



Quest® QoreStor™ 7.5.0

User Guide



© 2025 Quest Software Inc. ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software Inc.

The information in this document is provided in connection with Quest Software products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest Software products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST SOFTWARE ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST SOFTWARE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest Software makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest Software does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

Quest Software Inc.

Attn: LEGAL Dept

4 Polaris Way

Aliso Viejo, CA 92656

Refer to our Web site (<https://www.quest.com>) for regional and international office information.

Patents

Quest Software is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at <https://www.quest.com/legal>.

Trademarks

Quest, the Quest logo, and Join the Innovation are trademarks and registered trademarks of Quest Software Inc. For a complete list of Quest marks, visit <https://www.quest.com/legal/trademark-information.aspx>. All other trademarks and registered trademarks are property of their respective owners.

Legend

 **CAUTION: A CAUTION icon indicates potential damage to hardware or loss of data if instructions are not followed.**

 **IMPORTANT, NOTE, TIP, MOBILE, or VIDEO:** An information icon indicates supporting information.

QoreStor User Guide
Updated - January 2025
Version - 7.5.0

Contents

| | |
|--|-----------|
| Introducing Quest® QoreStor™ | 11 |
| Understanding the QoreStor documentation | 11 |
| Other information you may need | 12 |
| Information on compatible products | 12 |
| Source code availability | 13 |
| QoreStor data storage concepts | 14 |
| Data deduplication and compression | 14 |
| Encryption at rest | 14 |
| Streams and connections | 14 |
| Recycle Bin | 15 |
| Quotas | 15 |
| Secure connect | 15 |
| MultiConnect | 16 |
| Replication | 16 |
| Reverse replication | 17 |
| Reverse replication: alternate method | 18 |
| Performance Tier | 18 |
| Cloud Tier replication | 18 |
| Archive tier | 19 |
| Object interface (S3-compatible) | 20 |
| Object Direct Storage | 20 |
| Cloud Reader mode | 20 |
| Continuous data protection | 20 |
| Synthetic full backups | 21 |
| Instant Restore | 21 |
| Compatibility | 21 |
| Disaster Recovery from the cloud | 21 |
| Cloud tier disaster recovery | 21 |
| Archive tier disaster recovery | 22 |
| MSP Mode | 22 |
| Partial File System Checker (OFSCK) | 22 |
| Supported file system protocols | 23 |
| CIFS | 23 |
| CIFS ACL support | 23 |
| Making CIFS backup immutable for Veeam | 24 |
| NFS | 24 |
| Rapid NFS and Rapid CIFS benefits | 25 |
| Rapid NFS and Rapid CIFS | 25 |
| RDA with Quest's NetVault and vRanger | 25 |
| RDA with OST for QoreStor | 26 |

| | |
|--|-----------|
| Software components and operational guidelines | 27 |
| NDMP | 28 |
| iSCSI | 28 |
| Hardened Repository with Veeam | 28 |
| Anomaly detection | 29 |
| Anomaly detection categories | 29 |
| System-level | 29 |
| Container-level | 31 |
| Storage-group level | 31 |
| Retraining | 32 |
| Alerts | 32 |
| Reports | 32 |
| QoreStor processes | 32 |
| Accessing QoreStor | 34 |
| Networking prerequisites for QoreStor | 34 |
| Configurations required when using proxy servers | 34 |
| Logging onto the system GUI for the first time | 35 |
| Using the QoreStor UI | 35 |
| QoreStor UI Overview | 35 |
| Using the QoreStor command line interface | 37 |
| Using the QoreStor REST API | 37 |
| API Overview | 37 |
| Authentication | 38 |
| API Documentation | 38 |
| Understanding RDA immutability | 39 |
| Understanding EDM immutability | 39 |
| Configuring QoreStor settings | 40 |
| Licensing QoreStor | 40 |
| Evaluating QoreStor | 40 |
| Viewing your license configuration | 41 |
| Installing a license | 41 |
| NetVault Plus Data Copy licenses | 41 |
| Preparing QoreStor for a NetVault Plus Data Copy License | 41 |
| Installing a NetVault Plus Data Copy License | 42 |
| Configuring SAML | 42 |
| Accessing the SAML configuration settings in QoreStor | 43 |
| Registering an SSO user with the QoreStor Server | 43 |
| Logging in to the QoreStor UI using SAML | 44 |
| Configuring an SSL certificate for your QoreStor system | 44 |
| Installing an SSL certificate | 44 |
| Configuring Active Directory settings | 45 |

| | |
|--|----|
| Adding a login group to an ADS domain | 45 |
| Removing a login group | 46 |
| Securing QoreStor server root login | 46 |
| Enabling FIPS 140-2 support | 46 |
| Understanding system operation scheduling | 47 |
| Configuring cleaner schedules | 48 |
| Viewing cleaner status | 49 |
| Viewing cleaner statistics | 49 |
| Configuring Secure Connect | 50 |
| Enabling Secure Connect for OST and RDA plug-ins prior to 4.1.0.265 | 50 |
| Configuring Secure Connect properties | 51 |
| Managing Secure Connect with OST or RDA plug-in 4.1.0.265 or later | 52 |
| Checking Secure Connect status | 53 |
| Disabling Secure Connect | 53 |
| Enabling Secure Connect | 54 |
| Adding certificates for Secure Connect | 54 |
| Adding a Secure Connect certificate - Windows Client | 54 |
| Adding a Secure Connect certificate - Linux Client and QoreStor server | 55 |
| Enabling MultiConnect | 56 |
| Bandwidth throttling | 57 |
| Adding a throttling schedule from UI | 58 |
| Viewing/editing a throttling schedule | 58 |
| Removing/deleting a throttling schedule | 59 |
| Configuring and using Rapid NFS and Rapid CIFS | 59 |
| Rapid NFS and Rapid CIFS benefits | 59 |
| Best practices: Rapid NFS | 60 |
| Best practices: Rapid CIFS | 61 |
| Installing the Rapid NFS plug-in | 62 |
| Downloading the Rapid NFS Plug-in | 62 |
| Installing Basic Calculator | 63 |
| Enabling Secure Connect | 63 |
| Installing the plug-in | 63 |
| Installing the Rapid CIFS plug-in | 64 |
| Downloading the Rapid CIFS Plug-in | 64 |
| Enabling Secure Connect | 64 |
| Installing the plug-in | 64 |
| Uninstalling the Rapid NFS plug-in | 65 |
| Uninstalling the Rapid CIFS plug-in | 66 |
| Configuring and using VTL | 66 |
| Understanding VTL | 66 |
| Terminology | 66 |
| Supported virtual tape library access protocols | 67 |
| NDMP | 67 |
| iSCSI | 67 |

| | |
|---|-----------|
| VTL and QoreStor specifications | 67 |
| Guidelines for configuring VTL | 70 |
| Plan your Environment | 70 |
| Create Containers of Type VTL | 70 |
| Authentication/User Management Considerations | 71 |
| Verify the Tape Library Creation | 71 |
| Configure the Library in the DMA | 71 |
| Configuring and Using Encryption at Rest | 71 |
| Understanding Encryption at Rest | 72 |
| Encryption at Rest Terminology | 72 |
| Encryption at Rest and QoreStor Considerations | 72 |
| Understanding the encryption process | 73 |
| Configuring and using the Recycle Bin | 74 |
| Understanding Recycle Bin | 74 |
| Compatibility | 74 |
| Terminology | 74 |
| Guidelines for configuring Recycle Bin | 75 |
| Configuring Recycle Bin | 75 |
| Purging Recycle Bin files | 75 |
| Restoring files from Recycle Bin | 76 |
| Managing containers with Recycle Bin enabled | 76 |
| Container deletion | 76 |
| Container replication | 76 |
| Cloud tier and archive tier | 77 |
| Disaster recovery | 77 |
| Viewing Recycle Bin statistics | 77 |
| Configuring Cloud Reader | 77 |
| Understanding Cloud Reader mode | 77 |
| Operations allowed | 77 |
| Compatibility | 78 |
| Deploying Cloud Reader | 78 |
| Refreshing data in Cloud Reader | 78 |
| Configuring RDA immutability | 79 |
| Understanding RDA immutability | 79 |
| Managing containers with RDA immutability | 79 |
| Container deletion | 79 |
| Container replication | 79 |
| Cloud Tier and Archive Tier | 79 |
| Viewing immutable container statistics | 79 |
| Managing containers | 81 |
| Creating a container | 81 |
| Creating an OST or RDS connection type container | 82 |
| Configuring User Access Controls | 83 |
| Creating an NFS or CIFS connection type container | 84 |

| | |
|---|------------|
| Creating an Object Container | 86 |
| Adding an object container through the command line | 87 |
| Creating a VTL type container | 88 |
| Viewing VTL tape information | 91 |
| Creating an EDM connection type container | 92 |
| Viewing containers | 93 |
| Viewing containers in the GUI | 93 |
| Viewing containers via the CLI | 93 |
| Viewing container statistics | 94 |
| Displaying container statistics by using the CLI | 94 |
| Adding a cloud tiering policy | 95 |
| Using regular expressions | 97 |
| Example expressions | 97 |
| Limitations | 97 |
| Adding an Archive Tiering policy | 97 |
| Deleting a container | 98 |
| Deleting a container through the GUI | 99 |
| Deleting a container through the command line | 99 |
| Managing local storage | 100 |
| Viewing storage group information | 100 |
| Adding a storage group | 101 |
| Adding a storage group through the GUI | 102 |
| Adding a storage group through the command line | 103 |
| Modifying a storage group | 103 |
| Deleting a storage group | 104 |
| Deleting a storage group from the GUI | 105 |
| Deleting a storage group from the CLI | 105 |
| Configuring a Performance Tier | 105 |
| Adding a performance tier | 105 |
| Adding a performance tier through the command line | 106 |
| Editing a performance tier | 107 |
| Configuring Object Container | 108 |
| Creating an Object Container | 108 |
| Adding an object container through the command line | 109 |
| Creating a bucket | 109 |
| Editing bucket settings | 110 |
| Changing bucket retention | 111 |
| Managing Object container users | 112 |
| Updating Object container user | 112 |
| Deleting Object container user | 112 |
| Configuring additional storage | 113 |
| Guidelines for configuring additional storage | 113 |

| | |
|---|------------|
| Adding additional storage | 113 |
| Managing cloud storage | 115 |
| Understanding cloud and archive storage | 115 |
| Policy-based cloud storage | 115 |
| Direct-to-cloud storage | 116 |
| Limitations of direct-to-cloud containers | 116 |
| Configuring an RDS direct-to-cloud container | 117 |
| Configuring a direct-to-cloud container in the QoreStor UI | 117 |
| Configuring a direct-to-cloud container using the CLI | 117 |
| Cloud tiering | 117 |
| Adding a Cloud Tier through the GUI | 118 |
| Adding a Microsoft Azure cloud tier | 118 |
| Adding an Amazon S3 cloud tier | 119 |
| Adding a Wasabi S3 cloud tier | 121 |
| Adding an IBM S3 cloud tier | 123 |
| Adding a Google S3 cloud tier | 125 |
| Adding a Scality-Artasca-S3 cloud tier | 127 |
| Adding an S3 Compatible cloud tier | 128 |
| Adding a Backblaze S3 cloud tier | 130 |
| Adding a cloud tier through the command line | 131 |
| Creating a cloud tiering schedule | 131 |
| Editing a cloud tiering schedule | 132 |
| Deleting a cloud tier | 132 |
| Deleting a cloud tier from the GUI | 132 |
| Deleting a cloud tier from the CLI | 133 |
| Configuring a Cloud Archive Tier | 133 |
| Configuring required permissions to restore from Archive Tier | 134 |
| Modifying an Archive Tier after an upgrade | 138 |
| Adding an archive tier | 140 |
| Editing an archive tier restore mode using the command line interface | 142 |
| Deleting an archive tier | 144 |
| Deleting an archive tier from the GUI | 144 |
| Deleting an archive tier from the CLI | 144 |
| Creating an archive tiering schedule | 145 |
| Editing an archive tiering schedule | 145 |
| Restoring from an archive tier | 145 |
| Restoring files from RDS Container backups replicated to AWS S3 Glacier or Deep Archive | 146 |
| Restoring selective tapes of VTL backups replicated to AWS S3 Glacier or Deep Archive | 146 |
| Performing a disaster recovery from the cloud | 147 |
| Next steps | 149 |
| Manually restoring datastores from Amazon S3 Glacier | 150 |
| Managing replications | 152 |

| | |
|--|------------|
| Guidelines and prerequisites for replication | 152 |
| Adding replication relationships | 153 |
| Configuring replication schedules | 154 |
| Viewing replication information | 154 |
| Modifying replication relationships | 155 |
| Deleting replication relationships | 156 |
| Starting and stopping replication | 156 |
| Managing users | 158 |
| Viewing users | 158 |
| Viewing users through the GUI | 158 |
| Viewing users through the command line | 159 |
| Adding a user | 159 |
| Adding a user through the GUI | 159 |
| Adding a user through the command line | 160 |
| Modifying local user roles | 160 |
| Modifying a local user through the GUI | 161 |
| Modifying a user through the command line | 161 |
| Changing a password for a local user | 161 |
| Deleting a user | 162 |
| Deleting a user account through the GUI | 162 |
| Deleting a user account through the command line | 162 |
| Monitoring QoreStor | 163 |
| Using the Dashboard page | 163 |
| Viewing QoreStor statistics by using the CLI | 164 |
| Monitoring system alerts | 166 |
| Monitoring clients | 166 |
| Monitoring clients through the QoreStor GUI | 166 |
| Monitoring clients through the QoreStor CLI | 166 |
| Monitoring system events | 167 |
| Getting daily usage statistics from QoreStor | 167 |
| Managing QoreStor Remotely | 168 |
| Getting started with QorePortal | 168 |
| Registering QoreStor with Quest QorePortal | 168 |
| Enabling Remote Management | 169 |
| Viewing and using QorePortal | 169 |
| Support, maintenance, and troubleshooting | 171 |
| Using QoreStor Diagnostics | 171 |
| Viewing system diagnostic log files | 171 |

| | |
|---|------------|
| Understanding diagnostics collection | 171 |
| Generating a diagnostics log file | 172 |
| Downloading diagnostics log files | 173 |
| Deleting a Diagnostics Log File | 173 |
| Troubleshooting error conditions | 173 |
| Excluding QoreStor directories from antivirus scans | 174 |
| Security recommendations guide | 175 |
| About us | 1 |
| Technical support resources | 1 |

Introducing Quest® QoreStor™

Quest® QoreStor™ is a software-defined secondary storage platform based on Quest's proven DR Appliance's resilient deduplication and replication technologies. With QoreStor, you can break free of backup appliances and accelerate backup performance, reduce storage requirements and costs, and replicate safer and faster to the cloud for data archiving, disaster recovery and business continuity.

QoreStor supports all of the major backup software applications in use today and can lower your backup storage costs to as little as \$.16/GB while reducing your total cost of ownership. QoreStor achieves these results using patented Rapid technology as well as built-in, variable block-based deduplication and compression.

Lower costs and maximize the return on your IT investment by leveraging virtually any storage hardware, virtualization platform or cloud provider. QoreStor also supports many backup software solutions — so it's not just for Quest. Simple to deploy and easy to manage, QoreStor enables you to shrink replication time, improve data security and address compliance requirements.

QoreStor helps you to:

- Reduce on-premises and cloud storage costs with industry-leading deduplication and compression.
- Accelerate backup completion with protocol accelerators and dedupe.
- Shrink replication time by transmitting only changed data.
- Improve data security and comply with FIPS 140-2.
- Maximize return on investment for existing data protection technologies.
- Lower total cost of ownership through all-inclusive licensing.

QoreStor includes the following features:

- Hardware and software agnostic platform
- Next-generation storage dedupe engine
- Built-in protocol accelerators
- Support for a wide variety of data backup installations and environments.

Understanding the QoreStor documentation

The topics in this guide introduce and describe how to use the QoreStor web-based graphical user interface (GUI) to manage your system. It describes how to access the comprehensive system GUI and the associated features and capabilities, how to perform a wide variety of data storage and replication operations, how to manage the system, as well as how to manage the related storage and replication containers.

In addition to the QoreStor GUI, you can manage QoreStor by using a command-line interface (CLI). In some instances, the GUI provides additional features and options that are not available in the CLI and vice versa.

This documentation is written for an administrator.

i | **NOTE:** For information about the supported web browsers you can use with QoreStor, see the **System Requirements** chapter in the *QoreStor Installation Guide*.

Other information you may need

The following table lists the documentation available for QoreStor. The documents listed in this table are available on the Quest support website by selecting your specific QoreStor version at:

<http://support.quest.com/QoreStor>

Table 1: QoreStor documentation

| Document | Description |
|---|--|
| QoreStor Installation Guide | Provides information on installation and operation requirements, supported platforms as well as procedures for installing QoreStor. |
| QoreStor User Guide | Provides information on configuring and using QoreStor. |
| QoreStor Release Notes | Provides the latest information about new features and known issues with a specific product release. |
| QoreStor Command Line Reference Guide | Provides information about managing QoreStor data backup and replication operations using the QoreStor command line interface (CLI). |
| QoreStor Interoperability Guide | Provides information on supported infrastructure components. |
| QoreStor Virtual Machine Deployment Guide | Provides information on deploying the QoreStor virtual machine on VMware ESX or Microsoft Hyper-V. |
| Additional whitepapers | Instructions and best practices for configuring additional Quest and third-party applications to work with QoreStor. |

i | **NOTE:** Check for the latest documentation updates and release notes at <http://support.quest.com/qorestor>. Read the release notes first because they contain the most recently documented information about known issues with a specific product release.

Information on compatible products

QoreStor offers direct integration with Quest Software's NetVault[®] Backup and vRanger[®], as well as Veritas NetBackup and Backup Exec. For more information on those products refer to the documents below.

Table 2: Quest NetVault documentation

| Document | Description |
|-----------------------------|---|
| NetVault Installation Guide | Provides information about installing and upgrading the NetVault server and client software. |
| NetVault Administration | Describes how to configure and use NetVault to protect your data. This document also provides information on configuring QoreStor repositories and migrating NetVault SmartDisk data to the |

| Document | Description |
|------------------------|--|
| Guide | new QoreStor repository. |
| NetVault Release Notes | Provides the latest information about new features and known issues with a specific product release. |

i | **NOTE:** See the complete NetVault documentation at <https://support.quest.com/netvault>.

Table 3: Quest vRanger documentation

| Document | Description |
|------------------------------------|--|
| vRanger Installation/Upgrade Guide | This document provides information on supported platforms, system requirements, and instructions on installing and upgrading vRanger. |
| vRanger User Guide | This document provides information and procedures on configuring and using vRanger to protect virtual and physical environments. |
| vRanger Release Notes | This document details the issues resolved in this release, the known issues as of this release, and the third party components in vRanger. |

i | **NOTE:** See the complete vRanger documentation at <https://support.quest.com/vranger>.

Table 4: Veritas documentation

| Document | Description |
|---------------------|--|
| Veritas NetBackup | For information on Veritas NetBackup, refer to the NetBackup product documentation . |
| Veritas Backup Exec | For information on Veritas Backup Exec, refer to the Backup Exec product documentation . |

Source code availability

A portion of the QoreStor may contain or consist of open source software, which you can use under the terms and conditions of the specific license under which the open source software is distributed.

Under certain open source software licenses, you are also entitled to obtain the corresponding source files. For more information or to find the corresponding source files for respective programs, see the Quest website at opensource.quest.com.

QoreStor data storage concepts

Data deduplication and compression

The QoreStor design uses various data-reduction technologies, including advanced deduplication algorithms, in addition to the generic and custom compression solutions that prove effective across many differing file types. Data deduplication and compression are addressed in the following areas:

- **Deduplication** — This technology eliminates redundant copies of data and in the process it decreases disk capacity requirements and reduces the bandwidth needed for data transfer. Deduplication can be a major asset for companies that are dealing with increasing data volumes and require a means for optimizing their data protection.
- **Compression** — This technology reduces the size of data that is stored, protected, and transmitted. Compression helps companies improve their backup and recovery times while helping reduce infrastructure and network resource constraints.

In general, QoreStor offers advanced deduplication and compression capabilities to reduce the time and cost associated with backing up and restoring data. Based on deduplication and compression technology, QoreStor eliminates the need to maintain multiple copies of the same data. This lets customers keep more data online longer and reduce the need for tape backup dependency.

Using its deduplication and compression technology, QoreStor can help achieve an expected data reduction ratio of 15:1. Achieving this reduction in data means that you need fewer incremental storage operations to run and it provides you with a smaller backup footprint. By removing redundant data, QoreStor deliver fast reliable backup and restore functionality, reduce media usage and power and cooling requirements, and improve your overall data protection and retention costs.

For a complete list of supported management application, refer to the *QoreStor Interoperability Guide*.

Encryption at rest

Data that resides on QoreStor can be encrypted. When encryption is enabled, QoreStor uses the Industry standard FIPS 140-2 compliant 256-bit Advanced Encryption Standard (AES) encryption algorithm for encrypting and decrypting user data. The content encryption key is managed by the key manager, which operates in either a Static mode or an Internal mode. In Static mode, a global, fixed key is used to encrypt all data. In internal mode, key lifecycle management is performed in which the keys are periodically rotated. The minimum key rotation period before the content encryption key can be rotated and a new key is generated is 7 days. This rotation period is user-configurable and can be specified in days. A user-defined passphrase is used to generate a pass phrase key, which is used to encrypt the content encryption keys. It is mandatory to define a passphrase to enable encryption. The system supports up to a limit of 1023 different content encryption keys.

i | **NOTE:** Due to export regulations, the encryption at rest feature is not available in certain markets, and, therefore, may not be available in your locale.

Streams and connections

This topic describes the differences between data streams and application connections.

Streams refer to the number of files written at the same time to QoreStor. QoreStor tracks the number of files being written and assembles the data into 4-MB chunks before processing that section of the data. If the stream count is exceeded, the data is processed out of order and overall deduplication savings can be affected.

Connections are created by applications; and, within a single connection, there can be multiple streams depending on the application and the number backup jobs running in parallel over that single connection.

Recycle Bin

Recycle Bin temporarily holds deleted files. The QoreStor administrator is allowed to recover these deleted files in cases such as when backup files are deleted accidentally or maliciously. The retention time for the files in Recycle Bin can be configured on a container, and is based on the QoreCompliance Clock, an internal software clock, rather than the system clock. Once configured, files that are deleted from the container are not deleted permanently, rather they are retained in Recycle Bin for the configured duration. Files in Recycle Bin are not available for the backup applications until they are recovered. QoreStor automatically purges data from Recycle Bin based on the retention time of each file.

Quotas

QoreStor provides the option to configure quotas for Storage Groups and Performance Tiers. With quotas, you can define a limit for physical capacity usage. Upon reaching that limit, data ingests to that storage group or performance tier will be disallowed until sufficient space is recovered through cleaner operations.

Before considering implementing quotas, please take into account these limitations and recommendations:

- Quotas are not set by default.
- The minimum quota value is 100 GiB.
- Once set, quotas can be reduced, but cannot be set below the storage group's used capacity.
- For standard storage groups, the maximum quota value is restricted to the current physical capacity of the QoreStor system.
 - For performance tiers, the maximum quota value is restricted to the the size of the performance tier filesystem.

Secure connect

Secure Connect encompasses a set of client and server components that creates a secure channel for QoreStor communication with WAN-connected clients that is also resilient to WAN outages.

Secure Connect uses the TLS 1.2 standard with a 4096-bit RSA key. Certificates are created automatically for both client and server, but you can use you own certificates if you chose. When using Secure Connect (which is enabled by default when using the latest QoreStor and plug-in versions), the client opens a connection to the QoreStor server over port 9443. The client sends the actual QoreStor port number to the server, which then opens a local connection to that port enabling secure communication with the client. Configuration of Secure Connect ports is done through the client and server configuration file. See [Configuring Secure Connect](#) for more information.

Secure Connect also provides a method for resilient WAN connections. Packets processed by Secure Connect clients are assigned a unique identifier and are assigned to a temporary cache before being sent to the QoreStor server. When the packet is successfully delivered to the QoreStor Secure Connect server, the packet identifier is marked as delivered and acknowledged to the Secure Connect client. If the WAN connection is lost, the client and

server both continue to cache data packets. When the connection is restored, unacknowledged packets are re-sent and properly processed, avoiding data loss and process interruption.

MultiConnect

QoreStor MultiConnect establishes multiple connections between RDA and OST clients and the QoreStor server. On high-bandwidth networks, these connections can improve backup and replication performance. To obtain the most benefit from MultiConnect, both MultiConnect and SecureConnect should be enabled. Refer to the sections below:

- [Enabling Secure Connect for OST and RDA plug-ins prior to 4.1.0.265](#)
- [Enabling MultiConnect](#)

Replication

Replication is the process by which key data is saved from storage locations, with the goal of maintaining consistency between redundant resources in data storage environments. Data replication improves the level of fault-tolerance, which improves the reliability of maintaining saved data and permits accessibility to the same stored data.

QoreStor uses an active form of replication that lets you configure a primary-backup scheme. During replication, the system processes data storage requests from a specified source to a specified replica target, which acts as a replica of the original source data.

i **NOTE:** QoreStor includes version checking that limits replication only between other QoreStor instances or DR Series systems that run a compatible software release version. If versions are incompatible, the administrator is notified by an event.

Replicas are read-only and are updated with new or unique data during scheduled or manual replications. QoreStor can be considered to act as a form of a storage replication process in which the backup and deduplication data is replicated in real-time or via a scheduled window in a network environment. In a replication relationship between two or three QoreStor instances or DR Series systems, this means that a relationship exists between a number of systems. One system acts as the source and the other as a replica.

Replication is done at the container level and is one directional from source to replica; however, since replication is done at the container level you can set up various containers to meet your specific replication requirements for your specific workflow. This form of replication is supported for the CIFS, NFS, Rapid CIFS, and Rapid NFS, RDA, and EDM protocols and is fully handled by QoreStor.

i **NOTE:** Using DMA managed copy for object container is recommended.

QoreStor supports replication seeding, which provides the ability to create a local seed and place it in a remote system. The seed backup is a process on the source QoreStor system, which collects all of the unique data chunks from the containers and stores them on the target device. This is helpful if you have a new replication target to set up, the amount of data to be replicated is very large, and the network bandwidth is low.

i **NOTE:** The storage capacity of the target QoreStor system is directly affected by the number of source systems writing to its containers, and by the amount being written by each of the source systems.

If the source and target systems are in different Active Directory (AD) domains, then the data that resides on the target system may not be accessible. When AD is used to perform authentication for QoreStor systems, the AD information is saved with the file. This can act to restrict user access to the data based on the type of AD permissions that are in place.

i | **NOTE:** This same authentication information is replicated to the target QoreStor system when you have replication configured. To prevent domain access issues, ensure that both the target and source systems reside in the same Active Directory domain.

Reverse replication

The concept of reverse replication is not a supported operation on QoreStor. This is because replica containers are always in a R-O (read-only) mode on QoreStor, thus making write operations a non-supported operation.

Alternate ways to retrieve data

Under very specific conditions, it could be possible for replica containers to support a type of write operation whose sole function is to restore data from an archival target. For example, data could be replicated back to the remote site where a data management application (DMA), or backup software, is connected to allow this data to be restored directly.

This specific type of case applies only to configurations where data is backed up from a remote location to a local container, and then replicated over a WAN to a replica container that is backed up to tape. The data needs to be restored from the tape backup to the original location; first back to QoreStor replica container, and then back to the original source location of the data on the other side of the WAN link.

i | **NOTE:** If you choose to use this alternate workaround method, you must set up a new data storage unit in your DMA, and import the images before a restore to the original location can occur.

To leverage this type of deduplication across the WAN, complete the following:

1. Make sure that the replication operation has completed (between source and target).
2. Delete the current replication relationship, and re-create a replication relationship (reversing the source and target roles).
3. Restore data to the original source container (now the target).
4. Make sure that the replication operation has completed.
5. Delete the replication relationship and re-create a replication relationship (restoring original source and target destinations).

Under this scenario, a fraction of the data to be recovered is sent across the WAN link. This could speed up a remote restore significantly. However, there are some downsides to this type of scenario:

- If step 1 is not followed correctly, any changes not fully replicated are lost.
- During steps 2 and 3, any data that is written to the original QoreStor source container may be lost.
- During step 4, if the data is not fully replicated back before the switch is made, it may be lost.

Alternatively, you could still support this type of effort by completing the following:

1. Create a new container on the target QoreStor instance.
2. Set up replication from this container back to the source QoreStor system container.
3. Set up a new disk storage unit in the DMA and make sure that the DMA is aware of any new images.
4. Import the old images back into the DMA from the target QoreStor instance (the original source location).
5. Use a new disk storage unit in the DMA, and then restore the data back to the original client.

Reverse replication: alternate method

For an alternate method of reverse replication, complete the following steps:

1. Create a new container on the target QoreStor instance.
2. Set up replication from this container back to the source QoreStor container.
3. Set up a new disk storage unit in your Data Management Application (DMA) and make sure that the DMA is aware of any new images.
4. Import the old images back into the DMA from the target QoreStor instance (the original source location).
5. Use a new disk storage unit in the DMA, and then restore the data back to the original client.

Performance Tier

In situations where certain workloads have requirements for faster recovery, QoreStor allows you to write these workloads to a performance tier to enable faster read back unaffected by activity in other QoreStor Storage Groups. By utilizing a performance tier, you are able to maximize the value of higher-performing storage by ensuring that only the most critical workloads are written to it.

To create a QoreStor performance tier, you must create a physical volume comprised of high-performing storage (such as SSD) and then create a QoreStor storage group mapped to that volume. Containers created on that storage group will write and read from high-performance storage exclusively and will be isolated from read activity on other volumes.

Cloud Tier replication

QoreStor's cloud tier feature enables QoreStor data to be quickly and easily accelerated to the cloud tier. Using your existing data management applications (DMAs) and any supported protocol, files can be written to a QoreStor container and replicated to your cloud tier according to easily defined policies.

QoreStor provides a policy engine that allows you to set idle time and on-premises retention criteria to be used in identifying which files are most suited for replication to the cloud. Policies are defined at the container level and apply to all files within that container. Using the QoreStor Cloud Policy, you can replicate files based on:

- **Idle time** - replicate stable files idle for more than the selected number of hours.
- **File extensions** - replicate files that **match** or **do not match** names in a list of extensions.
- **Regular expressions** - include or exclude files based on their match to configured regular expressions.
- **File locations** - replicated files in a list of directories, or all files except those in a list of directories.

In addition, there is an **On-Prem Retention Age** policy that allows you to specify how long a copy of a file is kept after it has been replicated to the cloud tier. Once a file has been replicated, the file on the QoreStor server becomes a stub, meaning it exists in the namespace but the data exists only in the cloud tier. Once a file has been stubbed and moved to the cloud tier, that file can no longer be edited.

With these policy options, you are able to configure cloud tier replication to meet one of three use cases:

- **Data replication** - this creates a direct copy of the backup data stored in QoreStor on the cloud tier.
- **Extended hold** - this offloads older and less frequently accessed data to the cloud tier for long-term

archiving.

- Replicate and extend - this provides both a direct replica and a long-term archive.

CAUTION: When storing backups of a server using continuous data protection (CDP), if you delete a cloud tier that contains CDP backups, the backups that are older than the on-prem retention age become unrecoverable. A segment file may be present for these backups, but the associated files may be stubbed and become unlinked during cloud tier deletion, which prevents you from restoring the backup.

Any data that is transferred from the QoreStor instance by the cloud tier is encrypted with zero knowledge encryption. The encryption keys are solely owned by you. If the encryption keys are placed in the cloud tier, a passphrase is used to encrypt those keys and that passphrase is only known to you. For added security, QoreStor obfuscates metadata and data store objects that are stored in the cloud tier.

Starting with QoreStor release 7.2.1, up to two Cloud Tiers can be configured in an instance of QoreStor. You may use storage from different cloud vendors. Each QoreStor container can be mapped to only one Cloud Tier. Both Cloud Tiers can be active simultaneously and used for replicating data from a different container to cloud storage. For copying or migrating data from one Cloud Tier to another, please contact Quest Support for information on a suitable method for your environment.

Starting QoreStor 7.3.0, data tiered out to cloud providers to be locked in the cloud storage, which provides immutability in the cloud.

If AWS is the chosen cloud provider, then utilizing storage class options would directly store the objects in AWS's respective Storage Class thereby reducing the storage costs associated with AWS.

For more information about working with cloud tiers, see [Cloud tiering](#).

Archive tier

QoreStor's archive tier feature enables QoreStor data to be quickly and easily archived to long-term Amazon S3 Glacier or Amazon S3 Glacier Deep Archive storage. In addition, by choosing intelligent tiering for metadata, QoreStor Archive Tier metadata gets stored directly in AWS Intelligent Tiering Storage Class thereby reducing the costs associated with AWS. Using your existing data management applications (DMAs) and a supported protocol (Object(S3), VTL or RDS), files can be written to a QoreStor container and migrated to your archive tier according to easily defined policies. QoreStor provides a policy engine that allows you to set file age and on-premises retention criteria to be used in identifying which files are most suited for replication to the cloud. Policies are defined at the container level and apply to all files within that container.

Using the QoreStor Cloud Policy, you can replicate files based on:

- **Idle time** - replicate stable files idle for more than the selected number of hours.
- **File extensions** - replicate files that match or do not match names in a list of extensions.
- **Regular expressions** - include or exclude files based on their match to configured regular expressions.
- **File locations** - replicated files in a list of directories, or all files except those in a list of directories.

Any data that is archived from the QoreStor instance by the archive tier is encrypted with zero knowledge encryption. The encryption keys are solely owned by you. If the encryption keys are placed in the archive tier, a passphrase is used to encrypt those keys and that passphrase is only known to you. For added security, QoreStor obfuscates metadata and data store objects that are stored in the archive tier.

Data stored in the archive tier is not available for immediate recovery. When a recovery is initiated, the data stays in the archive tier while a copy is made in S3 standard storage and kept for an amount of time specified by the **archive_retention_in_warm** parameter. Although recovery times may vary, the general expectations for recovery times are:

- Amazon S3 Glacier storage: 3-5 hours
- Amazon S3 Glacier Deep Archive: within 12 hours

Object interface (S3-compatible)

QoreStor's Object container provides an object storage interface which enables customers to write Object data(S3 format) directly to QoreStor. This allows solutions that leverage an S3-based connection to send data directly to a QoreStor instance instead of Amazon S3 with the added benefits of deduplication, encryption, replication and network optimized data transfer.

Starting with QoreStor release 7.2.1, multiple Object containers can be created and used for backups. Each of them has a different set of users, policies, and buckets. Note that, Object container created before QoreStor release 7.2.1 will have the container name as ObjectContainer and the Storage Group name as ObjectStorageGroup.

Object container data can be replicated to another container using continuous replication. The object data and metadata including user information are replicated in that target container.

Object container is compatible with AWS S3 API, with certain limitations like custom user policies. Locking with Compliance and Governance modes and Object Versioning is supported. For overriding Governance mode lock settings, please contact Quest Support.

This container can be configured with a cloud tiering policy to seamlessly move data to long-term storage. QoreStor allows you to configure multiple buckets within your object container namespace, each with different locking and retention settings.

i | **NOTE:** Recycle Bin, Archive Tier, and Replication are not supported with Object container.

Object Direct Storage

Object storage organizes files and their associated metadata into objects, which are then stored in a flat address space. Object storage provides greater insight to data usage, improved scalability, and lower storage costs when compared to traditional file or block storage. QoreStor can be installed in object-direct mode, which utilizes object storage for the main data repository, while metadata is written to higher performing storage. For more information, refer to the topics **QoreStor Installation Modes** and **Installing QoreStor with Object Direct Storage** in the *QoreStor Installation Guide*.

i | **NOTE:** When QoreStor is installed in an Object Direct configuration, if the backend object storage is not accessible QoreStor will go into Manual Intervention mode.

Cloud Reader mode

Cloud Reader mode allows a temporary QoreStor instance to read the backup data stored in a cloud tier from another instance of QoreStor. It helps with restoring backups for auditing the data in the cloud and for testing disaster recovery using cloud data.

Continuous data protection

Continuous data protection (CDP) occurs when backup software takes frequent snapshots of a protected virtual machine (VM) and saves the backups to backup storage, where they can later be recovered. QoreStor supports

CDP backups over RDA protocol. Backup software can take advantage of the deduplication and compression capabilities that QoreStor offers, and the optimized network transfers of the RDA protocol.

Starting from NetVault 13.1 backup software uses the QoreStor CDP feature for its continuous backup option.

i NOTE: While features such as immutability, optimized copy of backups, and Cloud Tier are supported with CDP backups, the following features are not supported with CDP:

- Recycle Bin
- Archive Tier
- Container replication

Synthetic full backups

CDP begins with a full backup of the protected VM and is followed by a series of incremental backups for subsequent changes. When QoreStor rolls up the incremental backups with the initial full backup, it generates a complete and current image of the protected VM called a synthetic full backup. In the backup software, you can determine how often to create these synthetic backups. For more information, see the *NetVault Administrator's Guide*.

Instant Restore

With CDP backups, QoreStor lets you recover data using the Instant Restore feature. Backed-up files present as a temporary datastore over Linux network file sharing (NFS) protocol. NetVault then mounts the backed-up files on an ESX host, which lets the snapshot of the VM boot directly from QoreStor. You can then use NetVault to migrate the VM to another ESX host. For more information, see the *NetVault Administrator's Guide*.

Compatibility

QoreStor supports the CDP feature with NetVault 13.1 and later. CDP is compatible with the following NetVault plug-ins:

- NetVault Plug-in for VMware 13.1 and later
- Plug-in for Microsoft 365 13.2.5 and later.

Disaster Recovery from the cloud

When configured with a Cloud Tier or Archive Tier, QoreStor provides the ability to capture QoreStor configuration information as well as storage group and container data. With this information, it is possible to re-create a lost or failed QoreStor server using the data in the cloud. Please note that the data which is not cloud replicated and is still present in the local disk is not recoverable in case of a failure with the primary QoreStor.

