proceed with your device upgrades.

- Select **Settings** → **System Software Upgrade**. You see your devices listed (see Figure 5-65). The screen captures that are provided are for demonstration purposes only and show the fresh installation scenario, when devices are not in a pool yet.

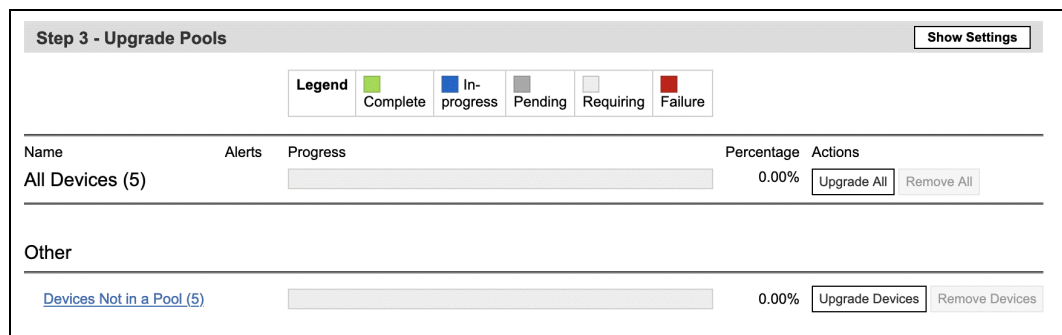


Figure 5-65 List of devices to be upgraded

2. Click the **Upgrade** to process all your devices until all of them are upgraded. At the end of the upgrade, all your devices should move to **Complete** status as seen on Figure 5-66.

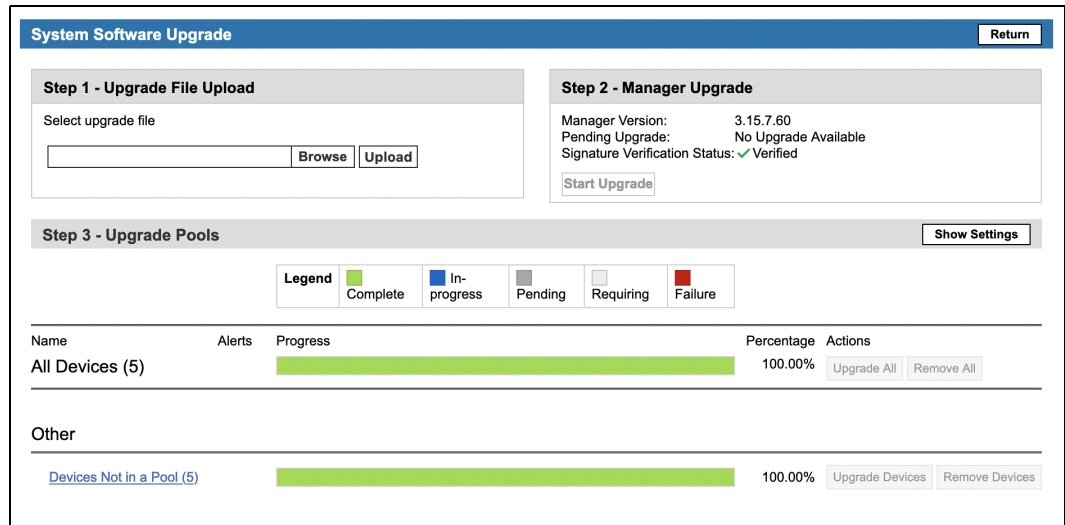


Figure 5-66 All devices are upgraded to 3.15.7.00

5.11 Configuring multiple Manager devices

To have continued visibility into system operation and provisioning capabilities if one Manager device fails, add a second Manager device to the system.

When both Manager devices are running, configuration changes can be performed concurrently, and events can be observed on either Manager device.

A maximum of two Manager devices may run simultaneously within the system.

5.11.1 Adding a second Manager device

Before you begin, ensure that the following prerequisites are in place:

1. The existing Manager device must be associated with a management vault.
1. The existing Manager device must be “in sync” with its associated management vault.
1. In the Management Vault Options subsection of the Configure Management Vault page, **System Configuration** must be enabled.
1. Ensure that there is only one Manager device in the system. The system supports a maximum of two Manager devices.
1. The Manager device that you want to add to the system must be newly re-imaged with a software version that matches the version of the existing Manager device.

About this task

When a system employs multiple Manager devices, certain activities might be unavailable on one Manager device during normal operation. Activities such as upgrade, log collection, and email alert forwarding can be performed only by one Manager at a time. For example, at one point, email alerts might be sent by Manager A, while at another point, email alerts are sent by Manager B.

In systems with multiple Manager devices, the associated management vault must be available to make configuration changes to the system. In an outage affecting the management vault, the outage must be resolved before configuration changes can be made. If the outage cannot be resolved before configuration changes must be made, then the procedure in 5.11.2, “Multiple Manager teardown” on page 156 must be followed, which returns the system to single Manager device operation and allows configuration changes without management vault availability.

Procedure

To add a second Manager device, complete the following steps:

1. On the **Settings** tab, select **System** and then **Multiple Manager**, as shown in Figure 5-67.

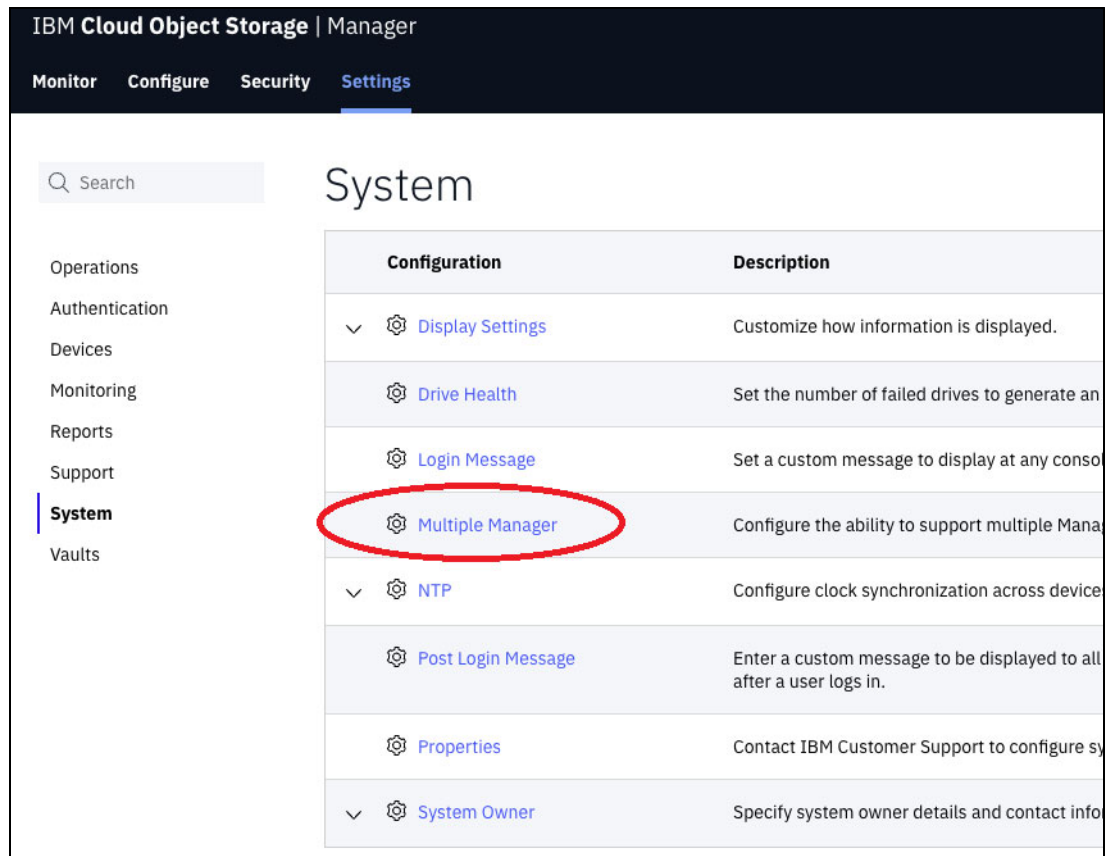


Figure 5-67 Clicking Multiple Manager

2. Select **Enable multiple Managers within the system**, as shown in Figure 5-68.

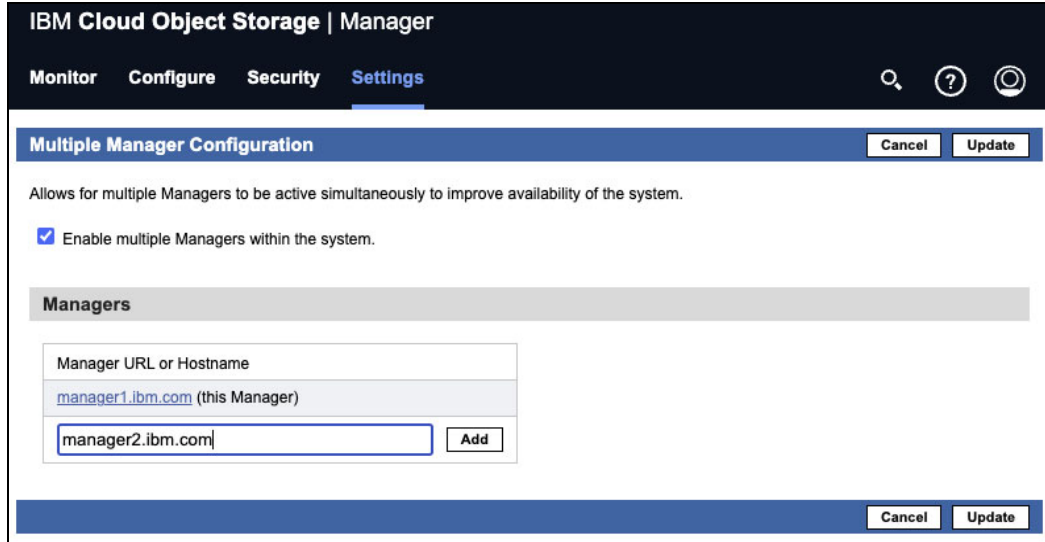


Figure 5-68 Selecting Enable multiple Managers within the system

3. Click **Update**. A Managers section opens.
4. In the Manager URL or Hostname text box, enter the IP address or hostname of the Manager device that you want to add to the system.
5. Click **Add**.
6. If the Manager device that you are adding has different credentials than the defaults, a New Manager Credentials dialog box opens. Enter the username and password and click **Save**.

Results

Note: The SSL connection that is used to communicate with the new Manager device is allowed to proceed even for connections that are considered insecure. Use caution when entering the username and password of the new Manager device.

When the new Manager device is successfully added, a message appears:

The new Manager device has been added successfully.

Note: In larger or older systems, this process can take a long time, and the time is proportional to the size of the Manager database.

5.11.2 Multiple Manager teardown

A Manager device can be removed only if it is not the Manager device performing the removal.

- ▶ Manager device A can remove Manager device B.
- ▶ Manager device B can remove Manager device A.
- ▶ Manager device A cannot remove Manager device A.
- ▶ Manager device B cannot remove Manager device B.

Procedure

To remove a Manager device, complete the following steps:

1. Click the **Configure** tab.
2. Select **Devices** → **Managers** in the navigation window.
3. Click the link of the second Manager device to display the Manager: {device-name} page.
4. Click **Remove** on the Configure Device page to display the Remove Device: {device-name} page, as shown in Figure 5-69.

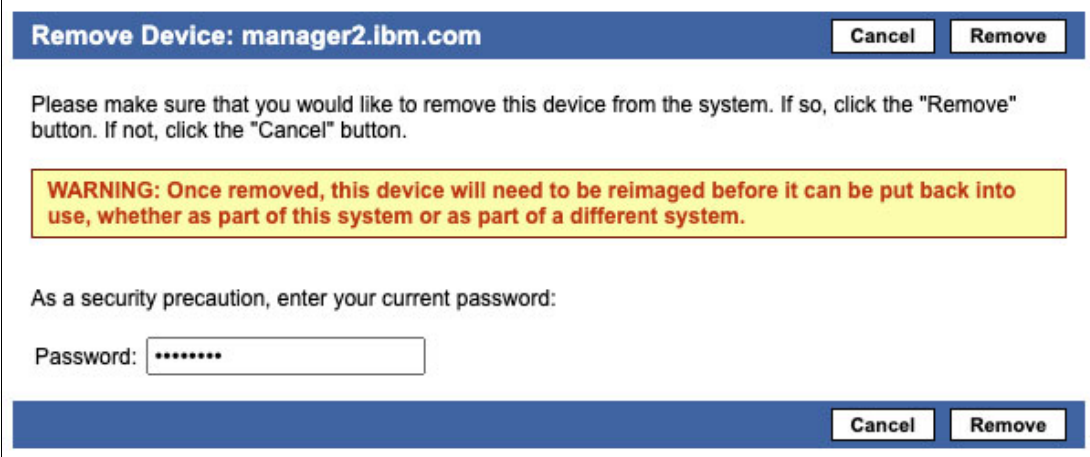


Figure 5-69 Clicking Remove

5. Enter your password, if required, into the Password field.
6. Click **Remove** to remove the device.
7. On the **Settings** tab, select **System** and then **Multiple Manager**.
8. Clear **Enable multiple Managers within the system**, as shown in Figure 5-70.

