

Isilon

OneFS CloudPools

Version 8.1.0

Administration Guide

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

Published May 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

Chapter 1	Introduction to this guide	7
	About this guide.....	8
	Where to go for support.....	8
Chapter 2	Setting up CloudPools	9
	CloudPools overview.....	10
	CloudPools concepts.....	10
	Accessing CloudPools with IsilonSD Edge.....	12
	Supported cloud providers.....	12
	EMC Isilon.....	12
	EMC ECS Appliance.....	13
	Virtustream Storage Cloud.....	13
	Amazon S3.....	13
	Microsoft Azure.....	14
	CloudPools setup and management.....	14
	Activating a CloudPools software license.....	14
	Configuring network proxy servers with CloudPools.....	15
	Managing cloud storage accounts.....	17
	Managing CloudPools.....	21
	Managing CloudPools settings.....	24
Chapter 3	Managing CloudPools policies	27
	CloudPools file processing.....	28
	Archiving files with file pool policies.....	28
	Retrieving file data from the cloud.....	33
	Managing cloud policies.....	34
	Create a file pool policy for cloud storage (Web UI).....	34
	Create a file pool policy for cloud storage (CLI).....	36
	Modify cloud attributes in a file pool policy (Web UI).....	36
	Modify cloud attributes in a file pool policy (CLI).....	37
	List file pool policies (CLI).....	37
	View details of a file pool policy (CLI).....	37
	Apply a file pool policy to a specified file or path (CLI).....	38
	Archive files directly to the cloud (CLI).....	38
Chapter 4	Managing CloudPools with other OneFS functions	39
	Compression and encryption of cloud data.....	40
	CloudPools protocol support.....	40
	NFS inline access.....	40
	SMB inline access.....	41
	Other protocols supporting inline access.....	41
	SynclQ interoperability.....	41
	SynclQ policies.....	41
	CloudPools cloud data retention time.....	42
	Managing SmartLink files that have been replicated.....	42
	SynclQ deep copy.....	43
	Configuring access to cloud data from a secondary cluster.....	44
	NDMP backup and restore of cloud data.....	46

	Checking the version of SmartLink files.....	46
	CloudPools and snapshots.....	47
	CloudPools and SmartLock.....	47
	CloudPools and SmartQuotas.....	48
	CloudPools and SmartDedupe.....	48
Chapter 5	CloudPools tips and troubleshooting	49
	CloudPools best practices.....	50
	Use time stamps for cloud data archival and recall.....	50
	CloudPools archiving and file size.....	50
	Create exclusive accounts for CloudPools purposes.....	50
	Managing cloud jobs.....	50
	View a list of cloud jobs (CLI).....	51
	View a cloud job (CLI).....	51
	Pause a cloud job (CLI).....	51
	Resume a paused cloud job (CLI).....	52
	Cancel a cloud job (CLI).....	52
	CloudPools troubleshooting.....	52
	CloudPools limitations and expected behaviors.....	52
	CloudPools logs.....	54
	Troubleshooting CloudPools.....	55
Chapter 6	CloudPools CLI commands	57
	CloudPools command reference.....	58
	isi cloud access add.....	58
	isi cloud access list.....	58
	isi cloud access remove.....	60
	isi cloud access view.....	60
	isi cloud accounts create.....	61
	isi cloud accounts delete.....	62
	isi cloud accounts list.....	63
	isi cloud accounts modify.....	64
	isi cloud accounts view.....	65
	isi cloud archive.....	66
	isi cloud jobs cancel.....	67
	isi cloud jobs files list.....	67
	isi cloud jobs list.....	68
	isi cloud jobs pause.....	69
	isi cloud jobs resume.....	70
	isi cloud jobs view.....	70
	isi cloud pools create.....	70
	isi cloud pools delete.....	71
	isi cloud pools list.....	72
	isi cloud pools modify.....	73
	isi cloud pools view.....	74
	isi cloud proxies create.....	74
	isi cloud proxies delete.....	75
	isi cloud proxies list.....	76
	isi cloud proxies modify.....	77
	isi cloud proxies view.....	78
	isi cloud recall.....	78
	isi cloud restore_coi.....	79
	isi cloud settings modify.....	80
	isi cloud settings regenerate-encryption-key.....	81

isi cloud settings view..... 82

CONTENTS

CHAPTER 1

Introduction to this guide

This section contains the following topics:

- [About this guide](#)..... 8
- [Where to go for support](#)..... 8

About this guide

This guide describes CloudPools, a licensed software module that works with Isilon's OneFS operating system. This guide describes how the CloudPools interface provides access to OneFS cloud configuration, operation, and management.

Your suggestions help us to improve the accuracy, organization, and overall quality of the documentation. Send your feedback to <https://www.research.net/s/isi-docfeedback>. If you cannot provide feedback through the URL, send an email message to docfeedback@isilon.com.

Where to go for support

If you have any questions about EMC Isilon products, contact EMC Isilon Technical Support.

Online Support	<ul style="list-style-type: none"> • Live Chat • Create a Service Request
Telephone Support	<ul style="list-style-type: none"> • United States: 1-800-SVC-4EMC (1-800-782-4362) • Canada: 1-800-543-4782 • Worldwide: 1-508-497-7901 • Local phone numbers for a specific country are available at <i>EMC Customer Support Centers</i>.
Support registration or access	For questions specific to EMC Online Support registration or access, email support@emc.com .
Isilon Info Hubs	For the list of Isilon info hubs, see the Isilon Info Hubs page on the EMC Isilon Community Network. Isilon info hubs organize Isilon documentation, videos, blogs, and user-contributed content into topic areas, making it easy to find content about subjects that interest you.

Support for IsilonSD Edge

If you are running a free version of IsilonSD Edge, support is available through the [Isilon Community Network](#). If you purchased one or more IsilonSD Edge licenses, support is available through , provided you have a valid support contract for the product.

CHAPTER 2

Setting up CloudPools

This section provides conceptual information about, and procedures for, setting up CloudPools:

- [CloudPools overview](#) 10
- [Supported cloud providers](#) 12
- [CloudPools setup and management](#) 14

CloudPools overview

CloudPools extends the capabilities of OneFS by enabling you to specify data to be moved to lower-cost cloud storage. CloudPools can seamlessly connect to EMC-based cloud storage systems and to popular third-party providers, Amazon S3 and Microsoft Azure.

CloudPools is a licensed module built on the SmartPools file pool policy framework, which gives you granular control of file storage on your cluster. CloudPools extends this file storage control to one or more cloud repositories, which act as additional tiers of OneFS storage.

Prior to the introduction of CloudPools, SmartPools enabled the grouping of nodes into storage pools called node pools, and the classification of node pools as different storage tiers. SmartPools includes a policy framework that allows you to segregate files into logical groups called file pools, and to store those file pools in specific storage tiers.

CloudPools expands the SmartPools framework by treating a cloud repository as an additional storage tier. This enables you to move older or seldom-used data to cloud storage and free up space on your cluster.

File pool policies

As with SmartPools, you define files to be stored in the cloud by creating file pool policies. These policies use file matching criteria to determine which file pools are to be moved to the cloud.

File pool policies are applied when the SmartPools system job runs, by default on a daily basis. For each policy, all matched files on the cluster are handled according to policy specifications.

When files match a file pool policy that contains CloudPools actions, OneFS moves the file data of matched files to the cloud. Only metadata and a proxy file remain on the cluster, thus freeing up storage space.

SmartLink files

Although file data is moved to remote storage, the files remain visible in the OneFS file system. CloudPools accomplishes this by retaining a local SmartLink file, which is a pointer to the location of data in the cloud. You can read, write, archive, and recall files from the cloud as needed.

When a user accesses a cluster and views the OneFS file system through a supported protocol (SMB, NFS, Swift, or HDFS), SmartLink files appear to be the original files. When the user opens a SmartLink file, OneFS automatically retrieves and caches as much data as needed from the cloud. This operation is called inline access. Any modifications the user makes to a file during inline access are updated in the file data stored in the cloud.

In addition to inline access, CloudPools also provides a CLI command to enable full recall of files from the cloud, in which case the SmartLink files are replaced by the actual files.

CloudPools concepts

CloudPools is a licensed module that enables you to move file data on your Isilon cluster to the cloud, and to access or recall these files when needed. Taking advantage

of CloudPools requires you to configure cloud storage accounts and file pool policies that specify cloud storage targets.

You can configure CloudPools to move files to the cloud automatically, based on file pool policies. You can also use a OneFS command to archive individual files to, or recall files from, the cloud.

CloudPools uses a similar workflow to OneFS SmartPools. To store files in the cloud, you must have at least one account with a cloud storage provider. In addition, you must configure OneFS for cloud storage, and create file pool policies that identify and move files to the cloud.

When the SmartPools job runs, typically once a day, file pool policies are executed, and matched files are sent to the cloud storage target. To access file data stored in the cloud, you can open its related SmartLink file through any supported protocol (SMB, NFS, Swift, or HDFS). This is referred to as inline access. To fully recall a file from the cloud, you can issue an `isi cloud recall` command from the OneFS command-line interface.

Following are descriptions of key CloudPools concepts:

Archive

The CloudPools process of moving file data to the cloud. This process involves extracting the data from the file and placing it in one or more cloud objects. CloudPools then moves these objects to cloud storage, and leaves in place on the local cluster a representative file called a SmartLink.

Recall

The CloudPools process of reversing the archival process. When you recall a file from the cloud, CloudPools replaces the SmartLink file by restoring the original file data on OneFS and removing the cloud objects from cloud storage.

SmartLink file

For every file archived to the cloud, OneFS maintains an associated SmartLink file on the cluster. A SmartLink file contains metadata and map information so the data in the cloud can be accessed or fully recalled. If allowed by a SmartLink file's archiving policy, accessing the SmartLink file on the cluster automatically retrieves and caches data from the cloud. Like other files, SmartLink files can be backed up through NDMP or synchronized to other clusters with SyncIQ. When SmartLink files are retrieved from a backup or SyncIQ operation, CloudPools maintains their links to related file data in the cloud.

File pool policies

File pool policies are the essential control mechanism for both SmartPools and CloudPools. OneFS runs all file pool policies on a regular basis. Each file pool policy specifies the files to be managed, actions to take on the files, protection levels, and I/O optimization settings.

If CloudPools has been enabled, each file pool policy can also contain cloud-specific parameters that specify the remote cloud account to archive files to, and how to handle files prior to archiving them. Moreover, a policy can also specify SmartPools targets that specify where to store the locally-retained SmartLink files related to the stored cloud data.

Cloud provider accounts

Making use of cloud storage requires you to set up one or more accounts with a cloud provider. The types of cloud storage that are currently supported are EMC Isilon, EMC ECS Appliance, Virtustream Storage Cloud, Amazon S3, and

Microsoft Azure. The account information from the cloud provider must match the information you use when configuring cloud accounts on your Isilon cluster.

Cloud storage accounts

A cloud storage account is a OneFS entity that defines access to a specific cloud provider account. The cloud storage account settings must match the account credentials provided by the cloud provider.

CloudPool

A CloudPool is a OneFS entity that contains a single cloud storage account and provides a conduit between OneFS and the cloud storage repository. Creating a CloudPool requires the availability of at least one cloud storage account. The cloud storage account must be of the same type as the CloudPool.

Inline access

CloudPools enables users connecting to a cluster through supported protocols to access cloud data by opening associated SmartLink files. This process is referred to as inline access. To the user connecting to OneFS through a supported protocol, a SmartLink file appears to be the original file. When the user opens a SmartLink file, CloudPools retrieves and caches cloud data locally. The user can view and edit the file as usual. CloudPools automatically retrieves and sends any updated file data back to the cloud so that the cloud contains the latest version.

Note

CloudPools offers inline access as a user convenience. However, CloudPools is designed mainly as an archival solution, and is not intended for storing data that is frequently updated. Such data should be left on the local cluster until it stabilizes and is ready for archival.

Accessing CloudPools with IsilonSD Edge

The CloudPools software module is available only with a purchased license of IsilonSD Edge. It is not packaged with the free license of this product.

Supported cloud providers

With CloudPools, OneFS supports these cloud providers: EMC Isilon, EMC ECS Appliance, Virtustream Storage Cloud, Amazon S3, and Microsoft Windows Azure.

EMC Isilon

CloudPools enables an EMC Isilon cluster to function as a cloud storage provider.

In this scenario, a secondary EMC Isilon cluster provides a private cloud solution. The secondary cluster archive files from your primary cluster, and is managed in your corporate data center.

To act as a cloud storage provider, an EMC Isilon cluster uses a set of APIs that include the capabilities to configure CloudPools policies, define cloud storage accounts, and retrieve cloud storage usage reports. These APIs are known collectively as the Isilon Platform API, and are described in the *OneFS 8.0 API Reference*.

The secondary cluster must be running the same version of OneFS as the primary cluster and must have active SmartPools and CloudPools licenses.

To configure a secondary EMC Isilon cluster as the cloud storage repository, you need to complete several tasks:

- On the secondary cluster, log on with system administrator privileges, and create a new user.
- On the secondary cluster, create a role with access to Console, Platform API, HTTP, License, Namespace Traverse, and Namespace Access privileges, and make the new user a member of this role.
- On the secondary cluster, log on as the new user, and create the directory where cloud data should be stored. For example: `/ifs/data/HQ-Archive`.
- On the primary cluster, set up the EMC Isilon cloud storage account, specifying the new user's credentials and the appropriate URI for the secondary cluster. Because the secondary cluster is within your corporate network, the URI would look similar to the following example:

```
https://10.1.210.310:8080/namespace/ifs/data/HQ-Archive
```

- On the primary cluster, create a CloudPool that contains the EMC Isilon cloud storage account.

EMC ECS Appliance

CloudPools supports EMC ECS (Elastic Cloud Service) appliance as a cloud provider.

ECS is a complete software-defined cloud storage platform deployed on a turn-key appliance from EMC. It supports the storage, manipulation, and analysis of unstructured data on a massive scale.

The ECS appliance is specifically designed to support mobile, cloud, big data, and next-generation applications. As an appliance, it is simple to install and deploy with support for multi-tenancy, self-service access, usage metering, on-demand cloud storage-as-a-service and dynamic application provisioning.

