

EMC® CloudArray®

Version 7.1

Administrator Guide

REVISION 1

Copyright © 2016-2017 Dell Inc. or its subsidiaries. All rights reserved.

Published June 2017

Dell believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS-IS.” DELL MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. USE, COPYING, AND DISTRIBUTION OF ANY DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners.
Published in the USA.

EMC Corporation
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.EMC.com

CONTENTS

Tables		7
	Preface	9
	Revision history.....	12
Chapter 1	Before You Begin	13
	Overview.....	14
	Before using CloudArray.....	14
	Supported web browsers.....	14
	Network port requirements.....	14
	Memory sharing.....	15
Chapter 2	Logging in to CloudArray	17
	Initial login and setup.....	18
	Initial login and setup without a portal connection.....	18
Chapter 3	CloudArray Dashboard	21
	About the CloudArray dashboard.....	22
	CloudArray Dashboard statistics.....	22
Chapter 4	Volumes	25
	CloudArray volumes.....	26
	Viewing volume details.....	26
	Creating and mapping a volume.....	26
	Formatting a volume in Windows.....	27
	Unmapping a volume.....	30
	Deleting a volume.....	31
	Expanding volume and share capacity.....	31
	Migrating a volume from one cache policy to another.....	32
	Migrating a volume from one cloud provider to another.....	33
Chapter 5	File Shares	35
	Enabling CIFS and NFS file sharing.....	36
	Creating a CIFS share.....	36
	CIFS security.....	37
	Setting CIFS security with Microsoft Active Directory.....	37
	Setting CIFS security with CloudArray management.....	38
	NFS security.....	43
	Setting NFS security with AUTH_SYS.....	44
	Setting NFS security with Kerberos V5.....	45
Chapter 6	Deduplication	47
	Deduplication.....	48
	System requirements for deduplication.....	48
	Deduplication configuration.....	48

Chapter 7	Cloud Providers	49
	Cloud providers.....	50
	Configuring remote cloud providers.....	50
	Viewing cloud provider details.....	50
	Configuring NFS as a local cloud provider.....	51
	Deleting cloud providers.....	51
Chapter 8	Caches	53
	Cache overview.....	54
	Viewing cache details.....	54
	Adding cache sources.....	54
	Viewing cache source details.....	55
	Allocating cache.....	56
	Deduplication cache sizing.....	57
	Expanding existing cache.....	58
	Deleting caches.....	59
Chapter 9	Host iSCSI	61
	iSCSI clients.....	62
	Viewing iSCSI client details.....	62
	Creating an iSCSI client.....	62
	Deleting iSCSI clients.....	62
	Configuring iSCSI for VMware vSphere Hypervisor (ESXi).....	63
	Configuring iSCSI for Windows.....	67
	Configuring iSCSI for Linux.....	69
	Configuring iSCSI for SUSE Linux.....	69
	Configuring iSCSI for HP-UX.....	70
Chapter 10	Snapshots	71
	Snapshots.....	72
	Creating a snapshot.....	72
	Scheduling a snapshot.....	72
	Enabling and disabling snapshot scheduling.....	74
	Exposing a snapshot.....	75
	Unexposing a snapshot.....	75
	Mapping a snapshot to a host.....	76
	Deleting a snapshot.....	77
Chapter 11	Network Tools	79
	Throttling network bandwidth.....	80
	Disabling the Bandwidth Throttler.....	82
	Configuring Cloud Performance Optimizer.....	82
	Configuring the default gateway.....	83
	Configuring DNS servers.....	83
	Configuring the hostname.....	83
	Modifying network adapter settings.....	84
	Configuring NTP servers.....	84
	Verifying network settings.....	85
	Enabling a network proxy.....	85
Chapter 12	Administration	87
	Backing up CloudArray.....	88

Changing your CloudArray password.....	88
Updating CloudArray software.....	88
Updating CloudArray without a portal connection.....	89
CloudArray portal settings.....	89
Disabling the CloudArray portal.....	90
Enabling the CloudArray portal.....	90
Disaster recovery testing.....	91
Configuring the Primary CloudArray.....	91
Configuring the DR Test CloudArray.....	93
Provisioning policies.....	94
Viewing provisioning policy details.....	95
Restoring CloudArray configurations.....	95
Selecting the configuration to restore.....	95
Restoring a CloudArray configuration.....	96
Settings.....	97
Managing SSL certificates.....	97
Collecting support data.....	98
Changing the time zone.....	99
User Management.....	99
Adding a new user account.....	99
Managing settings for a user account.....	100
Installing a custom SSL certificate.....	101
Updating the CloudArray license.....	102
Utilities.....	102
Executing CLI commands.....	102
Executing utilities.....	107
 Chapter 13	
Alerting	109
Registering with EMC Secure Remote Services	110
Configuring CloudArray portal alerts.....	110
Configuring SNMP traps.....	110
 Chapter 14	
Reporting	113
Hardware details.....	114

TABLES

1	Typographical conventions used in this content.....	10
2	Revision history.....	12

Preface

As part of an effort to improve its product lines, EMC periodically releases revisions of its software and hardware. Therefore, some functions described in this document might not be supported by all versions of the software or hardware currently in use. The product release notes provide the most up-to-date information on product features.

Contact your EMC technical support professional if a product does not function properly or does not function as described in this document.

Note

This document was accurate at publication time. Go to **Downloads & Documents** in the EMC CloudArray portal (www.cloudarray.com) to ensure that you are using the latest version of this document.

Purpose

This document is an illustrated tutorial on CloudArray operations. It provides CloudArray administrators with a fundamental understanding of the product.

Related documentation

The following documents provide additional information about CloudArray. Online help is available from the CloudArray graphical user interface. All other documents are available on the EMC CloudArray portal (<https://www.cloudarray.com>).

EMC CloudArray Online Help

Describes how to how to perform tasks with the CloudArray user interface.

EMC CloudArray Release Notes

Describes new features and known and fixed issues in the release.

EMC CloudArray Physical Appliance Installation Guide

Describes how to install and configure the CloudArray physical appliance.

EMC CloudArray Virtual Edition Installation Guide

Describes how to install and configure the CloudArray virtual machine.

EMC CloudArray Best Practices

Provides best practices for CloudArray implementation and usage.

EMC CloudArray Customer Maintenance Guide

Describes how to replace components in a CloudArray Physical Appliance.

Special notice conventions used in this document

EMC uses the following conventions for special notices:



Indicates a hazardous situation which, if not avoided, will result in death or serious injury.

⚠ WARNING

Indicates a hazardous situation which, if not avoided, could result in death or serious injury.

⚠ CAUTION

Indicates a hazardous situation which, if not avoided, could result in minor or moderate injury.

NOTICE

Addresses practices not related to personal injury.

Note

Presents information that is important, but not hazard-related.

Typographical conventions

EMC uses the following type style conventions in this document:

Table 1 Typographical conventions used in this content

Bold	Used for names of interface elements, such as names of windows, dialog boxes, buttons, fields, tab names, key names, and menu paths (what the user specifically selects or clicks)
<i>Italic</i>	Used for full titles of publications referenced in text
Monospace	Used for: <ul style="list-style-type: none"> • System code • System output, such as an error message or script • Pathnames, filenames, prompts, and syntax • Commands and options
<i>Monospace italic</i>	Used for variables
Monospace bold	Used for user input
[]	Square brackets enclose optional values
	Vertical bar indicates alternate selections - the bar means “or”
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate nonessential information omitted from the example

Where to get help

EMC support, product, and licensing information can be obtained as follows:

Product information

For documentation, release notes, software updates, or information about EMC products, go to EMC Online Support at <https://support.emc.com> or the CloudArray portal at <https://www.cloudarray.com>.

Technical support

For technical support, go to EMC Online Support <https://support.emc.com> and click **Service Center**. To open a service request, you must have a valid support agreement. Please contact your EMC sales representative for details about obtaining a valid support agreement or with questions about your account.

Your comments

Your suggestions help us improve the accuracy, organization, and overall quality of the documentation. Send your comments and feedback to:
VMAXContentFeedback@emc.com

Revision history

The following table presents the revision history of this document.

Table 2 Revision history

Revision	Description and/or change	Date released
1.1	Minor corrections and updates	Oct 14, 2016
1.0	Initial release	Mar 31, 2016

CHAPTER 1

Before You Begin

Topics include:

- [Overview](#) 14
- [Before using CloudArray](#) 14
- [Supported web browsers](#) 14
- [Network port requirements](#) 14
- [Memory sharing](#) 15

Overview

EMC® CloudArray® provides cloud-integrated storage that extends high-performance storage arrays with cost-effective cloud capacity. By providing access to a private or public cloud storage tier through standard interfaces, CloudArray technology simplifies storage management for inactive data and offsite protection.

Designed to combine the resource efficiency of the cloud with traditional, on-premises storage, CloudArray enables you to scale your storage area networks (SAN) and network-attached storage (NAS) with on-demand cloud capacity. You can easily adjust for future data growth by expanding existing cloud volumes or creating new ones. CloudArray's policy-driven caching capabilities ensure the appropriate level of accessibility and performance based on the data stored. In the background, CloudArray encrypts and compresses the data before sending it to the cloud.

Before using CloudArray

Before you log in to CloudArray:

- Review the *EMC CloudArray Release Notes*.
- For a virtual machine, complete the installation procedures described in the *EMC CloudArray Virtual Edition Installation Guide*.
- For a physical appliance, complete the installation procedures described in the *EMC CloudArray Physical Appliance Installation Guide*.

Supported web browsers

- Google Chrome v49 or higher
- Microsoft Edge v25 or higher
- Internet Explorer v11 or higher
- Mozilla Firefox v45 or higher
- Safari for Mac v6.1.6 or higher

Network port requirements

CloudArray allows network traffic on the following TCP/IP ports:

- 80 (HTTP)
- 111 (NFS Portmapper TCP/UDP)
- 137, 138, 139 (CIFS)
- 443 (HTTPS)
- 445 (CIFS)
- 662 (NFS Status TCP/UDP)
- 875 (NFS rquotad TCP/UDP)
- 892 (NFS mountd TCP/UDP)
- 2049 (NFS nfsd TCP/UDP)

- 3260 (iSCSI)
- 8080 (HTTP)
- 32803 (NFS nlockmgr TCP)
- 41022 (EMC Support)

Memory sharing

Memory for a CloudArray virtual machine must be provisioned for its exclusive use.

CloudArray is a memory-intensive application. Substantial performance degradation is likely if other virtual machines share the memory allocated to CloudArray.

CHAPTER 2

Logging in to CloudArray

This chapter describes how to log in to the CloudArray interface.

Topics include:

- [Initial login and setup](#)..... 18
- [Initial login and setup without a portal connection](#)..... 18

Initial login and setup

Complete the following steps to log in to the CloudArray user interface.

Note

If you are not connected to the CloudArray portal, see [Initial login and setup \(no portal connection\)](#) for instructions.

Procedure

1. Open any supported browser and navigate to the IP address of your CloudArray.
2. Click **Setup** to configure CloudArray for the first time.

A **Welcome** panel appears. Have your CloudArray License information and your service provider account credentials available before continuing.

If you do not have a CloudArray license, go to <https://cloudarray.com/signup/showSignupForm.action> and complete the registration form to get a trial license. After completing the form, enter your email address and password in step 4.

NOTICE

Do not reuse your CloudArray license. Reuse violates your license agreement with EMC to use a particular CloudArray license for only one CloudArray instance. If more than one instance is using the license, you risk loss of the ability to restore your CloudArray from backup.

3. Click **Next**.

A **License Details** panel appears.

4. Enter the Email Address, Password, and License Key from your cloudarray.com portal registration and click **Next**.

An **Administrator Account** panel appears.

5. Enter the Username and Password for your CloudArray login and click **Next**. Store these credentials in a safe place for future access.
6. Read the End User License Agreement and click **Accept EULA** to continue.
7. Click **User Interface**.

The login panel appears.

8. Enter your username and password and click **Finish**.

The CloudArray user interface opens in **Dashboard** view. [About the CloudArray dashboard](#) describes the interface.

9. Continue the initial setup by configuring cloud providers, caches and volumes.

Initial login and setup without a portal connection

If you are not connected to the CloudArray portal, complete the following steps to log in to the CloudArray user interface.

Procedure

1. From a PC with internet access, open any supported browser and navigate to the IP address of your CloudArray.
2. Click **Setup** to configure CloudArray for the first time.

A **Welcome** panel appears. Have your CloudArray License information and your service provider account credentials available before continuing.

If you do not have a CloudArray license, go to <https://cloudarray.com/signup/showSignupForm.action> and complete the registration form to get a trial license.

NOTICE

Do not reuse your CloudArray license. Reuse violates your license agreement with EMC to use a particular CloudArray license for only one CloudArray instance. If more than one instance is using the license, you risk loss of the ability to restore your CloudArray from backup.

3. Click **CloudArray Portal**.
4. In the **Using The CloudArray Portal** panel, deselect **Enable CloudArray Portal**.
A message on the panel indicates CloudArray will not use the portal.
5. Click **CloudArray License**.
6. In the **Upload Your License** panel, click the **Click to Select CloudArray License File** link.
7. Click **Download License File**.
8. Read the End User License Agreement and click **Accept EULA** to continue.
9. Click **User Interface**.
The login panel appears.
10. Enter your username and password and click **Finish**.
The CloudArray user interface opens in **Dashboard** view. [About the CloudArray dashboard](#) describes the interface.
11. Continue the initial setup by configuring cloud providers, caches and volumes.

CHAPTER 3

CloudArray Dashboard

This chapter describes the CloudArray dashboard.
Topics include:

- [About the CloudArray dashboard](#)..... 22
- [CloudArray Dashboard statistics](#)..... 22

About the CloudArray dashboard

The CloudArray interface initially opens in Dashboard view. The various panels in the Dashboard allow you to view current cache activity and monitor real-time and historical statistics for cache usage, cloud provider data transfer, and volume data transfer.

The CloudArray interface is comprised of the following major elements:

- **Main menu**
The EMC CloudArray main menu is at the left side of the screen. Click menu items to move directly to the corresponding top-level panel and view related sub-menu items.

Note

If you reduce the window size, the main menu moves to the top of the screen.

- **Status indicators**
The status indicators at the top of the screen report the status of CloudArray components. Hover over the green or red indicator light to see more information about the status. Click the status indicators to move directly to the corresponding top-level panel for each item.

CloudArray Dashboard statistics

You can view statistics in historical and real-time mode. Historical mode shows aggregated statistics in 15-minute, hourly, or daily intervals based on the time period chosen. The time period limit is 30 days. Real-time mode shows statistics in 15-second intervals over the last 15 minutes. In real-time mode, aggregate statistics as well as statistics on individual items are available.

To view dashboard statistics, choose **Dashboard** from the CloudArray main menu. Depending on your configuration, the Dashboard view contains the following panels (you may need to scroll down to see them):

Cache Overview

Use the **Showing Cache Details For** drop-down menu to show statistics for individual caches in real-time mode. There are two sub-panels:

- **Current Cache Activity:** Cache pages can be in one of four states :
 - **Dirty:** The cache page contains valid data that needs to be replicated to the cloud.
 - **Clean:** The cache page contains valid data that also exists in the cloud.
 - **Free:** No valid data is available. The cache page is not mapped to any part of the volume.
 - **Used:** The combined amount of clean and dirty cache.

When a host sends data into the CloudArray cache, the replication process starts and continues until all dirty (not yet replicated) data is written to the cloud. This panel shows the amount of free, dirty and clean data in all caches, the total size of all caches, and the read-hit rate of all caches. A second progress bar shows the progress of any replication activity. If there is no dirty data reported, all data written to the volume has been synchronized with the cloud.

- **Cache Usage:** This panel graphs the total size of all caches, the total amount of cache used, and the total amount of free, used and dirty cache.

iSCSI Client Data Transfer

This panel graphs reads and writes in MiB/s across all volumes. Use the **Showing Volume iSCSI Details** drop-down menu to show statistics for individual volumes in real-time mode.

Deduplication Overview

If you have enabled deduplication on a cloud provider, this panel identifies the cache and type of deduplication used, and provides information on the deduplication ratio. There are two sub-panels:

- **Deduplication Ratio:** This panel shows the number of duplicated segments and the number of unique segments. The ratio is the total number of segments divided by the number of unique segments. For example, if there are 6388 duplicated segments and 1759 unique segments, the ratio is 4.63:1 $((6388+1759)/1759)$.
- **Deduplication Ratio Over Time:** This panel shows the deduplication ratio over the last 15 minutes.

Cloud Data Transfer

This panel graphs reads and writes in MiB/s across all cloud providers. Use the **Showing Cloud Provider Details For** drop-down menu to show statistics for individual cloud providers in real-time mode.

Share Data Transfer

This panel graphs reads and writes in MiB/s for a particular share. Use the **Select Share** drop-down menu to show statistics for individual shares in real-time mode.

CHAPTER 4

Volumes

This chapter describes how to create and manage volumes in CloudArray.

Topics include:

- [CloudArray volumes](#)..... 26
- [Viewing volume details](#)..... 26
- [Creating and mapping a volume](#)..... 26
- [Formatting a volume in Windows](#)..... 27
- [Unmapping a volume](#)..... 30
- [Deleting a volume](#)..... 31
- [Expanding volume and share capacity](#)..... 31
- [Migrating a volume from one cache policy to another](#)..... 32
- [Migrating a volume from one cloud provider to another](#)..... 33

CloudArray volumes

To view configured volumes, choose **Volumes** from the CloudArray main menu. From this screen you can:

- [View details about a volume](#)
- [Create a new volume](#)
- [View details about the cloud provider associated with a volume](#)
- [Delete a volume](#)

Viewing volume details

You can view detailed information about each CloudArray volume.

Procedure

1. Choose **Volumes** from the CloudArray main menu.
2. In the **Volume Name** column, click the name of the volume you want to view.

Creating and mapping a volume

To create a new iSCSI volume, complete the following steps:

Procedure

1. Choose **Volumes** from the CloudArray main menu.
2. In the **Volumes** panel, click **Create New Volume**.

The **Create New Volume** panel appears.

3. Complete the following fields in the panel:

Volume Name: Enter name of the volume.

Volume Capacity: Enter a numerical value and choose the unit from the drop-down list.

Select The Cloud Provider For This Volume: Select the radio button next to the appropriate cloud provider in the list.

If desired, enable deduplication for the volume.

Note

To enable deduplication on a volume, deduplication must be enabled on the cloud provider associated with the volume.

Select The Cache For This Volume: Select the radio button next to the appropriate cache in the list.

Select Frontend For This Volume: Map the volume to an iSCSI client now or postpone this action until later.

- To map the volume to an iSCSI client, use the drop-down list to choose the iSCSI client to which this volume will be mapped.

Note

If you choose to postpone the action, you can complete it later by choosing either the **Map To iSCSI Client** button in the **Volumes Details** panel.

- Click **Create Volume** to create the volume mapped to the host you selected.

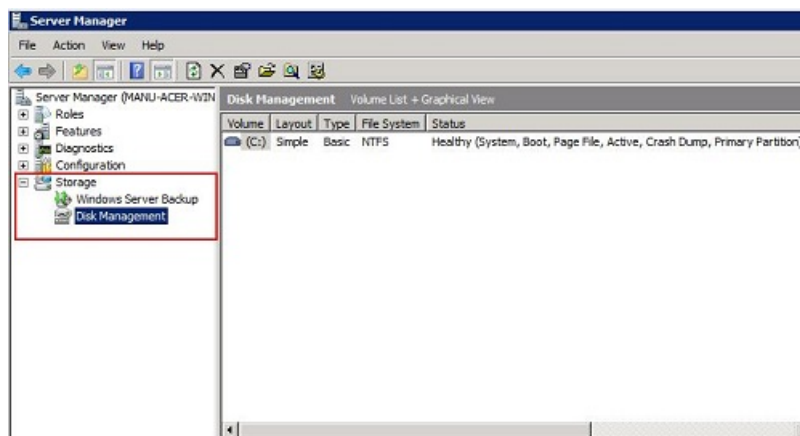
Formatting a volume in Windows

To detect, initialize, and format a new iSCSI volume in Windows, complete the following steps:

Procedure

- Right-click **My Computer** and choose **Manage**.

The following window opens.



- Expand the **Storage** category and choose **Disk Management**.
- Right-click **Disk Management** and choose **Rescan Disks**.
Windows checks for new volumes and should detect a new volume. For example:



- Right-click the text labeled **Unknown** and choose **Online**.



The disk status changes from **Offline** to **Not Initialized**.

- Right-click and choose **Initialize Disk**.