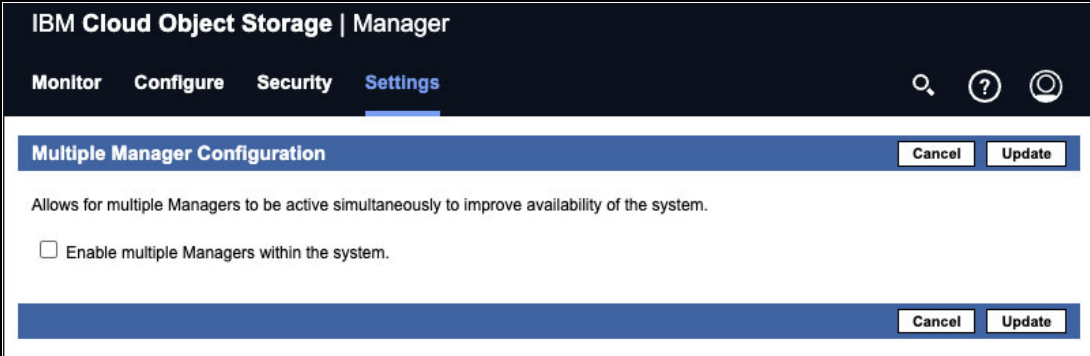


Figure 5-70 Clearing Enable multiple Managers within the system

9. Click **Update**.

Results

After performing the teardown procedure, shut down or re-image the removed Manager device. Allowing the device to continue running can result in unwanted side effects, for example, sending notifications, attempting to communicate with other devices, generating backups, and causing general confusion.

5.12 Multi-factor authentication

Multi-factor authentication (MFA) and integrating infrastructure into a company-wide authentication system is becoming a prevalent requirement. IBM Cloud Object Storage supports MFA against an OpenID Connect (OIDC) provider. This feature can be used to further secure authentication for IBM Cloud Object Storage administrators. When enabled, the IBM Cloud Object Storage Manager supports sign-on through OIDC for the GUI, and OAuth 2.0 bearer tokens for REST API access. The second factor is chosen by the OIDC provider itself, and the factor depends on company-wide guidelines.

Sign-on providers may choose to implement all or part of the OIDC framework. The only official integration is with [IBM Security Verfiy](#), through which it is possible to connect to other OIDC providers.

Note: When Manager REST API access is attempted, the provisioned users cannot call the Manager REST API by using HTTP basic authentication; instead, they must acquire a JSON Web Token (JWT) from the configured OIDC provider. Opaque tokens are not supported on the REST API.

Claims issuer, sub, name, email, and email_verified must be present in all JWTs or available from the configured UserInfo endpoint. Otherwise, authentication might fail. To ensure that all claims are in the tokens, contact a OIDC provider administrator.

For more information and example codes for acquiring an access token, see [Acquire access token to call Manager REST API](#).

The MFA feature can be enabled by using the IBM Cloud Object Storage Manager GUI. As a best practice, review and enter the information together with an admin of the identity provider solution.

To enable MFA, complete the following steps:

1. Click the **Settings** tab
2. Click **Authentication**.
3. Select **OIDC** (see Figure 5-71). For more information, see [Configuring OpenID Connect](#).

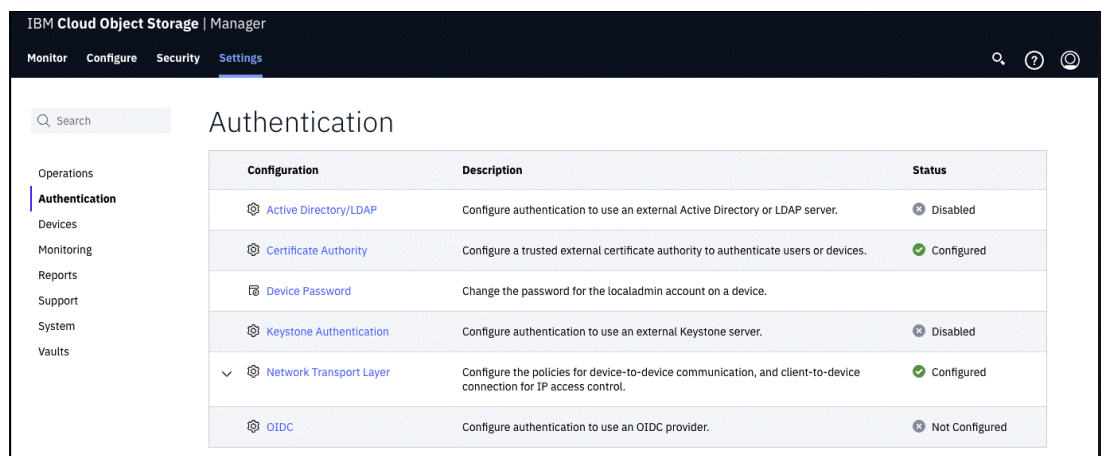


Figure 5-71 Authentication section

- (optional) To discover a configuration automatically, enter the issuer URL and click **Discover**, as shown in Figure 5-72. When the Manager can communicate with the OIDC provider, it detects and parses metadata at the issuer URL with the following path appended, according to OIDC standards:

`/.well-known/openid-configuration`

Because this step is optional, you can skip it to manually configure OIDC. Afterward, the collected metadata is presented, and the values are carried over to the next window where the adopted metadata can be manually modified.

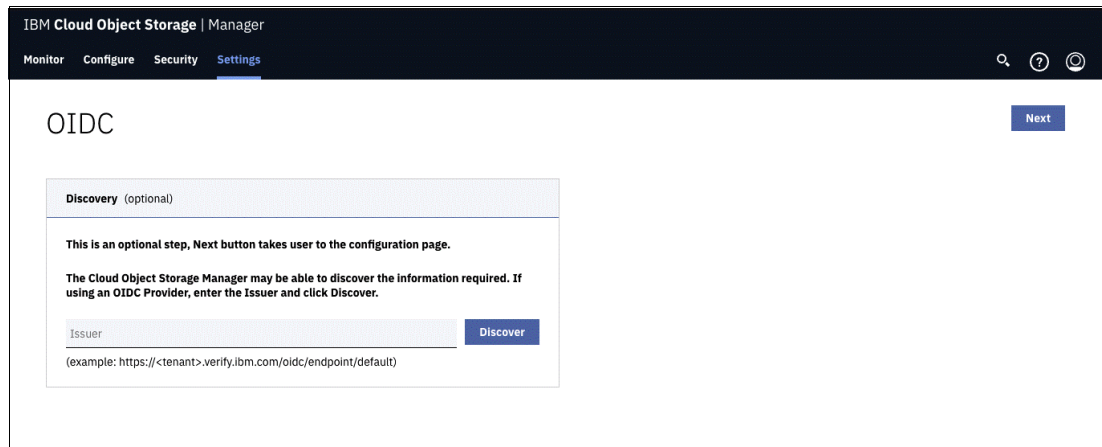


Figure 5-72 *OIDC configuration*

- Click **Enable OIDC Authentication**, as shown in Figure 5-73.

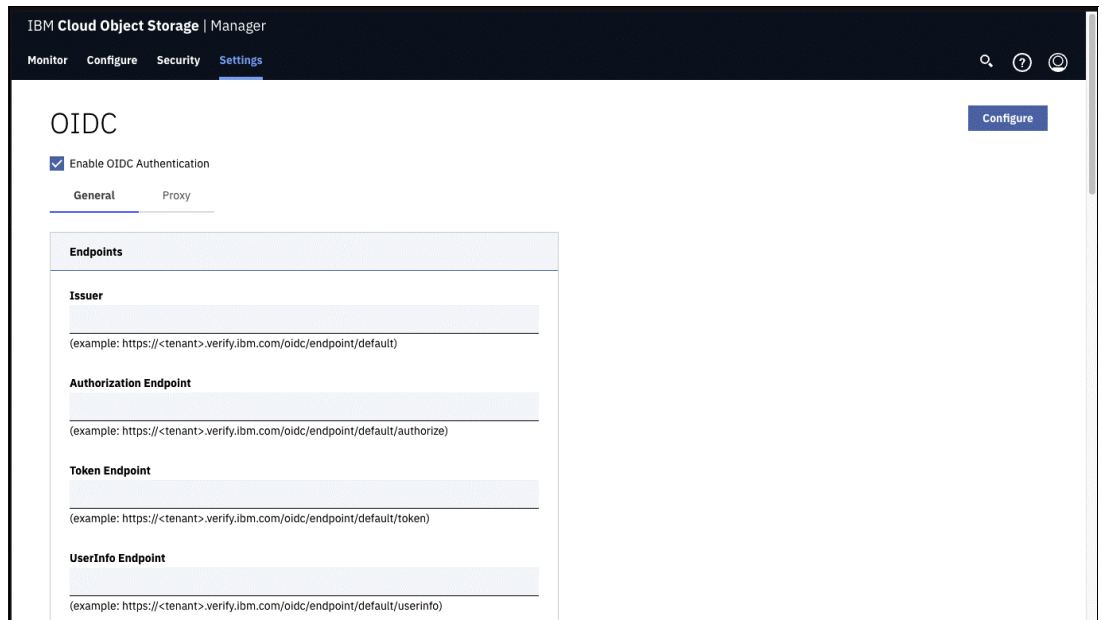


Figure 5-73 *Enable OIDC Authentication window*

6. Enter the values that match the application configuration on the OIDC provider side. If the discovery was successful, some values already might be completed. On this form, the following information must be specified:
 - a. Endpoints:
 - Issuer
 - Authorization Endpoint
 - Token Endpoint
 - UserInfo Endpoint
 - JWKS URI
 - b. Token settings:
 - Group Claim
 - Audience
 - ACR (optional)
 - c. Client Authentication:
 - Client Authentication Scheme
 - Client ID
 - Client Secret
 - Key ID
 - Private Key
 - CA Certificate (optional)
7. (optional) Click the **Proxy** tab to enter a proxy that is used for all communication with the OIDC provider, and then complete the following steps:
 - a. Click **Enable Proxy**.
 - b. Enter the proxy URL and include the scheme (HTTP/HTTPS), domain, and port. If no port is selected, the default HTTP port 80 or HTTPS port 443 is used.
 - c. If the proxy requires authentication, click **Use Authentication** and provide the credentials (username and password).
 - d. If the proxy must be set up by using HTTPS (also called a secure web proxy), a trusted certificate authority might be specified in the **CA Certificate** dialog box. The input to this dialog box can include a certificate chain to a root CA if needed.

Figure 5-74 shows the completed configuration for this step.

OIDC Configure

Enable OIDC Authentication

General **Proxy**

Enable Proxy

Proxy Configuration

URL

(example: https://sampleurl.com:443)

Authentication

Use Authentication

Username

Password

Figure 5-74 Proxy configuration

8. Click **Configure**.
9. Log out of the Manager and verify the new authentication method by logging in by clicking **OIDC Auth Provider**, as shown in Figure 5-75. You are forwarded to the OIDC provider page, where users can log in.

IBM Cloud Object Storage | Manager

IBM Cloud Object Storage Manager

Username:

Password:

Go!

Alternative Login

OIDC Auth Provider >

Figure 5-75 OIDC Auth Provider

5.13 Enabling Write Once Read Many capabilities

The Write Once Read Many (WORM) feature enables IBM Cloud Object Storage to participate in an organization's cyber-resiliency strategy. Whether the demand for WORM capabilities originates in legislation or internal requirements, IBM Cloud Object Storage can provide WORM protection to vaults and objects. IBM Cloud Object Storage supports IBM retention vaults and S3 Object Lock.

5.13.1 IBM retention vaults

IBM retention vaults underwent a compliance assessment to support the following regulatory requirements:

- ▶ Securities and Exchange Commission (SEC) Rule 17a-4(f). For more information, see [this web page](#).
 - SEC 17a-4(f)(2)(ii)(A): Protect data from deletion and overwriting.
 - SEC 17a-4(f)(2)(ii)(B): Automatically verify that the storage system properly stored the data.
 - SEC 17a-4(f)(2)(ii)(C): Manage retention periods for the objects.
 - SEC 17a-4(f)(2)(ii)(D): Download indexes and records.
 - SEC 17a-4(f)(2)(iii/v): Store duplicate copies and provide audit capabilities.
- ▶ Financial Industry Regulatory Authority (FINRA) Rule 4511, which references the requirements of SEC Rule 17a-4(f).
- ▶ Commodity Futures Trading Commission (CFTC) Rule 1.31(b)-(c).

To enable WORM on IBM Cloud Object Storage vaults and set up retention vaults, complete the following steps:

1. Log in to the Manager GUI.
2. Click the **Settings** tab, then **Vaults**, and finally **Protection**, as shown in Figure 5-76.

