



Reference

NetApp Backup and Recovery

NetApp
March 09, 2026

Table of Contents

- Reference 1
 - Policies in SnapCenter compared to those in NetApp Backup and Recovery 1
 - Schedule tiers 1
 - Multiple policies in SnapCenter with the same schedule tier 1
 - Imported SnapCenter daily schedules 1
 - Imported SnapCenter hourly schedules 2
 - Log retention from SnapCenter policies 2
 - Log backup retention 2
 - Retention count from SnapCenter policies 2
 - SnapMirror labels from SnapCenter policies 3
 - NetApp Backup and Recovery Identity and Access Management (IAM) roles 3
 - Restore NetApp Backup and Recovery configuration data in a dark site 3
 - Restore NetApp Backup and Recovery data to a new Console agent 3
 - Supported AWS archive storage tiers with NetApp Backup and Recovery 8
 - Supported S3 archival storage classes for NetApp Backup and Recovery 9
 - Restore data from archival storage 9
 - Supported Azure archive access tiers with NetApp Backup and Recovery 10
 - Supported Azure Blob access tiers for NetApp Backup and Recovery 10
 - Restore data from archival storage 10
 - Supported Google archive storage tiers with NetApp Backup and Recovery 11
 - Supported Google archival storage classes for NetApp Backup and Recovery 11
 - Restore data from archival storage 12

Reference

Policies in SnapCenter compared to those in NetApp Backup and Recovery

There are some differences between policies used in SnapCenter and those used in NetApp Backup and Recovery that might impact what you see after importing resources and policies from SnapCenter.

Schedule tiers

SnapCenter uses the following schedule tiers:

- **Hourly:** Multiple hours and minutes with any hours (0-23) and any minutes (0-60).
- **Daily:** Option to repeat every set number of days, for example, every 3 days.
- **Weekly:** Sunday to Monday, with an option to perform a snapshot on Day 1 of the week or on multiple days of the week.
- **Monthly:** January to December, with an option to perform on specific or multiple days each month, for example, the 7th.

NetApp Backup and Recovery uses the following schedule tiers, which are slightly different:

- **Hourly:** Performs snapshots only on 15-minute intervals, for example, 1 hour or 15-minute intervals less than 60.
- **Daily:** Hours of the day (0-23) with start time for example at 10:00 AM with an option to perform every so many hours.
- **Weekly:** Day of the week (Sunday to Monday) with an option to perform on 1 day or multiple days. This is the same as SnapCenter.
- **Monthly:** Dates of the month (0-30) with a starting time on multiple dates of the month.
- **Yearly:** Monthly. This matches SnapCenter's monthly.

Multiple policies in SnapCenter with the same schedule tier

You can assign multiple policies with the same schedule tier to a resource in SnapCenter. However, NetApp Backup and Recovery does not support multiple policies on a resource that uses the same schedule tier.

Example: If you use three policies (for Data, Log, and Log of snapshots) in SnapCenter, after migration from SnapCenter, NetApp Backup and Recovery uses a single policy instead of all three.

Imported SnapCenter daily schedules

NetApp Backup and Recovery adjusts the SnapCenter schedules as follows:

- If the SnapCenter schedule is set to less than or equal to 7 days, NetApp Backup and Recovery sets the schedule to weekly. Some snapshots are skipped during the week.

Example: If you have a SnapCenter daily policy with a repeating interval of every 3 days starting on Monday, NetApp Backup and Recovery sets the schedule to weekly on Monday, Thursday, and Sunday.

Some days will be skipped because it is not exactly every 3 days.

- If the SnapCenter schedule is set to greater than 7 days, NetApp Backup and Recovery sets the schedule to monthly. Some snapshots will be skipped during the month.

Example: If you have a SnapCenter daily policy with a repeating interval of every 10 days starting on the 2nd of the month, NetApp Backup and Recovery, after migration, sets the schedule to monthly on the 2nd, 12th, and 22nd day of the month. NetApp Backup and Recovery skips some days in the next month.

Imported SnapCenter hourly schedules

SnapCenter hourly policies with repeating intervals greater than one hour are converted to a daily policy in NetApp Backup and Recovery.

Any hourly policy with repeating intervals that are not a factor of 24 (for example 5, 7, etc) will skip some snapshots in a day.