6. Choose the disk to be initialized and select a partition style.



Note

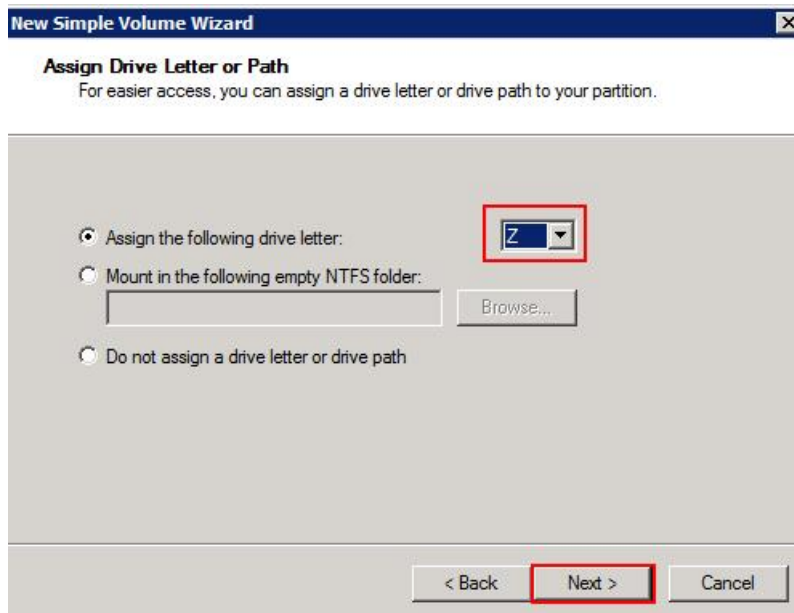
When creating a volume less than 2TB that may require expansion beyond 2TB, use the GPT partition style. The MBR partition style does not support partitions bigger than 2TB and converting from MBR to GPT is not supported in Windows without data loss.

7. Click **OK**.
- The disk should now have a status of **Online**.
8. Right-click the **Unallocated** label and choose **New Simple Volume**.

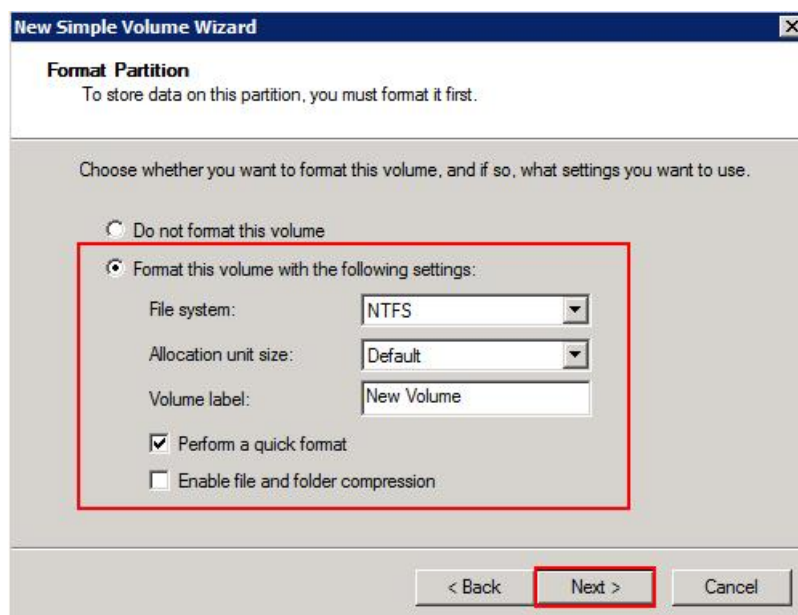


Windows displays the **New Simple Volume Wizard**.

9. Click **Next** twice and in the **Assign Drive Letter or Path** panel, choose a drive letter from the drop-down list and click **Next**.



10. In the **Format Partition** panel, choose the **File system** type from the drop-down list. Ensure that the **Perform a Quick Format** option is checked and click **Next**.



Note

If you plan to expand your file system at a later date and you are near a size boundary, select the next highest allocation unit size (or cluster size in NTFS parlance). For example, if you are creating a 10TB volume, the default unit size is 4k. If you expect to expand the volume to 20TB, you should use an 8k unit size.

11. Click **Finish**.

Windows begins to format the volume. The following displays once the process completes:

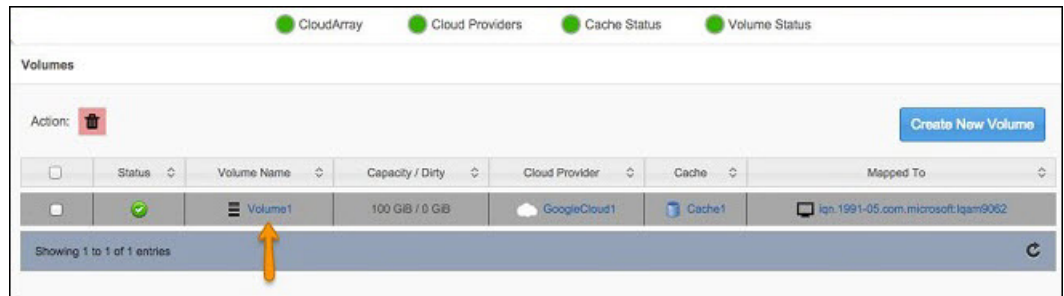


The volume should now be listed under **My Computer**. You can use this volume for host I/O.

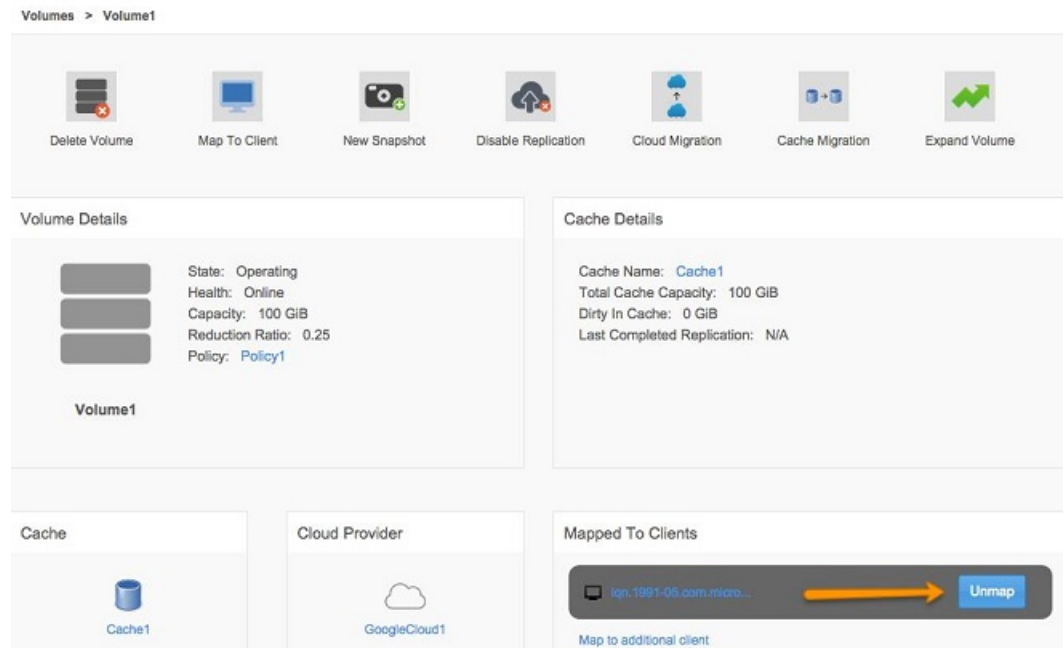
Unmapping a volume

Procedure

1. From the CloudArray main menu, select **Volumes** and then click the name of the volume to unmap.



2. From the selected Volume panel, Click **Unmap** from the **Mapped to Clients** section.



The volume is now unmapped and will disappear from the host upon rescanning with the host's volume management utility.

Deleting a volume

Only unmapped volumes can be deleted. Before deleting a CloudArray volume, ensure that you have unmapped it.

NOTICE

Volume deletion is an irreversible process. All data for the volume will be deleted from the cloud provider and CloudArray's cache.

Procedure

1. From the CloudArray main menu, select **Volumes** and then click the name of the volume to be deleted.
 2. Click **Delete Volume**.
 3. When you are prompted for confirmation, click **Yes** to permanently remove all volume and resident data from CloudArray and the cloud provider.
-

Note

Deletion of a CloudArray volume ensures that all data for that volume is removed from the cloud. This is why volumes may remain in the "deleting" state for an extended period.

Expanding volume and share capacity

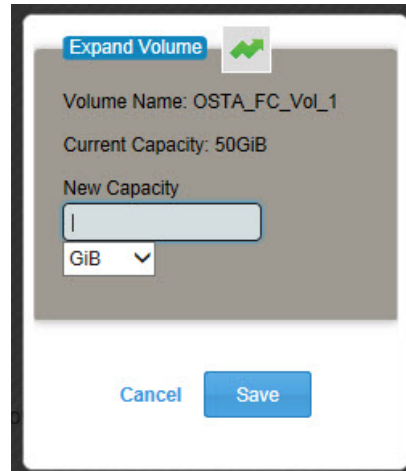
Procedure

1. From the CloudArray main menu, click **Volumes**.
2. Select the desired volume or volume underlying the share you wish to expand.
3. In the resulting panel, click **Expand Volume** (or **Expand Share**, as applicable).

The following panel illustrates volume expansion. Expanding a share is done similarly.



4. Enter the new capacity for the volume and click **Save**.



Note

For iSCSI volumes, use the utilities pertinent to the operating system running on the iSCSI host mapped to that volume to take advantage of the volume's additional capacity.

When expanding a CIFS or NFS share, the process is similar to that shown above. The share will expand and offer the increased capacity to the clients accessing it. The **Expand Share** button is accessible from both the **Volumes** and the **Shares** tabs.

Migrating a volume from one cache policy to another

Before you begin

Before starting a migration, the following requirements must be met:

- Only volumes can be migrated. CIFS and NFS shares cannot be migrated.
- The volume must not be mapped to a host.
- The old and new caches must have the same page size.
- The policies involved must both use the same cloud provider.

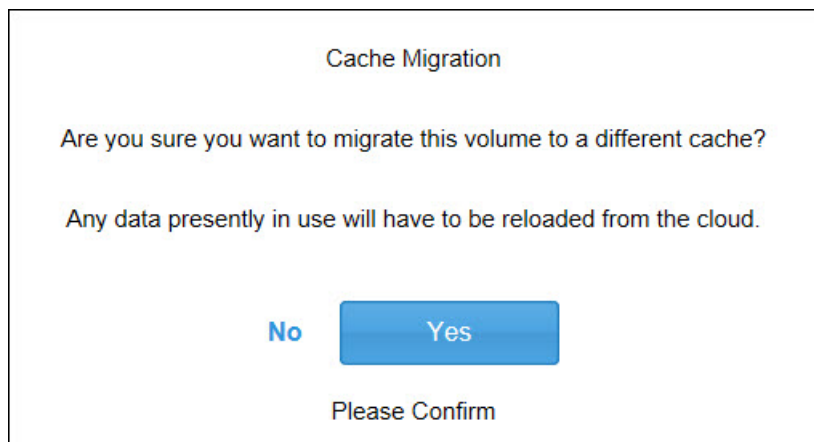
You can migrate a volume from one cache policy to another dynamically (without host IO interruption).

Procedure

1. From the CloudArray main menu, click **Volumes** and click the volume name to migrate.
2. Click the **Cache Migration** icon.



You are prompted to confirm.



3. Click **Yes** to continue.
4. Choose the cache to migrate to, then click **OK** to apply the changes.

Migrating a volume from one cloud provider to another

Before you begin

Before starting the migration, the following requirements must be met:

- The cache must be large enough to hold the entire contents of all volumes using it.
- The source and target cloud providers must have the same encryption and compression attributes (turned on or off).
- The volume being migrated must not be an exposed snapshot and must not contain any exposed snapshots.
- The target cloud provider must be online.

CAUTION

Migrating to a new cloud provider requires restarting CloudArray. You must schedule downtime for this operation.

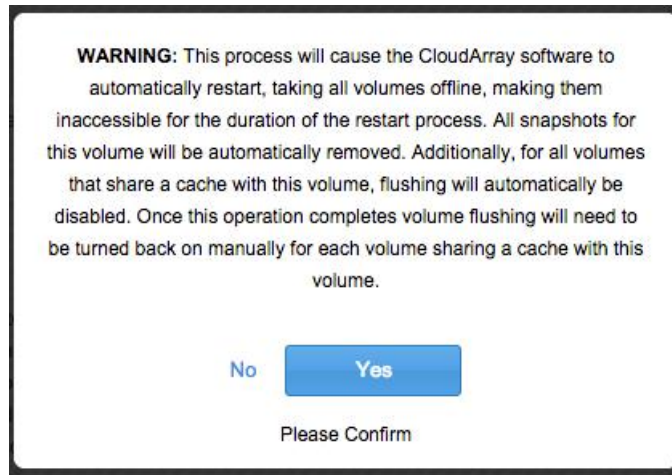
You can migrate a CloudArray volume or share volume from one cloud provider to another. This process involves using the cache as an intermediate storage facility for copying data from the source cloud provider before replicating to the new target cloud provider.

Procedure

1. Create a target cloud provider, making sure encryption and compression attributes match those of the source cloud provider.
2. Create a target provisioning policy containing the new target cloud provider as well as the *same* cache that the volume being migrated currently uses.
3. From the CloudArray main menu, click **Volumes** and click the volume name to migrate.
4. Click the **Cloud Migration** icon.

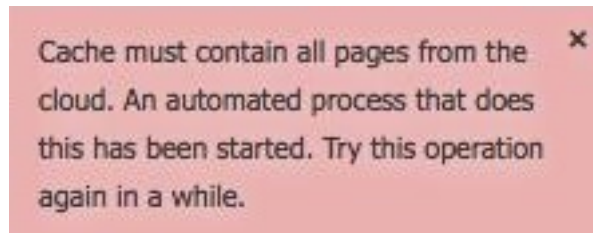


5. A warning dialog states that CloudArray software will be restarted and describes other implications related to the migration. Click Yes to proceed.



6. Select the cloud provider to migrate the volume to, then select **OK** to start the process.

The migration process involves ensuring that the cache contains the entire volume data, including all meta-data. If needed, this process copies data from the source cloud to the cache. This part of the operation takes time to complete, therefore you may be presented with the following message:



This message indicates that the copy process has been started. Once the process completes, the migration operation will need to be re-initiated. After re-initiating the process, another message indicates that the volume is migrating.

7. When the migration completes, the CloudArray software restarts. Check that the operation completed as expected by doing the following:
 - Go to the **Volume** details page for the volume just migrated and ensure that the cloud provider is now the target cloud provider and not the original source cloud provider.
 - Bring the volume online on the host and perform data verification.

Once this is done you can re-enable replication to cloud from the **Volume** details page for the volume that was migrated. Once all data has been flushed to the cloud and data has been verified from the host side, the original source provisioning policy and cloud provider can be removed. This removes the data from the original cloud bucket.

CHAPTER 5

File Shares

This chapter describes how to create CloudArray file shares and share them using either Common Internet File System (CIFS) or Network File System (NFS) protocol. Windows operating systems use CIFS for file sharing. Unix and Linux operating systems use NFS.

Topics include:

- [Enabling CIFS and NFS file sharing](#)..... 36
- [Creating a CIFS share](#)..... 36
- [CIFS security](#)..... 37
- [NFS security](#)..... 43

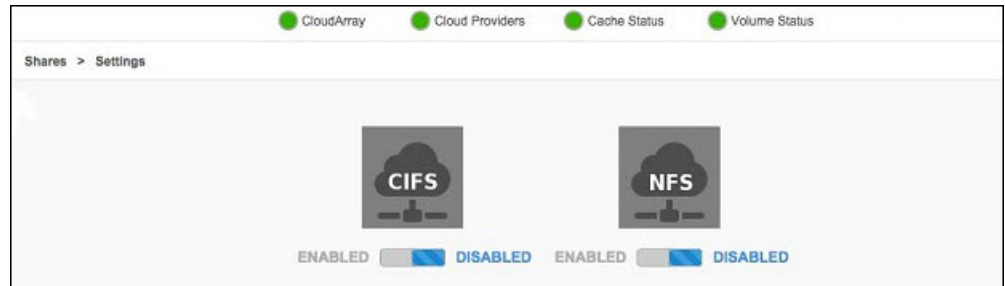
Enabling CIFS and NFS file sharing

If you did not enable CIFS or NFS through the initial Configuration Wizard, you can enable it by following these steps.

Procedure

1. Choose **Shares > Settings** from the CloudArray main menu.

The **Settings** panel appears.



2. Set the CIFS or NFS slider to **ENABLED**.
3. To open the configuration screen, do one of the following:
 - Click the **CIFS** icon. The **Settings** tab opens, where you can choose to use Microsoft Active Directory to authenticate CloudArray users.
 - Click the **NFS** icon. The **Settings** tab opens, where you can select either the AUTH_SYS or Kerberos security method. AUTH_SYS is the default setting.

Creating a CIFS share

Procedure

1. Choose **Shares > Settings** from the CloudArray main menu.
2. In the **Settings** panel, set the **CIFS** slider to **ENABLED**.
3. Choose **Shares > Settings > CIFS** from the CloudArray main menu.

The **CIFS** panel appears.

4. Click the **Settings** tab if it is not already selected.
5. Complete the fields and click **Save** to create the CIFS share.
6. Select the **Shares** tab to configure the CIFS share.
7. Click **New Share** to create the CIFS share.

The **New CIFS Share** panel appears.

8. Complete the **Share Name** and **Capacity** fields and select the **Cache** and **Cloud Provider** from the drop down boxes. If desired, select local storage or enable deduplication on the share.

Shares with a 16 TiB or less capacity will be general purpose shares by default. Check the **Expandable Greater than 16 TiB** checkbox to make them expandable at a future time. This changes the share type to a large archive share. Shares with a 17 TiB or larger capacity will automatically be large archive shares.

9. Click **Save** when done.

The new share appears in the **Selected CIFS Share** panel.

CIFS security

CloudArray supports two methods for authenticating users and groups on its shares:

- [Microsoft Active Directory integration](#)
- [CloudArray-managed CIFS security](#)

Note

You can only use one method. The method you choose applies to all shares.

Setting CIFS security with Microsoft Active Directory

Complete the following steps to authenticate CloudArray users and groups through Microsoft Active Directory.

Procedure

1. Set the **Use Active Directory** toggle control in the **Shares > CIFS > Settings** panel to **Yes**.

Note

When **Use Active Directory = Yes**, the **Users** and **Groups** tabs disappear, since Active Directory controls access to the share.

2. Enter the Primary Domain Controller info.

Name: The name that will appear in Active Directory

Description: Free text field

Workgroup: Free text field

Primary Domain Controller FQDN: Fully qualified domain name

3. Click **Lookup** to auto-populate the **Domain** and **Controllers** fields.

Note

The DNS servers used by CloudArray must have a valid entry for the domain controller server. This must match the domain controller identification.

4. Complete the following fields:

WINS (Optional): Your WINS server

Domain Administrator: Primary server's admin login name

Domain Password: Primary server's administrator password

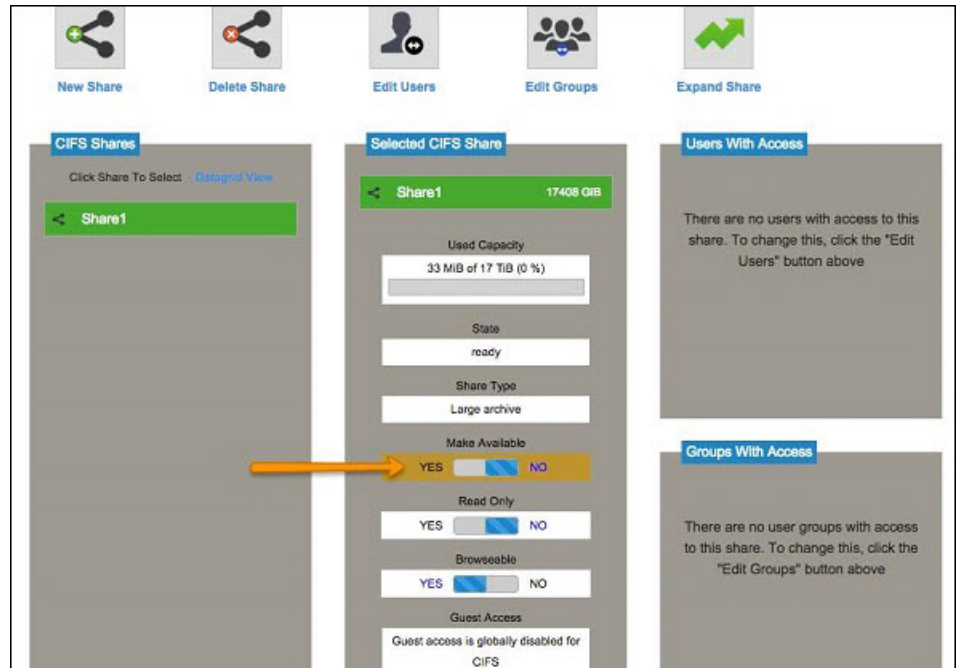
Guest Access: Select this checkbox if you want to allow guest access to the CIFS share

Enable Access Based Enumeration For CIFS Shares: Select this checkbox if you want to use Access Based Enumeration to prevent directories from being visible to users without permission to access them.

Note

To cache these security credentials, select the **Allow CloudArray to cache these credentials** checkbox.

- Click **Save** to join the CloudArray appliance to the domain.
If a connection is successful you will see a green joined indicator for the share.
- On the **Selected CIFS Share** tab, click the white box next to YES under **Make Available** to give network access to the share.