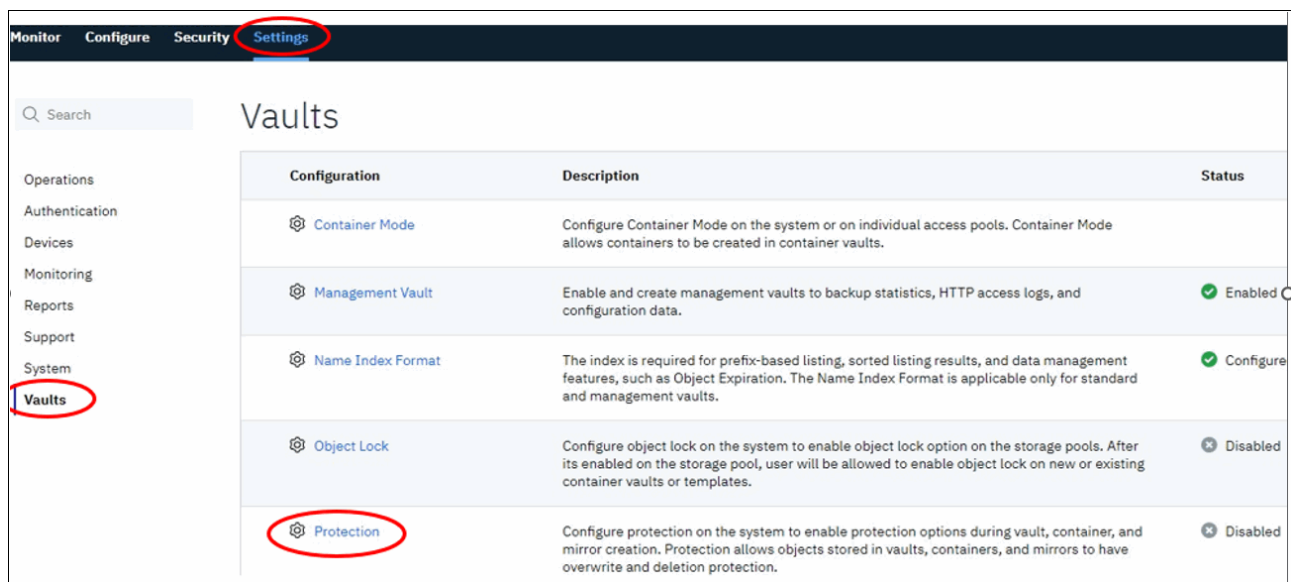


Figure 5-76 IBM Cloud Object Storage Manager GUI settings tab

- When you click **Protection**, you see the **Allow vault protection** option. By default, it is disabled. To enable IBM retention vaults, select the checkbox to enable vault protection. After selecting the checkbox, a Warning dialog box opens. The option **Allow** appears to be disabled, but it is selectable. After you click **Allow**, new options appear under vault protection, as shown in Figure 5-77.

Note: After enabling protection vaults, you cannot disable vault protection if there are any vaults with retention enabled that have data in them.

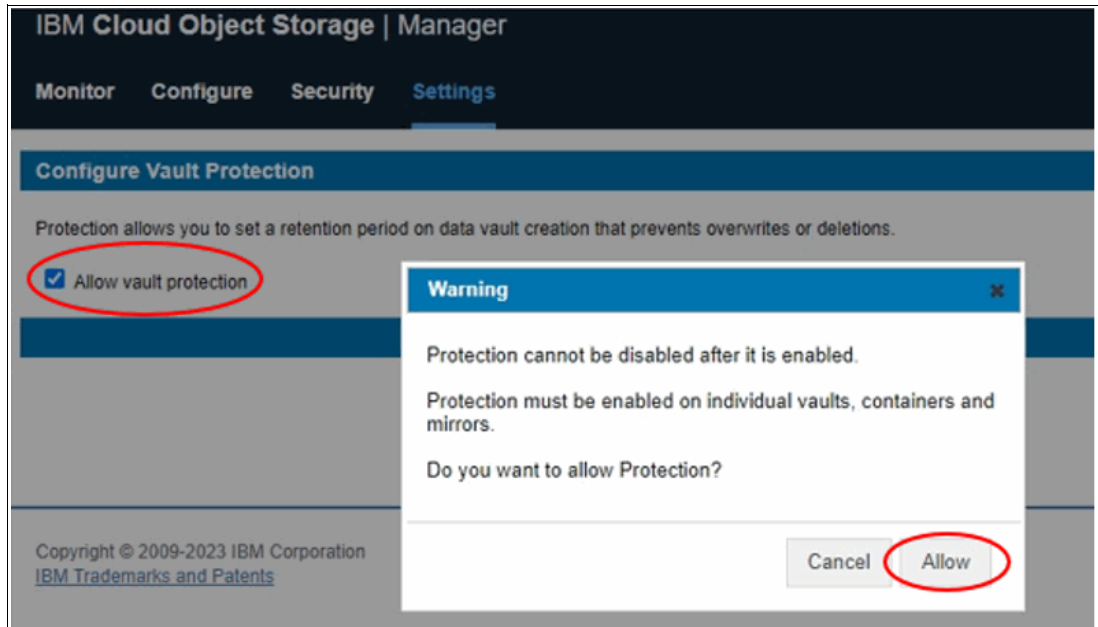


Figure 5-77 Enabling vault protection

- In the Advance vault protection settings window, you can set the following retention settings:
 - ▶ **Enabling permanent retention (Legal holds)**
 - ▶ **System minimum duration**
 - ▶ **System Maximum duration**
 - ▶ **System default retention duration**

Figure 5-78 shows these settings.

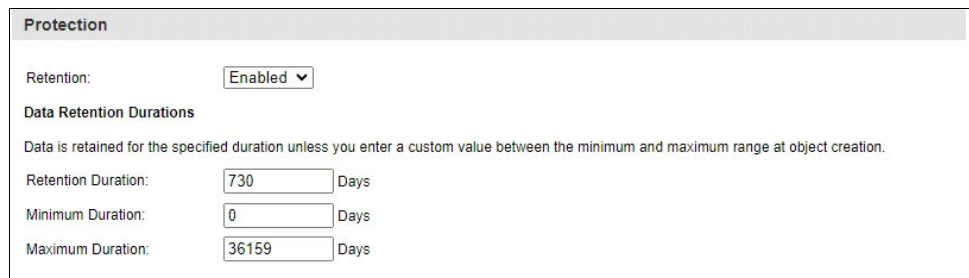


Figure 5-78 Advance vault protection settings

- After you enter the required system retention durations, select **Update** to enable system-wide retention parameters.

- Now, you must enable retention on a vault level. To do so, in the Manager GUI, click **Configure**, then **Vault**, and then **Create VAULT**.

Note: If you have not configured a container or vault mode in the previous steps, you might be required to do so now.

- In the vault creation window, as part of the vault setup process, you can enable retention if needed and choose minimum, maximum, and default durations, as shown in Figure 5-79. By default, the retention duration is prepopulated with the retention settings that were enabled system-wide in step 5 on page 163. However, those settings can be changed if they do not violate the retention durations that were set up system-wide.

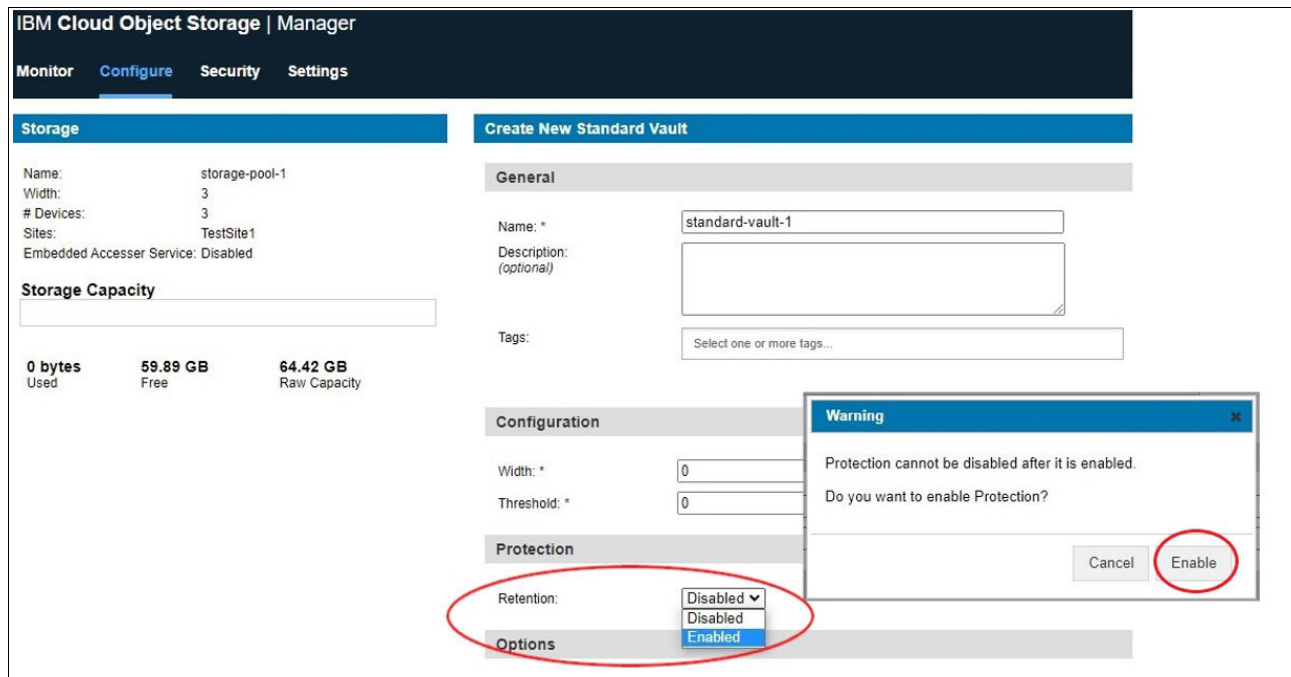


Figure 5-79 New vault creation and enabling vault protection

- After enabling vault protection, complete the creation process of the vault.

5.13.2 S3 Object Lock

To enable S3 Object Lock, IBM Cloud Object Storage System must be in Container Mode, and you must enable Container Mode in three different areas in the system:

- ▶ System level
- ▶ Pool level
- ▶ Container level

Unlike IBM retention vaults, there are no retention durations to configure because those parameters are dedicated by the application layer.

Complete the following steps:

1. In the Manager GUI, click the **Settings** tab, click **Vaults**, and then click **Object Lock**, as shown in Figure 5-80. By default, Object Lock is disabled.

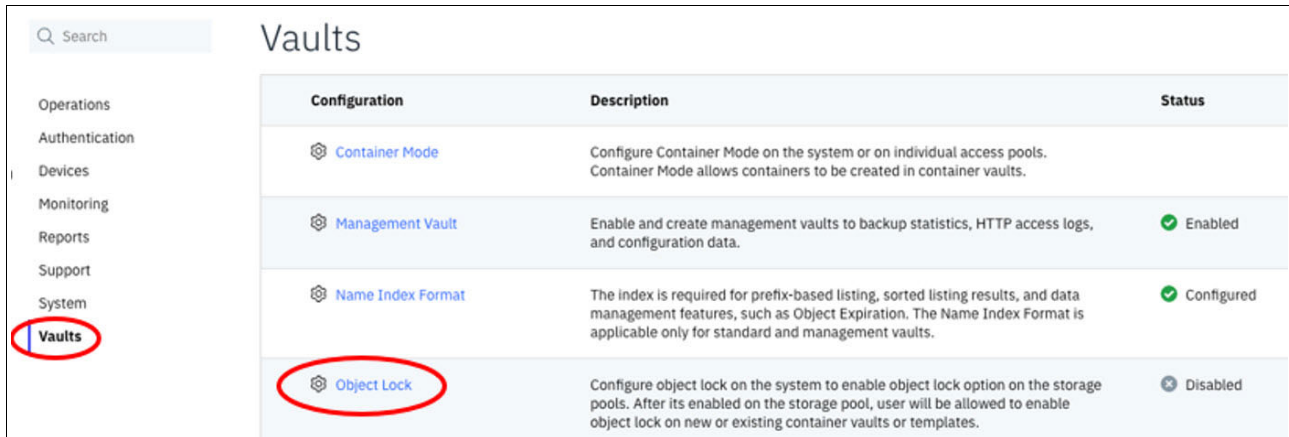


Figure 5-80 Object Lock setting under Vaults

2. Click **Object Lock configuration**, move the slider from **Disabled** to **Enabled**, and click **Save**.
3. Click the **Configure** tab, and then click **Storage pools**.
4. In the Storage pool window, click **Configure** in the Object Lock window, as shown in Figure 5-81.