Example: If you have a SnapCenter hourly policy with a repeating interval every 5 hours starting at 1:00 AM, NetApp Backup and Recovery (after migration) will set the schedule to daily with 5-hour intervals at 1:00 AM, 6:00 AM, 11:00 AM, 4:00 PM, and 9:00 PM. Some hours will be skipped, after 9:00 PM it should be 2:00 AM to repeat after every 5 hours, but it will be always 1:00 AM.

Log retention from SnapCenter policies

If you have a resource in SnapCenter with multiple policies, NetApp Backup and Recovery uses the following priority order to assign the log retention value:

- For "Full backup with log backup policy" plus "log-only" policies in SnapCenter, NetApp Backup and Recovery uses the log-only policy retention value.
- For "Full backup with log only" and "Full and Log" policies in SnapCenter, NetApp Backup and Recovery uses the log-only retention value.
- For "Full backup and log" plus "Full backup" in SnapCenter, NetApp Backup and Recovery uses the "Full backup and log" retention value.
- If you have only a full backup in SnapCenter, NetApp Backup and Recovery does not enable the log backup.

Log backup retention

SnapCenter supports multiple retention values for policies on a resource. NetApp Backup and Recovery supports only one retention value per resource.

Retention count from SnapCenter policies

If you have a resource with secondary protection enabled in SnapCenter with multiple source volumes, multiple destination volumes, and multiple SnapMirror relationships, NetApp Backup and Recovery uses only the first policy's retention count.

Example: If you have a SnapCenter policy with a retention count of 5 and another policy with a retention count of 10, NetApp Backup and Recovery uses the retention count of 5.

SnapMirror labels from SnapCenter policies

SnapCenter keeps SnapMirror labels for each policy after migration, even if the tier changes.

Example: An hourly policy from SnapCenter might change to daily in NetApp Backup and Recovery. However, the SnapMirror labels remain the same after migration.

NetApp Backup and Recovery Identity and Access Management (IAM) roles

NetApp Backup and Recovery employs Identity and Access Management (IAM) to govern the access that each user has to specific features and actions.

To learn about IAM roles that are specific to NetApp Backup and Recovery, refer to [NetApp Backup and Recovery roles in NetApp Console](#).

Restore NetApp Backup and Recovery configuration data in a dark site

When using NetApp Backup and Recovery in a site with no internet access, known as *private mode*, the NetApp Backup and Recovery configuration data is backed up to the StorageGRID or ONTAP S3 bucket where your backups are being stored. If you have an issue with the Console agent host system, you can deploy a new Console agent and restore the critical NetApp Backup and Recovery data.



This procedure applies only to ONTAP volume data.

When you use NetApp Backup and Recovery in a SaaS environment with the Console agent deployed at your cloud provider or on your own internet-connected host, the system backs up and protects all important configuration data in the cloud. If you have an issue with the Console agent, create a new Console agent and add your systems. The backup details are automatically restored.

There are two types of data that are backed up:

- NetApp Backup and Recovery database - contains a listing of all the volumes, backup files, backup policies, and configuration information.
- Indexed Catalog files - contains detailed indexes that are used for Search & Restore functionality that make your searches very quick and efficient when looking for volume data that you want to restore.

This data is backed up once per day at midnight, and a maximum of 7 copies of each file are retained. If the Console agent is managing multiple on-premises ONTAP systems, the NetApp Backup and Recovery files are stored in the bucket of the system that was activated first.



No volume data is ever included in the NetApp Backup and Recovery database or Indexed Catalog files.

Restore NetApp Backup and Recovery data to a new Console agent

If your on-premises Console agent stops working, you'll need to install a new Console agent, and then restore

the NetApp Backup and Recovery data to the new Console agent.

You'll need to perform the following tasks to return your NetApp Backup and Recovery system to a working state:

- Install a new Console agent
- Restore the NetApp Backup and Recovery database
- Restore the Indexed Catalog files
- Rediscover all of your on-prem ONTAP systems and StorageGRID systems to the NetApp Console UI

After you check that your system is working, create new backup files.

What you'll need

You'll need to access the most recent database and index backups from the StorageGRID or ONTAP S3 bucket where your backup files are being stored:

- NetApp Backup and Recovery MySQL database file

This file is located in the following location in the bucket `netapp-backup-<GUID>/mysql_backup/`, and it is named `CBS_DB_Backup_<day>_<month>_<year>.sql`.