Virtustream Storage Cloud

CloudPools supports Virtustream Storage Cloud as a cloud provider.

Virtustream Storage Cloud (VSC) is a managed cloud computing service from EMC that offers benefits associated both with a dedicated private cloud and a public multi-tenant cloud.

VSC enables enterprises to run complex, mission-critical applications with full cloud agility, economy, and automation, and to achieve enterprise-class service-level agreements for both application performance and availability. In addition, Virtustream provides a foundation for meeting national and industry-level security, compliance and auditing requirements.

Virtustream Storage Cloud can be managed in your own data center, or provided as a solution from a third party.

Amazon S3

CloudPools can be configured to store data on Amazon S3 (Simple Storage System), a public cloud provider.

When you configure CloudPools to use Amazon S3 for cloud storage, in addition to URI, username, and passkey, you must specify the following additional attributes.

- S3 Account ID

- S3 Telemetry Reporting Bucket
- S3 Storage Region

When you first establish an account with Amazon S3, the cloud provider gives you an account ID and allows you to choose a storage region. Amazon S3 offers multiple storage regions in the U.S. and other regions of the world.

Note

CloudPools supports Amazon Web Services Signature V2 to authenticate queries to its cloud storage. CloudPools does not support Signature V4.

To work with CloudPools, you must also identify an S3 telemetry reporting bucket. This is where billing reports are stored on Amazon S3. This bucket must be accessible to CloudPools.

To set up an S3 telemetry reporting bucket, navigate to **Billing & Cost Management** preferences in the S3 console. There you can indicate that you want to receive billing reports, and specify the bucket to which these reports should be saved. Use this bucket name as the telemetry reporting bucket when setting up an S3 cloud storage account in CloudPools.

Microsoft Azure

You can configure CloudPools to store data on Microsoft Azure, a public cloud provider

When you first establish an account with Microsoft Azure, you create a user name, and Microsoft provides you with a URI and a passkey. When you configure CloudPools to use Azure, you must specify the same URI, username and passkey.

CloudPools setup and management

Setting up and managing CloudPools includes activating licenses, configuring network proxies, and managing accounts and storage pools.

Activating a CloudPools software license

You can activate a CloudPools license from either the web interface or the CLI. Running CloudPools also requires the activation of a SmartPools license.

For complete information about obtaining and activating OneFS licenses, see the *Isilon OneFS Web Administration Guide*.

Upload the updated license file (Web UI)

After you receive an updated license file from EMC Software Licensing Central (SLC), upload the updated file to your cluster.

Procedure

1. Click **Cluster Management > Licensing**.
2. In the **Upload and activate a signed license file** area, click **Browse** and select your signed license file.
3. Click **Upload and Activate**.

Upload the updated license file (CLI)

After you receive an updated license file from EMC Software Licensing Central (SLC), upload the updated file to your cluster.

Procedure

1. Run the `isi license add` command.

The following command adds the `/ifs/local` license file to the cluster:

```
isi license add --path /ifs/local
```

Configuring network proxy servers with CloudPools

You can configure CloudPools so that data that is archived to, or recalled from, a public cloud provider is routed through a proxy server.

By default, CloudPools communicates directly with the designated cloud provider. If the cloud provider is private, such as another Isilon cluster or an EMC ECS appliance running on the same corporate network, the default communication protocol might be acceptable.

However, if CloudPools is archiving data to a public cloud provider, such as Amazon S3 or Microsoft Azure, communication happening directly through the public Internet might violate security policies that are established by some organizations.

In a typical configuration, the Isilon cluster is installed in a data center behind one or more firewalls. Ports that would enable communication to the public Internet often are closed. To enable CloudPools to archive data to a public cloud provider, CloudPools can be configured to work with a proxy server.

CloudPools works with proxy servers running the following protocols:

- SOCKS v4
- SOCKS v5
- HTTP

Configuration on the CloudPools side includes creating the network proxy and connecting the network proxy to a cloud storage account. Both SOCKS v5 and HTTP can be configured with or without authentication. SOCKS v4 does not support authentication.

From OneFS, you can also list network proxies, view network proxy properties, modify proxy settings, and delete proxies. Except for connecting the network proxy to a cloud storage account, you must use the CLI to run all other proxy server commands.

Create a network proxy (CLI)

You can create a network proxy to redirect CloudPools traffic to and from a public cloud provider. CloudPools supports proxy servers running the SOCKS v4, SOCKS v5, and HTTP protocols.

Before you begin

The proxy server should be online and ready to accept a connection from an Isilon cluster.

Procedure

1. Run the `isi cloud proxies create` command.

The following command creates a proxy object named `myproxy1` and links it to a specific proxy server URL, proxy type, and port:

```
isi cloud proxies create myproxy1 10.99.58.250 socks_5 1080
```

Results

When you later create or modify a cloud storage account, the `myproxy1` network proxy is available. At that time, when select the proxy, and save the changes, CloudPools verifies that the proxy server connection can be made.

View a list of network proxies (CLI)

You can view a list of existing network proxies in CloudPools.

Before you begin

You or someone in your organization must first have created network proxies using the `isi cloud proxies create` command.

Procedure

1. Run the `isi cloud proxies list` command.

The command displays a list of proxy names, hosts, and types.

View network proxy properties (CLI)

You can view the properties of a network proxy.

Before you begin

You or someone in your organization must have created a network proxy using the `isi cloud proxies create` command.

Procedure

1. Run the `isi cloud proxies view` command.

The following command displays the properties of a proxy named `myproxy1`:

```
isi cloud proxies view myproxy1
```

Properties shown include ID, name, host, type, and port.

Modify a network proxy (CLI)

You can modify the properties of an existing network proxy in CloudPools.

Before you begin

You or someone in your organization must have created the network proxy using the `isi cloud proxies create` command.

Procedure

1. Run the `isi cloud proxies modify` command.

The following command adds a user name and password necessary to connect to a network proxy:

```
isi cloud proxies modify myproxy1 --username cloud1 --
password @xy16+RZ20
```

Results

You can now add the network proxy to a cloud storage account.

Delete a network proxy (CLI)

You can delete an existing network proxy in CloudPools. However, if the proxy is connected to a cloud storage account, you cannot delete the proxy.

Before you begin

You or someone in your organization must have created the network proxy using the `isi cloud proxies create` command.

Procedure

1. Run the `isi cloud proxies delete` command.

The following command deletes the proxy named `myproxy1`:

```
isi cloud proxies delete myproxy1
```

OneFS asks you to confirm the deletion:

```
Are you sure? (yes/no) :
```

2. Type `yes` and press ENTER.

Results

If the proxy is already connected to a cloud storage account in CloudPools, OneFS prevents you from deleting the proxy. Otherwise, the proxy is deleted.

Managing cloud storage accounts

A cloud storage account provides OneFS with the information it needs to connect to the remote cloud storage provider.

You can create and edit one or more cloud storage accounts in OneFS.

Create cloud storage accounts (Web UI)

You define cloud storage accounts on your OneFS cluster as an essential part of CloudPools configuration. The account username, password, and URI that you used to establish an account with your cloud provider is required. You can also specify a proxy server to redirect CloudPools archive and retrieval traffic to and from a public cloud provider.

Procedure

1. Click **File System > Storage Pools > CloudPools**.
2. Click **+ Create a Cloud Storage Account**.
3. In the **Create a Cloud Storage Account** dialog box, perform these actions:

- a. In the **Name or Alias** field, enter a name for the account.
 - b. In the **Type** drop-down menu, select a type of cloud account. Choices are EMC Isilon, EMC ECS Appliance, Virtustream Storage Cloud, Microsoft Azure, and Amazon S3.
 - c. In the **URI** field, enter a fully qualified URI for the account. The URI must use the HTTPS protocol, and match the URI used to set up the account with your cloud provider.
 - d. In the **User Name** field, enter the account user name, which must be the same as the user name provided to or by the cloud provider.
 - e. In the **Key** field, enter the account password. The password must be the same as the password that you provided to the cloud provider, or the key that the cloud provider issued to you.
 - f. If you have already defined one or more network proxies, and want to use one for this cloud account, select the name from the **Proxy** drop-down box.
 - g. If you are creating an Amazon S3 account, you must also specify an S3 Account ID, S3 Telemetry Reporting Bucket, and an S3 Storage Region.
4. Click the **Connect Account** button.

The **Create a Cloud Storage Account** dialog box closes, and the new cloud account appears in the **Cloud Storage Accounts** list. A cloud account is represented by an orange user icon. The Name, Type, State, Username, and URI associated with the account is displayed.

Create a cloud storage account (CLI)

You create cloud storage accounts to enable CloudPools to archive files to cloud storage. The account username, password, and URI that you used to establish an account with your cloud provider is required. You can also specify a proxy server to redirect CloudPools archive and retrieval traffic to and from a public cloud provider.

Before you begin

Before creating a cloud storage account, you must establish an account with a cloud provider, such as EMC Elastic Cloud Storage Service. When you create a cloud storage account in OneFS, the system attempts to connect to the cloud provider using the credentials you provide. Also, if you want to specify a proxy server with the cloud storage account, you must have already created the proxy server by means of the `isi cloud proxies create` command.

OneFS enforces the following requirements for cloud storage accounts.

- Each cloud storage account can only belong to a single CloudPool storage container.
- A cloud storage account must be of the same type (EMC Isilon, EMC ECS Appliance, Virtustream Storage Cloud, Amazon S3, or Microsoft Azure) as the CloudPool.

To create a cloud storage account:

Procedure

1. Run the `isi cloud accounts create` command.

The following command creates a Microsoft Azure cloud storage account. Note that this type of account requires a key provided by the cloud provider.

```
isi cloud accounts create --name=c-acct1 --type=azure
--uri=https://admin2.blob.core.windows.net --account-
username=adml
--key=!$P@$c0de998==
```

The following command creates the same account, while specifying a proxy server.

```
isi cloud accounts create --name=c-acct1 --type=azure
--uri=https://admin2.blob.core.windows.net --account-
username=adml
--key=!$P@$c0de998== --proxy myproxy1
```

After you finish

After the cloud storage account successfully connects to the cloud provider, you must add the cloud storage account to a CloudPool in OneFS. OneFS is then able to archive files to the cloud.

Edit a cloud storage account (Web UI)

You can edit some of the settings of an existing cloud storage account.

Procedure

1. Click **File System > Storage Pools > CloudPools**.
2. In the **Cloud Storage Accounts** list, click the **View/Edit** button to the right of the account that you want to edit.
3. In the **View Cloud Storage Account Details** dialog box, click the **Edit Account** button.
4. In the **Edit Cloud Storage Account Details** dialog box, perform any of these actions:
 - a. In the **Name or Alias** field, enter a new name for the account. You cannot change the type of account.
 - b. In the **URI** field, enter a fully qualified URI for the account. The URI must use the HTTPS protocol, and match the URI used to set up the account with your cloud provider.
 - c. In the **Username** field, enter the account username, which must be the same as the username provided to the cloud provider.
 - d. In the **Key** field, enter the account password. The password must be the same as the password that you provided to the cloud provider, or the key that the cloud provider issued to you.
 - e. If you want to use a different proxy server for this cloud account, select the name of the new proxy from the **Proxy** drop-down box.
 - f. If you are editing an Amazon S3 account, you can also specify a new **S3 Account ID** and **S3 Telemetry Reporting Bucket**. You cannot change the **S3 Storage Region**.
5. Click the **Save Changes** button.

CloudPools validates that your cloud data is still accessible. Otherwise, it alerts you and does not save the changes.

Modify a cloud storage account (CLI)

You can modify the information associated with a cloud storage account.

To modify a cloud storage account, you must specify the account name. You can run the `isi cloud accounts list` command to list cloud storage accounts.

Procedure

1. Run the `isi cloud accounts modify` command.

This sample command changes the name of the cloud storage account `CloudAcct3` to `CloudAcct5` and specifies a proxy server through which communications with the public cloud provider are to be managed.

```
isi cloud accounts modify CloudAcct3 --name=CloudAcct5 --
proxy myproxy1
```

Delete a cloud storage account (CLI)

You can delete a cloud storage account. However, proceed with extreme caution, as deleting an account can result in loss of access to cloud data.

You can stop archiving data to a cloud storage account without deleting it by running the `isi cloud pools modify` command and removing the account from its parent CloudPool. Previously archived files remain in cloud storage, and SmartLink files on the local cluster still point to the cloud data.

Procedure

1. Run the `isi cloud accounts delete` command.

The following command deletes the cloud storage account `OldRecords`.

```
isi cloud accounts delete OldRecords --acknowledge yes
```

In this case, OneFS responds with the following message:

```
*****
WARNING: Deleting an account is extremely dangerous.
Continuing with this operation will result in a permanent loss
of data.
Type 'confirm delete data' to proceed. Press enter to cancel:
```

2. Type the confirmation string `confirm delete data`, then press ENTER.

The cloud storage account is deleted. Although cloud data remains with your cloud provider, the data will no longer be accessible in OneFS.

List cloud storage accounts (CLI)

You can list all cloud storage accounts created on your cluster.

Procedure

1. Run the `isi cloud accounts list` command.

A sample list of cloud accounts is shown with name, type, account username, URI, and status.

View a cloud storage account (CLI)

You can view detailed information about a cloud storage account.

Procedure

1. Run the `isi cloud accounts view` command.

The following command displays account information for the `CloudAcct3` account.

```
isi cloud accounts view CloudAcct3
```

Output from the command displays the properties of the cloud storage account, including account name, type, and more.

Managing CloudPools

A CloudPool contains one or more cloud storage accounts and enables OneFS to use cloud storage as simply another tier of storage available to the cluster.

You can create, view, edit, and monitor CloudPools.

Create a CloudPool (Web UI)

You can create a CloudPool and add a cloud storage account.

Procedure

1. Click **File System > Storage Pools > CloudPools**.
2. Click the **+ Create a CloudPool** button.
3. In the **Create a CloudPool** dialog box, in the **Name** field, enter a name for the CloudPool. The name must be unique on your cluster.
4. From the **Type** drop-down menu, select a type of cloud account, one of EMC Isilon, EMC ECS Appliance, Virtustream Storage Cloud, Microsoft Azure, or Amazon S3.
5. Enter a vendor and description for the CloudPool.
6. From the **Account in CloudPool** list, select the cloud storage account that this CloudPool should contain. The list is empty until you select a value from the **Type** list. The **Account in CloudPool** list then shows only those cloud storage accounts that match that type, for example, Azure.
7. Click **Create a CloudPool**.