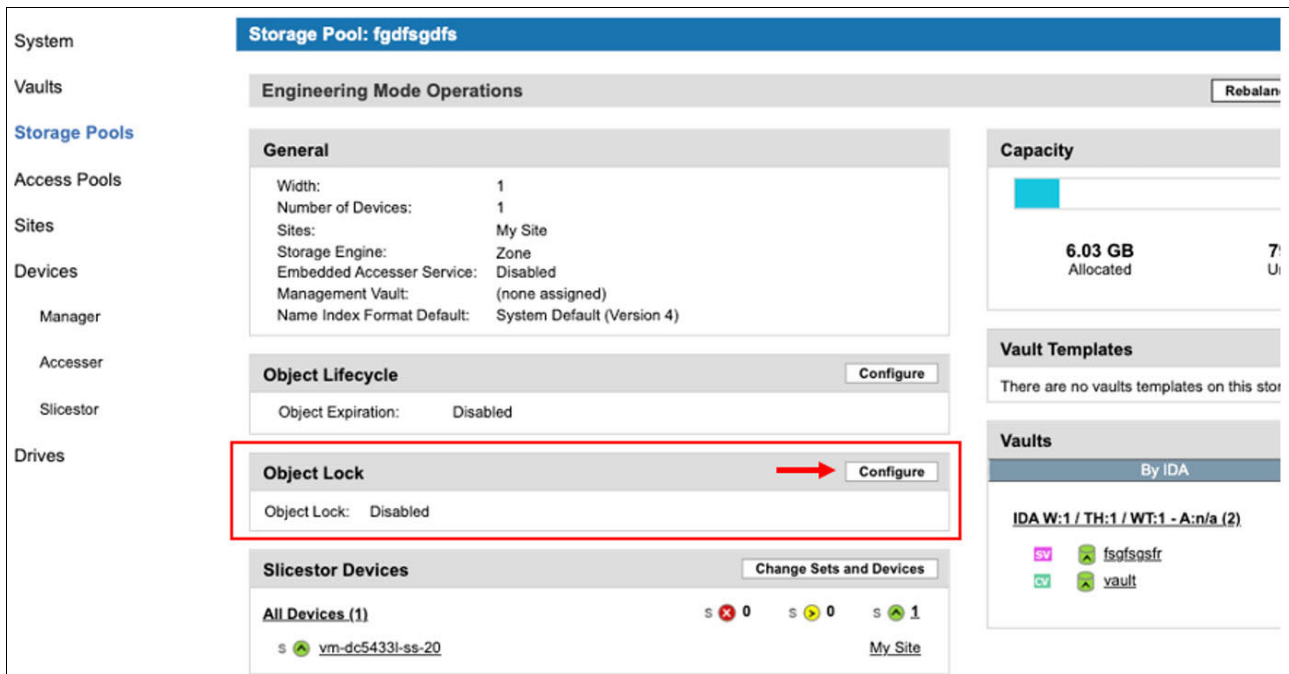


Figure 5-81 Object Lock tile under Storage pools

5. In the Object Lock window, click **Enable**, as shown in Figure 5-82.

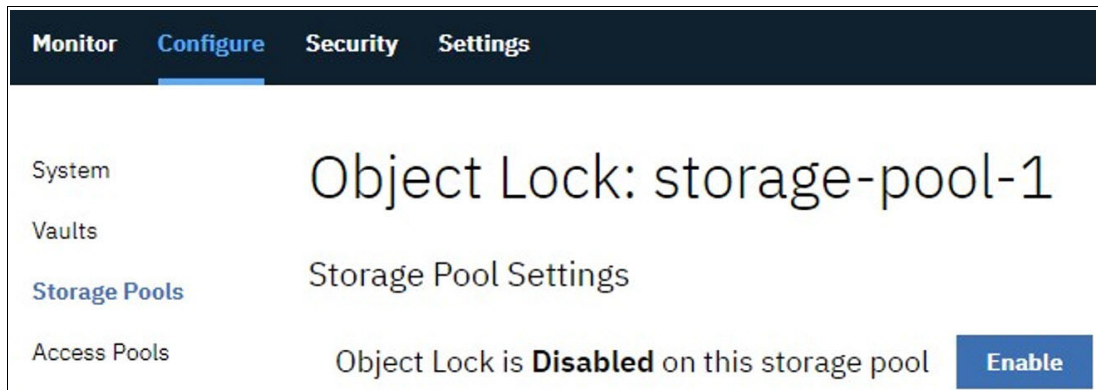


Figure 5-82 Storage pool Object Lock configuration

6. After enabling the Object Lock, a listing of all empty vaults with no Object Lock appears, as shown in Figure 5-83. From here, you can enable Object Lock on any of them.

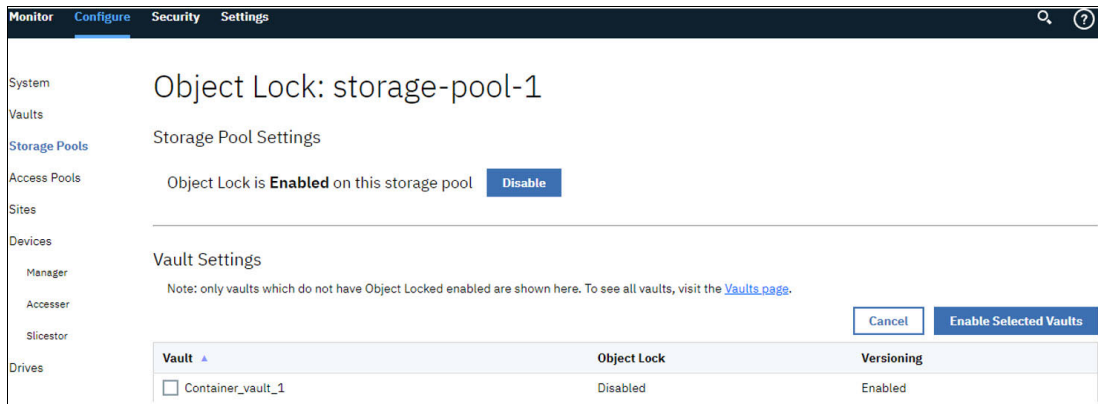


Figure 5-83 Storage pool Object Lock configuration enabled

6. To enable Object Lock on the available vaults, select the checkbox for the vault and click **Enable Selected Vaults**.
7. After the Object Lock setting is enabled on the pool level, any vault that is created in that pool has Object Lock enabled by default.

Note: For more information about and current limitations for Object Lock, see [Object Lock - Feature Description Document](#).



Scalability

This chapter describes the scalability options that are available for IBM Cloud Object Storage.

This chapter includes the following topics:

- ▶ 6.1, “Scaling an IBM Cloud Object Storage System” on page 168
- ▶ 6.2, “Scaling for performance” on page 168
- ▶ 6.3, “Scaling for capacity” on page 171

6.1 Scaling an IBM Cloud Object Storage System

Requirements for the IBM Cloud Object Storage System can change over time, and it might become necessary to adjust the current configuration. This change can be in terms of performance or capacity.

6.1.1 Non-disruptive upgrade

One key principal of IBM Cloud Object Storage is the upgrade process without any disruption of service. This principal is true for software upgrades and for hardware changes.

IBM Cloud Object Storage ensures that enough redundancies are available during the upgrade process. All tools that are needed to complete the upgrades are included in the IBM Cloud Object Storage software. All necessary update steps are guided through the GUI.

In the following sections, the available scaling options are described.

6.2 Scaling for performance

Over time, the performance requirements for IBM Cloud Object Storage can change. This change can occur with new use cases being deployed on IBM Cloud Object Storage, software changes, or temporary events, such as audits, where the workload is directed at the IBM Cloud Object Storage System.

For more information about performance considerations, see Chapter 2, “Planning and sizing an IBM Cloud Object Storage System” on page 23.

To determine where the performance bottleneck is, reviewing the Accesser node utilization and network performance can help. The Accesser node utilization can be viewed per Accesser node in **Monitor** → **Devices** → **Accesser** → **Choose Accesser** and then scroll down to the **Performance** section. A high CPU usage over a longer time, as shown in Figure 6-1, indicates that the Accesser-Layer might be the bottleneck.

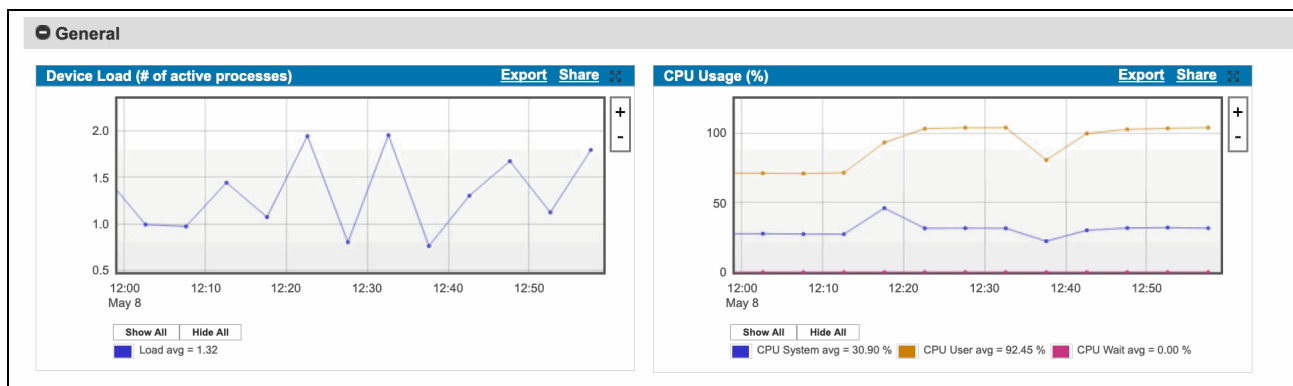


Figure 6-1 Accesser node utilization

The network statistics are shown below the Accesser node usage. Network performance should be checked to identify possible network bottlenecks.

Figure 6-2 shows the network performance monitoring per network port.



Figure 6-2 Network performance

Important: If an IBM Cloud Object Storage System is not performing as expected, contact IBM Support.

6.2.1 Adding Accesser nodes

Accesser nodes can be added at any time to the system without service interruptions. To realize the needed performance, all Accesser node types (physical, virtual, and container) can be added to the system. For more information about performance considerations, see 2.2, “Performance planning” on page 37.

Accesser nodes must be set up as described in 5.5.2, “Configuring the Accesser appliance” on page 102. After the initial configuration, the Accesser node contacts the Manager and asks to be added to the IBM Cloud Object Storage System.

During the acceptance into the system, an Accesser node can be added to an access pool, as shown in Figure 6-3.

Figure 6-3 Adding an Accesser node into IBM Cloud Object Storage

The newly added Accesser node pulls all configurations that belong to the specified access pool, and is ready to accept workload within seconds. On redistribution of the workload, overall performance of the Accesser layer is increased.

Tip: If a load balancer (LB) is used, add the new Accesser node to its configuration. Alternatively, reconfigure the application to use the newly added Accesser node.

6.2.2 Removing Accesser nodes

The removal of an Accesser node is done by way of the GUI in the **Configure** tab. Before removing an Accesser node, all workload that is directed to that Accesser node must be stopped. If an application uses only this Accesser node, it cannot write or read. To remove the Accesser node, select **Configure** → **Devices** → **Accesser** → **Select the Accesser node that is to be removed** → **Remove**, as shown in Figure 6-4.

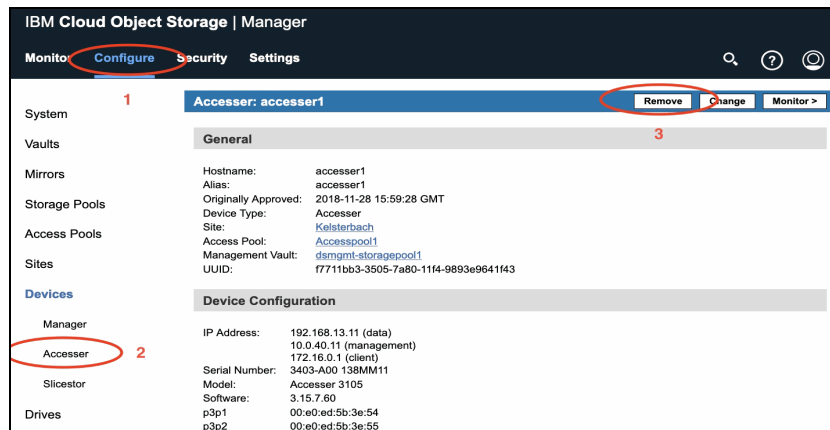


Figure 6-4 Removing an Accesser node

After this step is complete, the account password is required to accept the change. When the Accesser node is removed, it cannot be reaproved into the system before it is reimaged.

6.2.3 Automating performance with Docker and Kubernetes

The performance of the Accesser layer can be automated with the use of Docker containers and Kubernetes as a container management platform. Kubernetes offers the ability to autoscale containers based on their CPU usage. A basic tutorial is available at [this website](#).

Autoscaling for containers can be used to scale performance in the IBM Cloud Object Storage Accesser layer for uncertain workloads or variable workloads throughout the day. For example, if the workload drops overnight, static Accesser nodes can be more expensive than they need to be. But if the workload spikes, the system should perform to ensure a good client experience.

An Accesser image can be used to autoscale for performance. The Kubernetes autoscaling feature ensures that enough containers spin up when needed. When the CPU load is higher than the wanted value (for example, 70%), it automatically starts a new Accesser node.

For more information about how to use Accesser nodes as Docker containers, see 4.6, “Appliance Docker Containers” on page 83. For IBM Cloud Object Storage, the following steps are needed to use that new Accesser node:

- ▶ Approval of the Accesser node by the IBM Cloud Object Storage Manager.
- ▶ Adding the Accesser node to an access pool.

These steps can be automated through the Cloud Object Storage Management APIs. Another step is needed for removing Accesser nodes as the performance requirements drop. In this case, the Accesser node must be removed from the IBM Cloud Object Storage Manager. This process can also be done by a script that uses the IBM Cloud Object Storage Management APIs.

The Manager API reference is available at [IBM Documentation](#).

Tip: The autoscaling Accesser nodes can be added to hardware Accesser nodes. The hardware Accesser nodes then handle the base traffic, while the Docker Accesser nodes scale up to the needed performance.

6.3 Scaling for capacity

IBM Cloud Object Storage is built to scale without service disruptions, from a few terabytes up to an exabyte range and beyond.

Tip: Consider a storage expansion if the system reaches its limit in about 6 months.

For scaling capacity in IBM Cloud Object Storage, the following options are described next:

- ▶ Adding a device set to a storage pool
- ▶ Replacing a device set
- ▶ Removing a device set
- ▶ Adding a storage pool
- ▶ Planning for scalability

6.3.1 Adding a device set to a storage pool

IBM Cloud Object Storage uses the concept of device sets, as described in 1.1.1, “Key concepts and terminology” on page 2. For more information about configuring storage pools, see [IBM Documentation](#).