Cloud tier disaster recovery

In the event of QoreStor server failure, a recovery can be initiated on a new, licensed QoreStor server to restore from the previous configuration stored in the cloud. At a high-level, a recovery will go through the following steps:

- Connect to your cloud provider with configured credentials and passphrase.

i | **IMPORTANT:** The required passphrase is the passphrase used when creating the Cloud Tier. Without the Cloud Tier passphrase recovery cannot proceed.

- The new QoreStor will read the stored configuration once it is connected to the cloud .
- Rebuilds the cloud-replicated containers as well as other storage groups and containers.
- Runs the filesystem check and brings the QoreStor to a consistent state.

i | **NOTE:** It is recommended to use the same version of QoreStor when performing disaster recovery from cloud-tier data as the previous version of the source QoreStor.

For more information refer to [Performing a disaster recovery from the cloud](#).

Archive tier disaster recovery

When performing data recovery from an Archive Tier, you must first restore all datastores to standard AWS S3 storage using the AWS Management Console. For more information refer to [Manual Restore from Glacier](#) and [Restoring from archive tier](#).

MSP Mode

QoreStor MSP feature enables backup service providers to host the backup data of multiple customers in a QoreStor system without sharing the post-deduplication data between them. It enables support for the following:

- Up to 32 Storage Groups
- Up to 128 Containers
- Up to 32 Cloud Tiers

MSP mode can be enabled on Enterprise Plus instances.

MSP mode can be enabled using `system --msp enable`.

Refer QuestQoreStor Command Line Reference Guide for details.

Partial File System Checker (OFSCK)

On MSP-enabled systems, the 'Partial ofscck' feature enables the filesystem checker on a specific storage group while allowing other storage groups to remain in read-write mode. The specified storage group undergoing ofscck will be set to read-only mode.

Partial filesystem checker (ofscck) can be run from maintenance CLI with the option `--partial`.

(Ex: `maintenance --filesystem --start_scan --storage_group DefaultGroup --partial`. This will run ofscck on DefaultGroup and other storage groups are in read-write mode).

See the Quest QoreStor Command Line Reference Guide for details.

Only one storage group is allowed to run for this ofscck at a time. Additionally, only local storage groups are supported; cloud storage groups are not supported.

While this ofscck is run on a specific storage group, all config updates that requires this storage group are disabled.

(Ex: Adding containers to this storage group)

Supported file system protocols

QoreStor supports the following file system protocols. The Rapid Data Access (RDA) protocols below provide a logical disk interface that can be used with network storage devices to store data and support data storage operation.

- Network File System (NFS)
- Common Internet File System (CIFS)
- Rapid Data Access (RDA)
 - Rapid NFS
 - Rapid CIFS
 - RDA with OpenStorage Technology (OST)
 - RDA with NetVault
 - RDA with vRanger
- The virtual tape library (VTL) tape access protocols:
 - Network Data Management Protocol (NDMP)
 - Internet Small Computer System Interface (iSCSI)
- Other supported access protocols
 - Object (S3)
 - Veeam Enhanced Data Mover(EDM)

CIFS

The Common Internet File System (CIFS) remote file access protocol is supported by QoreStor, and is also known as a Server Message Block (SMB). SMB occurs more commonly than the Network File System (NFS) protocol on systems that run the Microsoft Windows operating system. CIFS allows programs to request files or services on remote computers.

CIFS also uses the client-server programming model, whereby the client requests access to a file or passes a message to a program running on the server. Servers review all requested actions and return a response. CIFS is a public (or open) variation of the SMB that was originally developed and used by Microsoft.

i | **NOTE:** QoreStor currently supports version 2.0 and 3.0 of the Server Message Block (SMB).

i | **NOTE:** For complete details on CIFS feature restrictions, see the *QoreStor Interoperability Guide*, at support.quest.com/qorestor.

CIFS ACL support

QoreStor supports the use of access control lists (ACLs) for CIFS and share-level permissions. By definition, an ACL is simply a list of permissions that can be associated with any network resource.

Each ACL can contain access control entries (ACEs) that define or describe the permissions for an individual user or a group of users. An ACL can consist of zero (meaning that all users have access) or a number of ACEs that define specific permissions on a per-user or per-group basis.

i | **NOTE:** If an ACE list is empty (meaning that it contains zero entries), this means that all access requests will be granted.

An ACL describes the entities that are allowed to access a specific resource. ACLs are a built-in access control mechanism in the Windows operating systems.

i | **NOTE:** QoreStor supports setting up share-level permissions for a CIFS share using a Microsoft Windows administrative tool. Share-level permissions let you control access to shares. For more information, see [Configuring share-level security for CIFS shares](#).

Making CIFS backup immutable for Veeam

QoreStor 7.5.0 onwards can protect Veeam backups over CIFS from unintentional changes. This can be achieved using UI option while creating a container or using CLI option “-veeam_overwrite_protect”. This option works with CIFS and Rapid CIFS protocols. Veeam customers can achieve immutability using this option along with enabling Recycle Bin. Once this option is enabled, it cannot be undone or disabled.

Due to its chain-level operation and lack of merges, Veeam forward incremental flow is the only application for which it is employed. Thus, backup files ending in .vbk and .vib, both complete and incremental, will only be unchangeable.

Following a backup and a predetermined amount of time that has passed since the file's last write activity, except delete and expiry, it is automatically rendered immutable for file operations like write, append, truncate, etc. It is safeguarded by a recycle bin for deletion. Therefore, as soon as this new option is activated on the CIFS container, the recycle bin is instantly enabled. All file operations are permitted on the system until the cool-off period.

The same is cascaded to the target QoreStor if replication is set up with this option. Therefore, the same data files become immutable after the cool-off period when replication is stopped. Only on-premises data is immutable for cloud replication.

Also, during disaster recovery, the data is not locked on the CIFS container. So backups are immutable post cool-off period.

NFS

The Network File System (NFS) is a file system protocol that is designated to be a file server standard, and its protocol uses the Remote Procedure Call (RPC) method of communication between computers. Clients can access files via the network similar to the way that local storage is accessed.

NFS is a client-server application in which a client can view, store, and update files on a remote system just like they are working on a local system. System or Network Administrators can mount all or a portion of a file system, and the file system (or portion) that is mounted can be accessed using the privileges assigned to each file.

i | **NOTE:** If you want to do a mount on AIX, you must set the `nfs_use_reserved_ports` and `portcheck` parameters first. The parameters cannot be set to 0. For example: `root@aixhost1 / # nfso -po portcheck=1 root@aixhost1 / # nfso -po nfs_use_reserved_ports=1`

Rapid NFS and Rapid CIFS benefits

When Rapid NFS and Rapid CIFS are used with QoreStor they offer the following benefits:

- Reduce network utilization and DMA backup time
 - Chunk data and perform hash computation on the client; transfer chunked hash files on the back-end
 - Reduce the amount of data that must be written across the wire
- Improve performance
- Support DMAs such as CommVault, EMC Networker, and Tivoli Storage Manager. For the current list of supported DMAs, see the *QoreStor Interoperability Guide*.
- Compatible with existing NFS and CIFS clients — just need to install a plug-in (driver) on the client
 - Can use Rapid NFS and Rapid CIFS to accelerate I/O operations on any client — including a client that uses home-grown backup scripts
 - Can service multiple and concurrent media server backups

Rapid NFS and Rapid CIFS

Rapid NFS and Rapid CIFS enable write operation acceleration on clients that use QoreStor replication and NFS or CIFS file system protocols. Similar to OST and RDS, these accelerators allow for better coordination and integration between QoreStor backup, restore, and optimized deduplication operations with Data Management Applications (DMAs) such as CommVault, EMC Networker, and Tivoli Storage Manager. For the current list of qualified DMAs, see the *QoreStor Interoperability Guide*.

Rapid NFS is a new client file system type that ensures that only unique data is written to QoreStor. It uses user space components and file system in user space (FUSE) to accomplish this. Metadata operations such as file creates and permission changes go through the standard NFS protocol, whereas write operations go through RDNFS.

Rapid CIFS is a Windows-certified filter driver that also ensures that only unique data is written to QoreStor.

All chunking and hash computations are done at the media or client server level.

Rapid NFS and Rapid CIFS require you to install a plug-in on the client or media server, depending on your DMA and configuration. For details, see "Configuring and Using Rapid NFS and Rapid CIFS".

RDA with Quest's NetVault and vRanger

Rapid Data Access (RDA) with NetVault and with vRanger provides the logical disk interface that can be used with network storage devices. QoreStor requires the NetVault Plug-in *for Rapid Data Access* to integrate its data storage operations with NetVault and vRanger. The plug-in is installed by default on the NetVault and vRanger servers and QoreStor when the latest software updates are installed. Using the Plug-in *for Rapid Data Access*, NetVault can take full advantage of key QoreStor and DR Series system features like data deduplication and managed replication.

When the Plug-in *for Rapid Data Access* is used with QoreStor, it offers the following benefits:

- RDA with NetVault and RDA with vRanger protocols provides faster and improved data transfers:
 - Focus is on backups with minimal overhead
 - Accommodates larger data transfer sizes
 - Provides throughput that is better than CIFS or NFS
- RDA and data management application (DMA) integration:
 - NetVault-to-media server software communication
 - QoreStor storage capabilities can be used without extensive changes to NetVault or vRanger.
 - Backup and replication operations are simplified by using built-in DMA policies
 - RDA provides an immutability feature, which NetVault plans to support soon. For compatible versions, see the *Quest QoreStor Interoperability Guide*.
- QoreStor RDA ports and write operations:
 - Control channel uses TCP port 10011
 - Data channel uses TCP port 11000
 - Optimized write operations enable client-side deduplication
- Replication operations between QoreStor systems:
 - No configuration is required on the source or target QoreStor systems
 - Replication is file-based, not container-based
 - Replication is triggered and managed by the backup application
 - QoreStor transfers the data file (not the media server)
 - After duplication completes, QoreStor notifies the DMA to update its catalog (acknowledging the second backup). This makes the DMA aware of the replication location. Restores from either the source or replication target can be used directly from the DMA.
 - Supports different retention policies between source and replica
 - Replication is set up in the DMA itself, not QoreStor

RDA with OST for QoreStor

OpenStorage Technology (OST), by Veritas, provides a logical disk interface for use with network storage devices. QoreStor can use OST via QoreStor plug-in software to integrate its data storage operations with NetBackup and Backup Exec.

RDA with OST allows for better coordination and tighter integration between QoreStor system backup, restore, and optimized duplication operations and data management applications (DMAs). For a list of the supported applications, see the *QoreStor Interoperability Guide*.

Integration is done via a RDA with OST plug-in developed for QoreStor, through which data management applications can control when the backup images are created, duplicated, and deleted. The major benefit of RDA with OST is that it allows the deduplication operations to happen on the client side so that network traffic can be reduced.

The RDA with OST plug-in allows data management applications to take full advantage of such QoreStor features as data deduplication, replication, and energy efficiency. QoreStor systems can access the OpenStorage API code through the plug-in, which can be installed on the media server platform choice you make (Windows or Linux). The

OST protocol allows the supported backup applications to communicate directly with QoreStor and determine whether a specific chunk of data already exists on the system. This process means that if the data already exists, only the pointers need to be updated on QoreStor, and the duplicate chunk of data does not need to be transferred to the system. This process provides two benefits: it improves the overall backup speed, and also reduces the network load.

When RDA with OST is used with QoreStor, it offers the following benefits:

- OST protocol provides faster and improved data transfers:
 - Focused on backups with minimal overhead
 - Accommodates larger data transfer sizes
 - Provides throughput that is significantly better than CIFS or NFS
- RDA with OST and DMA integration:
 - OpenStorage API enables the DMA-to-media server software communications
 - QoreStor storage capabilities can be used without extensive changes to DMAs
 - Backup and replication operations are simplified by using built-in DMA policies
- QoreStor and RDA with OST:
 - Control channel uses TCP port 10011
 - Data channel uses TCP port 11000
 - Optimized write operations enable client-side deduplication

Software components and operational guidelines

To better coordinate and integrate OpenStorage Technology (OST) with QoreStor data storage operations, the following guidelines list the required components and supported operations. For details on the supported operating systems and data management application (DMA) versions, see the *QoreStor Interoperability Guide*.

QoreStor licensing is all-inclusive, so that no additional licensing is required to use OST or the optimized duplication capability. The OST plug-in that gets installed on a supported Linux or Windows media server platform is a free download. However, Veritas NetBackup requires that you purchase an OpenStorage Disk Option license. Similarly, Veritas Backup Exec requires that you purchase the Deduplication Option to enable the OST feature.

- OST Media Server Component:
 - An OST server component resides on the QoreStor server.
 - For Linux media server installations, use the Linux OST plug-in and the Red Hat Package Manager (RPM) installer
 - For Windows media server installations, use the Windows OST plug-in and the Microsoft (MSI) installer
- Windows-based OST plug-in
- Linux-based 64-bit OST plug-in
- Supported OpenStorage (OST) protocol:
 - Version 9
 - Version 10

- Supported Veritas DMAs
 - NetBackup
 - Backup Exec
- Supported OST operations
 - Backup (Passthrough writes and Optimized writes)
 - Restore
 - Replication
 - Auto Image Replication (AIR)

NDMP

The Network Data Management protocol (NDMP) is used to control data backup and recovery between primary and secondary storage in a network environment. For example, a NAS server (Filer) can talk to a tape drive for the purposes of a backup.

You can use the protocol with a centralized data management application (DMA) to back up data on file servers running on different platforms to tape drives or tape libraries located elsewhere within the network. The protocol separates the data path from the control path and minimizes demands on network resources. With NDMP, a network file server can communicate directly to a network-attached tape drive or virtual tape library (VTL) for backup or recovery.

The QoreStor VTL container type is designed to work seamlessly with the NDMP protocol.

iSCSI

iSCSI or **Internet Small Computer System Interface** is an Internet Protocol (IP)-based storage networking standard for storage subsystems. It is a carrier protocol for SCSI. SCSI commands are sent over IP networks by using iSCSI. It also facilitates data transfers over intranets and to manage storage over long distances. iSCSI can be used to transmit data over LANs or WANs.

In iSCSI, clients are called *initiators* and SCSI storage devices are *targets*. The protocol allows an *initiator* to send SCSI commands (*CDBs*) to the *targets* on remote servers. It is a storage area network (SAN) protocol, allowing organizations to consolidate storage into data center storage arrays while providing hosts (such as database and web servers) with the illusion of locally attached disks. Unlike traditional Fibre Channel, which requires different cabling, iSCSI can be run over long distances using existing network infrastructure.

iSCSI is a low-cost alternative to Fibre Channel, which requires dedicated infrastructure except in FCoE (Fibre Channel over Ethernet). Note that the performance of an iSCSI SAN deployment can be degraded if not operated on a dedicated network or subnet

The VTL container type is designed to work seamlessly with the iSCSI protocol. For details, see [Creating a VTL type container](#).

Hardened Repository with Veeam

QoreStor supports integration with a hardened repository through an EDM connection. This EDM connection is utilized by Veeam DMA, which is compatible with hardened repository functionality. A hardened repository ensures filesystem-level immutability, providing robust protection for stored files.

Replication, cloud replication, recycle bin, and cloud locking are supported for EDM containers.

QoreStor does not support fast clones. The QoreStor install mode determines how many EDM containers are supported. For further information, refer to the QoreStor Interoperability Guide.

With an EDM connection, QoreStor supports a repository without immutability.

All QoreStor platforms and Veeam workflows that make use of the Veeam hardened repository support the EDM connection.

To learn more about the Hardened Repository workflows and OS compatibility with EDM, refer to the Veeam documentation.

For more information, please refer to the Veeam EDM documentation.

Anomaly detection

An anomaly is any data point or suspicious event that stands out from the baseline/expected pattern. When data unexpectedly deviates from the established dataset, it can show an early sign of system malfunction, breaches, or Backup configuration changes.

e.g. unexpected data deletions, modifications, or excess data insertions.

Anomalies don't always signify an issue but they are all worth investigating to better understand why a deviation occurred and if that anomaly is a valid point as compared to baseline or training set.

Importance

- Identify ransomware attacks sooner
- Limit downtime
- Limit/detect early data loss
- Better understanding of changes in the environment

Challenges

Anomaly detection is only valuable if it can find true anomalies that means training the system before it can be useful. Otherwise, the system can relay an excessive number of alerts /anomalies beyond what one could feasibly investigate.

Retraining anomaly detection system helps in re-establishing a new baseline.

Anomaly detection approach

To detect anomalies, required data is collected periodically. This data is used for training the model and detect anomalies.

As of now QoreStor collects required data with the interval of 5-minutes.

Once the QoreStor is installed or upgraded to 7.5.0, the data collection starts. Data collected in the first 3 months (90 days) is used for getting a baseline. Once a baseline is established, anomalies are predicted using those baseline. After that every 30 days, the baseline is re-established with the last 90 days of data.

Anomaly detection categories

Anomalies detection is categorized in the following categories: System-level, Container-level, Storage-group level.

System-level

Anomalies : Following anomalies are detected:

- System-level log in authentication failures /anomalies
- QoreStor UI authentication failures /anomalies
- Protocol (OST/RDS) authentication failures /anomalies
- OS audit process stopped – this anomaly is reported as soon as the OSAudit process is either stopped or stopped/paused logging to audit files due to low space or any other issue
- Report high or excess system load average
- diskio anomalies reported for repository/metadata/enclosure filesystems:
 - Report excess number of IO operations executed
 - Excess average wait time being observed

Authentication-related anomalies do not need training as thresholds are used. For instance, if “root” user login authentication fails 3 times or more from a host “abc”, then it is reported as an anomaly.

The anomaly report shows the expected value and observed /current value with date range.

Load average

During the training period, the load average value is collected every 15 minutes. Once the training period is over, the maximum load average value is computed. The maximum value is chosen so that it occurs a certain number of times repeatedly (it need not be consecutive). We call this the trained maximum value.

Reporting

Please check if the current load average value exceeds the trained maximum value for a minimum of 30 minutes consecutively. If it exceeds, report it as a load-average anomaly.

Diskio

During the training, the total number of I/O operations and average wait time are collected periodically for all the Qs consumed by file systems like repository/metadata/enclosure. Once the training is over, the maximum total I/O operations are noted. The average wait time is trained using a linear regression model and is saved for later reference.

Detection /Reporting

Once the training is over, compare the current total number of I/O operations executing with the maximum training I/O count reported. If it exceeds the maximum count, report an anomaly.

Similarly, compute current average wait times and compare them with trained/expected average wait times from the linear regression model. If the average wait time exceeds 15 minutes or more in the last 1 hour, report it as an anomaly.

Diskio report is generated with the following anomaly description:

- Filesystem path on which anomalies reported
- Expected total number of I/O operations and current total I/O operations OR expected average wait time and current average wait time date range.

CLI: ‘system’ CLI can be used to configure the above anomaly detections. Please refer the QoreStor CLI Reference guide for more information.

Report: To determine the System level anomalies, the following entities are shown in the report.

- Client-name – Client from which failed authentication happened
- User name – Username used in authentication

- Failed count – Number of failed attempts
- Failed start/end time - The period of failed attempts occurred.

Container-level

Qorestor detects anomalies related to data ingest, data overwrites and data expiry at container level. For this, corresponding data is collected at regular intervals like the following metrics on that container.

Ingest and Overwrite: Detects Backup pattern and data size.

- Number of bytes ingested onto this container across clients within regular intervals
- The number of bytes overwritten across clients within regular intervals

Expiry

Files-deleted – Number of files/images deleted across clients. Internally tracks the total sizes of all deleted files. Data collected over 30-minutes of interval is used for anomaly detection.

i NOTE: : Even if the containers are removed, anomalies can be queried through CLI or UI.

CLI : Container CLI can be used to tune/set anomaly detection metrics. Please refer to the QoreStor CLI Reference guide for more information.

Anomaly settings applied at the Storage Group level are automatically applied to all containers in it unless explicitly disabled/turned off at the individual container level.

Report : The container level anomaly report shows the following anomaly types :

- Ingest – Shows bytes-ingested and corresponding savings (which is not inline with the training period/dataset)
- Overwrite – Total bytes overwritten from backup (not expected as per training set)
- Expiry – number of files deleted and the sum of all file sizes (not expected as per training set)
- Start/end time – The time of anomaly occurred in the container

Storage-group level

At the storage group level, savings anomalies are detected. Savings are further classified into the following sub-categories:

- Savings – dedupe - If total post-dedupe bytes are outside of the training range
- Savings – compression - If total post-compression bytes are outside of the training range

CLI: storage_group CLI can be used to set anomaly detection metrics at the storage group level. Please refer to the storage_group command in the QoreStorCommand Line Reference Guide Reference guide.

Report: The storage group anomaly report shows the following metrics:

- Anomaly type – Savings
- Anomaly sub-type – deduplication or compression

To decide if this is an anomaly, following parameters are used:

- Current value of deduplication or compression bytes
- Minimum value and maximum value expected i.e., expected range
- Difference with range i.e., with nearest minimum or maximum value

Retraining

Automatic retraining: Once first-time training is completed (after 3 months of data), periodically every one month, retraining happens with the last 3 months of data. This is to update recent data for the training period. This helps in tuning the baseline. This happens for both containers and storage groups.

Manual retraining: ocamltrain CLI can be used to do on-demand retraining. Please refer to the QoreStor Command Line Reference Guide guide for more information.

Alerts

Following alerts/events are raised related to anomaly detection :

- When stats collection is not happening
- When anomaly detection service is not running
- When OS-audit stops running, the os-authentication anomaly detection is enabled at the system level

Reports

Anomaly reports are shown in QoreStor UI or can be queried using ocamlreport CLI. Emails can also be configured using the following CLI to send anomalies as and when detected.

```
/opt/qorestor/bin/email_anomalies --configure
```

Please refer the QoreStorCommand Line Reference Guide guide for more details.

QoreStor processes

The table below describes the processes that QoreStor installs and runs on the QoreStor server.

Table 5: QoreStor processes

| Process | Description |
|-------------|---|
| influxd | This process helps in storing the data gathered by anomaly detection |
| minio | Server for Object container |
| nhm | Node Health Monitoring service - maintains a database for alerts and events |
| oca_idm_eda | An event and data aggregator service for events and alerts. |

| | |
|-----------------------|---|
| ocaagent_charts | Periodically (every hour) sends QoreStor chart related data to the Global View Cloud. |
| ocaagent_diagnostics | Periodically queries database for keep-alive commands. Responsible only for handling diagnostic upload command. |
| ocaagent_keepalive | Queries the Global View Cloud for keep-alive commands (diagnostic upload and portal unregistration) and runs them. This query occurs once a minute. |
| ocaagent_managebutton | Allows for remote control of QoreStor through Global View Cloud. When enabled, waits for remote management RPC commands from Global View Cloud, invokes them via API on QoreStor, and sends results back to the Global View Cloud. |
| ocaagent_registration | Handles Global View Cloud registration, operates on a DB table responsible for holding registration status, and performs registration/unregistration requests. |
| ocaagent_stats | Periodically(every 30 minutes, or when significant changes occur) sends statistic data (storage groups, containers, analytics, system information etc.) to the Global View Cloud. |
| ocaconfigsvc | QoreStor configuration service to manage the storage groups, containers, and replication links and handle requests from the QoreStor CLI and UI. |
| ocafsd | QoreStor file system service to export containers, process data and manage the data on the disk. |
| ocahttp | QoreStor HTTP server to service requests for QoreStor web UI and REST API methods. Uses ocafsd, ocaconfigsvc and r3 database (graphs). |
| ocaml | QoreStor anomaly detection service, periodically checks collected data and identifies anomalies across systems, containers, and storage-groups using the ML model. It is also responsible for generating ML training models and retraining. |
| ocamonitor | QoreStor process for periodically collecting stats from ocafsd, ocaconfigsvc and Linux system monitoring tools and populating an internal (r3) database. |
| ocardslogwriter | Process captures logging from all the processes, writes the data to the corresponding logs, and rotates the logs. |
| policy_mgr | This process helps with identifying files matching with the policy for cloud replication and initiating cloud replication and on-prem space retention for them. |
| sc_server | A Secure Connect server that handles secure connects from OST and RDA plug-ins |
| smbd | This process handles the CIFS share I/O activity |
| vtllibrary | This process handles the SCSI medium changer device |
| vtltape | This process handles SCSI tape device |
| watcher | Watcher process monitors other QoreStor processes. The watcher process will respawn other processes as necessary to keep the service online. |
| watcher_spawn | This process monitors watcher process and spawns if necessary. |

Accessing QoreStor

You can interact with QoreStor using one of the methods below:

- The QoreStor GUI, accessible in a web browser using the URL `https://<YourQoreStorServerName>:5233`.
- The QoreStor command line interface (CLI). Refer to the *QoreStor Command Line Reference Guide* for more information.

In the system GUI, you can configure your system as well as create and manage containers, which store your backup and deduplicated data. A data container is a shared file system that is imported using a client, and is accessible via file system or tape access protocols. For details, see [Supported file system protocols](#). The system GUI also provides real-time summary information for monitoring the status of the data capacity, storage savings, and the throughput of your data containers.

Networking prerequisites for QoreStor

Before you can start using QoreStor, ensure that you have satisfied the following networking prerequisites:

- **Network:** An active network is available using Ethernet cables and connections.
- **Replication ports:** the replication service in QoreStor requires that enabled fixed ports be configured to support replication operations that are to be performed across firewalls (TCP ports 9904, 9911, 9915, and 9916).
 - **NOTE:** For more information about replication ports, see [Managing replications](#), and for more information about system ports, see the *QoreStor Installation Guide*.
- **Proxy servers:** when using a proxy server, some additional configurations are required. Refer to [Configurations required when using proxy servers](#).

Configurations required when using proxy servers

If you have configured your QoreStor instance to use one or both of the Linux exports below, the QoreStor Watcher service will not be able to function properly as watcher communications will be sent to the proxy.

```
export http_proxy=http://<hostname>:<port>
export https_proxy=http://<hostname>:<port>
```

To ensure proper operation, you must also include the **no_proxy** export.

```
export no_proxy="localhost, 127.0.0.1"
```

i **NOTE:** In order for these configurations to persist after a system reboot, they must be entered into `/etc/environment`.

Logging onto the system GUI for the first time

To log on to the QoreStor GUI for the first time, complete the following steps

1. In a supported web browser, enter:
 - `https://<YourQoreStorServerName>:5233`
2. In the **Username** field, type **admin**, and in the Password field, type **St0r@ge!** and then click Log in or press **<Enter>**. You will be notified to change the password if you are still using the default password.

NOTE: Please change the admin password before logging in.

Your logon username is displayed at the top of the page in the right corner.

Using the QoreStor UI


QoreStor offers a web-based user interface that you can use to configure, manage, and monitor your QoreStor system. You can use the QoreStor UI to perform tasks such as:

- Configure containers, storage groups, and cloud storage groups
- Add and manage replications
- View system information and monitor performance
- Manage user accounts
- Generate and download diagnostic bundles

You can access the QoreStor UI from any supported browser. Refer to [Logging onto the system GUI for the first time](#) for more information.

QoreStor UI Overview

The QoreStor UI consists of the Header pane, the Navigation pane, a Status pane, and the Operations pane.



The panes provide the functions and options described below:

| Pane | Description |
|------|-------------|
|------|-------------|

| | |
|-------------|--|
| Header pane | <p>Provides the following items:</p> <ul style="list-style-type: none"> • Contact Us icon - information for contacting Quest sales and Support. • QoreStor Alerts - lists any Alerts on the QoreStor system. • Current user - displays the current user account, provides the option to log out, and provides the option to switch the UI to the Light theme. |
|-------------|--|

| | |
|-------------|---|
| Status pane | <p>Provides the following items:</p> <ul style="list-style-type: none"> • Version - displays the version of QoreStor • System Status - displays the status of the QoreStor. |
|-------------|---|

| | |
|-----------------|--|
| Navigation pane | <p>Provides navigation options to:</p> <ul style="list-style-type: none"> • Dashboard • Containers • Local Storage • Cloud Storage • Replications • System • Diagnostics • Users • Events |
|-----------------|--|

- Management

Operations pane Displays the data and dialogs appropriate for the chosen navigation option.

Using the QoreStor command line interface

QoreStor includes a custom shell implementation that simplifies command line access. As a superuser, you can use the command `/opt/qorestor/bin/setup_qs_user <username>` to configure an existing shell user with sudo privileges to run QoreStor CLI commands. You can access the QoreStor command line through a remote access program (such as PuTTY).

i | **NOTE:** The user created with `/opt/qorestor/bin/setup_qs_user<username>` cannot run any other commands with sudo, only QoreStor CLI commands.

To access the QoreStor command line

1. Using a remote access program, connect to your QoreStor server.
2. Log in using the superuser credentials or the credentials of the user configured for QoreStor CLI access using `setup_qs_user`.
3. Enter the desired command at the prompt using one of the following command styles:
 - If logged in as the superuser, use the full path; for example, `/opt/qorestor/bin/system -show` or `/opt/qorestor/bin/qs_help`.
 - If logged in as a user configured by `setup_qs_user`, type `system -show` or `help`.

i | **NOTE:** Refer to the *QoreStor Command Line Reference Guide* for more information on the QoreStor CLI.

Using the QoreStor REST API

QoreStor provides a REST API in order to provide an efficient method for application management or integration. For more information on using the REST API, see the sections below.

API Overview

All API requests should be sent to **https://<QoreStor host>:5233/api**. With the exception of the Authentication request, all requests should carry the HTTP header `Authorization: Bearer <token>`.

API calls will generate one of the response types below. Refer to the API documentation for the specific request for more information on expected responses.

- 200 - OK
- 400 - Bad request

- 401 - Unauthorized
- 404 - Not found
- 500 - Internal server error

Authentication

Authentication to the API is accomplished using the username and password for an existing QoreStor account. Once logged in, a token will be generated that can be used to authenticate API commands.

To use API authentication:

1. Run the **login** request using the username and password for the desired QoreStor account as well as the IP or domain name of your QoreStor server.


```
curl -X POST "https://QoreStor.host.com:5233/api/auth/login" -H "accept: application/json" -H "Content-Type: application/json" -d "{ \"name\": \"TestUser\", \"password\": \"TestUserPassword\"}"
```
2. In the request response, find and record the value for **JWT_Token**.
3. For subsequent API requests, ensure that the header contains the value **Authorization: Bearer <token>**. You may also use this token to authorize API requests via the documentation page. To do so,
 - At the top of the API documentation page, click **Authorize**.
 - In the **Value** field, enter the token.
 - Click **Authorize**.

i | **NOTE:** The JWT_Token included in the response remains valid for a duration of 20 minutes. If a follow-up request is made after 20 minutes from the last response, a new log in request must be submitted.

API Documentation

The detailed documentation for the QoreStor API is hosted on your QoreStor server. To view the documentation, enter the URL below into a browser.

https://<QoreStor Server>:5233//api/docs/index.html

i | **NOTE:** In the example above, replace <QoreStor Server> with the FQDN or IP address of your QoreStor server.

The API documentation contains the following:

- An **Authorize** button to enable authentication.
- A list of available API resource categories. Each list is expandable to reveal the resources in each category.
- Each resource listing can be expanded to view the resource description, parameters and responses, and a **Try it Out** button that can be used to run the API call with entered parameters.

To view and test an API

i | **NOTE:** The example below uses the **GET containers** request. The steps are similar for other requests.

1. Before testing the API, authenticate as described in [Authentication](#).
2. In the resource list, expand **Containers**.
3. Find the Get Containers request (GET/v1/Containers) and click to expand.
4. Click **Try it Out**. Note that the parameters fields are now enabled.
 **NOTE:** Parameters for Get Containers are optional. Some requests have required parameters.
5. Optionally, enter the parameters as described below:
 - **Storage Group** - The name of the storage group to which the containers belong
 - **Type** - required container type (cloud, link, or tier)
6. Click **Execute**.

Understanding RDA immutability

RDA immutability provides protection from overwrites and deletes on backup files. This technology is present by default for RDS containers, but not all backups are protected by default. Backup applications sending data using the RDA protocol to QoreStor define whether the backup data should be immutable and the time period the data will remain immutable. After you set immutability on the backup data, you cannot modify or delete the backup data from the RDA container until the immutable time period expires. For the current list of supported backup applications that take advantage of this technology, see the *Quest QoreStor Interoperability Guide*.

Understanding EDM immutability

EDM (hardened repository) immutability protects backup files by preventing overwrites and deletions. Veeam sets up an immutability period for backup files that are written to a hardened repository. The files are protected since they are unchangeable during this time.

Furthermore, the same immutability period is applied if the repository is linked to the cloud, locking the backup files both on-prem and in the cloud.

Configuring QoreStor settings

In the QoreStor GUI, you can easily view and configure system settings such as, active directory, system date and time, expansion shelf enclosures, licenses, networking, schedules for system operations, SSL certificates, storage groups, and users.

Licensing QoreStor

QoreStor offers a backend capacity licensing model to allow for simple integration with other Quest Data Protection products.

- **Standalone license** - QoreStor is licensed by the amount of backend capacity required. Standalone licenses are available as either **perpetual** licenses (with no expiration), or **term** licenses, which expire after a specified period of time.

Perpetual and term QoreStor licenses are additive, meaning that if you purchase a 5 TB license now, and a 10 TB license in the future, you will have 15 TB total capacity. However, all term licenses installed on a QoreStor server must share the same expiration date.

i | **NOTE:** Licenses for QoreStor are specific to the QoreStor server. When installing a license, the System ID for your QoreStor server is required. You can obtain the System ID with the command **system --show | grep "System ID"**

Evaluating QoreStor

QoreStor offers two methods for evaluation:

- **Default installation** - If no license is installed, QoreStor defaults to a no-cost, 1 TB capacity installation supported by the QoreStor Community. This option requires no license and does not expire. If a license is applied to a server running in this mode, the free 1 TB is **not** added to the purchased license capacity.
- **Full capacity trial** - available on the Quest Software Trial site, which provides a 30-day evaluation license for up to 360 TB or 512 TB (if MSP support is enabled) and access to Quest Support. After the evaluation period has expired, the QoreStor server will operate in Manual Intervention mode until a license is applied. To use QoreStor beyond that time frame, you will need to purchase a perpetual standalone license. If a longer trial period is required, please contact Quest Sales.

If you have purchased a standalone license, you can install it using the **system --license** command, as described in the *QoreStor Command Line Reference Guide*.

i | **NOTE:** When ordering a license, the System ID for your QoreStor server is required. You can obtain the System ID with the command **system --show | grep "System ID"**

Viewing your license configuration

You can view QoreStor license configuration information through the QoreStor GUI or the command line interface.

To view the current license configuration in the GUI

1. In the navigation menu, click **System** to expand the menu, then click **License**.
2. Information about the configured license is displayed.

To view the current license configuration via the command line, use the command

```
system --show [--license] [--verbose]
```

Installing a license

You can add a license to QoreStor through either the QoreStor GUI or the command line interface.

To install a license:

1. In the navigation menu, click **System** to expand the menu, then click **License**.
2. Click **Upload License**.
3. Click **Upload License** and select the license file.
4. Click **Apply**.

i **NOTE:** You may also install a QoreStor license from the command line interface using the command:

```
system --license [--show] [--verbose] [--validate] [--file <path>] [--add] [--file <path>]
```

Refer to the *QoreStor Command Line Reference Guide* for more information.

NetVault Plus Data Copy licenses

NetVault Plus Data Copy licenses are term licenses for QoreStor installations with limited protocol support. Data Copy licenses can only be installed on QoreStor systems that have RDA protocol support, but do not support CIFS, NFS, Object (S3), OST, or VTL (NDMP, iSCSI) connections. Data Copy licenses are functionally identical to QoreStor term licenses. These licenses are specific to an individual QoreStor server (based on system ID) and they have defined expiry dates.

Preparing QoreStor for a NetVault Plus Data Copy License

Prerequisite:

Before installing a Data Copy license ensure that the QoreStor is placed in Data-Copy mode.

To place a QoreStor in Data-Copy mode,

To place a QoreStor as a superuser, execute the following command:

```
/opt/qorestor/bin/data-copy-mode --enable
```

i **NOTE:** The above command will not enable Data-Copy mode if there are existing CIFS, NFS, OST, NDMP, iSCSI, EDM, or Object (S3) connections. These connections will need to be removed before Data Copy mode can be enabled.

As part of enabling Data-Copy mode, the command will remove support for the above protocols and restart QoreStor services.

To verify that QoreStor is in Data-Copy mode, use the `system --show` command to display the status of Data-Copy mode:

```
# system --show
System Name           : qorestor-server
Current Time          : Fri Jan  3 20:11:14 2025 UTC
System ID             : fb9b26426aa9d8a506c2e953a50f94c4
Product Name         : QoreStor
Version              : 7.5.0
Build                : 212
Repository location  : /QSdata/ocaroot
Metadata location    : /QSmetadata/qs_metadata
Dictionary type      : Cloud-Optimized
Data-Copy mode       : Enabled
Anomaly Detection    : Enabled
Anomaly Detection Interval (min) : 30
Anomaly Detection Metric : os-auth,ui-auth,proto-auth,load-avg,diskio
System State         : Operational Mode
Reason               : Filesystem is fully operational for I/O.
Configuration Server : RUNNING Jan  3 20:09:50
Filesystem Server    : RUNNING Jan  3 20:09:50
Health Monitor       : RUNNING Jan  3 20:09:46
Filesystem Checker   : STOPPED
SecureConnect Server : RUNNING Jan  3 20:09:49
UI                   : RUNNING Jan  3 20:09:50
Policy Manager Daemon : RUNNING Jan  3 20:10:41
Disaster Recovery Daemon : STOPPED
Anomaly Detection Service : RUNNING Jan  3 20:09:50
```

Installing a NetVault Plus Data Copy License

Refer the above section for the steps to install a Data Copy license.

i **NOTE:** Data Copy licenses can only be installed on QoreStor systems that are already in Data-Copy mode. Conversely, QoreStor perpetual or term licenses cannot be installed on QoreStor systems that are in Data-Copy mode.