- From a Windows computer that is a member of the domain, log in as a domain administrator and access the newly-created CloudArray share.
 - Using that computer, change the permissions on the share and any folders created within to suit the security needs for that share.
-

Note

If you use CloudArray to configure shares and set up users and groups and then join that CloudArray to an Active Directory domain, CloudArray will no longer use its local configuration. Instead, share access will be governed by the Active Directory domain access control and permissions defined for each share. Use Active Directory instead of the CloudArray GUI to manage access to the shares.

Setting CIFS security with CloudArray management

Complete the following steps to manage access to the share and its contents.

Procedure

- Set the **Use Active Directory** toggle control in the **Shares > CIFS > Settings** panel to **No**.

Either allow Guest Access by checking the option for **Allow Guest Access** or complete the following steps to specify users and or groups which will have access to this share.

2. Click the **Users** tab and then click **New User**.

The following panel appears:

3. Complete the following fields:

Name: The username for the user who will have access to this share

Password: The password assigned to the user

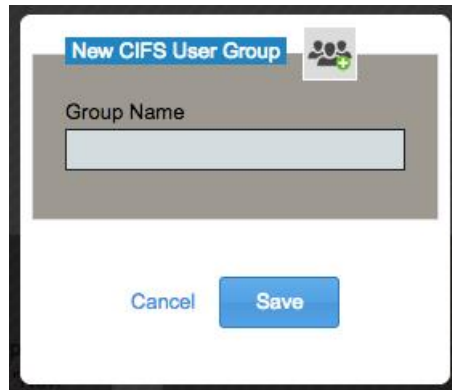
Confirm Password

4. Click **Save** to create the user.

Optionally you can assign this user to a group for easier management of the share.

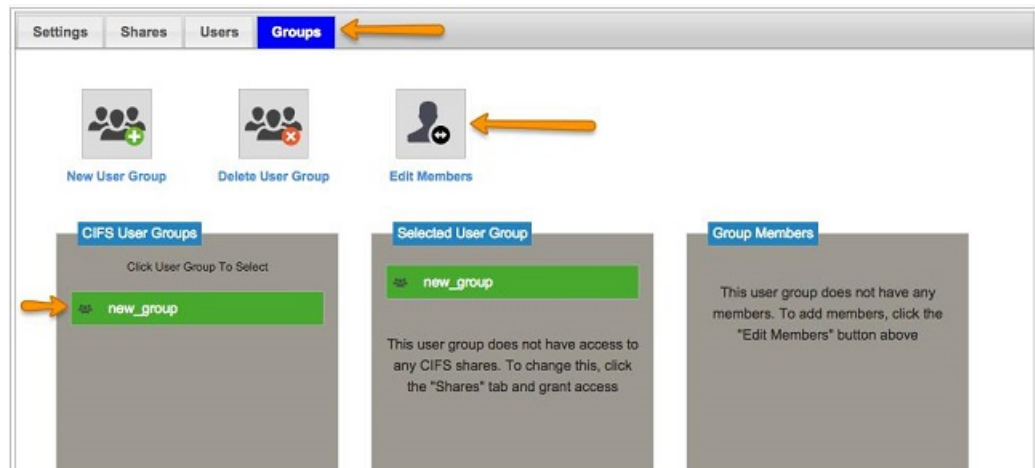
- Click the **Groups** tab and then click **New User Group**.

The following panel appears:



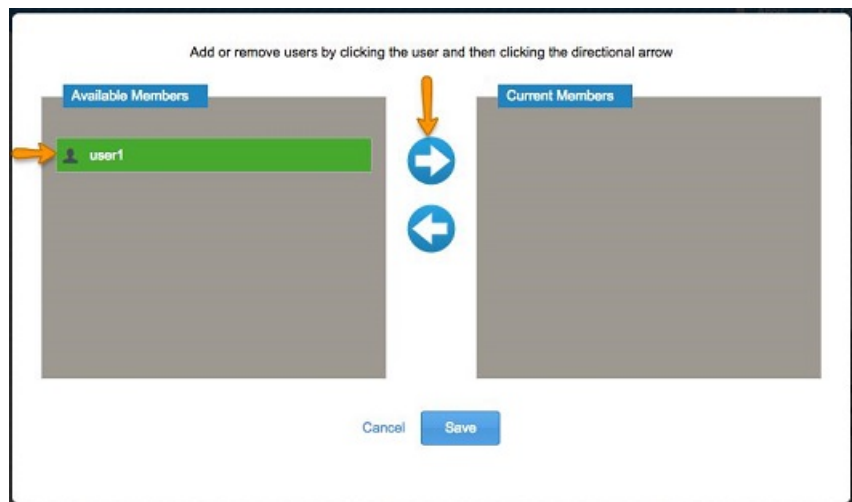
A dialog box titled "New CIFS User Group" with a group icon. It contains a text input field labeled "Group Name". At the bottom are "Cancel" and "Save" buttons.

- Enter a name for the group and click **Save**.
- To assign group members, select the group in the **CIFS User Groups** panel and then click **Edit Members**.



The interface shows tabs for Settings, Shares, Users, and Groups. The Groups tab is active. Below the tabs are three icons: "New User Group", "Delete User Group", and "Edit Members". Below these are three panels: "CIFS User Groups" (with a list containing "new_group"), "Selected User Group" (showing "new_group" and a message about access), and "Group Members" (showing a message that the group has no members). Orange arrows point to the Groups tab, the Edit Members icon, and the "new_group" entry in the CIFS User Groups list.

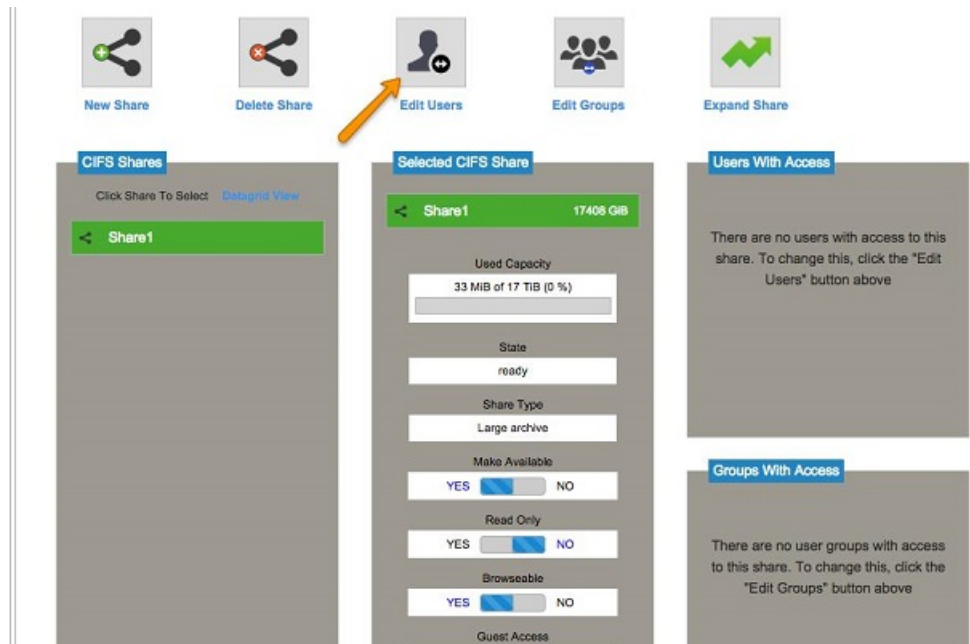
- Choose the user from the **Available Members** list and click the right directional arrow to add the user to the **Current Members** list.



A dialog box titled "Add or remove users by clicking the user and then clicking the directional arrow". It contains two panels: "Available Members" (with a list containing "user1") and "Current Members" (empty). Between the panels are two blue circular buttons with white arrows pointing right and left. At the bottom are "Cancel" and "Save" buttons. Orange arrows point to the "user1" entry and the right-pointing arrow button.

- Click **Save** when done.

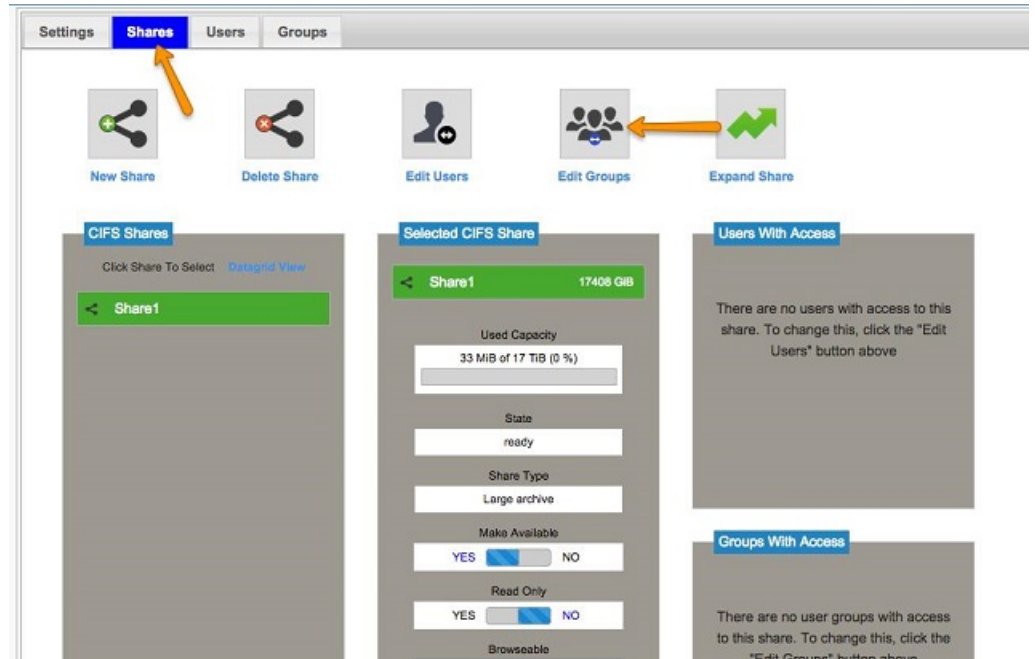
10. To associate users or groups with the CIFS share, add a user by clicking the **Shares** tab and then clicking **Edit Users**.



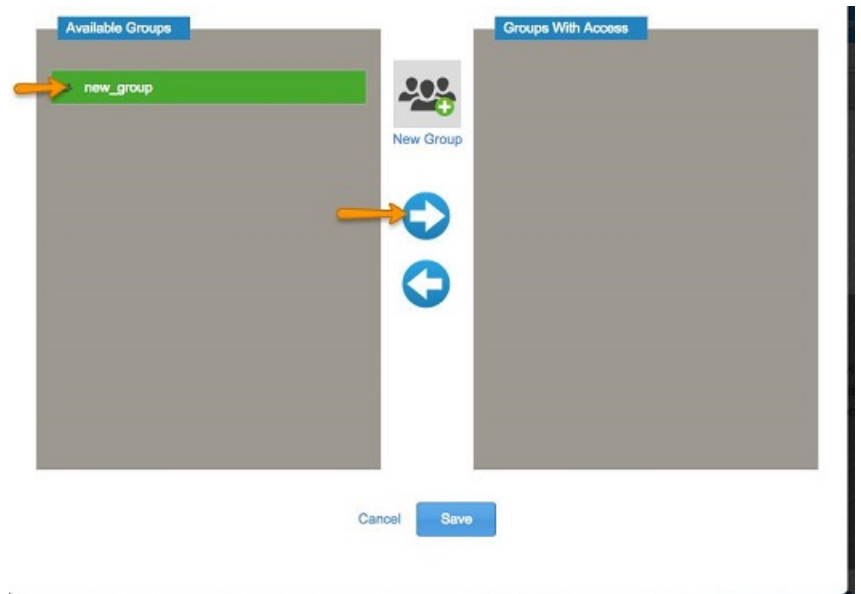
11. Choose a User from the **Available Users** list and click the right directional arrow to add the user to the **Users with Access** list.



12. Click **Save** when done.
13. To add a group, click the **Shares** tab and then click **Edit Groups**.



14. Choose a group from the **Available Groups** list and click the right directional arrow to add the group to the **Groups With Access** list.



15. Click **Save** when done.
16. Configure the following options as desired on the **Shares** tab for individual share levels.

Selected CIFS Share

Share1 17408 GiB

Used Capacity
33 MiB of 17 TiB (0 %)

State
ready

Share Type
Large archive

Make Available
YES ☒ NO

Read Only
YES ☒ NO

Browseable
YES ☒ NO

Guest Access
Guest access is globally disabled for CIFS
[Change Settings](#)

Make Available: This option is turned off by default on a new share. It allows a user to set up the appropriate permissions on the share and then enable sharing using this option. To allow computers to use the share, set this option to **Yes**.

Read Only: When set to **Yes**, this option removes the write permission from the share.

Browseable: When set to **Yes**, this option allows a user to browse for the share, without specifying an absolute path.

Guest Access: **Change Settings** gives you the option to allow Guest Access.

NFS security

NFS supports several different methods of security. CloudArray supports the AUTH_SYS and Kerberos v5 security methods. These security methods are global for all NFS shares, meaning that all NFS shares use the same security method.

- **AUTH_SYS** is a traditional UNIX authentication method where users self-identify using their UNIX user identifiers and group identifiers (UIDs and GIDs). This method is supported by NFSv3 and NFSv4 clients.
- **Kerberos V5** is a network authentication protocol which relies on a trusted host (KDC) to authenticate the identity of other hosts on an insecure network. This method is supported by NFSv4 clients.

Setting NFS security with AUTH_SYS

Note

If you are using NFS security with Kerberos V5, see [Setting NFS security with Kerberos V5](#) on page 45.

Procedure

1. Choose **Shares > Settings** from the CloudArray main menu.
2. In the **Settings** panel, set the **NFS** slider to **ENABLED**.
3. Choose **Shares > Settings > NFS** from the CloudArray main menu.
The **NFS** panel appears.
4. Click the **Settings** tab if it is not already selected.
5. In the **NFS** panel, select **AUTH_SYS** as the security type by toggling the blue selection slider to the right.
6. Click the **Shares** tab to create the share.
7. Click **New Share** to create the NFS share.
The **New NFS Share** panel appears.
8. Complete the **Share Name** and **Capacity** fields and select the **Cache** and **Cloud Provider** from the drop down boxes. If desired, select local storage or enable deduplication on the share.

Shares with a 16 TiB or less capacity will be general purpose shares by default. Check the **Expandable Greater than 16 TiB** checkbox to make them expandable at a future time. This changes the share type to a large archive share. Shares with a 17 TiB or larger capacity will automatically be large archive shares.
9. Click **Save** when done.
10. In the **NFS Shares** column, select the share.
11. In the **Selected NFS Share** column, move the **Make Available** slider to **YES** to make the share available to CloudArray.
12. Select the **Clients** tab to add a client that can be used to access NFSv3 shares.
13. Click **New Client**.
14. Enter a hostname, an IP address, or IP with subnet range of the host, then click **Save**.
15. To assign access to the client, select the **Shares** tab.
16. Select the **Edit Clients** icon.
17. Select the **Available Clients** from the column on the left.
18. Use the right directional arrow to add the client to the **Clients with Access** column.
19. Click **Save** when done.

Setting NFS security with Kerberos V5

Procedure

1. Choose **Shares > Settings** from the CloudArray main menu.
2. In the **Settings** panel, set the **NFS** slider to **ENABLED**.
3. Choose **Shares > Settings > NFS** from the CloudArray main menu.
The **NFS** panel appears.
4. Click the **Settings** tab if it is not already selected.
5. In the **NFS** panel, select **Kerberos** as the security type by toggling the blue selection slider to the right.

Note

NFSv4 clients are required when using this security method.

6. Complete the fields under **Kerberos Settings** to specify your KDC server.
Realm: This is your Kerberos realm name, which is usually your domain name in all capital letters. For example, the domain name engineering.org would have ENGINEERING.ORG as the corresponding realm name.
Master Key Distribution Center: This is the fully qualified domain name (FQDN) of the master Key Distribution Center (KDC) for your realm. For example, kdc.engineering.org.
All Key Distribution Centers: This is a comma-separated list of FQDNs for all of the KDCs for your realm. If you only have one KDC, this field should contain the master KDC.
Domain: The DNS domain name. For example, engineering.org
Domain Administrator: The administrative account for the realm. For example, john/admin.
Domain Password: The password for the domain administrator.

Note

To cache these security credentials, select the **Allow CloudArray to cache these credentials** checkbox.

7. Click **Save**.
8. Select the **Shares** tab.
9. Click **New Share** to configure the share.
10. Complete the **Share Name** and **Capacity** fields and select the **Cache** and **Cloud Provider** from the drop down boxes. If desired, select local storage or enable deduplication on the share.

Shares with a 16 TiB or less capacity will be general purpose shares by default. Check the **Expandable Greater than 16 TiB** checkbox to make them expandable at a future time. This changes the share type to a large archive share. Shares with a 17 TiB or larger capacity will automatically be large archive shares.

11. Click **Save** when done.
12. In the **NFS Shares** column, select the share.
13. In the **Selected NFS Share** column, move the **Make Available** slider to **YES** to make the share available to CloudArray.
14. Select the desired **Security** option.

The default option is **krb5** for Kerberos V5. The following list describes the Kerberos security levels, from least secure with lowest complexity to highest security and complexity with higher computing overhead.

- **krb5**: Authenticates using Kerberos V5 instead of local UNIX UIDs and GIDs.
- **krb5i**: Adds to krb5 integrity checking of NFS operations using secure checksums to prevent data tampering.
- **krb5p**: Adds to krb5i securing the connection by encrypting NFS traffic to prevent traffic sniffing.

CHAPTER 6

Deduplication

This chapter describes deduplication.

Topics include:

- [Deduplication](#)..... 48
- [System requirements for deduplication](#)..... 48
- [Deduplication configuration](#)..... 48

Deduplication

CloudArray deduplication efficiently utilizes storage space by reducing data size, reducing the amount of space required to store the data. CloudArray deduplication:

- Provides up to a 10x space savings
- Functions across volumes to maximize savings
- Is limited to a single cloud provider in a CloudArray instance

Deduplication reduces data size by the following process:

1. Data is written to the front-end cache.
2. During replication, data is reduced by the deduplication engine. Unique segments are stored along with metadata in the deduplication cache.
3. The data is compressed and encrypted, then the reduced pages are replicated to the cloud. The unique segments and metadata are also replicated in a single transaction.

System requirements for deduplication

Minimum requirements for a CloudArray virtual machine with deduplication enabled are:

- 4 CPU cores
- 8G of RAM plus 1G of RAM per terabyte of total front-end and back-end cache

Best practices include:

- The front-end cache size should be a minimum of 10% of the volume size.
- All post-deduplicated (reduced size) data should fit in the deduplication cache. For information on deduplication cache sizing, see [Deduplication cache sizing](#) on page 57.

Note

Deduplication is not supported on CloudArray providing storage to VMAX All Flash, VMAX3, VPLEX or DLM.

Deduplication configuration

You configure deduplication by first enabling it on a cloud provider. (See [Configuring remote cloud providers](#) on page 50.) You then select the cloud provider with deduplication enabled when you create volumes and shares.

Once enabled, deduplication can not be disabled.

Deduplication requires the use of a separate deduplication cache. For information on how to size the cache, see [Deduplication cache sizing](#) on page 57.

CHAPTER 7

Cloud Providers

You can choose from several remote cloud providers or use your NFS share as your local cloud.

Topics include:

• Cloud providers	50
• Configuring remote cloud providers	50
• Viewing cloud provider details	50
• Configuring NFS as a local cloud provider	51
• Deleting cloud providers	51

Cloud providers

To view configured cloud providers, choose **Cloud Providers** from the CloudArray main menu. From this screen you can:

- [View details about a configured cloud provider](#)
- [Configure a new cloud provider](#)
- [Delete a cloud provider](#)

Configuring remote cloud providers

You can choose from several remote cloud providers or use your [NFS share](#) as your local cloud.

CloudArray supports using multiple providers within the same CloudArray instance. You can create multiple providers by repeating the following procedure.

Procedure

1. Choose **Cloud Providers** from the CloudArray main menu
The **Cloud Providers** panel appears.
2. Click **Configure New Cloud Provider**.
3. Select the desired cloud provider from the drop-down menu and click **Continue**.
4. On the **Configure Cloud Provider** panel, provide the required information for your cloud storage account.
5. If desired, enable data encryption and compression.

CloudArray uses multi-layered AES-256 bit encryption for data protection and zlib for data compression.

6. If desired, enable deduplication.

Deduplication can only be enabled on one cloud provider in a CloudArray instance.

Note

Once enabled, deduplication can not be disabled.

7. If you are enabling deduplication, select the cache to use for deduplication.

Deduplication requires a separate cache. For information on how to size the cache, see [Deduplication cache sizing](#) on page 57.

8. If desired, enable data encryption for the deduplication data.
9. Click **Save Cloud Provider**.

CloudArray attempts to attach the cloud provider. If authentication is successful, the provider is added to the list of configured providers in the **Cloud Providers** panel.

Viewing cloud provider details

You can view detailed information about each cloud provider. If deduplication is enabled on the cloud provider, the panel includes information on the deduplication

ratio and the ratio over time. (See [CloudArray Dashboard statistics](#) on page 22 for information on the **Deduplication Overview** panel.)

