

Dell PowerScale: CloudPools and Amazon Web Services

Architectural Overview, Considerations, and Best Practices

December 2025

H17747.9

White Paper

Abstract

This white paper provides an overview of Dell PowerScale CloudPools software in OneFS. It describes its policy-based capabilities that can reduce storage costs and optimize storage by automatically moving infrequently accessed data to Amazon Web Services (AWS).

Copyright

The information in this publication is provided as is. Dell Inc. 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 software described in this publication requires an applicable software license.

Copyright © 2019-2025 Dell Inc. or its subsidiaries. Published in the USA December 2025 H17747.9.

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

Contents

Executive summary.....4

CloudPools solution architectural overview6

CloudPools 2.012

Best practices for PowerScale storage and AWS21

Reporting28

Commands and troubleshooting29

Appendix A: Step-by-step configuration example33

References.....50

Executive summary

Overview

This white paper describes about how Dell PowerScale CloudPools in OneFS 9.0 integrates with Amazon Web Services (AWS) and it covers the following topics:

- CloudPools solution architectural overview
- CloudPools 2.0 introduction with a focus on the following improvements:
 - AWS signature v4 authentication support
 - Commercial Cloud Services (C2S) support
 - Dell PowerScale NDMP and Dell PowerScale SyncIQ support
 - Non-disruptive upgrade (NDU) support
 - Snapshot efficiency
 - Sparse files handling
 - Quota management
 - Anti-virus integration
 - WORM integration
- General considerations and best practices for a CloudPools implementation
- CloudPools reporting, commands, and troubleshooting

Audience

This white paper is intended for experienced system administrators, storage administrators, and solution architects interested in learning how CloudPools works and understanding the CloudPools solution architecture, considerations, and best practices.

This guide assumes the reader has a working knowledge of the following:

- Network-attached storage (NAS) systems
- Dell PowerScale scale-out storage architecture and Dell PowerScale OneFS operating system
- AWS

The reader should also be familiar with PowerScale and AWS documentation resources including the following:

- Dell OneFS release notes, available on [Dell Support](#), containing important information about resolved and known issues
- [Dell PowerScale OneFS Best Practices](#)
- [Amazon Web Services \(AWS\)](#)

Revisions

Date	Part number/ revision	Description
April 2019		Initial release
October 2019		Updated snapshot efficiency

Date	Part number/ revision	Description
June 2020		Updated best practices
October 2020		Updated CloudPools operations
April 2021		Updated best practices
October 2021		Updated performance
April 2022	H17747.6	Updated reporting
May 2023	H17747.7	Updated best practices
July 2023	H17747.8	Document template update
December 2025	H17747.9	Updated for Amazon VPC endpoint support

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on this document. Contact the Dell Technologies team by [email](#).

Author: Jason He (Jason.He@dell.com)

Note: For links to other documentation for this topic, see the [PowerScale Info Hub](#).

CloudPools solution architectural overview

The CloudPools feature of OneFS allows tiering cold or infrequently accessed data to lower-cost cloud storage. It is built on the Dell PowerScale OneFS SmartPools file pool policy framework, which provides granular control of file placement on a PowerScale cluster.

CloudPools extends the PowerScale namespace to the public cloud, AWS, as illustrated in [Figure 1](#). It allows applications and users to seamlessly retain access the data through the same network path and protocols regardless of where the file data physically resides.




Figure 1. CloudPools solution overview

Note: A SmartPools license and a CloudPools license are required on each node of the PowerScale cluster. A minimum of Dell Isilon OneFS version 8.0.0 is required for CloudPools 1.0, and Dell Isilon OneFS version 8.2.0 for CloudPools 2.0.

Policies are defined on the PowerScale cluster and drive the tiering of data. Clients can access the archived data through various protocols including SMB, NFS, HDFS, and S3.

PowerScale

This section describes key CloudPools concepts including the following:

- SmartPools
- SmartLink files
- File pool policies

SmartPools

SmartPools is the OneFS data tiering framework, of which CloudPools is an extension. SmartPools alone tiers data between different node types within a PowerScale cluster. CloudPools also adds to tier data outside of a PowerScale cluster.

SmartLink files

Although file data is moved to cloud storage, the files remain visible in OneFS. After file data has been archived to the cloud storage, the file is truncated to an 8 KB file. The 8 KB file is called a SmartLink file or stub file. Each SmartLink file contains a data cache and a map. The data cache is used to retain a portion of the file data locally, and the map points to all cloud objects.

Figure 2 shows the contents of a SmartLink file and the mapping to cloud objects.




Figure 2. SmartLink file

File pool policies

Both CloudPools and SmartPools use the file pool policy engine to define which data on a cluster should live on which tier or be archived to a cloud storage target. The SmartPools and CloudPools job has a customizable schedule that runs once a day by default. If files match the criteria specified in a file pool policy, the content of those files is moved to cloud storage during the job execution. A SmartLink file is left behind on the PowerScale cluster that contains information about where to retrieve the data. In CloudPools 1.0, the SmartLink file sometimes referred to as a stub, which is a unique construct that does not behave like a normal file. In CloudPools 2.0, the SmartLink file is an actual file that contains pointers to the CloudPool target where the data resides.

This section describes the key options when configuring a file pool policy, which includes the following:

- Encryption
- Compression
- File matching criteria
- Local data cache
- Data retention

Encryption

CloudPools provides an option to encrypt data before it is sent to the cloud storage. It leverages the PowerScale key management module for data encryption and uses AES-256 as the encryption algorithm. The benefit of encryption is that only encrypted data is being sent over the network.

Compression

CloudPools provides an option to compress data before it is sent to the cloud storage. It implements block level compression using the zlib compression library. CloudPools does not compress data that is already compressed.

File-matching criteria

When files match a file pool policy, CloudPools moves the file data to the cloud storage. File matching criteria enable defining a logical group of files as a file pool for CloudPools. It defines which data should be archived to cloud storage.

File matching criteria include the following:

- File name
- Path
- File type
- File attribute
- Modified
- Accessed
- Metadata changed
- Created
- Size

Any number of file matching criteria can be added to refine a file pool policy for CloudPools.

Local data cache

Caching is used to support local reading and writing of SmartLink files. It reduces bandwidth costs by eliminating repeated fetching of file data for repeated reads and writes to optimize performance.

Note: The data cache is used for temporarily caching file data from the cloud storage on PowerScale disk storage for files that have been moved off cluster by CloudPools.

The local data cache is always the authoritative source for data. CloudPools looks for data in the local data cache first. If the file being accessed is not in the local data cache, CloudPools fetches the data from the cloud. CloudPools writes the updated file data in the local cache first and periodically sends the updated file data to the cloud.

CloudPools provides the following configurable data cache settings:

- **Cache expiration:** This option is used to specify the number of days until OneFS purges expired cache information in SmartLink files. The default value is one day.
- **Writeback frequency:** This option is used to specify the interval at which OneFS writes the data stored in the cache of SmartLink files to the cloud. The default value is nine hours.
- **Cache read ahead:** This option is used to specify the cache read ahead strategy for cloud objects (partial or full). The default value is partial.

- **Accessibility:** This option is used to specify how data is cached in SmartLink files when a user or application accesses a SmartLink file on the PowerScale cluster. Values are cached (default) and no cache.

Data retention

Data retention is a concept used to determine how long to keep cloud objects on the cloud storage. There are three different retention periods:

- **Cloud data retention period:** This option is used to specify the length of time cloud objects are retained after the files have been fully recalled or deleted. The default value is one week.
- **Incremental backup retention period for NDMP incremental backup and SyncIQ:** This option is used to specify the length of time that CloudPools retains cloud objects referenced by a SmartLink file. And SyncIQ replicates the SmartLink file or NDMP backs up the SmartLink file using an incremental NDMP backup. The default value is five years.
- **Full backup retention period for NDMP only:** This option is used to specify the length of time that OneFS retains cloud data referenced by a SmartLink file. And NDMP backs up the SmartLink file using a full NDMP backup. The default value is five years.

Note: If more than one period applies to a file, the longest period is applied.

AWS

This section describes the following cloud objects in AWS:

- Cloud metadata object
- Cloud data object

Cloud metadata object

A cloud metadata object (CMO) is a CloudPools object in AWS that is used for supportability purposes.

Cloud data object

A cloud data object (CDO) is a CloudPools object that stores file data in AWS. File data is split into 2 MB chunks to optimize performance before sending it to AWS. The chunk is called a CDO. If file data is less than the chunk size, the CDO size is equal to the size of the file data.

Note: The chunk size is 1 MB in CloudPools 1.0 and versions prior to OneFS 8.2.0.

CloudPools operations

This section describes the workflow of CloudPools operations:

- Archive
- Recall
- Read
- Update

Archive

The archive operation is the CloudPools process of moving file data from the local PowerScale cluster to cloud storage. Files are archived either using the SmartPools Job or from the command line. The CloudPools archive process can be paused or resumed. See the section [Commands](#) for details.

Figure 3 shows the workflow of the CloudPools archive.




Figure 3. Archive workflow

More workflow details include the following:

- The file pool policy in Step 1 (see the section [File pool policies](#)) specifies a cloud target and cloud-specific parameters. Example policies include the following:
 - [Encryption](#)
 - [Compression](#)
 - [Local data cache](#)
 - [Data retention](#)
- When chunks are sent from the PowerScale cluster to AWS in Step 3, a checksum is applied for each chunk to ensure data integrity.