The dialog box closes and, in the **CloudPools** list, the new CloudPool is displayed along with its type, state, vendor, and description.

Create a CloudPool (CLI)

You can create a CloudPool and add a cloud storage account.

A CloudPool is the mechanism that connects a cloud storage account to OneFS. When you create a CloudPool, OneFS enforces two requirements:

- The CloudPool may contain only one cloud storage account.

- The cloud storage account must be of the same type as the CloudPool. For example, an Azure CloudPool may only contain an Azure cloud storage account.

Procedure

1. Run the `isi cloud pools create` command.

When you create a CloudPool, you need to provide a unique name, the CloudPool type (`isilon`, `ecs`, `virtustream`, `azure`, or `s3`), and one cloud account. The following command creates an Azure-based CloudPool:

```
isi cloud pools create cp_az azure csa_azure1 --vendor
Microsoft
```

Results

You can view the result of this operation by running the `isi cloud pools view` command with the ID (name) of the CloudPool that you created, as shown in the following example:

```
isi cloud pools view cp_az
```

The output displays the CloudPool ID, name, type, and other properties.

View information about a CloudPool (Web UI)

You can view information about a CloudPool, including the cloud storage account, vendor, type, and description.

Procedure

1. Click **File System > Storage Pools > CloudPools**.

In the **CloudPools** list, each CloudPool is represented by a blue cloud. Associated cloud accounts are listed below each CloudPool, and represented by an orange user icon. The type, vendor, and description associated with each CloudPool is displayed.

2. To further view the settings of a CloudPool, click **View/Edit** to the right of the CloudPool.

The **View Cloud Storage Pool Details** dialog box displays information about the CloudPool.

3. Click **Close** to close the dialog box.

View information about a CloudPool (CLI)

You can view information about a CloudPool, including the cloud storage account, vendor, type, and description.

Before you begin

The CloudPool must already have been created.

Procedure

1. Run the `isi cloud pools view` command.

The following command provides information on the CloudPool named `cah_s3_cp`.

```
isi cloud pools view cah_s3_cp
```

The output of this command displays the ID, name, type, and other CloudPool properties.

Modify a CloudPool (Web UI)

You can modify a CloudPool, changing the name, the account it contains, the cloud vendor, and the description.

Procedure

1. Click **File System > Storage Pools > CloudPools**.

In the **CloudPools** list, each CloudPool is represented by a blue cloud icon. The cloud account associated with each CloudPool is listed and represented by an orange user icon. The type, vendor, and description are also displayed.

2. Click **View/Edit** to the right of the CloudPool that you want to modify.

The **View Cloud Storage Pool Details** dialog box appears.

3. Click the **Edit CloudPool** button.

The **Edit CloudPool Details** dialog box appears.

4. Modify the name, vendor, or description fields, as intended.
5. From the Account in CloudPool drop-down list, select a different account of the same type.
6. Click the **Save changes** button.
7. In the **View Cloud Storage Pool Details** dialog box, click **Close**.

Results

Any changes that you made to the CloudPool are reflected in the **CloudPools** list.

Modify a CloudPool (CLI)

You can modify a CloudPool, changing the name, the account it contains, the cloud vendor, and the description.

Before you begin

To determine the available CloudPools on your system, run the `isi cloud pools view` command.

Procedure

1. Run the `isi cloud pools modify` command.

The following command modifies a CloudPool named `c_pool_azure`, removing its cloud storage account

```
isi cloud pool modify c_pool_azure --remove-accounts c_acct2
--description "Secondary archive"
```

Delete a CloudPool (CLI)

You can delete a CloudPool. However, you should proceed with caution. CloudPools provide the mechanism to connect OneFS to your cloud storage accounts. If you delete a CloudPool, the associated cloud storage account is no longer accessible.

Before you begin

Run the `isi cloud pools list` command to display the names of the CloudPools on your cluster. Run the `isi cloud pools view` command, along with the name to get information about a CloudPool.

Procedure

1. Run the `isi cloud pools delete` command.

The following command deletes the CloudPool named `c_pool_azure`.

```
isi cloud pools delete c_pool_azure
```

OneFS asks you to confirm the deletion, as follows:

```
Are you sure? (yes/[no]):
```

2. Type `yes` and press ENTER.

The CloudPool is deleted.

Monitoring CloudPools (Web UI)

You can monitor the health of CloudPools configured on your cluster.

Procedure

1. Click **File System > Storage Pools > Summary**.
2. In the **Status** list, check the status for CloudPools.

Status conditions for CloudPools are `Good` or `Needs Attention`. A status of `Needs Attention` appears when a CloudPool cannot connect to the remote cloud provider. This could indicate issues with the Internet connection or with the cloud provider. If you confirm that your Internet connection is good, contact your cloud provider for help.

Managing CloudPools settings

You can manage CloudPools default settings, including snapshot archival, encryption, compression, cache settings, data retention settings, and the ability to regenerate an encryption key. An encryption key should only be regenerated if you suspect the existing key has been compromised.

View cloud settings (CLI)

You can view the top-level settings for CloudPools.

Procedure

1. Run the `isi cloud settings view` command.

The command displays CloudPools settings such as accessibility, cache expiration, whether compression and encryption is enabled, and so on.

Modify default cloud settings (CLI)

You can modify default CloudPools settings.

You can use the `isi cloud settings view` command to display current default settings.

Procedure

1. Run the `isi cloud settings modify` command.

The following command disables archival of files that have snapshot versions.

```
isi cloud settings modify --default-archive-snapshot-files=no
```

The following command enables both encryption and compression of cloud data:

```
isi cloud settings modify --default-encryption-enabled=yes
--default-compression-enabled=yes
```

Generate a new master encryption key (CLI)

You can generate a new master encryption key. The key is used to encrypt data and is stored with cloud data objects.

Before you begin

Only generate a new master encryption key if you believe the existing key has been compromised.

Procedure

1. Run the `isi cloud settings regenerate-encryption-key` command.

The following command generates a new encryption key in verbose mode.

```
isi cloud settings regenerate-encryption-key --verbose
```

In verbose mode, the system confirms the process:

```
Encryption key has been regenerated
```


CHAPTER 3

Managing CloudPools policies

This section describes how to define and maintain the file pool policies that archive files to cloud storage.

- [CloudPools file processing](#)..... 28
- [Managing cloud policies](#)..... 34

CloudPools file processing

CloudPools archives file data to the cloud, and enables you to access or fully recall this data whenever needed.

You can create file pool policies that identify the files to be archived to the cloud. When a file pool policy that contains cloud actions is run, CloudPools moves file data to the cloud and stores it in specialized cloud data objects, collectively referred to as cloud data. File data can be encrypted and compressed before it is archived to the cloud.

In place of each file that is archived, CloudPools retains a local proxy called a SmartLink file. SmartLink files include special metadata and maps to the actual file data in the cloud.

When a user browses OneFS, typically through an SMB connection or NFS export, SmartLink files appear in place as the files they link to. When a user opens a SmartLink file, a process referred to as inline access, CloudPools retrieves file data from the cloud and caches it locally.

The user can view the file, and CloudPools continues caching as much of the data as needed by the application. If the user modifies and saves the file, the changes are also held in cache. Periodically, CloudPools scans SmartLink files for pending data changes and writes them to the appropriate objects in the cloud. In this way, the archived data is kept up to date.

You can also recall archived files from the cloud. When you do, SmartLink files are fully replaced by the recalled files.

Like any file in OneFS, SmartLink files are controlled either by the default file pool policy or by parameters included in a custom file pool policy. If you configure additional file pool policies, these policies have priority over the default file pool policy.

File pool policies contain instructions that determine how OneFS manages files across a cluster and in the cloud.

Because SmartLink files produced by CloudPools are retained on the cluster, OneFS applies file pool policies to these files, as well.

When file pool policies run, the system compares each file on the system with each file pool policy. A file can match only some aspects of a custom file pool policy (for example, SSD strategy and snapshot configuration). In this case, those aspects of file handling are governed by the custom file pool policy, and all other aspects are governed by the default file pool policy.

Refer to the *SmartPools* section in the *OneFS Web Administration Guide* for additional information about file pool policies.

Archiving files with file pool policies

You can configure a file pool policy to identify the files you want to archive to the cloud and the CloudPools actions to apply to these files.

Specifying a file pool policy, you can archive files using either the OneFS web administration interface or the command-line interface. A file pool policy that archives files to the cloud must specify the following information:

- **Files to manage:** These can be files of a certain type, files in a specified path, or files that match specified criteria, such as size, creation date, or last modified date.

- CloudPools actions: The cloud storage pool to send file data to, and whether the data should be compressed or encrypted.

Sample policies with CloudPools actions

Each file pool policy identifies a set of files and the CloudPools actions to apply to the file pool. You can identify files to be archived based on multiple criteria, including file type, size, directory path, time of file creation, time of last file access, and time of last file modification.

File-matching criteria in a file pool policy enable you to define a logical group of files referred to as a file pool. After defining a file pool, you specify CloudPools actions to perform on the files, including the cloud storage target, compression, and encryption.

For example, you might define file pool policies that specify files to be archived based on criteria similar to the following:

- Files of <type>, last accessed before <date>
- Files older than <date>, last accessed after <date>, and of <type>
- Files in <directory> that are older than <date>
- Files marked with <custom attribute>, that are older than <date>

You can specify file-matching criteria on a per-policy basis. Each file pool policy allows you to combine multiple criteria using AND statements and OR statements, providing significant flexibility and control for your workflow.

Combining cloud and local storage policy actions

You can specify both cloud and a local storage actions in the same file pool policy. The cloud actions are applied to the data of matching files, while the local actions apply to the SmartLink files that are created in place.

SmartPools settings can determine the target storage pool or tier, file protection level, I/O optimization, and data access optimization. The SmartLink files are processed according to the specified SmartPools parameters. If some settings are not specified in the custom file pool policy, the default file pool policy settings are applied to the SmartLink files.

About file pool policy order

OneFS compares all files to file pool policies in order. The first custom policy that matches a file controls how that file is handled. All other custom file pool policies in the ordered list are ignored. For any of the attributes that the matching custom policy does not specify, the value from the default policy is applied.

This makes the order of file pool policies important. If two or more file pool policies would match the same file, you must ensure that the policy order delivers your preferred file handling instructions.

After a file match with a file pool policy occurs, the system uses the settings in the matching policy to store and protect the file. However, a matching policy might not specify all settings for the match file. In this case, the default policy is used for those settings not specified in the custom policy. For each file stored on the OneFS cluster, the system needs to determine the following:

- Requested protection level
- Data storage target for local data cache
- SSD strategy for metadata and data
- Protection level for local data cache

- Configuration for snapshots
- SmartCache setting
- L3 cache setting
- Data access pattern
- CloudPools actions (if any)

If no custom policy matches a file, the default policy specifies all storage settings for the file. The default policy, in effect, matches all files not matched by any other SmartPools policy. For this reason, the default policy is the last in the file pool policy list, and always the last policy the system applies.

Files that have been archived to the cloud are always governed by the original policy.

File pool policy cloud archive parameters

CloudPools provides a specific set of file pool parameters that support archiving files to the cloud. The following table lists and describes these parameters.

Web admin parameter	CLI parameter	Description	Usage notes
CloudPool Storage Target	cloud-pool	An Isilon administrative container for a cloud storage account.	Each CloudPool can contain only one cloud storage account with a cloud provider. You must create a cloud storage account before creating and configuring a CloudPool. A CloudPool and its contained cloud storage account must be of the same type: EMC Isilon, EMC ECS Appliance, Virtustream Storage Cloud, Amazon S3, or Microsoft Azure.
Encrypt data before transfer	cloud-encryption-enabled	Specifies whether CloudPools encrypts data prior to archiving it The default value is <code>disabled</code>	Specifies whether data is encrypted prior to archiving to the cloud. Cloud data is decrypted when accessed or recalled.
Compress data before transfer	cloud-compression-enabled	Specifies whether CloudPools compresses data prior to archiving it. The default value is <code>disabled</code>	Specifies whether data is compressed prior to archiving to the cloud. Cloud data is decompressed when accessed or recalled.
Cloud Data Retention Period	cloud-data-retention	The length of time cloud files are retained after the files have been fully recalled. The default value is <code>1 week</code> .	Specifies how long cloud objects are retained after a SmartLink file has been replaced by the recalled file. When this happens, CloudPools cleans up local resources allocated for the SmartLink files, and also removes the associated cloud objects. This work is performed weekly by the cloud objects garbage collector job.