Procedure

1. Choose **Cloud Providers** from the CloudArray main menu.
2. In the **Cloud Provider** column, click the name of the cloud provider you want to view.

Configuring NFS as a local cloud provider

CAUTION

NFS is not supported as a production cloud option.

Procedure

1. Click **Cloud Providers** in the CloudArray main menu or from the status indicators at the top of the screen.
2. Click **Configure New Cloud Provider**.
3. To use your NFS share as your local cloud, select **NFS** from the cloud provider drop-down list and click **Continue**.
4. Provide the following information to mount your NFS share as a cloud provider:
 - a. **Cloud Provider Name:** Enter text to describe the provider, for example, My NFS Share.
 - b. **Server:** Enter the FQDN or IP address of your NFS server, for example, nfs1.engineering.org
 - c. **Path:** Enter the share and directory path, for example, share1/cloudarray100.
 - The share and directory path must already exist on your NFS server.
 - Your NFS server must maintain sufficient free space for each share used by your CloudArray.
 - The path you specify should be an empty directory.
 - d. **Options:** Enter mounting options.

Note

This field should only be used by experienced users to define custom mounting options for NFS. The default options set by CloudArray automatically are: auto, rw, and intr.

Deleting cloud providers

You can delete one or more cloud providers from the **Cloud Providers** screen. Alternatively, you can delete an individual cloud provider from its details screen.

Procedure

1. Choose **Cloud Providers** from the CloudArray main menu.

The **Cloud Providers** panel appears.

2. Do one of the following:

- Click the checkbox next to each cloud provider you want to delete, then click the trash barrel **Action** button to delete the caches.
- In the **Cloud Provider** column, click the name of the cloud provider you want to delete, then click **Delete**.

The cloud provider is removed from the list on the **Cloud Providers** panel.

CHAPTER 8

Caches

CloudArray addresses bandwidth and latency issues typically associated with cloud storage by taking advantage of local storage, called cache. This chapter describes how to configure caches.

Topics include:

• Cache overview	54
• Viewing cache details	54
• Adding cache sources	54
• Viewing cache source details	55
• Allocating cache	56
• Expanding existing cache	58
• Deleting caches	59

Cache overview

CloudArray's policy-driven cache ensures the proper level of accessibility and performance based on the data stored. The cache is local on the appliance and delivers high performance while asynchronously replicating data to the cloud. Each cache can be sized and assigned a policy to support a percentage of your data based on current needs.

All data written to volumes and shares is stored in the cache for some period of time. Once replicated to the cloud, CloudArray removes the data from the cache as space is needed.

As the percentage of cache pages that are unreplicated approaches 100%, the cache throttles access to CloudArray to postpone reaching the 100% mark. When all cache pages are unreplicated, all volumes using that cache are taken offline until a sufficient number of pages have been replicated to the cloud. To prevent a cache bottleneck, you should appropriate enough cache space to avoid the cache becoming more than 80% unreplicated.

Note

The CloudArray virtual edition ships with a 25GB trial cache which is prepared for use automatically via the Trial Configuration Wizard. For production use, CloudArray benefits from a larger cache.

Cache size

The amount of cache CloudArray needs is dependent on the amount and type of data you are moving to the cloud. You can configure cache in three ways:

- **Fully cached** - The size of the cache is equal to or greater than the aggregate size of all volumes using the cache. This solution provides two copies of the data, one local and one in the cloud. It also provides optimal performance because data will not need to be read from the cloud under most circumstances.

Viewing cache details

You can view detailed information about each CloudArray cache.

Procedure

1. Choose **Cache Management > Caches** from the CloudArray main menu.
2. In the **Cache Name** column, click the name of the cache you want to view.

Results

The **Cache Details** panel displays the state of the cache, its size, and information on what volumes are using the cache.

Adding cache sources

To increase cache capacity, first add one or more virtual disks to CloudArray from the hypervisor. After attaching the virtual disks, follow the steps below to scan for the new storage and create a new cache or expand the default cache.

Procedure

1. In your hypervisor, add one or more virtual disks to your CloudArray virtual or physical appliance.
2. Choose **Cache Management** > **Cache Sources** from the CloudArray main menu.

A panel similar to the following appears:

Cache Source Name				Scan For Additional Cache Sources	
Cache Source Name	Path	Capacity / Raw Available	Caches Using		
Pool_1	/dev/sdc	48428 GiB / 48328 GiB	Cache1 (100 GiB)		
Pool_2	/dev/sdb	37052 GiB / 37052 GiB			
Showing 1 to 2 of 2 entries					

3. Click the **Scan for Additional Cache Sources** button.

This operation probes for new devices on the CloudArray virtual or physical appliance and creates a pool for each new device it finds. Once the scan completes, the results display in a pop up window.



For example, in the image below, two new devices were added and assigned the names "Pool_1 and Pool_2".

Cache Source Name				Scan For Additional Cache Sources	
Cache Source Name	Path	Capacity / Raw Available	Caches Using		
Pool_1	/dev/sdc	48428 GiB / 48328 GiB	Cache1 (100 GiB)		
Pool_2	/dev/sdb	37052 GiB / 37052 GiB			
Showing 1 to 2 of 2 entries					

Use these newly added cache sources to add a new cache or expand an existing cache.

Viewing cache source details

You can view detailed information about each cache source.

Procedure

1. Choose **Cache Management** > **Cache Sources** from the CloudArray main menu.
2. In the **Cache Source Name** column, click the name of the cache source you want to view.

Results

The **Cache Source Details** panel displays information about the cache source, including what caches are using it.

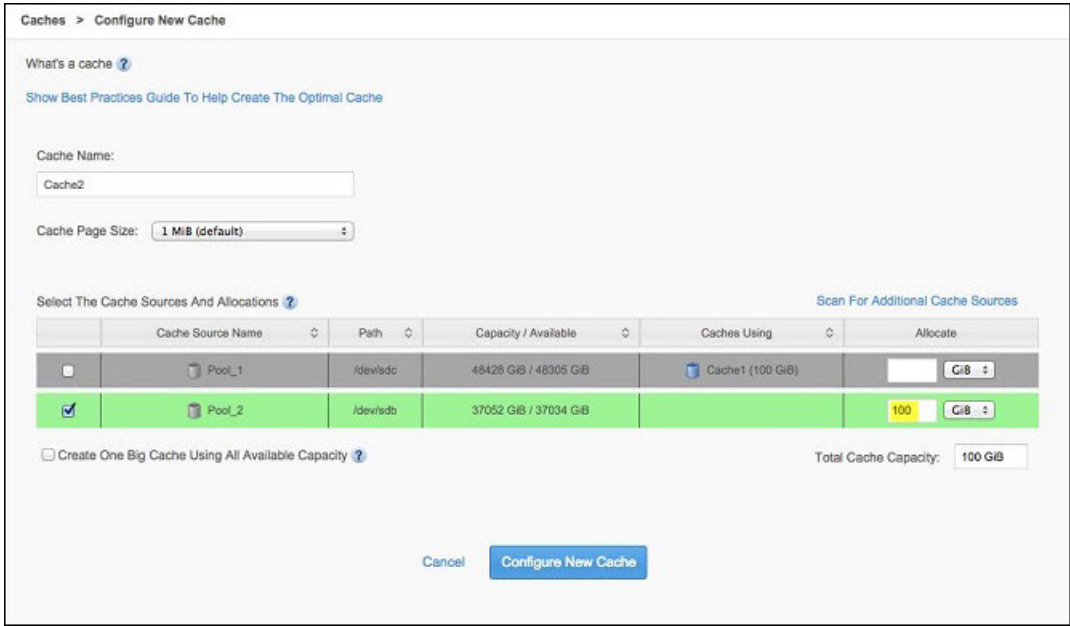
Allocating cache

Procedure

- 1. Choose **Cache Management > Caches** from the CloudArray main menu.
- 2. Click the **Configure New Cache** button.



The **Configure New Cache** panel appears.



- 3. Type a name in the **Cache Name** text field.

Note

Special characters are not allowed.

- 4. Check the desired cache source and in the **Allocate** text box, indicate how much of the available cache source to allocate to your new cache. Enter values equal to or smaller than the available capacity of each cache source you want to use for the new cache. You can select full or partial availability from one or more cache sources.

Note

You should plan for occasional network or cloud provider outages that will require the data to remain in cache for an extended period of time. EMC recommends that you over-provision the cache storage to avoid a high ratio of unreplicated pages. A good rule of thumb would be to configure the cache to be approximately twice the size of your anticipated cache.

5. Alternately, select the **Create One Big Cache Using All Available Capacity** checkbox. Use this option only when you are sure that you will use a single cache.
 6. Select a page size from the **Cache Page Size** drop-down menu. The default value of 1MB is designed to provide best performance in most use cases. Tune the cache page size based on the cloud provider you are using.
-

Note

Cache page size is a significant factor in CloudArray performance. Larger page sizes work well for applications that perform large sequential writes to the volume. Most backup applications are in this category. Smaller page sizes are preferred when small random reads are needed. This is typically seen when the CloudArray volume is used for a file system with generic user data. Larger pages also have smaller CloudArray RAM requirements. Using a page size smaller than 1MB is rarely appropriate. Cache pages greater than 1MB require an internal sector size of 4KB rather than the typical 512 bytes. The larger sector size is not recommended for transactional workloads such as database or typical file system applications.

Note

VMware vSphere Hypervisor (ESXi) does not support a cache sector size of 4096 bytes. To create CloudArray iSCSI volumes with a sector size of 512 bytes for ESXi systems, use the default cache page size (1 MiB).

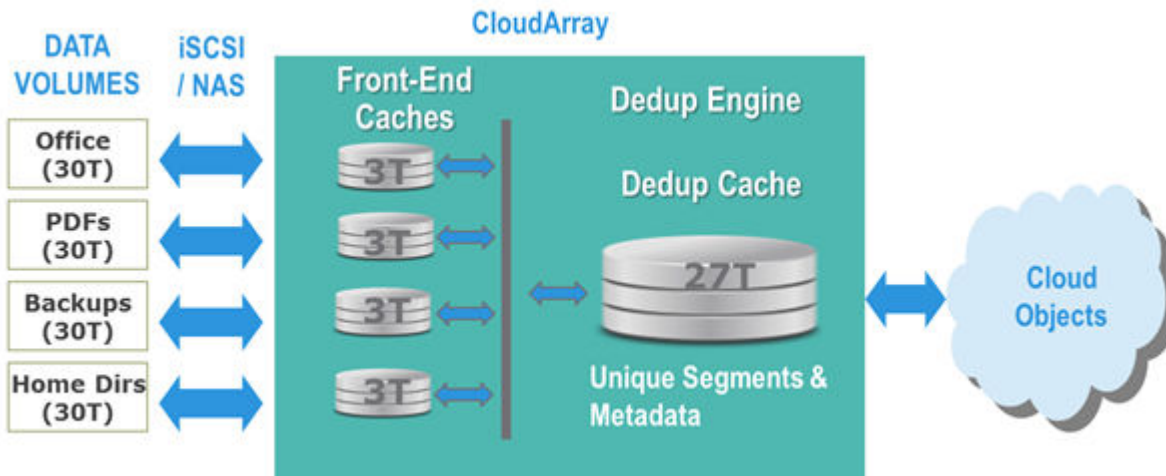
Deduplication cache sizing

Creating appropriate cache sizes requires understanding the volume of data that will be deduplicated. EMC recommends that the front-end cache be at least 10% of the volume size, and that all post-deduplicated (reduced) data fit in the deduplication cache to minimize the performance impact of fetching deduplicated data from the cloud.

The following example illustrates how to calculate cache size based on the size of the data volumes.

Example 1 Calculating cache size

The following figure represents the data and caches in this example:

Example 1 Calculating cache size (continued)

The example uses the following assumptions:

- The front-end cache size is 10% of the volume size
- The deduplication cache is able to contain 90% of the post-deduplicated data
- The deduplication ratio is 4:1

In this example there are four 30 TB volumes with a separate front-end cache configured for each volume. Each front-end cache is 10% of the volume size, so 3 TB each for 12 TB total.

The deduplication cache needs to be able to contain 90% of the post-deduplicated data. At a 4:1 ratio, the 120 TB total data in the volumes will be reduced to 30 TB by the deduplication process ($120 / 4$). To contain 90% of the 30 TB of post-deduplicated data, the deduplication cache needs to be 27 TB.

Deduplication requires 8 GB of RAM plus 1 GB of RAM per terabyte of cache. In this example, the CloudArray virtual machine should have at least 47 GB of RAM ($8 + 12 + 27$).

Expanding existing cache

You can add additional storage to your cache if more space is needed.

Procedure

1. Choose **Cache Management > Caches** from the CloudArray main menu.
2. In the **Cache Name** column, click the name of the cache you want to expand.
The **Cache Details** panel appears.
3. Under **Cache Volumes**, click **Expand Cache**.
The **Expand Cache** panel appears.
4. Enter values equal to or smaller than the available capacity of each cache source you want to use to expand the cache. Optionally, check **Use All Available Capacity**.

Note

You can choose to use less than the total capacity of the cache sources to expand the cache. For example, you may wish to allocate a portion for the existing caches and use the rest to create one or more additional caches.

Deleting caches

You can delete caches from the **Caches** screen.

Procedure

1. Choose **Cache Management > Caches** from the CloudArray main menu.
The **Caches** panel appears.
2. Click the checkbox next to each cache you want to delete.
3. Click the trash barrel **Action** button to delete the caches.

CHAPTER 9

Host iSCSI

Applications can interact with CloudArray via iSCSI. This chapter describes how to configure host-side iSCSI.

Topics include:

• iSCSI clients	62
• Viewing iSCSI client details	62
• Creating an iSCSI client	62
• Deleting iSCSI clients	62
• Configuring iSCSI for VMware vSphere Hypervisor (ESXi)	63
• Configuring iSCSI for Windows	67
• Configuring iSCSI for Linux	69
• Configuring iSCSI for SUSE Linux	69
• Configuring iSCSI for HP-UX	70

iSCSI clients

To view configured iSCSI clients, choose **iSCSI Clients** from the CloudArray main menu. From this screen you can:

- [View details about an iSCSI client](#)
- [Create a new iSCSI client](#)
- [Delete an iSCSI client](#)

You can configure host-side iSCSI for the following site types:

- [VMware vSphere Hypervisor \(ESXi\)](#)
- [Windows](#)
- [CentOS and RHEL Linux](#)
- [SUSE Linux](#)
- [HPUX](#)

Viewing iSCSI client details

You can view detailed information about each iSCSI client.

Procedure

1. Choose **iSCSI Clients** from the CloudArray main menu.
2. In the **Client Name** column, click the name of the client you want to view.

Creating an iSCSI client

To create a new iSCSI client, complete the following steps:

Procedure

1. From the CloudArray main menu, select **iSCSI Clients**, then click **Configure New Client**.

The **Configure New Client** panel appears.

2. Complete the following fields in the panel:

Client Name: Enter the name of the iSCSI client.

IQN: Enter the IQN of each iSCSI client you wish to connect to CloudArray.

Note

CloudArray does not auto-detect iSCSI clients attempting to attach.

3. Click **Configure New Client**.

The new client is created and added to the **Clients** list.

Deleting iSCSI clients

You can delete iSCSI clients from the **Clients** screen.

Procedure

1. Choose **iSCSI Clients** from the CloudArray main menu.
The **Clients** panel appears.
2. Click the checkbox next to each client you want to delete.
3. Click the trash barrel **Action** button to delete the clients.

Configuring iSCSI for VMware vSphere Hypervisor (ESXi)

This section describes how to create a VMware datastore on a CloudArray iSCSI volume on the following VMware platforms:

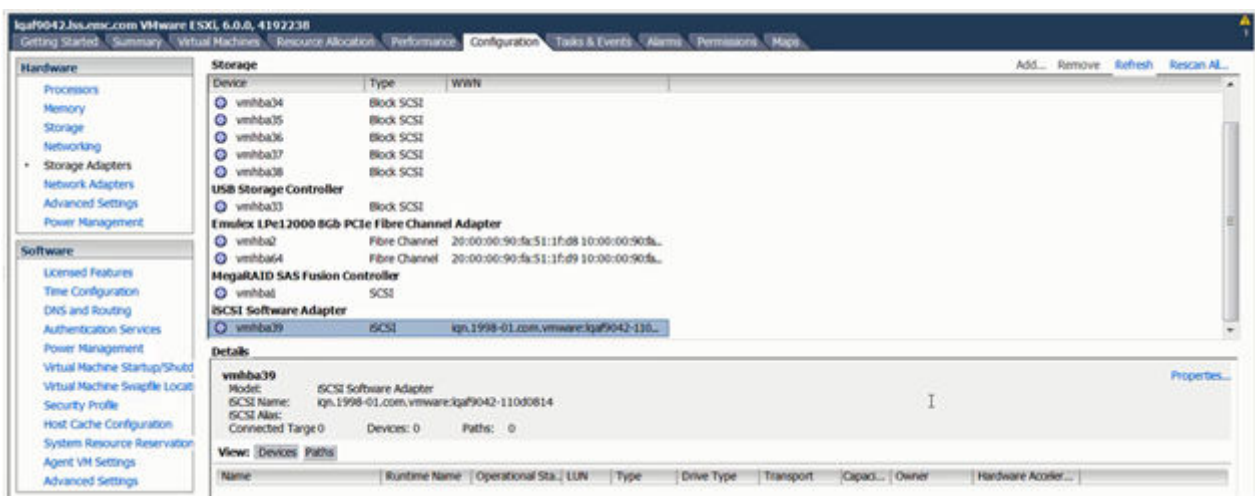
- VMware vSphere Hypervisor 6.5
- VMware vSphere Hypervisor 6.0

This procedure involves these high level phases:

1. Configure the iSCSI client.
2. Create a VMware datastore on the iSCSI volume.

Procedure

1. Install the iSCSI virtual adapter for VMware:
 - a. In the VMware vSphere client, select the VMware ESXi Host into which you are installing the VMware iSCSI Adapter.
 - b. Under Hardware, select **Storage Adapters**.
 - c. Click **Configure** tab.
 - d. Click **Add** and select **Software iSCSI Adapter**.
 - e. Select the iSCSI adapter from the list and click **OK**.
 - f. Obtain the iSCSI Qualified Name for the iSCSI Software Adapter, using the **Properties** control, locating the **iSCSI Name** field for the VMware ESXi Host.



2. Add an iSCSI Client to CloudArray:
 - a. From the CloudArray main menu, select **iSCSI Clients**, then click **Configure New Client**.

The **Configure New Client** panel appears.

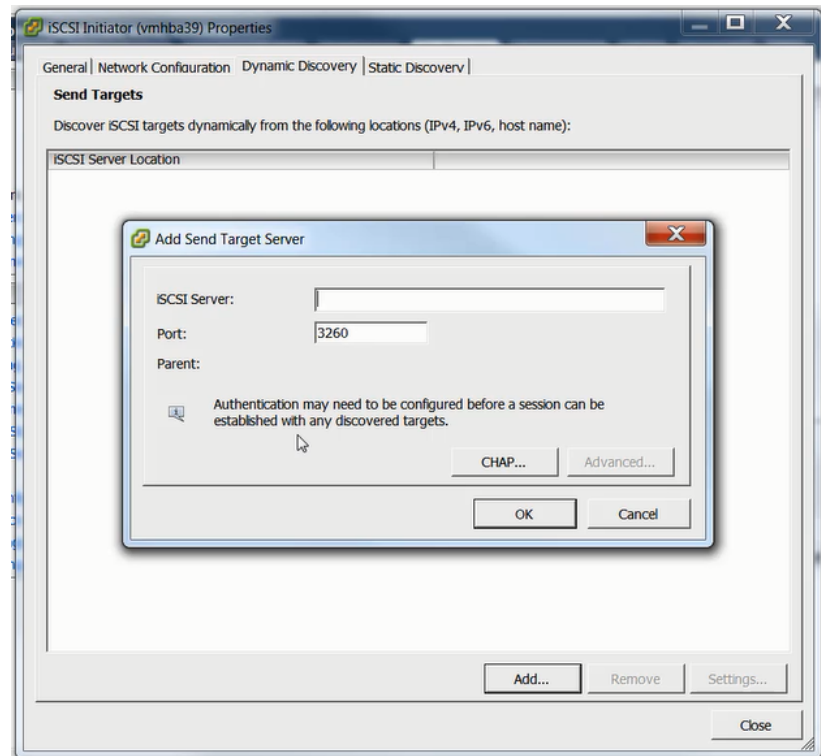
- b. Enter a logical name for the client in **Client Name**.

Note

A best practice is to use the hostname of the client in the **Client Name** field.

- c. Enter the iSCSI Qualified Name for the VMware ESXi Host into **IQN**.
 - d. Click **Configure New Client**.
3. Discover the CloudArray iSCSI Client on VMware:
 - a. In the VMware vSphere client, select the iSCSI Software Adapter.
 - b. Click **Properties**.
 - c. Click the **Dynamic Discovery** control.
 - d. Click **Add**.

The **Add Send Target Server** page appears.



- e. Enter the IP address for the CloudArray (or the FQDN for it) in the **iSCSI Server** field.
- f. Enter 3260 in **Port**.
- g. Click **OK**.
- h. Click **Close**.