Configuring SAML

SAML 2.0-based authentication is how QoreStor supports single sign-on (SSO) capabilities through an external identity provider (IdP). QoreStor supports only service provider-initiated login. QoreStor SAML configuration is compatible with the following IdPs:

- Azure AD
- OneLogin
- Okta

Accessing the SAML configuration settings in QoreStor

To access the SAML configuration settings in QoreStor

1. To display the Users configuration page, in the navigation menu, click **Users**.
2. Click **Configure SAML**.
3. Provide the URLs displayed on the SAML Configuration page to your chosen IdP:
 - **Service Provider EntityId**. This URL serves as a unique identifier that represents the particular QoreStor server to the IdP.
 - **Service Provider MetaData URL**. This URL provides an endpoint on the QoreStor Server that QoreStor uses to furnish keys and additional SAML endpoints to the IdP.
 - **Service Provider ACS (Assertion Consumer Services) URL**. This URL provides an endpoint on the QoreStor Server that the IdP uses to send its authentication response after the IdP has successfully authenticated a user attempting to login to QoreStor. The IdP uses this URL to redirect the user's browser session back to the QoreStor Server to allow authenticated access to the QoreStor Server.
4. To complete the SAML configuration, obtain the Identity Provider Metadata URL from the IdP and enter it in the text box.
This URL is required. It contains the URL of the metadata endpoint provided by the IP to furnish keys and additional SAML endpoints to the QoreStor Server.
5. Click **Save**.
6. Restart the QoreStor UI using the following command:


```
/opt/qorestor/bin/storage-server-services/ocaui restart
```

i NOTE: For SAML redirection to succeed, the QoreStor hostname should be resolvable from the Client browser.

For more information about configuring SAML for QoreStor, see the *QoreStor SAML Configuration Guide*.

Registering an SSO user with the QoreStor Server

To register an SSO user with the QoreStor Server

1. To display the Users configuration page, in the navigation menu, click **Users**.
2. To display the Add User sidebar, click **Add User**.
3. From the **Authentication Type** drop-down menu, select **SAML** as the authentication type for the new user.

4. For **User (Email)**, enter the email address of the user that was configured in the IdP for access to the QoreStor Server.
The value provided must be of a valid email address format; for example, john.doe@example.com).
5. From the **Roles** drop-down menu, select either the **Administrator** role or **Monitor** role.
6. Optionally, enter the following details:
 - **Full Name**
 - **Phone**
 - **Description**
7. To add the registered user account with the QoreStor Server, click **Save**.
The QoreStor UI Login page now includes a **SAML Sign In** option.

Logging in to the QoreStor UI using SAML

To log in to the QoreStor UI using SAML

1. From the QoreStor UI Login page, click **SAML Sign In**.
The browser redirects you to the configured IdP's login page.
2. On the IdP login page, enter the required IdP user credentials.
3. Optionally, if you configured MFA for your account with the IdP, enter the requested second factor token or provide approval through a mobile application.
After the IdP authenticates the credentials, the browser session redirects to the QoreStor UI. If the authentication was successful, the QoreStor UI login page briefly displays before the browser session redirects to the QoreStor UI Dashboard.
You are now logged in to the QoreStor Server.

Configuring an SSL certificate for your QoreStor system

For additional security, you can replace the self-signed, factory-installed certificate with another SSL certificate, for example, with one that is signed by a third-party CA. Once you have obtained your signed certificate and private key, you can install them by using the QoreStor UI or CLI. Only one certificate can be installed on a QoreStor system at any given point in time. The same certificate will be used for HTTPS access to object containers.

Installing an SSL certificate

To install an SSL certificate, complete the following steps:

1. In the navigation menu, click **Dashboard**.
2. In the System Information pane at the bottom, click **SSL Certificate**.
3. Click **Upload Certificate**.
4. Select the SSL certificate on your system that you want to install.

5. Click **Upload Key** and select your private key.
6. Click **Upload**.

Configuring Active Directory settings

You can easily join the QoreStor to your Microsoft Active Directory Services (ADS) domain. This topic describes how to configure Active Directory (AD) settings for the QoreStor system, which requires that you direct your QoreStor system to join or leave a domain that contains a Microsoft Active Directory Service (ADS). Instructions are provided below to join an ADS domain or to leave an ADS domain. When you join QoreStor to an ADS domain, this disables the Network Time Protocol (NTP) service and instead uses the domain-based time service.

To configure QoreStor for a domain using ADS, complete the following steps:

1. In the left navigation menu, click **System > Active Directory**.
2. Click **Join Domain**.
3. Enter the following AD logon information:
 - **Domain**—Enter a fully qualified domain name for the ADS; for example, AD12.acme.com. *(This is a required field.)*
 - NOTE:** Supported domain names are limited to 64 characters in length and can only consist of a combination of A-Z, a-z, 0-9, and three special characters: a dash (-), a period (.), and an underscore (_).
 - **Organization**—Enter a valid organizational name that meets the organization name guidelines for the ADS. *(This is an optional field.)*
 - **Username**—Enter a valid user name that meets the user name guidelines for the ADS. *(This is a required field.)*
 - NOTE:** Supported user names are limited to 64 characters in length and can only consist of a combination of A-Z, a-z, 0-9, and three special characters: a dash (-), a period (.), and an underscore (_).
 - **Password**—Enter a valid password that meets the password guidelines for the ADS. *(This is a required field.)*
4. Click **Join**.
5. To leave a domain, find the domain on the Active Directory page and click **Leave Domain**.
 - a. In the Leave Active Directory pane, enter the username and password for the ADS domain.
 - b. Click **Leave**.

Adding a login group to an ADS domain

NOTE: This is an optional configuration, and only necessary if Domain users in a login group are expected to authenticate to the QoreStor UI.

After you configure your QoreStor within the same ADS domain, you must ensure that a login group exists and add it to the domain. Adding a login group is only possible when the QoreStor system is joined to a domain. Also, you must be logged in as a domain user that is part of an enabled login group.

To add a login group in an ADS domain, complete the following steps:

1. On the left navigation menu, click **System > Active Directory**.
2. On the **Active Directory** page, find the domain pane and click **Add Login Group**.
3. In the **Add Login Group** pane, type the name of the login group including the domain name; for example, *Domain\Domain Admins*. If your login group name contains spaces, you must not enclose it in quotation marks. (This differs from the equivalent CLI command.)
4. Click **Add Login Group** to add the login group.

i **NOTE:** Changes made to the login group take effect on the next log in attempt (unlike Windows ADS, no active checking is done on the group).

Removing a login group

To remove a configured ADS login group from the QoreStor system, complete the following steps.

1. On the left navigation menu, click **System > Active Directory**.
2. On the **Active Directory** page, find the domain pane and click **Remove Login Group**.
3. When prompted to confirm, click **OK**.

i **NOTE:** Changes made to the login group take effect on the next log in attempt (unlike Windows ADS, no active checking is done on the group).

Securing QoreStor server root login

QoreStor UI provides a way to disable "root" user login over SSH to the QoreStor server. This is provided in QoreStor UI for convenience to secure the system.

After you log in to the QoreStor UI, navigate to the **System -> Configuration** page. Check the status of the SSH service. If it is enabled, you can disable root login over SSH.

■ WARNING: Ensure that there is an alternate way to re-enable root login (over SSH) in case the QoreStor UI is not accessible. This is required to ensure that the root is not permanently locked out of the server. Another way to access the QoreStor server in such condition could be to log in via the system console, or via SSH as a different user that has superuser privileges either by default, or with sudo.

i **NOTE:** When the status of 'root login over SSH' is shown as 'Unknown' it means the SSH access is either restricted (i.e. password authentication is prohibited) or QoreStor is not able to determine the status of the SSH server.

Enabling FIPS 140-2 support

FIPS 140-2, when installed on systems running RHEL/Oracle Linux/AlmaLinux/Rocky Linux 9.x, QoreStor can make use of FIPS-compliant implementations of cryptographic routines provided by the OS vendor.

i | **NOTE:** Currently, QoreStor running on RHEL/Oracle Linux/AlmaLinux/Rocky Linux 8.x systems will not be impacted by turning on system-wide support for FIPS. Despite using techniques permitted by FIPS 140-2, QoreStor encryption does not employ FIPS-compliant implementations of those algorithms.

Enabling support for FIPS 140-2:

To enable FIPS system-wide support, execute the following commands with the Superuser credentials:

```
fips-mode-setup --enable
update-crypto-policies --set FIPS
reboot
```

After a successful reboot, the QoreStor will automatically start using the FIPS-compliant implementations.

To check if the QoreStor is operating in FIPS mode, examine the `/var/log/oca/ocafsd.log` file and check for the following entry in the log:

```
NOTICE: FIPS status: enabled
```

Limitations with FIPS 140-2

CIFS, NFS, Object Containers, NDMP, and iSCSI protocols are not supported when the QoreStor is executing in a FIPS mode. Additionally, while cloud-tier and archive-tier support is available in FIPS mode, currently only Amazon-S3 and Azure are supported as cloud providers.

Understanding system operation scheduling

By scheduling system operations, you can optimize your system resources and achieve the best possible QoreStor performance. The most important thing to remember when scheduling critical QoreStor operations is that you want to ensure that you perform each of these operations at a time when it will not overlap or interfere with the running of any of the other key system operations.

You should carefully plan and schedule time periods in which to perform the following critical system operations:

- Data ingests (which are dependent upon your usage of your DMA(s))
- Replication
- System cleaner (space reclamation)

i | **NOTE:** Performing scheduled disk space reclamation operations are recommended as a method for recovering disk space from system containers in which files were deleted as a result of deduplication.

In QoreStor, the main goal in planning and scheduling operations should be to run the Cleaner and Replication operations at times when they do not overlap or interfere with other important system operations. You want to make sure that by properly scheduling and planning, your system can perform each of these key operations independent of the other.

The best practice is to run these two operations during non-standard business hours, so that they do not conflict with any of your other backup or ingest operations.

i | **NOTE:** By default, QoreStor is configured to run Cleaner operations daily between 1:00 P.M and 6:00 P.M.

The **Cleaner** schedule can be viewed and edited on the **Cleaner** page of the QoreStor GUI (**Local Storage > Cleaner**), or via the QoreStor command line interface, using the **schedule** command:

```
schedule --show --cleaner
```

The **Replication** schedule can be viewed and edited via the **Replications** page of the QoreStor GUI, or via the QoreStor command line interface using the **schedule** command:

```
schedule --show --replication [--name] <name>
```

i | **NOTE:** For more information on the **schedule** command, refer to the *QoreStor Command Line Reference Guide*.

Configuring cleaner schedules

Performing scheduled disk space reclamation operations are recommended as a method for recovering disk space from the system. The recommended method is to schedule a time when you can run the Cleaner on your QoreStor system with no other planned processes running.


Even if no Cleaner schedule is set, but the system detects that there is disk space that can be reclaimed, the Cleaner process runs. However, the Cleaner will not start until the following conditions are met:

- it detects that there are no active data ingests,
- that two minutes of system idle time have elapsed since the last data file ingest was completed,
- and that the Replication process is not running (the Cleaner process runs as a lower system priority operation than the Replication process).

To schedule cleaner operations on your system using the GUI:

1. On the navigation menu, click **Local Storage** to expand the menu, then click **Cleaner**.
2. On the Cleaner page, click **Edit Schedule**.

The schedule lists a **Start Time** and **End Time** for each day of the week.



3. For each day of the week, click the time selector field. Select the **From** and **To** times to configure a window during which replication can run. Click **Set**.
4. Click **Save Schedule**.
5. Click **Submit**

i | **NOTE:** Even if no Cleaner schedule is set, but the system detects that there is disk space that can be reclaimed, the Cleaner process runs. However, the Cleaner will not start until the following conditions are met: it detects that there are no active data ingests, that two minutes of system idle time have elapsed since the last data file ingest was completed, and that the Replication process is not running

(the Cleaner process runs as a lower system priority operation than the Replication process).

To schedule cleaner operations on your system using the CLI:

i **NOTE:** The procedure below is a summary. Please refer to the *QoreStor Command Line Reference Guide* for detailed information on accessing the command line interface for your system as well as using the QoreStor commands.

i **NOTE:** Running the Cleaner while ingesting data reduces system performance. Ensure that you schedule the Cleaner to run when backup or replication is not in progress.

1. Access the QoreStor command line interface.
2. Use the QoreStor command line interface (CLI) to create and delete the cleaner schedule. The available commands are:

```
schedule --add --day <Day of the Week> --start_time <HH:MM> --stop_time  
<HH:MM> --cleaner  
schedule --delete --day <Day of the Week> --cleaner
```

For full details on running the cleaner schedule commands, help is available by entering:

```
schedule --help
```

Viewing cleaner status

On the **Cleaner** page, you can also view graphs showing Cleaner runtime and bytes processed. You may also use the **stats --system** command to view the cleaner status via the QoreStor command line interface.

Viewing cleaner statistics

To view additional detailed cleaner statistics, you can use the QoreStor CLI **stats --cleaner** command to show the following categories of statistics:

- Last Run Files Processed (number of files processed by Cleaner)
- Last Run Bytes Processed (number of bytes processed by Cleaner)
- Last Run Bytes Reclaimed (number of bytes reclaimed by the Cleaner)
- Last Run Start Time (indicates date and time last Cleaner process started)
- Last Run End Time (indicates date and time last Cleaner process ended)
- Last Run Time To Completion(s) (indicates the number of times that Cleaner process has successfully completed)
- Current Run Start Time (indicates date and time current Cleaner process started)
- Current Run Files Processed (number of files processed by current Cleaner process)
- Current Run Bytes Processed (number of bytes processed by current Cleaner process)
- Current Run Bytes Reclaimed (number of bytes reclaimed by the current Cleaner processed)
- Current Run Phase 1 Start Time (indicates date and time for start of current Cleaner process phase 1)

- Current Run Phase 1 Records Processed (lists the number of data records processed in current Cleaner process phase 1)
- Current Run Phase 1 End Time (indicates date and time for end of current Cleaner process phase 1)
- Current Run Phase 2 Start Time (indicates date and time for start of current Cleaner process phase 2)
- Current Run Phase 2 Records Processed (lists the number of data records processed in current Cleaner process phase 2)
- Current Run Phase 2 End Time (indicates date and time for end of current Cleaner process phase 2)
- Current Run Phase 3 Start Time (indicates date and time for start of current Cleaner process phase 3)
- Current Run Phase 3 Records Processed (lists the number of data records processed in current Cleaner process phase 3)
- Current Run Phase 3 End Time (indicates date and time for end of current Cleaner process phase 3)
- Current Run Phase 4 Start Time (indicates date and time for start of current Cleaner process phase 4)
- Current Run Phase 4 Records Processed (lists the number of data records processed in current Cleaner process phase 4)
- Current Run Phase 4 End Time (indicates date and time for end of current Cleaner process phase 4)

For more information about QoreStor CLI commands, see the *QoreStor Command Line Reference Guide*.

Configuring Secure Connect

The sections below contain information necessary for the proper configuration of Secure Connect. The procedures for configuring Secure Connect differ depending on your plug-in version.

- For OST and RDA plug-in version 4.1.0.265 or later
 - [Managing Secure Connect with OST or RDA plug-in 4.1.0.265 or later](#)
- For OST and RDA plug-in versions prior to 4.1.0.265
 - [Enabling Secure Connect for OST and RDA plug-ins prior to 4.1.0.265](#)
 - [Configuring Secure Connect properties](#)
- For Rapid NFS plug-in version 4.0.3310.0 or later
 - [Installing the Rapid NFS plug-in](#)
- For Rapid CIFS plug-in version 4.0.3233.1 or later
 - [Installing the Rapid CIFS plug-in](#)
- [Adding certificates for Secure Connect](#)

Enabling Secure Connect for OST and RDA plug-ins prior to 4.1.0.265

i **IMPORTANT:** The procedure below is for plug-in versions prior to 4.1.0.265. To enable or disable Secure Connect on plug-in version 4.1.0.265, refer to [Managing Secure Connect with OST or RDA plug-in 4.1.0.265 or later](#).

Secure Connect is enabled through the use of environmental variables on the client machine. No configuration is required on the QoreStor server.

To enable Secure Connect on a Windows client

1. On the client server, press **Win+R** to open the **Run** window.
2. Type **sysdm.cpl** and click **OK**.
3. Click the **Advanced** tab, then **Environment Variables**.
4. In the **System Variables** section, click **New**.
5. In the **Variable name** field, enter **SECURE_CONNECT**.
6. In the Variable value field, enter one of the following:
 - **0** - disables Secure Connect
 - **1** - Secure Connect is enabled, but QoreStor will failback to an unsecured connection if the Secure Connect server is unavailable.
 - **2** - Secure Connect is enabled. Connection will fail if Secure Connect server is unavailable.
7. Click **OK**, then **OK**.

i **IMPORTANT:** After enabling Secure Connect, you will need to change the BypassPorts configuration in the **sc_client.properties** file. Refer to [Configuring Secure Connect properties](#) for information.

i **IMPORTANT:** After enabling Secure Connect, you must restart the DMA application services.

To enable Secure Connect on a Linux client

1. At the command prompt on the client machine, enter the following command

```
echo 'export SECURE_CONNECT=<0|1|2>' >> ~/.bashrc
```

Where:

- **0** - disables Secure Connect
 - **1** - Secure Connect is enabled, but QoreStor will failback to an unsecured connection if the Secure Connect server is unavailable.
 - **2** - Secure Connect is enabled. Connection will fail if Secure Connect server is unavailable.
2. Log out of the QoreStor system, then log in.

i **IMPORTANT:** After enabling Secure Connect, you will need change the BypassPorts configuration in the **sc_client.properties** file. Refer to [Configuring Secure Connect properties](#) for information.

i **IMPORTANT:** After enabling Secure Connect, you must restart the DMA application services.

Configuring Secure Connect properties

Before using Secure Connect, ensure that the default port configuration is appropriate for your environment. The ports used by Secure Connect are:

- 9443 - this is the listening port. The Secure Connect server listens for connection requests on this port.
- 10011, 11000 and 9920 - These are the standard Secure Connect communication ports.

By default, the Secure Connect ports are bypassed, which will cause Secure Connect to failback to a normal, unsecured connection. Before using Secure Connect, the **BypassPorts** setting must be set to 0 to enable full communication.

Secure Connect properties can be configured through the **sc_client.properties** file located in the client installation directory.

To configure Secure Connect

1. In the client installation directory, open the **sc_client.properties** file with a text editor.
The default installation directory differs depending on the client type and the OS of the client machine. For example,
 - The RDA client on a Windows machine installs to `C:\Program Files\Quest\RDA\dynlib`
 - The NetVault on a Linux server installs to `/usr/netvault/dynlib/sc_client.properties`
2. Find the entry shown below


```
# A list of ports to be excluded from SecureConnect. Example: 9904,9921-9999,10011
#BypassPorts = 0
BypassPorts = 9920,10011,11000
```
3. Do one of the following:
 - Comment out the line **BypassPorts = 9920, 10011, 11000** by adding a # to the front, then remove the # from `BypassPorts = 0`
 - Delete the listed ports (9920, 10011, 11000) and replace with 0.
4. Save the file.

Managing Secure Connect with OST or RDA plug-in 4.1.0.265 or later

Unless manually disabled, Secure Connect is always running on the QoreStor server. Starting with QoreStor plug-in version 4.1.0.265, Secure Connect is enabled by default on the client machine. Review the sections below for the procedures to check Secure Connect status or disable and enable Secure Connect.

The commands below can be run both on the QoreStor server and the client machines. In both cases, the **sc_manager** command must be run from the directory that includes the **sc.client.properties** file. By default this is:

- On the QoreStor server `/opt/qorestor/bin`

i | **NOTE:** When configuring Secure Connect on the QoreStor server, the changes made are applicable only for container or optimized replication in which the QoreStor instance is a source.

- For client machines, this is the plug-in installation directory. For example, for NetVault:
 - Linux clients - `/usr/local/ocarda`
 - Windows clients - `C:\Program Files\Quest\RDA\Dynlib`

i | **IMPORTANT:** The procedures below use the **sc_manager** command which must be run by the **root** account.

Checking Secure Connect status

To check the status of Secure Connect

1. Run the command **sc_manager status** according to one of the methods below:

- Run **sc_manager** from the directory containing **sc_client.properties**.

```
sc_manager status
```

- Run **sc_manager** from any directory, using the path to the **sc_client.properties**.

```
/opt/qorestor/bin/sc_manager status --property /opt/qorestor/bin/sc_client.properties
```

```
SecureConnect.enabled = true
```

2. The status of Secure Connect will be displayed:

```
SecureConnect.enabled = true
```

Disabling Secure Connect

To disable Secure Connect

1. Run the **sc_manager disable** command as described below:

- Run **sc_manager** from the directory containing **sc_client.properties**.

```
sc_manager disable
```

- Run **sc_manager** from any directory, using the path to the **sc_client.properties**.

```
/opt/qorestor/bin/sc_manager disable --property /opt/qorestor/bin/sc_client.properties
```

2. The status of Secure Connect will be displayed as confirmation:

```
SecureConnect.enabled = false
```

3. After disabling Secure Connect, a service restart must be performed.

- If you disabled Secure Connect on the QoreStor server, restart the QoreStor services using the commands:

```
/opt/qorestor/bin/ctrlrpc -p 9901 node.stop  
/opt/qorestor/bin/ctrlrpc -p 9901 node.start
```

- If you disabled Secure Connect on the client machine, services of the DMA application on that machine need to be restarted.

Enabling Secure Connect

To enable Secure Connect

1. Run the `sc_manager enable` command as described below:

- Run **sc_manager** from the directory containing **sc_client.properties**.

```
sc_manager enable
```

- Run **sc_manager** from any directory, using the path to the **sc_client.properties**.

```
/opt/qorestor/bin/sc_manager enable --property /opt/qorestor/bin/sc_client.properties
```

2. The status of Secure Connect will be displayed as confirmation:

```
SecureConnect.enabled = true
```

3. After enabling Secure Connect, a service restart must be performed.

- If you enabled Secure Connect on the QoreStor server, restart the QoreStor services using the commands:

```
/opt/qorestor/bin/ctrlrpc -p 9901 node.stop
```

```
/opt/qorestor/bin/ctrlrpc -p 9901 node.start
```

- If you enabled Secure Connect on the client machine, services of the DMA application on that machine need to be restarted.

Adding certificates for Secure Connect

The QoreStor Secure Connect feature requires custom certificates on both the client and QoreStor server machine.



NOTE: The certificates on both the client machine and QoreStor server must be from the same certificate authority.

Adding a Secure Connect certificate - Windows Client

1. Prepare custom certificates chain and install them to the certificate store using the Microsoft Management Console (MMC) **Certificates** snap-in.
 - a. Install the **Root** certificate to **Trusted Root Certification Authorities**.
 - b. If necessary, install the **Intermediate** certificate to **Intermediate Certification Authorities**.
 - c. Install the **Server** certificate to **Personal**.
2. In the client installation directory, open the **sc_client.properties** file with a text editor.

3. Edit the entries below:

- **openssl.client.caConfig** - The path to the file of the trusted root certificate or directory containing the trusted root certificates chain. For specifying directory path, be sure that directory contains certificates in the PEM format and symbolic links to the certificate files, created by the **c_rehash** utility.
- **openssl.client.certificateFile** - The path to the file containing the server's or client's certificate in PEM format.
- **openssl.client.privateKeyFile** - The path to the file containing the private key for the certificate in PEM format.

Example in case there is the chain of 3 certificates (root, intermediate, server), private key and they are located in the same directory as **sc_client.dll**:

- `openssl.client.caConfig = ${application.configDir}`
- `openssl.client.certificateFile = ${application.configDir}server-certificate-name.pem`
- `openssl.client.privateKeyFile = ${application.configDir}privat-key-name.key`

Example in case there is the chain of 2 certificates (root, server), private key and they are located at `C:\certificates`:

- `openssl.client.caConfig = C:\certificates\root-certificate-name.pem`
- `openssl.client.certificateFile = C:\certificates\server-certificate-name.pem`
- `openssl.client.privateKeyFile = C:\certificates\privat-key-name.key`

4. Make **c_rehash** for the certificates:

- a. Download perl from <https://www.activestate.com/ActivePerl>.
- b. Download the perl script **c_rehash**, stored inside OpenSSL (<https://wiki.openssl.org/index.php/Binaries>)
- c. Set the new **openssl** environment variable with the path to openssl.
- d. Run the command prompt.
- e. Use **perl.exe** with **path_to_the_c_rehash** and **path_to_the_cert_dir** arguments (e.g. `perl.exe C:\<path to the c_rehash> C:\<path to the certificates directory>`)

5. When Secure Connect is used with any DMA - restart DMA services.

i | **NOTE:** If certificate validation fails, the connection between client and server will fail back to a normal connection.

Adding a Secure Connect certificate - Linux Client and QoreStor server

1. Prepare custom certificates chain
2. Place the certificate to be trusted (in PEM format) in `/etc/pki/ca-trust/source/anchors/` and run `sudo update-ca-trust` at the prompt.
If the certificate is in OpenSSL's extended BEGIN TRUSTED CERTIFICATE format, place it in `/etc/pki/ca-trust/source` and run `sudo update-ca-trust`.

3. Make **c_rehash** for the certificates:
 - a. Install the **openssl-perl** package.
 - b. Run `c_rehash <path-to-the-folder-with-certificates>`.
4. In the client installation directory, open the **sc_client.properties** file with a text editor.
5. Edit the entries below:
 - a. **openssl.client.caConfig** - The path to the file of the trusted root certificate or directory containing the trusted root certificates chain. For specifying directory path, be sure that directory contains certificates in the PEM format and symbolic links to the certificate files, created by the `c_rehash` utility.
 - b. **openssl.client.certificateFile** - The path to the file containing the server's or client's certificate in PEM format.
 - c. **openssl.client.privateKeyFile** - The path to the file containing the private key for the certificate in PEM format.

Example in case there is the chain of 3 certificates (root, intermediate, server), private key and they are located in the same directory with `sc_client.so`, server side:

 - `openssl.server.caConfig = ${application.configDir}`
 - `openssl.server.certificateFile = ${application.configDir}server-certificate-name.pem`
 - `openssl.server.privateKeyFile = ${application.configDir}privat-key-name.key`

Example in case there is the chain of 2 certificates (root, server), private key and they are located at `/usr/certificates` on the client machine:

 - `openssl.client.caConfig = /usr/certificates/root-certificate-name.pem`
 - `openssl.client.certificateFile = /usr/certificates/server-certificate-name.pem`
 - `openssl.client.privateKeyFile = /usr/certificates/privat-key-name.key`
6. When Secure Connect is used with any DMA - restart DMA services.

i | **NOTE:** If certificate validation fails, the connection between client and server will fail back to a normal connection.

Enabling MultiConnect

Before using MultiConnect, ensure that the default port configuration is appropriate for your environment. The port used by MultiConnect is:

- 11000 - This is the standard MultiConnect communication port for backup.
- 9920 - This is the standard MultiConnect communication port for managed replication.

To enable Secure Connect on a Windows client

1. On the client server, press **Win+R** to open the **Run** window.
2. Type **sysdm.cpl** and click **OK**.

3. Click the **Advanced** tab, then **Environment Variables**.
4. In the **System Variables** section, click **New**.
5. In the **Variable name** field, enter **REMOTE_CLNT_MAX_CONNS**.
6. In the **Variable value** field, enter one of the following:
 - **4** - establishes 4 connections.
 - **8** - establishes 8 connections.
 - **16** - establishes 16 connections.
7. Click **OK**, then **OK**.
8. Restart the DMA services for the change to take effect.

To enable MultiConnect on a Linux client

1. At the command prompt on the client machine, enter the following command

```
echo 'export REMOTE_CLNT_MAX_CONNS=<4|8|16>' >> /etc/profile
```

Where:
 - **4** - establishes 4 connections.
 - **8** - establishes 8 connections.
 - **16** - establishes 16 connections.
2. Restart the DMA services for the change to take effect.

To enable MultiConnect between QoreStor servers

1. At the command prompt on the source machine, enter the following command

```
echo 'export REPL_CLNT_MAX_CONNS=16' >> /etc/oca/oca.cfg
```
2. Restart the **ocards** service for the change to take effect.

Bandwidth throttling

Bandwidth throttling, also known as traffic shaping or bandwidth management, refers to the deliberate slowing down or limiting of the internet speed. This technique is used to ensure network stability, prevent congestion, and distribute bandwidth resources more equitably. For example, bulk data backup may be provided the limited bandwidth to make way for time-critical application traffic.

QoreStor offers a bandwidth throttling tool between a pair of QoreStor systems for replication and op-dup. If enabled, it works on the TCP ports: 9904, 9911, 9915, and 9916. You can use the bandwidth throttling tool through the user interface (GUI) or with the command line interface (CLI). Bandwidth throttling applies to both, container replication and managed replication for RDA and OST protocol containers.

Adding a throttling schedule from UI

On the QoreStor GUI, the Bandwidth throttling page displays current information about bandwidth throttling schedules and default bandwidth set for the QoreStor system.

To add a throttling schedule from the UI

Execute the following steps to add a throttling schedule from the UI:

1. In the navigation menu, click **System** to expand the menu, then click **Bandwidth throttling**. The configured bandwidth throttling schedule is displayed.
2. Click **Add a Target System with Throttling** to add a new target with a schedule.
3. In the pop-up window, enter the following details:
 - a. **Target System**: Enter the target machine's IP, host name, or FQDN.
 - b. **Out of Schedule Bandwidth**: The default value is "Full Speed." To specify a custom value, select the **Custom** option and enter the desired value in mbps for out-of-schedule bandwidth.
 - c. Click **Save**.
 - d. Optionally, to add schedules with specific days and times, click **Add Schedule Event** and enter/select the desired values:
 - a. Select the **Start Day** (note: the week starts on Sunday and ends on Saturday).
 - b. Enter the **Start Time** (00:00 to 23:59)
 - c. Select the **End Day**.
 - d. Enter the **End Time**.
 - e. Enter the **Bandwidth** in Mbps.
 - f. Click **Add**.
Once successfully added, the schedule entry will appear in the grid. You can add multiple schedules by repeating the above steps.
 - g. Optionally, click **Remove** to delete an entry from the schedules or click **View/Edit** to view/modify the schedule entry.

i | **NOTE:** End Day and End Time should be ahead of Start Day and Start time.

CLI: Use `qs_bw_throttle` CLI to configure, view, or modify bandwidth throttling for replication and op-dup. Please refer the QoreStor CLI Reference guide for more information.

Viewing/editing a throttling schedule

To view/edit a throttling schedule

Execute the following steps to view/edit a throttling schedule:

1. In the navigation menu, click **System** to expand the menu, then click **Bandwidth throttling**. The configured bandwidth throttling schedule is displayed.
2. Click **View** on the schedule entry that you wish to view the entered values.
3. Click **Edit** on the schedule entry that you wish to modify.

4. In the pop-up window modify the **Out of Schedule Bandwidth** values if you wish to modify the entry.
5. To update the schedules:
 - a. To delete a schedule entry, click **Remove**.
 - b. To add a new schedule, click **Add Schedule Event**.

Removing/deleting a throttling schedule

To delete a throttling schedule

Execute the following steps to delete a throttling schedule:

1. In the navigation menu, click **System** to expand the menu, then click **Bandwidth throttling**. The configured bandwidth throttling schedule is displayed.
2. Click **Remove** on the schedule entry that you wish to delete.
3. Confirm the action to delete the entry.

Configuring and using Rapid NFS and Rapid CIFS

Rapid NFS and Rapid CIFS enable write operation acceleration on clients that use NFS and CIFS file system protocols. Similar to OST and RDS, these accelerators allow for better coordination and integration between QoreStor backup, restore, and optimized duplication operations with Data Management Applications (DMAs) such as CommVault, EMC Networker, and Tivoli Storage Manager. For the current list of supported DMAs, see the *QoreStor Interoperability Guide*.

Rapid NFS is a new client file system type that ensures that only unique data is written to QoreStor. It uses user space components and file system in user space (FUSE) to accomplish this. Metadata operations such as file creates and permission changes go through the standard NFS protocol, whereas write operations go through Rapid NFS.

Rapid CIFS is a Windows-certified filter driver that also ensures that only unique data is written to QoreStor. All chunking and hash computations are done at the client level.

i **NOTE:** The supported DMAs listed in the *QoreStor Interoperability Guide* are the DMAs that have been **tested and qualified** with Rapid NFS and Rapid CIFS. You can use Rapid NFS and Rapid CIFS with other DMAs, but those products have not been tested and qualified with Rapid NFS or Rapid CIFS.

Rapid NFS and Rapid CIFS benefits

When Rapid NFS and Rapid CIFS are used with QoreStor they offer the following benefits:

- Reduce network utilization and DMA backup time
 - Chunk data and perform hash computation on the client; transfer chunked hash files on the back-end
 - Reduce the amount of data that must be written across the wire
- Improve performance

- Support DMAs such as CommVault, EMC Networker, and Tivoli Storage Manager. For the current list of supported DMAs, see the *QoreStor Interoperability Guide*.
- Compatible with existing NFS and CIFS clients — just need to install a plug-in (driver) on the client
 - Can use Rapid NFS and Rapid CIFS to accelerate I/O operations on any client — including a client that uses home-grown backup scripts
 - Can service multiple and concurrent media server backups

Best practices: Rapid NFS

This topic introduces some recommended best practices for using Rapid NFS operations with QoreStor.

- Containers must be of type NFS/CIFS

RDA containers cannot use Rapid NFS. If you have existing NFS/CIFS containers, you do not need to create new containers to use Rapid NFS; you can install the plug-in (driver) to existing clients.

- The Rapid NFS plug-in (driver) must be installed on client systems

After the plug-in is installed, write operations will go through Rapid NFS while metadata operations such as file creates and permission changes will go through the standard NFS protocol. Rapid NFS can be disabled by uninstalling the plug-in.

- Markers must be set on the client, not in the QoreStor GUI
- If you are using a DMA that supports a marker, should explicitly set it. Your containers should have the marker type of None until you set the marker using the Mount command on the client (after installing the Rapid NFS plug-in).

- For existing containers, re-set the marker by doing the following:

For example, if you wanted to set the CommVault marker (cv):

```
mount -t rdnfs 10.222.322.190:/containers/backup /mnt/backup -o
marker=cv
```

Mount command usage:

```
rdnfs [nfs mount point] [roach mount point] -o marker=[marker]
```

where:

nfs mount point = Already mounted nfs mountpoint

roach mount point = A new mount point

marker = appassure, arcserve, auto, cv, dump, hdm, hpdp, nw, or tsm

- Your QoreStor system must meet the minimum configuration

Rapid NFS is available on a QoreStor system and a client with a minimum of 4 CPU cores running at a minimum of 4 GHz cumulative processing power and 2 GB memory. Kernels must be 2.6.14 or later. For a list of supported operating systems, see the *QoreStor Interoperability Guide*. If you update your operating system, you must update your Rapid NFS plug-in as well. Updates are available on the Quest Support site.
- Rapid NFS is stateful

If the QoreStor system goes down, the connection will end. DMAs will restart from the last checkpoint.
- Rapid NFS and passthrough mode

If Rapid NFS mode fails for any reason, QoreStor falls back to regular NFS mode automatically.

- Rapid NFS performance considerations
When using Rapid NFS on your client, Quest recommends that you do not run other protocols to the QoreStor system in parallel, as this will adversely affect your overall performance.
- Rapid NFS acceleration constraints
 - Rapid NFS does not support:
 - Direct I/O memory
 - Mapped files
 - File path size greater than 4096 characters
 - File write locks across clients

i | **NOTE:** If the client and server do not have the same times, the times seen will not match typical NFS behavior due to the nature of file system in user space (FUSE).

Best practices: Rapid CIFS

This topic introduces some recommended best practices for using Rapid CIFS operations with QoreStor.

- Containers must be of type NFS/CIFS
 - RDA containers cannot use Rapid CIFS. If you have existing NFS/CIFS containers, you do not need to create new containers to use Rapid CIFS; you can install the plug-in (driver) to existing clients.
- The Rapid CIFS plug-in (driver) must be installed on client systems
After the plug-in is installed, write operations will go through Rapid CIFS while metadata operations such as file creates and permission changes will go through the standard CIFS protocol. Rapid CIFS can be disabled by uninstalling the plug-in.
- Your QoreStor system must meet the minimum configuration
Rapid CIFS is available with a QoreStor system and a client with a minimum of 4 CPU cores running at a minimum of 4 GHz cumulative processing power and 2 GB memory. For a list of supported operating systems, see the *QoreStor Interoperability Guide*.
If you update your operating system, you must update your Rapid CIFS plug-in as well. Updates are available on the Quest Support site.
- Rapid CIFS is stateful
If the QoreStor system goes down, the connection will end. DMAs will restart from the last checkpoint.
- Rapid CIFS and passthrough mode
If Rapid CIFS mode fails for any reason, the QoreStor system falls back to regular CIFS mode automatically.

- Rapid CIFS acceleration constraints

Rapid CIFS does not support:

- NAS functionality
 - Optlocks (but supported if a single client is writing)
 - Byte-range locks
- Optimization of very small files (less than 10 MB). File size can be adjusted using configuration settings.
- FILE_NO_IMMEDIATE_BUFFERING and FILEWRITE_THROUGH operations (sent via CIFS only).
- File path size greater than 4096 characters

Installing the Rapid NFS plug-in

The QoreStor NFS plug-in must be installed on to the media server type you choose (for supported operating systems and DMAs, see the *QoreStor Interoperability Guide*). The plug-in software enables integration between QoreStor data storage operations and the supported data management applications (DMAs). Before you install, make sure you adhere to the Best Practices covered in another topic in this chapter.