Web admin parameter	CLI parameter	Description	Usage notes
			<p>Note</p> <p>The system removes (garbage-collects) cloud objects when their SmartLink files and all local references to them have been removed. If a SmartLink file has been backed up and the original SmartLink file is subsequently deleted, associated cloud objects are deleted only after the retention time of the backed-up SmartLink file has expired.</p>
Incremental Backup Retention Period for NDMP Incremental Backup and SyncIQ	cloud-incremental-backup-retention	<p>Specifies the length of time that OneFS retains cloud data referenced by a SmartLink file that has been replicated by SyncIQ or an incremental NDMP backup.</p> <p>The default value is 5 years.</p>	<p>If a SmartLink file has been backed up and the original SmartLink file is subsequently deleted, associated cloud objects are deleted only after the retention time of the backed-up SmartLink file has expired.</p>
Full Backup Retention Period for NDMP Only	cloud-full-backup-retention	<p>Specifies the length of time that OneFS retains cloud data referenced by a SmartLink file that has been backed up by a full NDMP backup.</p> <p>The default value is 5 years.</p>	<p>If a SmartLink file has been backed up and the original SmartLink file is subsequently deleted, associated cloud objects are deleted only after the original retention time, or a longer incremental or full backup retention period, has expired.</p>
Writeback Frequency	cloud-writeback-frequency	<p>Specifies the interval at which the system writes the data stored in the cache of SmartLink files to the cloud.</p> <p>The default value is 9 hours</p>	<p>Specifies how often SmartLink files modified on the cluster are written to their associated cloud data objects.</p>
Accessibility	cloud-accessibility	<p>Specifies how data is cached in SmartLink files when a user or application accesses a SmartLink file on the cluster. Values are <code>cached</code> and <code>no-cache</code>.</p> <p>The default value is <code>cached</code></p>	<p>Determines whether cloud data is cached when a file is accessed on the local cluster.</p> <p>cached</p> <p>When <code>cached</code> is selected, accessed cloud data is cached to the SmartLink file on read or write access.</p> <p>no-cache</p> <p>When <code>no-cache</code> is selected, the system does not cache data in the SmartLink files on read access, but passes it through to the local accessing application. If</p>

Web admin parameter	CLI parameter	Description	Usage notes
			you write to data accessed when this setting applies, the system caches your changes. Choose <code>no-cache</code> if you want to limit the use of cluster resources.
Cache Read Ahead	<code>cloud-readahead</code>	Specifies the cache readahead strategy for cloud files (one of <code>partial</code> or <code>full</code>) The default value is <code>partial</code>	Specifies whether cloud data is fully or partially recalled when you access a SmartLink file on the cluster. If <code>partial</code> is specified, the system only recalls the file blocks needed when a SmartLink file is accessed. If <code>full</code> is specified, all cloud data is fully cached when the SmartLink file is accessed.
Cache Expiration	<code>cloud-cache-expiration</code>	Specifies the number of days until the system purges expired cache information in SmartLink files. The default value is <code>1 day</code> .	Specifies how long the system retains cloud data that has been recalled in the cache of associated SmartLink files. The system purges the SmartLink file cache of data that has not been accessed for the number of days specified.
Archive Files with Snapshots	<code>cloud-archive-snapshot-files</code>	Specifies whether the policy should archive files with snapshots. The default is on.	Snapshots capture files at a moment in time. If a policy matches a file that is included in a snapshot, the file can be archived or not depending on this setting.

File matching options for cloud archival policies

Each file pool policy must provide match criteria to identify the files to archive and the cloud target where the files should be stored.

The following table describes the match criteria to use when creating file pool policies.

Match criteria		Description
Web admin interface	Command line interface	
Filename	<code>--name</code>	Includes or excludes files based on the file name. You can specify whether to include or exclude full or partial names that contain specific text. Wildcard characters are supported.
Path	<code>--path</code>	Includes or excludes files based on the file path. You can specify whether to include or exclude full or partial paths that contain specified text. You can also include the wildcard characters <code>*</code> , <code>?</code> , and <code>[]</code> .
File Type	<code>--file-type</code>	Includes or excludes files based on one of the following file-system object types: <ul style="list-style-type: none"> Regular file Directory

Match criteria		Description
Web admin interface	Command line interface	
		<ul style="list-style-type: none"> Other
File Attribute	--custom-attribute	Includes or excludes files based on a custom user-defined attribute.
Modified	--changed-time	Includes or excludes files based on when the file was last modified. You can specify a relative date and time, such as "older than 2 weeks," or a specific date and time, such as "before January 1, 2012." Time settings are based on a 24-hour clock.
Accessed	--accessed-time	<p>Includes or excludes files based on when the file was last accessed. You can specify a relative date and time, such as "older than 2 weeks," or a specific date and time, such as "before January 1, 2012." Time settings are based on a 24-hour clock.</p> <hr/> <p>Note</p> <p>Because it affects performance, access time tracking as a file pool policy criterion is disabled by default.</p> <hr/>
Metadata Changed	--metadata-changed-time	<p>Includes or excludes files based on when the file metadata was last modified. This option is available only if the global access-time-tracking option of the cluster is enabled.</p> <p>You can specify a relative date and time, such as "older than 2 weeks," or a specific date and time, such as "before January 1, 2012." Time settings are based on a 24-hour clock.</p>
Created	--birth-time	Includes or excludes files based on when the file was created. You can specify a relative date and time, such as "older than 2 weeks," or a specific date and time, such as "before January 1, 2012." Time settings are based on a 24-hour clock.
Size	--size	<p>Includes or excludes files based on their size.</p> <hr/> <p>Note</p> <p>File sizes are represented in multiples of 1024, not 1000.</p> <hr/>

Retrieving file data from the cloud

You can retrieve file data from the cloud either by inline access through a supported protocol (SMB, NFS, Swift, or HDFS), or by fully recalling files.

Inline access of cloud data

Users can retrieve file data in the cloud by accessing a SmartLink file on the local cluster through a supported protocol. This method is referred to as inline access.

When the user opens a SmartLink file, for example, through an SMB share, CloudPools retrieves and locally caches file data from the cloud. The amount of data that is cached is determined by the CloudPools **Cache Read Ahead** setting.

If the user makes changes to the file, CloudPools maintains those changes in cache and periodically updates file data to the cloud so that the latest version is always archived.

Recalling files from the cloud

You can fully recall a file from cloud storage. In this case, CloudPools restores the full file to the cluster and overwrites its associated SmartLink file. As part of a daily maintenance routine, CloudPools also completely removes the recalled file data from the cloud.

You can recall files from cloud storage only with the CLI command `isi cloud recall`. You can recall files individually by name or by specifying a fully recursive directory path.

Note

When you use the `isi cloud recall` command to recall a file from cloud storage, the full file is restored to its original directory. If the file pool policy that originally archived the file to the cloud is still in effect, the next time the SmartPools job runs, the recalled file is archived to the cloud again. If you do not want the recalled file to be re-archived, you can move the file to a different directory that would not be affected by the file pool policy, or you can modify or delete the policy.

Managing cloud policies

CloudPools takes advantage of the SmartPools infrastructure, and applies file pool policies to determine which files are to be archived to the cloud.

Consequently, you must activate both a SmartPools and a CloudPools license to store data in the cloud.

By defining file pool policies, you can have OneFS automatically archive files to the cloud when they match certain characteristics, such as age, size, type, or location.

File pool policies are both for SmartPools and CloudPools purposes. A file pool policy can specify a local storage target, a cloud storage target, or both. If you create a policy that specifies both local and cloud targets, the policy moves file data to the cloud, and applies the local settings to the SmartLink files retained on the local cluster.

If the purpose of a file pool policy is to move files to a local node pool or tier, do not configure a cloud target. Conversely, if the purpose of a policy is to archive files to the cloud, configuring a local target, while allowed, is not necessary. In this case, the system uses the settings of the default file pool policy to store the local SmartLink files.

Create a file pool policy for cloud storage (Web UI)

You can create file pool policies that specify CloudPools actions to be applied to selected files.

Procedure

1. Click **File System > Storage Pools > File Pool Policies**.
2. Click the **+ Create a File Pool Policy** button.

The **Create a File Pool Policy** dialog box displays.

3. Enter a policy name and, optionally, a description.

4. In the **Select Files to Manage** area, use the pull-down menus to specify the file selection criteria for cloud storage. The criteria you specify are used by OneFS to determine the files to archive. The criteria you specify for file selection can include the following attributes, combined with Boolean operators:
 - Filename
 - Path
 - File Type
 - File Attribute
 - Modified
 - Accessed
 - Metadata Changed
 - Created
 - Size
5. In the **Apply CloudPools Actions to Selected Files** area, select **Move to cloud storage**.
6. In the **CloudPool Storage Target** drop-down menu, select an existing CloudPool, and specify whether to encrypt and compress data before it is archived to the cloud.
7. Click **Show Advanced CloudPool Settings** to specify additional cloud storage options, as described in the following table:

Setting	Description
Data Retention Settings	<ul style="list-style-type: none"> • Cloud Data Retention Period: Specifies how long cloud objects are retained after a SmartLink file has been deleted. When a SmartLink file is deleted on the local cluster, CloudPools cleans up local resources allocated for the SmartLink files, and also removes the associated cloud objects. This work is performed weekly by the cloud objects garbage collector job. • Incremental Backup Retention Period for NDMP Incremental Backup and SyncIQ: Specifies the backup retention period for SmartLink files created with an incremental NDMP backup policy or a SyncIQ policy. This value is only relevant for incremental NDMP backups and SyncIQ replication. • Full Backup Retention Period for NDMP Only: Specifies the backup retention period for SmartLink files created with a full NDMP backup policy. This value is only relevant for full NDMP backups.
Accessibility and Cache Settings	<ul style="list-style-type: none"> • Writeback Frequency: Specifies how often SmartLink files edited on the Isilon cluster are written to their associated cloud data objects. • Accessibility: Specifies whether or not to cache archived data locally. Local caching enables faster access of cloud data, but reduces space saving on your cluster. • Cache Read Ahead: Specifies whether cloud data is fully or partially cached when you access a SmartLink file on the

Setting	Description
	<p>local file system. If the policy specifies partial caching, the system only caches the blocks needed when a file is accessed. If the policy specifies full caching, cloud data is fully cached when the SmartLink file is accessed.</p> <ul style="list-style-type: none"> • Cache Expiration: Specifies the amount of time cloud data that has been cached is retained in the local cache of associated SmartLink files. The system purges the SmartLink cache of data that has not been accessed for the number of days specified.
Archive files with snapshots	<p>Specifies that CloudPools is to archive files that have one or more snapshot versions. The latest versions of such files are archived, and SmartLink files are left in their place on the local cluster. Because the full snapshot versions of files are retained on the cluster, archiving these files does not save space on the cluster.</p>

8. Click **Create Policy**.

The file pool policy appears under File Pool Policies in the **File Pool Policies** window.

Results

The next time the SmartPools system job is run, the file pool policy performs the specified actions.

Create a file pool policy for cloud storage (CLI)

You can create file pool policies that specify CloudPools actions to be applied to selected files.

Procedure

1. Run the `isi filepool policies create` command.

The following command creates a file pool policy with the name `archive` and the CloudPool storage target, `S3_pool`. The command also specifies a file-matching pattern to archive all files in a directory path that have not been accessed after November 30, 2014.

```
isi filepool policies create archive --cloud-pool=S3_pool
--begin-filter --name="*.*" --and --path="/ifs/home/users"
--and --accessed-time=2014-11-30 --operator=lt --end-filter
```

Modify cloud attributes in a file pool policy (Web UI)

You can modify a file pool policy. Each file pool policy for cloud archival specifies a file-matching pattern and the actions to perform on the matched files (file pool).

Procedure

1. Click **File System > Storage Pools > File Pool Policies**.

The **File Pool Policies** page appears.

2. In the File Pool Policies list, next to the file pool policy you intend to modify, click **View/Edit**.

The **View File Pool Policy Details** dialog box appears.

3. Click **Edit Policy**.

The **Edit File Pool Policy Details** dialog box appears.

4. Make your changes in the appropriate areas and click **Save Changes**.

Results

Changes to the file pool policy are applied the next time the SmartPools system job runs.

Modify cloud attributes in a file pool policy (CLI)

You can modify a file pool policy. Each file pool policy for cloud archival specifies a file-matching pattern and the actions to perform on the matched files (file pool).

You can run the `isi filepool policies list` command to list available file pool policies.

Procedure

1. Run the `isi filepool policies modify` command.

The following example modifies the file-matching pattern in a file pool policy named `my_policy`.

```
isi filepool policies modify my_policy --begin-filter
--name="*.jpg" --and --accessed-time=2013-08-01 --operator=lt
--end-filter
```

List file pool policies (CLI)

You can list all file pool policies stored in OneFS.

Procedure

1. Run the `isi filepool policies list` command.

View details of a file pool policy (CLI)

You can display detailed information about a file pool policy.

To list all available file pool policies, you can run the `isi filepool policies list` command.

Procedure

1. Run the `isi filepool policies view` command.

The following command displays information about the policy `my_policy`, including status, associated CloudPool, whether encryption and compression are enabled, and more.

```
isi filepool policies view my_policy
```

Apply a file pool policy to a specified file or path (CLI)

You can apply a file pool policy to specified files or directories manually, rather than waiting for the SmartPools job to run.

For `isi filepool apply` to execute, the file or directory specified must match one of the defined file pool policies.

Procedure

1. Run the `isi filepool apply` command.

The following command applies the appropriate file pool policy to all files and subdirectories in a given path.

```
isi filepool apply --path=/ifs/data/images --recurse
```

Archive files directly to the cloud (CLI)

You can archive specific files directly to the cloud. To enable this, CloudPools must match these files to an existing file pool policy.

Before you begin

A custom file pool policy that matches the specified file or files and points to cloud storage must already exist on your system.

Procedure

1. Run the `isi cloud archive` command.

The following command specifies a directory and all of its subdirectories and files to be archived:

```
isi cloud archive /ifs/data/shared/images/*. * --recursive yes
```

CHAPTER 4

Managing CloudPools with other OneFS functions

CloudPools is designed to work seamlessly with other OneFS functions, including data encryption and compression, SMB and NFS support, SyncIQ, snapshots, and NDMP backup and recovery. This section includes the following topics:

- [Compression and encryption of cloud data](#)..... 40
- [CloudPools protocol support](#).....40
- [SyncIQ interoperability](#)..... 41
- [NDMP backup and restore of cloud data](#).....46
- [CloudPools and snapshots](#).....47
- [CloudPools and SmartLock](#)..... 47
- [CloudPools and SmartQuotas](#)..... 48
- [CloudPools and SmartDedupe](#).....48

Compression and encryption of cloud data

You can specify compression and encryption of data that is moved to the cloud.

With CloudPools, you can enable compression and encryption on a per-policy basis. Both encryption and compression are disabled by default.

Files encrypted or compressed when stored in the cloud are automatically decrypted and decompressed when data is cached (inline access) or the file is recalled from the cloud to local storage.

CloudPools uses a master encryption key to encrypt the data encryption keys. Encryption applies to both the SmartLink file and the file data archived to the cloud. Both the SmartLink file and the archived data include encrypted copies of the data encryption keys. After a file is encrypted, it can only be decrypted by recalling it.

CloudPools keeps track of the encryption status of SmartLink files in snapshots and referenced data in the cloud. If SmartLink files in snapshots are unencrypted and refer to unencrypted cloud objects, the SmartLink files in the snapshots remain unencrypted even if you create a new CloudPools policy that encrypts the latest version of the file.