Note

The VMware vSphere client may ask you to re-scan the host bus adapter to locate volumes. You have not yet created any iSCSI volumes in CloudArray. Respond No to the query.

4. Map a CloudArray Volume to the iSCSI Client:
 - a. From the CloudArray main menu, select the **VOLUMES** menu.
 - b. Click the **Create New Volume** button.

The **Create New Volume** page appears.

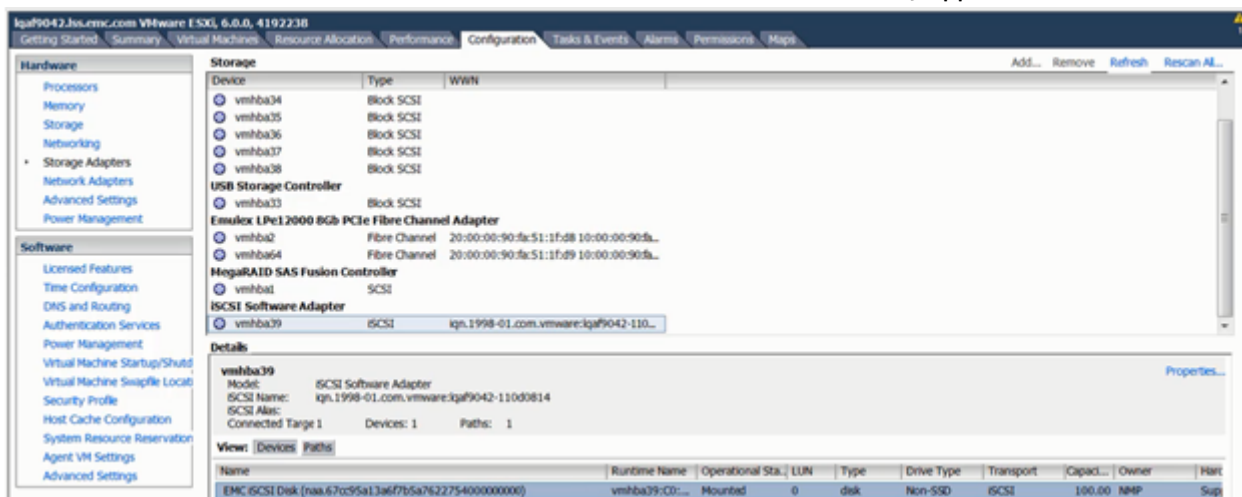
 - c. Set **Volume Name** and **Volume Capacity**.
 - d. Select a Cloud Provider and a Cache for the new Volume.

Note

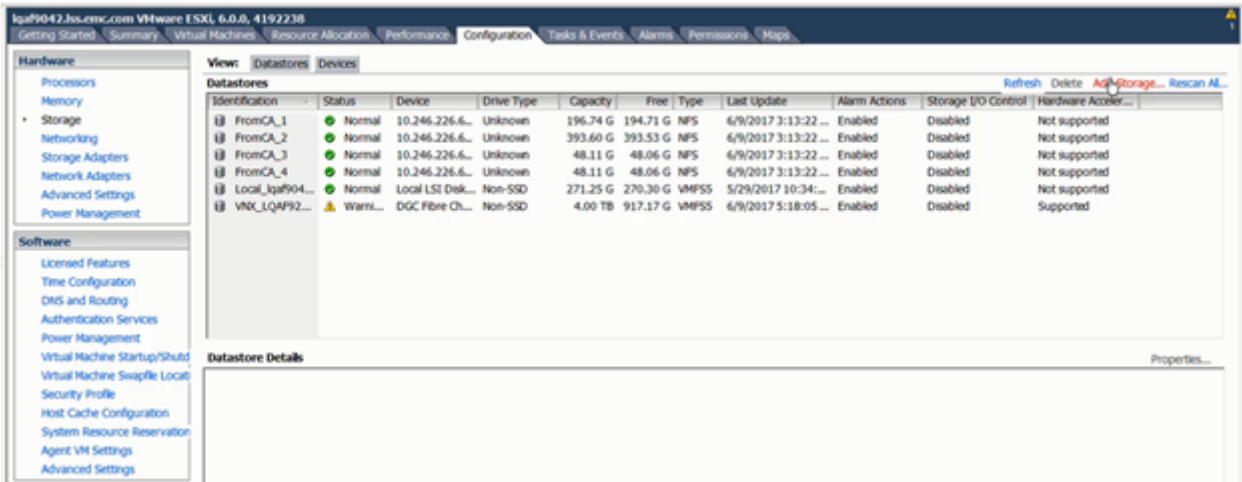
VMware vSphere Hypervisor (ESXi) does not support a cache sector size of 4096 bytes. To create CloudArray iSCSI volumes with a sector size of 512 bytes for ESXi systems, use the default cache page size (1 MiB).

- e. In the section at the bottom under **Select Frontend For This Volume**, click **Map This Volume to An iSCSI Client**.
 - f. In the **Select iSCSI Client** list, choose your iSCSI client for the VMware ESXi Host.
 - g. Click **Create Volume** to create the volume mapped to the host you selected.
5. Re-scan the CloudArray Volume on VMware:
 - a. In the VMware vSphere client, select the **iSCSI Software Adapter**, and
 - b. Click the **Rescan Storage** icon or whatever re-scan control that the client has.

The iSCSI disk volume from the CloudArray appears as a device.



6. Define the CloudArray Volume as a Datastore in VMware:
 - a. In the VMware vSphere client, select the **Storage** control.



b. Click the **Add Storage** control.

The **Add Storage** dialog appears.

c. For **Storage Type**, select **Disk/LUN** (or **VMFS** if that is the option in the VMware vSphere client), then click **Next**.

d. For **Select Disk/LUN**, select the iSCSI disk that the VMware vSphere client discovered on the Cloud Array, then click **Next**.

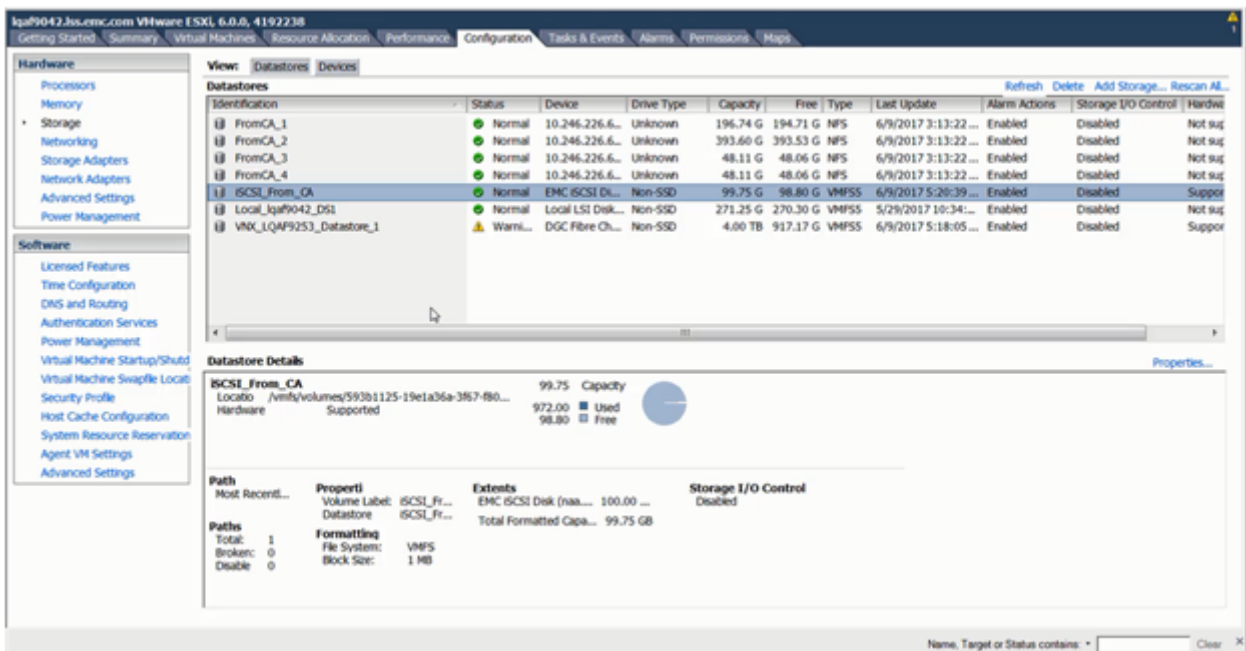
e. Accept the **Current Disk Layout**, then click **Next**.

f. For **Properties**, assign a datastore name, then click **Next**.

g. Specify the maximum file size and capacity for this datastore, then click **Next**.

h. Review the disk layout and then click **Finish**.

The VMware vSphere client uses the settings to create the datastore and displays it among the available datastores in the host.



Configuring iSCSI for Windows

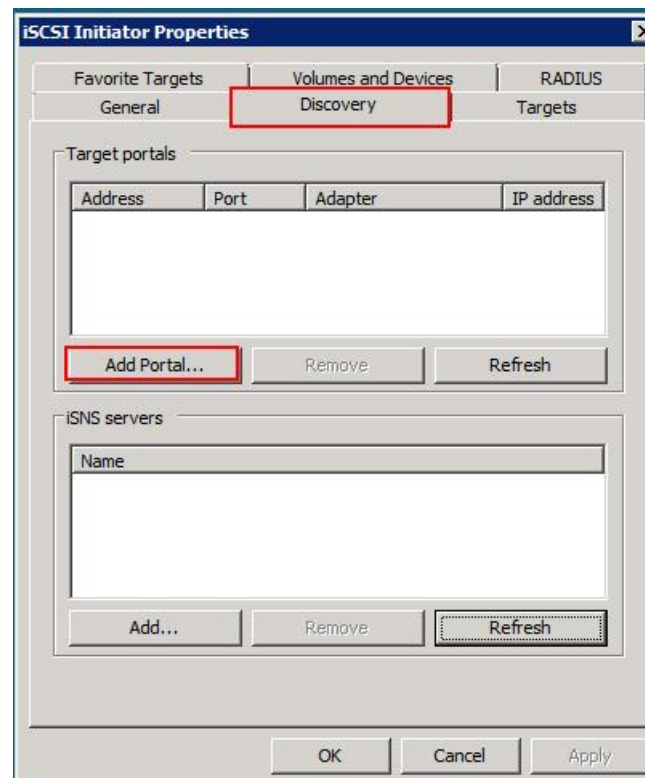
You can configure host-side iSCSI for the following Windows platforms:

- Windows Server 2012 R2
- Windows Server 2012
- Windows Server 2008 R2
- Windows Server 2008
- Windows 10
- Windows 8
- Windows 7

Procedure

1. On the Windows host, open the **Control Panel**.
2. Select **Administrative Tools > iSCSI Initiator**.

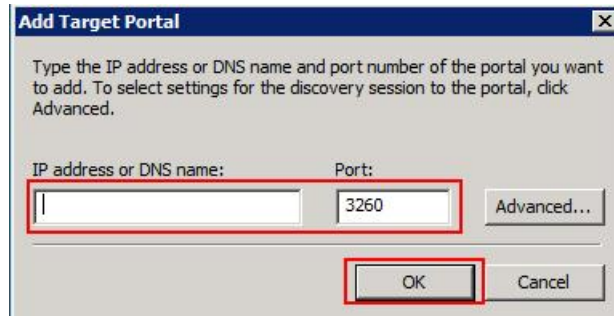
The following panel appears. Note that your panel may vary depending on the version of Windows you are using.



Note

If there is no iSCSI initiator menu item in your Control Panel, download a free copy from [Microsoft](#). Download either the 32 or 64-bit version based on the version of Windows you are using and run the installer, accepting the default settings.

3. Select the **Discovery** tab and then click **Add Portal**.
The following panel appears:



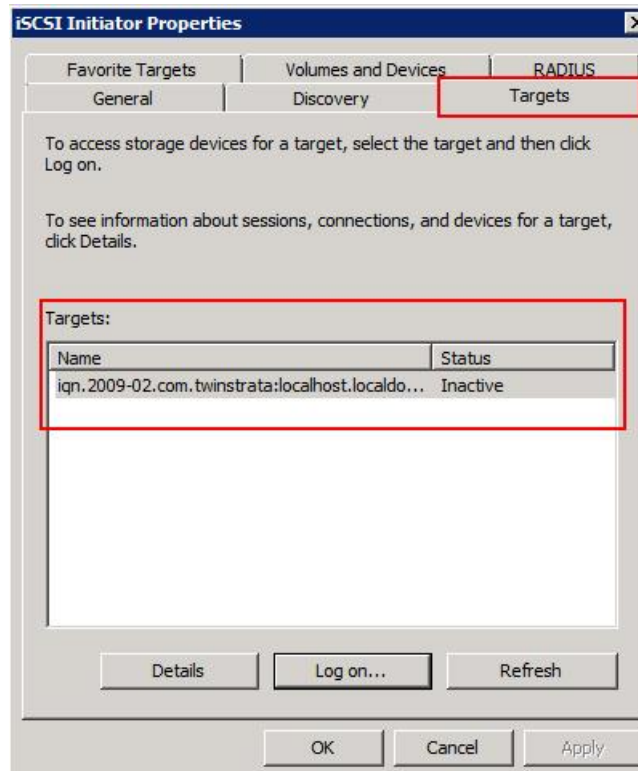
Add Target Portal

Type the IP address or DNS name and port number of the portal you want to add. To select settings for the discovery session to the portal, click Advanced.

IP address or DNS name: Port:

Advanced... OK Cancel

4. Enter the IP address of the VM instance running CloudArray. Leave the port as default (3260) and click **OK**.
5. Select the **Targets** tab.
The iSCSI target that you added in the previous step is listed as Inactive.



iSCSI Initiator Properties

Favorite Targets | Volumes and Devices | **RADIUS**

General | Discovery | **Targets**

To access storage devices for a target, select the target and then click Log on.

To see information about sessions, connections, and devices for a target, click Details.

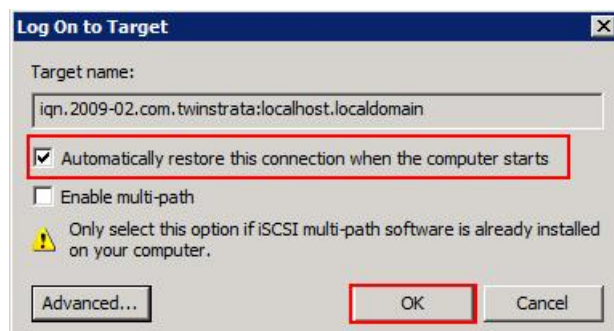
Targets:

Name	Status
iqn.2009-02.com.twinstrata:localhost.localdo...	Inactive

Details Log on... Refresh

OK Cancel Apply

6. Click **Log on**.
The following panel displays.



Log On to Target

Target name:

☒ Automatically restore this connection when the computer starts

☐ Enable multi-path

Only select this option if iSCSI multi-path software is already installed on your computer.

Advanced... OK Cancel

7. Check the option for **Automatically restore ...** and click **OK**.

The target status changes from Inactive to Connected. Host-side iSCSI configuration is now complete.

Configuring iSCSI for Linux

Complete the following steps to configure host-side iSCSI for the following Linux platforms:

- CentOS (Community Enterprise Operating System) version 7, 6
- RHEL (Red Hat Enterprise Linux)

Procedure

1. Log in to the CentOS or RHEL host with root privileges.
2. To begin configuring the Linux initiator, install the `iscsi-initiator-utils` packages (if not already installed) by typing:

```
yum install iscsi-initiator-utils
```

This starts the package download.

3. Once the package has been downloaded and installed, start the iSCSI daemon by typing:

```
service iscsid start
```

4. Use the `iscsiadm` utility to discover and login into CloudArray, by typing:

```
iscsiadm -m discovery -t sendtargets -p x.x.x.x
```

Where x.x.x.x is the IP address of your CloudArray.

5. Provision a volume to this host using the CloudArray main menu.
6. Once a volume has been provisioned, re-scan for the volume by typing:

```
service iscsi restart
```

7. Run:

```
fdisk -l
```

This output should list a new volume on which you can:

- Create a partition, using `fdisk` command
- Format, using `mkfs` command
- Mount, using the `mount` command

Configuring iSCSI for SUSE Linux

Complete the following steps to configure host-side iSCSI for SUSE Linux sites.

Procedure

1. Log in to the SUSE host with root privileges.
2. To install iSCSI initiator, type the following commands:

```
$ zypper install open-iscsi
```

```
$ rcopen-iscsi start
```

```
$ iscsiadm -m discovery -t sendtargets -p <IP address of CloudArray>
```

```
$ iscsiadm -m discovery -t sendtargets -p <IP address of CloudArray>
```

3. Provision a volume to this host using the CloudArray main menu.
4. Run the following command to re-scan for the volume(s) mapped:

```
rcopen-iscsi restart
```

Configuring iSCSI for HP-UX

Before you begin

This procedure assumes the HP-UX iSCSI initiator software is already installed.

Complete the following steps to configure host-side iSCSI for HP-UX sites.

Procedure

1. Log in to the HP-UX host with root privileges.
2. After you have downloaded the `iscsi-00_B.11.31.01_HP-UX_B.11.31_IA+PA.depot` file, use the `mv` command to move it to the `/tmp` directory on your system.
3. On a stand-alone system, run the following command to install the product:

```
# swinstall -x autoreboot=true -s
```

```
/tmp/iscsi-00_B.11.31.01_HP-UX_B.11.31_IA+PA.depot iscsi-00
```

4. Add the path for the `iscsiutil` executable program and other iSCSI executable programs to the root path.

```
# PATH=$PATH:/opt/iscsi/bin
```

Note

Add the previous string to the `/.profile` file to avoid manually updating the `PATH` environment variable each time.

5. Add a discovery target with the following command.

```
# iscsiutil -a -I <IP address of CloudArray>
```

Note

The HP-UX iSCSI software initiator does not support IPv6 addresses. Do not configure IPv6 addresses as a target IP.

6. Provision a volume to this host using the CloudArray main menu.
7. Check the disk information.

```
diskinfo -v <disk>
```

CHAPTER 10

Snapshots

CloudArray provides space-efficient snapshots and advanced scheduling and retention functionality to protect data in the cloud.

Topics include

- [Snapshots](#)..... 72
- [Creating a snapshot](#)..... 72
- [Scheduling a snapshot](#)..... 72
- [Enabling and disabling snapshot scheduling](#)..... 74
- [Exposing a snapshot](#)..... 75
- [Unexposing a snapshot](#)..... 75
- [Mapping a snapshot to a host](#)..... 76
- [Deleting a snapshot](#)..... 77

Snapshots

CloudArray's snapshot feature allows you to reduce or even eliminate dependence on separate backup software. CloudArray snapshots:

- Provide capacity savings over traditional full backups
- Reduce space requirements and storage costs
- Provide continuous, offsite data protection in the cloud

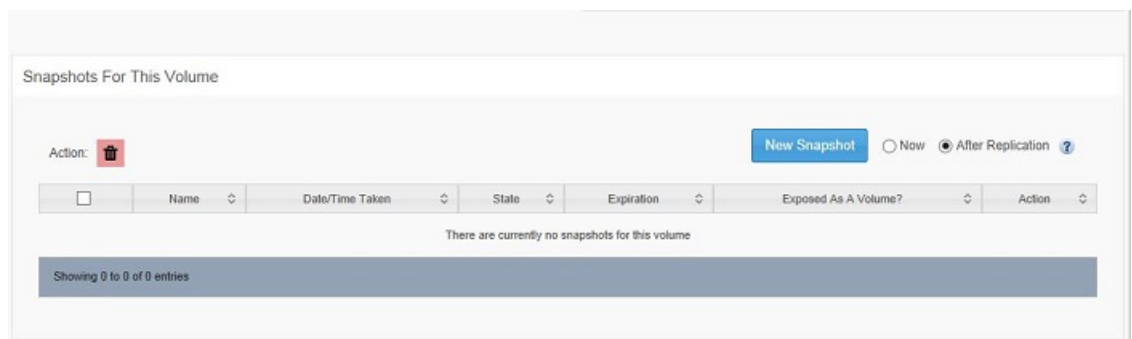
Each snapshot creates a reference to the data in the cloud for a volume. When first taken, the snapshot merely references the data in the cloud (it is not a copy of the data). Over time, as the data in the volume changes, the snapshot and volume diverge so that they no longer represent the same physical data in the cloud.

CloudArray supports snapshots with retention policies, enabling you to predetermine how often to take and how long to keep scheduled snapshots. When the retention policy for a particular snapshot expires, any data referenced in the snapshot that is not otherwise used (either by other snapshots or by current data volumes and shares) is removed from the cloud.

Creating a snapshot

Procedure

1. From the CloudArray main menu, select **Volumes** and then click on the **Volume Name** for which you want to create a snapshot.
2. In the **Snapshots For This Volume** portion of the panel, select either **Now** to take the snapshot immediately or **After Replication** to take the snapshot after the cache for volume has been replicated completely to the cloud.



3. Click **New Snapshot**.

After a snapshot has been taken, it is listed in the volume's snapshot details section. The snapshot name is a timestamp that indicates when the data was replicated to the cloud. If you have multiple snapshots, click the column headings to determine the sorting order.