The plug-in must be installed on the designated Linux-based media server in the following directory, **/usr/opensv/lib/**. The plug-in is installed using a self-extracting installer that installs the Rapid NFS plug-in and all of its related components. The installer supports the following modes, with the default being Help (-h):

- Help (-h)
- Install (-install)
- Upgrade (-upgrade)
- Uninstall (-uninstall)
- Force (-force)

```
$> ./QuestRapidNFS-xxxxxx-xxxxxx-x86_64.bin -help
Quest plug-in installer/uninstaller
usage: QuestRapidNFS-xxxxxx-xxxxxx-x86_64.bin [ -h ] [ -install ] [ -uninstall ]
-h                : Displays help
-install         : Installs the plug-in
-upgrade        : Upgrades the plug-in
-uninstall       : Uninstalls the plug-in
-force          : Forces the installation of the plug-in
```

Downloading the Rapid NFS Plug-in

You can download the plug-in installer from the Quest website:

- Go to support.quest.com/qorestor, select your specific QoreStor version, and then navigate to Software Downloads.
- Locate the Rapid NFS plug-in and download it to your system.

Installing Basic Calculator

Download and install Basic Calculator (bc). If you install the plug-in without it, you receive the following error: bc not found, please install it.

Enabling Secure Connect

Starting with the Rapid NFS plug-in version 4.0.3.3, the plug-in installation script includes the option to install and enable Secure Connect. Secure Connect is comprised of a set of client and server components that create a resilient and secure communication channel for use by WAN-connected clients. The server components are installed and enabled by default during the installation of the QoreStor server.

If you have WAN-connected Rapid NFS clients and choose to install Secure Connect, you can do so by adding the **sc** parameter to the mount command, as described in step 7 in the installation procedure.

Installing the plug-in

To install the Rapid NFS plug-in, follow the steps below.

i | **NOTE:** The plug-in needs to be installed on client systems to support client-side deduplication.

1. Download `QuestRapidNFS-4.0.3320.1-x86_64.bin.gz` from the website.

2. Unzip the package using the following `gunzip` command:

```
gunzip QuestRapidNFS-4.0.3320.1-x86_64.bin.gz
```

3. Change the permission of the binary package to allow it to run:

```
chmod +x QuestRapidNFS-4.0.3320.1-x86_64.bin
```

4. Install the Rapid NFS package. Before installing, remove the stale NFS entry.

```
QuestRapidNFS-4.0.3320.1-x86_64.bin -install
```

5. Load the file system in user space (FUSE) module, if not already loaded:

```
modprobe fuse
```

6. Create a directory on the client. For example:

```
mkdir /mnt/backup
```

7. Mount Rapid NFS as a file system type using the mount command. For example:

```
mount -t rdnfs 10.0.0.1:/containers/backup /mnt/backup
```

To enable Secure Connect, mount Rapid NFS using the **sc** flag, as shown below:

```
mount -t rdnfs 10.0.0.1:/containers/backup /mnt/backup sc,user=backup_
user,password=<backup_user password>
```

If you are using a DMA that supports a marker, set the marker by using **-o** in the mount command. For example, if you wanted to set the CommVault marker (**cv**):

```
mount -t rdnfs 10.0.0.1:/containers/backup /mnt/backup -o marker=cv
```

i | **NOTE:** If you want to do a mount on AIX, you must set the `nfs_use_reserved_ports` and `portcheck` parameters first. The parameters cannot be set to 0. For example: `root@aixhost1 / # nfso -po portcheck=1 root@aixhost1 / # nfso -po nfs_use_reserved_ports=1`

To ensure that the plug-in is running successfully, check the log file at: `tail -f /var/log/oca/rdnfs.log`.

Installing the Rapid CIFS plug-in

The Rapid CIFS plug-in must be installed on to the media server type you choose (for supported operating systems and DMAs, see the *QoreStor Interoperability Guide*). The plug-in software enables integration between QoreStor data storage operations and the supported data management applications (DMAs). Before you install, make sure you adhere to the Best Practices covered in another topic in this chapter.

Downloading the Rapid CIFS Plug-in

You can download the plug-in installer from the Quest website as follows

- Go to support.quest.com/qorestor, select your specific QoreStor version, and then navigate to Software Downloads.
- Locate the Rapid CIFS plug-in and download it to your system.

After it is downloaded, follow the steps below to run the plug-in installer to install the plug-in on your designated media server.

Enabling Secure Connect

Secure Connect is comprised of a set of client and server components that create a resilient and secure communication channel for use by WAN-connected clients. The server components are installed and enabled by default during the installation of the QoreStor server. Starting with the Rapid CIFS plug-in version 4.0.3.2, the plug-in installer includes the option to install and enable Secure Connect.

Installing the plug-in

To install the Rapid CIFS plug-in, follow the steps below.

i | **NOTE:** The plug-in needs to be installed on client systems to support client-side deduplication.

1. On the media server, map a network share to your CIFS-enabled container.
2. Download the plug-in installer from the website, as detailed previously.
3. Open a command prompt with the “Run as Administrator” option selected. To do this using the Windows Start menu, click **Start** → **All Programs** → **Accessories**. Right-click **Command Prompt** and select **Run as Administrator**.
This gives all the required privileges to install/copy the driver files to the Windows drivers folder.
4. Run `QuestRapidCIFS.4.0.4.4002.exe`.
5. When prompted, select **Install a Secure Connection** to install Secure Connect.
Secure Connect ensures backups to QoreStor are completed in spite of slow or unreliable WAN connections. When selecting this option, you will be prompted for additional configuration information required to make Secure Connect function correctly.
6. At the **Configure Secure Connect** page, complete the dialog using the IP or fully qualified domain name (FQDN) of the QoreStor server and appropriate credentials.
The user account used must have the CIFS role assigned.

i **NOTE:** The **Configure Secure Connect** operation is not supported for CIFS containers on DR Series appliances. DR Series appliances do not contain the required Secure Connect server component. This operation is available only for QoreStor instances. For CIFS containers on DR Series appliances, connect using the Rapid CIFS plug-in and the DR Series host name.

i **NOTE:** The value entered in the **Display Name** field must **not** match the hostname of the QoreStor server.

To ensure that the plug-in is running successfully, check the Windows Event log file.

Uninstalling the Rapid NFS plug-in

Use the following procedure to remove the Rapid NFS plug-in from a Linux-based media server. After you uninstall the plug-in, Rapid NFS will be disabled and “inactive” will be shown next to **NFS Write Accelerator** on the **NFS Connection Configuration** pane on the **Container Statistics** page.

i **NOTE:** You should retain the Rapid NFS plug-in installer on the media server in case you need to use it to reinstall the plug-in. It is usually located in `/opt/oca/DR-series/RDNFS/scripts`.

To uninstall the Rapid NFS plug-in on Linux:

1. Stop the Data Management Application (DMA) backup service before using the -uninstall option.
The Rapid NFS plug-in installer returns an error if the DMA service is running when attempting to uninstall the plug-in.
2. Run the Rapid NFS plug-in installer (usually located in `/opt/Quest/QoreStor/RDNFS/scripts`) with the -uninstall option, which uninstalls the plug-in, using the following command:

```
$> ./QuestRapidNFS-xxxxx-x86_64.bin -uninstall
```

i **NOTE:** You must stop the DMA service before uninstalling the Rapid NFS plug-in (you are also required to use the Rapid NFS plug-in installer to uninstall the plug-in).

Uninstalling the Rapid CIFS plug-in

Use the following standard Microsoft Windows uninstall process to remove the Rapid CIFS plug-in from a Windows-based media server. After you uninstall the plug-in, Rapid CIFS will be disabled and “inactive” will be shown next to CIFS Write Accelerator on the CIFS Connection Configuration pane on the Container Statistics page.

Alternatively, if you want to disable (but not uninstall) the plug-in, you can run the following Rapid CIFS utility command. The utility is located in *C:\Program Files\Quest\Rapid CIFS*.

```
rdcifscctl.exe driver -d
```

NOTE: Replace this text with a description of a feature that is noteworthy.

To uninstall the Rapid CIFS plug-in on Windows:

1. Click **Start**, and click **Control Panel**.
2. Under **Programs and Features**, click **Uninstall a program**.
3. Locate the Rapid CIFS plug-in in the listed of installed programs, right-click, and select **Uninstall**.
4. Click **Yes** to uninstall the Rapid CIFS plug-in.

Configuring and using VTL

This topic introduces Virtual Tape Libraries (VTLs) and related concepts and tasks. Refer to the subsequent topics and procedures in this section for more information.

Understanding VTL

A Virtual Tape Library (VTL) is an emulation of a physical tape library on a disk-based deduplication and compression system such as QoreStor. The tape library is exposed to a Data Management Application (DMA) as if it is a physical library with tape drives and cartridges, which the application uses for backup. Because a VTL completely emulates a standard library, the introduction of virtual tape is seamless and transparent to existing tape backup/recovery applications. The management of the library, including the drives and tapes, is done by the DMA using SCSI commands. For details on the applications supported, see the *QoreStor Interoperability Guide*.

Terminology

This topic introduces and briefly defines some basic VTL terminology used throughout the QoreStor documentation.

| Term | Description |
|------------|--|
| Library | A library is an emulation of a physical tape library and shares the same characteristics such as media changer, tape drives, and slots (cartridge slots). |
| Tape Drive | A Tape drive is a logical unit which is part of the emulated library. The media or cartridge is loaded in the Tape drives to be accessed by the Data Management application. |

| Term | Description |
|------------------------|---|
| Tapes/Media/Cartridges | Tapes are represented as files and are units within the VTL where data is actually written. Tapes are loaded into a Tape Drive before being accessed. |
| Slots | Tapes are parked in Slots before they are retrieved by the data management application for access. |

Supported virtual tape library access protocols

QoreStor supports the following virtual tape library (VTL) tape access protocols.

- Network Data Management Protocol (NDMP)
- Internet Small Computer System Interface (iSCSI)

NDMP

The Network Data Management protocol (NDMP) is used to control data backup and recovery between primary and secondary storage in a network environment. For example, a NAS server (Filer) can talk to a tape drive for the purposes of a backup.

You can use the protocol with a centralized data management application (DMA) to back up data on file servers running on different platforms to tape drives or tape libraries located elsewhere within the network. The protocol separates the data path from the control path and minimizes demands on network resources. With NDMP, a network file server can communicate directly to a network-attached tape drive or virtual tape library (VTL) for backup or recovery.

The QoreStor VTL container type is designed to work seamlessly with the NDMP protocol.

iSCSI

iSCSI or **Internet Small Computer System Interface** is an Internet Protocol (IP)-based storage networking standard for storage subsystems. It is a carrier protocol for SCSI. SCSI commands are sent over IP networks by using iSCSI. It also facilitates data transfers over intranets and to manage storage over long distances. iSCSI can be used to transmit data over LANs or WANs.

In iSCSI, clients are called *initiators* and SCSI storage devices are *targets*. The protocol allows an *initiator* to send SCSI commands (*CDBs*) to the *targets* on remote servers. It is a storage area network (SAN) protocol, allowing organizations to consolidate storage into data center storage arrays while providing hosts (such as database and web servers) with the illusion of locally attached disks. Unlike traditional Fibre Channel, which requires different cabling, iSCSI can be run over long distances using existing network infrastructure.

iSCSI is a low-cost alternative to Fibre Channel, which requires dedicated infrastructure except in FCoE (Fibre Channel over Ethernet). Note that the performance of an iSCSI SAN deployment can be degraded if not operated on a dedicated network or subnet.

The VTL container type is designed to work seamlessly with the iSCSI protocol. For details, see [Creating a VTL type container](#).

VTL and QoreStor specifications

This topic describes key specifications of VTL support in QoreStor.

- **Supported VTL Types** — QoreStor supports two types of virtual tape libraries.
 - Standard emulation of StorageTek L700 library
 - Quest OEM version of the StorageTek L700 library

i **NOTE:** The Quest type VTL is supported only with VeritasBackup Exec and Netbackup data management applications (DMAs).

i **NOTE:** Refer to the documentation for your specific QoreStor version, which includes DMA best practices whitepapers and the latest *QoreStor Interoperability Guide*, for a complete list of the supported DMAs. Visit the following site and select your specific QoreStor to download documentation:
support.quest.com/qorestor.

- **Number of Tape Drives** — Each tape library contains 10 tape drives of the type IBM-LTO-4 ('ULT3580-TD4')

- **Tapes or Media Sizes**— Each library initially is created with 60 slots housing 60 tape media of the default size of 800GiB, which is the equivalent of an LTO4 tape. You can add additional tapes to the library as needed by editing the container in the GUI or by using the following CLI command:

```
vtl --update_carts --name <name> --add --no_of_tapes <number>
```

i | **NOTE:** For more information about using the CLI, see the *QoreStor Command Line Interface Reference Guide*.

A library can only contain tapes of the same size. For example, if the library is originally created with 10 tapes of size 10GB, additional tapes of size 10GB can only be added.

The table below details the tape size and capacity configurations supported by each QoreStor installation type.

Table 6: Supported tape configurations per installation mode

| Tape | Enterprise installation | | Enterprise Plus installation | | Standard installation | | Cloud-Optimized installation | |
|-----------------|-------------------------|-------------------------------|------------------------------|-------------------------------|-----------------------|-------------------------------|------------------------------|-------------------------------|
| | Size | Max number of slots supported | Size | Max number of slots supported | Size | Max number of slots supported | Size | Max number of slots supported |
| LTO-4/ LTO-8 | 800GB | 2000 | 800GB | 2000 | 800GB | 1000 | 800GB | 500 |
| LTO-4/ LTO-8 | 400GB | 4000 | 400GB | 4000 | 400GB | 2000 | 400GB | 1000 |
| LTO-4/ LTO-8 | 200GB | 8000 | 200GB | 8000 | 200GB | 4000 | 200GB | 2000 |
| LTO-4/ LTO-8 | 100GB | 10000 | 100GB | 10000 | 100GB | 5000 | 100GB | 2500 |
| LTO-4/ LTO-8 | 50GB | 10000 | 50GB | 10000 | 50GB | 5000 | 50GB | 2500 |
| LTO-4/ LTO-8 | 10GB | 10000 | 10GB | 10000 | 10GB | 5000 | 10GB | 2500 |

i | **NOTE:** VTL on 9.x OS releases uses LTO-8 tapes.

- **Maximum Number of DMAs or Initiators Supported** — A tape library can be accessed by one DMA or iSCSI initiator at a time.

Guidelines for configuring VTL

The overall steps and recommended guidelines for using and configuring a virtual tape library (VLT) with QoreStor are described below.

Plan your Environment

Determine the following before creating a container of type VTL.

- Identify the Data Management Application (DMA) that you will be using to back up data. Refer to the *QoreStor Interoperability Guide* for a complete list of the supported DMAs.
- For the NDMP protocol, determine the filer that will be backed up using NDMP. Refer to the *QoreStor Interoperability Guide* for a list of the supported Filers and Operating systems.
- For the iSCSI protocol, determine the iSCSI initiator's properties – This is the DMA IP, hostname or IQN of the software initiator on the operating system.
- Assess the estimated size of full and incremental backups and retention periods.

i **NOTE:** The size of the full and incremental backups will determine the tape capacity size that you set. You should use a larger tape size for full backups and a smaller size for incremental backups that have smaller retention periods. Note that faster expiration periods of incremental backups residing on smaller tapes results in the release of space back to the system for future backups.

Create Containers of Type VTL

- Determine the VTL library type (NDMP or iSCSI) that you should be using as per the suggested type in the best practices guide of your preferred DMA.
Refer to the QoreStor documentation, which includes DMA best practices whitepapers for your specific QoreStor version at support.quest.com/qorestor.
- When creating the container in the GUI or by using the CLI, you will need to set the connection type of either NDMP or iSCSI. You need to provide either the DMA IP/hostname for NDMP or the IP/hostname or IQN for an iSCSI connection type.
Refer to the topics, [Creating a container](#) and [Creating a VTL type container](#), for detailed instructions about creating containers. Refer to the *QoreStor Command Line Interface Guide* for details about the CLI commands for creating containers.

Authentication/User Management Considerations

- You can use the following commands to view user information for the iSCSI user: `iscsi_user`, and NDMP user: `ndmp_user`.
 - `iscsi --show`
 - `ndmp --show`

Refer to the *QoreStor Command Line Reference Guide* for more details about using these commands.

- For iSCSI, you need to set the system-wide CHAP account for the QoreStor system. You can add this user on the Users page in the QoreStor GUI. See the topic, [Adding a user](#), for instructions for adding an iSCSI user and password.
- For NDMP, you can set the password for `ndmp_user` by using the Users page in the QoreStor GUI. These credentials are needed for configuring the NDMP-VTL in the DMA. See the topic, [Adding a user](#), for instructions for adding an NDMP user and password.

Verify the Tape Library Creation

You can easily check that the library has been created and is available for use by using the following commands.

- Check the container properties by executing the following command:

```
container --show -verbose
```

 - Upon initial addition of the connection, the NDMP/iSCSI connection status shows as ‘Added’. At this time, the library is not officially created.
 - After a few minutes, the NDMP/iSCSI connection status changes to “Available”. This status indicates that the library is online, and the tape drives and media is available for usage.
- To check the status of the virtual tape library and all the tapes in the library, you can run one of the following commands:
 - `vtl -show`
 - `vtl --show --name <container_name> --verbose`

Configure the Library in the DMA

See the QoreStor documentation, which includes DMA best practices whitepapers for your specific QoreStor version at:

support.quest.com/qorestor.

Configuring and Using Encryption at Rest

This chapter introduces the concept of Encryption at Rest as used by QoreStor as well as related concepts and tasks.

i | **NOTE:** Due to export regulations, the encryption at rest feature is not available in certain markets, and, therefore, may not be available in your locale.

Understanding Encryption at Rest

Data that resides in QoreStor can be encrypted. When encryption is enabled, QoreStor uses the Industry standard FIPS 140-2 compliant 256-bit Advanced Encryption Standard (AES) encryption algorithm for encrypting and decrypting user data. The content encryption key derives from the passphrase, which you assign to a specified storage group, and is managed by the key manager, which operates in either a **Static** mode or an **Internal** mode. In **Static** mode, a global, fixed key is used to encrypt all data. In **internal** mode, key lifecycle management is performed in which the keys are periodically rotated. The minimum key rotation period before the content encryption key can be rotated and a new key is generated is 7 days. This rotation period is user-configurable and can be specified in days.

A user-defined passphrase is used to generate a pass phrase key, which is used to encrypt the content encryption keys. It is mandatory to define a passphrase to enable encryption. The system supports up to a limit of 1023 different content encryption keys. All streams of a data-store are encrypted or re-encrypted with the same content encryption key. QoreStor statistics report the amount of data encrypted and decrypted bytes consistently.

Encryption at Rest Terminology

This topic introduces and briefly defines some basic encryption at rest terminology used in QoreStor documentation.

| Term | Description |
|------------------------|---|
| Passphrase | A passphrase is a sequence of words or other text used to control access to data, similar to a password in usage, but is generally longer for added security. The QoreStor passphrase is user-defined and is used to generate a passphrase key that encrypts the file in which the content encryption keys are kept. The passphrase is a human readable key, which can be up to 255 bytes in length. It is mandatory to define a passphrase to enable encryption. |
| Content encryption key | The key used to encrypt the data. The content encryption key is managed by the key manager, which operates in either a static mode or an internal mode. The system supports up to a limit of 1023 different content encryption keys. |
| Key management mode | The mode of key lifecycle management as either static or internal. |
| Static mode | A mode of key management in which a fixed key is used to encrypt all data and is global for each storage group, which lets you configure static mode for one storage group and internal mode for another storage group. |
| Internal mode | A mode of key lifecycle management in which the keys are periodically generated and rotated. The minimum key rotation period before the content encryption key can be rotated and a new key is generated is 7 days. This rotation period is user-configurable and can be specified in days. |

Encryption at Rest and QoreStor Considerations

This topic describes key features and considerations of using Encryption at Rest in QoreStor.

- **Key Management** — In internal mode there is a maximum limit of 1023 keys. By default when encryption is enabled on the system, the key rotation period is set to 30 days. Users can later change the key rotation period from 7 days to 70 years, while configuring internal mode of encryption.
- **Performance Impacts** — Encryption should have minimal to zero impact on both backup and restore workflows. It should also have no impact on the replication workflows.
- **Replication** — Encryption must be enabled on both the source and target QoreStor systems to store encrypted data on the systems. This means that encrypted data on the source does not automatically imply that when it is replicated to the target it will be encrypted unless encryption is explicitly turned 'ON' on the target QoreStor system.
- **Security Considerations for Passphrase and Key Management** —
 - A passphrase is very important part of the encryption process on the QoreStor system as the passphrase is used to encrypt the content encryption key or keys. If the passphrase is compromised or lost, the administrator should change it immediately so that the content encryption keys do not become vulnerable.
 - The administrator should closely consider security requirements to drive the decision for selecting the mode of key management for the QoreStor system.
 - The Internal mode is more secure than the Static mode since the keys are periodically changed. Key rotation can be set to 7 days minimum.
 - Key modes can be changed at any time during the lifetime of the QoreStor system; however, changing the key mode is a significant operation to undertake as all encrypted data must be re-encrypted.
 - Content encryption keys are stored in their encrypted form in a primary keystore, which is maintained on the same enclosure as the data-stores. For redundancy purposes, a backup copy of the primary keystore is stored on the system in the root partition, separate from the data-store partitions.

Understanding the encryption process

The overall steps for how Encryption at Rest is enabled and used in QoreStor are described below.

1. Enabling encryption.

Encryption is disabled by default on QoreStor. An administrator can enable encryption by using the GUI or CLI.

Encryption is set at the storage group level.

2. Setting a passphrase and setting the mode.

When defining encryption for a storage group, a passphrase is set. This passphrase is used to encrypt the content encryption keys, which adds a second layer of security to the key management. At this time, the mode is also set. The default key management mode is "internal" mode, in which key rotation happens periodically as specified by the set key rotation period.

3. Encryption process.

After encryption is enabled, the data in the storage group that gets backed up is encrypted and is kept encrypted until it is expired and cleaned by the system cleaner. Note that the encryption process is irreversible.

4. Encryption of pre-existing data.

Any pre-existing data will also be encrypted using the currently set mode of key management. This encryption occurs as part of the system cleaner process. Encryption is scheduled as the last action item in the cleaner workflow. You must launch the cleaner manually using the maintenance command to reclaim space. It then encrypts all pre-existing unencrypted data. The cleaner can also be scheduled as per the existing pre-defined cleaner schedule.

i **NOTE:** The cleaner can take some time to start the encryption process if the system is nearing full system capacity. Encryption starts only after the cleaner processes data slated for cleaning and the related logs. This ensures that space reclamation is prioritized when free space is low and also ensures that data stores are not redundantly encrypted.

Refer to the *QoreStor Command Line Interface Reference Guide* for information about the CLI commands used for encryption.

Configuring and using the Recycle Bin

This topic introduces the Recycle Bin feature and related concepts and tasks. Refer to the subsequent topics and procedures in this section for more information.

Understanding Recycle Bin

Recycle Bin is a process whereby any data that is deleted from a compromised backup solution can be retrieved from an immutable Recycle Bin by a QoreStor administrator. Data in the Recycle Bin is not visible through client protocols such as NFS and CIFS. You can set a retention period of between 7 and 30 days after which the files are automatically removed permanently from QoreStor.

Recycle Bin is a property of a container in QoreStor. Recycle Bin can be enabled on a container when you create the container or any time later. Only files that are deleted after enabling Recycle Bin are retained in it.

Recycle Bin activation is permanent. After you enable Recycle Bin for a container, it cannot be removed from the container configuration.

You cannot delete a container until the container and Recycle Bin associated with it are empty.

i **NOTE:** When you delete a file, you cannot create a file with the same name and location until the deleted file is permanently removed from Recycle Bin.

NOTE: Enabling Recycle Bin does not allow Linux hardlinks to be used.

Compatibility

Recycle Bin is supported with NFS, CIFS, RDA, and OST containers. For more information, see the *Quest QoreStor Interoperability Guide*.

Terminology

This topic introduces and briefly defines some basic Recycle Bin terminology used throughout the QoreStor documentation.

| Term | Description |
|--------------------------------|---|
| Recycle Bin | A hidden space where QoreStor temporarily stores deleted files. |
| Retention time | Duration by which QoreStor keeps a deleted file in Recycle Bin before permanently purging it. |
| Purging files from Recycle Bin | Permanently deleting files from Recycle Bin. |
| Recycle Bin restore | Recovering or bringing back files in Recycle Bin to their original location. |
| QoreCompliance Clock | An internal software clock that QoreStor uses to determine retention time and immutability for the Recycle Bin feature. |

Guidelines for configuring Recycle Bin

Refer to the following important notes and guidelines for understanding Recycle Bin configuration in QoreStor.

- Identify the Data Management Application (DMA) that you will be using to back up data. Refer to the *QoreStor Interoperability Guide* for a complete list of the supported DMAs for Recycle Bin.
- Depending on the backup size, deduplication and compression ratios, and Recycle Bin retention time, Recycle Bin occupies additional storage space. No additional QoreStor license is required for Recycle Bin space. Provision storage to use all of the license capacity.
- Recycle Bin retention time is a time period during which you can retrieve any of the deleted files. Shorter the retention time means less additional storage space required, but also less of a time window to detect an attack or accidental deletion and respond.

Configuring Recycle Bin

You can enable Recycle Bin on a container while adding a container or later by editing the properties of a container. For more information, see [Creating a container](#) and [Editing container settings](#).

i | **NOTE:** Recycle Bin supports NFS, CIFS, RDA, and OST containers.

You can also change the retention time for a container, which affects files that are deleted from then onwards. The retention time does not include when QoreStor services are not running. You can set a retention time when you add a container or at a later time. A new retention time applies to the files that are deleted from the moment you set the new retention time.

Purging Recycle Bin files

Files in the Recycle Bin are permanently deleted after their retention time is over. This cleanup process happens automatically on a daily basis, and it cannot be forced or expedited. Once the files are deleted from Recycle Bin, they cannot be recovered.

Restoring files from Recycle Bin

If files are accidentally or maliciously deleted, you can restore them from Recycle Bin by identifying the time period during which the files were deleted, optionally viewing the list of files that were deleted during that period, and then initiating a restore from Recycle Bin. Once restore is complete, you can read the files from client systems or the backup software. Restore time is usually short, depending on the number of files being restored.

i | **NOTE:** In the restore window, it is recommended to not expire any backups, as it could cause file deletes to fail during that time. Retrying them after restore operation would succeed.

If backups were deleted from a backup application after restoring the files from Recycle Bin, some backup applications have to rescan the backup target or device to repopulate their catalog.

To restore files from Recycle Bin

1. In the QoreStor GUI, click **Container**.
2. Next to the name of the appropriate container, click the **Details** menu, and then click **Recycle Bin**.
3. To receive a list of files that are in Recycle Bin, click **Report**. You can filter files based on their deletion times.
4. To restore the files, click **Restore**.
 - To restore all of the files, click **Full restore**.
 - To restore files that were deleted during a specific time, click **Date Range**, and then select the start and end times.

For usage details on the `container --recycle_bin` command, see the *Quest QoreStor Command Line Reference Guide*.

Managing containers with Recycle Bin enabled

When working with containers that have Recycle Bin enabled, take the following conditions into consideration.

Container deletion

When deleting a container that has Recycle Bin enabled, you cannot delete the container until Recycle Bin is empty. Be sure to delete or expire all backups from the container and wait for the files to expire from Recycle Bin. When Recycle Bin is empty and the container has no more files, delete the container.

Container replication

Recycle Bin can be configured on the source container only. The same configurations are automatically applied on the peer or replica container. Replication transfers all of the files in the container, including the Recycle Bin files. Restoring files from Recycle Bin is allowed on the source container only. The restored files appear on the replica container as well.

If the target container has files in Recycle Bin, then replication is not allowed. Doing so could prevent replication from being re-added between a previously configured replication pair. Quest recommends using a new container on the target QoreStor in the same Storage Group as the replica container.

Cloud tier and archive tier

Cloud data also maintains Recycle Bin files. A retention period in a cloud is the same as an on-premises retention period for each file. Files in Recycle Bin must be restored to before they can be restored from Archive storage (Glacier) and later readback.

Disaster recovery

During disaster recovery using cloud data, the files in Recycle Bin are also made available on the recovered QoreStor. The Recycle Bin retention time for all of the files is reset to the last retention time set on the container, which means that all of the files in Recycle Bin are retained for the same number of days from the recovery time.

Viewing Recycle Bin statistics

The number of files in Recycle Bin, their size, and estimated post compression bytes can be seen for each Recycle Bin enabled container.

To view Recycle Bin statistics

1. In the QoreStor GUI, click **Containers**.
2. Next the name of the appropriate container, click the **Details** menu, and then click **Recycle Bin**. The available statistics for Recycle Bin display on the page.

For the usage of the `stats --container` command, see the *Quest QoreStor Command Line Reference Guide*.

Configuring Cloud Reader

This topic introduces the Cloud Reader feature and related concepts and tasks. Refer to the subsequent topics and procedures in this section for more information.

Understanding Cloud Reader mode

A QoreStor instance can be configured to read the backups uploaded to a Cloud Tier from a different QoreStor instance. The primary QoreStor, where the Cloud Tier bucket is configured, writes and manages the data in the bucket while another QoreStor instance in Cloud Reader mode can access the same data in Read-Only mode. Any number of Cloud Reader instances can access the cloud bucket data. The data in the cloud is only read by this Cloud Reader instance. A snapshot of the files is made available to the Cloud Reader.

Use cases for cloud reader include Testing and Development and for validating QoreStor disaster recovery.

Operations allowed

Cloud Reader mode is a restrictive mode, where only few QoreStor configuration changes are allowed. For example, new storage groups and containers cannot be created. Only the following actions are allowed:

- Adding QoreStor to Active Directory domain
- Adding Domain Group
- Password changes for the users
- Addition of license
- Diagnostics generation

i **NOTE:** If a QoreStor instance is only used to read the data in the cloud, then no license is required to be added.

NOTE: Any new data written to the Cloud Reader instance are saved in local storage and not propagated to Cloud Tier.

Compatibility

Only RDA protocol is supported. NetVault and BridgeHead backup solutions are supported. For the list of supported DMAs, see the *QuestQoreStor Interoperability Guide*.

Deploying Cloud Reader

Install QoreStor with similar mode as the QoreStor that is writing data to the cloud tier bucket. Use the minimal required CPU, memory, and storage as per the *Quest QoreStor Interoperability Guide*.

After you complete the installation, change the mode by providing the cloud connection details. For details about the maintenance `--disaster_recovery` and `--quick_ro_recovery` commands, see the *Quest QoreStor Command Line Reference Guide*. To refresh the data at any time later, repeat the same command. The existing data is replaced with the file space from cloud bucket.

Example command:

```
maintenance --disaster_recovery --cloud_string
"DefaultEndpointsProtocol=https;AccountName=qorestortest;AccountKey=8Tt7/ysHSGSBSW7
FG1Vr2+27xgccskbUWf9GLlGEPeMHYfmVxl+fTg1XYpA==;EndpointSuffix=core.windows.net" --
container_name cloud_container1 --cloud_provider_type AZURE --passphrase qqq --
logfile /tmp/t1 --quick_ro_recovery yes
```

Quest recommends that the cloud reader QoreStor instance be in the cloud, but it also works from on-premises.

i **NOTE:** The QoreStor UI does not support initiating Cloud Reader mode. After you deploy Cloud Reader and it is operational, you can use the QoreStor UI for monitoring and configuration changes.

Refreshing data in Cloud Reader

When you run the maintenance command with the `--quick_ro_recovery`, the QoreStor instance only sees what was in cloud bucket at the time the command was done. Later backup data can become expired from the primary QoreStor instance and the Cloud Tier, and new backup data uploads. To see the new file space view, run the same maintenance command. It erases the current view and retrieves the latest view from the cloud.

Configuring RDA immutability

This topic introduces the RDA immutability feature and related concepts and tasks. Refer to the subsequent topics and procedures in this section for more information.

Understanding RDA immutability

RDA immutability provides protection from overwrites and deletes on backup files. This technology is present by default for RDS containers, but not all backups are protected by default. Backup applications sending data using the RDA protocol to QoreStor define whether the backup data should be immutable and the time period the data will remain immutable. After you set immutability on the backup data, you cannot modify or delete the backup data from the RDA container until the immutable time period expires. For the current list of supported backup applications that take advantage of this technology, see the *Quest QoreStor Interoperability Guide*.

Managing containers with RDA immutability

When working with containers that have RDA immutability, take the following conditions into consideration.

Container deletion

Containers with immutability and cloud locked files cannot be deleted until the immutability on all the backup data expires.

Container replication

When container replication is set up between two RDA containers and when replicated backup data includes an immutability time period, the remaining immutability time period is set on the backup data replicated to target containers.

i | **NOTE:** It is not possible to add a container with immutable backup data as a replication target.

In case of secondary copy or duplicate operations performed by backup applications over RDA, backup applications set immutability time period on the target backup data.

Cloud Tier and Archive Tier

QoreStor maintains the immutability time period on backup data when it transfers them to a cloud tier or archive tier. During a disaster recovery from a cloud tier or archive tier, the immutable time period is again set on the backup data, making it immutable on new QoreStor instances as well..

Viewing immutable container statistics

The container statistics for an RDA container show the current number of immutable files. This immutable statistic does not automatically update when the immutability time period expires for some backup data, which means you must update the statistics and check the current immutable files after that time.

To view immutable container statistics

1. In the command line interface, to update the statistics on the immutable files, run the following command:

```
container --update --name <name> --immutable_file_stat
```

2. To check the latest immutable file count in the container, run the following command:

```
stats --container --name <container_name>
```

Managing containers

In QoreStor, data is stored in containers, which are stored in storage groups. Some containers function like a shared file system. These types of containers can be assigned a specific connection type, for example, NFS/CIFS or RDA (including both OST and RDS clients). These containers are then accessed via NFS, CIFS, and RDA protocols.

Each container includes an optional Recycle Bin. When enabled, the Recycle Bin holds deleted files for a specified number of days before completing the deletion. During this time, the files are still available for recovery. Enabling the Recycle Bin for a container not only provides a safety net for accidentally deleted files, but also protects files from ransomware that attempts to destroy your data.

In QoreStor you can manage your storage groups and data containers, including viewing storage groups and containers, creating new storage groups and containers, modifying or deleting them, moving data into containers, and viewing current statistics. Management for containers can be done either through the GUI or the command line.

i **NOTE:** If only the DefaultGroup storage group exists on your system, all containers you create are automatically added to that group. You can create custom storage groups, and then when you create a container, you can specify that it be added to the custom storage group. For more information about storage groups, see the topic, "Managing Storage Groups."

Creating a container

For more information about storage groups, see [Managing local storage](#)

i **NOTE:** QoreStor does not support container names that begin with a number.

Containers can be accessed using the following connection types:

- **NFS**
- **CIFS**
- **Quest Rapid Data Storage (RDS)**
- **Object (S3 compatible)**
- **Veritas OpenStorage Technology(OST)**
- **Virtual tape library (VTL)**
- **EDM**

Refer to the sections below for instructions on creating containers:

- [Creating an OST or RDS connection type container](#)
- [Creating an NFS or CIFS connection type container](#)

- [Creating an Object Container](#)
- [Creating a VTL type container](#)
- [Creating an EDM connection type container](#)

Creating an OST or RDS connection type container

To create an OST or RDS connection type container, follow these steps:

1. In the navigation menu, click **Containers**. The **Containers** page is displayed.
2. Click **Add Container**.
3. For **Protocol**, select **Rapid Data Storage (RDS)** or **Veritas OpenStorage (OST)** as appropriate.
4. For the container **Name**, type the name of the container, and then click **Next**.
Container names cannot exceed 32 characters in length, must start with a letter, and can be composed of any combination of the following characters:
 - A-Z (uppercase letters)
 - a-z (lowercase letters)
 - 0-9 (numbers). Do not start a container name with a number.
 - dash (-) or underscore (_) special characters

i | **NOTE:** QoreStor does not support the use of the following special characters in container names: /, #, or @.
5. In the **Storage Group** drop-down, select the Storage Group for this container.
6. Click **Next**.
7. On the **User Access Control** page, select the appropriate permissions for the displayed user accounts. Refer to [Configuring User Access Controls](#) for more information.
8. If you selected **RDS**, **LSU Capacity** is set to **Unlimited** by default. If you selected **Veritas OpenStorage (OST)**, for **LSU Capacity** select one of the following options allowed per container:
 - **Unlimited** — To define the allowed amount of incoming raw data per container (based on the physical capacity of the container).
 - **Quota:** To define a set limit in Gibibytes (GiB) for incoming raw data allowed per container.
9. Click **Next**.
10. Optionally, select **Recycle Bin**, and then enter the number of days you want files to remain in the Recycle Bin before deleting. For more information, see [Managing containers](#).

i | **NOTE:** Enabling the Recycle Bin is an irreversible step. Once it is enabled, you cannot disable it on a container.
For information about which versions of NetVault support the Recycle Bin and RDA immutability, see the *Quest QoreStor Interoperability Guide*.
11. Click **Next**.
12. Click **Finish**.



NOTE: To add a container through the command line, use the command:

```
container --add --name <name> [--group_name <name>]
```

Refer to the *QoreStor Command Line Reference Guide* for more information.

Configuring User Access Controls

QoreStor allows you to specify user access controls for individual RDA containers. User access controls allows for multiple containers of the same type while ensuring that access to each container is isolated to specific users. User access controls (UACs) can be set at the container level to assign RWD (read-write-delete) or RW (read-write) permissions on that container to individual user accounts. Data within the containers can be access or deleted by users with both permission types, but the container can only be deleted by a user with RWD permissions.

The **backup_user** account is assigned by default to RDA containers, and is granted RWD permissions. Additional user access can be configured through the GUI or CLI.



NOTE: User Access Controls is currently only supported on RDA containers.

Requirements for using User Access Controls

- The user account to which access is going to be assigned should be created before creating the container or configuring UAC.

