



NetApp Backup and Recovery documentation

NetApp Backup and Recovery

NetApp

February 12, 2026

This PDF was generated from <https://docs.netapp.com/us-en/data-services-backup-recovery/index.html> on February 12, 2026. Always check docs.netapp.com for the latest.

Table of Contents

NetApp Backup and Recovery documentation	1
Release notes	2
What's new in NetApp Backup and Recovery	2
09 February 2026	2
19 January 2026	3
08 December 2025	4
06 October 2025	4
25 August 2025	6
12 August 2025	7
28 July 2025	9
14 July 2025	10
09 June 2025	11
13 May 2025	12
16 April 2025	13
17 March 2025	14
21 February 2025	15
13 February 2025	16
22 November 2024	16
27 September 2024	17
Known limitations with NetApp Backup and Recovery for ONTAP volumes	17
Replication limitations for ONTAP volumes	18
Backup-to-object limitations for ONTAP volumes	18
Restore limitations for ONTAP volumes	20
Known limitations with NetApp Backup and Recovery for Microsoft SQL Server workloads	20
Clone lifecycle support	20
Standard deployment mode only	21
Windows cluster name restriction	21
SnapCenter migration issues	21
Limited support for virtualization management software	22
Known limitations with NetApp Backup and Recovery for VMware workloads	22
Known limitations with NetApp Backup and Recovery for Hyper-V workloads	23
Unsupported actions	23
Known limitations with NetApp Backup and Recovery for KVM workloads	23
Unsupported actions	23
Unsupported configurations	24
Troubleshooting notes	24
Known limitations with NetApp Backup and Recovery for Oracle Database workloads	24
Get started	25
Learn about NetApp Backup and Recovery	25
What you can do with NetApp Backup and Recovery	25
Benefits of using NetApp Backup and Recovery	26
Cost	27
Licensing	28

Supported workloads, systems, and backup targets	29
How NetApp Backup and Recovery works	29
Terms that might help you with NetApp Backup and Recovery	30
NetApp Backup and Recovery prerequisites	31
Prerequisite for ONTAP 9.8 and later	31
Prerequisites for backups to object storage	31
Requirements for protecting Microsoft SQL Server workloads	31
Requirements for protecting VMware workloads	32
Requirements for protecting KVM workloads	33
Requirements for protecting Oracle Database workloads	33
Requirements for protecting Kubernetes applications	34
Requirements for protecting Hyper-V workloads	34
In NetApp Console	35
Set up licensing for NetApp Backup and Recovery	36
30-day free trial	36
Use a NetApp Backup and Recovery PAYGO subscription	37
Use an annual contract	37
Use a NetApp Backup and Recovery BYOL license	38
Set up security certificates for StorageGRID and ONTAP in NetApp Backup and Recovery	39
Create a security certificate for StorageGRID	39
Create a security certificate for ONTAP	43
Create a certificate for both ONTAP and StorageGRID	46
Set up backup destinations before you use NetApp Backup and Recovery	46
Prepare the backup destination	47
Set up S3 permissions	47
Log in to NetApp Backup and Recovery	50
Discover offsite backup targets in NetApp Backup and Recovery	50
Discover a backup target	51
Add a bucket for a backup target	52
Change credentials for a backup target	53
Switch to different NetApp Backup and Recovery workloads	54
Switch to a different workload	54
Configure NetApp Backup and Recovery settings	54
Add credentials for host resources	54
Maintain VMware vCenter settings	56
Import and manage SnapCenter host resources	57
Add a KVM management platform	58
Configure log directories in snapshots for Windows hosts	59
Create an execution hook template	59
Set up role-based access control in NetApp Backup and Recovery	60
Related information	61
Use NetApp Backup and Recovery	62
View protection health on the NetApp Backup and Recovery Dashboard	62
View the Protection summary	62
View the Job summary	62