- Indexed Catalog backup zip file

This file is located in the following location in the bucket `netapp-backup-<GUID>/catalog_backup/`, and it is named `Indexed_Catalog_DB_Backup_<db_name>_<day>_<month>_<year>.zip`.

Install a new Console agent on a new on-premises Linux host

When installing a new Console agent, download the same software version as the original agent. Changes to the NetApp Backup and Recovery database may cause newer software versions to not work with old database backups. You can [upgrade the Console agent software to the most current version after restoring the Backup database](#).

1. [Install the Console agent on a new on-premises Linux host](#)
2. Log into the Console using the admin user credentials that you just created.

Restore the NetApp Backup and Recovery database

1. Copy the MySQL backup from the backup location to the new Console agent host. We'll use the example file name "CBS_DB_Backup_23_05_2023.sql" below.
2. Copy the backup into the MySQL docker container using one of the following commands, depending on whether you are using a Docker or Podman container:

```
docker cp CBS_DB_Backup_23_05_2023.sql ds_mysql_1:/. 
```

```
podman cp CBS_DB_Backup_23_05_2023.sql ds_mysql_1:/. 
```

3. Enter the MySQL container shell using one of the following commands, depending on whether you are

using a Docker or Podman container:

```
docker exec -it ds_mysql_1 sh
```

```
podman exec -it ds_mysql_1 sh
```

4. In the container shell, deploy the "env".
5. You'll need the MySQL DB password, so copy the value of the key "MYSQL_ROOT_PASSWORD".
6. Restore the NetApp Backup and Recovery MySQL DB using the following command:

```
mysql -u root -p cloud_backup < CBS_DB_Backup_23_05_2023.sql
```

7. Verify that the NetApp Backup and Recovery MySQL DB has been restored correctly using the following SQL commands:

```
mysql -u root -p cloud_backup
```

8. Enter the password.

```
mysql> show tables;  
mysql> select * from volume;
```

9. Ensure that the volumes that are shown are the same as those that existed in your original environment.

Restore the Indexed Catalog files

1. Copy the Indexed Catalog backup zip file (we'll use the example file name "Indexed_Catalog_DB_Backup_catalogdb1_23_05_2023.zip") from the backup location to the new Console agent host in the "/opt/application/netapp/cbs" folder.
2. Unzip the "Indexed_Catalog_DB_Backup_catalogdb1_23_05_2023.zip" file using the following command:

```
unzip Indexed_Catalog_DB_Backup_catalogdb1_23_05_2023.zip -d catalogdb1
```

3. Run the **ls** command to make sure that the folder "catalogdb1" has been created with the subfolders "changes" and "snapshots" underneath.

Discover your ONTAP clusters and StorageGRID systems

1. [Discover all the on-prem ONTAP systems](#) that were available in your previous environment. This includes the ONTAP system you have used as an S3 server.
2. [Discover your StorageGRID systems](#).


```

curl -X POST 'http://10.193.192.202/account/account-
DARKSITE1/providers/cloudmanager_cbs/api/v1/sg/credentials/working-
environment/OnPremWorkingEnvironment-pMtZND0M' \
> --header 'authorization: Bearer
eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCIsImtpZCI6IjJlMGFiZjRiIn0eyJzdWIiOiJvY
2NtYXV0aHwxIiwiaWF0IjpbImh0dHBzOi8vYXBpLmNsb3VkLm5ldGFwcC5jb20iXSwiaHR0c
DovL2Nsbn3Vklm5ldGFwcC5jb20vZnVsbF9uYW11IjoiYWRTaW4iLCJodHRwOi8vY2xvdWQub
mV0YXBwLmNvbS9lbWFPbCI6ImFkbWluQG5ldGFwcC5jb20iLCJzY29wZSI6Im9wZW5pZCBwc
m9maWx1IiwiaWF0IjoxNjcyNzIyNzEzNDQzMTMsImIzcyI6Imh0dHA6L
y9vY2NtYXV0aDo4NDIwLyJ9X_cQF8xttD0-S7sU2uph2cdu_kN-
fLWpdJJX98HODwPpVUitLcxV28_sQhuopjWobozPelNISf7KvMqcoXc5kLDyX-
yE0fH9gr4XgkdsWjWcNvw2rRkFzjHpWrETgfqAMkZcAukV4DHuxogHWh6-
DggB1NgPZT8A_szHinud5W0HJ9c4AaT0zC-
sp81GaqMahPf0KcFVybBL4krOewgKHGfo_7ma_4mF39B1LCj7Vc2XvUd0wCaJvDMjwp19-
KbZqmmBX9vDnYp7SSxC1hHJRdstcFgJLdJHtowweNH2829KsjEGBTtcBd08SvIDtctNH_GAx
wSgMT3zUfwaOimPw' \
> --header 'x-agent-id: vB_1xShPpBtUosjD7wfB1LIhqDgIPA0wclients' \
> -d '
> { "storage-server" : "sr630ip15.rtp.eng.netapp.com:10443", "access-
key": "2ZMYOAVAS5E70MCNH9", "secret-password":
"uk/6ikd4LjlXQOFnzSzP/T0zR4ZQlG0w1xgWsB" }'