To configure User Access Controls

1. In the navigation menu, click **Containers**. The **Containers** page is displayed.
2. In the list of containers, find the container for which you add user access controls, and then click **User Access Control**.
3. The **backup_user** account is listed by default. To configure other accounts with permissions, use the **Search** field to find the account. Click the desired account to add it to the users list.
4. For each listed account, select the appropriate permissions. Options are:
 - **Read/Write** - gives the account read and write permissions on the container.
 - **Read/Write/Delete** - gives the account read, write, and delete permissions on the container.
5. Click **Finish**.



NOTE: To add user access controls to a container through the command line, use the command:

```
container --add --name <name> [--group_name <name>]container --add_uac --name <name> --user <user name> --mode <RW|RWD>
```

Refer to the *QoreStor Command Line Reference Guide* for more information.

Creating an NFS or CIFS connection type container

To add an NFS or a CIFS connection type container, complete the following steps:

1. In the navigation menu, click **Containers**. The **Containers** page is displayed.
2. Click **Add Container**.
3. For **Protocol**, select **NAS (NFS, CIFS)**.
4. For the container **Name**, type the name of the container, and then click **Next**.
Container names cannot exceed 32 characters in length, must start with a letter, and can be composed of any combination of the following characters:
 - A-Z (uppercase letters)
 - a-z (lowercase letters)
 - 0-9 (numbers). Do not start a container name with a number.
 - dash (-) or underscore (_) special characters

i | **NOTE:** QoreStor does not support the use of the following special characters in container names: /, #, or @.
5. In the **Storage Group** drop-down, select the Storage Group for this container.
6. Click **Next**.
7. For **Marker Type**, select the appropriate marker that supports your Data Management Application (DMA).
 - **Auto** — Automatically detects CommVault, Tivoli Storage Manager (TSM), ARCserve, and HP Data Protector marker types. In addition, select this option if you need to support EMC Networker 2.0.
 - **ARCserve**—Supports the ARCserve marker.
 - **BridgeHead** — Supports the BridgeHead HDM marker.
 - **CommVault**—Supports the CommVault marker.
 - **HP DataProtector**—Supports the HP Data Protector marker.
 - **Networker** — Supports EMC Networker 3.0. If you need to support EMC Networker 2.0, select Auto.
 - **Time Navigator**—Supports the Time Navigator marker.
 - **TSM**—Supports the TSM marker.
 - **Unix Dump** — Supports the Amanda marker, among others.

i | **IMPORTANT:** Improper marker selection can result in non-optimal savings. As a best practice, if you have only one type of DMA with traffic directed to a container, it is best to select the marker type that supports your DMA (for example, **BridgeHead**, **Auto**, or another). Conversely, as a best practice, if you have traffic from a DMA that is not one of the supported marker types, it is best to disable marker detection for the container by selecting the **None** marker type.
8. For **Access Protocols**, select **NFS** and **CIFS** as appropriate.
(Use **NFS** to back up UNIX or LINUX clients. Use **CIFS** to back up Windows clients.)
9. Click **Next**.

10. If you selected **NFS** as the connection type, configure NFS access as follows. For CIFS connections, proceed to **step 11**.
- **NFS Options** — Defines the type of access to the container. Select one of the following options.
 - **Read Write Access** — To allow read-write access to the container.
 - **Read Only Access** — To allow read-only access.
 - **Root Mapping**— Select one of the following options from the drop-down list to define the user level you want mapped to this container.
 - **Root** — to specify a remote user with root access to read, write, and access files on the system.
 - **Nobody** — to specify a user on the system without root access permissions.
 - **Administrator** — to specify the system administrator.
 - **NFS Client Access** — Define the NFS client(s) that can access the NFS container or manage the clients that can access this container by selecting one of the following options.
 - **Open (allow all clients)** — To allow open access for all clients to the NFS container you create. (Select this option *only* if you want to enable access for all clients to this NFS container.)
 - **Create Client Access List** — To define specific clients that can access the NFS container. In the Client FQDN or IP text box, type the IP address (or FQDN hostname) and click the Add icon. The “added” client appears in the Allow Clients list box. (To delete an existing client from this list box, select the IP address (or FQDN hostname) of the client you want to delete, and click the Delete icon. The “deleted” client disappears from the list box.)

11. If you selected **CIFS** as the connection type, configure CIFS access as follows.
 - Select the **Veeam Overwrite Protection** checkbox. A consent warning window is displayed. If you provide a consent, the NFS selection will be automatically cleared (if previously selected).

i NOTE: The Recycle Bin option is automatically selected with Overwrite Protection and cannot be disabled by the user. Once selected, the Overwrite Protection in Edit cannot be disabled.
 - **Client Access** — Define the CIFS client(s) that can access the container or manage the clients that can access this container by selecting one of the following options.
 - **Open (allow all clients)** — To allow open access for all clients to the container you create. (Select this option *only* if you want to enable access for all clients to this container.)
 - **Create Client Access List** — To define specific clients that can access the container. In the Client FQDN or IP text box, type the IP address (or FQDN hostname) and click the Add icon. The “added” client appears in the Allow Clients list box. (To delete an existing client from this list box, select the IP address (or FQDN hostname) of the client you want to delete, and click the Delete icon. The “deleted” client disappears from the list box.)

i NOTE: The QoreStor administrator that manages the system has a different set of privileges than does the CIFS administrator user. Only the QoreStor administrator can change the password for the CIFS administrator user. To change the password that allows access for the CIFS administrator user, use the authenticate --set --user administrator commands. For more information, see the *QoreStor Command Line Reference Guide*.
12. Click **Next**.
13. Optionally, select **Recycle Bin**, and then enter the number of days you want files to remain in the Recycle Bin before deleting. For more information, see [Managing containers](#).

i NOTE: Enabling the Recycle Bin is an irreversible step. Once it is enabled, you cannot disable it on a container.
14. Click **Next**.

A Configuration Summary of the options you selected for creating the container appears.
15. Click **Finish**.

Creating an Object Container

Adding an object container can be accomplished through the QoreStor UI or via the **object_container** command in the QoreStor CLI. Refer to the *QoreStor Command Line Reference Guide* for more information on the **object_container** command.

- i NOTE:** QoreStor object container does not support object lifecycle management, which means transitioning storage classes or server side expiration of objects is not supported. User policies are limited to predefined readwrite, writeonly, and readonly.

To create an object container

1. In the navigation menu, click **Containers**.
2. On the **Containers** pane, click **Add Container**. The Add Container dialog will be displayed.
3. In the Protocol field, select **Object (S3 Compatible)**.
4. In the **Storage Group** drop-down, select the required storage group for this container.
5. Click **Next**.
6. Optionally, select **Use HTTP instead of HTTPS**. To use an HTTP connection, you must also follow the steps below:
 - a. On the QoreStor server, copy the aws.conf file to a new location:

i **NOTE:** The QoreStor implementation of object storage uses a self-signed certificate. If your data management application requires third party certificates, you must use HTTP to connect to the object container.

7. Click **Next**.
8. Review the summary and click **Finish**.

When the process is completed the object container is added to the QoreStor. For Object container created prior to QoreStor release 7.2.1 you will see the storage group **ObjectContainer** and the container **ObjectStorageGroup** added to the **Storage Groups** and **Container** pages, respectively. See the topics below for information on working with object storage.

- [Creating a bucket](#)
- [Managing Object container users](#)
- [Editing bucket settings](#)

Adding an object container through the command line

To add an object container, complete the following steps.

1. Access the QoreStor CLI. Refer to [Accessing the CLI commands on a Linux QoreStor server](#) for more information.
2. Add a Object container:

```
object_container --add --name <container name> [--group <storage group name>]
```

Refer to the *QoreStor Command Line Reference Guide* for more information.
3. Get end-point details of it:

```
object_container --show --name <container name>
```

4. Create user for this container. This user name is used as Access key and user's password is used as Secret key while accessing Object container from the client systems (backup clients):

```
object_container --user-add --name <name> --user-name <user name>
```

i | **IMPORTANT:** The User's name is used as Access Key and the user's password is used as Secret Key while connecting to QoreStor from the S3 clients.

To see the S3 endpoint, use the command `object_container --show --endpoint --name <name of container>`

The endpoint is displayed in the format **https://<QoreStor IP address>:<port>**

Make sure the port is allowed for access through the firewall.

5. Set access policy for the user. Use <Policy name> as "readwrite" to allow the user to backup and restore data.

```
object_container --policy-set --name <name> --policy-name <Policy name> --user-name <user name>
```

6. Create bucket for use in the backup application. Optionally add locking support:

```
object_container --bkt-add --name <name> --bkt-name <bucket name> [--enable-object-lock] [--enable-object-versioning]
```

7. Configure the backup application with the endpoint, access key, secret key, and bucket name.

Creating a VTL type container

i | **NOTE:** For more information on using VTL containers, see [Configuring and using VTL](#).

To create a virtual tape library (VTL) type container, complete the following steps.

i | **NOTE:** The number of supported VTL containers varies depending on the QoreStor installation mode. Refer to the *QoreStor Interoperability Guide* for more information.

1. In the navigation menu, click **Containers**. The **Containers** page is displayed.
2. Click **Add Container**.

3. For **Name**, type the name of the container.

i **NOTE:** QoreStor does not support spaces or the following special characters in container names: /, #, or @. VTL container names cannot exceed 32 characters in length, must start with a letter, and can be composed of any combination of the following characters:

- A-Z (uppercase letters)
- a-z (lowercase letters)
- 0-9 (numbers). (Do not start a container name with a number.)
- underscore (_) special characters
- hyphen (-) special character

i **NOTE:** iSCSI VTL containers do not support the following characters:

- ASCII CONTROL CHARACTERS and SPACE through ,
- ASCII /
- ASCII ; through @
- ASCII [through `
- ASCII { through DEL

4. For **Protocol**, select **Virtual Tape Library (VTL)**.

5. Click **Next**.

6. For **Robot Model**, select the type of virtual tape library for the VTL container.

- STK L700—This is the standard emulation of the StorageTek L700 library.
- QUEST DR_L700 - This is a Quest OEM version of StorageTek L700 library.

i **NOTE:** The Quest version of the VTL is supported only with Symantec Backup Exec and Netbackup data management applications (DMAs).

7. For **Tape Size**, select the size of the tapes for your tape library from one of the following options.

- 800 GB
- 400 GB
- 200 GB
- 100 GB
- 50 GB
- 10 GB

i **NOTE:** Creating a VTL container type creates a tape library of type Storage Tek L700 with 10 tape drives of type IBM Ultrium LTO-4 and 60 tape slots holding 60 tapes. Additional tapes can be added as required. For more information, see [VTL and QoreStor specifications](#).

8. For **Access Protocol**, select one of the following options. Each protocol has different configuration requirements, as listed below.

- **NDMP**

- Enter DMA's **FQDN or IP address** that will access the VTL container.
- For **Marker Type**, select the appropriate marker that supports your DMA from the options below:
 - **None** — Disables marker detection for the container.
 - **Auto** — Automatically detects CommVault, Tivoli Storage Manager (TSM), ARCserve, and HP Data Protector marker types. In addition, select this option if you need to support EMC Networker 2.0.
 - **ARCserve**—Supports the ARCserve marker.
 - **BridgeHead** — Supports the BridgeHead HDM marker.
 - **CommVault**—Supports the CommVault marker.
 - **HP DataProtector**—Supports the HP Data Protector marker.
 - **Acronis** —Supports the Acronis marker
 - **Networker** — Supports EMC Networker 3.0. If you need to support EMC Networker 2.0, select Auto.
 - **TSM**—Supports the TSM marker.
 - **Unix Dump** — Supports the Amanda marker, among others.

- **iSCSI**

- Enter the **FQDN, IQN, or IP address** of the iSCSI initiator that can access the VTL container.
- For **Marker Type**, select the appropriate marker that supports your DMA from one of the following options:

i **NOTE:** Improper marker selection can result in non-optimal savings. As a best practice, if you have only one type of DMA with traffic directed to a container, it is best to select the marker type that supports your DMA. Conversely, as a best practice, if you have traffic from a DMA that is not one of the supported marker types, it is best to disable marker detection for the container by selecting the None marker type.

- **None** — Disables marker detection for the container.
- **Auto** — Automatically detects CommVault, Tivoli Storage Manager (TSM), ARCserve, and HP Data Protector marker types. In addition, select this option if you need to support EMC Networker 2.0.
- **ARCserve**—Supports the ARCserve marker.
- **BridgeHead** — Supports the BridgeHead HDM marker.
- **CommVault**—Supports the CommVault marker.
- **HP DataProtector**—Supports the HP Data Protector marker.
- **Acronis** —Supports the Acronis marker
- **Networker** — Supports EMC Networker 3.0. If you need to support EMC Networker 2.0, select Auto.
- **TSM**—Supports the TSM marker.
- **Unix Dump** — Supports the Amanda marker, among others.

- **No Access.**

i **NOTE:** QoreStor allows you to create a VTL container type without configuring it with a specific protocol (that is, by selecting **No Access**). You can configure the container at a later date.

9. Click **Next**.

A summary of the options you selected for creating the container appears.

10. Click **Finish**.

Viewing VTL tape information

Once you have created a virtual tape library (VTL) type container, you can view the detailed tape information of the VTL. This includes information about the vendor and model information for medium changer and tape drives. To view VTL information, complete the following steps.

1. In the navigation menu, click **Containers**. The **Containers** page is displayed.
2. In the list of containers, find the container for which you want to view statistics, and then click the  (ellipsis icon) in the upper-right corner of the container's information pane.
3. Click **Details**.

4. You can view the following information:

- Container Details
 - Number of Tape Drives
 - Library ID
 - Tape Size
 - is OEM
 - Container Path
 - Marker
 - Created On
- Connection
 - Enabled
 - Status
 - Cartridges Available

Creating an EDM connection type container

To create an EDM connection type container, complete the following steps.

1. In the navigation menu, click **Containers**. The **Containers** page is displayed.
2. Click **Add Container**.
3. For the container **Name**, type the name of the container.

i **NOTE:** QoreStor does not support spaces or the following special characters in container names: /, #, or @. VTL container names cannot exceed 32 characters in length, must start with a letter, and can be composed of any combination of the following characters:

- A-Z (uppercase letters)
- a-z (lowercase letters)
- 0-9 (numbers). (Do not start a container name with a number.)
- underscore (_) special characters
- hyphen (-) special character

4. For **Protocol**, select **Enhanced Data Mover (EDM)**.
5. Click **Next**.
6. Select the required storage group from the **Storage Group** dropdown.
7. Click **Next**.
8. Optionally, select **Recycle Bin**, and enter the number of days you want files to remain in the Recycle Bin before deleting.

For more information, see [Managing containers](#).

i **NOTE:** Once the Recycle Bin option is enabled, it cannot be undone or disabled on the containers.

9. Click **Next**.
10. Click **Finish**.

Viewing containers

You can easily view a list of containers in your QoreStor instance on the Storage Containers page, or by using the **container** command in the QoreStor CLI.

Viewing containers in the GUI

1. In the navigation menu, click **Containers**. The **Containers** page is displayed.
2. You can view the following columns of information about the containers.
 - **Container**—The name of the container.
 - **Path**—The path for the container.
 - **Marker**—The marker type that supports your Data Management Application (DMA).
 - **Connection**—The connection type/access protocol for the container:
 - NFS
 - CIFS
 - RDA
 - OST
 - **Cloud Tiering Policy**—The status of whether the container is connected to a cloud container. If a cloud container link is not configured, **Enable Cloud Tiering Policy** will be displayed.
 - **Archive Tiering Policy**—The status of whether the container is connected to an archive tier. If an archive tier link is not configured, **Enable Archive Tiering Policy** will be displayed.
 - **Ellipsis icon** —Provides additional options as listed below:
 - Details
 - Enable Cloud Tiering Policy
 - Enable Archive Tiering Policy
 - Edit
 - Delete

Viewing containers via the CLI

To view the list of containers, follow these steps.

1. Access the QoreStor CLI. Refer to [Accessing the CLI commands on a Linux QoreStor server](#) for more information.
2. To view the containers on this QoreStor instance, use the command

```
container --show [--name <name>] [--verbose]
```

Refer to the QoreStor Command Line Reference Guide for more information.

Viewing container statistics

In the QoreStor GUI, you can view statistics about a selected container. All statistics displayed represent specific information about the backup data, throughput, replication, marker type, and connection type for the selected container. The displayed statistics will vary depending upon the connection type used by the specified container.

To display container statistics for a selected container, complete the following steps.

1. In the navigation menu, click **Containers**. The **Containers** page is displayed.
2. In the list of containers, find the container for which you want to view statistics, and then click the  (ellipsis icon) in the upper-right corner of the container's information pane.
3. Click **Details**.
4. The **Container details** page contains the following sections:
 - In the **Active Files**, **Active Bytes**, and **Throughput** charts, you can view current statistics for the container.
The **Active Files** chart displays the number of active files ingested based on time (in minutes), and the **Active Bytes** displays the number of active bytes ingested based on time (in minutes).
The **Throughput** chart displays the number of read data in Mebibytes/per second (MiB/s) based on time (in minutes), and the number of write data in MiB/s based on time (in minutes).
 **NOTE:** The values in the **Active Files**, **Active Bytes**, and **Throughput** charts refresh automatically every 15 seconds.
 - In the **Connection** pane, you can view information about the configured connection type for the selected container. The type of information displayed can be different depending on the connection type. For example, for an RDS container, the following information is displayed:
 - Type
 - Enabled
 - Status
 - Capacity
 - In the **Inbound Stats** and **Outbound Stats** panes, you can view detailed information on the inbound throughput, and outbound throughput.
 - The **Client Stats** pane lists ingest and network information.
 - The **Recycle Bin** pane shows whether this feature is enabled, the length of its retention period, the number of files it contains, and the number of logical bytes it is using. From here, you can also generate a report or restore the contents of the Recycle Bin.

Displaying container statistics by using the CLI

An alternate method for viewing container statistics is by using the QoreStor CLI command: `stats --container --name <container name>`

This command shows the following information:

```

Container Name           :<name of the container>
Container ID             :<ID associated with container>
Total Inodes             :<total number of data structures in container>
Read Throughput         :<read throughput rate in Mebibytes or MiB/s for container>
Write Throughput        :<write throughput rate in MiB/s for container>
Current Files           :<current number of files in container>
Current Files Stubbed   :<current number of stubbed files in container>
Current Bytes           :<current number of ingested bytes in container>
Cleaner Status          :<current space reclamation process status for selected container>
Current Recycle Bin Files :<when enabled, current number of files in Recycle Bin>
Current Recycle Bin Logical Bytes :<when enabled, current number of logical bytes in
Recycle Bin>
Current Immutable Files :<current number of files for which immutability is set, RDA
containers only>
Current Immutable Logical Bytes:<current number of immutable logical Bytes>

```

For more information on QoreStor CLI commands, see the *QoreStor System Command Line Reference Guide*.

Adding a cloud tiering policy

Once a cloud tier is created, you must enable a cloud tiering policy for a container before that container can be replicated to the cloud. Refer to [Cloud tiering](#) for more information.

1. In the navigation menu, click **Containers**. The **Containers** page is displayed.
2. In the list of containers, find the correct container and then click **Enable Cloud Tiering Policy**.

3. In the **Enable Cloud Tiering Policy** page, enter the following information:

- If you have more than one Cloud Tier, select the required Cloud Tier that you wish to use for this container from the dropdown.
- **Cloud Policy**-- allows you to configure the amount of time that files must be idle before being sent to the cloud as well as how long to keep a local copy of the files. The configured times should be less than the backup expiry or retention time configured in the backup application. When the backup is expired by the backup application, the files corresponding to the backup are removed from local storage and cloud storage.
 - **Idle time before cloud migration**--Specify the number of hours files must be idle before being sent to the cloud.
 - **On-Prem Retention Age**--Specify the number of hours files will be kept locally after they are sent to the cloud.
- **Enable Cloud Locking** - If Cloud locking is enabled for this Cloud Tier, the container can use Cloud Locking. Based on the cloud provider being used there are different options. For AWS and other S3-Comptaible cloud providers, either of the two locking types can be chosen - **Compliance mode**, **Governance mode**. For Azure, the lock type is internally set.

i **NOTE:** Cloud Locking is currently applicable for RDA and Object container types. For the RDA container, the objects in the cloud are locked for the same immutability period as set on the backup images by the backup application. The locking mode is as per the user choice. For the Object container, the period and mode are both as set on the backup images by the backup application.

- Click **Advanced Options** to view the following:

i **NOTE:** The Advanced Options settings are not supported for object containers.

- **Folder Paths**--files can be included or excluded based on their location.
 - **Include**--only files in the listed directories will be included for replication.
 - **Exclude**-- all files except those in the listed directories will be replicated.
- **File Extensions/Regular Expressions**--files can be included or excluded based on their extension or by regular expressions. Refer to [Using regular expressions](#) for more information.
 - **Include**--only files with the specified extensions will be included for replication.
 - **Exclude**--all files except those with the specified extensions will be replicated.
 - **Include Regular Expression**-- includes only files matching the entered regular expression.
 - **Exclude Regular Expression**-- excludes only files matching the entered regular expression.
- **Stub Exclude Dir/Extensions/Regular Expressions (Optional)** lets you exclude stubs based on their directory, extension, or regular expression. Refer to [Using regular expressions](#) for more information.
 - **Stub Exclude Directory List.** Excludes all stubs in the specified directory.
 - **Stub Exclude File Extension List.** Excludes all stubs with the specified file extension.
 - **Stub Exclude File Regular Expression.** Excludes only stubs matching the entered regular expression.

4. Click **Enable**.

i **NOTE:** To add a cloud tiering policy through the command line, use the command **container --cloud_policy**. Refer to the *QoreStor Command Line Reference Guide* for more information.

Using regular expressions

QoreStor provides an option to select files for inclusion (or exclusion) based file names, character sets, and extensions using regular expressions.

Example expressions

- To include all the files with filename starting with "BackupImage" and having extension .tar, .gz, or .zip.
`^Backupimage.*(zip|tar|gz)`
- To perform a simple match of backup images with IMG | HDR | META
`(IMG|HDR|META)`
- Support for complex regex patterns using character class, meta sequence, and range
 - To match files with filenames starting with **[abn]** and having file extension .txt, or.log
`^[abn]\S*. (txt|log)`
 - To match file names with only alphanumeric characters. For example, the expression below matches *test1*, *f123file* etc.,
`[[:alnum:]]+`
 - Archive series of log files. For example, the expression below matches the first five logs of **ocafsd** & **aws** logs
`ocafsd[1-5].log|aws[1-5].log`

Limitations

The current implementation of regular expressions does not support the use of global modifiers such as **/g** or **/i**. In order to perform case-insensitive matching, a combination of character classes may be used. For example, the expression below can be used to find all combinations of quest, QUEST, QuEST, QUES, etc.

`[Qq][Uu][Ee][Ss][Tt]`

Adding an Archive Tiering policy

Once an archive tier is created, you must enable an archive tiering policy for a container before that container can be archived to the cloud. Refer to [Cloud tiering](#) for more information.

Archive tiering is available for RDA and VTL container types.

To add an archive tiering policy

1. In the navigation menu, click **Containers**. The **Containers** page is displayed.
2. On the **Containers** pane, find the container you wish to archive. Click the ellipses icon, and click **Enable Archive Tiering Policy**.

3. In the **Enable Archive Tiering Policy** page, enter the following information:
 - **Archive Policy** allows you to configure the amount of time that files must be idle before being sent to the cloud as well as how long to keep a local copy of the files.
 - **Idle time before archive migration**. Specifies the number of hours files must be idle before being archived.
 - **On-Prem Retention Age**. Specifies the number of hours files will be kept locally after they are archived.
4. **Enable Cloud locking** : If Cloud locking is enabled for this Archive Tier, the container can use Cloud Locking. Either of the two locking types can be chosen - Compliance mode or Governance mode.
5. Click **Advanced Options** to view the following:
 - **Folder Paths** lets you to select files to be included or excluded based on their location.
 - **Include**. Only files in the listed directories will be included for replication.
 - **Exclude**. All files except those in the listed directories will be replicated.
 - **File Extensions/Regular Expressions** files can be included or excluded based on their extension or by regular expressions. Refer to [Using regular expressions](#) for more information.
 - **Include**. Only files with the specified extensions will be included for replication.
 - **Exclude**. All files except those with the specified extensions will be replicated.
 - **Include Regular Expression**. Includes only files matching the entered regular expression.
 - **Exclude Regular Expression**. Excludes only files matching the entered regular expression.
 - **Stub Exclude Dir/Extensions/Regular Expressions (Optional)** lets you exclude stubs based on their directory, extension, or regular expression. Refer to [Using regular expressions](#) for more information.
 - **Stub Exclude Directory List**. Excludes all stubs in the specified directory.
 - **Stub Exclude File Extension List**. Excludes all stubs with the specified file extension.
 - **Stub Exclude File Regular Expression**. Excludes only stubs matching the entered regular expression.
6. Click **Enable**.

Deleting a container

There are two ways you can delete a container in QoreStor: through the user interface (GUI) or through the command line interface (CLI). The following procedures assume that the containers contain data.

CAUTION: Before deleting a container, you should first carefully consider whether you need to preserve the data in the container. Before deleting any QoreStor container that contains important data, you should take steps to preserve this data using another means of long-term retention. Once a container is deleted, the important data cannot be retrieved.

NOTE: When deleting a container that has Recycle Bin enabled, you cannot delete the container until Recycle Bin is empty. Be sure to delete or expire all backups from the container and wait for the files to expire from Recycle Bin. When Recycle Bin is empty and the container has no more files, delete the container.

i | **NOTE:** No buckets should be present in the Object container for it to be deleted.

Deleting a container through the GUI

i | **NOTE:** When there are files in the container, deleting the container through the GUI fails. If you want to delete a container and the files within it, use the Command Line Interface.

i | **NOTE:** In case a container has an immutable file or locked file in the cloud, deleting the container is not allowed. The stat can be updated by executing the following command with the respective options and the delete operation can be retried.

```
container --update --name <container name>
  [--immutable_file_stat]
  [--locked_file_stat]
```

To delete a container through the GUI

1. On the navigation menu, click **Storage Groups**. Select the storage group that contains the container you want to delete. (If you are only using the DefaultGroup storage group you do not need to select a group.)
2. In the list of containers, select the container you want to delete, and then click the **Delete** icon.
3. In the **Warning** dialog box, click **Yes** to confirm the deletion.

Deleting a container through the command line

To delete a container through the command line

1. Access the QoreStor CLI. Refer to [Accessing the CLI commands on a Linux QoreStor server](#) for more information.
2. Delete a container with the command

```
container --delete --name <name> [--delete_files]
```

Refer to the *QoreStor Command Line Reference Guide* for more information.

Managing local storage

To organize your data, you can easily create storage groups and create containers within those storage groups on your QoreStor system. A storage group allows you to create separate storage policies for different data groups and the different capacities utilized on a single QoreStor. You can also create and organize storage groups for the different organizations in your enterprise, such as Engineering, Sales, Finance, and so on.

After initialization, QoreStor contains a single default storage group, named DefaultGroup.

Refer to these important notes about storage groups.

- Only administrator users can create storage groups.
- Data/containers cannot be moved between storage groups.
- Deduplication is defined at the storage group level and is not global to the appliance.
- Encryption is defined separately for each storage group.
- Compression is defined separately for each storage group.
- The system cleaner cannot be run on a single storage group; only at the system level.
- Before you can delete a storage group, you must delete all containers contained in that storage group.
- A filesystem scan can be run on a single storage group.

To view the Storage Groups page, on the left navigation menu, click **Local Storage > Storage Groups** in the QoreStor navigation pane.

i | **NOTE:** This section refers to using QoreStor with local storage. Please see [Managing cloud storage](#) for information on cloud and archive tiers as well as object storage.

Viewing storage group information

In the QoreStor GUI you can easily view all of your storage groups on the Storage Groups page.

To view storage groups, complete the following steps.

1. In the navigation menu, click **Local Storage** to expand the menu, then **click Storage Groups**.
2. On the **Storage Groups** page you can view the following columns of information about your storage groups.
 - **Name**—Displays the name of the Storage Group.
 - **Encryption**—Displays whether Encryption is turned On or Off.
 - **Compression**—Displays the compression type as either Fast or Best.
 - **Containers**—Displays the number of containers in this storage group. You can click this number, which links to the Containers page for the storage group.

3. In the list of storage groups, find the storage group for which you want to view statistics, and then click the  (ellipsis icon) in the upper-right corner of the information pane.
4. Click **Details**. A page for the selected storage group is displayed, showing the following:
 - **Summary**— Provides summary data for the storage group, including:
 - Name
 - Compression Mode
 - Number of Containers
 - Encryption status
 - Storage group metadata
 - Record Cleaner—
 - **Throughput**—Displays the number of read data in Mebibytes/per second (MiB/s) based on time (in minutes), and the number of write data in MiB/s based on time (in minutes)
 - **Statistics**—Lists the following statistics for the storage group.
 - Cleaner Status
 - Total Inodes
 - Current Files
 - Current Bytes
 - Decrypted Bytes
 - Post Compression Bytes
 - Post Dedupe Bytes
 - Post Encryption Bytes
 - Compression Savings
 - Dedupe Savings
 - Total Savings
 - **Processed Cleaner**—
 - **Log Cleaner**—

Adding a storage group

You can add a storage group to QoreStor through either the QoreStor GUI or the command line interface. In both cases, when creating a storage group you define the name and compression level.

Adding a storage group through the GUI

To add a storage group, complete the following steps.

1. In the navigation menu, click **Local Storage** to expand the menu, then click **Storage Groups**.
2. Click **Add Storage Group**. The Add a Storage Group pane is displayed.
3. In the **Name** field, enter a name for the storage group.
4. Select a **Compression Type** from the drop-down list:
 - **Fast** — Results in shorter backup time, but with less space savings.
 - **Best** — Provides the highest space savings, but with a longer backup time.
5. Optionally, configure a **Quota** by entering an amount and setting the unit (GiB or TiB). If no value is set, the quota will be unlimited.

i **NOTE:** Quotas define the limit for limit for physical capacity usage. Once reached, data ingests are not allowed until space is recovered through savepoint expiration and cleaner operations.

6. To apply encryption, select **Encryption** and enter the following:
 - **Passphrase** — the passphrase is user-defined and is used to generate a passphrase key that encrypts the file in which the content encryption keys are kept. The passphrase is a human readable key, which can be up to 255 bytes in length. It is mandatory to define a passphrase to enable encryption.

i **NOTE:** It is mandatory to define a passphrase to enable encryption. If the passphrase is compromised or lost, the administrator should change it immediately so that the content encryption keys do not become vulnerable.

- **Confirm Passphrase** — re-enter the passphrase used above.
- **Encryption Mode** — Select either **static** or **internal**.
 - **static** - A global mode of key management in which a fixed key is used to encrypt all data.
 - **internal** - A mode of key lifecycle management in which the keys are periodically generated and rotated. The minimum key rotation period before the content encryption key can be rotated and a new key is generated is 7 days. This rotation period is user-configurable and can be specified in days.

i **NOTE:** Refer to [Configuring and Using Encryption at Rest](#) for more information about encryption.

i **NOTE:** Due to export regulations, the encryption at rest feature is not available in certain markets, and, therefore, may not be available in your locale. For more information about recommended guidelines for encryption, see [Understanding Encryption at Rest](#)

i **NOTE:** After encryption is enabled, all of the data that is backed up is encrypted and is kept encrypted until it is expired and cleaned by the system cleaner. **Note that encryption is an irreversible process.**

7. Click **Add**

Adding a storage group through the command line

To add a storage group, complete the following steps.

1. Access the QoreStor CLI. Refer to [Accessing the CLI commands on a Linux QoreStor server](#) for more information.
2. Add a new storage group using the command

```
storage_group --add --name <name> [--compression_mode <fast|best>] [--quota <Quota value in GiB or TiB>]
```

Refer to the *QoreStor Command Line Reference Guide* for more information

3. To apply encryption to the data in this storage group, use the command:

```
storage_group --encryption --name <name> [--set <ON | OFF>] [--mode <static|internal> <--interval <7 days to 70 years>]
```

For more information, refer to the **Storage Group commands** section of the *QoreStor Command Line Reference Guide*.

- i** **NOTE:** Due to export regulations, the encryption at rest feature is not available in certain markets, and, therefore, may not be available in your locale. For more information about recommended guidelines for encryption, see [Understanding Encryption at Rest](#)
- i** **NOTE:** It is mandatory to define a passphrase to enable encryption. If the passphrase is compromised or lost, the administrator should change it immediately so that the content encryption keys do not become vulnerable.
- i** **NOTE:** After encryption is enabled, all of the data that is backed up is encrypted and is kept encrypted until it is expired and cleaned by the system cleaner. **Note that encryption is an irreversible process.**

Modifying a storage group

To modify a storage group via the user interface, complete the following steps

1. In the navigation menu, click **Local Storage** to expand the menu, then click **Storage Groups**.
2. Find the desired storage group from the displayed list. Click the  (ellipsis icon) in the upper-right corner of the storage group's information pane.
3. Click **Edit**.

4. For Storage Optimization, select a **Compression Type** from the drop-down list as needed:
 - **Fast**— Results in shorter backup time, but with less space savings.
 - **Best** — Provides the highest space savings, but with a longer backup time.
 - Optionally, configure a **Quota** by entering an amount and setting the unit (GiB or TiB). If no value is set, the quota will be unlimited. If a quota is currently set, you may increase or decrease the quota value (provided the decreased quota is not less than the current used capacity on the storage group). Once a quota is set, it can be reset to the default value of unlimited.

i **NOTE:** Quotas define the limit for physical capacity usage. Once reached, data ingests are not allowed until space is recovered through savepoint expiration and cleaner operations.

5. You can modify the following **Encryption** settings:

i **NOTE:** For more information about recommended guidelines for setting up encryption, see the topic, [Configuring and Using Encryption at Rest](#).

i **NOTE:** Due to export regulations, the encryption at rest feature is not available in certain markets, and, therefore, may not be available in your locale.

- **Encryption**—Next to **Encryption**, select or clear as needed.
- **Old Passphrase**—Enter the current passphrase you want to change.
- **New Passphrase**—Enter the new passphrase to be used to encrypt content encryption keys. (The passphrase string can take up to 255 characters. And, alphanumeric and special characters can be entered as part of the passphrase string.)
- **Confirm Passphrase**—Re-enter the encryption passphrase.
- **Encryption Mode**—Select the mode of key lifecycle management from one of the following options:
 - **Static**—A global, fixed key is used to encrypt all data.
 - **Internal**—Content encryption keys are generated and rotated on a specified period of days.

6. Click **Update**.

To modify a storage group via the CLI, complete the following steps

1. Access the QoreStor CLI. Refer to [Accessing the CLI commands on a Linux QoreStor server](#) for more information.
2. Modify your storage group using the commands below. Refer to the *QoreStor Command Line Reference Guide* for more information.

```
storage_group --show [--name <name>] [--verbose]
storage_group --update --name <name> [--compression_mode <fast|best>] [--quota
<Quota value in GiB or TiB>]
storage_group --encryption --name <name> [--set <ON | OFF>] [--mode
<static|internal> <--interval <7 days to 70 years>]
storage_group --setpassphrase --name <name>
```

Deleting a storage group

Before you can delete a storage group, you must first delete the containers in the storage group. See [Deleting a container](#) for more information.

Deleting a storage group from the GUI

To delete a storage group, complete the following steps.

1. In the navigation menu, click **Local Storage** to expand the menu, then click **Storage Groups**.
2. Find the desired storage group from the displayed list. Click the  (ellipsis icon) in the upper-right corner of the storage group's information pane.
3. Click **Delete**.
 **NOTE:** You cannot delete the DefaultGroup storage group.
4. When prompted to confirm, click **Delete**.

Deleting a storage group from the CLI

Access the QoreStor CLI. Refer to [Accessing the CLI commands on a Linux QoreStor server](#) for more information.

1. Delete your storage group using the command below. Refer to the *QoreStor Command Line Reference Guide* for more information.

```
storage_group --delete --name <name>
```

Configuring a Performance Tier

Before configuring a performance tier, ensure that you have reviewed the requirements and limitations below.

- Only one volume can be mapped to a performance tier.
- Your performance tier volume must be mapped to high-performance storage.
- Your performance tier volume must be mapped to an XFS file system.
- An existing QoreStor repository may not be mapped to a performance tier.
- Once a performance tier has been created, the storage path cannot be changed.
- Only one performance tier can be added to QoreStor.
- Currently, it is not possible to delete the performance tier.
- The mount location for the high-performance storage must be added to your QoreStor system's file systems table (fstab) to mount the storage automatically when the server reboots.
Add mount location to */etc/fstab* using the **noatime** and **dirsync** mount options

```
echo '<storage path> /perf01 xfs defaults,noatime,dirsync 0 0' >> /etc/fstab
```

Adding a performance tier

Adding a performance tier can be accomplished through the QoreStor UI or via the **performance_tier** command in the QoreStor CLI. Refer to the *QoreStor Command Line Reference Guide* for more information on the **performance_tier** command.

To add a performance tier

Before adding a performance tier, ensure that the storage is mounted to the QoreStor server as an XFS filesystem. Refer to [Guidelines for configuring additional storage](#) for additional guidelines and requirements.

1. In the navigation menu, click **Local Storage** to expand the menu, then click **Performance Tier**.
2. Click **Add Performance Tier**.
3. Enter the mount path for the performance tier volume.
4. Optionally, click **Test** to ensure this storage meets the performance requirements. You will be prompted that the test can take several minutes, click **Confirm** to continue.
5. Click **Add**.
When a performance tier is added, QoreStor will also create the corresponding storage group "PerformanceTier" mapped to the performance tier storage.
6. A warning will be displayed notifying you that a service restart may be required. Click **Confirm** to continue.
i | **NOTE:** QoreStor services will be restarted. You will need to log into the UI once the restart is complete.
7. Once the performance tier is created, you can enable encryption. Refer to [Editing a performance tier](#).

Adding a performance tier through the command line

To add a performance tier, complete the following steps.