OneFS stores the master encryption key in the local key management system. You can generate a new version of the key if you believe the key has been compromised. If regenerated, the new master key secures new data written to the cloud. Previously written data is secured by the old data encryption keys, resident in the local SmartLink files.

Note

CloudPools works seamlessly with nodes that are equipped with self-encrypting drives (SEDs). CloudPools can apply encryption to files that are archived to cloud storage. Similarly, any SmartLink files left on SEDs are handled like any other file.

CloudPools protocol support

CloudPools supports inline access of cloud data through SMB, NFS, and other file system protocols. Consequently, users who access files on an Isilon cluster from other systems can also access files stored in the cloud.

NFS inline access

CloudPools enables access of SmartLink files from NFS exports.

When a user connects to a cluster through an NFS export, and browses the file system, SmartLink files appear to be the original files. When the user opens a SmartLink file, CloudPools retrieves and caches the original file data from the cloud. Depending on the **Cache Read Ahead** setting, either a portion of the file data, or the entire file, is cached.

If the user modifies the file, CloudPools stores the changes in the cache and periodically writes the changes back to the cloud. In this way, cloud data is kept fully up to date.

SMB inline access

CloudPools enables access of SmartLink files from SMB shares.

When a user connects to a cluster through an SMB share, and browses the file system, SmartLink files appear to be the original files. When the user opens a SmartLink file, CloudPools retrieves and caches the original file data from the cloud. Depending on the **Cache Read Ahead** setting, either a portion of the file data, or the entire file, is cached.

If the user modifies the file, CloudPools caches the changes and periodically writes the changes back to the cloud. In this way, the cloud data is kept fully up to date.

Other protocols supporting inline access

Cloud data can be accessed programmatically through Swift and HDFS by accessing the associated SmartLink files.

We recommend that you consult with your Isilon representative for more information on these use cases.

SyncIQ interoperability

SyncIQ enables you to synchronize data from your Isilon primary (source) cluster to a secondary (target) cluster. If your primary cluster becomes unavailable, you can fail over to the secondary cluster, and users can continue to access data, including data stored in the cloud.

During SyncIQ replication, all files, including SmartLink files, are copied from the source cluster to the target cluster. Users given access to the target cluster through supported protocols can retrieve cloud data or fully recall files from the cloud. In these cases, CloudPools retrieves and caches data (inline access) or recalls the full file exactly as it would from the original source cluster.

Unless you specifically grant cloud write access to the secondary cluster, CloudPools stores any changes to SmartLink files in the local cache, which is limited only by available space on the cluster.

SyncIQ policies

CloudPools supports SyncIQ replication of SmartLink files to one or more target clusters. SyncIQ can also be used to restore backed up SmartLink files to their original (source) cluster.

The two types of SyncIQ policies are synchronization policies and copy policies. These policies can be run manually, or configured to run automatically, based on policy settings.

CloudPools supports both types of SyncIQ policy. When SyncIQ replicates SmartLink files to a target cluster, secondary information associated with a SmartLink file, such as local cache state and unsynchronized cache data, is also replicated.

If your source (primary) cluster goes down or is unavailable for any reason, and you fail over to your secondary cluster, users can continue to access SmartLink files and, therefore, cloud data, as they would normally.

If the failover is temporary and you plan to restore your source cluster to full operation, you do not need to enable cloud write access on the secondary cluster. Any

changes that users make to SmartLink files are stored in the local cache, which is limited only by the amount of free space on your cluster.

When you fail back to your source cluster, and restore updated SmartLink files, only then will CloudPools write the cached modifications back to the cloud.

⚠ CAUTION

If your failover is long-term or permanent, you can grant cloud write access to the secondary cluster by using the `isi cloud access add` command in the CLI. However, you should only ever allow cloud write access on one cluster at a time. Otherwise, corruption of cloud data can occur.

CloudPools cloud data retention time

In CloudPools, cloud data is governed by a concept referred to as retention. Retention time defines an absolute time for cloud data to remain in cloud storage after the related SmartLink file has been deleted.

When you use CloudPools to archive a file from your cluster to cloud storage, a SmartLink file is created on the Isilon cluster in place of the archived file. As long as data from a file remains in the cloud, the SmartLink file remains in place to represent and point to the cloud data.

Because a SmartLink file can be backed up to tape through NDMP and replicated to another Isilon cluster through SyncIQ, more than one SmartLink file can be pointing to the same cloud data at the same time.

However, importantly, cloud data can be deleted even if there is a related SmartLink file on another cluster that was restored by a SyncIQ or NDMP process. Cloud data is maintained until the minimum duration of its retention times on all clusters has expired.

Retention time is affected by the following archive policy retention periods:

- **Cloud Data Retention Period** specifies the retention time of cloud data beyond the time when an associated local Smartlink file is deleted. The default setting is one week.
- **Backup Retention Period for NDMP Incremental Backup and SyncIQ** specifies the retention time of cloud data whose SmartLink file has been backed up by an incremental NDMP backup, or replicated by a SyncIQ operation. If a local SmartLink file is deleted, the SmartLink file copy can be restored, and cloud data can still be accessed. The default setting is five years.
- **Full Backup Retention Period for NDMP Only** is the retention time of cloud data whose SmartLink file has been backed up by a full NDMP backup only. If a local SmartLink file is deleted, the SmartLink file copy can be restored from the backup, and cloud data can still be accessed. The default setting is five years.

Managing SmartLink files that have been replicated

If you modify or delete a SmartLink file that has been replicated in a SyncIQ operation, CloudPools can modify or delete its associated cloud objects.

If you modify a SmartLink file on the primary cluster, changes are cached and, depending on the **Writeback Frequency** setting, are periodically written back to the cloud. In this way, cloud data is always kept up to date.

If you modify a SmartLink file on a secondary cluster, because the primary cluster is temporarily unavailable, changes will remain in cache. When you fail back to the

primary cluster, only then are changes able to be written back to the cloud according to the **Writeback Frequency** setting.

Similarly, when a SmartLink file that has been replicated in a SyncIQ operation is deleted, CloudPools appropriately manages the associated cloud data. Two retention periods can affect the cloud objects associated with a SmartLink file that has been replicated: the **Cloud Data Retention Period** and the **Incremental Backup Retention Period for NDMP Incremental Backup and SyncIQ**.

When a SmartLink file has been replicated to a secondary cluster, and is then deleted from the primary cluster, CloudPools uses both the **Cloud Data Retention Period** and the **Incremental Backup Retention Period for NDMP Incremental Backup and SyncIQ** settings to determine when the associated cloud objects should be deleted. CloudPools uses the longer of the two durations to determine when to delete cloud data.

For example, if the longer of the two retention periods is the **Incremental Backup Retention Period for NDMP Incremental Backup and SyncIQ** setting, then CloudPools uses that setting to determine when to delete cloud data after its associated SmartLink file is deleted.

If you delete a SmartLink file on a secondary cluster, because the primary cluster is temporarily unavailable, the deleted state will remain in cache. When you fail back to the primary cluster, CloudPools deletes the SmartLink file, and uses the retention settings to determine when to delete the associated cloud data.

SyncIQ deep copy

You can create a SyncIQ policy that restores full files rather than SmartLink files when copying data from the primary (source) cluster to a secondary (target) cluster.

When you create a SyncIQ policy, you can modify the **Deep Copy for CloudPools** setting. The default setting is `Deny`, which means that, during a SyncIQ operation, SmartLink files are replicated to the target cluster.

Alternatively, you can select either the `Allow` or `Force` option for deep copy. When you select `Allow`, SyncIQ still replicates SmartLink files to the target cluster unless there is a SmartLink version mismatch, in which case the full file data is retrieved from the cloud and replicated.

When you specify `Force` for deep copy, CloudPools retrieves and copies full file data from the cloud for all SmartLink files affected by the SyncIQ policy, and replicates the full files to the target cluster.

Note

A SyncIQ operation that forces deep copy can take significantly more time and consume more system resources. We recommend that you not specify deep copy unless you have a specific reason to do so. For example, if you are backing up data from the primary cluster to a secondary cluster that is running an older (pre-8.0) version of OneFS, then you should use deep copy. If you are unsure whether to use deep copy, contact your EMC Isilon representative for guidance.

Configuring access to cloud data from a secondary cluster

You can make cloud data available on a secondary cluster if your primary cluster becomes unavailable.

To configure such access, you need to have replicated the primary cluster's data on a secondary cluster using SyncIQ. Alternatively, you can restore an NDMP backup of the data to a secondary cluster.

The secondary cluster must have active SyncIQ, SmartPools, and CloudPools licenses.

With SyncIQ, when failover to a secondary cluster is required, two use cases are supported: short-term failover versus long-term failover.

In the short-term failover use case, the intention is to restore and failback to the primary cluster as quickly as possible. The secondary cluster is a temporary solution, enabling users to open SmartLink files from supported protocols and access cloud data as usual. Instead of writing any changes back to the cloud, however, CloudPools caches these changes locally in the SmartLink files on the secondary cluster. After the primary cluster is restored to service, changes from the secondary cluster must be written back to the primary cluster to enable any cached data in SmartLink files to be written back to cloud storage.

In a long-term failover situation, in which the primary cluster is to be out of service for an extended period, or is to be decommissioned entirely, other considerations become important. In this scenario, because only one cluster can have write access to cloud storage, you need to transfer write access to the failover cluster. From a CloudPools perspective in this scenario, the failover cluster effectively becomes the primary cluster.

With the NDMP approach, however, the short-term failover scenario is less practical. The secondary cluster should be given cloud write access to enable any cached modifications to SmartLink files to be written back to cloud storage. The alternative would be to somehow write modified SmartLink files back to the primary cluster after it is restored to service, but this might be more time-consuming.

CAUTION

Never allow write access to cloud data from more than one cluster at a time because it can result in data corruption.

List available clusters (CLI)

You can list the GUIDs of clusters that are accessible for SyncIQ failover or restore operations.

Available clusters are the primary cluster and any other clusters to which data has been replicated with SyncIQ or restored with NDMP.

Procedure

1. Run the `isi cloud access list` command.

OneFS displays a list of the GUIDs of eligible clusters.

Allow access to cloud data on a secondary cluster (CLI)

You can enable write access of SmartLink files and their cloud data on a secondary (target) cluster. Data from the primary (source) cluster must have been replicated to or restored on the target cluster by a SyncIQ or NDMP process.

Before you begin

Run the `isi cloud access list` command to display available target clusters. The target cluster must also have both a SmartPools and CloudPools license.

CAUTION

This capability is offered to work around cases where the primary cluster is going to be unavailable for an extended period. Never allow write access to cloud data from more than one cluster at a time because it can result in data corruption. Before allowing another cluster to have cloud write access, make sure that cloud write access is removed from the primary cluster, and that the primary cluster is offline.

Procedure

1. Run the `isi cloud access add` command.

The following command enables a secondary cluster, identified by GUID, to have write access to cloud data.

```
isi cloud access add ab9dd991261e11e382240800200c9a66
```

You can also specify the secondary cluster by name:

```
isi cloud access add isi2
```

If you know when the primary cluster will be restored to service, you can set an expiration date. The following command adds an expiration date of December 1, 2015:

```
isi cloud access add isi2 --expiration-date 12012015
```

Remove access to data from a source cluster (CLI)

You can remove previously granted access to SmartLink files replicated from a source cluster or restored from an NDMP backup. This operation must be done on the target cluster.

Before you begin

CAUTION

Never allow write access to cloud data from more than one cluster at a time because it can result in data corruption. Removing access to SmartLink files on a target (secondary) cluster is typically done when the source (primary) cluster has been or is about to be restored to service.

Procedure

1. Run the `isi cloud access remove` command on the target cluster. The following command removes SmartLink file access from a target cluster by specifying its GUID.

```
isi cloud access remove ab9dd991261e11e382240800200c9a66
```

NDMP backup and restore of cloud data

You can perform NDMP backup and restore operations on data that has been archived to the cloud.

Backup and restore capabilities with CloudPools data include:

- Archive SmartLink files when backing up from a cluster
- Restore data, including SmartLink files, to the same cluster
- Restore data, including SmartLink files, to another cluster

With NDMP backup, by default, CloudPools supports backup of SmartLink files only. No cloud data is included in the backup. Secondary information such as account information, local cache state, and unsynchronized cache data associated with the SmartLink file is also backed up.

However, you can force NDMP backup to store a full copy of file data rather than SmartLink files. This is sometimes referred to as the deep copy option. You specify deep copy by setting the `BACKUP_OPTIONS` environment variable to `0x00000100`.

In CloudPools settings, you can set three retention periods that affect backed up SmartLink files and their associated cloud data:

- Full Backup Retention Period for NDMP takes effect when the SmartLink file is backed up as part of a full backup. The default is five years.
- Incremental Backup Retention Period for Incremental NDMP Backup and SyncIQ takes effect when a SmartLink file is backed up as part of an incremental backup. The default is five years.
- Cloud Data Retention Period defines the duration that data in the cloud is kept when its related SmartLink file is deleted. The default is one week.

CloudPools ensures the validity of a backed-up SmartLink file within the cloud data retention period. It is important for you to set the retention periods appropriately to ensure that when the SmartLink file is restored from tape, it remains valid. CloudPools disallows restoring invalid SmartLink files.

To check whether a backed-up SmartLink file is still valid, CloudPools checks the retention periods stored on tape for the file. If the retention time is past the restore time, CloudPools prevents NDMP from restoring the SmartLink file.

CloudPools also makes sure that the account under which the SmartLink files were originally created has not been deleted. If it has, both NDMP backup and restore of SmartLink files will fail.

Checking the version of SmartLink files

During an NDMP backup session, version data for CloudPools SmartLink files is included in the backup stream.

When restoring data, a version check is performed on the SmartLink files. If the version check determines that the SmartLink files are incompatible with the operating

system version running on the target cluster, the NDMP restore session does not restore the SmartLink files to the target cluster and reports the version incompatibilities in the NDMP log.

CloudPools and snapshots

With CloudPools, you can control whether files that have snapshot versions are archived to the cloud.

The SnapshotIQ, SyncIQ, FSAnalyze, and NDMP Backup functions create point-in-time snapshots of directories in OneFS. Even as files are modified, the snapshot versions are maintained. As part of file matching, CloudPools can include files that have snapshot versions.