Recall

The recall operation is the CloudPools process of reversing the archive process. It replaces the SmartLink file by restoring the original file data on the PowerScale cluster and removing the cloud objects in AWS. The recall process can only be performed using the command line. The CloudPools recall process can be paused or resumed. See the section [Commands](#) for detailed instructions on commands.

Figure 4 shows the workflow of CloudPools recall.




Figure 4. Recall workflow

Read

The read operation is the CloudPools process of client data access, known as inline access. When a client opens a file for read, the blocks will be added to the cache in the associated SmartLink file by default. The cache can be disabled by setting the accessibility. For more detail, see the section [File pool policies](#).

Figure 5 shows the workflow of CloudPools read by default.




Figure 5. Read workflow

Starting from OneFS 9.1.0.0, cloud object cache is introduced to enhance CloudPools functions for communicating with cloud. In step 1, OneFS looks for data in the object

cache first and OneFS retrieves data from the object cache if the data is already in the object cache. Cloud object cache reduces the number of requests to AWS when reading a file.

Prior to OneFS 9.1.0.0, OneFS looks for data in the local data cache first in step 1. It moves to step 3 if the data is already in the local data cache.

Note: Cloud object cache is per node. Each node maintains its own object cache on the cluster.

Update

The update operation is the CloudPools process that occurs when clients update data. When clients change to a SmartLink file, CloudPools first writes the changes in the data local cache and then periodically sends the updated file data to AWS. The space used by the cache is temporary and configurable. For more information, refer to the section [File pool policies](#).




Figure 6. Update workflow

CloudPools 2.0

CloudPools 2.0 is the next generation of CloudPools, released in OneFS 8.2.0. This section describes the following improvements in CloudPools 2.0:

- AWS signature v4 authentication support
- Commercial Cloud Services (C2S) support
- NDMP and SyncIQ support
- Non-disruptive upgrade (NDU) support

- Snapshot efficiency
- Sparse files handling
- Quota management
- Anti-virus integration
- WORM integration

AWS signature v4 authentication support

CloudPools 2.0 supports AWS signature version 4 (V4) with signature version 2 (V2). V4 provides an extra level of security for authentication with the enhanced algorithm and no action is required from end users. For more information about V4, see the article [Authenticating Requests: AWS Signature V4](#).

CloudPools 2.0 handles the compatibility of SyncIQ for data replication and NDMP for data backup and restore. When the source and target PowerScale clusters use different authentication versions, consider the following points for CloudPools features:

- With **SyncIQ**, when the source PowerScale cluster is running OneFS 8.2.0 and the target PowerScale cluster is running a version of OneFS prior to OneFS 8.2.0:
 - If the CloudPools cloud storage account is using V2 or V4 on the source PowerScale cluster, V2 is used on the target PowerScale cluster.
- With **NDMP**, when files are restored from tape to the target PowerScale cluster:
 - If the CloudPools cloud storage account is using V4 on the target PowerScale cluster, V4 is used.
 - If the CloudPools cloud storage account is using V2 on the target PowerScale cluster, V2 is used.
- With **NDU**, when upgrading OneFS to version 8.2.0:
 - Once the PowerScale cluster is COMMITTED to OneFS 8.2.0, it automatically begins using V4.
 - CloudPools cloud storage accounts cannot use V4 in the UPGRADED state if the version prior to the OneFS 8.2.0 upgrade did not support V4.

Note: A patch can be installed on OneFS 8.1.2 to support AWS signature V4 authentication. Contact your Dell representative if you plan to install the patch.

Commercial Cloud Services support

CloudPools supports Commercial Cloud Services (C2S) in OneFS 8.2.0. An administrator can create or manage CloudPools C2S accounts using the WebUI or CLI.

C2S is a private instantiation of the AWS commercial cloud. The C2S Access Portal (CAP) is a centralized application developed to grant controlled access to the C2S Management Console and C2S APIs for government users and applications. The CAP uses X.509 client certificate authentication and a federated role-based access control model. Temporary credentials are used with [AWS Security Token Service \(STS\)](#) for federated users.

CloudPools and C2S integrate with token-based CAP access and bring the benefits of security for federal customers.

NDMP and SyncIQ support

When the CloudPools version differs between the source cluster and the target PowerScale cluster, the CloudPools cross-version compatibility is handled.

NDMP and SyncIQ provide two types of copy or backup: shallow copy and deep copy. For more information about NDMP and SyncIQ protection, refer to the white paper [High Availability and Data Protection with Dell PowerScale Scale-out NAS](#).

- **Shallow copy (SC)/backup:** Replicates or backs up SmartLink files to the target PowerScale cluster or tape as SmartLink files without file data.
- **Deep copy (DC)/backup:** Replicates or backs up SmartLink files to the target PowerScale cluster or tape as regular files or unarchived files. The backup or replication will be slower than normal. Disk space will be consumed on the target cluster for replicating data.

Table 1 shows the CloudPools and OneFS mapping information. CloudPools 2.0 is released along with OneFS 8.2.0. CloudPools 1.0 is running in OneFS 8.0.x or 8.1.x.

Table 1. CloudPools and OneFS mapping information

OneFS version	CloudPools version
OneFS 8.0.x/OneFS 8.1.x	CloudPools 1.0
OneFS 8.2.0 or higher	CloudPools 2.0

Table 2 shows the NDMP and SyncIQ supported use cases when running a different version of CloudPools on the source and target clusters. As noted below, if CloudPools 2.0 is running on the source PowerScale cluster and CloudPools 1.0 is running on the target PowerScale cluster, shallow copies are not allowed.

Table 2. NDMP and SyncIQ Supported use cases with CloudPools 2.0

Source	Target	SC NDMP	DC NDMP	SC SyncIQ replication	DC SyncIQ replication
CloudPools 1.0	CloudPools 2.0	Support	Support	Support	Support
CloudPools 2.0	CloudPools 1.0	No Support	Support	No Support	Support

Nondisruptive upgrade support

When a cluster that has been using CloudPools 1.0 is upgraded to OneFS 8.2.0 or higher, a new CHANGEOVER process is initiated automatically after the upgrade commit. The process ensures a smooth transition from CloudPools 1.0 to CloudPools 2.0. CloudPools 2.0 is ready to use once the upgrade state is committed. For more information about upgrade states, refer to the white paper [PowerScale Non-Disruptive Upgrade \(NDU\) Best Practices](#).

Snapshot efficiency

Prior to OneFS 8.2.0, CloudPools 1.0 supported archiving files with existing snapshots. However, CloudPools 1.0 had a limitation when archiving files that have existing snapshots: the copy-on-writes (CoW) process copied the entire contents of the file into the snapshot. Archiving files with existing snapshots therefore did not save space on the PowerScale cluster until the previously CoW-created snapshots expired. CloudPools 1.0 offered an option (Uncheck Archive files with snapshots in WebUI) to skip such files with snapshots. A user might have not chosen to archive files with snapshots if the previously

CoW-created snapshots had long retentions. This case is to avoid creating another copy on cloud storage where the retention period meant it would persist on PowerScale storage anyway.

CloudPools 2.0 eliminates CoW on the primary data source PowerScale cluster when archiving files with snapshots to the cloud. The file data is only stored in the cloud storage, which saves space on the PowerScale cluster. For more information about data CoW for snapshots, refer to the white paper [Data Protection with Dell PowerScale SnapshotIQ](#).

However, CloudPools 2.0 does not archive files on the target cluster in a SyncIQ relationship. In an environment with long snapshot retentions and an expectation that the same snapshots are maintained in both clusters. It is possible for storage usage on a target cluster to grow larger than the storage on the primary cluster which has CloudPools enabled. For space efficiency, a user with requirements for long snapshot retentions on two clusters in a SyncIQ relationship might choose to use natively tiered PowerScale archive storage, rather than CloudPools.

SnapshotIQ can take read-only, point-in-time copies of any directory or subdirectory within OneFS. A file in one directory can be either a regular file or a SmartLink file before creating a snapshot. A regular file can be truncated to a SmartLink file after archiving its file data to the cloud. A SmartLink file can be converted to a regular file after recalling its file data to the PowerScale cluster. When a snapshot is taken, it preserves the exact state of a file system at that instant. A file in the snapshot directory (/ifs/.snapshot) is a SmartLink file if the same file in the source directory is a SmartLink file. A file in the snapshot directory is a regular file if the same file in the source directory is a regular file. The earlier version of data can be accessed later in the snapshot directory.

The following scenarios address CloudPools 2.0 and snapshots. HEAD is the current version of a SmartLink file in the source directory.

- The file is already a SmartLink file in the source directory before creating a snapshot.
 - Scenario 1: Update HEAD
 - Scenario 2: Update HEAD multiple times and a new snapshot is created between multiple updates
 - Scenario 3: Read file data from a snapshot
- The file is still a regular file in the source directory before creating a snapshot. Then, the regular file is archived to the cloud after a snapshot creation.
 - Scenario 4: Update HEAD
 - Scenario 5: Read file data from a snapshot

Scenario 1

When updating HEAD (SmartLink files in snapshot), a new SmartLink is generated for HEAD when updating HEAD and write-back to the cloud. Cache for HEAD will be empty once its own cache expires. For the workflow of updating a SmartLink file, refer to the section [Update](#). The original version SmartLink file is still used for the next snapshot of HEAD. This scenario does not cause the snapshot space to grow. [Figure 7](#) shows the process of scenario 1 to update HEAD when SmartLink files are in the snapshot directory.