1. Access the QoreStor CLI. Refer to [Accessing the CLI commands on a Linux QoreStor server](#) for more information.
2. Add a performance tier using the command

```
performance_tier --add --path <enclosure filesystem dir> [--compression_mode <fast|best>] [--quota <Quota value in GiB or TiB>]
```

Refer to the *QoreStor Command Line Reference Guide* for more information

3. To apply encryption to the data in this performance tier, use the command:

```
performance_tier --encryption [--set <ON | OFF>] [--mode <static|internal> <--interval <7 days to 70 years>]
```

For more information, refer to the *QoreStor Command Line Reference Guide*.

- i** | **NOTE:** Due to export regulations, the encryption at rest feature is not available in certain markets, and, therefore, may not be available in your locale. For more information about recommended guidelines for encryption, see [Understanding Encryption at Rest](#)
- i** | **NOTE:** It is mandatory to define a passphrase to enable encryption. If the passphrase is compromised or lost, the administrator should change it immediately so that the content encryption keys do not become vulnerable.
- i** | **NOTE:** After encryption is enabled, all of the data that is backed up is encrypted and is kept encrypted until it is expired and cleaned by the system cleaner. **Note that encryption is an irreversible process.**

Editing a performance tier

To modify a performance tier via the user interface, complete the following steps

1. In the navigation menu, click **Local Storage** to expand the menu, then click **Performance Tier**
 2. Click **Edit Performance Tier**.
 3. On the **Edit Performance Tier** page, select or modify the options below, as appropriate:
 - **Unlimited Quota**— This option is selected by default. To set a quota, de-select this option and enter a numeric value (greater than 100) in the **Quota** field. Select the value format (either GiB or TiB). Refer to [Quotas](#) for more information. Once a quota is set, it can be reset to the default value of unlimited.
 - **Encryption**— Select this option to enable encryption on the performance tier.
 - **Old Passphrase**—Enter the current passphrase you want to change.
 - **New Passphrase**—Enter the new passphrase to be used to encrypt content encryption keys. (The passphrase string can take up to 255 characters. And, alphanumeric and special characters can be entered as part of the passphrase string.)
 - **Confirm Passphrase**—Re-enter the encryption passphrase.
 - **Encryption Mode**— Select the mode of key lifecycle management from one of the following options:
 - **Static**— A global, fixed key is used to encrypt all data.
 - **Internal**— Content encryption keys are generated and rotated on a specified period of days.
 - **Key Rotation**— Displays the number of key rotation interval days as N/A, or the number that was set for Internal Encryption Mode. If you selected Internal as the mode of key management, select the number of days for key rotation when a new key is to be generated. This option is available only for **Internal** encryption mode.
- NOTE:** After encryption is enabled, all of the data that is backed up is encrypted and is kept encrypted until it is expired and cleaned by the system cleaner. Encryption is an irreversible process.
4. Click **Update**.

To modify a performance tier via the CLI, complete the following steps

1. Access the QoreStor CLI. Refer to [Accessing the CLI commands on a Linux QoreStor server](#) for more information.
2. Modify your performance tier using the command below. Refer to the *QoreStor Command Line Reference Guide* for more information.

```
performance_tier --update [--compression_mode <fast|best>] [--quota <Quota value in GiB or TiB>]
```

```
performance_tier --encryption [--set <ON | OFF>] [--mode < static | internal >] [--interval <7 days to 70 years>]
```

Configuring Object Container

QoreStor's Object container provides an object storage interface which enables customers to write Object data(S3 format) directly to QoreStor. This allows solutions that leverage an S3-based connection to send data directly to a QoreStor instance instead of Amazon S3 with the added benefits of deduplication, encryption, replication and network optimized data transfer.

Object storage is configured by adding a container with the Object (S3 Compatible) protocol.

Creating an Object Container

Adding an object container can be accomplished through the QoreStor UI or via the **object_container** command in the QoreStor CLI. Refer to the *QoreStor Command Line Reference Guide* for more information on the **object_container** command.

i **NOTE:** QoreStor object container does not support object lifecycle management, which means transitioning storage classes or server side expiration of objects is not supported. User policies are limited to predefined readwrite, writeonly, and readonly.

To create an object container

1. In the navigation menu, click **Containers**.
2. On the **Containers** pane, click **Add Container**. The Add Container dialog will be displayed.
3. In the Protocol field, select **Object (S3 Compatible)**.
4. In the **Storage Group** drop-down, select the required storage group for this container.
5. Click **Next**.
6. Optionally, select **Use HTTP instead of HTTPS**. To use an HTTP connection, you must also follow the steps below:
 - a. On the QoreStor server, copy the aws.conf file to a new location:

i **NOTE:** The QoreStor implementation of object storage uses a self-signed certificate. If your data management application requires third party certificates, you must use HTTP to connect to the object container.

7. Click **Next**.
8. Review the summary and click **Finish**.

When the process is completed the object container is added to the QoreStor. For Object container created prior to QoreStor release 7.2.1 you will see the storage group **ObjectContainer** and the container **ObjectStorageGroup** added to the **Storage Groups** and **Container** pages, respectively. See the topics below for information on working with object storage.

- [Creating a bucket](#)
- [Managing Object container users](#)
- [Editing bucket settings](#)

Adding an object container through the command line

To add an object container, complete the following steps.

1. Access the QoreStor CLI. Refer to [Accessing the CLI commands on a Linux QoreStor server](#) for more information.
2. Add a Object container:

```
object_container --add --name <container name> [--group <storage group name>]
```

Refer to the *QoreStor Command Line Reference Guide* for more information.
3. Get end-point details of it:

```
object_container --show --name <container name>
```
4. Create user for this container. This user name is used as Access key and user's password is used as Secret key while accessing Object container from the client systems (backup clients):

```
object_container --user-add --name <name> --user-name <user name>
```

i **IMPORTANT:**The User's name is used as Access Key and the user's password is used as Secret Key while connecting to QoreStor from the S3 clients.
To see the S3 endpoint, use the command `object_container --show --endpoint --name <name of container>`
The endpoint is displayed in the format `https://<QoreStor IP address>:<port>`
Make sure the port is allowed for access through the firewall.
5. Set access policy for the user. Use <Policy name> as "readwrite" to allow the user to backup and restore data.

```
object_container --policy-set --name <name> --policy-name <Policy name> --user-name <user name>
```
6. Create bucket for use in the backup application. Optionally add locking support:

```
object_container --bkt-add --name <name> --bkt-name <bucket name> [--enable-object-lock] [--enable-object-versioning]
```
7. Configure the backup application with the endpoint, access key, secret key, and bucket name.

Creating a bucket

In S3 compatible storage, buckets are organizational containers that store objects. When creating a bucket, you have the option to enable or disable object locking, and select one of the available Object Locking modes. Object locking settings apply to all objects in the bucket.

- **Governance** mode - prevents users without the appropriate permissions from overwriting or deleting an object version or altering its lock settings. With governance mode, you protect objects against being deleted by most users, but you can still grant some users permission to alter the retention settings or delete the object if necessary. You can also use governance mode to test retention-period settings before creating a compliance-mode retention period.
- **Compliance** mode - prevents objects from being overwritten or deleted by any user during the specified lock period. When an object is locked in compliance mode, its retention mode can't be changed, and its retention period can't be shortened. Compliance mode ensures that an object version can't be overwritten or deleted for the duration of the retention period.
- **None** - no restrictions are applied.

i | **NOTE:** QoreStor supports a maximum of 1000 buckets. The bucket **default-bucket** is created automatically when the object container is created.

To create a bucket

1. In the navigation menu, click **Containers**.
2. On the **Containers** pane, find the object storage container that you want to edit. Click the ellipses icon, and click **Edit**.
3. Click **Create bucket**.
4. Enter a **Name** for your bucket.
5. Optionally, select **Object Locking** and configure
 - **Locking Mode** - select between **Compliance** and **Governance**.
 - **Locking Duration** - select the number and format (days or years) to specify the time that the object lock will be active.

i | **IMPORTANT:** The **Object Locking** status of a bucket cannot be changed once the bucket is created. To ensure flexibility in the future, you may set the object locking status to **enabled**, but the locking mode to **None**. If the locking mode is set to **disabled**, you will not be able to edit the bucket settings in the future

6. Click **Save**.

Editing bucket settings

The settings for existing buckets can be edited on the **Object Storage** page. When changing bucket settings using this procedure, the changes are applied to new objects. To change retention settings on existing objects, refer to [Changing bucket retention](#).

To edit bucket settings

1. On the **Containers** pane, find the object storage container that you want to edit. Click the ellipses icon, and click **Edit**.
2. In the Buckets section, find the desired bucket. Right-click on the ellipsis icon and click **Bucket Settings**.

3. In the **Locking Mode** section, select from one of the options below:
 - **Compliance** - prevents objects from being overwritten or deleted by any user during the specified lock period. When an object is locked in compliance mode, its retention mode can't be changed, and its retention period can't be shortened. Compliance mode ensures that an object version can't be overwritten or deleted for the duration of the retention period.
 - **Governing** - prevents users without the appropriate permissions from overwriting or deleting an object version or altering its lock settings. With governance mode, you protect objects against being deleted by most users, but you can still grant some users permission to alter the retention settings or delete the object if necessary. You can also use governance mode to test retention-period settings before creating a compliance-mode retention period.
 - **None** - provides no object locking.

i | **IMPORTANT:** Changes in the locking mode are applied to files written to the container after the change was applied. Existing files are locked according to the locking method in place at the time they were added to the container.

4. In the **Locking Duration** field, select the number and format (days or years) to specify the time that the object lock will be active.
5. Click **Update**.

Changing bucket retention

To change retention settings for existing objects, follow the procedures below. To change the default bucket settings for new objects, refer to [Editing bucket settings](#)

To change bucket retention settings

1. On the **Containers** pane, find the object storage container that you want to edit. Click the ellipses icon, and click **Edit**.
2. In the Buckets section, find the desired bucket. Right-click on the ellipsis icon and click **Change Retention**.
3. [Optional] In the **Prefix** section, enter the folder prefix. This will limit retention changes to objects in the specified folder.
4. In the **Locking Mode** section, select from one of the options below:
 - **Compliance** - prevents objects from being overwritten or deleted by any user during the specified lock period. When an object is locked in compliance mode, its retention mode can't be changed, and its retention period can't be shortened. Compliance mode ensures that an object version can't be overwritten or deleted for the duration of the retention period.
 - **Governing** - prevents users without the appropriate permissions from overwriting or deleting an object version or altering its lock settings. With governance mode, you protect objects against being deleted by most users, but you can still grant some users permission to alter the retention settings or delete the object if necessary. You can also use governance mode to test retention-period settings before creating a compliance-mode retention period.
 - **None** - provides no object locking.
5. In the **Locking Duration** field, select the number and format (days or years) to specify the time that the object lock will be active.
6. Click **Update**.

Managing Object container users

This dialog allows you to manage the Object container users. Each Object container user is managed separately and is different from the other QoreStor users.

Currently, one of the predefined policies can be assigned to each user. User name and password are to be used in place of Access Key and Secret Key respectively when configuring access to this container from the DMA.

To add an Object container user

1. In the navigation menu, click **Containers**.
2. On the **Containers** pane, find the object storage container that you want to update. Click the ellipses icon, and click **Edit**.
3. Click **Users**.
4. Click **Add User**.
5. To add a user, enter user name, Secret key (password) and select a policy from the dropdown. The available policies are:
 - **readwrite**
 - **readonly**
 - **writeonly**
6. Click **Save**.

i | **NOTE:** You may also configure object container user policies using the QoreStor CLI, Refer to the QoreStor Command Line Reference Guide for usage details on the **object_container** command option.

Updating Object container user

This dialog allows you to update a user's secret key or policy.

To update

1. In the navigation menu, click **Containers**.
2. On the **Containers** pane, find the object storage container that you want to update. Click the ellipses icon, and click **Edit**.
3. Select the **User** from the current list of users and click the ellipses icon under **Actions**. Select **Edit**.
4. Update **Secret Key** or **Policy** as required and click **Save**.

Deleting Object container user

This dialog allows you to delete an existing Object container user.

To delete the user

1. In the navigation menu, click **Containers**.
2. On the **Containers** pane, find the object storage container that you want to update. Click the ellipses icon, and click **Edit**.
3. Select the **User** from the current list of users and click the ellipses icon under **Actions**. Select **Delete**.
4. Confirm the removal in the dialog window.

Configuring additional storage

QoreStor allows for the configuration of additional storage enclosures in order to increase capacity.

Guidelines for configuring additional storage

Refer to the following important notes and guidelines for understanding storage expansion in QoreStor.

- For Standard and Cloud-Optimizeed installation modes, QoreStor supports a maximum of 5 locations (1 internal plus 4 expansion). When installed in Large mode, QoreStor supports a maximum of 7 locations.
- QoreStor can only add a mounted filesystem path. QoreStor will not create a filesystem if one is not present.
 - The mounted filesystem should be XFS type.
 - Mount path cannot be read-only.
 - The mount path should not be already configured enclosure path.
- When adding storage, the filesystem service will be restarted.
- Storage can be added without being licensed, but an appropriately sized license is required to make the storage usable by QoreStor.
 - Licenses can be added in multiples of Terabytes from 1 TB to the maximum capacity for your QoreStor installation mode.

Adding additional storage

Additional storage can be added through the QoreStor UI or via the **system --storage** command in the QoreStor CLI. Refer to the *QoreStor Command Line Reference Guide* for more information on the **system --storage** command.

i | **IMPORTANT:** Adding storage requires QoreStor services to be restarted. This will take the QoreStor server offline for several minutes.

To add storage

Before adding storage, ensure that the storage is mounted to the QoreStor server as an XFS filesystem. Refer to [Guidelines for configuring additional storage](#) for additional guidelines and requirements.

1. In the navigation menu, click **Local Storage** to expand the menu, then click **Storage Groups**.
2. Click **Add Storage Location**.

3. Enter the mount path for the additional storage. Click **Save**.
4. A warning will be displayed notifying you that a service restart may be required. Click **Confirm** to continue.

Managing cloud storage

QoreStor enables multiple options for extending your on-prem environment to the cloud. Following table shows supported storage capacities of Cloud Tier and Archive Tier across different QoreStor installation modes.

Table 7: Cloud Tier and Archive Tier capacity matrix

| Install Mode | Maximum Allowed or Used Capacity | | |
|-----------------|--|------------|--------------|
| | Local storage or Object Direct storage | Cloud Tier | Archive Tier |
| Cloud optimized | 43 TB | 129 TB | NA |
| Standard | 150 TB | 450 TB | 1500 TB |
| Enterprise | 360 TB | 1080 TB | 3600 TB |
| Enterprise-plus | 512 TB | 1536 TB | 5120 TB |

Please refer to the sections below for additional information.

- [Understanding cloud and archive storage](#)
- [Cloud tiering](#)
- [Configuring a Cloud Archive Tier](#)
- [Performing a disaster recovery from the cloud](#)

Understanding cloud and archive storage

There are two ways QoreStor lets you use cloud storage containers: as secondary, policy-based storage or as primary storage with optimized deduplication.

Policy-based cloud storage

Policy-based cloud storage is available to any compatible backup software. It requires that you configure a schedule for when QoreStor automatically transfers backups to cloud storage.

There are three purposes for using policy-based cloud storage:

- Replication
- Tiering

- Archiving

For replication and tiering, the cloud storage group remains online, keeping the data accessible. For archiving, QoreStor uses offline Glacier storage, which is designed to store data for longer terms.

For more information about creating a cloud or archive tier, see [Cloud tiering](#).

Direct-to-cloud storage

Direct-to-cloud storage is the feature that QoreStor offers which lets you store backed up data directly to a cloud container. The container can be in either the cloud storage group or the archive storage group. This feature is available only for NetVault backups. After you create the RDS container, you can enter the name of this container in the NetVault GUI as your primary storage location. QoreStor then automatically saves the backups to the cloud, performing optimized deduplication them during the process. Because the backups go directly to the container, QoreStor does not use the policy manager.

When you back up data from NetVault to a direct cloud container, the container behaves the same as any other QoreStor RDS container. The data is easily accessible, which lets you perform the following operations:

- Listing
- Editing
- Backup
- Restore

For more information about creating a direct-to-cloud container, see [Configuring an RDS direct-to-cloud container](#).

Limitations of direct-to-cloud containers

Direct cloud containers include the following limitations:

- Only backup applications using the RDA container type are compatible with direct-to-cloud containers. It is recommended to use them for storing copies of backups.
- **You can access and control the direct cloud container only through QoreStor.** Although only NetVault is compatible with the direct cloud container feature, the container does not appear as a repository in the NetVault UI and NetVault is not aware of its existence. The backups still pass through QoreStor to reach their destination.
- **Archive tiers do not function as direct cloud containers.** Verification is required for direct cloud containers, but data in Glacier containers is offline and inaccessible.
- **Quest discourages using direct cloud containers as sources.** When replicating, QoreStor recommends that you always use the RDS direct container as a target rather than a source repository.
- Cloud locking is not supported with direct to cloud containers.
- Direct cloud containers cannot be used with container continuous replication.

Configuring an RDS direct-to-cloud container

You can configure an RDS container as a direct-to-cloud container using either the QoreStor UI or the command line interface (CLI).

Configuring a direct-to-cloud container in the QoreStor UI

To add a direct cloud container in the QoreStor UI

1. In the navigation menu, click **Containers**.
2. On the Containers page, click **Add Container**.
3. In the Add Container window, select the following options:
 - a. For Protocol, select **Quest Rapid Data Storage (RDS)**.
 - b. For Name, enter the name of the direct cloud container.
 - c. For Storage Group, select one of the following options:
 - **DefaultCloudTier - Online Cloud Storage**
 - **DefaultCloudArchiveTier - Offline Long-term Cloud Storage**
 - d. Click **Next**.
4. On the next page, for Protocol RDS Capacity, select **Unlimited**, and then click **Next**.
5. On the next page, for User, enter the name of the user you want to access the container, and then click **Next**.
6. On the Summary page, confirm your selections, and then click **Finish**.
The container appears on the Cloud Storage page and available to use as storage from the NetVault UI.

Configuring a direct-to-cloud container using the CLI

To configure a direct-to-cloud container using the CLI

1. In the CLI, if you want to create a container in the cloud storage group, enter the following command:

```
container --add --name K<name> --group_name DefaultCloudTier
```

2. If you want to create a container in the archive storage group, enter the following command:

```
container --add --name K<name> --group_name DefaultCloudArchiveTier
```

The container appears on the Cloud Storage page of the QoreStor UI and available to use as storage from the NetVault UI.

Cloud tiering

You can add a cloud tier to QoreStor through either the QoreStor GUI or the command line interface. For more information on Cloud Storage Groups and cloud replication, refer to the following topics:

i | **NOTE:** Starting from the 7.4.1 installation, the Cloud tier can be added optionally. The following section is applicable only if the cloud tier option is selected.

- Adding a Cloud Tier through the GUI
- Adding a cloud tiering policy
- Editing a cloud tiering policy
- Creating a cloud tiering schedule
- Editing a cloud tiering schedule
- Deleting a cloud tier
- Cloud Tier replication

Adding a Cloud Tier through the GUI

i | **IMPORTANT:** Your cloud storage must be configured prior to completing this procedure.

The procedures for adding a cloud tier differ slightly based on the cloud storage provider. Review the procedure for your cloud storage provider:

- Adding a Microsoft Azure cloud tier
- Adding an Amazon S3 cloud tier
- Adding a Wasabi S3 cloud tier
- Adding an IBM S3 cloud tier
- Adding a Google S3 cloud tier
- Adding a Scality-Artasca-S3 cloud tier
- Adding a Backblaze S3 cloud tier
- Adding an S3 Compatible cloud tier

Adding a Microsoft Azure cloud tier

To add a Microsoft Azure cloud storage group, complete the following steps:

1. In the navigation menu, click **Cloud Storage** to expand the menu, then click **Cloud Tier**.
2. Click **Configure**.
3. In the **Cloud Provider** drop-down, select **Azure Blob**.
4. Provide a container name. This is the name of your existing Azure container.
5. Enter your **Cloud Connection String**.

i | **NOTE:** The Connection String can be found in your **Azure** portal under **Storage Accounts >Access keys > Connection string**.

6. **Enable Cloud Locking** : Select this option to achieve immutability of data tiered to the cloud.

i | **NOTE:** Cloud Locking works with RDA and Object containers only. Data that is set as immutable on the

local storage is also set as immutable when tiered to the cloud.

Enabling locking at the Cloud Tier is required so the containers that are tiering data to this Cloud Tier can choose to enable locking as part of their cloud tiering policy.

7. To apply encryption, in the **Cloud Tier Encryption** section enter the following:
 - **Passphrase** — the passphrase is user-defined and is used to generate a passphrase key that encrypts the file in which the content encryption keys are kept. The passphrase is a human readable key, which can be up to 255 bytes in length. It is mandatory to define a passphrase to enable encryption.
 - **IMPORTANT:** It is mandatory to define a passphrase to enable encryption. If the passphrase is compromised or lost, the administrator should change it immediately so that the content encryption keys do not become vulnerable. If this passphrase is lost or forgotten, data in the cloud will be unrecoverable.
 - **Confirm Passphrase** — re-enter the passphrase used above.
8. Click **Configure**. A Cloud Storage Group will be created.
9. To enable replication to the cloud, you must link a local container to the cloud using the procedures in [Adding a cloud tiering policy](#).

Adding an Amazon S3 cloud tier

To add a cloud storage group, complete the following steps

1. In the navigation menu, click **Cloud Storage** to expand the menu, then click **Cloud Tier**.
2. In the Cloud pane, click **Configure** to add a cloud tier.
3. In the **Cloud Provider** drop-down, select **AWS S3**.
4. Provide the name for your S3 bucket.

5. Enter your **Connection String** using one of the two methods below:

- **Default** - this option will compile your connection string into the correct format using the inputs below.
 - **Access key** - The access key is typically 20 upper-case English characters
 - **Secret key** - The secret key is generated automatically by AWS. It is typically 40 characters, including mixed upper and lower-case and special symbols.
 - **Region** - The region specifies the Amazon-specific region in which you want to deploy your backup solution. You can obtain your region code from https://docs.aws.amazon.com/general/latest/gr/rande.html#s3_region.
- **Custom** - this option allows you to enter your connection string with additional parameters.
- Your connection string uses the following syntax:

```
"accesskey=<ABDCEWERS>;secretkey=< >; loglevel=warn; region=<aws-region>;"
```

Please note the following:

- a. The **access key** is typically 20 upper-case English characters
- b. The **secret key** is generated automatically by AWS. It is typically 40 characters, including mixed upper and lower-case and special symbols.
- c. The **region** specifies the Amazon-specific region in which you want to deploy your backup solution. You can obtain your region code from https://docs.aws.amazon.com/general/latest/gr/rande.html#s3_region.

An example of a connection string with this syntax follows. Logically, each connection string is unique.

```
accesskey=AKIARERFUCFODHFJUCWK;secretkey=p+8/T+o5WeZkX11QbuPazHX1IdWbwgFplxuVl08J;loglevel=warn;region=eu-central-1;
```

6. **Enable Cloud Locking:** Select this option to achieve immutability of data tiered to the cloud.

i NOTE: Cloud Locking works with RDA and Object containers only. Data that is set as immutable on the local storage is also set as immutable when tiered to the cloud.

Enabling locking at the Cloud Tier is required so the containers that are tiering data to this Cloud Tier can choose to enable locking as part of their cloud tiering policy.

AWS storage classes can be chosen from the following:

- S3 Standard
- S3 Intelligent-Tiering
- S3 Standard-Infrequent Access
- S3 One Zone-Infrequent Access
- S3 Glacier Instant Retrieval

Please note, apart from S3 Standard, other storage classes provide further cost savings for data that is infrequently accessed.

7. To apply encryption, in the **Cloud TierEncryption** section enter the following:
 - **Passphrase** — the passphrase is user-defined and is used to generate a passphrase key that encrypts the file in which the content encryption keys are kept. The passphrase is a human readable key, which can be up to 255 bytes in length. It is mandatory to define a passphrase to enable encryption.
 - **IMPORTANT:** It is mandatory to define a passphrase to enable encryption. If the passphrase is compromised or lost, the administrator should change it immediately so that the content encryption keys do not become vulnerable. If this passphrase is lost or forgotten, data in the cloud will be unrecoverable.
 - **Confirm Passphrase** — re-enter the passphrase used above.
8. Click **Configure**. A Cloud Storage Group will be created.
9. To enable replication to the cloud, you must link a local container to the cloud using the procedures in [Adding a cloud tiering policy](#).

Adding a Wasabi S3 cloud tier

To add a cloud storage group, complete the following steps:

1. In the navigation menu, click **Cloud Storage** to expand the menu, then click **Cloud Tier**.
2. In the Cloud pane, click **Configure** to add a cloud tier.
3. In the **Cloud Provider** drop-down, select **Wasabi S3**.
4. Provide a container name. This is the existing name of your container in your cloud platform.

5. Enter your **Connection String** using one of the two methods below:

- **Default** - this option will compile your connection string into the correct format using the inputs below.
 - **Access key** - The access key is typically 20 upper-case English characters
 - **Secret key** - The secret key is generated automatically by Wasabi. It is typically 40 characters, including mixed upper and lower-case and special symbols.
 - **Region** - The region specifies the Wasabi-specific region in which you want to deploy your backup solution. You can obtain your region code from <https://wasabi.com/help/docs/>.
 - **Endpoint** - If you are using VPC endpoints, enter the correct endpoint information.
- **Custom** - this option allows you to enter your connection string with additional parameters.
 - Your connection string uses the following syntax:
"accesskey=<ABDCEWERS>;secretkey=< >; loglevel=warn;endpoint=
https://s3.<region>.wasabisys.com;"
 - Wasabi S3 endpoint based on region can be obtained from the following URL:
<https://docs.wasabi.com/docs/what-are-the-service-urls-for-wasabi-different-storage-regions>

Please note the following:

- a. The **access key** is typically 20 upper-case English characters
- b. The **secret key** is generated automatically by Wasabi. It is typically 40 characters, including mixed upper and lower-case and special symbols.
- c. The **region** specifies the Wasabi-specific region in which you want to deploy your backup solution. You can obtain your region code from <https://wasabi.com/help/docs/>.

An example of a connection string with this syntax follows. Logically, each connection string is unique.

```
accesskey=AKIARERFUCFODHFJUCWK;secretkey=p+8/T+o5WeZkX11QbuPazHX1Id  
WbwgFplxuVl08J;loglevel=warn;region=eu-central-1;
```

6. **Enable Cloud Locking:** Select this option to achieve immutability of data tiered to the cloud.

i NOTE: Cloud Locking works with RDA and Object containers only. Data that is set as immutable on the local storage is also set as immutable when tiered to the cloud.

Enabling locking at the Cloud Tier is required so the containers that are tiering data to this Cloud Tier can choose to enable locking as part of their cloud tiering policy.

7. To apply encryption, select **Encryption** and enter the following:
 - **Passphrase** — the passphrase is user-defined and is used to generate a passphrase key that encrypts the file in which the content encryption keys are kept. The passphrase is a human readable key, which can be up to 255 bytes in length. It is mandatory to define a passphrase to enable encryption.

i **IMPORTANT:** It is mandatory to define a passphrase to enable encryption. If the passphrase is compromised or lost, the administrator should change it immediately so that the content encryption keys do not become vulnerable. If this passphrase is lost or forgotten, data in the cloud will be unrecoverable.
 - **Confirm Passphrase** — re-enter the passphrase used above.
8. Click **Configure**. A Cloud Storage Group will be created.
9. To enable replication to the cloud, you must link a local container to the cloud using the procedures in [Adding a cloud tiering policy](#).

Adding an IBM S3 cloud tier

To add a cloud storage group, complete the following steps:

1. In the navigation menu, click **Cloud Storage** to expand the menu, then click **Cloud Tier**.
2. In the Cloud pane, click **Configure** to add a cloud tier.
3. In the **Cloud Provider** drop-down, select **IBM S3**.
4. Provide a container name. This is the existing name of your container in your cloud platform.

5. Enter your **Connection String** using one of the two methods below:

- **Default** - this option will compile your connection string into the correct format using the inputs below.
 - **Access key** - The access key is typically 20 upper-case English characters
 - **Secret key** - The secret key is generated automatically by IBM. It is typically 40 characters, including mixed upper and lower-case and special symbols.
 - **Region** - The region specifies the IBM-specific region in which you want to deploy your backup solution. You can obtain your region code from <https://cloud.ibm.com/docs/cloud-object-storage?topic=cloud-object-storage-endpoints>.
 - **Endpoint** - If you are using VPC endpoints, enter the correct endpoint information.
- **Custom** - this option allows you to enter your connection string with additional parameters.
 - Your connection string uses the following syntax:
"accesskey=<ABDCEWERS>;secretkey=< >; loglevel=warn; region=<ibmS3-region>;endpoint=<IBM-S3 endpoint url>"
 - You can obtain region code and endpoint information from <https://cloud.ibm.com/docs/cloud-object-storage?topic=cloud-object-storage-endpoints>.
Please note the following:
 - a. The **access key** is typically 20 upper-case English characters
 - b. The **secret key** is generated automatically by IBM. It is typically 40 characters, including mixed upper and lower-case and special symbols.
 - c. The **region** specifies the IBM-specific region in which you want to deploy your backup solution. You can obtain your region code from <https://cloud.ibm.com/docs/cloud-object-storage?topic=cloud-object-storage-endpoints>.

An example of a connection string with this syntax follows. Logically, each connection string is unique.

```
accesskey=AKIARERFUCFODHFJUCWK;secretkey=p+8/T+o5WeZkX11QbuPazHX1Id  
WbwgFplxuVl08J;loglevel=warn;region=eu-central-1;
```

6. **Enable Cloud Locking:** Select this option to achieve immutability of data tiered to the cloud.

i NOTE: Cloud Locking works with RDA and Object containers only. Data that is set as immutable on the local storage is also set as immutable when tiered to the cloud.

Enabling locking at the Cloud Tier is required so the containers that are tiering data to this Cloud Tier can choose to enable locking as part of their cloud tiering policy.

7. To apply encryption, select **Encryption** and enter the following:
 - **Passphrase** — the passphrase is user-defined and is used to generate a passphrase key that encrypts the file in which the content encryption keys are kept. The passphrase is a human readable key, which can be up to 255 bytes in length. It is mandatory to define a passphrase to enable encryption.
 - i** **IMPORTANT:** It is mandatory to define a passphrase to enable encryption. If the passphrase is compromised or lost, the administrator should change it immediately so that the content encryption keys do not become vulnerable. If this passphrase is lost or forgotten, data in the cloud will be unrecoverable.
 - **Confirm Passphrase** — re-enter the passphrase used above.
8. Click **Configure**. A Cloud Storage Group will be created.
9. To enable replication to the cloud, you must link a local container to the cloud using the procedures in [Adding a cloud tiering policy](#).

Adding a Google S3 cloud tier

i **NOTE:** For Google S3 cloud tier to work correctly, first create a default project in the Google cloud platform for interoperable access. For more information, see [Migrating from Amazon S3 to Cloud Storage | Google Cloud](#).

To add a cloud storage group, complete the following steps:

1. In the navigation menu, click **Cloud Storage** to expand the menu, then click **Cloud Tier**.
2. In the Cloud pane, click **Configure** to add a cloud tier.
3. In the **Cloud Provider** drop-down, select **Google S3**.
4. Provide a container name. This is the existing name of your container in your cloud platform.

5. Enter your **Connection String** using one of the two methods below:

- **Default** - this option will compile your connection string into the correct format using the inputs below.
 - **Access key** - The access key is typically 20 upper-case English characters
 - **Secret key** - The secret key is generated automatically by Google. It is typically 40 characters, including mixed upper and lower-case and special symbols.
 - **Region** - The region specifies the Google-specific region in which you want to deploy your backup solution. You can obtain your region code from <https://cloud.google.com/compute/docs/regions-zones>.
 - **Endpoint** - If you are using VPC endpoints, enter the correct endpoint information.
- **Custom** - this option allows you to enter your connection string with additional parameters.
 - Your connection string uses the following syntax:
"accesskey=<ABDCEWERS>;secretkey=<SECRETKEY>;loglevel=warn;endpoint=storage.googleapis.com;region=<google-S3-region>;"
 - You can obtain your region code from <https://cloud.google.com/compute/docs/regions-zones>

Please note the following:

- a. The **access key** is typically 20 upper-case English characters
- b. The **secret key** is generated automatically by Google. It is typically 40 characters, including mixed upper and lower-case and special symbols.
- c. The **region** specifies the Google-specific region in which you want to deploy your backup solution. You can obtain your region code from <https://cloud.google.com/compute/docs/regions-zones>.

An example of a connection string with this syntax follows. Logically, each connection string is unique.

```
accesskey=AKIARERFUCFODHFJUCWK;secretkey=p+8/T+o5WeZkX11QbuPazHX1Id  
WbwgFplxuVl08J;loglevel=warn;region=eu-central-1;
```

6. To apply encryption, select **Encryption** and enter the following:

- **Passphrase** — the passphrase is user-defined and is used to generate a passphrase key that encrypts the file in which the content encryption keys are kept. The passphrase is a human readable key, which can be up to 255 bytes in length. It is mandatory to define a passphrase to enable encryption.

i | **IMPORTANT:** It is mandatory to define a passphrase to enable encryption. If the passphrase is compromised or lost, the administrator should change it immediately so that the content encryption keys do not become vulnerable. If this passphrase is lost or forgotten, data in the cloud will be unrecoverable.

- **Confirm Passphrase** — re-enter the passphrase used above.

7. Click **Configure**. A Cloud Storage Group will be created.

8. To enable replication to the cloud, you must link a local container to the cloud using the procedures in [Adding a cloud tiering policy](#).

i | **IMPORTANT:** Enabling bucket-level object lock retention is not recommended.



NOTE: Objects that were deleted in QoreStor versions earlier than 7.5.0 need to be removed via a lifecycle policy, whereas deletions conducted after updating to 7.5.0 version will function as intended.

Adding a Scality-Artasca-S3 cloud tier

To add a *Scality-Artasca-S3* cloud storage group, complete the following steps:

1. In the navigation menu, click **Cloud Storage** to expand the menu, then click **Cloud Tier**.
2. Click **Configure**.
3. In the **Cloud Provider** drop-down, select **Scality-Artasca-S3**.
4. Provide a container name. This is the name of your existing Azure container.
5. Enter your **Connection String** using one of the two methods below:
 - Default** - this option will compile your connection string into the correct format using the inputs below.
 - **Access key** - The access key is typically 20 upper-case English characters.
 - **Secret key** - These secret keys are generated automatically by the cloud provider. It is typically 40 characters, including mixed upper and lower-case and special symbols.
 - **Region** - The region specifies the region in which you want to deploy your backup solution. To obtain the region code refer to vendor documentation.
 - **Endpoint** - If you are using VPC endpoints, enter the correct endpoint information.
 - Custom** - this option allows you to enter your connection string with additional parameters.
 - Your connection string uses the following syntax:

```
"accesskey=<ABDCEWERS>;secretkey=< >; loglevel=warn; region=<cloud-provider-region>;endpoint=<S3 cloud storage endpoint url>"
```

Please note the following:

- a. The **access key** is typically 20 upper-case English characters
- b. The **secret key** is generated automatically by the cloud provider. It is typically 40 characters, including mixed upper and lower-case and special symbols.
- c. The **region** specifies the region in which you want to deploy your backup solution. Refer to vendor documentation for more details.

An example of a connection string with syntax is as follows. Logically, each connection string is unique.

```
accesskey=AKIARERFUCFODHFJUCWK;secretkey=p+8/T+o5WeZkX11QbuPazHX1Id  
WbwgFplxuVl08J;loglevel=warn;region=us-east-1;
```

6. **Enable Cloud Locking** : Select this option to achieve immutability of data tiered to the cloud.



NOTE: Cloud Locking works with RDA and Object containers only. Data that is set as immutable on the local storage is also set as immutable when tiered to the cloud.

Enabling locking at the Cloud Tier is required so the containers that are tiering data to this Cloud Tier can choose to enable locking as part of their cloud tiering policy.

7. To apply encryption, in the **Cloud Tier Encryption** section enter the following:
 - **Passphrase** — the passphrase is user-defined and is used to generate a passphrase key that encrypts the file in which the content encryption keys are kept. The passphrase is a human readable key, which can be up to 255 bytes in length. It is mandatory to define a passphrase to enable encryption.
 - **IMPORTANT:** It is mandatory to define a passphrase to enable encryption. If the passphrase is compromised or lost, the administrator should change it immediately so that the content encryption keys do not become vulnerable. If this passphrase is lost or forgotten, data in the cloud will be unrecoverable.
 - **Confirm Passphrase** — re-enter the passphrase used above.
8. Click **Configure**. A Cloud Storage Group will be created.
9. To enable replication to the cloud, you must link a local container to the cloud using the procedures in [Adding a cloud tiering policy](#).

Adding an S3 Compatible cloud tier

To add a cloud storage group, complete the following steps

1. In the navigation menu, click **Cloud Storage** to expand the menu, then click **Cloud Tier**.
2. In the Cloud pane, click **Configure** to add a cloud tier.
3. In the **Cloud Provider** drop-down, select **S3 Compatible**.
4. Provide a container name. This is the existing name of your container in your cloud platform.

5. Enter your **Connection String** using one of the two methods below:

- **Default** - this option will compile your connection string into the correct format using the inputs below.
 - **Access key** - The access key is typically 20 upper-case English characters
 - **Secret key** - The secret key is generated automatically by the cloud provider. It is typically 40 characters, including mixed upper and lower-case and special symbols.
 - **Region** - The region specifies the provider-specific region in which you want to deploy your backup solution. To obtain your region code, see the documentation for your vendor.
 - **Endpoint** - If you are using VPC endpoints, enter the correct endpoint information.
- **Custom** - this option allows you to enter your connection string with additional parameters.
 - Your connection string uses the following syntax:
"accesskey=<ABDCEWERS>;secretkey=< >; loglevel=warn; region=<cloud-provider-region>;endpoint=<S3 cloud storage endpoint url>"
Please note the following:
 - a. The **access key** is typically 20 upper-case English characters
 - b. The **secret key** is generated automatically by the cloud provider. It is typically 40 characters, including mixed upper and lower-case and special symbols.
 - c. The **region** specifies the provider-specific region in which you want to deploy your backup solution. To obtain your region code, see the documentation for your vendor.