View the Restore summary	63
Create and manage policies to govern backups in NetApp Backup and Recovery	63
View policies	63
Create a policy	64
Edit a policy	69
Delete a policy	70
Protect ONTAP volume workloads	70
Protect your ONTAP volume data using NetApp Backup and Recovery	70
Plan your protection journey with NetApp Backup and Recovery	79
Manage backup policies for ONTAP volumes with NetApp Backup and Recovery	85
Backup-to-object policy options in NetApp Backup and Recovery	89
Manage backup-to-object storage options in NetApp Backup and Recovery Advanced Settings	97
Back up Cloud Volumes ONTAP data to Amazon S3 with NetApp Backup and Recovery	100
Back up Cloud Volumes ONTAP data to Azure Blob storage with NetApp Backup and Recovery	109
Back up Cloud Volumes ONTAP data to Google Cloud Storage with NetApp Backup and Recovery	118
Back up on-premises ONTAP data to Amazon S3 with NetApp Backup and Recovery	128
Back up on-premises ONTAP data to Azure Blob storage with NetApp Backup and Recovery	141
Back up on-premises ONTAP data to Google Cloud Storage with NetApp Backup and Recovery	151
Back up on-premises ONTAP data to ONTAP S3 with NetApp Backup and Recovery	161
Back up on-premises ONTAP data to StorageGRID with NetApp Backup and Recovery	171
Migrate volumes using SnapMirror to Cloud Resync in NetApp Backup and Recovery	180
Restore NetApp Backup and Recovery configuration data in a dark site	184
Manage backups for your ONTAP systems with NetApp Backup and Recovery	189
Restore from ONTAP backups	198
Protect Microsoft SQL Server workloads	213
Protect Microsoft SQL workloads using NetApp Backup and Recovery overview	213
Prerequisites for importing from the Plug-in service into NetApp Backup and Recovery	215
Discover Microsoft SQL Server workloads and optionally import from SnapCenter in NetApp Backup and Recovery	218
Back up Microsoft SQL Server workloads with NetApp Backup and Recovery	222
Restore Microsoft SQL Server workloads with NetApp Backup and Recovery	225
Clone Microsoft SQL Server workloads using NetApp Backup and Recovery	230
Manage Microsoft SQL Server inventory with NetApp Backup and Recovery	233
Manage Microsoft SQL Server snapshots with NetApp Backup and Recovery	239
Create reports for Microsoft SQL Server workloads in NetApp Backup and Recovery	240
Protect VMware workloads	240
Protect VMware workloads with NetApp Backup and Recovery overview	240
Discover VMware workloads with NetApp Backup and Recovery	241
Create and manage protection groups for VMware workloads with NetApp Backup and Recovery	244
Back up VMware workloads with NetApp Backup and Recovery	246
Restore VMware workloads	247
Protect KVM workloads (Preview)	257
Protect KVM workloads overview	257
Discover KVM workloads in NetApp Backup and Recovery	258
Create and manage protection groups for KVM workloads with NetApp Backup and Recovery	259

Back up KVM workloads with NetApp Backup and Recovery	260
Restore KVM virtual machines with NetApp Backup and Recovery	261
Protect Hyper-V workloads	263
Protect Hyper-V workloads overview	263
Discover Hyper-V workloads in NetApp Backup and Recovery	264
Create and manage protection groups for Hyper-V workloads with NetApp Backup and Recovery	265
Back up Hyper-V workloads with NetApp Backup and Recovery	266
Restore Hyper-V workloads with NetApp Backup and Recovery	267
Protect Oracle Database workloads (Preview)	269
Protect Oracle Database workloads overview	269
Discover Oracle Database workloads in NetApp Backup and Recovery	270
Create and manage protection groups for Oracle Database workloads with NetApp Backup and Recovery	271
Back up Oracle Database workloads using NetApp Backup and Recovery	272
Restore Oracle databases with NetApp Backup and Recovery	273
Mount and unmount Oracle database recovery points with NetApp Backup and Recovery	276
Protect Kubernetes workloads (Preview)	277
Manage Kubernetes workloads overview	277
Discover Kubernetes workloads in NetApp Backup and Recovery	278
Add and protect Kubernetes applications	280
Restore Kubernetes applications	289
Manage Kubernetes clusters	298
Manage Kubernetes applications	299
Manage NetApp Backup and Recovery execution hook templates for Kubernetes workloads	300
Monitor jobs in NetApp Backup and Recovery	303
View job status on the Job Monitor	303
Review retention (backup lifecycle) jobs	305
Review backup and restore alerts in the NetApp Console Notification Center	305
Review operation activity in Console Timeline	307
Restart NetApp Backup and Recovery	307
Automate with NetApp Backup and Recovery REST APIs	309
API reference	309
Getting started	309
Example using the APIs	311
Reference	314
Policies in SnapCenter compared to those in NetApp Backup and Recovery	314
Schedule tiers	314
Multiple policies in SnapCenter with the same schedule tier	314
Imported SnapCenter daily schedules	314
Imported SnapCenter hourly schedules	315
Log retention from SnapCenter policies	315
Log backup retention	315
Retention count from SnapCenter policies	315
SnapMirror labels from SnapCenter policies	316
NetApp Backup and Recovery Identity and Access Management (IAM) roles	316