CloudPools archives the latest versions of those files to the cloud, and creates local SmartLink files in place of the archived files. As long as the snapshots remain, however, OneFS also maintains the full-data versions of these files in snapshots.

The default CloudPools setting is to allow files with snapshot versions to be archived, but you can change the default setting.

Note

Allowing archival of files that have snapshot versions does not result in space savings on your local cluster until the snapshots are deleted. Some functions, such as FSAnalyze, SyncIQ, and NDMP Backup, automatically delete snapshots when they are no longer needed. However, snapshots created with SnapshotIQ are not automatically deleted. To free up space on your cluster, we recommend that you regularly delete older snapshots that you no longer need.

CloudPools also supports SnapRevert for SmartLink files. For example, suppose that CloudPools archived a directory named `/ifs/data/images` to the cloud. The files in the `images` directory would be replaced with SmartLink files.

If you create a SnapRevert domain for the directory, and run the SnapRevert job, the CloudPools archival process is reversed, and the original files are restored to the directory. CloudPools removes any cloud data that was created as part of the original archive process.

CloudPools and SmartLock

Normally, CloudPools archives files to a cloud storage repository and leaves SmartLink files in their place on local storage. However, CloudPools is not designed to be used with SmartLock domains, in which files are committed to a WORM (write-once, read-many) state.

When you create file pool policies to archive certain files to cloud storage, make sure that these policies avoid SmartLock domains and the WORM files they include.

If you inadvertently attempt to archive WORM files to the cloud, CloudPools will generate a "file system read-only" error. Although it is possible to archive a file to the cloud, then add the resulting SmartLink file to a WORM domain, with or without committing it to a WORM state, this is not recommended.

CloudPools and SmartQuotas

The administrator can enforce storage limits for users with SmartQuotas. In this case, users should be aware that recalling data from the cloud could potentially cause them to exceed those limits.

When CloudPools archives files to cloud storage, CloudPools creates SmartLink files on local storage in place of the archived files. SmartLink files typically take up considerably less storage space than the archived files they replace.

When users recall archived files from the cloud, the full files replace the SmartLink files in local storage. This could potentially cause users to exceed their quotas. For example, suppose a user's quota is 500 MB, and files older than six months are archived to the cloud. This saves the user 250 MB of space, as the SmartLink files take up relatively little local storage space. In the meantime, the user has added more files and now has 400 MB of data in local storage. Should the user recall files from the cloud that would take up more than 100 MB of storage, the user would exceed the quota.

As a storage administrator, you should make your users aware of this possibility and how best to mitigate the issue.

CloudPools and SmartDedupe

SmartDedupe scans the OneFS file system for files that contain identical blocks of data. If SmartDedupe finds duplicate blocks, SmartDedupe moves a single copy of the blocks to a hidden file called a shadow store. SmartDedupe then deletes the duplicate blocks from the original files and replaces the blocks with pointers to the shadow store.

CloudPools interacts as follows with SmartDedupe:

- If a file pool policy specifies that de-duplicated files should be archived to cloud storage, CloudPools archives those de-duplicated files and leaves SmartLink files in their place in local storage.
- When an archived file that had been de-duplicated is recalled from the cloud, the SmartLink file is replaced and the recalled file placed back in local storage is no longer de-duplicated.
- SmartDedupe does not de-duplicate SmartLink files.

CHAPTER 5

CloudPools tips and troubleshooting

This section provides best practices and other advanced information about CloudPools.

- [CloudPools best practices](#).....50
- [Managing cloud jobs](#)..... 50
- [CloudPools troubleshooting](#)..... 52

CloudPools best practices

For best results using CloudPools, follow these best practices.

Use time stamps for cloud data archival and recall

Use time matching patterns (creation, modification, last access) when you archive data to and recall data from the cloud. This enables more efficient archival and recall operations, therefore better performance.

When you create a file pool policy for archiving data to the cloud, several of the file-matching criteria involve time:

- Created
- Accessed
- Modified

Therefore, you can specify file-matching criteria that specify when the files were created, when files were last accessed, or when they were last modified.

CloudPools can also more efficiently recall files based on time stamps.

CloudPools archiving and file size

You can gain the most benefit from CloudPools, in terms of freeing up storage space on your cluster, by archiving larger files. Archiving small files provides less, if any, benefit.

One of the benefits of archiving files to the cloud with CloudPools is how quickly you can recall these files when needed.

To enable fast recall, CloudPools creates a SmartLink file for every file whose data is archived to the cloud. SmartLink files each contain a map to the data in the cloud, meta data, and cache space. SmartLink files are generally small in size, but can grow if data is cached through inline access.

Therefore, if you archive small files to the cloud, SmartLink files are left in their place on the cluster, and could approach, or even exceed, the size of the original file.

Create exclusive accounts for CloudPools purposes

You should create an account with your cloud provider that is exclusively for CloudPools use. This prevents conflicts that might lead to data corruption or loss.

If your organization accesses cloud provider accounts outside of OneFS CloudPools operation, users must be careful not to in any way access or change data archived by CloudPools. Any such data access or modification would likely corrupt the data and compromise data retrieval and recall from CloudPools.

To prevent this, create an account in CloudPools that is exclusively for CloudPools use. Use entirely separate accounts for other cloud applications with your cloud provider.

Managing cloud jobs

You can monitor and manage two types of cloud jobs: system jobs that are always running in the background, and manual jobs that are created with the `isi cloud jobs archive` and `isi cloud jobs recall` commands. OneFS enables you to

monitor the status of both job types, and to monitor and manage your manual archive and recall jobs, as needed.

View a list of cloud jobs (CLI)

You can list all CloudPools jobs. Both CloudPools system jobs and manual jobs are listed.

CloudPools system jobs are always running to service caching and clean-up (garbage collection) processes. CloudPools manual jobs include archive jobs specified in file pool policies, and recall jobs started from the OneFS command-line interface. Each job is listed by ID, description, state, and type.

Procedure

1. Run the `isi cloud jobs list` command.

Output from the command lists CloudPools job ID, description, status, and type.

View a cloud job (CLI)

You can view information about a CloudPools job.

Before you begin

You need to know the ID of the job you want to view. You can run the `isi cloud jobs list` command to see the IDs for all cloud jobs.

Procedure

1. Run the `isi cloud jobs view` command.

The following command views information about a job with the ID of 63.

```
isi cloud jobs view 63
```

Pause a cloud job (CLI)

You can pause a running CloudPools job. This operation is typically done only for troubleshooting purposes.

Before you begin

To pause a job, you need to know the ID of the job. Run the `isi cloud jobs list` command to see a list of all cloud job IDs.

Procedure

1. Run the `isi cloud jobs pause` command.

The following command pauses a job with the ID of 63.

```
isi cloud jobs pause 63
```

This command pauses all running archive jobs:

```
isi cloud jobs pause archive
```

Note

Only currently running archive jobs are paused. Any subsequent jobs that are kicked off by a file pool policy, or manually through `isi cloud archive` are not paused and will run.

Resume a paused cloud job (CLI)

You can resume a cloud job that has been paused.

Before you begin

To resume a job, you need to know the ID of the job. Run the `isi cloud jobs list` command to see a list of all cloud job IDs.

Procedure

1. Run the `isi cloud jobs resume` command.

The following command resumes a job with the ID of 63.

```
isi cloud jobs resume 63
```

Cancel a cloud job (CLI)

You can cancel a running CloudPools job.

Before you begin

To cancel a job, you need to know the ID of the job. Run the `isi cloud jobs list` command to see a list of all cloud job IDs.

Procedure

1. Run the `isi cloud jobs cancel` command.

The following command cancels a job with an ID of 63.

```
isi cloud jobs cancel 63
```

CloudPools troubleshooting

If you encounter problems using CloudPools, refer to the information provided in this section before contacting customer support.

CloudPools limitations and expected behaviors

During normal CloudPools operation, you should be aware of the following limitations and expected behaviors.

Rolling upgrade before CloudPools usage

If you are performing a rolling upgrade to the new OneFS version, make sure the upgrade is fully complete before activating CloudPools.

Cloud storage account deletion

Warning: Do not delete a cloud storage account that is in use by archived files. This can lead to data being lost or unavailable for the archived files that use that

account. Any attempt to open SmartLink files associated with a deleted account will fail with I/O error messages. In addition, NDMP backup and restore and SyncIQ failover and failback will fail when a cloud storage account has been deleted. If, through inline access, an NFS or SMB user attempts to open a SmartLink file, and receives an I/O error, this can mean that the related cloud storage account has been deleted. We recommend trying inline access of other SmartLink files in the same CloudPool. If the same error is generated for those files, the cloud storage account has been deleted and data is lost. If the other SmartLink files are accessible, the SmartLink file that generated the error might be corrupted. Either way, you should contact EMC Isilon Support for assistance.

Accessing SmartLink files

You can view and modify cloud data by accessing SmartLink files through supported protocols (NFS, SMB, Swift, and HDFS). If you attempt to access a SmartLink file directly on the cluster, through a command running on the cluster, or through an unsupported protocol, an I/O error is generated. Local access for SmartLink files is blocked for these commands: `tar`, `gzip`, `scp`, `AVscan`, and Job Engine jobs. Cloning of a SmartLink file is also prevented and will return an EINVAL error.

SmartLink file timestamps can change

Opening a SmartLink file through a supported protocol can change the timestamp data. When a file is first archived, and the SmartLink file is created in its place, the ctime timestamp stays the same as the original file's timestamp. However, the first time the SmartLink file is opened (inline access), the ctime timestamp changes as a cache component is added to the file. Also, if an archived file is fully recalled, its ctime and mtime timestamps change.

Inline access can appear to convert a SmartLink file to a regular file

When a user accesses a SmartLink file on the Isilon cluster from a supported protocol, the file opens in an application on the client computer. During this process, called inline access, most applications support the creation of a CloudPools cache from which users can view and, if desired, modify archived data. With inline access, the SmartLink file remains intact on the cluster, and any modifications that the user makes to file data are stored in the cache and updated to the cloud.

However, some applications do not support inline access. Instead, these applications create a new copy of the original file apart from the SmartLink file. The new file, containing all original file data, is given a new logical I-node (LIN) number and timestamps that differ from the file that was originally archived. This behavior has been observed in only a few programs, including Microsoft Office applications. In these cases, since an entirely new file is created, the original SmartLink file and its associated data in the cloud is tagged for removal (garbage collection).

If the new file meets the criteria of the file pool policy that archived the original file to the cloud, the new file is archived to the cloud the next time the SmartPools job runs, and a new SmartLink file is created in its place on the local cluster. If the new file does not meet the policy criteria, the full file remains on the cluster.

For best results using CloudPools, we recommend that you avoid archiving files that are still being actively modified by your users.

Copying a SmartLink file from one cluster to another can result in a corrupted SmartLink file

If you use scp (Secure Copy) or other copy commands to copy a SmartLink file from one cluster to another, the resulting SmartLink file is corrupted.

Running Antivirus on SmartLink files generates errors

We recommend that you avoid creating policies that run Antivirus on directories containing SmartLink files. Running Antivirus on SmartLink files generates error messages. If you do run Antivirus on such directories, you can safely ignore these messages.

Other issues accessing SmartLink files

If the node does not have quorum, or the file system is read-only, or `/ifs` is not available, then you might encounter issues accessing SmartLink files.

Client-based tools and SmartLink files

If you run an SMB or NFS client-based tool such as AVScan (anti-virus scan) or a backup application, file data in the cloud is fully cached back to the SmartLink files. This can result in heavy network usage and increased service provider costs, and would also negate space saving on your cluster.

Expired SmartLink files

Expired SmartLink files are not restored using NDMP and do not synch back using SyncIQ. A SmartLink file on an NDMP backup or on a SyncIQ secondary (target) cluster is expired when the original SmartLink file has been deleted from the primary (source) cluster, or the original file data in the cloud has been fully recalled.

Recall can be interrupted

When a full cache is in process (that is, someone performed an inline access of a SmartLink file from an SMB share or NFS export), recall of the same file can fail. When this happens, the full cache is allowed to complete first, and the user should retry the recall after caching is completed.

ADS files

CloudPools does not archive and recall ADS (alternate data stream) files.

SMB Oplock

SMB Oplock (lease/notification) does not work in cases where you create a file with the SUPERCEDE flag, and the file already exists and is archived.

CloudPools logs

You can access CloudPools logs to view activity and troubleshoot problems.

The following logs are available in OneFS for CloudPools operation.

Type	Name	Path
Client cluster-side logs	Cpool daemon	<code>/var/log/isi_cpool_d.log</code>
	Job Engine	<code>/var/log/isi_job_d.log</code>
	SMB and NFS I/O	<code>/var/log/isi_cpool_io_d.log</code> and <code>/var/log/lwiod.log</code>
	Provisioning	<code>/var/log/isi_papi_d.log</code>

Type	Name	Path
	NDMP	/var/log/isi_ndmp_d.log
	SyncIQ	/var/log/isi_migrate.log
	Messages	/var/log/messages
Platform API cloud-side logs	Platform API (RAN)	/var/log/isi_object_d.log
	HTTPd apache	/var/log/apache2/webui_httpd_error.log and /var/log/apache2/ webui_httpd_access.log
	Session authentication	
	Messages	/var/log/messages

Note

Make sure that the client cluster-side time is accurate to within 15 minutes of the cloud provider.

Troubleshooting CloudPools

This section describes other troubleshooting items for CloudPools administration and operation.

Cloud storage account cannot connect to the cloud

In the OneFS, if a cloud storage account is shown in the web administration interface with a red `Needs Attention` icon, or in the CLI interface with an `Unreachable` state, this usually indicates that the cluster has lost Internet connectivity or the service provider's cloud storage facility is offline. Ensure that the cluster has Internet connectivity. If it does, contact your service provider for help.

Determining if a file is a SmartLink file

To determine if a file has been archived to the cloud, you can check whether the local version on the cluster is a SmartLink file. Run the `isi get -D` command as in the following example:

```
isi get -D koala.jpg | grep Stubbed:
```

The output would be as follows if the specified file was a SmartLink (stub) file:

```
* Stubbed: True
```

If the file is not a SmartLink file, the output would be `False`.