An example of a connection string with this syntax follows. Logically, each connection string is unique.

```
accesskey=AKIARERFUCFODHFJUCWK;secretkey=p+8/T+o5WeZkX11QbuPazHX1Id  
WbwgFp1xuV108J;loglevel=warn;region=eu-central-1;
```

6. **Enable Cloud Locking**: Select this option to achieve immutability of data tiered to the cloud.

i NOTE: Cloud Locking works with RDA and Object containers only. Data that is set as immutable on the local storage is also set as immutable when tiered to the cloud.

Enabling locking at the Cloud Tier is required so the containers that are tiering data to this Cloud Tier can choose to enable locking as part of their cloud tiering policy.

7. To apply encryption, select **Encryption** and enter the following:

- **Passphrase** — the passphrase is user-defined and is used to generate a passphrase key that encrypts the file in which the content encryption keys are kept. The passphrase is a human readable key, which can be up to 255 bytes in length. It is mandatory to define a passphrase to enable encryption.

i IMPORTANT: It is mandatory to define a passphrase to enable encryption. If the passphrase is compromised or lost, the administrator should change it immediately so that the content encryption keys do not become vulnerable. If this passphrase is lost or forgotten, data in the cloud will be unrecoverable.

- **Confirm Passphrase** — re-enter the passphrase used above.

8. Click **Configure**. A Cloud Storage Group will be created.

9. To enable replication to the cloud, you must link a local container to the cloud using the procedures in [Adding a cloud tiering policy](#).

Adding a Backblaze S3 cloud tier

To add a cloud storage group, complete the following steps

1. In the navigation menu, click **Cloud Storage** to expand the menu, then click **Cloud Tier**.
2. In the Cloud pane, click **Configure** to add a cloud tier.
3. In the **Cloud Provider** drop-down, select **Backblaze S3**.
4. Provide the **S3 Bucket** name in your cloud platform to which cloud replicated data to be written.
5. Enter your **Connection String** using one of the two methods below:
 - **Default** - This option will compile your connection string into the correct format using the inputs below.
 - **Access key** - Enter the "Application Key ID" created in the Backblaze Console. The access key is typically 20 upper-case English characters.

i **NOTE:** Please refer to the Backblaze Cloud provider documentation for creating Application Key ID and Application Key for your account at <https://help.backblaze.com/hc/en-us/articles/360047425453-Getting-Started-with-the-S3-Compatible-API>.

- **Secret key** - Enter the 'Application Key' which is Backblaze Cloud provider equivalent of the Secret Key. It is typically 40 characters, including mixed upper and lower-case and special symbols.
- **Endpoint** -. Provide the S3 endpoint URL for the region for which the S3 bucket is to be created.

- **Custom** - This option allows you to enter your connection string with additional parameters.
 - The BackBlaze connection string uses the following syntax:

```
"accesskey=<ABDCEWERS>;secretkey=< >; loglevel=warn;endpoint=https://s3.<region>.backblazeb2.com;"
```

Please note the following:

- a. The **access key**: Enter the "Application Key ID" created in the Backblaze Console. The access key is typically 20 upper-case English characters
- b. The **secret key**: Enter the 'Application Key' which is Backblaze Cloud provider equivalent of the Secret Key. The secret key is generated automatically by the cloud provider. It is typically 40 characters, including mixed upper and lower-case and special symbols.
- c. The **Endpoint** - Provide the S3 endpoint URL for the region for which the S3 bucket is to be created.

An example of a connection string with this syntax follows. Logically, each connection string is unique.

```
accesskey=AKIASAMPLEJUCWK;secretkey=p+SAMPLE1QbuPazHX1IdWbwgFplxuV108J;l  
oglevel=warn;endpoint=https://s3.us-west-004.backblazeb2.com
```

6. **Enable Cloud Locking:** Select this option to achieve immutability of data tiered to the cloud.

i **NOTE:** Cloud Locking works with RDA and Object containers only. Data that is set as immutable on the local storage is also set as immutable when tiered to the cloud.

Enabling locking at the Cloud Tier is required so the containers that are tiering data to this Cloud Tier can choose to enable locking as part of their cloud tiering policy.

7. To apply encryption, select **Encryption** and enter the following:
 - **Passphrase** — The passphrase is user-defined and is used to generate a passphrase key that encrypts the file in which the content encryption keys are kept. The passphrase is a human readable key, which can be up to 255 bytes in length.

i **IMPORTANT:** It is mandatory to define a passphrase to enable encryption. If the passphrase is compromised or lost, the administrator should change it immediately so that the content encryption keys do not become vulnerable. If this passphrase is lost or forgotten, data in the cloud will be unrecoverable.
 - **Confirm Passphrase** — Re-enter the passphrase used above.
8. Click **Configure**. A Cloud Storage Group will be created.
9. To enable replication to the cloud, you must link a local container to the cloud using the procedures in [Adding a cloud tiering policy](#).

Adding a cloud tier through the command line

To add a cloud tier, complete the following steps.

1. Access the QoreStor CLI. Refer to [Accessing the CLI commands on a Linux QoreStor server](#) for more information.
2. Add a new cloud tier using the command

```
cloud_tier --add --cloud_container <bucket name> --cloud_provider <AWS-S3|AZURE|Wasabi-S3|Backblaze-S3|Scality-Artasca-S3|Google-S3|IBM-S3|S3-Compatible> --name <cloud tier name>
```

i **NOTE:** You will be prompted to enter your Azure connection string or cloud provider access string after executing the command.

Refer to the *QoreStor Command Line Reference Guide* for more information on the **cloud_tier** command and available options.

Creating a cloud tiering schedule

After you create a cloud tier, you can schedule when you want data to transfer to the tier.

To add a cloud tiering schedule

1. In the navigation menu, expand **Cloud Storage**, and then click **Cloud Tier**. The **Cloud Tier** page displays.
2. Select the required **Cloud Tier** from the dropdown.

3. On the **Cloud Tier** page, to reveal the scheduling options, click **Schedule**.
4. For each day of the week, select a time range by completing the following steps:
 - a. Click the **From** box, select a time to begin the schedule, and then click **Set**.
 - b. Click the **To** box, select a time for the schedule to end, and then click **Set**.
5. When finished, click **Save**.
6. To hide the schedule options, click **Schedule**.

Editing a cloud tiering schedule

After you create a cloud tiering schedule, you can edit the schedule by completing the following steps.

To edit a cloud tiering schedule

1. In the navigation menu, expand **Cloud Storage**, and then click **Cloud Tiers**.
The **Cloud Tier** will be displayed.
2. Select the required **Cloud Tier**, click ellipsis icon (three dots) and go to the details page.
3. On the **Cloud Tier** page, to reveal the scheduling options, click **Schedule**.
4. For each time you want to change, do either of the following steps:
 - To delete a time, click the trash can symbol.
 - To change a time, click the box, select a new time, and then click **Set**.
5. When finished, click **Save**.
6. To hide the schedule options, click **Schedule**.

Deleting a cloud tier

Before deleting a cloud tier, review the details below:

- Only the files for whom the on-premises retention age applies and the file data which resides in the cloud will be removed locally.
- Data in the cloud bucket has to be deleted manually.
- Cloud policy settings on the source containers will be removed and the source container made available for cloud replication to a new cloud or archive tier.
- When the Cloud Lock feature is used, the lock duration should have expired for all backups tiered to the cloud for the Cloud Tier to be deleted.

Deleting a cloud tier from the GUI

To delete a cloud tier, complete the following steps.

1. In the navigation menu, click **Cloud Storage** to expand the menu, then select **Cloud Tier**.
2. Click **Delete**.

3. When prompted to confirm, click **Delete**.
4. In the **Passphrase** field, enter the passphrase used for Cloud Tier encryption. This provides validation that the person deleting the cloud tier has the appropriate authorization.
5. Review the containers linked to the cloud tier and confirm that data in these containers can be deleted. Any containers with managed replication configured must be deleted manually before the cloud tier can be deleted.
6. Click **Delete**.

Deleting a cloud tier from the CLI

Access the QoreStor CLI. Refer to [Accessing the CLI commands on a Linux QoreStor server](#) for more information.

1. Delete your cloud tier using the command below. Refer to the *QoreStor Command Line Reference Guide* for more information.

```
cloud_tier --delete
```

2. At the prompt, enter **y** for yes and press **[Enter]**.

Configuring a Cloud Archive Tier

Cloud Archive Tiers can be configured via the UI or via the **cloud_tier** command in the QoreStor CLI.

i **NOTE:** Starting from the 7.4.1 installation, the Archive tier can be added optionally. The following section is applicable only if the Archive tier option is selected.

Before configuring an archive tier, ensure the following requirements are met:

- Your cloud storage must be configured prior to configuring a cloud or archive tier.
- Archive tier is not supported when QoreStor is installed in Cloud Optimized mode or Object Direct Small mode.
- Permissions for your cloud storage must be correctly configured. Refer to [Configuring required permissions to restore from Archive Tier](#) for more information.
- Only RDA and VTL containers can be configured to tier data to Archive Tier.

i **NOTE:** QoreStor's archive tier functionality relies on Amazon S3 Glacier and/or Amazon S3 Glacier Deep Archive storage. Before configuring an archive tier, your cloud archive storage must be properly configured. Please refer to the Amazon S3 documents below for more information:

- [Getting Started with Amazon Simple Storage Service](#)
- [Amazon S3 Console User Guide](#)
- [S3 Batch Operations](#)
- [About AWS Lambda](#)

Configuring required permissions to restore from Archive Tier

For QoreStor to perform batch operations for restoring objects to Amazon S3 storage from Amazon S3 Glacier or Amazon S3 Glacier Deep Archive storage, you must configure an AWS IAM policy with the required permissions and then attach the policy to your AWS account used to access the for accessing AWS S3 storage.

i **NOTE:** When crating an archive tier after upgrading to QoreStor 7.1, the default mode of restores is Lambda. If you create the archive tier before upgrading to QoreStor 7.1, the upgrade automatically switches the restores from Batch operations to Lambda. To change this option, see [Editing an archive tier restore mode using the command line interface](#).

To configure required permissions to restore from Archive Tier

1. From the AWS console, go to the IAM dashboard.
2. On the IAM dashboard, go to the **Policies** page, and then click **Create Policy**.
3. On the Create policy page, click the JSON tab, and then copy and enter the text from the following JSON document:

i **NOTE:** Enter the "AWS Account ID" and the "S3 Archive Tier Bucket Name" as appropriate. Using "*" as a placeholder for the "S3 Archive Tier Bucket Name" may cause an unimportant warning, which you can ignore.

JSON Create Policy document

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "s3:GetObject",
        "s3:RestoreObject",
        "lambda:InvokeFunction"
      ],
      "Resource": [
        "arn:aws:lambda:*:<AWS Account ID>:function:*",
        "arn:aws:s3:::<S3 Archive Tier Bucket Name | *>/*"
      ]
    },
    {
      "Sid": "VisualEditor1",
      "Effect": "Allow",
      "Action": "s3:PutObject",
      "Resource": "arn:aws:s3:::<S3 Archive Tier Bucket Name | *>/batch/*"
    }
  ]
}
```

4. Note the name of the new policy for the next steps. For example, `GlacierTierRolePolicy`.

5. On the IAM dashboard **Roles** page, click **Create Role**.
6. Select a trusted entity, select **Custom trust policy**, and then copy and enter the following JSON document:

JSON Custom Trust Policy document

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "batchoperations.s3.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    },
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "lambda.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

7. Add permissions by searching and selecting the policy you created, and then click **Next**.
8. Give the new role a name and then note the ARN of the IAM Role for next steps. For example, `arn:aws:iam::<AWS Account ID>:role/GlacierTierRole`.
9. Return to the Policies page of the IAM dashboard and click **Create Policy**.

10. Select JSON for permissions, and then replace the JSON text with the following policy document and save it:

i **NOTE:** In the "Resource" portion, for "AWS Account ID" and "IAM Role Name," enter the specific Account ID and ARN of the role. Do not use the lambda function for batch restores. AWS requires you to use "*" in place of the bucket name.

JSON Create Policy document

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "lambda:CreateFunction",
        "iam:GetRole",
        "lambda:InvokeFunction",
        "lambda:GetFunction",
        "lambda:UpdateFunctionConfiguration",
        "s3:RestoreObject",
        "s3:CreateBucket",
        "lambda:GetFunctionConfiguration",
        "s3:ListBucket",
        "lambda:PutFunctionConcurrency",
        "lambda:UpdateFunctionCode",
        "s3:PutObject",
        "s3:GetObject",
        "iam:PassRole",
        "lambda:GetFunctionConcurrency",
```

```

        "lambda:DeleteFunction",
        "lambda:DeleteFunctionConcurrency",
        "s3:DeleteObject",
        "s3:DeleteBucket"
    ],
    "Resource": [
        "arn:aws:iam::<AWS Account ID>:role/<IAM Role Name>",
        "arn:aws:lambda:*:<AWS Account ID>:function:QorestorArchiveRestore",
        "arn:aws:s3::*"
    ]
},
{
    "Sid": "VisualEditor1",
    "Effect": "Allow",
    "Action": "s3:ListAllMyBuckets",
    "Resource": "*"
},
{
    "Sid": "VisualEditor2",
    "Effect": "Allow",
    "Action": [
        "s3:DescribeJob",
        "s3:UpdateJobPriority",
        "s3:UpdateJobStatus"
    ]
},

```

```

        "Resource": "arn:aws:s3:*:<AWS Account ID>:job/*"
    },
    {
        "Sid": "VisualEditor3",
        "Effect": "Allow",
        "Action": [
            "s3:ListJobs",
            "s3:CreateJob"
        ],
        "Resource": "arn:aws:s3:*:<AWS Account ID>:job/*"
    }
]
}

```

The policy creation is complete. Check that the permissions you entered are saved in the policy JSON document.

11. To create an IAM User for the archive tier, go to the **Users** page of the IAM dashboard, click **Add User**, and complete the following steps:
 - a. On the Add user page under Select AWS access type, to generate the `access_key` and `secret_key`, select **Programmatic access**.
 - b. On the Permissions page, select **Attach existing policy directly**, and then select the policy you created in Step 10 to attach to this user.
12. Following the directions in the remaining two tabs to finish creating the user.

i | **NOTE:** Be sure to download the `access_keys` for this user to use when creating an archive tier in QoreStor.

Modifying an Archive Tier after an upgrade

If you created an Archive Tier after an upgrade to QoreStor 7.1 or later release, then the default mode of restores is Lambda. If you created the Archive Tier prior to upgrading to QoreStor 7.1, then the upgrade process automatically switched the default restore mode from Batch operations to Lambda. To revert this change back to the Batch option, complete the following procedure in the CLI.

To modify an Archive Tier after an upgrade

1. In the CLI, use the following commands:

Restore mode change commands

```
cloud_tier --update
[--cloud_password]

[--cloud_archive]

[--archive_retention_in_warm <1 to 365 days>]

[--archive_role_arn <archive role arn>]

[--archive_restore_type <Batch|Lambda>]

[root@jayant-ol82-tst1 ~]
# cloud_tier --update --archive_role_arn
arn:aws:iam::177436582181:role/GlacierTierRole --archive_restore_type Lambda --
cloud_archive
Validating Role-arn string format for group name DefaultCloudArchiveTier ...
Role-arn string format is valid →We do basic format validation for the role ARN
string. We cannot validate permissions at the time of addition/update - AWS does
that during restore operation itself.
Archive Tier updated successfully.
[root@jayant-ol82-tst1 ~]
#
```

```

[root@jayant-ol82-tst1 ~]
# cloud_tier --show --verbose --cloud_archive
Cloud_tier Entry ID : 8
Cloud_tier Name : DefaultCloudArchiveTier
Cloud_tier Compression Type : Fast
Cloud_tier Encryption Set : On
Cloud_tier Encryption Type : Static
Cloud_tier Rotate Period : 0
Cloud_tier Passphrase set : Yes
Cloud_tier Type : Cloud
Cloud_tier Cloud container name : jayantcloud1
Cloud_tier Cloud provider name : AWS-S3
Cloud_tier Cloud archive service name : S3-Glacier
Cloud_tier Archive retention in warm : 2 days
Cloud_tier Archive role ARN string :
arn:aws:iam::177436582181:role/GlacierTierRole
Cloud_tier Archive Restore Type : Lambda Function
Cloud_tier Cloud connection string : loglevel=trace;region=us-east-1;
Cloud_tier Created On : Mon Aug 9 14:29:07 2021 PDT
Cloud_tier Created Bld : 24E2B069
Cloud_tier status : Online
Storage_group Is Storage Agent Group : No
DefaultCloudArchiveTier's Containers
-----
None

[root@jayant-ol82-tst1 ~]
#

```

2.

Adding an archive tier

To add an archive tier

1. In the navigation menu, click **Cloud Storage** to expand the menu, then click **Archive Tier**.
2. In the Archive Tier pane, click **Configure** to add a cloud tier.
3. In the **archive provider** drop-down, select **AWS S3**.
4. Provide the name for your S3 bucket.

5. Enter your **Connection String** using one of the two methods below:

- **Default** - this option will compile your connection string into the correct format using the inputs below.
 - **Access key** - The access key is typically 20 upper-case English characters
 - **Secret key** - The secret key is generated automatically by AWS. It is typically 40 characters, including mixed upper and lower-case and special symbols.
 - **Region** - The region specifies the Amazon-specific region in which you want to deploy your backup solution. Your region name can be obtained from https://docs.aws.amazon.com/general/latest/gr/rande.html#s3_region
- **Custom** - this option allows you to enter your connection string with additional parameters.
- Your connection string uses the following syntax:

```
"accesskey=<ABDCEWERS>;secretkey=< >; loglevel=warn; region=<aws-region>;"
```

Please note the following:

- a. The **access key** is typically 20 upper-case English characters
- b. The **secret key** is generated automatically by AWS. It is typically 40 characters, including mixed upper and lower-case and special symbols.
- c. The **region** specifies the Amazon-specific region in which you want to deploy your backup solution. Your region name can be obtained from https://docs.aws.amazon.com/general/latest/gr/rande.html#s3_region

An example of a connection string with this syntax follows. Logically, each connection string is unique.

```
accesskey=AKIARERFUCFODHFJUCWK;secretkey=p+8/T+o5WeZkX11QbuPazHX1IdWbwgFplxuVl08J;loglevel=warn;region=eu-central-1;
```

6. Enable **Use AWS intelligent Tiering for metadata** to get some cost savings by storing Archive Tier metadata in the AWS S3 Intelligent Tiering storage class.

7. To apply encryption, in the **Archive Tier Encryption** section enter the following:

- **Passphrase** — the passphrase is user-defined and is used to generate a passphrase key that encrypts the file in which the content encryption keys are kept. The passphrase is a human readable key, which can be up to 255 bytes in length. It is mandatory to define a passphrase to enable encryption.

i **IMPORTANT:** It is mandatory to define a passphrase to enable encryption. If the passphrase is compromised or lost, the administrator should change it immediately so that the content encryption keys do not become vulnerable. If this passphrase is lost or forgotten, data in the cloud will be unrecoverable.

- **Confirm Passphrase** — re-enter the passphrase used above.

8. In the **Archive Tier Options** section, enter the following:
 - **Archive Retention in Warm Cloud** - When restore operation succeeds, a temporary copy of the Glacier object is created in standard S3 storage. This setting specifies the number of days this temporary copy is held in S3 before it is deleted. Valid values are any integral values from 1 through 365.
 - **Archive Role ARN** - S3 must have permissions to perform Lambda and batch operations on behalf of the user. An IAM role must be created that has "Create Job", "Pass Role" and other permissions to access the buckets as well as perform the Lambda and batch operations. The account admin is expected to create such roles.
 - **NOTE:** For more information on required permissions and lambda and batch operations, refer to [Configuring required permissions to restore from Archive Tier](#) and the AWS documents [Granting permissions for Amazon S3 Batch Operations](#), [What is AWS Lambda?](#), and [The basics: S3 Batch Operations](#).
 - **Archive Service Name**- Select between **S3-Glacier** or **S3 Deep Archive**.
9. Click **Configure**. A Cloud Storage Group will be created.
10. To enable replication to the cloud, you must link a local container to the cloud using the procedures in [Adding a cloud tiering policy](#).

Editing an archive tier restore mode using the command line interface

If you create an archive tier after an upgrade to 7.1, the default mode of restores is Lambda. If the archive tier was been created prior to a 7.1 upgrade, the upgrade switches the restores from Batch operations to Lambda. You can change this option using the command line interface (CLI) or the user interface (UI).

To edit an archive tier restore mode using the command line interface

1. To change an archive tier that was created before upgrading to QoreStor 7.1. go to the CLI and enter the following commands:

Commands for editing restore mode

```
cloud_tier --update [--cloud_password]
[--cloud_archive]
[--archive_retention_in_warm <1 to 365 days>]
[--archive_role_arn <archive role arn>]
[--archive_restore_type <Batch|Lambda>]

[root@qorestor-ol82-tst1 ~]# cloud_tier --update --archive_role_arn
arn:aws:iam::177436582181:role/IAMLambdaOps_Restrictive --archive_restore_type
Lambda --cloud_archive
Validating Role-arn string format for group name DefaultCloudArchiveTier ...
Role-arn string format is valid →We do basic format validation for the role ARN
string. We cannot validate permissions at the time of addition/update - AWS does
that during restore operation itself.
Archive Tier updated successfully.
[root@jayant-ol82-tst1 ~]#

[root@qorestor-ol82-tst1 ~]# cloud_tier --show --verbose --cloud_archive
Cloud_tier Entry ID : 8
Cloud_tier Name : DefaultCloudArchiveTier
Cloud_tier Compression Type : Fast
Cloud_tier Encryption Set : On
Cloud_tier Encryption Type : Static
Cloud_tier Rotate Period : 0
Cloud_tier Passphrase set : Yes
Cloud_tier Type : Cloud
Cloud_tier Cloud container name : jayantcloud1
Cloud_tier Cloud provider name : AWS-S3
Cloud_tier Cloud archive service name : S3-Glacier
Cloud_tier Archive retention in warm : 2 days
Cloud_tier Archive role ARN string : arn:aws:iam::177436582181:role/IAMLambdaOps_
Restrictive
Cloud_tier Archive Restore Type : Lambda Function
Cloud_tier Cloud connection string : loglevel=trace;region=us-east-1;
Cloud_tier Created On : Mon Aug 9 14:29:07 2021 PDT
Cloud_tier Created Bld : 24E2B069
Cloud_tier status : Online
Storage_group Is Storage Agent Group : No
DefaultCloudArchiveTier's Containers
-----
None
```

2. To change the restore operations of an archive tier while adding the archive tier, go to the CLI and enter the following commands:

Changing Archive Tier restore operations after upgrade

```
cloud_tier --add --cloud_container <bucket name>
--cloud_provider <AWS-S3|AZURE|Backblaze S3|Wasabi-S3|Google-S3|IBM-S3|S3-
Compatible>
[--cloud_archive_service <S3-Glacier|S3-Deep-Archive>]
[--archive_retention_in_warm <1 to 365 days>]
[--archive_role_arn <archive role arn>]
[--archive_restore_type <Batch|Lambda>]

[root@jayant-ol82-tst1 ~]# cloud_tier --add --cloud_container jayantcloud1 --
cloud_provider AWS-S3 --cloud_archive_service S3-Glacier --archive_retention_in_
warm 2 --archive_role_arn arn:aws:iam::177436582181:role/IAMLambdaOps --archive_
restore_type=Lambda
```

Deleting an archive tier

Before deleting an archive tier, review the details below:

- The metadata for the files archived to the cloud will be removed locally. This makes those files unrecoverable.
- Data in the cloud bucket has to be deleted manually.
- Archive policy settings on the source containers are unaffected.

Deleting an archive tier from the GUI

To delete an archive tier, complete the following steps.

1. In the navigation menu, click **Cloud Storage** to expand the menu, then click **Archive Tier**.
2. Click **Delete**.
3. When prompted to confirm, click **Delete**.
4. In the **Passphrase** field, enter the passphrase used for Archive Tier encryption. This provides validation that the person deleting the archive tier has the appropriate authorization.
5. Review the containers linked to the archive tier and confirm that data in these containers can be deleted.
6. Click **Delete**.

Deleting an archive tier from the CLI

Access the QoreStor CLI. Refer to [Accessing the CLI commands on a Linux QoreStor server](#) for more information.

1. Delete your archive tier using the command below. Refer to the *QoreStor Command Line Reference Guide* for more information.

```
cloud_tier --delete --cloud_archive
```

2. At the prompt, enter **y** for yes and press **[Enter]**.

Creating an archive tiering schedule

After you create an archive tier, you can schedule when you want data to transfer to the tier.

To add an archive tiering schedule

1. In the navigation menu, expand **Cloud Storage**, and then click **Archive Tier**. The **Archive Tier** page displays.
2. On the **Archive Tier** page, to reveal the scheduling options, click **Schedule**.
3. For each day of the week, select a time range by completing the following steps:
 - a. Click the **From** box, select a time to begin the schedule, and then click **Set**.
 - b. Click the **To** box, select a time for the schedule to end, and then click **Set**.
4. When finished, click **Save**.
5. To hide the schedule options, click **Schedule**.

Editing an archive tiering schedule

After you create an archive tiering schedule, you can edit the schedule by completing the following steps.

To edit an archive tiering schedule

1. In the navigation menu, expand **Cloud Storage**, and then click **Archive Tier**. The **Archive Tier** page displays.
2. On the **Archive Tier** page, to reveal the scheduling options, click **Schedule**.
3. For each time you want to change, do either of the following steps:
 - To delete a time, click the trash can symbol.
 - To change a time, click the box, select a new time, and then click **Set**.
4. When finished, click **Save**.
5. To hide the schedule options, click **Schedule**.

Restoring from an archive tier

Depending on the container type, data is sent to the archive tier by different methods. For RDA and Object containers, data is archived based on the Archive Tiering Policy. For VTL containers, exporting the cartridge from the backup application will trigger the movement of cart data to cloud.

Restoring data from an archive tier differs in some ways from a standard restore process. There are two possible methods for restoring:

- Selectively restoring backups based on need, and
- Performing a full disaster recovery by completely restoring all data from Glacier to AWS S3.

In both cases, when restoring from an archive tier, no files are saved to on-prem storage. Instead, files are copied from the archive storage (Glacier or Amazon S3 Glacier Deep Archive) to warm AWS S3 storage for a period of time specified by the **Archive Retention in Warm Cloud** setting. When restoring from an archive tier, consider the following:

- Restoring from archive storage is a two-step operation. First, archive data is restored to standard AWS S3 storage, then the objects are read from there.
- There are two options for restoring from Glacier storage: Batch operations and Lambda with batch operations. According to AWS, "Lambda is a compute service that lets you run code without provisioning or managing servers." Using Lambda with batch operations can help avoid certain restore failures. For more information, see [What is AWS Lambda?](#) in the AWS documentation.
- Restored objects will be ready for readback after 4-6 hours for Amazon S3 Glacier (10-12 hours for Amazon S3 Glacier Deep Archive). No notification is issued when restored objects are available. You may view the status of restore operations in the AWS Console. Refer to the Amazon S3 document [Checking Archive Restore Status and Expiration Date](#) for more information. To perform a batch restore for disaster recovery purposes, refer to [Manually restoring datastores from Amazon S3 Glacier](#).
- When restoring objects from archive, you are charged for both the archive copy and the restore copy in warm storage. Use the **Archive Retention in Warm Cloud** value to minimize the duration objects are kept in warm storage.
- For restoring VTL cartridge data, the **command vtl --import** must be run on the QoreStor server.

Restoring files from RDS Container backups replicated to AWS S3 Glacier or Deep Archive

Backups written to an RDA container replicate to an archive tier and stubbed from on-premises storage based on the Archive Tiering Policy.

In the case of an RDA container, the process of restoring files from an archive tier differs based on direct memory access (DMA) . For more information, see the respective DMA setup guide on Quest QoreStor support portal.

Restoring selective tapes of VTL backups replicated to AWS S3 Glacier or Deep Archive

In the case of VTL containers, there are no individual backups replicated to cloud storage, but the entire cartridge is exported to the cloud. You can replicate VTL cartridges that are no longer required as on-premises storage to the cloud using VTL export.

For detailed instructions, refer to the respective DMA setup guide on the Quest Support Portal.

After you replicate and stub VTL tapes to an archive tier, to bring the data to **Warm Cloud** storage and restore it, use the following procedure.

To restore selective tapes of VTL backups replicated to AWS S3 Glacier or Deep Archive

1. To bring data to **Warm Cloud** storage and initiate a restore from AWS S3 Glacier to S3 Standard storage, use the following command:

```
vtl --restore --name <cont_name> --barcode <barcode_of-media>
```

i | **NOTE:** This process typically takes approximately four hours when performed from AWS S3 Glacier storage and as many as eight to 12 hours from Glacier Deep Archive.

2. To check the Glacier to Standard S3 restore status, use the following command:

```
vtl --show --verbose
```

The restore status appears at the end of the command output, as shown in the following example:

- Restore Status: [Cart Name] [Restore current State] [Restored Data available until]
- Restore status can be:
 - [Cart Name] [Restore current State] [Restored Data available until]
GV758Q001 **Restore process initiated** Tue 2020-03-31 09:23:50 EAT
 - [Cart Name] [Restore current State] [Restored Data available until]
GV758Q001 **Restore initiation successful** Tue 2020-03-31 09:23:50 EAT
 - [Cart Name] [Restore current State] [Restored Data available until]
GV758Q001 **Restore Process InProgress** Wed 2020-04-08 03:00:00 EAT
 - [Cart Name] [Restore current State] [Restored Data available until]
GV758Q001 **Restore SUCCESSFUL** Wed 2020-04-08 03:00:00 EAT

3. After the restore is done, import carts using the following command:

```
vtl --import_cart --name <container-name> --barcode <comma-separated-barcodes>
```

You can provide multiple carts during the VTL import operation.

4. To view the status of the selected carts moving from the cloud to the storage_slot of the VTL tape drive, use the following command:

```
vtl --show --verbose
```

5. To rescan the media and update the VTL, perform an Inventory Robot operation.

i | **NOTE:** Before you attempt to restore to a local disk, bring the newly added media online.

Performing a disaster recovery from the cloud

There are two ways to perform a recovery from the cloud, also known as a disaster recovery. You can recover data by creating a new QoreStor instance and transferring the data, or you can perform a quick recovery, which provides a read-only version of RDA container data of your current QoreStor instance. To recover your QoreStor configuration and cloud-replicated data from the cloud, perform the steps below. Before performing these steps, make sure you have the following:

- A functional, properly licensed QoreStor server.
- In the case of Archive Tier, to bring data to S3-Warm Cloud before proceeding with a disaster recovery operation, refer to [Manually restoring datastores from Amazon S3 Glacier](#).
- The connection string for your cloud storage account. This is different depending on your cloud provider. Refer to the appropriate section below for more information.
 - [Adding a Microsoft Azure cloud tier](#)
 - [Adding an Amazon S3 cloud tier](#)
 - [Adding a Wasabi S3 cloud tier](#)
 - [Adding an IBM S3 cloud tier](#)
 - [Adding a Google S3 cloud tier](#)
 - [Adding a Scality-Artesca-S3 cloud tier](#)
 - [Adding an S3 Compatible cloud tier](#)
 - [Adding a Backblaze S3 cloud tier](#)

To perform a disaster recovery from the cloud

1. On a newly installed QoreStor server, run the following recovery command using the definitions provided in the table:

```
maintenance --disaster_recovery --cloud_string <name> --container_name <name>
--cloud_provider_type <name> --passphrase <name> --logfile <name> --quick_ro_
recovery no
```

Table 8: Recovery command definitions

| Command option | Definition |
|------------------------------------|---|
| <code>--cloud_string</code> | Cloud connection string, to connect to the cloud bucket. |
| <code>--container_name</code> | Name of the cloud bucket from where data is to be recovered. Valid values are [a-z, 0-9, '-', '.]. |
| <code>--cloud_provider_type</code> | Name of the cloud service provider, such as <AWS-S3 Azure Wasabi-S3 Google-S3 IBM-S3 Scality-Artesca-S3 S3-Compatible>. |
| <code>--passphrase</code> | Passphrase used on original machine for encrypting the data in the cloud bucket. |
| <code>--logfile</code> | Log file path to capture the ongoing recovery activity. |

This will regenerate configuration data and populate the metadata from the cloud copy.

When completed, you will see the following message:

```
Filesystem disaster recovery started successfully.
```

Please see the `/var/log/oca/qsdr.log` and the logfile given in the command.

2. After the data recovery process is complete, perform a filesystem repair with the command

```
maintenance --filesystem --repair_now
```

When the file system repair is finished, the process is complete.

To perform a quick recovery

1. On a newly installed QoreStor server of the same install mode and configuration as your previous QoreStor server, run the following recovery command using the definitions provided in the table:

```
maintenance --disaster_recovery [--cloud_string <name>] [--container_name <name>] [--cloud_provider_type <name>] [--passphrase <name>] [--logfile <name>] [--quick_ro_recovery <[yes | no]>]
```

Table 9: Quick recovery command definitions

| Command option | Definition |
|------------------------------------|--|
| <code>--cloud_string</code> | Cloud connection string, to connect to the cloud bucket. |
| <code>--container_name</code> | Name of the cloud bucket from where data is to be recovered. Valid values are [a-z, 0-9, '-', '.]. |
| <code>--cloud_provider_type</code> | Name of the cloud service provider, such as <AWS-S3 Azure Wasabi-S3 Google-S3 IBM-S3 Scality-Artasca-S3 S3-Compatible>. |
| <code>--passphrase</code> | Passphrase used on original machine for encrypting the data in the cloud bucket. |
| <code>--logfile</code> | Log file path to capture the ongoing recovery activity. |
| <code>--quick_ro_recovery</code> | Fast disaster recovery, with data in cloud bucket accessible in RO mode only.--quick_ro_recovery quick read-only recovery. |

This will regenerate configuration data and populate the metadata from the cloud copy. When used with a newly deployed QoreStor instance, the `--quick_ro_recovery` command lets you attach a cloud bucket and conduct read-only restores without disturbing the cloud connection of the existing QoreStor.

When completed, you will see the following message:

```
Filesystem disaster recovery started successfully.
```

Please see the `/var/log/oca/qsdr.log` and the logfile given in the command.

Next steps

Depending on your configuration, there may be several steps required after recovering your QoreStor server. Some actions to consider are:

- If you are using QoreStor with NetVault, you will need to add the new QoreStor as target device and add the container.

- Depending on your DMA, you may need to reconfigure DMA or client connections to reference the new QoreStor server.
- Once a disaster recovery completes, the recovered source containers will be unencrypted. Before ingesting new data into the recovered containers, you must enable encryption on the recovered storage groups.

i | **NOTE:** The recovered containers contain only stub files. The data remains encrypted in the cloud tier.

Manually restoring datastores from Amazon S3 Glacier

Performing a disaster recovery from an archive tier, involves restoring all datastores to standard Amazon S3 storage, which includes generating the manifest of objects in the cloud bucket used for archive tier and then initiating a Batch Restore from the console.

i | **NOTE:** Before performing a recovery, confirm the status of various objects and ensure that they are in warm storage for the duration the readback of data is planned.

To manually restore datastores from Amazon S3 Glacier

1. Sign in to the AWS Management Console and open the Amazon S3 Console at <https://console.aws.amazon.com/s3>.
2. Under **Buckets**, select the name of the bucket for which you want to configure Amazon S3 inventory.
3. Click **Management**.
4. Under **Inventory configurations**, click **Create inventory configuration**.
5. For **Inventory configuration name**, enter a name.
6. Set the Inventory scope by entering the following details:
 - For prefix, enter **s3://<cloud_bucket_name>/cds**.
 - For object versions, select **Current versions only**.

7. Select the following options:

Table 10: Manifest options

| Option | Selection |
|-------------------------------|--|
| Report details | Select the location of the AWS account to which you want to save the reports. For example, This account . |
| Destination | <p>Complete one of the following options:</p> <ul style="list-style-type: none"> • Select the destination bucket where you want to save the reports, or • Select a different prefix under the same bucket. For example, s3://<cloud_bucket_name>/manifest. <p>i NOTE: The destination bucket must be in the same AWS Region as the bucket for which you are setting up the inventory.</p> |
| Destination bucket | Here you can see the Destination bucket permissions added to the destination bucket policy which allows Amazon S3 to place data in that bucket. For more information, see Creating a destination bucket policy in the documentation for Amazon S3. |
| Frequency | <p>Select Daily.</p> <p>i NOTE: The first report may take 24 to 48 hours to generate.</p> |
| Output format | Select CSV . |
| Status | Select Enable . |
| Server-side encryption | Select Disable . |
| Additional Fields | Do not select any options, as they are not included in the Batch Job submission. |

The manifest file generates.

8. To initiate the Batch Job for restoring datastores, complete the following steps:
- Navigate to **AWS S3 bucket > Management > Inventory Configurations**.
 - Select the latest manifest file.
 - In the AWS console, to initiate the batch job, click **Create Job from Manifest** option in AWS console.

Managing replications

In the QoreStor GUI, you can set up and manage data replication operations. Such replication operations include, creating new replication relationships, managing or deleting existing replication relationships, starting and stopping replication, and displaying current replication statistics.

Guidelines and prerequisites for replication

Refer to the following important notes and guidelines for understanding and using replication in QoreStor.