Scheduling a snapshot

You can schedule snapshots in advance for large file servers or application data that would otherwise require days to back up. You can also establish age-based retention policies for the snapshots.

Procedure

1. From the CloudArray main menu, select **Volumes** and then click on the **Volume Name** for which you want to schedule a snapshot.
2. Click **Show Snapshot Schedule** in the **Snapshot Scheduler** section of the panel.

Snapshots For This Volume

Action: [New Snapshot](#) ☒ Now ☐ After Replication [?](#)

<input type="checkbox"/>	Name	Date/Time Taken	State	Expiration	Exposed As A Volume?	Action
<input checked="" type="checkbox"/>	20141112T105207.025446	2014-Nov-12 10:52:07	operating	Never	NO	EXPOSE

Showing 1 to 1 of 1 entries

Snapshot Scheduler [Show Snapshot Schedule](#) Snapshot Scheduler is: **Enabled**

The Snapshot Scheduler window opens.

3. Double-click the time you want to start the snapshot.

How To Use Double-click the start time on the calendar to create a scheduled snapshot. Click existing events to modify/view
Times in CloudArray timezone EST (GMT-5)

Day Week Month Year **Snapshot Scheduler** Today

	Sun, November 9	Mon, November 10	Tue, November 11	Wed, November 12	Thu, November 13	Fri, November 14	Sat, November 15
00:00							
01:00							
02:00							
03:00							
04:00							
05:00							
06:00							
07:00							
08:00							
09:00							
10:00							
11:00							
12:00							
13:00							
14:00							
15:00							
16:00							
17:00							
18:00							
19:00							
20:00							
21:00							
22:00							
23:00							

The Snapshot information panel appears.

4. Fill in the details and click **Save** to schedule the snapshot.

Snapshot of Volume1

Snapshot Description

Snapshot of Volume1

When To Take The Snapshot

☒ Now
 ☐ After All Cache Data Has Been Replicated To The Cloud

Retention Policy

Retention Policy: 1 Month

Repeat Frequency

☒ Weekly
 ☐ Monthly
 ☐ Yearly

Repeat every 1 week(s) on the following days:
☐ Sunday
 ☐ Monday
 ☐ Tuesday
 ☒ Wednesday
 ☐ Thursday
 ☐ Friday
 ☐ Saturday

Snapshot Time

04:15 12 November 2014

Save

Cancel

Delete

The snapshot calendar updates to indicate the new snapshot.


Note

You can edit the snapshot by double-clicking the snapshot name in the calendar.

Enabling and disabling snapshot scheduling

Procedure

1. To enable or disable the snapshot scheduler, select **Administration > Settings** from the CloudArray main menu.
2. Under **Snapshot Scheduler**, click the icon to toggle the setting.



Exposing a snapshot

To access a snapshot, use the **Expose** action to mount it and present it as a new volume. Exposing a snapshot allows you to use the snapshot as an ordinary volume or share. In this way, you can use data from snapshots for detailed analysis without impacting production data.

Procedure

1. From the CloudArray main menu, select **Volumes** and then click on the **Volume Name** for which you want to expose a snapshot.
2. In the **Snapshots For This Volume** portion of the panel, select a snapshot by selecting the checkbox to the left of the snapshot name.
3. In the **Action** column, click **Expose**.
4. Select the cache to be associated with the exposed snapshot. Check **Fast Cache Reload** to pre-populate the cache with cloud data.
5. Click **OK**.

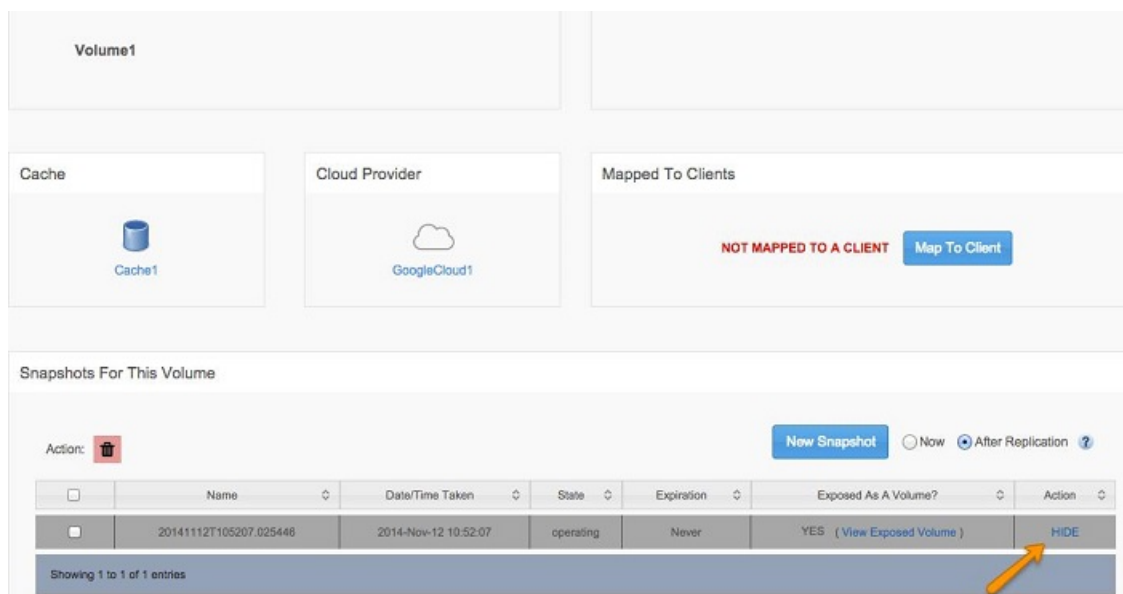
The snapshot is exposed.

Unexposing a snapshot

Use **HIDE** to hide the volume and unexpose a snapshot.

Procedure

1. Unmap the snapshot volume from its respective client.
See [Unmap a volume](#) for details.
2. From the CloudArray main menu, select **Volumes** and then click on the **Volume Name** for which you want to unexpose a snapshot.
3. In the **Snapshots For This Volume** portion of the panel, click **HIDE** to unexpose the snapshot.



The screenshot shows the CloudArray interface for a volume named 'Volume1'. The 'Cache' section shows 'Cache1'. The 'Cloud Provider' section shows 'GoogleCloud1'. The 'Mapped To Clients' section shows 'NOT MAPPED TO A CLIENT' with a 'Map To Client' button. The 'Snapshots For This Volume' section shows a table with one snapshot. The 'Expose' button is highlighted with an orange arrow.

	Name	Date/Time Taken	State	Expiration	Exposed As A Volume?	Action
<input type="checkbox"/>	20141112T105207.025446	2014-Nov-12 10:52:07	operating	Never	YES (View Exposed Volume)	EXPOSE

Showing 1 to 1 of 1 entries

Mapping a snapshot to a host

Procedure

1. From the CloudArray main menu, select **Volumes** and then click on the **Volume Name** for which you want to map a snapshot to a host.
2. In the **Snapshots For This Volume** portion of the panel, click the **View Exposed Volume** link to open the snapshot details page.

The screenshot shows the 'Snapshots For This Volume' section of the CloudArray interface. At the top, there are three panels: 'Cache' (Cache1), 'Cloud Provider' (GoogleCloud1), and 'Mapped To Clients' (NOT MAPPED TO A CLIENT, Map To Client). Below these is a table of snapshots. The table has columns: Action, Name, Date/Time Taken, State, Expiration, Exposed As A Volume?, and Action. The first row shows a snapshot named '20141112T105207.025446' with state 'operating' and 'Exposed As A Volume?' set to 'YES (View Exposed Volume)'. An orange arrow points to the 'View Exposed Volume' link in the 'Exposed As A Volume?' column.

Action	Name	Date/Time Taken	State	Expiration	Exposed As A Volume?	Action
<input checked="" type="checkbox"/>	20141112T105207.025446	2014-Nov-12 10:52:07	operating	Never	YES (View Exposed Volume)	HIDE

Showing 1 to 1 of 1 entries

3. Click the **Map to Client** button to map the snapshot to a host.

The screenshot shows the 'Map to Client' button highlighted with an orange arrow. The interface includes a top navigation bar with 'CloudArray', 'Cloud Providers', 'Cache Status', and 'Volume Status'. Below this is a breadcrumb trail 'Volumes > Volume1,20141112T105207.025446'. The main area has four buttons: 'Hide Snapshot', 'Map To Client', 'Disable Replication', and 'Cache Migration'. Below these are two panels: 'Volume Details' and 'Cache Details'. The 'Volume Details' panel shows 'Snapshot of Volume1', 'State: Operating', 'Health: Online', 'Capacity: 100 GiB', and 'Policy: Policy1'. The 'Cache Details' panel shows 'Cache Name: Cache1', 'Total Cache Capacity: 100 GiB', 'Dirty In Cache: 0 GiB', and 'Last Completed Replication: N/A'. At the bottom, there are three panels: 'Cache' (Cache1), 'Cloud Provider' (GoogleCloud1), and 'Mapped To Clients' (NOT MAPPED TO A CLIENT, Map To Client). An orange arrow points to the 'Map To Client' button.

Volume Details

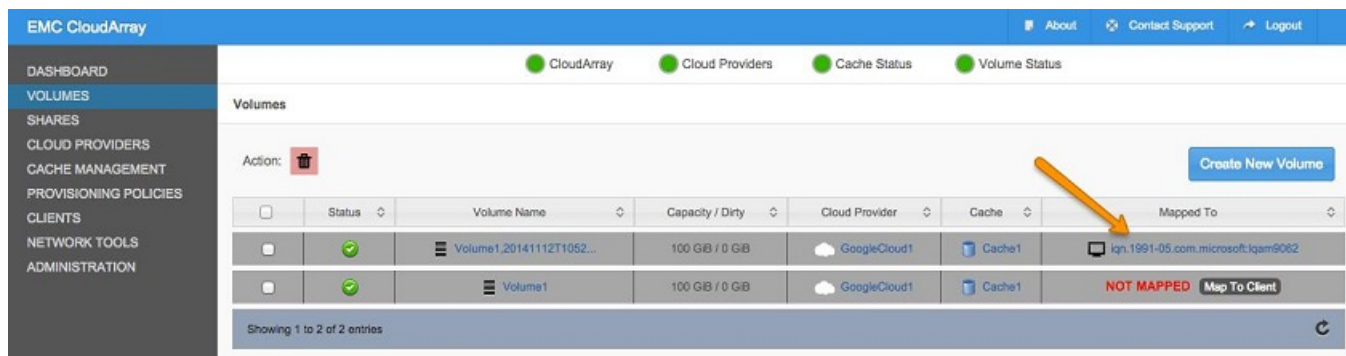
Snapshot of Volume1
 State: Operating
 Health: Online
 Capacity: 100 GiB
 Policy: Policy1

Cache Details

Cache Name: Cache1
 Total Cache Capacity: 100 GiB
 Dirty In Cache: 0 GiB
 Last Completed Replication: N/A

4. Choose the host from the drop-down list and click **Map Volume**.

- The exposed snapshot appears under the **Volumes** screen. Mouse over the volume name to see the complete snapshot name.

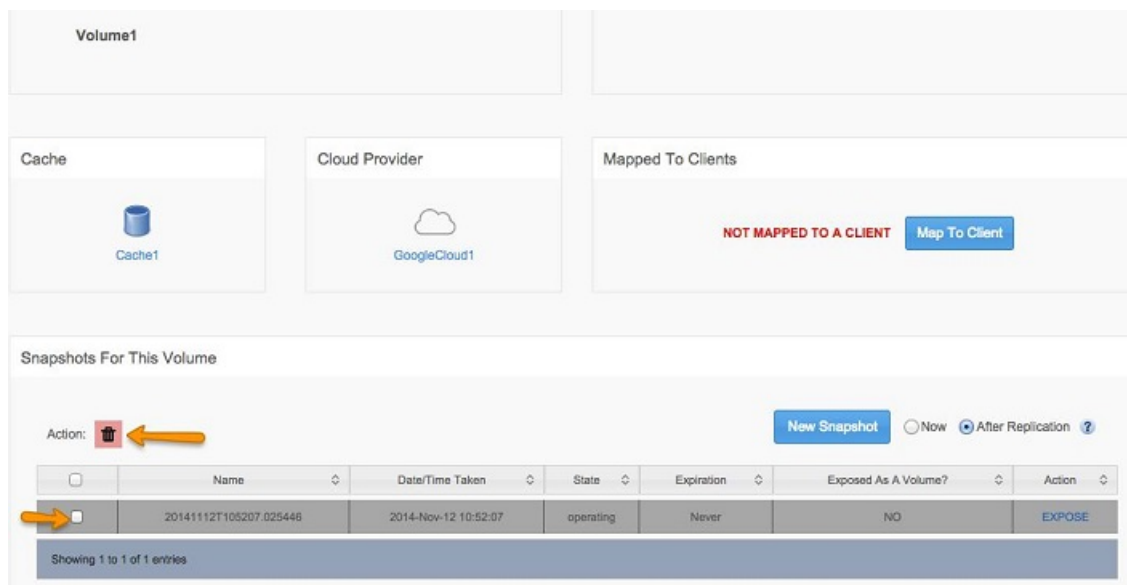


To access the volume from the client, re-scan for it with the host's volume management utility.

Deleting a snapshot

Procedure

- From the CloudArray main menu, select **Volumes** and then click on the **Volume Name** for which you want to delete a snapshot.
- Select one or more snapshots from the **Snapshots For This Volume** section of the panel.



- Click the **Trash** icon.

CHAPTER 11

Network Tools

This chapter describes CloudArray's network tools.

Topics include:

- [Throttling network bandwidth](#).....80
- [Disabling the Bandwidth Throttler](#).....82
- [Configuring Cloud Performance Optimizer](#).....82
- [Configuring the default gateway](#).....83
- [Configuring DNS servers](#).....83
- [Configuring the hostname](#).....83
- [Modifying network adapter settings](#).....84
- [Configuring NTP servers](#).....84
- [Verifying network settings](#).....85
- [Enabling a network proxy](#).....85

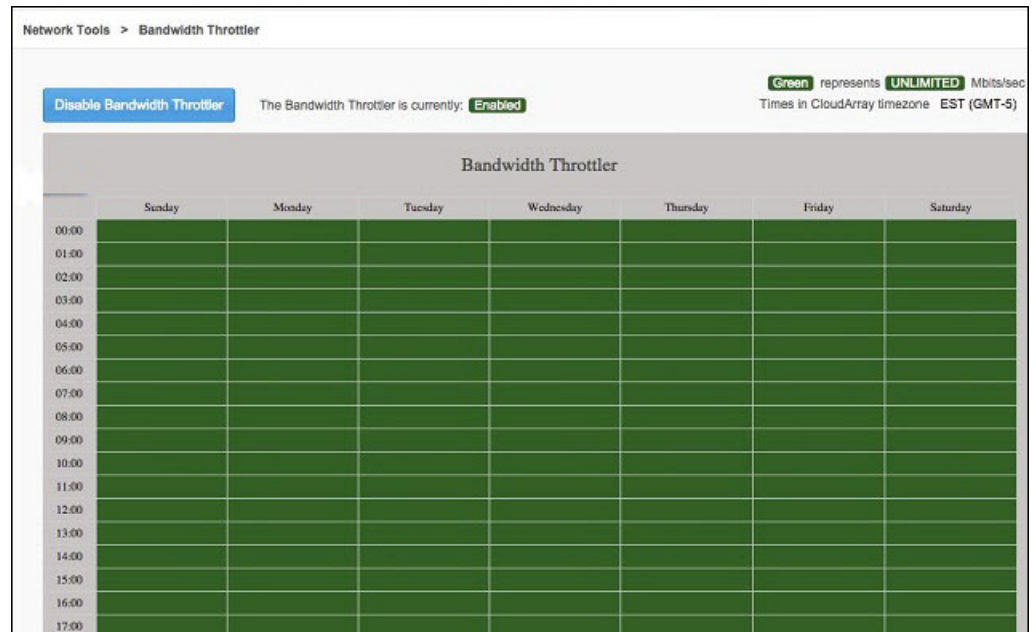
Throttling network bandwidth

The Bandwidth Throttler provides throttling and scheduling capability to regulate the replication of data from CloudArray to your cloud storage.

Procedure

1. Select **Network Tools > Bandwidth Throttler** from the CloudArray main menu.

A panel appears with a weekly calendar that has hourly slots. The green color indicates unlimited bandwidth is being consumed during those time ranges.



2. Click on a day and drag the mouse vertically to select a time range to which you intend to apply a throttle. During this process a blue box indicating the time slot and default bandwidth appears. Release the mouse.

A panel appears for the selected time range.

Bandwidth Settings

Maximum Allowed Bandwidth (Mbit/s)

125 Mbit/s

2 250 500 750 1000

Use left and right arrow keys on keyboard for precision selection

Schedule Repeats Weekly

Apply bandwidth settings to the following days:

☐ Sunday ☒ Monday ☒ Tuesday ☒ Wednesday ☒ Thursday ☒ Friday ☐ Saturday

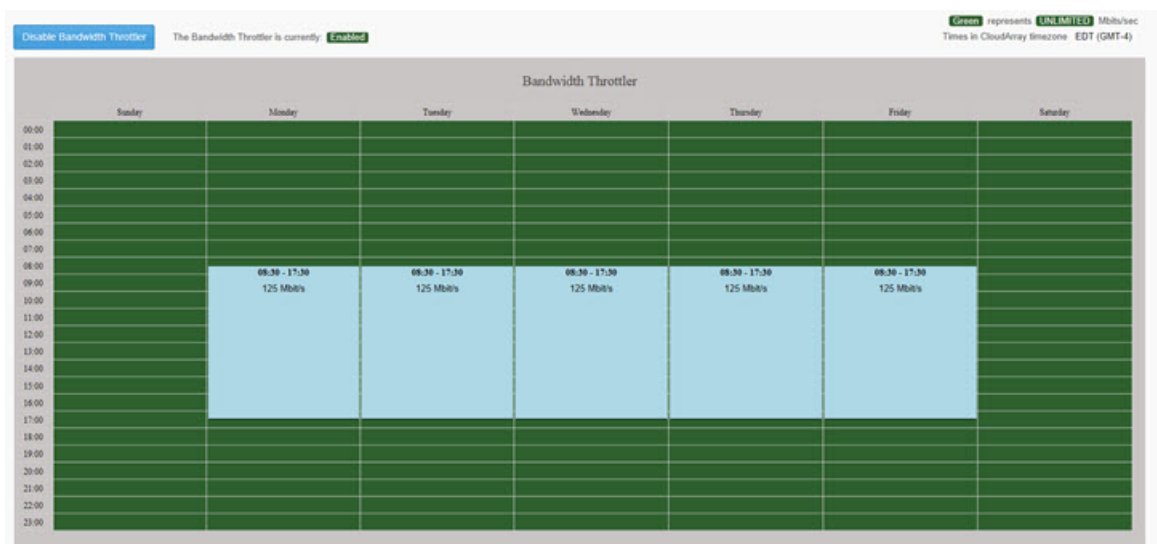
Time Range

08:30 – 17:30

Save **Cancel** **Delete**

- Set the bandwidth. The tool uses a slider to determine the allocated bandwidth. The default setting is 1000 Mbit/sec. Move the slider to the right to adjust the bandwidth.
- Select the days that the bandwidth will be in effect.
- Select the time range for the bandwidth.
- Click **Save** to apply the throttle.


Verify your settings. Bandwidth will be throttled during the indicated time ranges. You can apply additional throttling to other time ranges. During unmarked (green) time ranges, CloudArray uses the maximum available bandwidth for data transfer.



Disabling the Bandwidth Throttler

Procedure

1. Do one of the following to disable the Bandwidth Throttler:
 - From the CloudArray main menu, select **Network Tools > Bandwidth Throttler**, then click **Disable Bandwidth Throttler** at the top of the **Bandwidth Throttler** panel.
 - Select **Administration > Settings** from the CloudArray main menu. Under **Bandwidth Throttler**, click the icon to toggle the setting.



Configuring Cloud Performance Optimizer

The Cloud Performance Optimizer allows you to best utilize the bandwidth available to CloudArray. This is a global network setting that is not specific to any particular cloud provider.

Procedure

1. From the CloudArray main menu, select **Network Tools > Cloud Performance Optimizer**.

Network Tools > Cloud Performance Optimizer

Use this feature to configure CloudArray to best utilize the typical bandwidth available to CloudArray to replicate data to cloud storage.

This is a global setting. This is not a per cloud provider setting

☐ Less Than 5 Mbit/sec (default)
 ☐ Between 6 and 10 Mbit/sec
 ☐ Between 11 and 25 Mbit/sec
 ☒ Between 26 and 50 Mbit/sec
 ☐ Between 51 and 100 Mbit/sec
 ☐ Greater Than 100 Mbit/sec

Cancel
Save Changes

2. Select the uplink bandwidth available from the CloudArray to the cloud storage.

Note

This setting is directly correlated to the Bandwidth Throttler. If you throttle down the bandwidth, select a lower uplink bandwidth setting and vice versa. This will avoid congestion on the uplink.