In general, the newly added device set must support the same Information Dispersal Algorithms (IDAs) as the existing device sets. For example, in a 12/6/8 IDA in Standard Dispersal (SD) Mode, it must be 12 Slicestor nodes. In a 18/9/11 IDA that is running on 3 Slicestor nodes in Concentrated Dispersal (CD) Mode, 3 or 18 Slicestor nodes can be added.

Slicestor nodes in a storage pool can have different configurations in terms of model, number of disks, and disk capacities, than an existing device set. Having different configurations for Slicestor nodes offers the flexibility to include newer Slicestor models or third-party Slicestor nodes to the same system. The existing device set stays in place and protects investments in the infrastructure.

Note: Within a device set, Slicestor nodes must be identical.

Storage pool expansion

As an example, a storage pool with a single device set that uses Gen1 Slicestor nodes can exist, as shown in Figure 6-5. A vault is deployed on that storage pool.

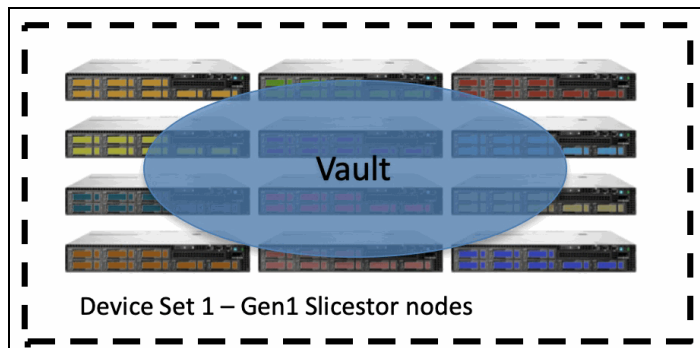


Figure 6-5 Single device set in a storage pool

Adding another device set with Gen2 Slicestor appliances to the same storage pool extends this storage pool, as shown in Figure 6-6.

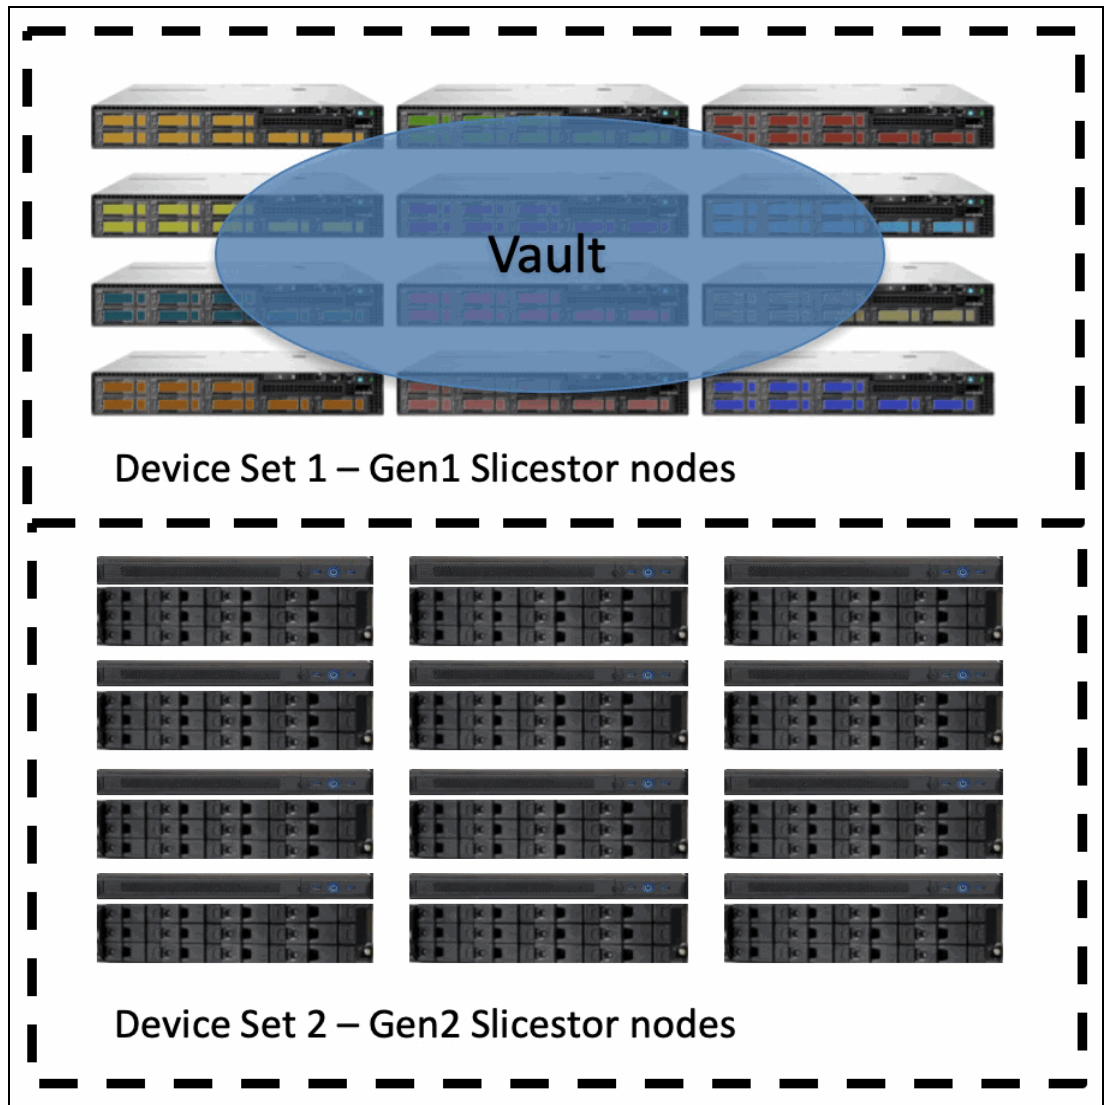


Figure 6-6 Expanding a storage pool

IBM Cloud Object Storage now automatically expands the vault across all device sets in the storage pool, as shown in Figure 6-7.




Figure 6-7 Expanded storage pool

The data on the Slicestor nodes is rebalanced between the device sets according to their capacity. Therefore, if different capacities are used in the device sets, the higher capacity Slicestor nodes receive more data than the low capacity nodes. With that configuration, an equal fill grade of each Slicestor node is achieved.

During this procedure, all read and write operations can continue on the vault.

Adding Slicestor nodes to a storage pool

Before this procedure is conducted, the hardware must be installed, the initial configuration must be completed as described in 5.5.3, “Configuring the Slicestor appliance” on page 104, and the extra Slicestor nodes must be approved in the IBM Cloud Object Storage System.

Complete the following steps to add the Slicestor nodes to a storage pool:

1. Select **Configure** → **Storage Pools** → **select storage pool where Slicestor nodes are to be added** → **Change Sets and Devices**, as shown in Figure 6-8.

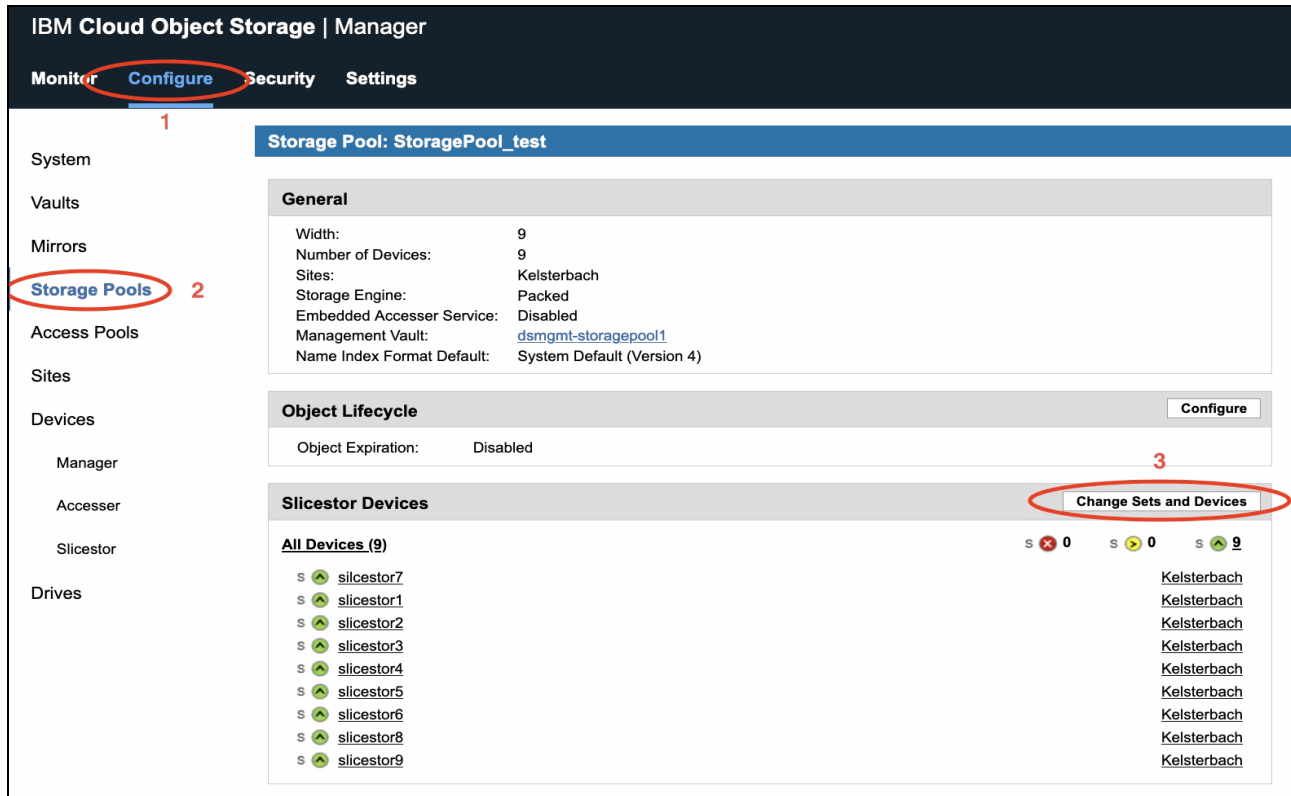


Figure 6-8 Changing Storage Pool device sets

2. Select **Configure Storage Pool Expansion**.

3. In the next window, select the width of the storage pool and select the devices that are to be added (see Figure 6-9). The width must be the same as the device set.

Note: The expansion of a CD Mode system can be done by adding the same physical width, or the same number of physical SliceStor nodes as the IDA width.

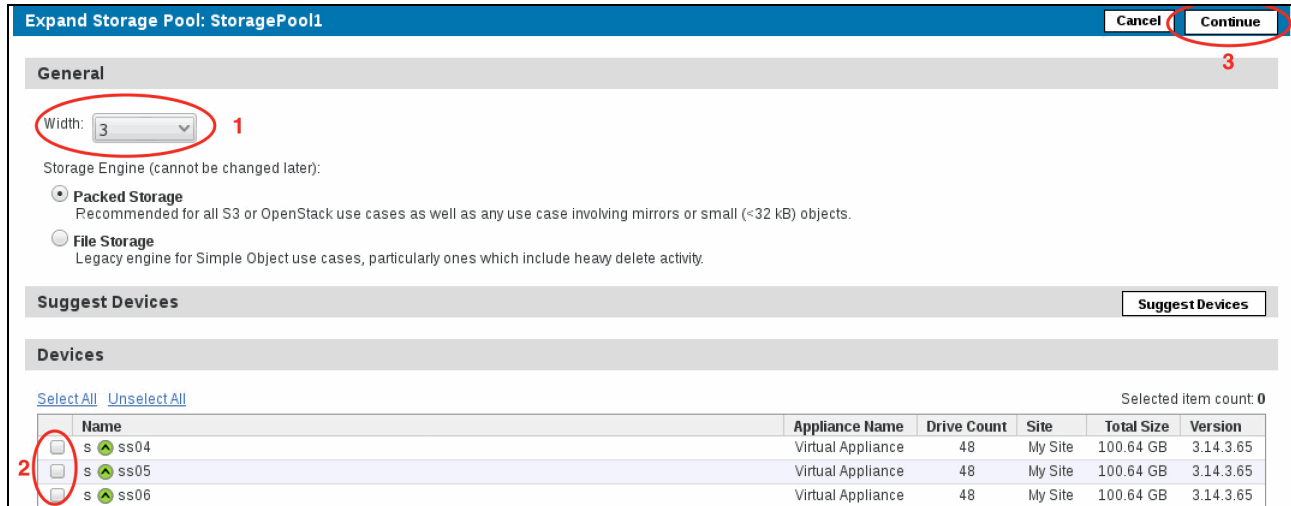


Figure 6-9 Adding SliceStor nodes to a storage pool

Note: Since Version 3.15.6, it is possible to add multiple device sets at the same time. To add multiple device sets, repeat the steps above until all new device sets are added. After that, proceed with the reallocation. IBM Cloud Object Storage will automatically rebalance across all new device sets.

4. An overview of the changes is shown. Click **Start Expansion** to start the expansion of the storage pool and rebalancing of capacities.

During the data reallocation, the progress can be tracked by selecting **Show Progress** in the storage pool overview, as shown in Figure 6-10 on page 177.