Restore NetApp Backup and Recovery configuration data in a dark site	316
Restore NetApp Backup and Recovery data to a new Console agent	316
Supported AWS archive storage tiers with NetApp Backup and Recovery	321
Supported S3 archival storage classes for NetApp Backup and Recovery	322
Restore data from archival storage	322
Supported Azure archive access tiers with NetApp Backup and Recovery	323
Supported Azure Blob access tiers for NetApp Backup and Recovery	323
Restore data from archival storage	323
Supported Google archive storage tiers with NetApp Backup and Recovery	324
Supported Google archival storage classes for NetApp Backup and Recovery	324
Restore data from archival storage	325
Legal notices	326
Copyright	326
Trademarks	326
Patents	326
Privacy policy	326
Open source	326

NetApp Backup and Recovery documentation

Release notes

What's new in NetApp Backup and Recovery

Learn what's new in NetApp Backup and Recovery.

09 February 2026

This NetApp Backup and Recovery release includes the following updates.

Microsoft Hyper-V workloads supported in General Availability (GA)

Microsoft Hyper-V workload support is now generally available (GA) in NetApp Backup and Recovery.

VMware workloads supported in General Availability (GA)

VMware workload support is now generally available (GA) in NetApp Backup and Recovery.

Kubernetes workloads enhancements

This release of Kubernetes workloads introduces the following enhanced capabilities:

- **CR workflow support:** You can now perform common protection tasks using CRs as well as the Backup and Recovery web UI.
- **Cluster migration:** You can now add existing Kubernetes clusters protected with Trident Protect to Backup and Recovery.
- **Alerting framework support:** You can now receive email and UI alerts for certain Kubernetes workload events.
- **Restore tab integration:** You can now access the Kubernetes workload restore actions from the Restore menu.
- **Support for the 3-2-1 fanout backup architecture:** You can now use a 3-2-1 fanout architecture in your protection policy when protecting Kubernetes workloads.

For details about protecting Kubernetes workloads, refer to [Protect Kubernetes workloads overview](#).

Oracle Database workloads enhancements

This release of Oracle Database workloads introduces the following enhanced capabilities:

- **Non-root user support:** Non-root users can now perform backup, restore, and clone operations, improving security and compliance.
- **Clone support:** Clone capabilities are now supported across primary and secondary NAS, SAN, and ASM environments using ASM library v2, enabling coordinated protection workflows.
- **Clone split support:** You can now split writable snapshots (clones) from their parent volumes, freeing up storage and enabling independent operations.
- **Backup and restore for object store:** Native backup and restore capabilities are now supported for object-based S3-compatible storage targets.
- **Clone Lifecycle Management (CLM):** Clone refresh operations are supported on primary storage.

- **Clone to alternate host:** You can now clone databases to a different host (for testing or analytics) from both primary and secondary storage.
- **ONTAP consistency group support:** ONTAP consistency groups are now supported, ensuring application-consistent snapshots across multiple volumes.
- Backup and Recovery now supports the following protection policy architectures for Oracle Database workloads:
 - 3-2-1 fanout
 - Disk to disk
 - Disk to object storage
 - Cascading
 - Local snapshot

For details about protecting Oracle Database workloads, refer to [Protect Oracle Database workloads overview](#).

19 January 2026

This NetApp Backup and Recovery release includes the following updates.

ONTAP volume workloads enhancements

This release of ONTAP volume workloads introduces the following enhanced capability:

Support for multiple buckets: (Private preview) Beginning with ONTAP 9.17.1 and newer, you can now protect the volumes within a system with up to 6 buckets per system across different cloud providers.

[Learn more about backing up ONTAP volume data using NetApp Backup and Recovery](#).

VMware workloads enhancements

This release of VMware workloads introduces the following enhanced capabilities:

- VMware workload support is now generally available (GA) in NetApp Backup and Recovery.
- You can now restore guest OS files and folders.

[Learn more about restoring guest files and folders](#).

Hyper-V workloads preview enhancements

This release of Hyper-V workloads introduces the following enhanced capabilities:

- You can now restore Hyper-V VM backups and snapshots to an alternate location. Use this feature to manage VM versions on different Hyper-V hosts.
- NetApp Backup and Recovery now supports Hyper-V virtual machines provisioned by System Center Virtual Machine Manager (SCVMM) and hosted on a CIFS share.
- You can now edit protection groups.