- **TCP Port Configuration**—If you plan to perform replication operations across a firewall, the replication service requires that the following fixed TCP ports be configured to support replication operations:
 - port 9904
 - port 9911
 - port 9915
 - port 9916
- **DMA and Domain Relationships** — To allow replication storage information to be viewed by a corresponding data management application (DMA), the target QoreStor system must reside in the same domain as the source QoreStor system in the replication relationship.
- **Replication Limits** — Refer to the *QoreStor Interoperability Guide* for details about the supported system limits for replication. For a definition of connections and streams, see [Streams and connections](#).
- **Version Checking** — The QoreStor software includes version checking that limits replication only between other QoreStor systems that run the same system software version. If versions are incompatible, the administrator will be notified by an event, and replication will not continue.
- **Storage Capacity and Number of Source Systems** — Be aware that the storage capacity of the target QoreStor system is directly affected by the number of source systems writing to its containers, and also by the amount being written by each of these source systems.
- **Bandwidth throttling** — Refer to the `qs_bw_throttle` command in the QoreStor Command Line Reference Guide for information regarding throttling bandwidth consumption for replication between QoreStor systems..
- **MTU Setting** — Primary and secondary replication targets should have the same network maximum transmission unit (MTU) setting.

Adding replication relationships

When configuring replication for object containers, replication must be added from the source system. Also once replication is added for Object container, the replica container is not accessible. Once replication is deleted then the replica becomes accessible over S3 to the clients.

i | **NOTE:** For DR-to-QoreStor replication, configure replication from QoreStor.

i | **NOTE:** Starting from the QoreStor 7.4.1 version, the replication between object containers is not supported.

To add a new replication relationship, complete the following steps.

1. In the navigation menu, click **Replications**.
2. Click **Add Replication**.
3. To define the **Source Container**, select the **Local** or **Remote** option.
 - If you select **Local**, select the local container from the drop-down list.
 - If you select **Remote**, configure the following settings:
 - **Username**—enter the username for the remote system.
 - **Password**—enter the password for the remote system.
 - **Remote Machine**—enter the domain name of the remote system.
 - Click **Retrieve Containers**.
 - **Select Remote Container**—Select the remote container from the drop-down list.
4. For **Encryption**, select one of the following encryption options to encrypt the data as it is replicated: **None**, **AES 128-bit**, or **AES 256-bit**.
5. Under **Target Container**, define the target replica container by configuring the following settings.
 - **Username**—enter the username for the remote system.

i | **NOTE:** The credentials used need to be either the admin or administrator account.
 - **Password**—enter the password for the remote system.
 - **Remote Machine**—enter the domain name of the remote system.
 - Click **Retrieve Remote Containers**.
 - Select the remote container from the drop-down list.
6. Click **Next**.
7. Review the summary and click **Finish**.

i | **NOTE:** For information about starting and stopping replication, see the topic, [Starting and stopping replication](#). For information about scheduling system operations such as replication, see [Understanding system operation scheduling](#).

Configuring replication schedules

By default, replication of any newly written files in the replicated container will occur when QoreStor detects three (3) minutes of idle time. If you wish to constrain replication activities to a specific schedule, you can configure replication schedules on a weekly basis for individual replication-enabled source containers.

i **NOTE:** It is recommended that you do not schedule the running of any Replication operations during the same time period when Cleaner or data ingest operations will be running. If you do not follow this recommendation, the time required to complete the system operations or system performance might be affected.

To configure replication schedules using the UI:

1. On the left navigation menu, click **Replications**.
The Replications page is displayed, listing all configured replication relationships.
2. The configured replications are displayed.
3. To configure or edit the replication schedule, click **Details**.
4. Click **Edit Schedule**.
5. For each day of the week, click the time selector field. Select the **From** and **To** times to configure a window during which replication can run. Click **Set**.
6. Click **Save Schedule**.

To configure replication schedules, complete the following steps.

1. Access the QoreStor command line interface.
2. Use the QoreStor command line interface (CLI) to create and delete the replication schedule. The available commands are:

```
schedule --add --day <Day of the Week> --start_time <HH:MM> --stop_time  
<HH:MM> --name --replication  
schedule --delete --day <Day of the Week> --name --replication
```

For full details on running the cleaner schedule commands, help is available by entering:

```
schedule --help
```

You can view replication details and status on the **Replications** page in the QoreStor GUI by selecting a replication and clicking to expand and view details.

Viewing replication information

In the QoreStor GUI, the Replication page displays current information about replication relationships for data containers in your QoreStor system.

To view replication information

1. In the navigation menu, click **Replications**.
2. The configured replications are displayed.
Using the  or  toggle in the upper-right corner of the Replications pane, you may switch between tile and grid views.
3. To view more detailed information and additional options, click the ellipsis icon  and click **Details**.
 - **Tile view** - the ellipsis icon is to the right of the source container. You may also click the **Details** button to the right of the configured replication.
 - **Grid View** - the ellipsis icon is in the Actions column.

The information below is displayed:

- Replication Source and Target Containers
- Replication Source and Target Systems
- Replication Schedule
- A Summary showing:
 - Cloud Enabled status
 - Cloud Storage Group
 - Peer Status
 - Estimated time to sync
 - Last sync time
 - Schedule status.
- Network Savings
- Replication Throughput
- Transfer
- Network Throughput

 **NOTE:** These statistics refresh every 30 seconds.

Modifying replication relationships

You can modify the following replication settings: encryption and remote container's IP address/host name settings. To modify settings for an existing replication relationship, complete the following steps.

 **CAUTION:** You should exercise caution when configuring the direction of replication for source and target containers. For example, target containers can have their contents deleted if they contain existing data.

 **NOTE:** Because you cannot modify an existing defined role (source or target replica) for a replication relationship, if necessary, you must delete the existing replication relationship, and then recreate a new relationship with the specific source and target roles that you want.

1. In the navigation menu, click **Replications**.
 2. The configured replications are displayed.
Using the  or  toggle in the upper-right corner of the Replications pane, you may switch between tile and grid views.
 3. To edit the replication, click the ellipsis icon  and click **Edit**.
 - **Tile view** - the ellipsis icon is to the right of the source container.
 - **Grid View** - the ellipsis icon is in the Actions column.
 4. Modify the settings/values for the **Source**, or **Target** containers as needed.
 - a. **For Remote System**, modify the IP address/host name and user logon credentials of the source remote system as needed.
 - b. Review the replication details, and then click **Save**.
-  **NOTE:** Replication needs to be stopped before the encryption settings can be modified.
5. Click **Save**.

Deleting replication relationships

To delete an existing replication relationship, complete the following steps:

1. In the navigation menu, click **Replications**.
 2. The configured replications are displayed.
Using the  or  toggle in the upper-right corner of the Replications pane, you may switch between tile and grid views.
To view more detailed information and additional options, click the ellipsis icon  and click **Delete**.
 - **Tile view** - the ellipsis icon is to the right of the source container.
 - **Grid View** - the ellipsis icon is in the Actions column.
 3. In the confirmation dialog box, click **Delete**.
-  **NOTE:** If the deletion fails, you can click **Force Delete** to force removal of the relationship.

Starting and stopping replication

To start or stop replication in an existing replication relationship, complete the following steps.

1. In the navigation menu, click **Replications**.

2. The configured replications are displayed.

Using the  or  toggle in the upper-right corner of the Replications pane, you may switch between tile and grid views.

To view more detailed information and additional options, click the ellipsis icon .

- **Tile view** - the ellipsis icon is to the right of the source container.
- **Grid View** - the ellipsis icon is in the Actions column.

3. To stop the replication process, click the **Stop** icon, and, in the confirmation dialog box, click **Yes** to stop replication.

4. To start the replication process, click the **Start** icon, and, in the confirmation dialog box, click **Yes** to start replication.

 **NOTE:** You can also set up replication schedules as needed. For more information see the topic, [Configuring replication schedules](#).

Managing users

QoreStor gives you the ability to define user roles and assign users to those roles, depending on the authentication type of the user. A local user can have more than one role. There are default user roles for the system as well.

QoreStor has the following types of local user roles: CIFS, OST, RDA, Secure Connect, Object, and Monitor. For the protocol specific user roles, the user is validated with the protocol credentials when the clients connect.

SAML users can have Administrator or Monitor roles. For more information about SAML authentication, see [Configuring SAML](#).

Refer to these important notes about user management in QoreStor.

- Excluding the admin, users can have multiple roles.
- However, OST users are exclusive user roles that can only be assigned to one user at a time.
- The **admin** user is a special default user; it cannot be deleted and no new administrator roles can be created.
- **Monitor** user can not have backup protocol roles like RDA.
- **Object** role is applicable only to Object containers created before QoreStor 7.2.1 release.
- The maximum number of users that can be created for the system is 64.

To view the Users page, on the left navigation menu, click **Users**.

Viewing users

You can view user accounts and settings through either the QoreStor GUI or the command line interface.

Viewing users through the GUI

To view users in the GUI, complete the following steps.

1. In the navigation menu, click **Users**.
2. On the **Users** page you can view the following columns of information about users of your system.
 - **Name**—Displays the name of the user.
 - **Role**—Displays the role(s) assigned to this user.

3. For each user, you have the available actions:
 - **Edit Roles** - Displays the roles currently assigned to user and allows you to change assigned roles.
 - **Change password** - Allows you to change the password for a user account.
 - **Remove** - Removes the user account

Viewing users through the command line

To view users using the command line:

1. Access the QoreStor CLI. Refer to [Accessing the CLI commands on a Linux QoreStor server](#) for more information.
2. View current user accounts with the command:

```
user --show [--users] [--logins] [--verbose] [--name <username>] [--roles  
<cifs|ost|rda|ndmp|iscsi|monitor|administrator|secure_connect|object>
```

Refer to the QoreStor Command Line Reference Guide for more information.

Adding a user

You can easily add users and assign them specific system roles by using either the QoreStor GUI or the command line interface. The system supports up to 64 users.

Adding a user through the GUI

To add a user through the GUI, complete the following steps.

1. In the left navigation menu, click **Users**.
2. On the right side of the page, click **Add User**. The **Add** pane is displayed.

3. Enter the following information.

- **Authentication Type**—Select either **Local**, for user identities managed by QoreStor; or **SAML**, for user identities managed by an external identity provider.
- **Name**—For users with the Local authentication type, enter a user name between 1 and 32 characters. This setting is required.
- **Password**—For users with the Local authentication type, enter a password between 8 and 16 characters. This setting is required.
- **Confirm Password**—Re-enter the password. This setting is required.
- **Roles**—Select from the following options available for your authentication type:
 - **RDA**—For local users, designates the user as an RDA protocol user.
 - **Secure connect**— For local users, designates the user as a Secure Connect user.
 - **CIFS**—For local users, designates the user as a CIFS protocol user
 - **Monitor**—Limits the user to read-only access in the QoreStor GUI.
 - **NDMP**—For local users, designates the user as an NDMP protocol user
 - **Object**—For local users, designates the user as an object protocol user
- **Full Name**—Enter a name for the user.
- **Email Address**—Enter an email address for the user.
- **Phone**—Enter a phone number for the user.
- **Description**—Enter a description for the user.

4. Click **Save**.

Adding a user through the command line

To add a user through the command line, complete the following steps:

1. Access the QoreStor CLI. Refer to [Accessing the CLI commands on a Linux QoreStor server](#) for more information.
2. Add a new user account with the command

```
user --add --name <user name>
```

3. To define roles for the new user, use the command

```
user --update --name <user name> [--add_roles  
<cifs|ost|rda|ndmp|iscsi|monitor|administrator|secure_connect|object>]
```

Refer to the *QoreStor Command Line Reference Guide* for more information.

Modifying local user roles

You can easily change the roles assigned to a local user through both the QoreStor GUI and the command line interface.

Modifying a local user through the GUI

i | **NOTE:** Modifying users through the QoreStor GUI applies only to users with the Local authentication type. It does not apply to SAML users. For more information, see [Configuring SAML](#).

To modify user roles through the GUI, complete the following steps.

1. In the left navigation menu, click **Users**.
2. In the list of users, find the user you want to modify. In the **Actions** column, click the ellipsis icon , then click **Edit**. The **Edit User** pane is displayed.
3. Edit the user's **Roles** as required, select from the following options.

i | **NOTE:** You can select more than one role for a user.

- **OST**—Designates the user as an OST protocol user.
- **RDA**—Designates the user as an RDA protocol user.
- **Secure connect** - Designates the user as a Secure Connect user.
- **Monitor**—Limits the user to read-only access in the QoreStor GUI.

4. Click **Save**.

Modifying a user through the command line

i | **NOTE:** Modifying users through the command line applies only to users with the Local authentication type. It does not apply to SAML users.

To modify a user's roles through the command line, complete the following steps:

1. Access the QoreStor CLI. Refer to [Accessing the CLI commands on a Linux QoreStor server](#) for more information.
2. Modify a user account with the command

```
user --update --name <user name> [--add_roles  
<cifs|ost|rda|ndmp|iscsi|monitor|administrator|secure_connect|object>] [--  
remove_roles <cifs|ost|rda|ndmp|iscsi|monitor|administrator|secure_  
connect|object>]
```

Refer to the *QoreStor Command Line Reference Guide* for more information.

Changing a password for a local user

To change a local user's password for logging in to QoreStor, including the administrator if you have proper permissions, complete the following steps.

i | **NOTE:** This procedure applies only to users with the Local authentication type. SAML users do not have passwords for QoreStor.

To change a password for a local user

1. In the left navigation menu, click **Users**. The Users page is displayed.
2. In the list of users, find the user you want to modify. In the **Actions** column, click the ellipsis icon , then click **Change Password**.
3. In the **Old password** field, type the current password for the user.
4. In the **New password** field, type the new password.
5. In **Confirm password**, retype the new password to confirm.
6. Click **Save**.

Deleting a user

You can delete a user account from both the QoreStor GUI and the command line interface.

 **NOTE:** You cannot delete the administrator user.

Deleting a user account through the GUI

To delete a user, complete the following steps:

1. In the navigation menu, click **Users**.
2. In the list of users, find the user you want to delete. In the **Actions** column, click the ellipsis icon , then click **Delete**.
3. When prompted to confirm, click **Remove**.

Deleting a user account through the command line

To delete a user, complete the following steps:

1. Access the QoreStor CLI. Refer to [Accessing the CLI commands on a Linux QoreStor server](#) for more information.
2. Delete a user account with the following command:

```
user --delete --name <user name>
```

Refer to the QoreStor Command Line Reference Guide for more information.

Monitoring QoreStor

This topic describes how you can monitor the current state of QoreStor operations on the **Dashboard** page.

Using the Dashboard page

The Dashboard page contains graphics that show key information about the current state of your QoreStor instance. This page automatically refreshes every 30 seconds.

To use the Dashboard page, follow these steps.

1. Click **Dashboard** in the navigation menu of the QoreStor GUI.
2. You can view the following graphs:
 - **Compression**—displays a total compression in percentage (combining both deduplication and compression) over a time period (for example, every hour, which is the default).
 - **Swap memory Usage**—displays the amount of swap memory used.
 - **Physical Capacity**—displays total used space, free space, and used and encrypted space in GBs and TBs.
 - **Network throughput**—displays performance information for the configured NICs.
 - **Throughput**—displays the throughput volume (reads and writes) in Mebibytes/second (MiB/s) based on time (for example, every hour, which is the default).
 - **Storage capacity**—displays the amount of used and free storage space.
 - System Information—displays information on the QoreStor system, including:
 - System State
 - Hostname
 - System ID
 - Dictionary Type
 - Number of Storage Groups
 - Number of Containers
 - Operating System

3. At the top of the Dashboard page, you can also view the System Summary section, which lists key information about the current QoreStor system, including:
 - **Cleaner status**—The current cleaner status as one of the following states:
 - **Pending**—displayed when there is any scheduled window set and the current time is outside the scheduled window for the Cleaner operation.
 - **Running**—displayed when the Cleaner operation is running during a scheduled window.
 - **Idle**—displayed only if there is no Cleaner operation running during a scheduled window.
 - **Total number of files in all containers**
 - **Current savings**
 - **Number of containers**
 - **Capacity used**
 - **Number of Storage Groups**
 - **Physical Capacity**
4. To change the time display in the graphs, click **History** within the graph you want to view, and then select the time increment to view.

Viewing QoreStor statistics by using the CLI

An alternate method for viewing the current QoreStor statistics is by using the QoreStor CLI command: `stats --system`. This command shows the following categories of system statistics:

- **Capacity Used** - Total capacity used by QoreStor in all of the mount points and enclosures configured for QoreStor.
- **Capacity Used in GB** - Total capacity in GB used by QoreStor in all of the mount points and enclosures configured for QoreStor.
- **Capacity Free** - Total free capacity in all of the mount points configured for QoreStor. If using QoreStor with Object Direct Storage, it shows Licensed Capacity - Capacity Used.
- **Capacity Free in GB** - Total free capacity in GB in all of the mount points configured for QoreStor. If using QoreStor with Object Direct Storage, it shows Licensed Capacity - Capacity Used.
- **Cloud Capacity Used** - Total capacity of cloud-replicated data (post-dedupe) by this QoreStor instance.
- **Cloud Capacity Used in GB** - Total capacity in GB of cloud-replicated data (post-dedupe) by this QoreStor instance.
- **Object Direct Capacity Used** - Total capacity used by object storage by this QoreStor instance with Object Direct Storage.
- **Object Direct Capacity Used in GB** - Total capacity in GB used by object storage by this QoreStor instance with Object Direct Storage.
- **Archive Capacity Used** - Total capacity used in Azure Glacier or Deep Glacier by this QoreStor instance.
- **Archive Capacity Used in GB** - Total capacity in GB used in Azure Glacier or Deep Glacier by this QoreStor instance.

- **Metadata Used** - Total capacity used by metadata partition, which is not included in licensed capacity.
- **Metadata Used in GB** - Total capacity in GB used by metadata partition, which is not included in licensed capacity.
- **Reserve Space** - Free capacity below which QoreStor goes into read-only mode.
- **Reserve Space in GB** - Free capacity below which QoreStor goes into read-only mode.
- **Total Capacity** - Physical capacity of all of the mount points configured for QoreStor. For QoreStor with Object Direct Storage, it is Licensed capacity.
- **Total Capacity in GB** - Physical capacity in GB of all of the mount points configured for QoreStor. For QoreStor with Object Direct Storage, it is Licensed capacity.
- **Licensed Capacity** - Total licensed capacity.
- **Licensed Capacity in GB** - Total licensed capacity in GB.
- **Read Throughput** - Throughput at which data is being read out of QoreStor during restore or replication.
- **Write Throughput** - Logical throughput at which data is being ingested, which is computed using total data written to QoreStor including duplicate data in a particular period of time.
- **Current Files** - Current number of files present on the system, which does not include files in the Cleaner laundry.
- **Current Files Stubbed** - Current number of files for which a cloud on-premises retention policy was applied and local storage was released. File data resides in cloud storage only.
- **Modified Files Uploaded to Cloud** - Current number of files (count only) that were updated after uploading to a cloud tier and are re-uploaded again to the cloud tier along with updates.
- **Modified Files Uploaded to Archive** - Current number of files (count only) that were updated after uploading to an archive tier and are re-uploaded again to the archive tier along with updates.
- **Current Bytes** - Current size of data protected on the system.
- **Current Stubbed Bytes** - Current logical bytes which are released after applying cloud on-premises retention policy.
- **Current Local Bytes** - Current logical bytes on on-premises disks.
- **Post Dedupe Bytes** - Current bytes in the system post deduplication.
- **Post Compression Bytes** - Current bytes in the system after compression (which is done after deduplication).
- **Post Encryption Bytes** - Current bytes after encryption. Encryption (if enabled) is done after compression.
- **Post Encryption Bytes in GiB** - Current bytes in GiB after encryption. Encryption (if enabled) is done after compression.
- **Cleaner Status** - Status of the space reclamation process.
- **Compression Status** - Status of the compression process.
- **Total Inodes** - current number of file and directory inodes.
- **Bytes decrypted** - Total number of bytes decrypted during data restore.
- **Dedupe Savings** - Current dedupe savings across the system (including cloud and archive tier) if configured any.

- **Compression Savings** - Current compression savings across system (including cloud and archive tier), if configured.
- **Total Savings** - Current dedupe and compression savings across system (including cloud and archive tier), if configured.
- **Local Dedupe Savings** - Current dedupe savings of system for any data residing locally on disk.
- **Local Compression Savings** - Current compression savings of system for any data residing locally on disk.
- **Local Total Savings** - Current dedupe and compression savings of system for any data residing locally on disk.
- **Current Recycle Bin Files** - Current number of files in Recycle Bin (for all the containers on which Recycle Bin was enabled).
- **Current Recycle Bin Logical Bytes** - Current number of logical bytes in Recycle Bin (for all the containers on which Recycle Bin was enabled).
- **Total Immutable Files** - Current number of files for which immutability is set across all the RDA containers.

For more information on QoreStor CLI commands, see the *QoreStor Command Line Reference Guide*.

Monitoring system alerts

You can easily view current system alerts and events in the QoreStor GUI.

- To view system alerts, click the **Alerts** icon  in the header pane. The Alerts drawer displays the active alerts.

i | **NOTE:** For a detailed list of possibly occurring alerts, see the topic, “QoreStor system alert and event messages,” in the “Support, maintenance, and troubleshooting” chapter of this guide.

Monitoring clients

You can easily view the current clients that are connected to the QoreStor system. Client information can be viewed through the System Configuration page in the QoreStor GUI, or using the QoreStor CLI. .

Monitoring clients through the QoreStor GUI

1. To view client information for QoreStor, on the navigation menu, click **System Configuration**. The total number of currently active clients for a particular type is displayed at the bottom of the page.
2. Client information is grouped by client type (RDA/OST, NFS/CIFS, NDMP, or iSCSI). To view client information, click the appropriate pane to view information for that client type.

Monitoring clients through the QoreStor CLI

You can monitor client status through the CLI by using either the **rda** or **ost** commands.

To monitor clients, complete the following steps:

1. Access the QoreStor CLI. Refer to [Accessing the CLI commands on a Linux QoreStor server](#) for more information.
2. View client information with the appropriate client command

```
stats -clients[--type <NFS|CIFS|OST|RDS|NDMP|ISCSI>]
```

Refer to the *QoreStor Command Line Reference Guide* for more information.

Monitoring system events

You can easily view and filter current system events in the QoreStor GUI. To monitor QoreStor events, follow these steps:

1. In the navigation menu, click **Events**.
The Events page displays a summary table of events listed by index number, severity (Informational, Warning, or Critical), timestamp of the event, and brief message describing the event.
2. To filter the list of events, do any of the following:
 - In the **Date Range** field, click to select a range of dates.
 - In the **Message** field, enter a word or string of words for which to search in Event messages. (The system is not case-sensitive)
 - In the Filter Events field, select the severity level for which to search as **Info**, **Warning**, or **Critical**.
 - Click **Filter**. Search results appear in the Events summary table.

Getting daily usage statistics from QoreStor

Daily reports can be configured to be delivered over email from the QoreStor server to see the usage at the system, storage group, and container levels. Replication and cloud tiering statistics are also included. The current and delta of statistics from the previous are part of the report.

To configure daily email reports, access QoreStor CL; execute the following command.

```
email_stats --configure
```

Refer to the *QoreStor Command Line Reference Guide* for more information.

Managing QoreStor Remotely

QoreStor leverages the Quest QorePortal to provide a detailed dashboard of all of the QoreStor and DR Series systems in your organization. With this dashboard, you can easily monitor and manage all of the QoreStor and DR Series systems in your enterprise through one view.

To view QoreStor and DR Series systems on QorePortal, you must register the system with the portal and enable remote management. See the topics below for more information:

- [Registering QoreStor with Quest QorePortal](#)

Getting started with QorePortal

Before registering your QoreStor system(s) with the QorePortal, you must have at least one user account and one organization. Follow the steps below to complete the initial configuration:

1. Access the QorePortal at <https://qoreportal.quest.com>.
2. If you have a Quest Support login, enter your user name and password. Otherwise, click **Sign up for a new account** and complete the dialog.
3. On QorePortal, under **Organizations**, click **Add**.
4. Enter a name for the Organization.
5. Click **Add**.

Registering QoreStor with Quest QorePortal

In order to register your QoreStor or DR Series system with QorePortal, you must have an account on the portal and at least one Organization configured.

To register QoreStor

1. On QorePortal, select the organization to which you want to add a QoreStor system.
2. Click **Register QoreStor**.
3. Click **Generate registration token**.
4. In the Register QoreStor dialog, click the copy icon to copy the token to your clipboard or other wise record the token.
5. On your QoreStor system navigation menu, click **Management**.
6. In the **Registration** section, read the registration rules and select **I have read and accept the rules**.
7. Click **Register**.
8. Paste or enter the token generated above. Click **Register**.

Enabling Remote Management

Once a QoreStor system is registered with the Global View Portal, you have the option of enabling remote management, which allows management and configuration of a QoreStor system through the Global View Portal.

To enable remote management

1. On your QoreStor system navigation menu, click **Management**.
2. In the **Remote Management** section, click **Enable**.

Viewing and using QorePortal

QorePortal displays a convenient view of the operating statistics for all of the QoreStor and DR Series systems that you have added. On this page, you can monitor the status of and easily navigate to the QoreStor and DR Series systems that you have added to QorePortal. Using the portal makes it easy to navigate to a different system in your enterprise without having to log out and log on by using new browser sessions.

To view and use QorePortal

1. To view **QorePortal**, enter the URL (<https://qoreportal.quest.com>) in a supported browser and log in.
2. In the **Organizations** list, click the desired organization. The QorePortal **Systems** page is displayed, showing a summary and a list of assets that have been added to QorePortal.

This list includes all of the systems in QorePortal and provides a high-level status. By default, assets are listed alphabetically by host name. You can sort the list by a clicking the column header, which toggles between ascending and descending order. This sort order is retained if you leave the page and return later. The following table describes the information displayed in the asset list.

| Column | Description |
|---------------|---|
| Host name | Lists the host name of the system |
| Service tag | For QoreStor systems, this column lists the System ID. For DR Series systems, it lists the Service Tag. |
| Product model | Lists the model of the asset. |
| System state | Lists the current state of the system (such as Operational Mode or Manual Intervention Required) |
| Alerts | Displays the alert count. You can click the number to navigate to the Alerts page. |
| Diags | Displays the number of generated diagnostic bundles available. |

3. To view more detailed information on a specific system, click **Details**. This action provides the following views:

| View | Description |
|----------------|---|
| Dashboard | Provides a condensed view of the QoreStor dashboard containing charts for Capacities , Storage savings , Throughput , and System usage . The charts can be configured to show data for the previous three hours, the previous day, week or month, or for a custom date range. |
| Storage Groups | Lists the configured storage groups and containers and key data points for each. |
| Alerts | Lists generated alerts |
| Diagnostics | Provides details about generated diagnostic bundles as well as a direct download link. |

- To manage a QoreStor system, find the system in QorePortal **Systems** page. Click **Manage**. A cloud-enabled version of the QoreStor UI will be displayed, providing management and configuration options as documented in this *User Guide*.

i | **NOTE:** Remote Management must be enabled on each QoreStor system before you can manage the system through QorePortal. Refer to [Enabling Remote Management](#) for information.

Support, maintenance, and troubleshooting

The QoreStor GUI provides various information and tools that can help you better understand the current state of your system and that provide basic, support, maintenance, and troubleshooting functionality.

Using QoreStor Diagnostics

In the QoreStor GUI, the **Diagnostics** page provides the ability to generate and view diagnostics bundles used by Quest Support to troubleshoot QoreStor problems.

Viewing system diagnostic log files

A QoreStor system diagnostics log file is a bundle that contains a variety of file types that record the latest system settings and saves them in a compressed .lzip file format.

In the QoreStor system GUI, the **Diagnostics** page allows you to generate diagnostic logs that capture the state of your system. You can also download these log files or delete them as needed.

To view the system diagnostics page, follow these steps.

1. In the navigation menu, click **Diagnostics**.
2. You can view the following columns of information on the **Diagnostics** page for the diagnostics logs that have been generated.

- **File name**—in this format, `<hostname>_<date>_<time>.lzip`, as in this example: **acme-sys-19_2012-10-12_13-51-40.lzip**

 **NOTE:** Diagnostic log file names are limited to 128 characters.

- **Size**—in Megabytes.
- **Path**—The location to which the diagnostics bundle is saved.
- **Download**—Use this icon to download the diagnostics bundle.

Understanding diagnostics collection

The Diagnostics function in the QoreStor system lets you collect and manage your system's diagnostic log file bundles. The Diagnostics function works by collecting all the system-related information that could help when diagnosing a problem or error condition in the system. Each diagnostic log file bundle provides:

- A current snapshot of system operations
- System-related information that assists in understanding system operations
- A record of system operations in case Technical Support needs to provide technical assistance

Diagnostics bundles are generated when a QoreStor CLI or GUI request is made by the administrator (and the default reason that is listed is admin-generated).

When the diagnostics log directory exceeds the maximum storage capacity, any log older than one hour is automatically deleted. QoreStor GUI lets you download and save diagnostics log files to other systems on your network. QoreStor also maintains a separate archive logs directory that collects other system-related information, and these archive logs are also automatically deleted when they exceed a maximum capacity. When you generate a diagnostics log file bundle, it contains all of the QoreStor information that you need when contacting Technical Support for technical assistance. . When a diagnostics log file bundle is generated, this process also collects all the previous auto-generated diagnostics and deletes them from the system.

The QoreStor GUI provides options to display existing diagnostics logs, generate new diagnostics logs, download and save copies of existing diagnostics logs, or delete existing diagnostics logs. The QoreStor CLI also provides the means for managing, generating, or downloading the diagnostics log files. For more information, see the *QoreStor Command Line Reference Guide*.

Generating a diagnostics log file

A QoreStor diagnostics log file is a bundle that contains a variety of file types that record the latest system settings, and saves them in a compressed .zip file format. When you generate a diagnostics log file bundle, it contains all of the QoreStor information that may be needed when contacting Technical Support for technical assistance. This also includes all the previous auto-generated diagnostics log files, which are then deleted from the QoreStor system.

Generating a diagnostics file through the GUI

To generate a diagnostics log file bundle for your system, complete the following steps:

1. In the left navigation menu, click **Diagnostics**.
2. On the right side of the page, click **Generate Diagnostics**.

Once completed, the new diagnostics log file resides at the top of the **File Name** column in the table. To verify, check its timestamp (using its date and time), to ensure this is the latest diagnostics file created.

Generating a diagnostics log file through the command line

To generate a diagnostics log file bundle for your system, complete the following steps:

1. Access your system's command line interface. QoreStor commands are located in the `/opt/qorestor/bin` directory. Either change directories to the appropriate location, or add the appropriate location to your system path.

2. Generate a diagnostics log file with the command

```
maintenance --diags [--collect [basic]]
```

i | **NOTE:** Running the command with the "basic" suboption collects diagnostics without core files.

Refer to the *QoreStor Command Line Reference Guide* for more information.

Downloading diagnostics log files

To download an existing diagnostics log file, complete the following steps:

1. In the left navigation menu, click **Diagnostics**.
2. In the list, select the diagnostics log file you want to download, and click the **Download** icon.
3. Download and save the file as needed.

Deleting a Diagnostics Log File

To delete an existing diagnostics log file from the Diagnostics summary table on the Diagnostics page, complete the following:

1. Select **Diagnostics**.
2. Click **Select** to select the diagnostics file you want to delete, and click **Delete**.
3. Click **OK** to delete the selected diagnostics log file (or click **Cancel** to display the **Diagnostics** page).

Troubleshooting error conditions

To troubleshoot error conditions that disrupt your normal QoreStor operations, complete the following:

1. Generate a QoreStor diagnostics log file bundle if one has not already been automatically created. For more information, see [Generating a diagnostics log file](#).
2. Check the system alert and system event messages to determine the current status of your QoreStor system.
3. Verify if the QoreStor system has recovered or whether it has entered into Maintenance mode.
4. If you cannot resolve the issue using the information in this QoreStor documentation, contact Quest Technical Support.

Excluding QoreStor directories from antivirus scans

Antivirus software can disable processes or cause files in the QoreStor server and corresponding repositories to be quarantined, causing a QoreStor system to go offline, go into maintenance mode, and initiate a filesystem scan in which there will be data loss. Antivirus software incorrectly identifies files in the datastore as viruses and quarantines or deletes them according to the antiviruses rule set. To avoid this issue, see the complete list of the processes and directories that you should exclude from antivirus scans at [Antivirus exclusions for QoreStor](#).

Security recommendations guide

The following table describes the recommendations Quest offers for specific security scenarios.

Table 11: Security recommendations

| Sr. No. | Asset | Recommendation |
|---------|---|--|
| 1 | Secure connect certificates | Use third-party signing certificates like DigiCert, SSL.com, etc. Refer to the QoreStor User Guide for instructions on using third party certificates. |
| 2 | Object Container Certificate | Use third-party signing certificate. Currently Object Container and QS UI use the same certificate. We recommend using different certificates for each service. |
| 3 | QS UI Certificate | Use third-party signing certificate that can be uploaded via UI Dashboard. Refer to the QoreStor User Guide for instructions on using third party certificates. |
| 4 | QoreStor default passwords | <p>The user should change the passwords immediately after installation. Minimum strength policies must be enforced at the time of changing passwords.</p> <p>Passwords to change:</p> <ul style="list-style-type: none"> • backup_user (default OST user) • UI admin password • CIFS admin password, if enabled <p>In addition, Cloud Tier and Archive Tier need passphrases at the time of creation of the storage groups. These passphrases must be treated like passwords from security and strength standpoint.</p> |
| 5 | Default port settings and firewall settings | <p>Quest recommends disabling the network ports that are not needed for customer use cases.</p> <ul style="list-style-type: none"> • Quest recommends enabling just the following ports: 9443 (secure connect), 22 (SSH) and 5233 (HTTPS) • QoreStor recommended EDM ports: 2500-3500 (Data) and 6160, 6162(Control). |

- Quest recommends disabling the following ports unless the customer is using the specific functionality: 80 (HTTP), 9000-9005 (Object container), 12000-12127 (RDA-NDMP), 9920, 10011, 11000 (OST/RDA without secure connect), 9904, 9911, 9915, 9916 (Replication), 111, 2049 (NFS), 138, 139, 445 (CIFS), 10000, 43000-43040 (NDMP) and 3260 (iSCSI)
- Customers can enable or disable ports using system firewall configuration. Alternatively, customers can use `fw_config`, a script provided by QoreStor, to manage the port settings. Below are some commands to open ports using `fw_config`:

To limit the set of open ports to a minimum set `/opt/qorestor/bin/fw_config -c sc`

{This implicitly includes the UI port and ssh which is enabled by the OS)

To enable ports used for RDCIFS or CIFS `/opt/qorestor/bin/fw_config -c sc,cifs`

To enable ports used for RDNFS or NFS `/opt/qorestor/bin/fw_config -c sc,nfs`

To enable ports used for the object container `/opt/qorestor/bin/fw_config -c sc,objstor`

To enable ports used for replication from a DR Appliance to the QoreStor server `/opt/qorestor/bin/fw_config -c sc,oca`

To enable ports used for iSCSI `/opt/qorestor/bin/fw_config -c sc,iscsi`

To enable ports used for VTL NDMP `/opt/qorestor/bin/fw_config -c sc,ndmp`

To enable ports used for EDM `/opt/qorestor/bin/fw_config -c sc,edm`

i **NOTE:** Ports can be combined if needed. For example, to enable ports for replication from a DR, and RDCIFS, you would use:

```
/opt/qorestor/bin/fw_config -c sc,cifs,oca
```

6

AWS least privileges

As a general rule, enable only the least set of permissions needed to perform operations on cloud objects.

- Bucket policies: Quest recommends setting RW permissions to users within the account and not give permissions to users outside the account.

| | | |
|----|--|---|
| | | <ul style="list-style-type: none"> • IAM Policies: Batch and Lambda operations use IAM policies to manage access and permissions. Please refer to the QoreStor User Guide for sample policies. |
| 7 | Azure and other SPs least privileges and | As a general rule, enable only the least set of permissions needed to perform operations on cloud objects. For storage buckets, Quest recommends setting RW permissions to users within the account and not give permissions to users outside the account |
| 8 | Network Security Group (NSG) port settings for Azure market place images | Please refer to Azure market place deployment guide for recommended NSG settings |
| 9 | UI log-in attempts | Quest recommends monitoring login attempts from UI using events. This will be useful to detect unauthorized login attempts to QoreStor via the UI. Refer to user guide for instructions on event monitoring. |
| 10 | Users logged into QoreStor | Monitor local users logged into the QoreStor server. Super users can check <code>/var/log/secure</code> for shell logins. |
| 11 | Access to external CIFS/NFS shares | Quest recommends restricting access to CIFS/NFS shares based on IP white-listing. Check QoreStor events for mount access to the shares. |
| 12 | Encryption at rest and replication channel encryption | Quest recommends encryption at rest and encryption of in-flight data (replication channel) using internal keys and SHA256 to secure the backup data. Please refer to the user guide for instructions on how to enable them |
| 13 | RDA immutability | QoreStor version 7.1 and later offers enhanced security using RDA Immutability, which is under integration by DMAs. Please refer to user guide for details on the feature and instructions to enable it. |
| 14 | Recycle Bin | QoreStor version 7.1 and later offers protection against ransomware attacks with Recycle Bin. Please refer to user guide for details on the feature and instructions to enable it. |

Quest provides software solutions for the rapidly-changing world of enterprise IT. We help simplify the challenges caused by data explosion, cloud expansion, hybrid datacenters, security threats, and regulatory requirements. We are a global provider to 130,000 companies across 100 countries, including 95% of the Fortune 500 and 90% of the Global 1000. Since 1987, we have built a portfolio of solutions that now includes database management, data protection, identity and access management, Microsoft platform management, and unified endpoint management. With Quest, organizations spend less time on IT administration and more time on business innovation. For more information, visit www.quest.com.

Technical support resources

Technical support is available to Quest customers with a valid maintenance contract and customers who have trial versions. You can access the Quest Support Portal at <https://support.quest.com>.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- Submit and manage a Service Request.
- View Knowledge Base articles.
- Sign up for product notifications.
- Download software and technical documentation.
- View how-to-videos.
- Engage in community discussions.
- Chat with support engineers online.
- View services to assist you with your product.