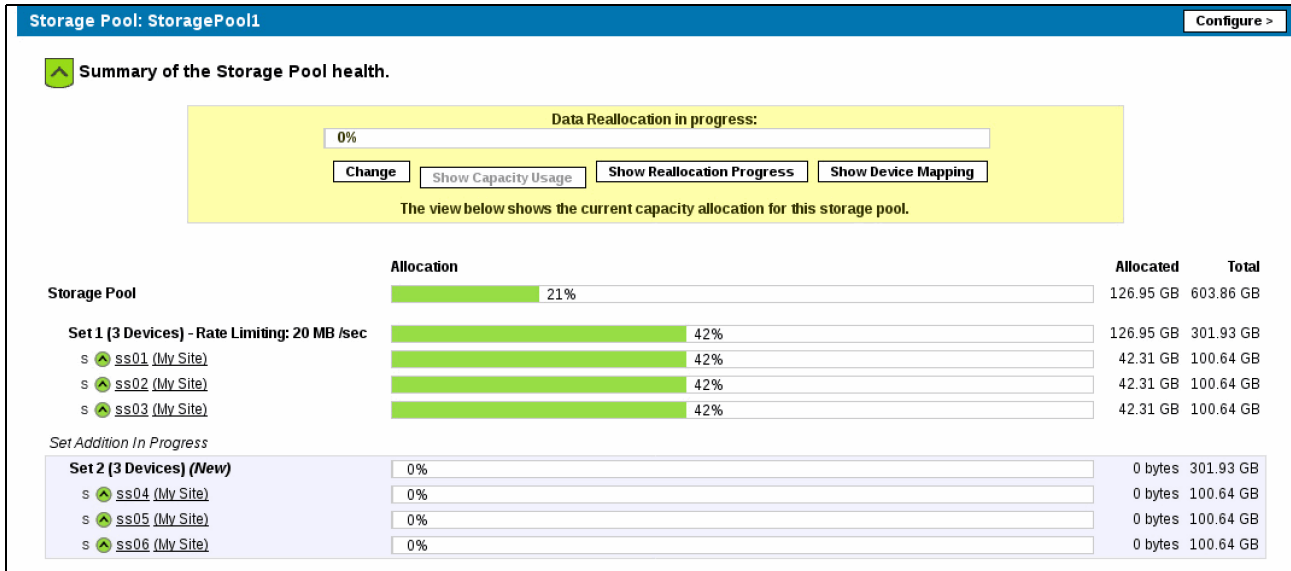


Figure 6-10 Reallocation of capacities

This process can take some time and can affect the overall performance and network traffic. The effect on performance can be controlled by setting a limit on the data transfer per device. The standard rate is set to 20 MBps per device. This limit always can be changed; for example, to speed up the process in “quiet” times.

- After the reallocation is done, all Slicestor devices in that storage pool have the same usage and the message that is shown in Figure 6-11 appears.




Figure 6-11 Finished reallocation

From this point on, all workload is distributed across all Slicestor nodes in the storage pool evenly.

Implications on availability

With more than one device set in a storage pool, there are other considerations with regards to availability of the IBM Cloud Object Storage System.

Slicestor nodes across the device sets are logically combined to represent one entity regarding the IDA. The first Slicestor node in device set 1 is bundled with the first Slicestor node in device set 2, the second Slicestor node in a device set with the second of the new device set, and so on.

Each device set features the same redundancies that are defined by the IDA so that the availability does not decrease while an IBM Cloud Object Storage System is scaled. In comparison, a system in which the ratio of redundancies to the rest of the data is reduced has less availability after scaling.

Implications on performance

While planning to expand the IBM Cloud Object Storage System storage pool with different hardware specifications, it is a best practice to review it with an IBM Cloud Object Storage Specialist or Support for guidance to avoid a potential unexpected performance degradation.

The key specifications with performance implications include number of data drives, capacity of data drives, number of CPUs (both physical and logical cores due to multi-threading), and amount of available memory.

Example: A customer is considering scaling IBM Cloud Object Storage System upwards by increasing overall capacity substantially by taking advantage of the latest large drive-count Slicestor appliances.

As the number of disks and disk capacity per Slicestor appliance scales up, the number of logical CPUs and amount of memory should scale up proportionally to ensure equivalent performance. Expanding with new Storage Pool Sets that have a lower CPU-to-disk ratio might result in a performance degradation for certain workloads.

6.3.2 Replacing a device set

IBM Cloud Object Storage can be scaled by replacing a device set with higher capacity Slicestor nodes; for example, during hardware refresh. The new device set can have different hardware in terms of model, number of disks, and disk sizes than the existing device set. In this way, it is possible to scale capacity by using the same number of Slicestor nodes as before.

For replacing a device set, the width must stay the same. In SD Mode, the same number of Slicestor nodes must be added. In CD Mode, device sets can be replaced by using the same physical width or the same number of physical Slicestor nodes as the IDA width. When replaced with an IDA width device set, IBM Cloud Object Storage considers this storage pool SD Mode after the old device set is removed. New vaults can be created with SD Mode IDAs.

IBM Cloud Object Storage copies all the data from the existing device set to the new device set so that read/write operations continue without disruption. After the data is copied to the new device set, the old device set is removed.

Procedure

The new Slicestor nodes are initialized and added to the system as described in 5.5, “Step 3: Configuring the appliance” on page 101. Complete the following steps:

1. Select **Configure** → **Storage Pools** → **Select storage pool where the device set is to be replaced** → **Change Sets and Devices**, as shown in Figure 6-12.

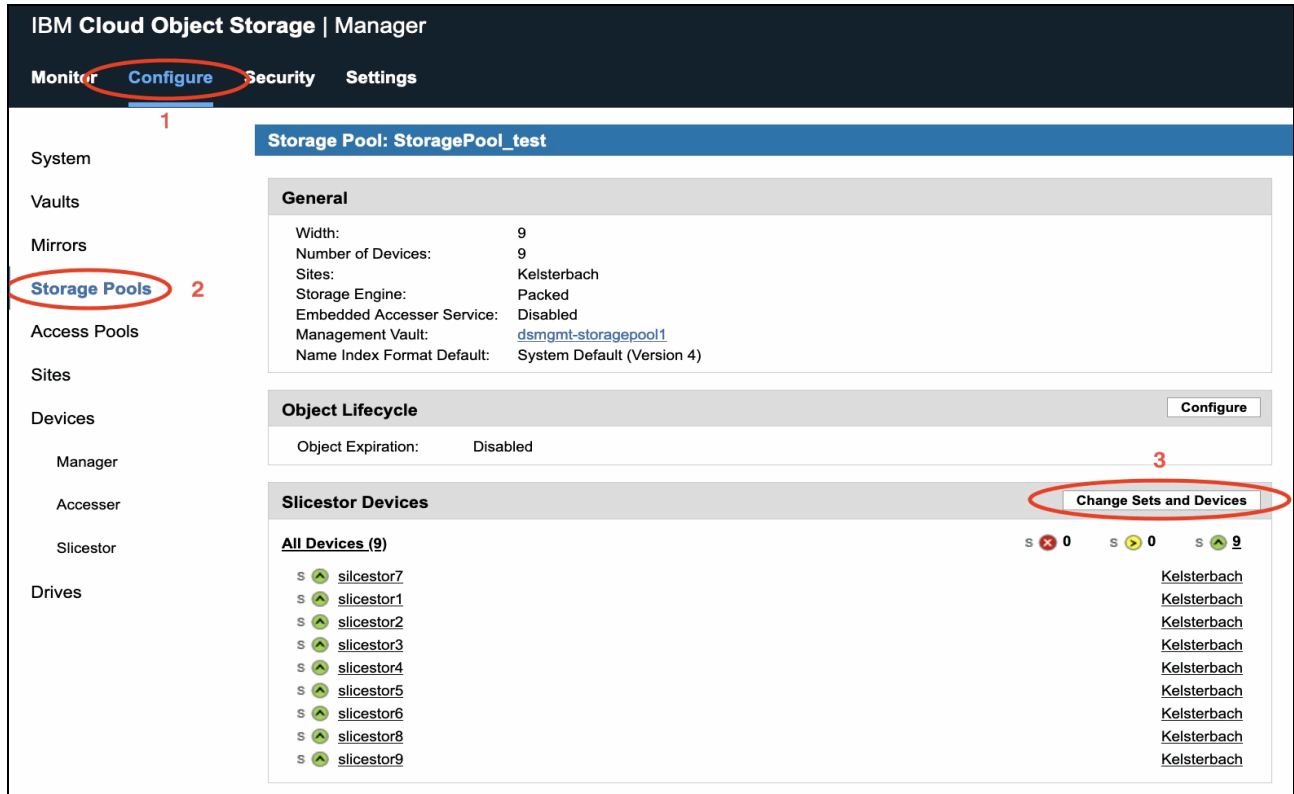


Figure 6-12 Changing storage pools

- In the next step, select **Configure Set Replacement** and in the next window, select a device set that is to be replaced, and the new Slicestor nodes that form the new device set, as shown in Figure 6-13.

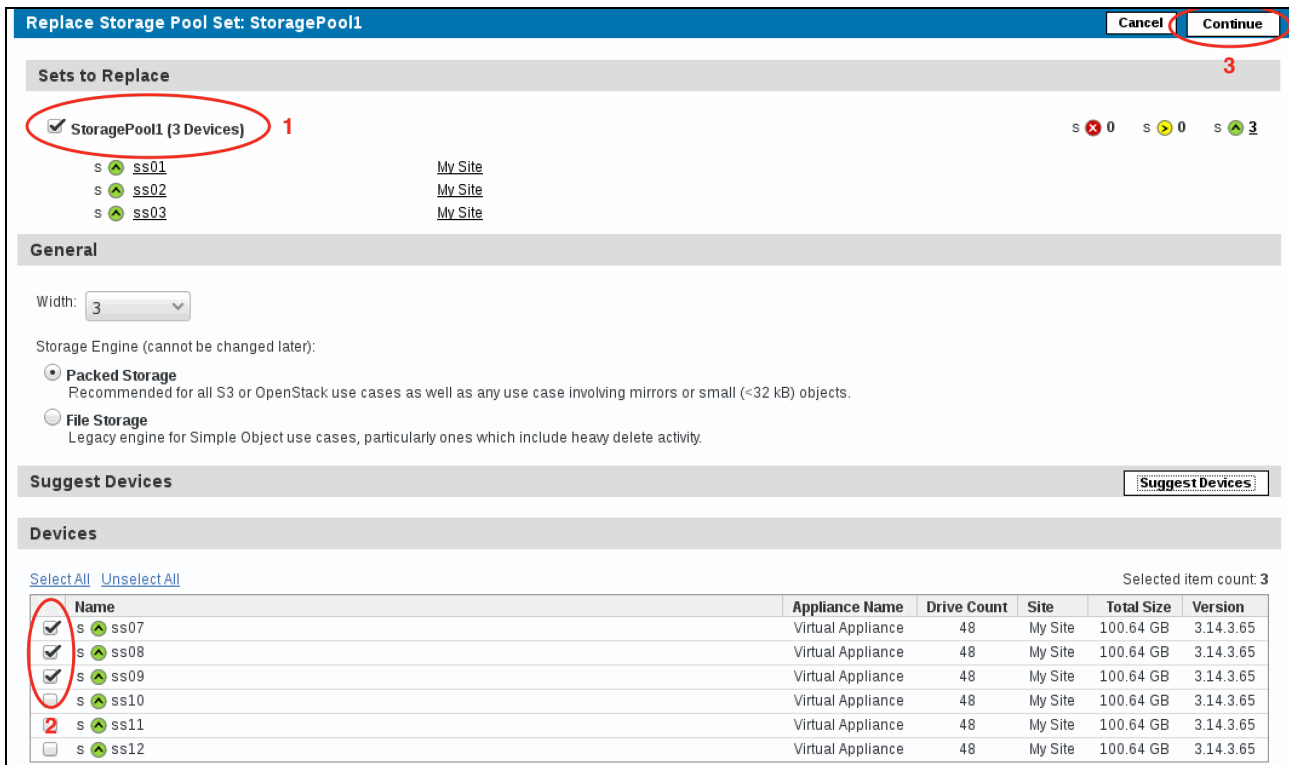


Figure 6-13 Replacing a device set

- An overview of the changes is shown. Click **Start Set Replacement** start the migration of the device set.
- IBM Cloud Object Storage shows the progress of the data reallocation on the device set when you select **Show Progress** in the Storage pool overview.

In larger environments, this process can take some time and affect the overall performance and network traffic. The effect on performance is controlled by setting a limit on the data transfer per device. This limit can be changed; for example, to speed up the process during quiet times.

The old device set is removed from the storage pool after all the data is moved. Slicestor nodes that were removed from a storage pool must be reimaged before they can be reused in an IBM Cloud Object Storage System.

6.3.3 Removing a device set

Removing a device set can be used for scaling down capacity or to reallocate Slicestor nodes to another storage pool. Before the device set is removed, IBM Cloud Object Storage moves the data to remaining device sets in a way that read and write operations continue without disruption.

When all data is moved from the selected device set, these Slicestor nodes are removed from the storage pool. The removed Slicestor nodes are still included in the monitoring section, but must be reimaged before they can be reused in an IBM Cloud Object Storage System.

Note: Removing a device set is possible only if the remaining device sets feature enough capacity to hold all the data.

Procedure

To remove a device set, complete the following steps:

1. Select **Configure** → **Storage Pools**.
2. Select the storage pool where the device set is to be removed.
3. Select **Change Sets and Devices**, as shown in Figure 6-14.