Figure 7. Scenario 1: Update HEAD when SmartLink files are in the snapshot directory

Scenario 2

This scenario describes updating HEAD multiple times, and a new snapshot is created between multiple updates (SmartLink files in snapshot). For example, a user updates HEAD (the first update) while a new (most recent) snapshot is created before the first update write-back is made to the cloud. Then, another user updates (the second update) HEAD again after the new (most recent) snapshot is created. Now there are two snapshots: one snapshot is the next snapshot of HEAD, the other is the most recent snapshot of HEAD.

When a snapshot is taken, it preserves the exact state of a file system at that instant. Data for the next snapshot of HEAD is the old data that is already archived to the cloud and its cache is empty. Data for the most recent snapshot is the new data and its cache is dirty before the new data write-back is made to the cloud. The new data contains old data with the first update. Data for HEAD is the latest data and its cache is dirty before the latest data write-back is made to the cloud. The latest data contains old data with the first update and the second update. A new version SmartLink is generated for the most recent snapshot after the new data write-back is made to the cloud (write-back in the snapshot). The new data contains old data with the first update. Also, a new version SmartLink is generated for HEAD after the latest data write-back is made to the cloud (write-back in HEAD). Cache for the most recent snapshot or HEAD becomes empty once its own cache expires. Now, all file data is only stored on the cloud and saves space on the PowerScale cluster. Users can read file data from its own SmartLink file at any time.

Figure 8 shows the process of scenario 2.




Figure 8. Scenario 2: Update HEAD multiple times and perform a write-back in the snapshot

Scenario 3

This scenario describes reading file data from a snapshot (SmartLink files in snapshot). The files in the next snapshot and HEAD use the same version of SmartLink file when not updating HEAD after the snapshot is created. This scenario is no different than reading the same file from HEAD or the next snapshot of HEAD. For the workflow of reading a SmartLink file, refer to the section [Read](#). The same local data cache is used when reading the same file from HEAD and the next snapshot of HEAD simultaneously. This scenario does not cause the snapshot space to grow. The file in the snapshot directory uses its version of SmartLink file when updating HEAD and performing a write-back to the cloud like in scenario 1 or scenario 2. Users can read earlier versions of file data in the snapshot directory. The snapshot space could grow temporarily for cache data, and the grown space is released once its own cache expires.

Scenario 4

In this scenario, when updating HEAD (regular files in snapshot). A SmartLink file is used for HEAD, and a regular file is used for the same file in the next snapshot of HEAD. A new SmartLink file is generated for HEAD when updating HEAD and performing a write-back to the cloud. The cache for HEAD is empty once its own cache expires. Meanwhile, OneFS enables the Block Allocation Manager Cache Manager (BCM) on the regular file in the next snapshot of HEAD. BCM contains the metadata of mapping to cloud objects for the regular file in the next snapshot of HEAD. This scenario does not cause the snapshot space to grow.

[Figure 9](#) shows scenario 4.




Figure 9. Scenario 4: Update HEAD when regular files are in the snapshot directory

Scenario 5

In this scenario, when reading file data from a snapshot (regular files in snapshot). File data is the same for HEAD (SmartLink file) and the same file (regular file) in the next snapshot of HEAD when not updating HEAD after the snapshot creation. File data is read from HEAD when reading the same file in the next snapshot of HEAD. This scenario does not cause the snapshot space to grow. The file in the next snapshot of HEAD is a regular file (enabled BCM). And the file has the earlier version of data when updating HEAD and performing a write-back to the cloud like in scenario 4. The earlier version of data is retrieved from the cloud by BCM. File data is stored on the PowerScale cluster when reading the earlier version of data from the regular file in the next snapshot of HEAD. The snapshot space grows, and the grown space is not released unless the snapshot is deleted.

Note: In OneFS 8.2.0, CloudPools 2.0 supports write-back in a snapshot. See the scenario 3 (section [Scenario 3](#)) for details. However, CloudPools 2.0 does not support archiving and recalling files in the snapshot directory. Consider the case when there is already file data in a snapshot on a cluster running a OneFS release prior to OneFS 8.2.0. That data takes up storage space on the PowerScale cluster, and then the cluster is upgraded to OneFS 8.2.0. Because CloudPools 2.0 does not support archiving files in snapshots to the cloud, the storage space for this snapshot cannot be released when the cluster is upgraded.

If SyncIQ or NDMP backs up the SmartLink files, the mapping file data should be retrieved from the cloud using the backup copy of the SmartLink file. If the backup retention has not expired, the CDOs of the mapping file data cannot be deleted even though the snapshot has been deleted. The reason is that the SmartLink file backup still references the CDOs of the mapping file data. When the backup retention period has expired and the CDOs of the mapping file data are no longer used, the CDOs of the mapping file data are deleted. For more information about data retention, refer to section

Data retention. If SyncIQ or NDMP does not back up SmartLink files, the CDOs of the mapping file data are deleted after the snapshot is deleted.

Users can revert a snapshot or access snapshot data through the snapshots directory (/ifs/.snapshot). The main methods for restoring data from a snapshot are as follows:

- Revert a snapshot through the SnapRevert job.
- Restore a file or directory using Microsoft® Shadow Copy Client on Windows® or cp command on Linux.
- Clone a file from a snapshot (CloudPools does not support cloning a file from a snapshot).

For details on restoring snapshot data, see the administration guide [OneFS 8.2.0 Web Administration Guide](#). CloudPools does not support cloning a file from a snapshot. The other two methods for restoring data from a snapshot in a CloudPools environment are described as follows.

When using the SnapRevert job to restore data from a snapshot, it reverts a directory back to the state it was in when a snapshot was taken. For example, there is a /ifs/test directory including a regular.txt regular file, and a smartlink.txt SmartLink file that has its file data archived to the cloud. A snap01 snapshot is created on the /ifs/test directory, and updates are made on the two files. The regular.txt file is then archived to the cloud, and it is truncated to a SmartLink file. Then, the SmartLink file smartlink.txt is recalled and it is converted to a regular file. If the snapshot snap01 is restored, it overwrites the files in directory /ifs/test. The regular.txt file reverts to a regular file, and the smartlink.txt reverts to a SmartLink file. The directory /ifs/test is reverted to the state it was in when snap01 was taken.

When using Microsoft Shadow Copy Client on Windows or the cp command on Linux, the file data is retrieved from the cloud through SmartLink files in a snapshot. This copy operation will create new regular files. That means extra space is required for the new regular files restored from a snapshot.

Sparse files handling

CloudPools 2.0 provides a new sparse file format to improve handling of empty blocks. With this improvement, sparse zeros are not in CloudPools operations, which reduce network utilization and saves space on the cloud target.

Note: No cloud objects are written when archiving full sparse files (fully empty blocks).

Quota management

In OneFS 8.2.0, quotas present actual space consumed on the PowerScale cluster.

For example, there is a directory or user quota of 500 GB and it is reporting 400 GB used. 200 GB of files are archived from the PowerScale cluster to cloud. Moving data to the cloud reduces the quota's measured node space consumption. In OneFS releases prior to 8.2.0, the amount of data that has been archived to the cloud frees the quota. And the quota shows 200 GB (400 GB to 200 GB) used out of 500 GB. That means the user or directory quota can exceed the set limit (500 GB). In OneFS 8.2.0, the application logical size integrated with CloudPools 2.0 measures the true capacity consumption even if data is archived from the PowerScale cluster to the cloud. And the quota shows 400 GB used

out of 500 GB through the application logical size. That means the user or directory quota cannot exceed the set limit of 500 GB.

For more information about the new SmartQuota reporting capabilities in OneFS 8.2.0, see the white paper [Storage Quota Management and Provisioning with Dell PowerScale SmartQuotas](#).

Anti-virus integration

In OneFS releases prior to OneFS 8.2.0, SmartLink files were skipped for anti-virus scanning.

In OneFS 8.2.0, CloudPools 2.0 provides a configurable option for anti-virus scanning of SmartLink files. The file data is retrieved from the cloud and cached on the cluster for the scan only if the option is enabled. The scan will be slower than normal. As shown in [Figure 10](#), the **Scan Cloudpool Files** option is configured and verified using the command line.

```
hop-isi-n-1# isi antivirus settings modify --scan-cloudpool-files=1
hop-isi-n-1# isi antivirus settings view
      Fail Open: Yes
      Glob Filters: -
Glob Filters Enabled: No
Glob Filters Include: No
      Path Prefixes: -
      Repair: Yes
      Report Expiry: 1Y
      Scan On Close: No
      Scan On Open: No
Scan Cloudpool Files: Yes
      Scan Size Maximum: 2.00G
      Service: No
      Quarantine: Yes
      Truncate: No
```

Figure 10. Enable Scan Cloudpool Files

Note: The Scan Cloudpool Files option is disabled by default, which means SmartLink files are skipped when scanning a directory which includes SmartLink files.

WORM integration

DellPowerScale SmartLock is an optional software feature of OneFS that enables SEC 17-a4 data compliance. In enterprise mode, individual directories can be set up as Write Once, Read Many (WORM) directories. And the data is immutable by everyone except the root account on the cluster once the files have been committed. A PowerScale cluster can also be set up in compliance mode where the root account on the cluster is removed. And no user can change or delete data in WORM-locked folders.

Prior to OneFS 8.2.0, SmartLink files are not allowed in both enterprise and compliance modes. In OneFS 8.2.0, details about CloudPools 2.0 and SmartLock integration are listed below:

- **Compliance mode:** SmartLink files are not allowed in compliance mode.
- **Enterprise mode:** SmartLink files are allowed in enterprise mode.
 - Enterprise mode can be enabled on a directory with SmartLink files.