```

Verify NetApp Backup and Recovery settings

1. Select each ONTAP system and click **View Backups** next to the Backup and recovery service in the right-panel.

You should see all backups created for your volumes.

2. From the Restore Dashboard, under the Search & Restore section, click **Indexing Settings**.

Make sure that the systems which had Indexed Cataloging enabled previously remain enabled.

3. From the Search & Restore page, run a few catalog searches to confirm that the Indexed Catalog restore has been completed successfully.

Supported AWS archive storage tiers with NetApp Backup and Recovery

NetApp Backup and Recovery supports two S3 archival storage classes and most regions.



To switch to and from NetApp Backup and Recovery UI versions, refer to [Switch to the previous NetApp Backup and Recovery UI](#).

Supported S3 archival storage classes for NetApp Backup and Recovery

When backup files are initially created they're stored in S3 *Standard* storage. This tier is optimized for storing data that's infrequently accessed; but that also allows you to access it immediately. After 30 days the backups transition to the S3 *Standard-Infrequent Access* storage class to save on costs.

If your source clusters are running ONTAP 9.10.1 or greater, you can choose to tier backups to either S3 *Glacier* or S3 *Glacier Deep Archive* storage after a certain number of days (typically more than 30 days) for further cost optimization. You can set this to "0" or to 1-999 days. If you set it to "0" days, you cannot change it later to 1-999 days.

Data in these tiers can't be accessed immediately when needed, and will require a higher retrieval cost, so you need to consider how often you may need to restore data from these archived backup files. See the section on this page about restoring data from archival storage.

- If you select no archive tier in your first backup policy when activating NetApp Backup and Recovery, then S3 *Glacier* will be your only archive option for future policies.
- If you select S3 *Glacier* in your first backup policy, then you can change to the S3 *Glacier Deep Archive* tier for future backup policies for that cluster.
- If you select S3 *Glacier Deep Archive* in your first backup policy, then that tier will be the only archive tier available for future backup policies for that cluster.

Note that when you configure NetApp Backup and Recovery with this type of lifecycle rule, you must not configure any lifecycle rules when setting up the bucket in your AWS account.

[Learn about S3 storage classes.](#)

Restore data from archival storage

While storing older backup files in archival storage is much less expensive than Standard or Standard-IA storage, accessing data from a backup file in archive storage for restore operations will take a longer amount of time and will cost more money.

How much does it cost to restore data from Amazon S3 Glacier and Amazon S3 Glacier Deep Archive?

There are 3 restore priorities you can choose when retrieving data from Amazon S3 Glacier, and 2 restore priorities when retrieving data from Amazon S3 Glacier Deep Archive. S3 Glacier Deep Archive costs less than S3 Glacier:

Archive Tier	Restore Priority & Cost		
	High	Standard	Low
S3 Glacier	Fastest retrieval, highest cost	Slower retrieval, lower cost	Slowest retrieval, lowest cost
S3 Glacier Deep Archive		Faster retrieval, higher cost	Slower retrieval, lowest cost

Each method has a different per-GB retrieval fee and per-request fee. For detailed S3 Glacier pricing by AWS Region, visit the [Amazon S3 pricing page](#).

How long will it take to restore my objects archived in Amazon S3 Glacier?

There are 2 parts that make up the total restore time:

- **Retrieval time:** The time to retrieve the backup file from archive and place it in Standard storage. This is sometimes called the "rehydration" time. The retrieval time is different depending on the restore priority you choose.

Archive Tier	Restore Priority & Retrieval Time		
	High	Standard	Low
S3 Glacier	3-5 minutes	3-5 hours	5-12 hours
S3 Glacier Deep Archive		12 hours	48 hours

- **Restore time:** The time to restore the data from the backup file in Standard storage. This time is no different than the typical restore operation directly from Standard storage - when not using an archival tier.