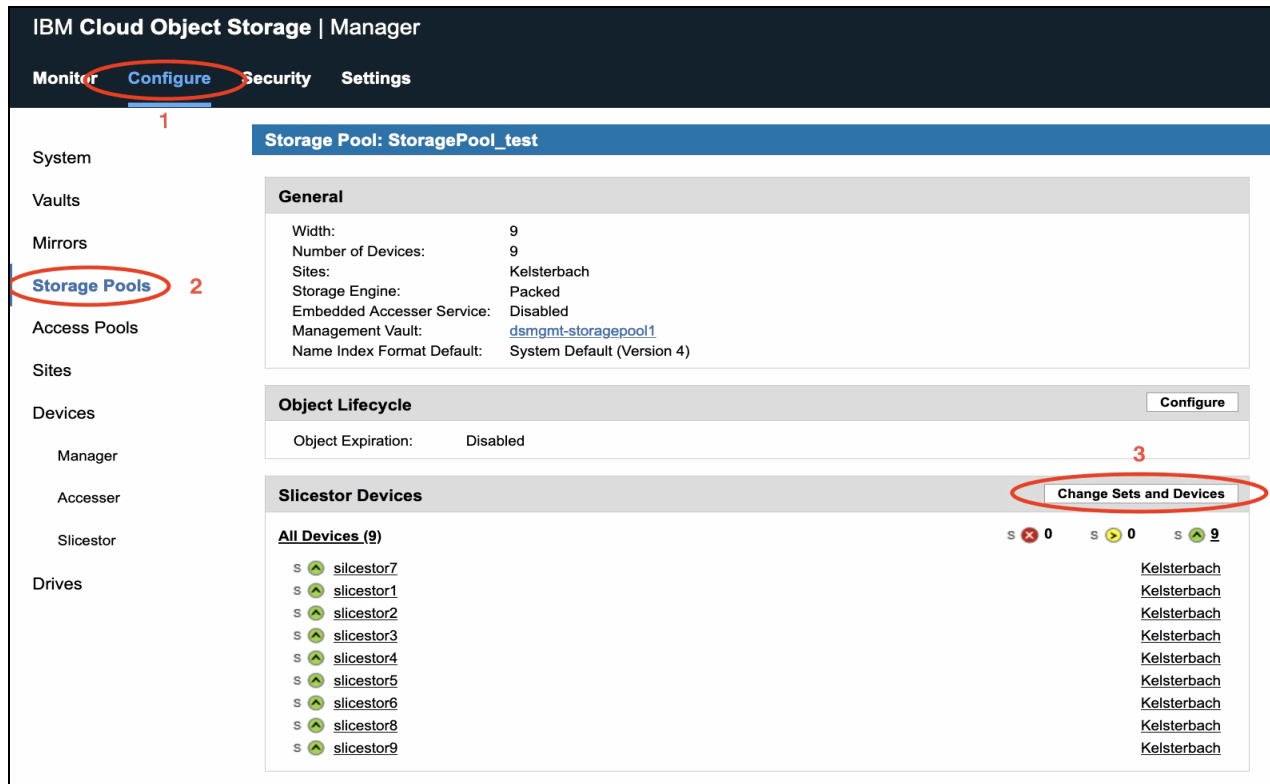


Figure 6-14 Changing storage pools

4. In the next window, select **Configure Set Removal**.

- Select the device set that is to be removed from the storage pool, as shown in Figure 6-15. Click **Continue**.

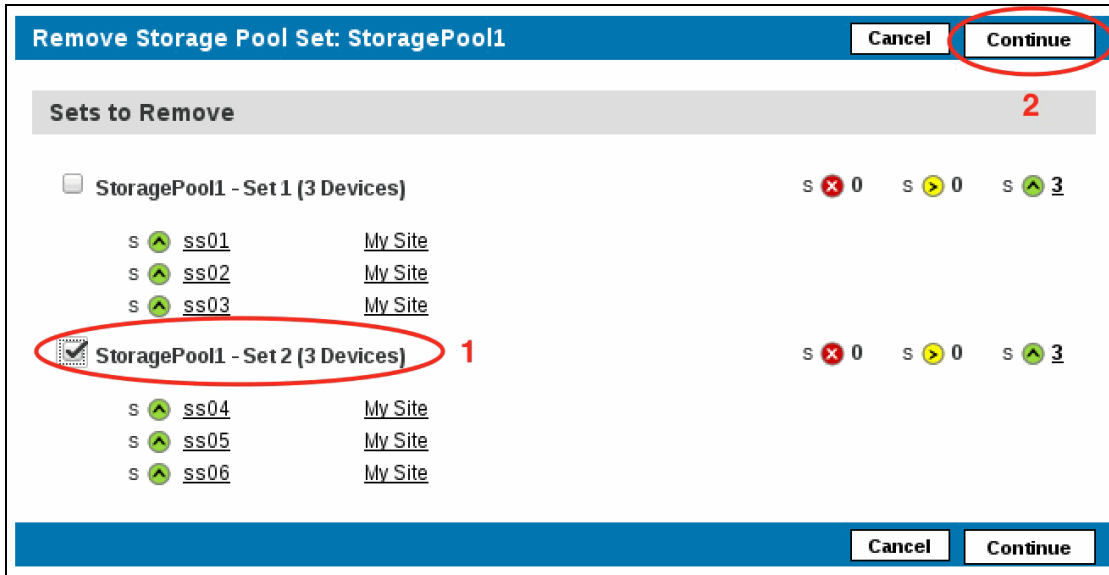


Figure 6-15 Selecting device set for removal

- An overview of the changes is shown. Click **Start Set Removal** to start the data evacuation from the device set.
- During data evacuation, the progress can be tracked by selecting **Show Progress** in the storage pool overview, as shown in Figure 6-16.

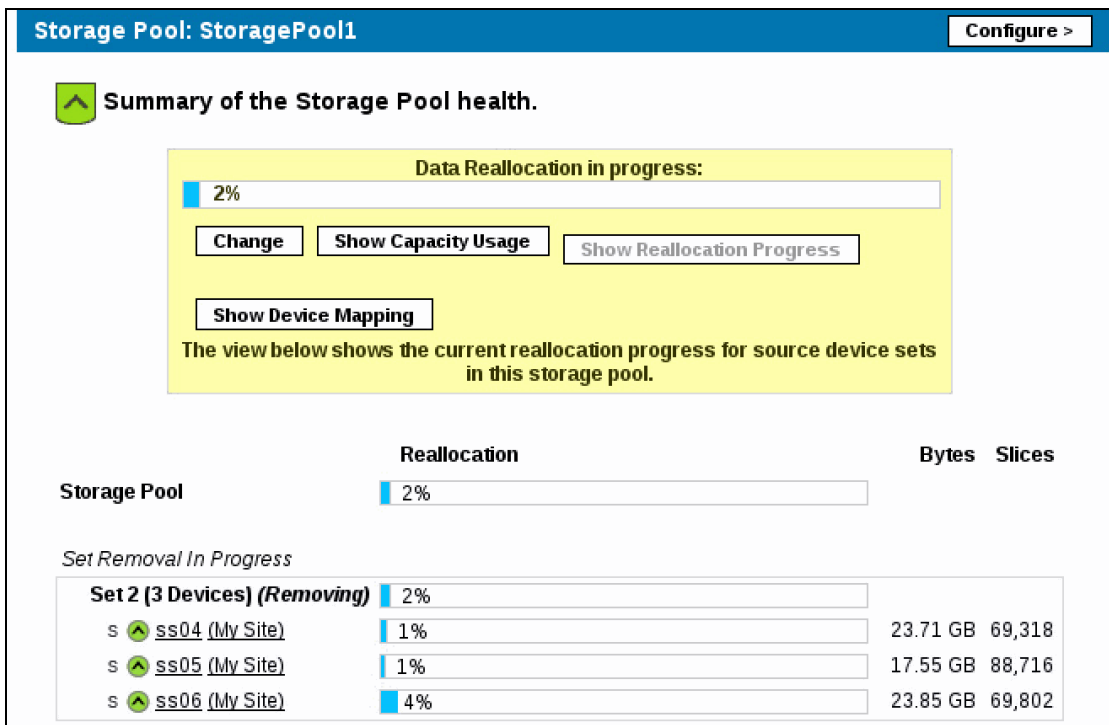


Figure 6-16 Progression of data reallocation

In larger environments, this process can take some time and affect the overall performance and network traffic. The effect on performance is controlled by setting a limit on the data transfer per device. This limit always can be changed; for example, to speed up the process during quiet times.

Removing Slicestor nodes

After all the data is moved off the device set, Slicestor nodes can be removed from the IBM Cloud Object Storage System by selecting **Configure** → **Slicestor**. Select the Slicestor node that is to be removed and select **Remove**, as shown in Figure 6-17.

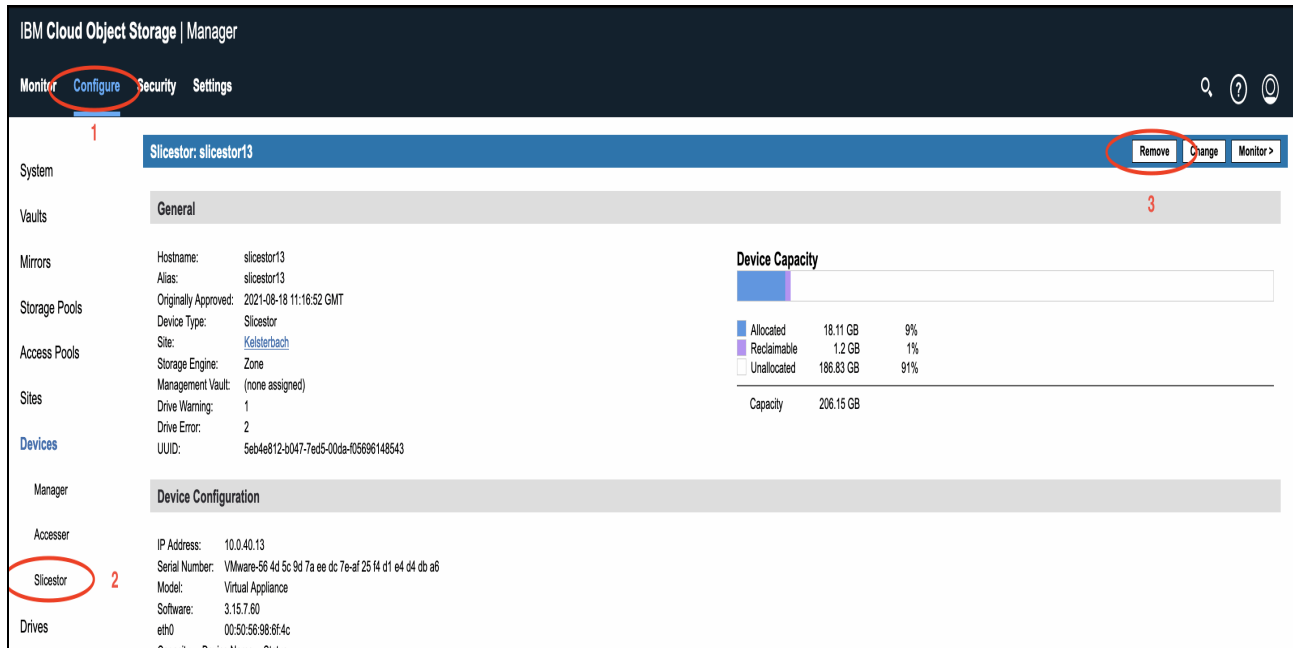


Figure 6-17 Removing a Slicestor node

Next, you are prompted for the account password to accept the removal of a Slicestor node.

6.3.4 Adding a storage pool

A storage pool can be added to an IBM Cloud Object Storage System to increase capacity of the overall system. This addition leads to a separate namespace; for example, for test namespaces. Adding a storage pool does *not* contribute to the storage for vaults because a vault is on one storage pool.

The new storage pool is independent of any storage pool IDAs locations and dispersal modes. For example, one storage pool can be run in CD Mode with only three Slicestor nodes on one site for software development, while another storage pool in that same IBM Cloud Object Storage System has 36 Slicestor nodes spread across three sites that are running in SD Mode for production use.

The procedure for adding a storage pool is similar to the first-time setup process. For more information, see 5.6.5, “Creating a storage pool” on page 114.

6.3.5 Planning for scalability

At the initial sizing of an IBM Cloud Object Storage System, the future data growth must be considered. If the capacity growth requirements are known for at least the first few years of production, the IDA must be chosen in a way that this expansion is considered because it stays the same during expansions. It might be higher cost per GB at the beginning, in favor of lower cost per GB overall. It is important to be aware of the *expansion factor*. The lower the expansion factor, the more storage-efficient the system.

For more information about capacity planning, see 2.1, “Planning for capacity” on page 24.

Example: A customer is considering starting with a small system in CD Mode. The IDA of 18/9/11 has an expansion factor of 2. If the customer wants 1.5 PB usable capacity, 3 PB of raw storage is needed. The customer can achieve this storage by using six Slicestor 53 nodes with 12 TB drives.

The customer’s requirement can also be met by using 15 Slicestor 53 nodes with 4 TB drives and an IDA of 15/8/10, which features an expansion factor of 1.875. It makes sense to choose this IDA if the customer wants to scale out soon and can take advantage of the lower expansion factor.



IBM Cloud Object Storage System File Access

This chapter provides an overview of IBM Cloud Object Storage File Access System.