- SmartLink files can be moved into an Enterprise mode directory which prevents modifying or deleting the SmartLink files.
 - SmartLink files can be recalled from the cloud to the PowerScale cluster once they are committed.

Best practices for PowerScale storage and AWS

This section focuses on the considerations and best practices for configuring PowerScale CloudPools and AWS.

PowerScale configuration

This section includes considerations and best practices for configuring PowerScale CloudPools.

CloudPools settings

CloudPools settings can be changed either on the CloudPools setting tab or on a per-file-pool policy from the OneFS WebUI. It is highly recommended to change these settings on a per-file-pool policy. The following list includes general considerations and best practices for CloudPools settings.

- **Encryption:** Encryption is an option that can be enabled either on the PowerScale cluster or on AWS. The recommendation is to enable encryption on the PowerScale cluster instead of on AWS. If the average CPU is high (greater than 70%) on the PowerScale cluster, the encryption can be enabled on AWS instead of on the PowerScale cluster. Encryption adds an additional load on the PowerScale cluster. Encryption can also impact the CloudPools archive and recall performance. For more information about protecting data using encryption on AWS, see the [AWS documentation](#).
- **Compression:** Compression is an option that can be enabled on the PowerScale cluster, in which file data is compressed before sending it to AWS. If network bandwidth is a concern, the recommendation is to enable compression on the PowerScale cluster to save network resources. Compression adds an additional load on the PowerScale cluster which means it might take more time to archive files from PowerScale storage to AWS.
- **Data retention:** The recommendation is to explicitly set the data retention for the file data being archived from the PowerScale cluster to AWS. If the SmartLink files are backed up with SyncIQ or NDMP, the data retention defines how long the cloud objects remain on AWS. When the retention period has passed, the PowerScale cluster sends a delete command to AWS. AWS marks the associated cloud objects for deletion. The delete process is asynchronous and the space is not reclaimed until garbage collection completes. This process is a low-priority background process, which may take days to fully reclaim the space, depending on how busy the system is.
- **Local data cache:** If the storage space is limited on the PowerScale cluster, the recommendation is to set lower values for the Writeback Frequency and Cache Expiration. This option reduces the time to keep file data in the local data cache and frees up storage space sooner on the PowerScale cluster.

File pool policy

File pool policies define what data will be archived from the PowerScale cluster to AWS. The considerations are listed below:

- Ensure that the priority of file pool policies is set appropriately. Multiple file pool policies can be created for the same cloud storage account. When the SmartPools job runs, it processes file pool policies in priority order.
- In terms of freeing up storage space on the PowerScale cluster, the recommendation is not to archive small files that are less than 32 KB in size.
- If the files need to be updated frequently, the recommendation is not to archive those files.
- OneFS supports a maximum of 128 file pool policies (SmartPools and CloudPools combined). The recommendation is not to exceed 30 file pool policies per PowerScale cluster.
- If the file pool policy is updated, it has no impact on the files already archived. It will only affect the files to be archived when the SmartPools job next runs.
- Archiving based on *accessed time*, rather than *modified* or *created* times, results in files that are used often, including applications, libraries, and scripts. Take care to exclude these types of files from being archived to the cloud, which would result in delays for clients or users loading these applications. One example is when you are archiving user home directories that contain files that are created once but accessed often.

Other considerations

More considerations include the following:

- **Deduplication:** CloudPools can archive deduped files from a PowerScale cluster to cloud storage. However, un-deduped files will be created when recalling those files from the cloud to the PowerScale cluster. For more information about deduplication within OneFS, see the white paper [Next Generation Storage Efficiency with Dell PowerScale SmartDedupe](#).
- **Small file storage efficiency (SFSE):** CloudPools and SFSE cannot work together. For PowerScale clusters using CloudPools, any SmartLink files cannot be containerized or packed. It is best practice to not archive small files that will be optimized using SFSE. The efficiencies gained from implementing SFSE for small files outweigh the storage advantages gained from archiving them to the cloud using CloudPools. For more information about the Small File Storage Efficiency feature of OneFS, see the white paper [Dell PowerScale OneFS Storage Efficiency](#).
- **Network proxy:** When a PowerScale cluster cannot connect to the CloudPool storage target directly, network proxy servers can be configured for an alternate path to connect to the cloud storage.
- **SmartConnect:** If user access SmartLink files regularly from a specific node, clogging the inline access path may impact client performance. You can configure PowerScale SmartConnect for load-balancing connections for the cluster. For more information about SmartConnect, see the white paper [Dell PowerScale Network Design Considerations](#).

- **Cloud storage account:** Do not delete a cloud storage account that is in use by archived files. Any attempt to open a SmartLink file associated with a deleted account will fail. In addition, NDMP backup and restore and SyncIQ failover and failback will fail when a cloud storage account has been deleted.
- **Cloud objects and data retention:** Cloud objects are crucial for SmartLink files. Any attempt to open a SmartLink file associated with deleted cloud objects will fail. OneFS checks data retention and the reference count for cloud objects before garbage collection. When data retention has expired and there is no reference count for cloud objects, cloud objects will be deleted through garbage collection. Data retention is a concept used to determine the Date of Death (DoD) setting for objects that support a SmartLink file. DoD is used to trigger garbage collection only if the reference count is zero for a file on the cluster only. The reference count is a concept used to determine whether cloud objects are associated with SmartLink files, including SmartLink files in the snapshots, SyncIQ backup, and NDMP backup. The considerations include:
 - Data retention periods include Cloud data retention period, Incremental backup retention period for NDMP incremental backup and SyncIQ, and Full backup retention period for NDMP only. If more than one period applies to a SmartLink file, the longest period is applied.
 - If a SmartLink file is unchanged through multiple SyncIQ backups or NDMP backups, its data retention will remain unchanged.
 - Data retention is set or updated on any event that changes the backed up version of a file or the state of the SmartLink file.
 - If a SmartLink file is changed and incrementally backed up, its data retention will be set by calculating the current time plus incremental backup retention period.
 - If a SmartLink file is recalled, the reference count will be removed, and its data retention will be set by calculating the current time plus cloud data retention period. Its cloud objects will be deleted through garbage collection after its data retention has expired.
 - If a SmartLink file is deleted, its data retention will be set by calculating the current time plus cloud data retention period. If cloud objects are still associated with snapshots, SyncIQ backup, or NDMP backup, its cloud objects will not be deleted through garbage collection after its data retention has expired.
- **OneFS upgrade (CloudPools 1.0 to CloudPools 2.0):** Before beginning the upgrade, it is recommended to check the OneFS CloudPools upgrade path shown in the following table.

Table 3. OneFS CloudPools upgrade path

Installed OneFS Version (CloudPools 1.0)	Upgrade to OneFS Version (CloudPools 2.0)			
	8.2.0	8.2.1 with May 2020 RUPs	8.2.2 with May 2020 RUPs	9.x
8.0.x or 8.1.x	Strongly discouraged	OK if needed but recommend 8.2.2	Strongly Recommended	Strongly Recommended

Note: Contact your Dell representative if you plan to upgrade OneFS to 8.2.0.

In a SyncIQ environment with unidirectional replication, the SyncIQ target cluster should be upgraded before the source cluster. The reason is that OneFS allows the CloudPools-1.0-formatted SmartLink files to be converted into CloudPools-2.0-formatted SmartLink files through a post-upgrade SmartLink conversion process. Otherwise, a SyncIQ policy needs to be reconfigured to deep copy but deep copy will cause archived file content to read from the cloud and be replicated. In a SyncIQ environment with bi-directional replication, it is recommended to disable SyncIQ on both source and target clusters and upgrade both source and target clusters simultaneously. You can then reenabling SyncIQ on both source and target clusters when the OneFS upgrades have been committed on both source and target clusters. Depending on the number of SmartLink files on the target DR cluster and the processing power of that cluster, the SmartLink conversion process can take considerable time.

Note: There is no need to stop SyncIQ and Snapshot during the upgrade in a SyncIQ environment with unidirectional replication. Because SyncIQ must resynchronize all converted stub files, it may take SyncIQ some time to catch up with all the changes.

To check the status of the SmartLink upgrade process, run the following command, substituting the appropriate job number.

```
# isi cloud job view 6
ID: 6
Description: Update SmartLink file formats
Effective State: running
Type: smartlink-upgrade
Operation State: running
Job State: running
Create Time: 2019-08-23T14:20:26
State Change Time: 2019-09-17T09:56:08
Completion Time: -
Job Engine Job: -
Job Engine State: -
Total Files: 21907433
Total Canceled: 0
Total Failed: 61
Total Pending: 318672
Total Staged: 0
Total Processing: 48
Total Succeeded: 21588652
```

Note: CloudPools recall jobs will not run while SmartLink upgrade or conversion is in progress.

For a Not All Nodes on Network (NANON) cluster, it is recommended to get the unconnected nodes connected to the network before starting the SmartLink conversion. Also, you need disable SnapDelete until the SmartLink conversion is completed.

AWS configuration

Before configuring PowerScale CloudPools on the PowerScale cluster, AWS needs to be configured properly. The following are the general considerations and best practices when configuring AWS for CloudPools.