Configuring the default gateway

Procedure

1. From the CloudArray main menu, select **Networks Tools > Default Gateway Configuration**.
2. To use DHCP, check **Use DHCP to determine gateway**.
3. Under **Default Network Interface Card**, select the interface to use from the drop-down list. You can also select **No Network Interface Selected**.
4. If you are not using DHCP, enter the IP address in the **Default Gateway IP Address** field.
5. Click **Save Changes** to apply the settings.

Configuring DNS servers

By default, CloudArray attempts to obtain the list of DNS servers from your DHCP server. Alternatively, you can specify up to three static DNS servers to be used by CloudArray.

Procedure

1. From the CloudArray main menu, select **Network Tools > DNS Configuration**.

2. To specify static DNS servers, select the **Specify My DNS Server** button.
3. Specify up to three static DNS servers to be used by CloudArray and click **Save**.

Configuring the hostname

You can customize the CloudArray hostname.

Procedure

1. From the CloudArray main menu, select **Networks Tools > Hostname Configuration**.

2. Enter a new hostname and confirm the name.
3. Click **Change**.

Modifying network adapter settings

Procedure

1. From the CloudArray main menu, select **Networks Tools > Interface Management**.

The **Interface Management** panel appears.

2. In the **Interface** column, click the name of an interface to modify it.

The **Network Interface Card Details** panel appears.

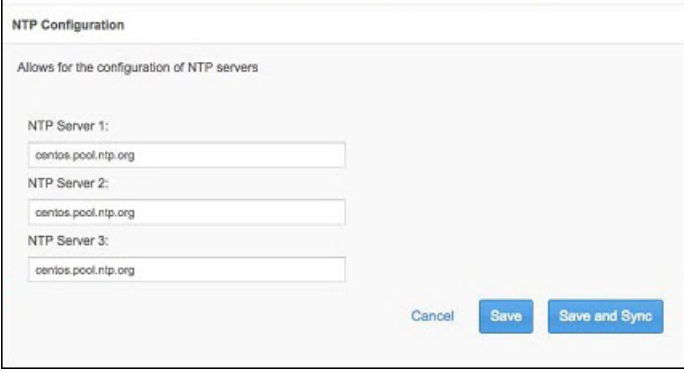
3. To set a static IP address and route, select **Edit** from the **IP Address** option.
4. Select the **Static** button.
5. Enter the **IP Address** and **Netmask** information. If an exit point is needed through your gateway to another subnet, then enter the static route information under the **Static Route Details** section.
6. To create a team, select **Edit** from the **Team** option. Check the **New Team** box, then select **activebackup** or **LACP** for the network adapter teaming runner.
7. To change the maximum transmission unit (MTU) size, select **Edit** from the **MTU** option.
 - a. Choose a pre-defined value or enter a custom value for your network.
8. When finished making changes to the screen, click **Save Changes** to apply them.

Configuring NTP servers

NTP configuration allows you to synchronize the CloudArray clock with an NTP server. You can specify up to three NTP servers.

Procedure

1. From the CloudArray main menu, select **Networks Tools > NTP Configuration**.



NTP Configuration

Allows for the configuration of NTP servers

NTP Server 1:
centos.pool.ntp.org

NTP Server 2:
centos.pool.ntp.org

NTP Server 3:
centos.pool.ntp.org

Cancel Save Save and Sync

2. Enter up to three NTP servers.
3. When finished, click **Save** or **Save and Sync**.

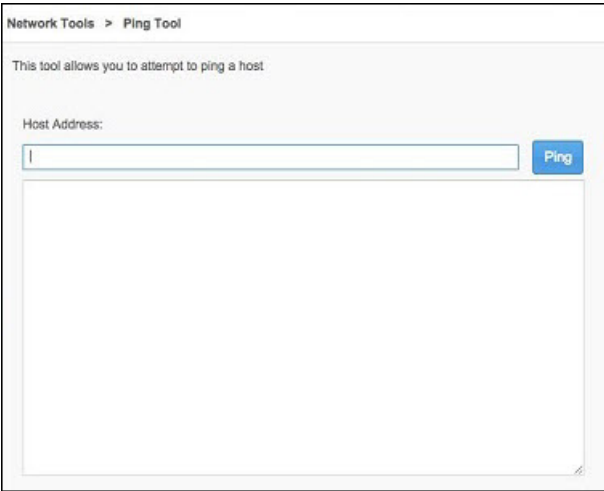
When you click **Save and Sync**, CloudArray simultaneously saves the configuration and synchronizes to the first NTP server.

Verifying network settings

The ping tool allows you to verify network settings by attempting to ping a host.

Procedure

1. From the CloudArray main menu, select **Networks Tools > Ping Tool**.



Network Tools > Ping Tool

This tool allows you to attempt to ping a host.

Host Address:
|

Ping

2. Enter the host address and click **Ping**.

The panel displays connectivity statistics.

Enabling a network proxy

Procedure

1. From the CloudArray main menu, select **Networks Tools > Proxy Management**.

Network Tools > Proxy Management

If a Web proxy is required for CloudArray to connect to servers via http(s), please enter the proxy details below.

☐ A Web proxy Is Required

Proxy:

Port:

☐ Requires Authentication

[Cancel](#) [Save Proxy Settings](#)

2. Check the **A Web proxy is Required** checkbox and enter the proxy name and port.

Note

A default port value of 3128 is pre-populated; you can change this value as needed.

3. If your proxy requires authentication, check the **Requires Authentication** checkbox and enter the username and password.
4. Click **Save Proxy Settings** to apply the changes.

CHAPTER 12

Administration

This chapter describes CloudArray's administration functions.

Topics include:

- [Backing up CloudArray](#)..... 88
- [Changing your CloudArray password](#)..... 88
- [Updating CloudArray software](#)..... 88
- [CloudArray portal settings](#)..... 89
- [Disaster recovery testing](#)..... 91
- [Provisioning policies](#)..... 94
- [Restoring CloudArray configurations](#)..... 95
- [Settings](#)..... 97
- [Managing SSL certificates](#)..... 97
- [Collecting support data](#)..... 98
- [Changing the time zone](#)..... 99
- [User Management](#)..... 99
- [Installing a custom SSL certificate](#)..... 101
- [Updating the CloudArray license](#)..... 102
- [Utilities](#)..... 102

Backing up CloudArray

CloudArray metadata is automatically and securely backed up to the CloudArray portal. Backups occur when an hour has elapsed since the last configuration change. In the event of a disaster, you can log into the CloudArray portal with your unique credentials and have immediate access to your metadata. Using this metadata, you can then instantly restore access to your data on a new appliance. You can also manually back up your CloudArray metadata and save it locally. Complete the following steps to perform a manual backup.

Procedure

1. Select **Administration** > **Backup CloudArray** from the CloudArray main menu.
2. Click **Run Backup**.
3. When the backup completes, save the file locally.

Changing your CloudArray password

Procedure

1. Choose **Administration** > **Change CloudArray Password** from the CloudArray main menu
2. Enter your current password and then enter and confirm a new password.
3. Click **Change Password**.

Updating CloudArray software

When you log in to CloudArray, the appliance checks whether a downloaded software update is present. If an update is available, you have the option to apply it.

Note


If you are not connected to the CloudArray portal, see [Update your CloudArray \(no portal connection\)](#) for update instructions.

Procedure

1. If an update is available, a notification displays to alert you.



2. To manually check for an update, select **Administration** from the CloudArray main menu and click the **Check for Updates** icon.



Note


The process of downloading and validating the update can take several minutes, depending on network speed and workload on the appliance.

Updating CloudArray without a portal connection

If you are not connected to the CloudArray portal, you can update your software manually.

Procedure

1. From a PC with internet access, log in to your CloudArray portal account at cloudarray.com.
2. Click **Downloads & Documents**.




3. From the **CloudArray Upgrade ISO** section, under **Actions**, select **Download** to download the software update file.
4. After the download completes, select **Administration** from the CloudArray main menu.

Note

Make sure the ISO image is accessible to the CloudArray. You can mount the ISO image from the BMC web interface or from the hypervisor's interface as a virtual device.

5. Click the **Check for Updates** icon.



Note

The process of downloading and validating the update can take several minutes, depending on network speed and workload on the appliance.

CloudArray portal settings

The CloudArray portal provides the following:

- Automatic license installation

- Configuration backups saved to the portal
- Alerts for CloudArray issues
- Automatic software updates

From the **CloudArray Portal Management** panel you can [enable](#) or [disable](#) the CloudArray connection to the portal.

Disabling the CloudArray portal

The CloudArray portal provides the following:

- Automatic license installation
- Configuration backups saved to the portal
- Alerts for CloudArray issues
- Automatic software updates

There may be times where security policies may not allow for portal operation and you need to disable it. To use CloudArray locally, you need to download your license from the portal and then upload it manually.

Note

[Enable the CloudArray portal](#) describes how to reconnect to the portal.

Procedure

1. Choose **Administration > CloudArray Portal Settings** from the CloudArray main menu.
2. Deselect the **Enable CloudArray Portal** checkbox.
The panel updates to indicate CloudArray will not use the portal.
3. Select and copy the CloudArray serial number displayed on the panel.
You will need this information in the next step.
4. Click the **CloudArray Portal** link to download your license.
 - a. Log in to your CloudArray portal account.
 - b. Click the **My Licenses** button.
 - c. In the **License Details** panel, click the **Details** link.
 - d. Click **Download License File**.
 - e. Paste the CloudArray serial number into the text box and click **Download**.
 - f. Click **Save** to save the file locally.
5. To upload your license file, select the **Click To Select CloudArray License File** link.
6. Locate your downloaded license file and click **Open**.
The license file info displays in the panel.
7. Click **Save Changes**.

Enabling the CloudArray portal

The CloudArray portal provides the following:

- Automatic license installation
- Configuration backups saved to the portal
- Alerts for CloudArray issues
- Automatic software updates

Complete the following steps if you are disconnected from the portal and want to enable CloudArray portal operations.

Procedure

1. Choose **Administration > CloudArray Portal Settings** from the CloudArray main menu.
2. Select the **Enable CloudArray Portal** checkbox.
3. Enter the your CloudArray username, password, and confirm the password, then click **Save Changes**.
4. If you changed your password on the CloudArray portal, also enter and confirm it in the **Update CloudArray Portal Details** panel.

Disaster recovery testing

With multiple CloudArray systems, you can create and launch a disaster recovery (DR) test scenario. You need to have two CloudArray systems available and running to perform the test:

- The [Primary CloudArray system](#) is where the DR test master backup file is created.
- The [DR Test CloudArray system](#) is installed using only license information and log in credentials and is restored using the backup configuration from the Primary CloudArray system.

Configuring the Primary CloudArray

Procedure

1. From the CloudArray main menu, select **Administration > Disaster Recovery Test**.

The **Configure Disaster Recovery Test Mode** panel appears:

Configure Disaster Recovery Test Mode

Disaster Recovery Test Mode Is Currently : **NOT ACTIVE** [Disaster Recovery Documentation](#) ?

Disaster Recovery Test is designed to allow users to test the disaster recovery functionality provided by CloudArray in a safe environment without affecting production data

☐ Select This Checkbox To Begin Activating Disaster Recovery Test Mode ?

2. Activate the disaster recovery test by selecting the checkbox.

The following panel appears:

CloudArray Cloud Providers Cache Status Volume Status

Configure Disaster Recovery Test Mode

Disaster Recovery Test Mode Is Currently : **NOT ACTIVE** [Disaster Recovery Documentation](#) ?

Disaster Recovery Test is designed to allow users to test the disaster recovery functionality provided by CloudArray in a safe environment without affecting production data

☒ Select This Checkbox To Begin Activating Disaster Recovery Test Mode ?

Select Volumes To Be Included ?

	When	Volume Name	Capacity / Dirty	Mapped To	Snapshot In Progress
<input type="checkbox"/>	Now	Volume1	100 GiB / 0 GiB		No

[Cancel](#) [Activate Disaster Recovery Test Mode](#)

3. Select the volumes to be used in the DR test by selecting the checkbox under **Select Volumes to be Included**.
4. From the **When** dropdown list, select **Now** to immediately take the backup snapshot or select **Flush** to take it after a flush operation completes.
5. Click **Activate Disaster Recovery Test Mode** to begin the DR test snapshot process.

You are prompted to confirm your volume selection.

Are you sure you have selected all of the volumes to include? Once the disaster recovery process is started, it can take several hours to complete if you have chosen to create any snapshots after replication.

[No](#) [Yes](#)

Please Confirm

6. Click **Yes** to confirm.
7. When the DR test snapshot completes, you are prompted to download the backup.

CloudArray Cloud Providers Cache Status Volume Status

Disaster recovery test mode is currently in progress. This is the Production CloudArray. (Status: ACTIVE) [Manage](#)

Configure Disaster Recovery Test Mode

Disaster Recovery Test Mode Is Currently : **ACTIVE** [Disaster Recovery Documentation](#) ?

Disaster Recovery Test is currently in progress. To terminate the test, use the 'terminate' button below. Please note that once clicked, your test snapshots will be deleted and test node will stop working.

Disaster Recovery Snapshot Progress ?

Volume	Date/Time Taken	State	Expiration
Volume1	2014-Nov-12 10:52:07	operating	Never

Showing 1 to 1 of 1 entries

[CLICK HERE TO DOWNLOAD DISASTER RECOVERY TEST BACKUP](#)

[Terminate Disaster Recovery Test Mode](#)

8. Click **DOWNLOAD DISASTER RECOVERY TEST BACKUP**.

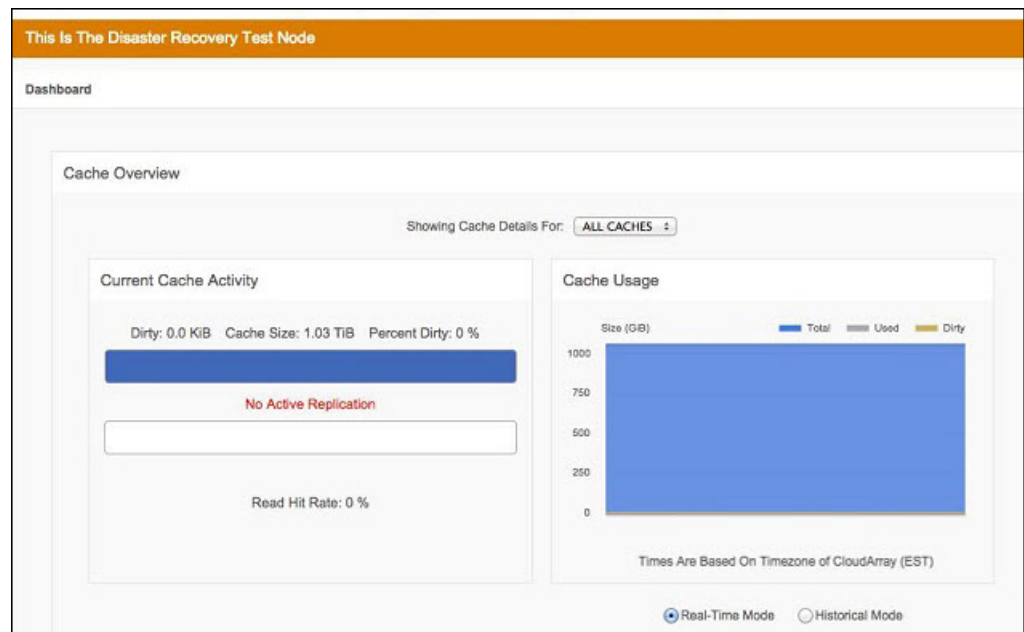
The `ca_backup_DR_TEST.tgz` file will download locally to your Primary CloudArray host PC.

Configuring the DR Test CloudArray

Procedure

1. Complete the steps described in [Restore a CloudArray configuration](#) to perform an unlicensed restore before setting up this DR Test CloudArray.
2. Once the restore completes, log in when prompted.

The **Disaster Recovery Test Node** panel appears with the restored configuration.



If you have administered Active Directory on the Primary CloudArray from where the DR Test Master backup file was created, then you must do the following before exposing a CIFS volume snapshot on the DR Test node:

- Change the hostname of the DR Test Node to something other than the name of the Primary CloudArray from where the DR Test master backup file was created.
 - Join the DR Test Node to the Active Directory domain prior to exposing the CIFS snapshot.
3. To expose a DR volume snapshot, select **Volumes** from the CloudArray main menu and select a volume name to be exposed.

The screenshot shows the 'This Is The Disaster Recovery Test Node' dashboard with the 'Volumes' section selected. It displays a table with the following data:

	Status	Volume Name	Capacity / Dirty	Cloud Provider	Cache	Mapped To
<input type="checkbox"/>	✔	Volume1	100 GiB / 0 GiB	GoogleCloud1	Cache1	NOT MAPPED Map To Client

Showing 1 to 1 of 1 entries

4. Check the snapshot name checkbox and select **Expose**.

Volume1

Cache: Cache1

Cloud Provider: GoogleCloud1

Mapped To Clients: NOT MAPPED TO A CLIENT [Map To Client](#)

Snapshots For This Volume

<input type="checkbox"/>	Name	Date/Time Taken	State	Expiration	DR Test	Exposed As A Volume?	Action
<input type="checkbox"/>	DRTEST.20141112T105207.025446	2014-Nov-12 10:52:07	operating	Never	YES	NO	EXPOSE

Showing 1 to 1 of 1 entries

5. You are prompted to select the cache. Select it and click **OK**.
6. Click **Map to Client** to map the exposed DR snapshot to a client.

This Is The Disaster Recovery Test Node

Volumes > Volume1,DRTEST.20141112T105207.025446

[Map To Client](#) [Cache Migration](#)

Volume Details

Snapshot of: Volume1
 State: Operating
 Health: Online
 Capacity: 100 GiB
 Policy: Policy1

Cache Details

Cache Name: Cache1
 Total Cache Capacity: 33.99 GiB
 Dirty In Cache: 0 GiB
 Last Completed Replication: N/A

Cache: Cache1

Cloud Provider: GoogleCloud1

Mapped To Clients: NOT MAPPED TO A CLIENT [Map To Client](#)

Provisioning policies

CloudArray provisioning policies are containers that show the relationship between a volume, cache and cloud provider. They are automatically created when you create a volume.

To view CloudArray provisioning policies, choose **Administration > Provisioning Policies** from the CloudArray main menu. The **Provisioning Policies** panel lists each policy and displays the cache and cloud provider associated with the policy.

Viewing provisioning policy details

You can view detailed information about each CloudArray provisioning policy.

Procedure

1. Choose **Administration > Provisioning Policies** from the CloudArray main menu

The **Provisioning Policies** panel appears.

2. In the **Policy** column, click the name of the policy you want to view.

Results

The **Provisioning Policies Details** panel displays the volumes using the policy and the cloud provider and cache associated with the policy.

Restoring CloudArray configurations

Using the **Restore CloudArray** option, you can apply a previously saved CloudArray configuration to a new CloudArray. You can also use this utility to return an existing appliance to an earlier configuration.

Note

This process only restores the configuration and not the version of software running on the CloudArray. The backup file also does not contain any data from your volumes or shares.

There are two ways to restore a CloudArray configuration:

- **Unlicensed restore:** apply a CloudArray configuration backup archive to a new CloudArray where you have not yet applied your license
- **Licensed restore:** apply a CloudArray configuration backup archive to a previously licensed CloudArray (this is not common)

NOTICE

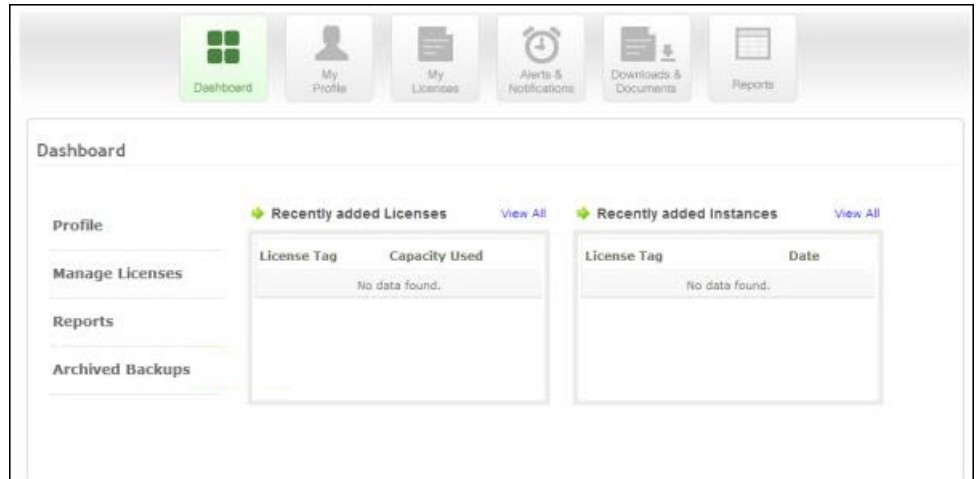
If the original CloudArray configuration had additional cache disks, the new CloudArray server **MUST** have at least the same number of cache disks. For example, if the original CloudArray had a 2000 GB additional disk, the CloudArray on which you perform the restore must also have an additional disk attached before proceeding.

Selecting the configuration to restore

Before performing the restore operation, you need to choose the backup file to restore. You can restore your CloudArray configuration from either an automatic or manual (local) backup file.