For more information about Amazon S3 Glacier and S3 Glacier Deep Archive retrieval options, refer to [the Amazon FAQ about these storage classes](#).

Supported Azure archive access tiers with NetApp Backup and Recovery

NetApp Backup and Recovery supports one Azure archival access tier and most regions.



To switch to and from NetApp Backup and Recovery UI versions, refer to [Switch to the previous NetApp Backup and Recovery UI](#).

Supported Azure Blob access tiers for NetApp Backup and Recovery

When backup files are initially created they're stored in the *Cool* access tier. This tier is optimized for storing data that's infrequently accessed; but when needed, can be accessed immediately.

If your source clusters are running ONTAP 9.10.1 or greater, you can choose to tier backups from *Cool* to *Azure Archive* storage after a certain number of days (typically more than 30 days) for further cost optimization. Data in this tier can't be accessed immediately when needed, and will require a higher retrieval cost, so you need to consider how often you may need to restore data from these archived backup files. See the section on this page about restoring data from archival storage.

Note that when you configure NetApp Backup and Recovery with this type of lifecycle rule, you must not configure any lifecycle rules when setting up the container in your Azure account.

[Learn about Azure Blob access tiers.](#)

Restore data from archival storage

While storing older backup files in archival storage is much less expensive than Cool storage, accessing data from a backup file in Azure Archive for restore operations will take a longer amount of time and will cost more money.

How much does it cost to restore data from Azure Archive?

There are two restore priorities you can choose when retrieving data from Azure Archive:

- **High:** Fastest retrieval, higher cost
- **Standard:** Slower retrieval, lower cost

Each method has a different per-GB retrieval fee and per-request fee. For detailed Azure Archive pricing by Azure Region, visit the [Azure pricing page](#).



The High priority is not supported when restoring data from Azure to StorageGRID systems.

How long will it take to restore my data archived in Azure Archive?

There are 2 parts that make up the restore time:

- **Retrieval time:** The time to retrieve the archived backup file from Azure Archive and place it in Cool storage. This is sometimes called the "rehydration" time. The retrieval time is different depending on the restore priority you choose:
 - **High:** < 1 hour
 - **Standard:** < 15 hours
- **Restore time:** The time to restore the data from the backup file in Cool storage. This time is no different than the typical restore operation directly from Cool storage - when not using an archival tier.

For more information about Azure Archive retrieval options, refer to [this Azure FAQ](#).

Supported Google archive storage tiers with NetApp Backup and Recovery

NetApp Backup and Recovery supports one Google archival storage class and most regions.



To switch to and from NetApp Backup and Recovery UI versions, refer to [Switch to the previous NetApp Backup and Recovery UI](#).

Supported Google archival storage classes for NetApp Backup and Recovery

When backup files are initially created they're stored in *Standard* storage. This tier is optimized for storing data that's infrequently accessed; but that also allows you to access it immediately.

If your on-prem cluster is using ONTAP 9.12.1 or greater, you can choose to tier older backups to *Archive* storage in the NetApp Backup and Recovery UI after a certain number of days (typically more than 30 days) for further cost optimization. Data in this tier will require a higher retrieval cost, so you need to consider how often you may need to restore data from these archived backup files. See the section on this page about restoring data from archival storage.

Note that when you configure NetApp Backup and Recovery with this type of lifecycle rule, you must not configure any lifecycle rules when setting up the bucket in your Google account.

[Learn about Google storage classes.](#)

Restore data from archival storage

While storing older backup files in Archive storage is much less expensive than Standard storage, accessing data from a backup file in Archive storage for restore operations will take a slightly longer amount of time and will cost more money.

How much does it cost to restore data from Google Archive?

For detailed Google Cloud Storage pricing by region, visit the [Google Cloud Storage pricing page](#).

How long will it take to restore my objects archived in Google Archive?

There are 2 parts that make up the total restore time:

- **Retrieval time:** The time to retrieve the backup file from Archive and place it in Standard storage. This is sometimes called the "rehydration" time. Unlike the "coldest" storage solutions provided by other cloud providers, your data is accessible within milliseconds.
- **Restore time:** The time to restore the data from the backup file in Standard storage. This time is no different than the typical restore operation directly from Standard storage - when not using an archival tier.

Copyright information

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.