- **URI for CloudPools:** The URI is region-specific. For example, the URI for region us-west-1 is <https://s3.us-west-1.amazonaws.com>. The mapping region needs to be set when configuring the CloudPools on the PowerScale cluster. For more details about AWS Regions and Endpoints, see the document [AWS Regions and Endpoints](#) on the AWS website. You need to ensure that the default URI is used for CloudPools. You cannot add any prefix in the URI, such as `https://prefix.s3.us-west-1.amazonaws.com`. Starting with OneFS 9.13, CloudPools now supports Amazon VPC endpoints. For more details about Amazon VPC, see the document [Amazon VPC](#).
- **CloudPools support of Amazon S3:** CloudPools supports S3 Standard, S3 Intelligent-Tiering, S3 Standard-IA, and S3 One Zone-IA. CloudPools does not support S3 Glacier Instant Retrieval, S3 Glacier Flexible Retrieval, S3 Glacier Deep Archive, and S3 Outposts. For more details about Amazon S3 storage classes, see the document [Amazon S3 Storage Classes](#).
- **Identity and access management (IAM):** An IAM user needs to be created with proper permissions before setting up CloudPools on a PowerScale cluster. CloudPools uses the IAM user to manage buckets and objects for CloudPools operations. The policy AmazonS3FullAccess needs to be attached to the IAM user. The secret key can be created following the process of [Managing Access Keys for IAM Users](#).
- **Cost and usage report:** The cost and usage report need to be enabled in AWS. See the document [Creating an AWS Cost and Usage Report](#) for details. You can specify an existing bucket or create one for the report. The bucket is used as the telemetry reporting bucket when configuring CloudPools and OneFS will generate XML files to track the usage data for AWS.
- **Commercial cloud services (C2S):** The client certificate must be granted authorization to access specific IAM roles within one or more C2S accounts. Interaction with C2S Access Portal requires an X.509 client certificate signed by an appropriate Government Certificate Authority (CA). Work with your government contact to obtain a client certificate.

Protecting SmartLink files

Because SmartLink files are the sole means to access file data stored in AWS, it is important to protect them from accidental deletion.

This section describes using PowerScale SyncIQ and NDMP to back up SmartLink files.

Note: SmartLink files cannot be backed up using a copy command, such as secure copy (scp).

SyncIQ

SyncIQ is CloudPools-aware, but consider the section [Snapshot efficiency](#) for guidance, especially where snapshot retention periods on the target cluster will be long.

SyncIQ policies support two types of data replication for CloudPools:

- **Shallow copy:** This option is used to replicate files as SmartLink files without file data from the source PowerScale cluster to target PowerScale cluster.
- **Deep copy:** This option is used to replicate files as regular files or unarchived files from the source PowerScale cluster to target PowerScale cluster.

For information about cross-version compatibility of CloudPools, refer to the section [NDMP and SyncIQ support](#).

SyncIQ, SmartPools, and CloudPools licenses are required on both the source and target PowerScale cluster. It is highly recommended to set up a scheduled SyncIQ backup of the SmartLink files. For more information about PowerScale SyncIQ, see the white paper [Dell PowerScale SyncIQ: Architecture, Configuration, and Considerations](#).

When SyncIQ replicates SmartLink files, it also replicates the local cache state and unsynchronized cache data from the source PowerScale cluster to the target PowerScale cluster. The following figure shows the SyncIQ replication when replicating a directory including SmartLink files and unarchived normal files. Both unidirectional and bi-directional replication are supported. Appendix A provides steps for failing over to a secondary PowerScale cluster and failing back to a primary PowerScale cluster.

Note: OneFS manages cloud access at the cluster level and does not support managing cloud access at the directory level. You need to remove cloud access on the source cluster and add cloud access on the target cluster when failing over a SyncIQ directory containing SmartLink files to a target cluster. If there are multiple CloudPools storage accounts, removing/adding cloud access will impact all CloudPools storage accounts on the source/target cluster.




Figure 11. SyncIQ replication

Note: If encryption is enabled in a file pool policy for CloudPools, SyncIQ also replicates all the relevant encryption keys to the secondary PowerScale cluster along with the SmartLink files.

NDMP

NDMP is also CloudPools-aware and supports three backup and restore methods for CloudPools:

- **DeepCopy:** This option is used to back up files as regular files or unarchived files. Files can only be restored as regular files.
- **ShallowCopy:** This option is used to back up files as SmartLink files without file data. Files can only be restored as SmartLink files.
- **ComboCopy:** This option is used to back up files as SmartLink files with file data. Files can be restored as regular files or SmartLink files.

For information about cross-version compatibility of CloudPools, refer to the section [NDMP and SyncIQ support](#).

It is possible to update the file data and send the updated data to the cloud storage. Multiple version SmartLink files can be backed up to tapes using NDMP, and multiple versions of CDOs are protected on AWS under the data retention setting. You can restore a specific version of a SmartLink file from tapes to a PowerScale cluster and continue to access (read or update) the file as before.

Note: If encryption is enabled in the file pool policy for CloudPools, NDMP also backs up all the relevant encryption keys to tapes along with the SmartLink files.

Performance

CloudPools is designed to move cold data from primary storage to the cloud. It is deliberately slow to ensure that it does not compete with processes that are performance sensitive such as SMB and NFS user activity. By default, CloudPools is using 10 threads per node which balances CloudPools CPU usage with other cluster functions. It is recommended to use the default number of threads for typical workloads. CloudPools does provide an option to modify the number of archive and recall threads. However, modifying the number of archive and recall threads can improve archive and recall performance but can also have significant impact on the CPU load of your system.

Note: Contact your Dell representative if you want to configure a higher number of threads.

CloudPools archive and recall performance depend on many factors, such as the network bandwidth between the PowerScale cluster and the cloud, available system resources, and file size. These performance considerations include:

- As the file size increases, the CloudPools archive and recall performance also increases. However, the effect on archive performance is minor when file size is greater than or equal to 10 MB. The effect on recall performance is negligible when file size is greater than or equal to 10 MB.
- As the thread counts increase, the CloudPools archive and recall performance also increases. However, the effect on archive and recall performance is negligible when the number of threads per PowerScale node is greater than or equal to 40.
- For a single large file, the effect on archive and recall performance is negligible regardless of the file size or number of threads per node. A single thread manages a single file transfer on a single node.
- Starting from OneFS 9.3.0.0, CloudPools creates Likewise sparks to drive the read of each CDO from the cloud. This enhancement can cache multiple chunks or CDOs of a stub file concurrently to improve CloudPools read and recall performance.

- The effect on archive and recall performance is negligible regardless of the number of file pool policies or jobs.
- With the expansion of PowerScale nodes, CloudPools archive and recall performance increases, but not linearly.
- For a single, heterogeneous cluster, tier 1 (fast) node pool has a minor impact on CloudPools archive performance and a large impact on CloudPools recall performance. The archive and recall performance are better when data is stored in the tier 1 node pool. The setting **Data Storage Target** of a file pool policy can determine the node pool for recall. However, the node pool cannot be changed for inline read. The node pool for a stub is used for inline read for this stub.
- Not All Nodes on Network (NANON) cluster has a large impact on CloudPools archive and recall performance.

Reporting

This section describes reporting for CloudPools network stats and includes the following topics:

- CloudPools network stats
- Query network stats by CloudPools account
- Query network stats by file pool policy
- Query history network stats

CloudPools network stats introduction

CloudPools network stats collect every network transaction and provide network activity statistics from connections to the cloud storage. The network activity statistics include bytes In, bytes Out, and the number of GET, PUT, and DELETE operations. CloudPools network stats are available in two categories:

- Per CloudPools account
- Per file pool policy


Note: CloudPools network stats do not provide file statistics, such as the file list being archived or recalled.

Query network stats by CloudPools account

Use the following command to check the CloudPools network stats by CloudPools account.

```
isi_test_cpool_stats -Q --accounts <account_name>
```

Figure 12 shows an example of current CloudPools network stats by CloudPools account.



Account Name	Bytes In	Bytes Out	Num Reads	Num Writes	Num Deletes
testaccount	4194896000	4194905034	4000	2001	8001

Figure 12. Network stats by CloudPools account

Query network stats by file pool policy

Use the following command to check the CloudPools network stats by file pool policy.

```
isi_test_cpool_stats -Q --policies <policy_name>
```

Figure 13 shows an example of current CloudPools network stats by file pool policy.

```
hop-isi-p-1# isi_test_cpool_stats -Q --policies testpolicy
```

Policy Name	Bytes In	Bytes Out	Num Reads	Num Writes
ecspolicy	4194896000	4194905034	4000	2001

Figure 13. Network stats by file pool policy

Note: The command output does not include the number of deletes by file pool policy.

Query history network stats

Use the following command to check the history CloudPools network stats.

```
isi_test_cpool_stats -q -s <number of seconds in the past to start stat query>
```

Use the s parameter to define the number of seconds in the past. For example, set it as 86,400 to query CloudPools network stats over the last day.

Figure 14 shows an example of CloudPools network stats over the last day.

```
hop-isi-p-1# isi_test_cpool_stats -q -s 86400
```

Account	bytes-in	bytes-out	gets	puts	deletes
testaccount	4194896000	4194905034	4000	2001	8001

Figure 14. Network stats last day

Use the following command to flush stats from memory to database and get the latest CloudPools history network stats.

```
isi_test_cpool_stats -f
```

Cloud statistics namespace with CloudPools

The cloud statistics namespace with CloudPools is added in OneFS 9.4.0.0. This feature leverages existing OneFS daemons and systems to track statistics about CloudPools activities. The statistics include bytes In, bytes Out, and the number of Reads, Writes, and Deletions. CloudPools statistics are available in two categories:

- Per CloudPools account
- Per file pool policy

Note: The cloud statistics namespace with CloudPools do not provide file statistics, such as the file list being archived or recalled.