Procedure

1. To restore a configuration using a file created during automatic backup, log in to your customer portal account located at cloudarray.com, click **Archived Backups**, and locally download a backup file from the list that appears.



2. To restore a configuration using a manual backup file, use the file that you saved locally. See [Back up CloudArray](#) for more information.

Restoring a CloudArray configuration

Procedure

1. For an unlicensed restore, do the following:
 - a. Import a new CloudArray image to the hypervisor and power on the VM.
 - b. Modify the new CloudArray Host name and IP address (if static).
 - c. Connect to the CloudArray using a supported web browser and click **Restore**.



- d. Go to step 3.
2. For a licensed restore, do the following:
 - a. Log into the licensed CloudArray where you wish to apply the back up.
 - b. From the CloudArray main menu, select **Administration > Restore CloudArray**.




3. Select the backup archive by clicking **Select CloudArray Backup Archive**.

4. Browse to the backup archive file you saved from your account on www.cloudarray.com or from running a manual backup. The file will be named *filename_backup.tgz* or something similar.

After you choose the backup file, the **Restore** button appears.

5. Click **Restore**.



Note


You may be prompted to enter the password for the administrator account of the CloudArray at the time the backup archive was created. This occurs if you are restoring the configuration to either an unlicensed CloudArray or a CloudArray where the current administrator name or password does not match those used when the backup archive was created.

6. A warning dialog informs you that the restore operation requires a CloudArray restart. Click **Yes** to continue.

CloudArray begins the restore process. Once the restore process starts, do not shut down CloudArray. If you close or refresh your browser you will lose the ability to track the restore progress and result. However, the restore process itself will proceed.

The restore process can take several minutes. The user interface will automatically update with the current status of the restore.

7. After CloudArray restarts, the following screen appears. The restore is now complete. Click the link to return to the login screen.



Settings

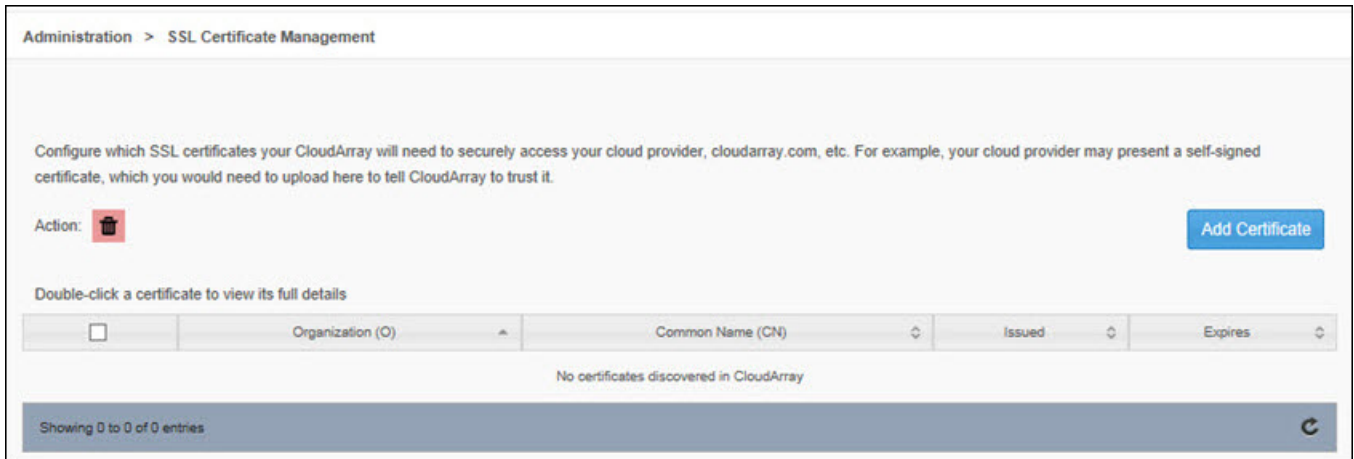
From the **Settings** panel you can toggle the [Snapshot Scheduler](#) and [Bandwidth Throttler](#) on and off.

Managing SSL certificates

If security certificates are required in your environment, CloudArray provides an interface for managing them.

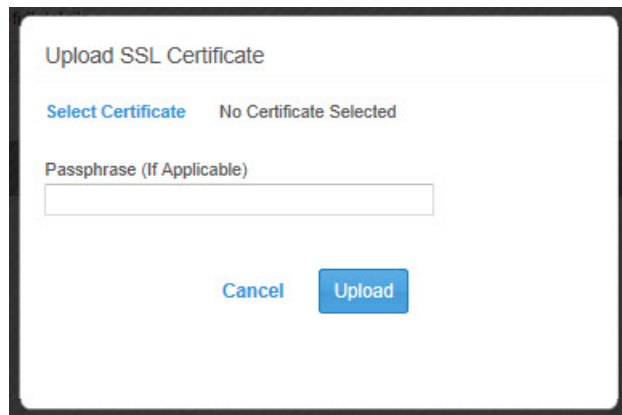
Procedure

1. From the CloudArray main menu, select **Administration > SSL Certificate Management**.



2. Click **Add Certificate**.

The **Upload SSL Certificate** panel appears.



3. Click **Select Certificate** and browse to the .pem formatted certificate.

Note

The .pks12 format is also supported; however, each certificate must be uploaded separately in its own format.

The certificate name displays in the panel.

4. Enter a passphrase if applicable and click **Upload**.

Collecting support data

CloudArray provides an easy way to collect all support-related data and send it for analysis to the EMC support team.

Procedure

1. From the CloudArray main menu, choose **Administration > Support Capture Utility**. Log in if prompted.

2. Click **Generate Support Capture** to continue.
3. Once the data has been collected, you will be prompted to save it locally.

Changing the time zone

Complete the following steps to change your local CloudArray time zone.

Procedure

1. Choose **Administration > Timezone Configuration** from the CloudArray main menu.

The following panel appears.

2. Choose the time zone from the drop-down menu and click **Update TimeZone**.
3. After completing the command, reboot the appliance to apply the change.

User Management

CloudArray provides tools to manage user accounts.

From the CloudArray main menu, select **Administration > User Management**. From that page, you can perform these operations:

- Create new user accounts
- Change the password for a user account
- Change the user role for a user account
- Delete a user account

Related tasks that you can perform:

- [Creating a new user account](#)
- [Managing settings and passwords for a user account](#)

Adding a new user account

CloudArray allows you to create user accounts through the **Administration** page, and allows you to set either a read-only or Administrator role for that user account.

Procedure

1. From the CloudArray main menu, select **Administration > User Management**.
2. Click **Create New User**.

Administration > undefined

Create User

Username

Password

Confirm Password

User Role

☐ Administrator
 ☒ Read-only

Cancel

Create User

3. Enter a the name of the user account in **Username**.
 4. Enter the password for the new user account in **Password** and **Confirm Password**.
 5. Choose a User Role:

Administrator: This role provides full access to view or change any settings in the CloudArray interface.

Read only: This role provides view-only access, and prevents the user from changing any settings in the CloudArray interface.
 6. Click **Create User**.
- The new user account appears in the list of users on the **Users** page.

Managing settings for a user account

CloudArray allows you to manage User Accounts through the **Administration** page, and allows you to change the password and User Role for the user account, and lets you delete the use account.

Procedure

1. From the CloudArray main menu, select **Administration > User Management**.
2. Click on a user account name from list of users on the **Users** page.

The **Update User** screen appears.

3. On the **Update User** screen, you can do any of the following:
 - Use the radio buttons to change the **User Role** for the user account
 - Click **Change Password** to change the password for the user account
 - Click **Delete User** to remove the user account from the CloudArray
4. Click **Update User**.

Installing a custom SSL certificate

CloudArray ships with a self-signed SSL certificate. However, you can use a separate .p12 file and follow these steps to install an SSL certificate signed by a trusted signing authority, such as Verisign, GoDaddy, Thawte, or others.

Procedure

1. From the CloudArray main menu, select **Administration > User Interface Custom SSL**.

2. Click **Select Certificate** to browse to the .p12 file you want to install.
3. Enter the password used during creation of the .p12 file and click **Upload**.

Note

A custom SSL certificate is not required to access your CloudArray securely. Although the connection is secure and encrypted, you may see a warning when accessing your CloudArray with the default SSL certificate in place. This is because the certificate is unknown to a public signing authority (Verisign, GoDaddy, Thawte, etc.).

After installing the certificate, ensure that your browser accesses your CloudArray using the domain name provided in the certificate. For example, if the certificate was signed with cloudarray.com as the domain name, this CloudArray might be accessed at cloudarray1.cloudarray.com.

Updating the CloudArray license

If you are connected to the CloudArray portal you can update your license directly from CloudArray.

Procedure

1. Choose **Administration > Update CloudArray License** from the CloudArray main menu.
2. Enter your CloudArray portal email address and password.
3. Click **Save Changes**.

Utilities

CloudArray provides several built-in utilities and CLI commands you can execute to automate administrative tasks.

- From the CloudArray main menu, select **Administration > Utilities**, then select either the [Execute Utility](#) or [Execute CLI](#) button.

Executing CLI commands

CloudArray provides command line interface (CLI) commands that allow you to perform the following tasks:

- [User management](#)
- [Provider management](#)
- [Cache management](#)
- [Pool management](#)
- [Volume management](#)
- [Client management](#)
- [System management](#)

Procedure

1. From the CloudArray main menu, select **Administration > Utilities**.
2. Check the **Execute CLI** button and type the command name.

Administration > Utilities

☐ Execute Utility ☒ Execute CLI

Enter CLI command:

Results:

3. Click **Go**.

The results of the command appear in the panel below the CLI command text box.

User management commands

Command	<code>list_users</code>
Explanation	Display a list of the available user accounts.
Command	<code>new_user --name name --password password</code>
Explanation	Create a new user account with name and cleartext password by the caller.
Command	<code>check_user --name name --password password</code>
Explanation	Check the validity of a user name/password combination.
Command	<code>change_password --name name --old old --new new [--force true false]</code>
Explanation	Change the user password. If the old password is not provided, the force flag must be set to true.
Command	<code>remove_user --name name</code>
Explanation	Remove an existing named user account.

Provider management commands

Command	<code>new_provider --pool pool --name name [options]</code>
Explanation	Create a new provider of specified type in named pool. Current types and options are: <code>s3 --key key --secret secret [--node node] [--secure true false] [--port port] [--location location] [--bucket bucket]</code> <code>atmos --token token --secret secret [--node node] [--secure true false] [--port port]</code>

```
synaptic --user userID --application applicationID --secret secret [--node node] [--secure true|false] [--port port]
```

Command	<code>remove_provider --pool pool --name name</code>
Explanation	Remove a provider from the cloud array. This action does not affect any volumes that were created using this provider's policy description.
Command	<code>provider_info --pool pool --name provider</code>
Explanation	Provide an xml description of the provider that was created and its current status, for example, up/down.
Command	<code>set_encryption --pool pool --provider provider -value true false [-keypair keypair]</code>
Explanation	Set the encryption for a given provider in a given pool to be on or off. If a keypair is not specified, generate and use a new one. The provider you wish to change this for must have no policies or volumes associated with it.
Command	<code>set_compression --pool pool --provider provider -value true false [-level 0 1 2 3 4 5 6 7 8 9]</code>
Explanation	Set compression for a given provider in a given pool to be on or off. The provider you wish to change this for must have no policies or volumes associated with it.

Cache management commands

Command	<code>new_cache --name name --pool pool [--page_size pageSize] [--block_size block_size]</code>
Explanation	Add a new cache policy to the system using the named pool. Optionally, set page and block size.
Command	<code>add_cache_volume --name name --cache cache --size size [--units blocks MB GB] [--pool pool]</code>
Explanation	Create a new volume based on the policies described in the cache. GB is the default size unit. Optionally, draws capacity from an alternate specified pool.
Command	<code>remove_cache --cache cache</code>
Explanation	Remove a cache policy.
Command	<code>list_caches [--raw]</code>
Explanation	List caches configured on the system.
Command	<code>flush_volume --name name</code>
Explanation	Force a flush on the named volume. This operation will remain pending until the flush completes.

Pool management commands

Command	<code>new_pool --name name [--cache true false]</code>
Explanation	Add a new pool to the system with the specified name. May have the attribute of being a pool for caches.
Command	<code>remove_pool --name name</code>

Explanation	Remove a pool from the system. Does not remove any volumes created using this pool as a prototype.
Command	<code>list_pools [--raw]</code>
Explanation	List pools configured on the system.

Volume management commands

Command	<code>list_volumes [--raw]</code>
Explanation	List volumes configured on the system.
Command	<code>new_policy --name name --pool pool --cache cache</code>
Explanation	Create a new volume policy as named, drawing capacity from the specified pool and cache, optionally dedicated.
Command	<code>remove_policy -name</code>
Explanation	Remove the named volume policy. Removing the policy does not affect any existing volumes.
Command	<code>new_volume --name name --policy policy --size size [--units blocks MB GB]</code>
Explanation	Create a new volume based on the policies described in the pool. Defaults to GB as the size unit.
Command	<code>remove_volume --name name</code>
Explanation	Remove the named volume from the system. Will also remove any mappings and/or snapshots based on this volume.
Command	<code>lock_volume --name name</code>
Explanation	Prevent any I/O from being processed by this volume.
Command	<code>unlock_volume --name name</code>
Explanation	Allow I/O processing on this volume.
Command	<code>new_snapshot --volume name --when now flush after_next</code>
Explanation	Create a snapshot of the named volume at the specified event. This can occur at three specified times: <ul style="list-style-type: none"> • <code>now</code> – will be consistent with the previous but not the current flush • <code>flush</code> – will be consistent with the current flush, if active • <code>after_next</code> – will always wait for the next flush to complete before snapping
Command	<code>expose_snapshot --volume name --snapshot name [--cache cache]</code>
Explanation	Create a read/write volume based on the named snapshot which draws its local storage from the specified cache. It does not map the volume to any clients.
Command	<code>hide_snapshot --volume name --snapshot name</code>
Explanation	Remove the specified snapshot volume from the system. It does not remove the snapshot data.
Command	<code>remove_snapshot --volume name --snapshot name</code>

Explanation	Remove the specified snapshot data from the system. Any volumes exposed by this snapshot will also be removed.
--------------------	--

Client management commands

Command `list_clients [--raw]`

Explanation List clients configured on the system.

Command `new_client --name name [--iqn name]`

Explanation Add a new named client to the system. If the iqn is specified, it will use that iqn as the identifier. If the iqn is not specified, it will attempt to match the specified name with the final characters of the iqns of incoming clients. Note that clients are automatically added in the discovery phase, so this command is an optional part of the process. It can be used to map volumes to a client before the client is attached.

Command `rename_client --old name --new name`

Explanation Change the external name of the client, which is the name that will be used by the UI.

Command `map_volume --volume name --client name [--option readonly|readwrite]`

Explanation Expose the volume to the client, optionally as a read-only disk. If the mapping already exists, permissions can be modified.

Command `unmap_volume --volume name --client name`

Explanation Remove the volume from the client. Does not check for IO in progress before removal.

Command `remove_client --name name`

Explanation Remove a named client from the system.

System management commands

Command `terminate`

Explanation Cease operations.

Command `version_info`

Explanation Report version and system information.

Command `set_logging --level 0|1|2|3|4|5|6|7|8|9`

Explanation Set the logging output to the specified level.

Command `set_license --owner owner --password password --license license [--user user --userpass password]`

Explanation Set the license for a cloudarray system. The owner, password, and license fields are passed on to the server to retrieve the actual license information. If a valid license key is not set on the system, then the user and password fields are not required. If the current system is licensed, a valid local user and password is necessary in order to change the license information.

Command `list_stats`

Explanation List the current system I/O statistics.

Command	<code>reset_log</code>
Explanation	Reset the system log files to empty.
Command	<code>license_info</code>
Explanation	Retrieve the current license status, including any restrictions and expirations.
Command	<code>set_tag --type volume client pool provider --name name --tag tag</code>
Explanation	Set a tag on a named volume, client, pool, or provider.
Command	<code>clear_tag --type volume client pool provider --name name --tag tag</code>
Explanation	Clear a tag on a named volume, client, pool, or provider.
Command	<code>get_key --name name [--passphrase passphrase]</code>
Explanation	Return a pkcs12 envelope containing the named key pair, optionally encrypted using the passphrase.
Command	<code>add_key --key key [--passphrase passphrase]</code>
Explanation	Takes a pkcs12 envelope containing a named key pair, optionally encrypted using the passphrase, and stores it as a tempest key pair.
Command	<code>list_security [--raw]</code>
Explanation	List the keys that are currently installed.
Command	<code>set_proxy [-h] [--proxy PROXY] [--port PORT] [--useProxy {true,false}] [--proxyUsername PROXYUSERNAME] [--proxyPassword PROXYPASSWORD]</code>
Explanation	Set proxy for CloudArray to use to reach cloud providers.

Executing utilities

Procedure

1. From the CloudArray main menu, select **Administration > Utilities**.
2. Select the **Execute Utility** button and choose a utility from the drop-down menu.

Administration > Utilities

☒ Execute Utility ☐ Execute CLI

Utility: Select a Utility Parameters (optional): Go

Results:

The following utilities are available:

- **cache_usage:** Displays the cache usage for each CloudArray cache disk resource.
- **check_share_health:** Checks health of CIFS shares and attempts to recover from common problems.
- **fill_volume_cache:** Operates in the background to reload a cache with all volume data from the cloud. Use this command when the cache has been lost or used to pre-populate the cache for an exposed snapshot. This command will not eject used pages from the cache. It attempts to copy as much data as will fit in cache starting at the beginning of the volume.
- **klist:** Lists the Kerberos principal and Kerberos tickets held in a credentials cache.
- **local_disk:** Cleans up a local disk (and associated metadata) used for cache storage.
- **reboot:** Reboots the system. Enter the --confirm parameter with this command.
- **reset_configuration:** Erases all configuration and the raw device headers on the cache disks, allowing them to be re-used. Enter the --confirm parameter with this command. Note that all dirty data in cache will be lost. If you do not have a backup archive, all data in the cloud will be orphaned. Be sure to either have a backup archive or be certain that you no longer need this CloudArray configuration.
- **set_read_ahead:** Supplements volume data pages read from the cloud due to a read cache miss with extra sequential pages so that subsequent sequential reads are pre-cached.
- **update_license:** Forces CloudArray to check for an updated license after renewal or capacity upgrade.
- **versions:** Retrieves the version numbers for various CloudArray components.

3. Click **Go**. The results display in the panel.

CHAPTER 13

Alerting

This chapter describes CloudArray alerting facilities.
Topics include:

- [Registering with EMC Secure Remote Services](#) 110
- [Configuring CloudArray portal alerts](#)..... 110
- [Configuring SNMP traps](#)..... 110

Registering with EMC Secure Remote Services

EMC Secure Remote Services (ESRS) is an IP-based automated connect home and remote support solution that provides a common point of access for remote support activities performed on EMC products. Complete the following steps to connect to the ESRS system.

Note

ESRS version 3.08 or higher is required.

Procedure

1. Choose **Alerting > EMC Secure Remote Services** from the CloudArray main menu.
2. Complete the text fields.
 - Enter your username and password.
 - If you have a CloudArray physical appliance, locate and enter the PSNT number listed on the tag attached to the appliance. If you have a CloudArray virtual edition, this field is dimmed.
 - The remaining fields are specific to the ESRS gateway on your site. Consult your ESRS documentation for details.
3. Click **Save Changes**.

Configuring CloudArray portal alerts

You can configure CloudArray to send alerts to the CloudArray portal.

Procedure

1. Choose **Alerting > Portal Alerts** from the CloudArray main menu.
2. Select the **Send Alerts to CloudArray Portal** checkbox.
3. Click **Save Changes**.
Alerts will now appear in the CloudArray portal.
4. To view alerts in the CloudArray portal, log in to the portal, then click the **Alerts & Notifications** button.
Any alerts appear in the **Manage Alerts and Notifications** panel.

Configuring SNMP traps

SNMP (Simple Network Management Protocol) traps are alerts generated from software to administrative managers. Complete the following steps to configure CloudArray to send SNMP traps to a Network Management System (NMS).

Procedure

1. Choose **Alerting > SNMP Configuration** from the CloudArray main menu.
2. Select the **Enable CloudArray To Send SNMP Traps** checkbox.
3. Enter your IP address in the **Network Management System IP Address** text field. Accept the default values for the remaining text fields.

4. To confirm that the panel is configured correctly, click **Send Test Trap** and verify that your SNMP receiver received the trap.
5. Click **Save Changes**.

CHAPTER 14

Reporting

This chapter describes hardware settings reported in CloudArray.

Note

This option is available if you are using a CloudArray physical appliance. The menu item is not present if you are using CloudArray virtual edition.

Topics include:

- [Hardware details](#)..... 114

Hardware details

Choose **Reporting > Hardware** from the CloudArray main menu to view hardware details. The **Hardware Details** panel provides information about the following:

- BMC (Baseboard Management Controller) chassis components
- BMC FRU (Field Replaceable Unit) details
- FRU power supplies
- Front panel
- HS backplane
- I/O module
- BMC sensor
- BMC alerts

This information can be used to help diagnose problems related to the CloudArray hardware.