This chapter includes the following topics:

- ▶ 7.1, “Introduction” on page 186
- ▶ 7.2, “Features” on page 186
- ▶ 7.3, “IBM Cloud Object Storage example use case” on page 187
- ▶ 7.4, “IBM Cloud Object Storage File Access deployment architecture” on page 188
- ▶ 7.5, “Conclusion” on page 189

7.1 Introduction

IBM Cloud Object Storage File Access is a software-defined offering that provides SMB and Network File System (NFS) protocol interfaces to applications to store, archive, and retrieve infrequently accessed files on IBM Cloud Object Storage.

The following lists the main IBM Cloud Object Storage File Access features:

- ▶ Security

All wide area network (WAN) transfers use Transport Level Security (TLS) protocol over the WAN, preventing unauthorized interception of data transfers.

- ▶ Active Directory Integration

IBM Cloud Object Storage File Access provides role-based access control, by using Active Directory.

- ▶ Infinite file capacity

Intelligent caching technology delivers unlimited file access to office users, with visibility to all organizational files centralized in the cloud. Files are dynamically cached from IBM Cloud Object Storage File Access Portal to the IBM Cloud Object Storage File Access Gateway.

The efficient caching technology includes incremental updates, data compression, block level deduplication, and simultaneous synchronization.

- ▶ Centralized management and monitoring

IBM Cloud Object Storage File Access Gateways can be managed from a single pane of glass, including usage monitoring, upgrades, and remote troubleshooting. SNMP monitoring is also available.

- ▶ File access

Files can be accessed by using SMB 2.x/3.x (Windows File Sharing) or NFS.

7.2 Features

The IBM Cloud Object Storage File Access Solution includes the following components:

- ▶ IBM Cloud Object Storage File Access Portal
- ▶ IBM Cloud Object Storage File Access Gateway

The IBM Cloud Object Storage File Access Portal is the management component of the offering, which enables the creation, delivery and management of the services mentioned below. The IBM Cloud Object Storage File Access Portal interacts with the IBM Cloud Object Storage File Access Gateways and efficiently handles file data exchange between these applications and users and the on-premises or public IBM Cloud Object Storage instance side. A centralized management console makes it possible to effectively manage a large number of connected IBM Cloud Object Storage File Access Gateways.

The IBM Cloud Object Storage File Access Gateway is the component that the application and other data sources are connected to, and allows LAN speed writes through SMB and NFS protocols, and is in charge of onboarding the data to IBM Cloud Object Storage instantly and seamlessly.

The IBM Cloud Object Storage File Access Gateway works in caching mode, which means that it has a dedicated local disk space to allow local LAN speed ingestion. The main storage is on the IBM Cloud Object Storage File Access Portal in the cloud with stubs that are saved on the IBM Cloud Object Storage File Access Gateway. A stub is a file with a tiny footprint that contains the metadata about the file, such as the file name, size, and modification date. Only the folder and file metadata and not the actual file content is saved locally. This configuration results in the cost of storage being lower.

Also, systems with many file changes, where only some of the files are required locally, don't over use bandwidth between the cloud and IBM Cloud Object Storage File Access Gateway. Only the required files are passed across the wire. When a user accesses a file stub, the file is opened without delay, where possible by streaming the file content from the cloud. After the download completes, the file is unstubbed. Any changes to the file are synced back to the IBM Cloud Object Storage File Access Portal. Folders that are always required can be pinned, in which case the files in the folders, and not the stubs, are stored on the IBM Cloud Object Storage File Access Gateway.

IBM Cloud Object Storage File Access Gateways are virtual appliances, which can be installed on any customer provided ESXi, Hyper-V, or KVM/OpenStack environment.

7.3 IBM Cloud Object Storage example use case

The IBM Cloud Object Storage File Access deployment configuration should be able to handle the load generated by a specific user environment. This section describes the IBM Cloud Object Storage File Access Archiving use case.

The focus of this use case is the ability to actively archive data from the edge/datacenter/any location that is sourcing data to be archived, into cloud storage to lower cost of archival data.

The following examples can demonstrate the capabilities of the use case:

- ▶ Consolidate archival data (for example, video, images, logs) from 10s to 1000s of remote edge locations.
- ▶ Store application backups in IBM Cloud Object Storage for retention (database backups, Application logs, and so on).
- ▶ Provide bursts of storage capacity to archival data.

The Archiving use case uses the direct to storage write capabilities to fully use the available network capacity to the archive tier storage target and gain fast and reliable ingestion of data to be archived.

7.4 IBM Cloud Object Storage File Access deployment architecture

An IBM Cloud Object Storage File Access installation comprises a cluster of one or more servers (VMs). The relevant roles are divided into IBM Cloud Object Storage File Access Portal VMs and IBM Cloud Object Storage File Access Gateway VMs.

7.4.1 IBM Cloud Object Storage File Access Portal Application Server

The application server (VM) is in charge of all the communication between the moving parts in the system. It receives all requests from IBM Cloud Object Storage File Access Gateways and administrators and communicates them to the database node. The application VM is stateless and is in active-active mode with at least one more application VM as part of a production implementation. This situation means that every production implementation has at least two application VMs.

7.4.2 Database Server

A main database server (VM) holds administrative information that is related to the various cloud storage services, and metadata for files stored on the IBM Cloud Object Storage storage infrastructure. Only one VM can host the main database node. The VM that hosts the main database node is called the master server and is fully replicated to a replica database for full HA and DR (HADR) capabilities. This situation means that every production implementation has two database VMs.

7.4.3 IBM Cloud Object Storage File Access Gateway

The IBM Cloud Object Storage File Access Gateway is the component that makes the SMB and NFS shares that the applications and users write to available. This configuration means that the gateway must be always on, and includes another DR IBM Cloud Object Storage File Access Gateway in a full production environment. Each ingestion site has at least two FA Gateways.

While the application is busy navigating the relevant requests to the other parts of the system, and the database is in charge of the application database and metadata of all the files that are archived, the IBM Cloud Object Storage File Access Gateway uses the Direct to Storage feature, allowing direct writes to the IBM Cloud Object Storage bucket. This feature dramatically increases data ingestion based on the available network between the IBM Cloud Object Storage File Access Gateway and the storage, with speeds of around 350 MBps write to IBM Cloud Object Storage for archiving workflows.

Figure 7-1 on page 189 shows a high-level architecture of this deployment use case.

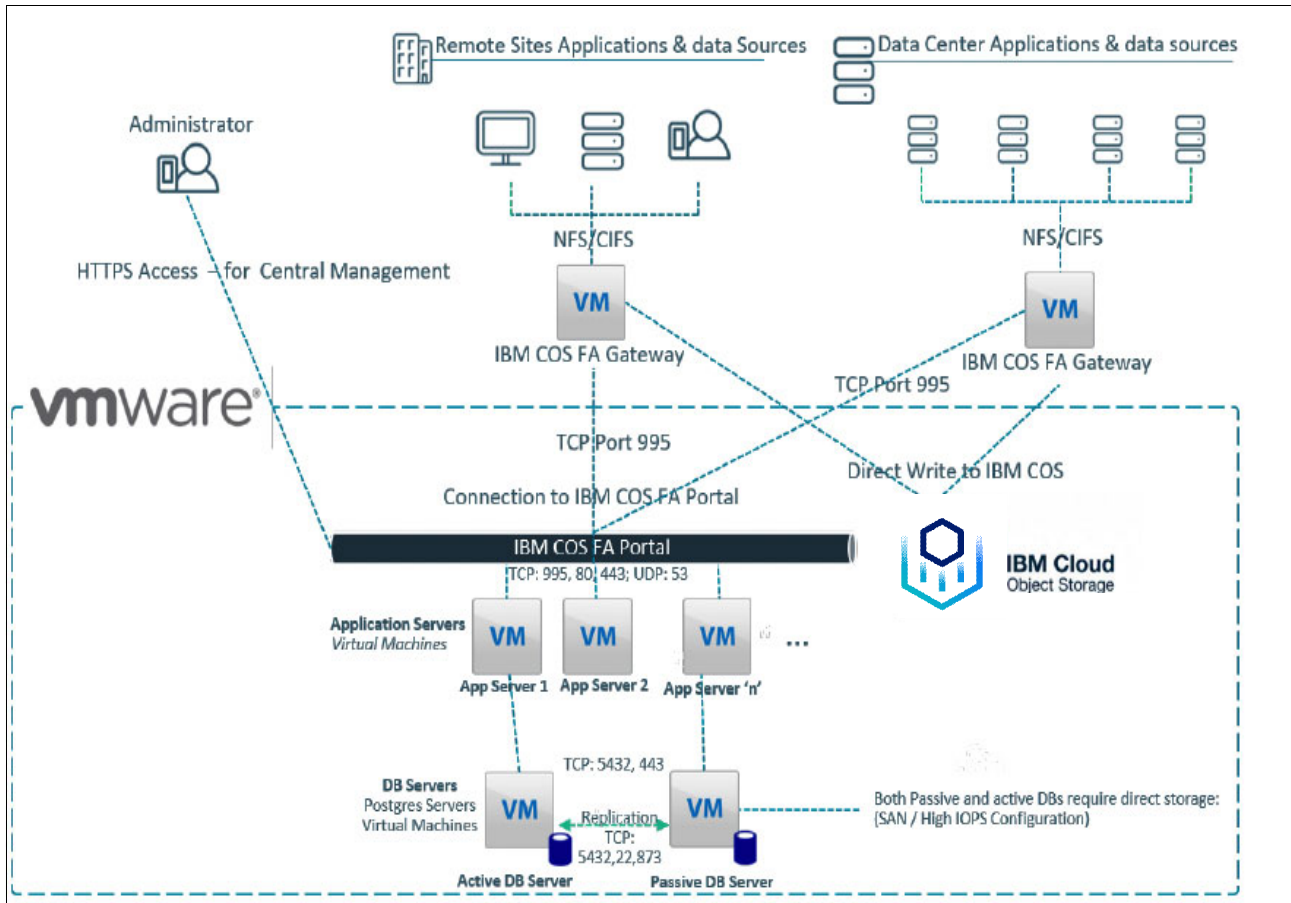


Figure 7-1 IBM Cloud Object Storage File Access deployment architecture

7.5 Conclusion

The IBM Cloud Object Storage File Access Portal was designed to scale from tens to hundreds and thousands of connected IBM Cloud Object Storage File Access Gateways and to support an easy-to-scale file system with PBs of data and more.

The IBM Cloud Object Storage File Access Portal can support both scale-up and scale-out deployment schemes: administrators may deploy the IBM Cloud Object Storage File Access Portal software on increasingly powerful compute platforms, thus scaling the deployment up. Alternatively, they can distribute the IBM Cloud Object Storage File Access Portal processes on multiple concurrent compute platforms, thus scaling out the deployment.

In addition, the file system is fully scalable by enlarging the database to accommodate data capacity growth.

Abbreviations and acronyms

AONT	All-or-Nothing Transform	PSS	Packed Slice Storage
BCDR	business continuity and disaster recovery	RDIMM	registered dual inline memory module
CAGR	compound annual growth rate	RT	read threshold
CD	Concentrated Dispersal	RU	rack unit
CFTC	Commodity Futures Trading Commission	SAS	serial-attached SCSI
CoD	Capacity on Demand	SD	Standard Dispersal
DNS	Domain Name System	SDS	software-defined storage
DSR	Direct Server Return	SEC	Securities and Exchange Commission
ECuRep	Enhanced Customer Data Repository	SFTP	Secure File Transfer Protocol
FEC	forward error correction	SLB	server load balancing
FINRA	Financial Industry Regulatory Authority	SoE	Systems of Engagement
FQDN	fully qualified domain name	SoR	Systems of Record
GSLB	global server load balancing	TCO	total cost of ownership
HA	high availability	TLS	Transport Level Security
HPC	high-performance computing	UPS	uninterruptible power supply
HTTPS	Hypertext Transfer Protocol Secure	WAN	wide area network
IBM	International Business Machines Corporation	WORM	Write Once Read Many
IDA	Information Dispersal Algorithm	WT	write threshold
IPMI	Intelligent Platform Management Interface	ZSS	Zone Slice Storage
ITIL	IT Infrastructure Library		
IoT	Internet of Things		
JWT	JSON Web Tokens		
LACP	Link Aggregation Control Protocol		
LB	load balancer		
MFA	Multi-factor authentication		
MTTDL	Mean Time To Data Loss		
MTU	maximum transmission unit		
NAS	network-attached storage		
NFS	Network File System		
NTP	Network Time Protocol		
OG	Object Generator		
OIDC	OpenID Connect		
OPS	operations per second		
OVA	Open Virtual Appliance		
PCM	power and cooling modules		
PMP	Project Management Professional		

Related publications

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

IBM Redbooks

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

- ▶ *Cloud Object Storage as a Service: IBM Cloud Object Storage from Theory to Practice - For developers, IT architects and IT specialists*, SG24-8385.
- ▶ *IBM Cloud Object Storage System Product Guide*, SG24-8439

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

ibm.com/redbooks

Online resources

The following website also is relevant as a further information source:

- ▶ IBM Cloud Object Storage System documentation:
<https://www.ibm.com/docs/en/coss/3.17.2>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Redbooks

IBM Cloud Object Storage System Product Guide

(0.2"spine)
0.17"->0.473"
90->249 pages



SG24-8439-03

ISBN 0738460133

Printed in U.S.A.

Get connected