You can use the `isi statistics cloud` command to view statistics about CloudPools activities. For more information about the `isi statistics cloud` command, see the document [PowerScale OneFS 9.4.0.0 CLI Command Reference](#).

Commands and troubleshooting

This section describes CloudPools commands and troubleshooting methodologies.

Commands

This CloudPools operations and job monitoring commands discussed in this section include the following:

- CloudPools archive
- CloudPools recall
- CloudPools monitoring

CloudPools archive

Run the following command to archive files from a PowerScale cluster to the cloud on demand.


```
isi cloud archive <file name> --recursive [true | false] --policy <policy name>
```

Parameters:

- <file name>: File name to be archived
- --recursive: Whether the archive should apply recursively to nested directories
- --policy: Policy name to be used with archiving

Run either of the following two commands to check whether the file is a SmartLink file or not, as shown in [Figure 15](#).

```
ls -loh <file name>
isi get -DD <file name> | grep -i smartlink
```



```
hop-isi-n-1# ls -loh test01.mp3
-rwx-----+ 1 root wheel uarch,inherit,writcache,wcinherit,ssmartlinked,shasntfsacl 339K Mar 14 0
5:20 test01.mp3
hop-isi-n-1# isi get -DD test01.mp3 | grep -i smartlink
* SmartLinked: True
hop-isi-n-1#
```

Figure 15. SmartLink file

CloudPools recall

Run the following command to recall files from the cloud to a PowerScale cluster on demand.

```
isi cloud recall <files> --recursive [true | false]
```

Parameters:

- <file name>: File name to be archived
- --recursive: Whether the archive should apply recursively to nested directories


CloudPools job monitoring

To check the CloudPools job status, use the following command.

```
isi cloud jobs list
```

To check the archive or recall file list status for a specific CloudPools job, use the following command. As shown in [Figure 16](#), the job id can be found using the command `isi cloud jobs list`.

```
isi cloud jobs files list <job id>
```



```
hop-isi-n-2# isi cloud jobs files list 219
Name                                     State
-----
/ifs/ecs/InsightIQ_b.4.1.2.7.zip completed
/ifs/ecs/a.pptx                        completed
-----
Total: 2
```

Figure 16. File list of specific CloudPools job

Note: The output of the prior command only shows the file name and state for specific CloudPools job.

To perform additional actions, run the following commands:

- Pause a CloudPools job:

```
isi cloud jobs pause <job id>
```

- Resume a paused CloudPools job:

```
isi cloud jobs resume <job id>
```

- Cancel a CloudPools job:

```
isi cloud jobs cancel <job id>
```

- Check the file list state of writing updated data to the cloud (job id is 1), which is an internal CloudPools job and always running:

```
isi cloud jobs files list 1
```

Note: The CloudPools system jobs should not be paused except temporarily for troubleshooting. No jobs should be left paused for an indefinite time.

Troubleshooting

This section describes various CloudPools troubleshooting methodologies, which include:

- CloudPools state
- CloudPools logs

CloudPools state

To check the CloudPools storage account state, use the following command:

```
isi cloud accounts view <cloudpools storage account name>
```

To check the CloudPool state, use the following command:

```
isi cloud pools view <cloud pool name>
```

To check the file pool policy state, use the following command:

```
isi filepool policies view <filepool policy name>
```

CloudPools logs

Check the CloudPools log if needed. The location of CloudPools log is as follows:

- Most normal daemon log is at /var/log/isi_cpool_d.log
- The log of IO to the cloud is at /var/log/isi_cpool_io_d.log
- Key management log is at /var/log/isi_km_d.log
- CloudPools job (Job Engine) log is at /var/log/isi_job_d.log

Appendix A: Step-by-step configuration example

This section describes a step-by-step configuration example for CloudPools and AWS and includes the following topics:

- AWS
- PowerScale configuration
- SmartLink files and cloud data protection

AWS configuration

This section describes the AWS configuration for CloudPools. S3 or C2S S3 can be used as the cloud target, which includes the following:

- S3
- C2S S3

The example AWS configuration is a general guide when AWS is used for CloudPools. It does not cover all details of AWS configuration for other use cases. Consult the [Amazon Web Services \(AWS\)](#) documentation for more details on AWS configuration.

S3

This section describes how to collect the information about S3 for CloudPools.

1. Ensure your AWS account is working properly.
2. Log in to the console of AWS at <http://aws.amazon.com> using your own username and password. Write down the URI and region to connect to Amazon S3. For example, the URI is <https://s3.us-west-1.amazonaws.com> and region is US-west-1.
3. Follow the process in the document [Creating an IAM User in Your AWS Account](#) on the AWS website to create an IAM user and assign proper permissions for CloudPools.
4. Follow the process in the document [Managing Access Keys for IAM Users](#) on the AWS website to create the secret key for the IAM user for CloudPools. As shown in [Figure 17](#), you can see the Access key ID and Secret access key.




Figure 17. Access Key

5. Go to **My Account > Account Settings** from the console of AWS and write down the Account ID as shown in Figure 18.




Figure 18. Account ID

6. Follow the process in the document [Creating an AWS Cost and Usage Report](#) on the AWS website to enable the cost and usage report. Write down the bucket name which will be used as the telemetry reporting bucket when configuring CloudPools.

Now all Amazon S3 information is gathered for CloudPools.

C2S S3

This section describes how to collect the information about C2S S3 for CloudPools.

1. Ensure your AWS account, and C2S S3 are working properly.
2. Log in to the console of AWS at <http://aws.amazon.com> using your own username and password. Write down the URI and region to connect to C2S S3.
3. Before configuring the CloudPools C2S account, the Certificate Authority (CA) certificate and C2S Access Portal (CAP) client certificate and private key need to be gathered on CAP server. You can find the three parts -----BEGIN CERTIFICATE-----<<CERT DATA TRUNCATED>>-----END CERTIFICATE----- for the CA certificate, CAP client certificate, and private key and save it into three files on OneFS. For example, the files are capCA.pem, capClientcert.pem, and capClientcert.key.

4. Run the following command to import CA certificate on OneFS.

```
isi certificate authority import --name <name> <CA
certificate path>
```

5. Run the following command to verify the CA certificate.

```
isi certificate authority list
```

6. Run the following command to import CAP client cert and private key on OneFS. This information will be used when configuring the CloudPools C2S account.

```
isi cloud certificates import --name <name> <CAP client cert
path> <CAP client key path>
```

7. Run the following command to verify the CAP client certificate and private key.

```
isi cloud certificates list
```

- Write down the credential provider URI, Agency, Mission, and Role found from the CAP server.

Now all Amazon C2S S3 information is gathered for CloudPools.

PowerScale configuration

This section describes the CloudPools configuration on a PowerScale cluster, which includes the following:

- Verify licensing
- Cloud storage account creation
- CloudPool creation
- File pool policy creation
- Run SmartPools job for CloudPools
- SyncIQ policy creation

Verify licensing

This section describes how to verify licensing on the PowerScale system.

- Log in to the OneFS WebUI and go to **Cluster Management > Licensing** as shown in [Figure 19](#).
- Verify that the CloudPools and SmartPools license status is **Activated**.




Figure 19. Verifying licenses

Cloud storage account for S3

This section describes how to create a cloud storage account for S3 on a PowerScale cluster.

- Log in to the OneFS WebUI and go to **File System > Storage Pools**. Click **CloudPools** as shown in [Figure 20](#).




Figure 20. CloudPools

2. Click the **+ Create a Cloud Storage Account** button from the **Create a Cloud Storage Account** page as shown in [Figure 21](#). The minimum information for CloudPools and Amazon S3 is as follows:
 - **Name or alias:** Type a name to identify the cloud storage account.
 - **Type:** Select Amazon S3.
 - **URI:** Type the URI to connect AWS. For example, URI is `https://s3.us-west-1.amazonaws.com`.
 - **User name (key):** Type the Access key ID of the specific IAM user for CloudPools.
 - **Key (secret key):** Type the secret access key of the specific IAM user for CloudPools.
 - **Account ID:** Type the Account ID gathered on AWS console.
 - **Telemetry reporting bucket:** Type the bucket name created on AWS console.
 - **Storage region:** Select the mapping region for the URI.

Create a cloud storage account

* = Required field

Help ?

Connection Settings

* Name or alias

s301

* Type

Amazon S3

Cloud account information

* URI

https://s3.us-west-1.amazonaws.com

* User name (key)

* Key (secret key)

.....

Proxy

No value

☐ Skip SSL certificate validation (not recommended)

* Account ID

* Telemetry reporting bucket

* Storage region

N. California (US West 1)

Cancel

Connect account

Figure 21. Create a cloud storage account

- Click **Connect account** to create a cloud storage account. This operation results in two buckets being created in Amazon S3. One bucket will start with a d as a container to store the CDOs, and the other will start with an m as a container to store the associated metadata.

Cloud storage account for C2S S3

This section describes how to create a cloud storage account for C2S S3 on a PowerScale cluster.

- Log in to the OneFS WebUI and go to **File System > Storage Pools**. Click **CloudPools** as shown in [Figure 20](#).
- Click the **+ Create a Cloud Storage Account** button from the **Create a Cloud Storage Account** page as shown in [Figure 22](#). The minimum information for CloudPools and Amazon C2S S3 is as follows:
 - Name or alias:** Type a name to identify the cloud storage account.
 - Type:** Select Amazon C2S S3.