CHAPTER 6

CloudPools CLI commands

This section provides a reference to all CloudPools commands in the OneFS command line interface:

- [CloudPools command reference](#)..... 58

CloudPools command reference

Use CloudPools commands to manage general settings, and to create and manage cloud accounts and cloud storage targets. This section provides reference information for each command. For any command, you can use the `--help` option to get a full listing of command options.

isi cloud access add

Adds cloud write access to the cluster.

Syntax

```
isi cloud access add <guid>
  [--expiration-date] <timestamp>
  [--verbose]
```

Options

<guid>

The reference number, or globally unique identifier (GUID), of the cloud account.

--expiration-date <timestamp>

The date and time at which write access to cloud data ends on this cluster. The timestamp format is *MMDDYY:hh:mm*. For example, *022016:12:00* specifies an expiration date and time of February 20, 2016 at 12:00 PM.

--verbose

Displays more detailed information.

Examples

The following example adds cloud write access to a cluster by specifying the cluster GUID and an expiration date:

```
isi cloud access add 000556bf1e82059801563f1ad44a8c155acf
--expiration-date 022016:12:00
```

OneFS displays a message indicating the cloud accounts and file pool policies to which the secondary cluster will have access, and requires confirmation. Type `yes`, and press ENTER to complete the process.

isi cloud access list

Displays a list of clusters on your network that have, or are eligible for, write access to cloud data. Available clusters are the primary cluster and any other clusters to which data has been replicated with SyncIQ or restored with NDMP.

Syntax

```
isi cloud access list
  [--limit] <integer>
  [--sort {name | guid | synced_from | state | accounts | policies}]
  [--descending]
```

```
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

--limit <integer>

Limits the number of eligible clusters displayed in the list.

--sort

Sort the list of eligible clusters according to the specified category. The following values are valid:

```
name
guid
synced_from
state
accounts
policies
```

--format

Outputs the list of eligible clusters in the specified format. The following values are valid:

```
table
json
csv
list
```

--descending

Outputs the list of eligible clusters in descending order according to the specified sort option.

--no-header

Displays table and CSV output without headers.

--no-footer

Displays table output without footers.

--verbose

Displays more detailed information.

isi cloud access remove

Removes cloud write access from the specified cluster.

Syntax

```
isi cloud access remove <guid>
  [--force]
  [--verbose]
```

Options

<guid>

The reference number, or globally unique identifier (GUID), of the cluster from which you want to remove cloud write access.

--force

Execute the command without requiring confirmation.

--verbose

Displays more detailed information.

Examples

The following example removes cloud write access from a cluster identified by a specified GUID:

```
isi cloud access remove 000556bf1e82059801563f1ad44a8c155acf
```

OneFS displays a message indicating the cloud accounts and file pool policies to which the cluster will no longer have access, and requires confirmation. Type **yes**, and press ENTER to complete the process.

isi cloud access view

View the details of a cluster with, or eligible for, write access to cloud data.

Syntax

```
~isi cloud access view <guid>
```

Options

<guid>

The reference number, or globally unique identifier (GUID), of the cluster.

isi cloud accounts create

Creates a cloud storage account that connects CloudPools to your cloud storage provider.

Syntax

```
isi cloud accounts create <name> <type> <uri> <account-username>
<key>
  [--enabled {yes | no}]
  [--account-id <string>]
  [--telemetry-bucket <string>]
  [--storage-region <string>]
  [--skip-ssl-validation {yes | no}]
  [--proxy <string>]
  [--force]
  [--verbose]
```

Options

<name>

The name of the cloud storage account.

<type>

The type of cloud storage account: one of `isilon`, `azure`, `s3`, `virtustream`, or `ecs`.

<uri>

The cloud account URI. This URI must match that provided to the cloud vendor.

<account-username>

The username for the cloud account. This name must be identical to the user name provided to the cloud vendor.

<key>

The cloud account access key or password. This information is provided by the cloud vendor.

--enabled {yes | no}

By default, when you create a cloud storage account, it is enabled. To disable the account on creation, you can use this setting with the `no` option.

--account-id <string>

This is a required Amazon S3-only setting. The account ID number provided by Amazon when you first establish an account with the vendor.

--telemetry-bucket <string>

This is a required Amazon S3-only setting. The telemetry bucket name that you specified when you first established an account with the vendor.

--storage-region <string>

This is a required Amazon S3-only setting. The storage region that you specified when you first established an account with the vendor. For example, `us-west-1`.

--skip-ssl-validation {yes | no}

Specifies whether to circumvent SSL certificate validation when connecting to a cloud provider's storage repository. Unless you specify this setting with a `yes` instruction, OneFS will attempt to perform SSL certificate validation when connecting. For security purposes, we recommend not enabling this setting. If you are connecting to cloud provider (for example, Isilon or ECS) that is within your corporate network, and you are having trouble connecting, you can skip SSL validation.

--proxy <string>

The network proxy through which CloudPools traffic to and from a public cloud provider should be redirected. The specified network proxy must already have been created with the `isi cloud proxies create` command.

--force

Execute the command without requiring confirmation.

--verbose

Displays more detailed information.

Examples

The following example creates a Microsoft Azure cloud account:

```
isi cloud accounts create my_azure azure https://
myazure.windows.net myuser dhgXJ9OAIahXvYmL
```

isi cloud accounts delete

Delete a cloud storage account. Caution: deleting an account can result in loss of access to cloud data.

Syntax

```
isi cloud accounts delete <id>
[--acknowledge <string>]
[--verbose]
```

Options

<id>

The name of the cloud account. You can use the `isi cloud accounts list` command to display the names of cloud accounts.

--acknowledge <string>

Enables the account deletion to proceed. This parameter is required. You must include a text string with the parameter, such as `yes`, `proceed`, or other string.

--verbose

Displays more detailed information.

Example

The following example deletes a Microsoft Azure cloud account:

```
isi cloud accounts delete my_azure --acknowledge yes
```

When you run the command, OneFS displays the following message and requires confirmation:

```
*****
WARNING: Deleting an account is extremely dangerous.
Continuing with this operation will result in a permanent loss of
data.
Type 'confirm delete data' to proceed. Press enter to cancel:
```

To proceed, type **confirm delete data**, and press ENTER.

isi cloud accounts list

List cloud accounts.

Syntax

```
isi cloud accounts list
  [--limit <integer>]
  [--sort {id | name | type | account_username | uri | state |
bucket}]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--limit <integer>

Limits the number of cloud accounts displayed in the list.

--sort

Sort the list of cloud accounts according to the specified category. The following values are valid:

```
id
name
type
account_username
uri
state
bucket
```

--format

Outputs the list of cloud accounts in the specified format. The following values are valid:

```
table
json
csv
list
```

--descending

Outputs the list of cloud accounts in descending order according to the specified sort option.

--no-header

Displays table and CSV output without headers.

--no-footer

Displays table output without footers.

--verbose

Displays more detailed information.

isi cloud accounts modify

Modify a cloud account.

Syntax

```
isi cloud accounts modify <id>
  [--name <string>]
  [--account-username <string>]
  [--key <string>]
  [--uri <string>]
  [--enabled {yes | no}]
  [--account-id <string>]
  [--telemetry-bucket <string>]
  [--storage-region <string>]
  [--proxy <string>]
  [--clear-proxy]
  [--skip-ssl-validation {yes | no}]
  [--verbose]
```

Options

<id>

The ID of the cloud account. In this case, the ID is the same as the cloud account name.

--name <string>

The name of the cloud account. In this case, the name is the same as the ID.

--account-username <string>

The username for the cloud account. This name must be identical to the user name provided to the cloud vendor.

--key <string>

The cloud account access key or password. This information is provided by the cloud vendor.

--uri <string>

The cloud account URI. This URI must match that provided to the cloud vendor.

--enabled {yes | no}

By default, when you create a cloud storage account, it is enabled. To disable the account on creation, you can use this setting with the `no` option.

--account-id <string>

This is a required Amazon S3-only setting. The account ID number provided by Amazon when you first establish an account with the vendor.

--telemetry-bucket <string>

This is a required Amazon S3-only setting. The telemetry bucket name that you specified when you first established an account with the vendor.

--storage-region <string>

This is a required Amazon S3-only setting. The storage region that you specified when you first established an account with the vendor. For example, `us-west-1`.

--skip-ssl-validation {yes | no}

Specifies whether to circumvent SSL certificate validation when connecting to a cloud provider's storage repository. Unless you specify this setting with a `yes` instruction, OneFS will attempt to perform SSL certificate validation when connecting. For security purposes, we recommend not enabling this setting. If you are connecting to a cloud provider (for example, RAN or ECS) that is inside your corporate network, and you are having trouble connecting, you can skip SSL validation.

--proxy <string>

The network proxy through which CloudPools traffic to and from a public cloud provider should be redirected. The specified network proxy must already have been created with the `isi cloud proxies create` command.

--clear-proxy

Removes the network proxy through which CloudPools traffic to and from a public cloud provider had been redirected. When you remove a proxy, CloudPools traffic would flow directly to the cloud provider.

--skip-account-check {yes | no}

If set to `yes`, CloudPools skips the validation step to determine that the cloud storage account is accessible.

--verbose

Displays more detailed information.

Example

The following example modifies a Microsoft Azure cloud account:

```
isi cloud accounts modify my_azure --uri https://
myazure.windows.net
--account-username myuser --key dhgXJ90AIahXvYmL
```

isi cloud accounts view

View the details of a cloud account.

Syntax

```
isi cloud accounts view <name>
```

Options

<name>

Specifies the name of the cloud account to view. You can use the `isi cloud accounts list` command to display a list of the names of available cloud accounts.

Example

The following example displays the details of an Amazon S3 cloud account named `my_S3`:

```
isi cloud accounts view my_s3
```

isi cloud archive

Queue one or more files to be archived to or recalled from the cloud. Specify files individually or by using a file matching pattern. For files to be archived, they must match the specified file pool policy, or any file pool policy with a cloud target.

Syntax

```
isi cloud archive <files>
[--recursive {yes | no}]
[--policy <string>]
[--verbose]
[--help]
```

Options

<files>

Specifies the files to archive or recall. Specify `--files` for each additional file to process. Alternatively, you can specify a file matching pattern such as `/ifs/data/archive/images/*.jpg`.

--recursive {yes | no}

Specifies whether the operation should apply recursively to nested directories in the file string.

--policy <string>

Specifies the file pool policy to apply to the specified files. If you specify one or more files to be archived and do not specify a policy, OneFS will compare the files with each configured file pool policy.

--verbose

Displays more detailed information.

Examples

The following example archives multiple files to the cloud according to a specific file pool policy:

```
isi cloud archive /ifs/data/images/big.jpg --file /ifs/data/huge.jpg
--policy my_policy
```

The following example archives an entire directory to the cloud. The operation must match an existing file pool policy to be successful.

```
isi cloud archive /ifs/data/images/*.* --recursive yes
```

The following example recalls files from the cloud:

```
isi cloud archive /ifs/data/images/*.* --type recall
```

isi cloud jobs cancel

Cancel a CloudPools job initiated manually with `isi cloud archive` or `isi cloud recall`). CloudPools system jobs (such as `cache-writeback`) cannot be canceled.

Syntax

```
isi cloud jobs cancel <id>
[--verbose]
```

Options

<id>

The ID for the cloud job. Run `isi cloud jobs list` to see a list of all manual and system jobs and their associated IDs.

--verbose

Displays more detailed information.

Example

This following example cancels a CloudPools job with the ID of 21.

```
isi cloud cancel 21
```

isi cloud jobs files list

Displays the list of files matched by the specified CloudPools job.

Syntax

```
isi cloud jobs files list <job-id>
[--limit <integer>]
[--sort {name | state}]
[--descending]
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

<job-id>

The ID of the job. To find the list of job IDs in CloudPools, run the `isi cloud jobs list` command.

--limit *<integer>*

Display no more than the specified number of items.

--sort {name | state}

Order results by the specified field. The default value is `name`.

--descending

Sort and present data in descending order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

--no-header

Displays table and CSV output without headers.

--no-footer

Displays table output without footers.

--verbose

Displays more detailed information.

Example

The following example displays a list of files associated with a specific cloud job:

```
isi cloud jobs files list 21
```

isi cloud jobs list

View the status of CloudPools jobs, including system, archive, and recall jobs.

Syntax

```
isi cloud jobs list
[--limit <integer>]
[--sort {id | job_state | operation_state | effective_state | type
| state_change_time | completion_time | create_time | description}]
[--descending ]
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

--limit *<integer>*

Display no more than the specified number of items.

--sort {id | job_state | operation_state | effective_state | type | state_change_time | completion_time | create_time | description}

Order results by this field. The default value is `id`. Note that, to sort on other than ID, description, effective state, and type, use the `--verbose` parameter with the command.

--descending

Sort and present data in descending order.

--format {table | json | csv | list}

Display output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

--no-header

Displays table and CSV output without headers.

--no-footer

Displays table output without footers.

--verbose

Displays more detailed information.

isi cloud jobs pause

Pause a cloud job. A paused job can be resumed with the `isi cloud jobs resume` command.

Syntax

```
isi cloud jobs pause <id>
[--verbose]
```

Options

id

The ID of the cloud job to pause. Use the `isi cloud jobs list` command to view the IDs of all cloud jobs. Although possible, we recommend that you not pause any of the CloudPools system jobs that run in the background and are critical for proper operation. These include cache-writeback, cache-invalidation, local-garbage-collection, and cloud-garbage-collection.

--verbose

Displays more detailed information.

Example

The following example pauses a cloud job with ID 19.

```
isi cloud jobs pause 19
```

isi cloud jobs resume

Resume a paused cloud job.

Syntax

```
isi cloud jobs resume <id>
[--verbose]
```

Options

<id>

The ID for the cloud job to resume. Use the `isi cloud jobs list` command to view a list of jobs and their associated IDs.

--verbose

Displays more detailed information.

Example

The following command resumes a paused job with an ID of 26:

```
isi cloud jobs resume 26
```

isi cloud jobs view

View the details of a cloud job.

Syntax

```
isi cloud jobs view <id>
```

Options

<id>

Specify the ID of the cloud job. Use the `isi cloud jobs list` command to view all jobs and their associated IDs.