- **Cloud account information** – URI: Type the URI to connect AWS. For example, URI is https://s3.amazonaws.com.
- **Storage region**: Select the mapping region for the URI.
- **Credential server information** – URI: Type the credential provider URI used in C2S.
- **Agency**: Type the credential provider agency used in C2S.
- **Mission**: Type the credential provider mission used in C2S.
- **Role**: Type the credential provider role used in C2S.
- **Certificate**: Type the credential provider certificate used in C2S.

Create a cloud storage account

* = Required field

Connection settings

* Name or alias
c2s_s3

* Type
Amazon C2S S3

Cloud account information

* URI
https://s3.amazonaws.com

Proxy
No value

☐ Skip SSL certificate validation (not recommended)

* Storage region
default

Credential server information

* URI
[Empty text box]

* Agency
[Empty text box]

* Mission
[Empty text box]

* Role
[Empty text box]

* Certificate
cap

Proxy
No value

Cancel Connect account

Figure 22. Create a cloud storage account

3. Click **Connect account** to create a cloud storage account. This operation results in two buckets being created in Amazon C2S S3. One bucket will start with a “d” as a

container to store the CDOs, and the other will start with an “m” as a container to store the associated metadata.

CloudPool for S3

This section describes how to create a CloudPool for S3 on a PowerScale cluster.

1. Log in to the OneFS WebUI and go to **File System > Storage Pools**. Click **CloudPools** as shown in Figure 20.
2. Click the **+ Create a CloudPool** button from the **Create a CloudPool** page as shown in Figure 23. The minimum information is as follows:
 - **Name:** Type a name to identify the CloudPool.
 - **Type:** Select Amazon S3.
 - **Account in CloudPool:** Select the cloud storage account.

Create a CloudPool Help ?

* = Required field

— CloudPool settings —

State
Enabled

State details

* Name
PoolS3

* Type
Amazon S3

Vendor

Description

* Account in CloudPool
s301

Cancel Create a CloudPool

Figure 23. Create a CloudPool

3. Click **Create a CloudPool** to create a CloudPool.

CloudPool for C2S S3

This section describes how to create a CloudPool for Amazon C2S S3 on a PowerScale cluster.

1. Log in to the OneFS WebUI and go to **File System > Storage Pools**. Click **CloudPools** as shown in Figure 20.
2. Click the **+ Create a CloudPool** button from the **Create a CloudPool** page as shown in Figure 24. The minimum information is as follows:
 - **Name:** Type a name to identify the CloudPool.

- **Type:** Select Amazon C2S S3.
- **Account in CloudPool:** Select the cloud storage account.

Create a CloudPool

* = Required field

— CloudPool settings —

State
Enabled

State details

* Name
PoolC2S

* Type
Amazon C2S S3

Vendor

Description

* Account in CloudPool
c2s_s3

Cancel Create a CloudPool

Figure 24. Create a CloudPool

3. Click **Create a CloudPool** to create a CloudPool.

File pool policy

This section describes how to create a file pool policy on a PowerScale cluster.

1. Log in to the OneFS WebUI and go to **File System > Storage Pools**. Click **File Pool Policies** as shown in Figure 25.

Dashboard Cluster management File system Data protection Access Protocols

Storage pools

Summary **File pool policies** SmartPools CloudPools SmartPools settings CloudPools settings

File pool policies + Create a file pool policy

Bulk actions

Order	Policy name	CloudPools state	Description	Actions
-------	-------------	------------------	-------------	---------

Figure 25. Create a file pool policy

2. Click **+ Create a File Pool Policy** from the **Create a file pool policy** page as shown in Figure 26 and Figure 27. The minimum information is as follows:
 - **Policy Name:** Type a name to identify the file pool policy.
 - **File Matching Criteria:** Define a logical group of files for CloudPools. See the section [File-matching criteria](#).

- **Move to cloud storage:** Select the specific CloudPool as the CloudPool storage target.
- **Data retention settings:** Set the data retention as your own. See the section [Data retention](#).

Create a file pool policy [Help](#) ?

* = Required field

— Description —

* Policy name
test

Description

— Select files to manage —

Specify criteria for determining which files will be managed by this policy:

File matching criteria

IF condition

Remove criteria Path Select operator Enter value Case sensitive

Type all or part of a path. Filenames and paths can contain any characters except backslash and space. Brackets [] are accepted. Note: '*' will never match '/'

Path

File type

+ Add an "Or" condition

File attribute

Modified

Accessed

Metadata changed

Created

Size

Apply SmartPools as default

Storage settings

☐ Move to storage pool or tier

Cancel Create policy

Figure 26. Create a file pool policy

Figure 27. Create a file pool policy (continued)

3. Click **Create policy** to create a file pool policy.

Run SmartPools job for CloudPools

This section describes how to run the SmartPools job for CloudPools on a PowerScale cluster.

1. Log in to the OneFS WebUI and go to **Cluster management > Job operations**. Click **Job types** as shown in Figure 28.

Figure 28. Job types

2. Select the **SmartPools** item and click **Edit** as shown in Figure 29.

Figure 29. SmartPools job

3. From the **Edit job type** details page as shown in [Figure 30](#), you can do the following:
 - Enable or disable the job
 - Set the priority of the job
 - Set the impact policy
 - Set the job schedule as manual or scheduled as your own

The screenshot shows the 'Edit job type details' window for a job type named 'SmartPools'. The window has a title bar with 'Edit job type details' and a 'Help' icon. The main content area is titled 'Job type details' and contains the following fields:

- Name:** SmartPools
- Description:** Enforce SmartPools file policies. This job requires a SmartPools license.
- Enable this job type:** A checked checkbox.
- Default priority:** A dropdown menu set to '6'.
- Default impact policy:** A dropdown menu set to 'LOW'.
- Schedule:** A section with two radio buttons: 'Manual' (unselected) and 'Scheduled' (selected). Below the 'Scheduled' radio button is a dropdown menu set to 'Daily'.
- Daily schedule:** A section with the following options:
 - Run policy every:** A dropdown menu set to '1' and a dropdown menu set to 'Day(s)'.
 - Run one policy per specified day:** A selected radio button.
 - Run policy at:** A dropdown menu set to '10:00 PM'.
 - Run multiple policies per specified day:** An unselected radio button.

At the bottom of the window are two buttons: 'Cancel' and 'Save changes'.

Figure 30. Edit job type details

4. Click **Start job** as shown in [Figure 29](#) to run the file pool policy to archive files from the PowerScale cluster to AWS. If you want to start a specific file pool policy job manually, refer to the section [Commands and troubleshooting](#).

SyncIQ policy

This section describes how to create a SyncIQ policy on a PowerScale cluster.

1. Log in to the OneFS WebUI and go to **Cluster Management > Licensing** as shown in [Figure 19](#). Verify that the CloudPools, SmartPools, and SyncIQ license status are **Activated**.
2. Go to **Data Protection > SyncIQ > Policies** and click the **+ Create a SyncIQ policy** button as shown in [Figure 31](#) and [Figure 32](#). The minimum information is as follows.
 - **Policy name:** Type a name to identify the policy name.

- **Source root directory:** Type the directory name from source PowerScale cluster you want to replicate to the target PowerScale cluster.
- **Target host:** Type the IP or name of the target PowerScale cluster.
- **Target directory:** Type the directory name from the target PowerScale cluster you want to store the data replicated from the source PowerScale cluster.
- **Deep copy for CloudPools:** Select the type you want to use.

The screenshot shows the 'Create SyncIQ policy' window. It has a title bar with 'Create SyncIQ policy' and a 'Help' icon. Below the title bar is a legend: '* = Required field'. The main content area is divided into sections. The 'Settings' section contains a 'Policy name' text box with 'CPPolicyDR' entered, a 'Description' text box, and a checked checkbox for 'Enable this policy'. Below this is the 'Action' section with radio buttons for 'Copy' and 'Synchronize' (selected). The 'Run job' section has radio buttons for 'Manually' (selected), 'On a schedule', 'Whenever the source is modified', and 'Whenever a snapshot of the source directory is taken'. The 'Source cluster' section contains a 'Source root directory' text box with '/ifs/ecs' entered and a 'Browse...' button. Red circles highlight the 'Policy name' and 'Source root directory' fields.

Figure 31. Create SyncIQ policy

Target cluster

Target host
10

* Target directory
/ifs/ecs

☐ Connect only to nodes within the SmartConnect zone on the target cluster

Target snapshots

☐ Enable capture of snapshots on the target cluster

Snapshot alias name
Default name: SIO-%(SrcCluster)-%(PolicyName)-latest
SIO-%(SrcCluster)-%(PolicyName)

Snapshot naming pattern
Default pattern: SIO-%(SrcCluster)-%(PolicyName)-%Y-%m-%d_%H-%M-%S
SIO-%(SrcCluster)-%(PolicyName)-%Y-%m-%d_%H-%M

Snapshot expiration
☒ Snapshots do not expire
☐ Snapshots expire after...

Advanced settings

Priority
Normal (default)

Log level
Notice

☒ Validate file integrity

☐ Prepare policy for accelerated failback performance

Keep reports for
1 Year

Record deletions on synchronization
☒ Do not record when a synchronization deletes files or directories
☐ Record when a synchronization deletes files or directories

Deep copy for CloudPools
Deny

Cancel Create policy

Figure 32. Create SyncIQ policy (continued)

3. Click **Create policy** to create a SyncIQ policy.

SmartLink files protection

This section describes an example of how to protect SmartLink files and cloud data. Ensure that you have already configured SyncIQ on the PowerScale clusters, which include the following:

- Fail over to the secondary PowerScale cluster
- Fail back to the primary PowerScale cluster

Fail over to the secondary PowerScale cluster

This section describes the steps required to fail over to the secondary PowerScale cluster.

1. Log in to the **secondary** OneFS WebUI and go to **Data Protection > SyncIQ**. Click **Local Targets** on the policy that you want to failover and select **More > Allow Writes** as shown in Figure 33. This operation will grant read/write access to the data on the primary PowerScale cluster being replicated to the secondary PowerScale cluster.

SyncIQ




Figure 33. Allow writes on secondary cluster

Note: If the primary PowerScale cluster is still online, stop all writes to the replication policy's directory.

2. Check and change cloud access. Log in to the PowerScale clusters using SSH. To identify the CloudPools GUID, use the commands `isi cloud access list` and `isi cloud access view <GUID>`. Figure 34 shows the cloud access status on the secondary PowerScale cluster.

```
hop-isi-p-1# isi cloud access list
Name      GUID                                     Synced From  State
-----
hop-isi-n 006016894ae21826755c5a15e4a547aba6bb hop-isi-n    not permitted
hop-isi-p 006048509dlc6325875cd003f35f88a983de (current)    permitted
-----
Total: 2
hop-isi-p-1# isi cloud access view 006016894ae21826755c5a15e4a547aba6bb
Name: hop-isi-n
GUID: 006016894ae21826755c5a15e4a547aba6bb
Synced From: hop-isi-n
State: not permitted
Accounts: testaccount
Policies: ecspolicy
```

Figure 34. Identify CloudPools GUID to be transferred

3. On the primary PowerScale cluster, remove the cloud write permission using the command `isi cloud access remove <GUID>` as shown in Figure 35. This operation disables the file pool policy, CloudPool, and cloud storage account on the primary PowerScale cluster.

```
hop-isi-n-1# isi cloud access remove 006016894ae21826755c5a15e4a547aba6bb
Removing access to 006016894ae21826755c5a15e4a547aba6bb will disable the following CloudPool accounts and FilePool policies:
    testaccount (CloudPool Account)
    ecspolicy (FilePool Policy)
Are you sure?? (yes/[no]): yes
hop-isi-n-1#
```

Figure 35. Remove Cloud write access on the primary PowerScale cluster

4. On the secondary PowerScale cluster, add the cloud write permission using the command `isi cloud access add <GUID>` as shown in Figure 36. This operation enables the file pool policy, CloudPool, and cloud storage account on the secondary PowerScale cluster.

```

hop-isi-p-1# isi cloud access add 006016894ae21826755c5a15e4a547aba6bb
Giving access to 006016894ae21826755c5a15e4a547aba6bb will enable the following CloudPool accounts and FilePool policies:
    testaccount (CloudPool Account)
    ecspolicy (FilePool Policy)
Are you sure?? (yes/[no]): yes
To ensure proper cleanup, a job must be run for each S3 enabled account to set an expiration date for all stale cloud files.
Failure to set an expiration date will cause leaked data in the cloud resulting in additional costs from cloud service providers.
Note that after the expiration date has passed, backups may no longer be able to restore deleted files.
Expiration dates can be set later using the 'isi cloud restore-coi' command.
To start expiration date jobs for applicable accounts, enter an expiration date now or 'default' to accept the default expiration date (2029-03-20): <date>/[default]/cancel):
hop-isi-p-1#

```

Figure 36. Add Cloud write access on the secondary PowerScale cluster

Note: It is important to not allow write access to the CloudPools from more than one PowerScale clusters.

Now the SyncIQ failover is complete.

Fail back to primary PowerScale cluster

This section describes the steps required to fail back to the primary PowerScale cluster.

1. Log in to the **primary** OneFS WebUI and go to **Data Protection > SyncIQ**. Click **Policies** on the policy that you want to failback and select **More > Resync-prep** as shown in Figure 37. This operation creates a SyncIQ replication mirror policy on the secondary PowerScale cluster.

SyncIQ




Figure 37. Resync prep SyncIQ policy on primary PowerScale cluster

2. Log in to the **secondary** OneFS WebUI and go to **Data Protection > SyncIQ > Policies**. On the replication mirror policy that you want to failover and select **More > Start Job** as shown in Figure 38. This operation will sync any changes that have been written to the secondary PowerScale cluster back to the primary PowerScale cluster.

SyncIQ




Figure 38. Sync data from secondary PowerScale cluster to primary PowerScale cluster

3. Log in to the primary OneFS WebUI and go to **Data Protection > SyncIQ**. Click **Local Targets** on the policy that you want to failover and select **More > Allow Writes** as shown [Figure 39](#). This operation will grant read/write access to the replication directory back to the primary PowerScale cluster and change the secondary PowerScale cluster's access to this directory as read-only.

SyncIQ




Figure 39. Allow writes on primary PowerScale cluster to SyncIQ replication directory

Note: If the secondary PowerScale cluster is still online, stop all writes to the replication policy's directory. Perform a final replication from the secondary PowerScale cluster to the primary PowerScale cluster to ensure both sites are synchronized.

4. Check and change the cloud access. Log in to the PowerScale clusters using SSH. To identify the CloudPools GUID, use the commands `isi cloud access list` and `isi cloud access view <GUID>`. [Figure 40](#) shows the cloud access status on the secondary PowerScale cluster.

```

hop-isi-p-1# isi cloud access list
Name          GUID                               Synced From  State
-----
hop-isi-n 006016894ae21826755c5a15e4a547aba6bb hop-isi-n    permitted
hop-isi-p 006048509dlc6325875cd003f35f88a983de (current)    permitted
-----
Total: 2
hop-isi-p-1# isi cloud access view 006016894ae21826755c5a15e4a547aba6bb
Name: hop-isi-n
GUID: 006016894ae21826755c5a15e4a547aba6bb
Synced From: hop-isi-n
State: permitted
Accounts: testaccount
Policies: ecspolicy
hop-isi-p-1#

```

Figure 40. Identify GUID for CloudPools account and file pool policy

- On the **secondary** PowerScale cluster, remove the cloud write permission using the command `isi cloud access remove <GUID>` as shown in Figure 41. This operation disables the file pool policy, CloudPool, and cloud storage account on the secondary PowerScale cluster.

```

hop-isi-p-1# isi cloud access remove 006016894ae21826755c5a15e4a547aba6bb
Removing access to 006016894ae21826755c5a15e4a547aba6bb will disable the following CloudPool accounts and FilePool policies:
testaccount (CloudPool Account)
ecspolicy (FilePool Policy)
Are you sure?? (yes/[no]): yes

```

Figure 41. Remove cloud write access on the secondary PowerScale cluster

- On the **primary** PowerScale cluster, add the cloud write permission using the command `isi cloud access add <GUID>` as shown in Figure 42. This operation enables the file pool policy, CloudPool, and cloud storage account on the primary PowerScale cluster.

```

hop-isi-n-1# isi cloud access add 006016894ae21826755c5a15e4a547aba6bb
Giving access to 006016894ae21826755c5a15e4a547aba6bb will enable the following CloudPool accounts and FilePool policies:
testaccount (CloudPool Account)
ecspolicy (FilePool Policy)
Are you sure?? (yes/[no]): yes
To ensure proper cleanup, a job must be run for each S3 enabled account to set an expiration date for all stale cloud files.
Failure to set an expiration date will cause leaked data in the cloud resulting in additional costs from cloud service providers.
Note that after the expiration date has passed, backups may no longer be able to restore deleted files.
Expiration dates can be set later using the 'isi cloud restore-coi' command.
To start expiration date jobs for applicable accounts, enter an expiration date now or 'default' to accept the default expiration date (2029-03-20): (<date>/[default]/cancel):
hop-isi-n-1#

```

Figure 42. Give the primary PowerScale cluster cloud write access

Note: It is important to not allow write access to the CloudPools from more than one PowerScale cluster.

- Log in to the **secondary** OneFS WebUI and go to **Data Protection > SyncIQ**. Click **Policies** on the policy that you want to failback and select **More > Resync-prep**. This operation will disable the SyncIQ replication mirror policy on the secondary PowerScale cluster and place the secondary PowerScale cluster back into read-only mode. In addition, this operation will enable the SyncIQ replication policy on the primary PowerScale cluster.

The SyncIQ failback is complete.

References

Dell Technologies documentation

The following Dell Technologies documentation provides other information related to this document. Access to these documents depends on your login credentials. If you do not have access to a document, contact your Dell Technologies representative.

- [OneFS CloudPools Administration Guide](#)
- [OneFS Technical Overview](#)
- [Next Generation Storage Efficiency with Dell PowerScale SmartDedupe](#)
- [Dell PowerScale OneFS Storage Efficiency](#)
- [Dell PowerScale SyncIQ: Architecture, Configuration, and Considerations](#)
- [High Availability and Data Protection with Dell PowerScale Scale-out NAS](#)
- [Storage Quota Management and Provisioning with Dell PowerScale SmartQuotas](#)
- [PowerScale Non-Disruptive Upgrade \(NDU\) Best Practices](#)
- [Data Protection with Dell PowerScale SnapshotIQ](#)
- [Dell PowerScale: Network Design Considerations](#)
- [Amazon Web Services \(AWS\)](#)

Amazon documentation

See also the following Amazon documentation.

- [Amazon Web Services \(AWS\)](#)