Example

The following command views the details of a job with the ID of 27:

```
isi cloud jobs view 27
```

isi cloud pools create

Create a CloudPool, which provides the connection between OneFS and a cloud storage account.

Syntax

```
isi cloud pools create <name> <type> <account>
[--description <string>]
```

```
[--vendor <string>]
[--verbose]
```

Options

<name>

The name of the CloudPool.

<type>

The type of account, one of `isilon`, `azure`, `s3`, `ecs`, or `virtustream`.

<account>

The name of the cloud storage account to which the CloudPool connects. The cloud storage account is required and must match the CloudPool type. Only one cloud storage account can be specified.

--description <string>

A description of the CloudPool.

--vendor <string>

The name of the vendor hosting the cloud storage account.

--verbose

Displays more detailed information.

Example

This following command creates a CloudPool containing a Microsoft Azure cloud storage account:

```
isi cloud pools create my_cp azure http://myazure.microsoft.com
--description="Financial records 2013" --vendor=Microsoft
```

isi cloud pools delete

Delete a CloudPool. Proceed with caution, however. If you delete a CloudPool, OneFS is no longer able to access the associated cloud storage account. If the CloudPool is referenced by a file pool policy, OneFS does not allow the CloudPool to be deleted.

Syntax

```
isi cloud pools delete <id>
[--force]
[--verbose]
```

Options

<id>

The name of the CloudPool. You can use the `isi cloud pools list` command to list existing CloudPools and their associated IDs.

--force

Deletes the account without asking for confirmation.

--verbose

Displays more detailed information.

Example 1 Example

The following command specifies a CloudPool to be deleted:

```
isi cloud pool delete my_azure_pool
```

When you press ENTER to run the command, OneFS asks for confirmation. Type **yes**, then press ENTER.

isi cloud pools list

Display a list of CloudPools.

Syntax

```
isi cloud pools list
[--limit <integer>]
[--sort {id | name | type | state | state_details | description
| vendor}]
[--descending ]
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

--limit <integer>

Displays no more than the specified number of items.

--sort {id | name | type | state | state_details | description | vendor}

Order results by this field. The default value is `id`, which, in this case, is the same as `name`. Unless you use the `--verbose` option, you can only sort on `name`, `type`, or `state`.

--descending

Sorts and presents data in descending order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

--no-header

Displays table and CSV output without headers.

--no-footer

Displays table output without footers.

--verbose

Displays more detailed information.

isi cloud pools modify

Modifies a CloudPool.

Syntax

```
isi cloud pools modify <id>
  [--name <string>]
  [--accounts <string>]
  [--add-accounts <string>]
  [--clear-accounts]
  [--remove-accounts <string>]
  [--description <string>]
  [--vendor <string>]
  [--verbose]
```

Options

<id>

The ID of the CloudPool. Run `isi cloud pools list` to view the IDs of all CloudPools.

--name <string>

Specify a new name for the CloudPool.

--account <string>

Specify the name of the cloud account to add to the CloudPool. Only one account per CloudPool is allowed.

--add-account <string>

Specify the name of a cloud account to add to the CloudPool. Only one account per CloudPool is allowed.

--remove-accounts <string>

Specify the name of the cloud account to remove from the CloudPool. You can only remove an account if you are adding a different account in the same command.

--description

Specify the name of the cloud account to remove from the CloudPool. You can only remove an account if you are adding a different account in the same command.

---vendor <string>

The name of the vendor hosting the cloud pool accounts.

--verbose

Displays more detailed information.

Examples

The following command adds a vendor name and description to an existing CloudPool:

```
isi cloud pools modify my_azure --vendor Microsoft
--description "preferred azure account"
```

The following command removes one cloud account from the CloudPool, and adds another cloud account:

```
isi cloud pools modify my_s3 --remove-accounts s3_acct_1
--add-accounts s3_acct_2
```

isi cloud pools view

View detailed information about a CloudPool.

Syntax

```
isi cloud pools view <id>
```

Options**<id>**

The ID of the cloud pool. Run the `isi cloud pool list` command to view all CloudPools and their associated IDs.

Example

The following command displays information about a CloudPool named `my_azure_pool`.

```
isi cloud pools view my_azure_pool
```

isi cloud proxies create

Creates a network proxy through which a cloud storage account can connect to a cloud storage provider.

Syntax

```
isi cloud proxies create <name> <host> <type> <port>
[--username <string>]
[--password <string>]
[--verbose]
```

Options**<name>**

The name of the network proxy. This can be any alphanumeric string, but should be a simple, recognizable name.

<host>

The DNS name or IP address of the proxy server. For example, `myproxy1.example.com` or `192.168.107.107`.

<type>

The proxy protocol type, one of `socks_4`, `socks_5`, or `http`.

<port>

The port number to communicate with the proxy server. The correct port number depends on the port opened up on the proxy server for communication with CloudPools.

--username <string>

The user name to authenticate with the SOCKS v5 or HTTP proxy server. Note that SOCKS v4 does not support authentication.

--password <string>

The password to authenticate with the SOCKS v5 or HTTP proxy server.

--verbose

Displays more detailed information.

Examples

The following example creates a network proxy to use with CloudPools:

```
isi cloud proxies create myproxy1 myprox1.example.com socks_5 1080
--username mycloudpools --password dhgXJ90AIahXvYmL
```

isi cloud proxies delete

Delete a network proxy in CloudPools. Note that CloudPools prevents deletion of a proxy that is attached to a cloud storage account.

Syntax

```
isi cloud proxies delete <name>
[--force]
[--verbose]
```

Options**<name>**

The name of the network proxy. You can use the `isi cloud proxies list` command to display the names of proxies.

--force

Enables the proxy deletion to proceed without confirmation.

--verbose

Displays more detailed information.

Example

The following example deletes a network proxy named `myproxy1`:

```
isi cloud accounts delete myproxy1
```

When you run the command, OneFS displays the following message and requires confirmation:

```
Are you sure? (yes/[no]):
```

To proceed, type **yes**, and press ENTER. If the proxy is attached to a cloud storage account, OneFS displays the following message:

```
Cannot delete proxy while used by accounts
```

isi cloud proxies list

Displays a list of network proxies created in CloudPools.

Syntax

```
isi cloud proxies list
[--limit <integer>]
[--sort {id | name | host | type | port}]
[--descending]
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

--limit <integer>

Limits the number of network proxies displayed in the list.

--sort

Sort the list of cloud proxies according to the specified category. The following values are valid:

```
id
name
host
type
port
```

--format

Outputs the list of network proxies in the specified format. The following values are valid:

```
table
json
csv
list
```

--descending

Outputs the list of network proxies in descending order according to the specified sort option.

--no-header

Displays table and CSV output without headers.

--no-footer

Displays table output without footers.

--verbose

Displays more detailed information.

Example

The following example creates a network proxy to use with CloudPools:

```
isi cloud proxies create myproxy1 myprox1.example.com socks_5 1080
--username mycloudpools --password dhgXJ90AIahXvYmL
```

isi cloud proxies modify

Modifies the properties of a network proxy.

Syntax

```
isi cloud proxies create <name>
[--name <string>]
[--host <string>]
[--type {socks_4 | socks_5 | http}]
[--port <integer>]
[--username <string>]
[--clear-username]
[--password <string>]
[--clear-password]
[--verbose]
```

Options

<name>

The current name of the network proxy.

--name <string>

The new name of the network proxy. This can be any alphanumeric string, but should be a simple, recognizable name.

--host <string>

The DNS name or IP address of the proxy server. For example, myproxy1.example.com or 192.168.107.107.

--type

The network proxy protocol, one of socks_4, socks_5, or http.

--port

The port number to communicate with the proxy server. The correct port number depends on the port opened up on the proxy server for communication with CloudPools.

--username <string>

The user name to authenticate with the SOCKS v5 or HTTP proxy server. Note that SOCKS v4 does not support authentication.

--clear-username

Clear the user name that was previously specified for proxy server authentication.

--password <string>

The password to authenticate with the SOCKS v5 or HTTP proxy server.

--clear-password

Clear the password that was previously specified for proxy server authentication.

--verbose

Displays more detailed information.

Examples

The following example modifies a network proxy in CloudPools:

```
isi cloud proxies modify myproxy1 --type socks_4 --clear-username --
clear-password
```

isi cloud proxies view

View the details of a network proxy created for CloudPools.

Syntax

```
isi cloud proxies view <name>
```

Options**<name>**

Specifies the name of the network proxy to view. You can use the `isi cloud proxies list` command to display a list of the available proxies.

Example

The following example displays the details of a network proxy named `myproxy1`:

```
isi cloud proxies view myproxy1
```

isi cloud recall

Specify one or more files to be recalled from the cloud. You can specify files individually or by using a file matching pattern. To make sure that the specified files are present in the cloud, OneFS scans the cluster for SmartLink files prior to performing the recall.

Syntax

```
isi cloud recall <files>
[--recursive {yes | no}]
[--verbose]
```

Options**<files>**

Specifies the files to recall. Specify `--files` for each additional file name.

`--recursive {yes | no}`

Specifies whether the recall should apply recursively to nested subdirectories.

`--verbose`

Displays more detailed information about the operation.

Examples

The following example recalls all files from the cloud for a directory and its subdirectories:

```
isi cloud recall /ifs/data/archives/archives2014/projects/*.*
--recursive yes
```

The command starts a cloud job. If you use the `--verbose` parameter, OneFS reports the job number, as in the following example:

```
Created job [29]
```

You can use the `isi cloud jobs view` command with the job number to see information about the job.

Note

When you use the `isi cloud recall` command to recall a file from cloud storage, the full file is restored to its original directory, and the associated SmartLink file is overwritten. If the file pool policy that originally archived the file to the cloud is still in effect, the next time the SmartPools job runs, the recalled file is archived to the cloud again. If you do not want the recalled file to be re-archived, you can move the file to a different directory that would not be affected by the file pool policy, or you can modify or delete the policy.

isi cloud restore_coi

Restores the cloud object index (COI) for a cloud storage account on the cluster. If you run the `isi cloud access add` command on a cluster, it automatically restores the COI.

Syntax

```
isi cloud restore_coi
[--account <string>]
[--expiration-date <timestamp>]
[--verbose]
```

Options

`--account <string>`

Specifies the name of the cloud storage account whose COI you intend to restore. By restoring the COI, you enable OneFS to not only read data from the cloud, but also to write data to the cloud.

`--expiration-date <timestamp>`

Specifies the expiration date for orphaned cloud data objects.

--verbose

Displays more detailed information about the operation.

Example

The following example restores the COI for a cloud storage account:

```
isi cloud restore_coi --account my_azure_acct
```

isi cloud settings modify

Controls archiving of snapshot files. By default, archiving of snapshots is enabled.

Syntax

```
isi cloud settings modify
[--default-accessibility {cached | no-cache}]
[--default-cache-expiration <duration>]
[--default-compression-enabled {yes | no}]
[--default-data-retention <duration>]
[--default-encryption-enabled {yes | no}]
[--default-full-backup-retention <duration>]
[--default-incremental-backup-retention <duration>]
[--default-read-ahead <string>]
[--default-writeback-frequency <duration>]
[--default-archive-snapshot-files {yes | no}]
[--verbose]
```

Options

--default-accessibility {cached | no-cache}

Specifies whether, when a SmartLink file is accessed, cloud data is incrementally downloaded (cached) as needed, or fully downloaded (not cached).

--default-cache-expiration <duration>

Specifies the minimum amount of time until the cache expires. A number followed by a unit of time is accepted. For example, a setting of 9H would specify a nine-hour duration. Similarly, a setting of 2D would specify a two-day duration.

--default-compression-enabled {yes | no}

Specifies whether data is to be compressed when archived to the cloud.

--default-data-retention <duration>

Specifies the minimum amount of time that cloud objects associated with a SmartLink file will be retained in the cloud after the SmartLink file is deleted from the cluster. A number followed by a unit of time is accepted. For example, a setting of 9H would specify a nine-hour duration. Similarly, a setting of 2D would specify a two-day duration.

--default-encryption-enabled {yes | no}

Specifies whether data is to be encrypted when archived to the cloud.

--default-full-backup-retention <duration>

Specifies the length of time that OneFS retains cloud data referenced by a SmartLink file that has been backed up by a full NDMP backup and is subsequently deleted. A number followed by a unit of time is accepted. For

example, a setting of 9H would specify a nine-hour duration. Similarly, a setting of 2D would specify a two-day duration.

--default-incremental-backup-retention <duration>

Specifies the length of time that OneFS retains cloud data referenced by a SmartLink file that has been backed up by an incremental NDMP backup, or replicated by a SyncIQ operation, and is subsequently deleted. A number followed by a unit of time is accepted. For example, a setting of 5Y would specify a five-year duration.

--default-read-ahead {partial | full}

Specifies the cache readahead strategy when SmartLink files are accessed. A partial strategy means that only the amount of data needed by the user is cached. A full strategy means that all file data will be cached when the user accesses a SmartLink file.

--default-writeback-frequency <duration>

Specifies the minimum amount of time to wait before OneFS updates cloud data with local changes. A number followed by a unit of time is accepted. For example, a setting of 9H would specify a nine-hour duration. Similarly, a setting of 2D would specify a two-day duration.

--default-archive-snapshot-files {yes | no}

Whether or not policies should archive files with snapshots.

--verbose

Displays more information about the operation.

Example

The following examples modifies several of the default CloudPools settings:

```
isi cloud settings modify --default-writeback-frequency 12H
--default-cache-expiration 9H --default-accessability no-cache
--default-encryption-enabled yes
```

isi cloud settings regenerate-encryption-key

Generates a new master encryption key for data to be archived to the cloud. Encrypted data already stored in the cloud requires and stores the previous encryption key.

Syntax

```
isi cloud settings regenerate-encryption-key
[--verbose]
```

Option

--verbose

Displays more detailed information.

isi cloud settings view

Display the current default settings in CloudPools. You can use the `isi cloud settings modify` command to change default settings.

Syntax

```
isi cloud settings view
```

Options

There are no options for this command.

Copyright © 2015 EMC Corporation. All rights reserved. Published in USA.

EMC 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. EMC Corporation 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 EMC software described in this publication requires an applicable software license.

EMC², EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries.

All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to EMC Online Support (<https://support.emc.com>).