In this release only, you cannot upgrade the NetApp plugins for Hyper-V or Windows using the **Upgrade** option in the Actions menu. Instead, remove each Hyper-V host and add it again to update the plugins.

[Learn more about restoring Hyper-V VMs with NetApp Backup and Recovery.](#)

KVM workloads preview enhancements

The KVM workloads preview now protects KVM hosts and virtual machines that are managed by Apache CloudStack.

For details about protecting KVM workloads, refer to [Protect KVM workloads overview](#).

08 December 2025

This NetApp Backup and Recovery release includes the following updates.

VMware workloads preview enhancements

The preview release of VMware workloads introduces the following enhanced capabilities:

- You can now restore backups and snapshots to an alternate location. This is useful if you want to manage versions of a VM on different VMware vCenter deployments, VMware ESXi hosts, or VMware datastores.

[Learn more about restoring VMware VMs with NetApp Backup and Recovery.](#)

- You can now restore specific VMware virtual disks (VMDK images) from either a primary or secondary location, enabling more granular control over restoring VM data.

[Learn more about restoring VMware virtual disks with NetApp Backup and Recovery.](#)

06 October 2025

This NetApp Backup and Recovery release includes the following updates.

BlueXP backup and recovery is now NetApp Backup and Recovery

BlueXP backup and recovery has been renamed to NetApp Backup and Recovery.

BlueXP is now NetApp Console

The NetApp Console, built on the enhanced and restructured BlueXP foundation, provides centralized management of NetApp storage and NetApp Data Services across on-premises and cloud environments at enterprise grade—delivering real-time insights, faster workflows, and simplified administration that is highly secure and compliant.

For details on what's changed, see the [NetApp Console release notes](#).

Hyper-V workload support as a private preview

This release of NetApp Backup and Recovery introduces support for discovering and managing Hyper-V workloads:

- Backup and restore VMs on standalone instances as well as failover cluster instances (FCI)
- Protect VMs stored on SMB3 shares
- Bulk protection at virtual machine level
- VM and crash consistent backups

- Restore VMs from primary, secondary, and object storage
- Search and restore VM backups

For details about protecting Hyper-V workloads, refer to [Protect Hyper-V workloads overview](#).

KVM workload support as a private preview

This release of NetApp Backup and Recovery introduces support for discovering and managing KVM workloads:

- Back up and restore qcow2 VM images stored on NFS shares
- Back up storage pools
- Bulk VM and storage pool protection using protection groups
- VM consistent and crash consistent VM backups
- Search and restore VM backups from primary, secondary, and object storage
- Guided process to back up and restore KVM-based VMs and VM data

For details about protecting KVM workloads, refer to [Protect KVM workloads overview](#).

Kubernetes preview enhancements

The preview release of Kubernetes workloads introduces the following enhanced capabilities:

- 3-2-1 Fan out backup architecture support
- Support for ONTAP S3 as a backup target
- New Kubernetes dashboard for easier management
- Enhanced role based access control (RBAC) configuration includes support for the following roles:
 - Backup and Recovery super admin
 - Backup and Recovery backup admin
 - Backup and Recovery restore admin
 - Backup and Recovery viewer
- Support for the SUSE Rancher Kubernetes distribution
- Multi-bucket support: You can now protect the volumes within a system with multiple buckets per system across different cloud providers

For details about protecting Kubernetes workloads, refer to [Protect Kubernetes workloads overview](#).

VMware preview enhancements

The preview release of VMware workloads introduces the following enhanced capabilities:

- Support for restoring from object storage
- NetApp Console Dashboard now displays VMware workload status information
- Role-based access control (RBAC) support
- Email alert and notification support for job events
- Support for backing up and restoring to NVMe-based storage

- Edit protection groups
- Edit protection policies

For details about protecting VMware workloads, refer to [Protect VMware workloads overview](#).

Oracle Database workload support as a private preview

This release of NetApp Backup and Recovery introduces support for discovering and managing Oracle Database workloads:

- Discover standalone Oracle databases
- Create protection policies for data only or data and log backups
- Protect Oracle databases with a 3-2-1 backup scheme
- Configure backup retention
- Mount and unmount ARCHIVELOG backups
- Virtualized databases
- Search and restore database backups
- Oracle dashboard support

For details about protecting Oracle Database workloads, refer to [Protect Oracle Database workloads overview](#).

ONTAP volume workload enhancements

This release of ONTAP volume workloads introduces the following enhanced capability:

Beginning with ONTAP 9.17.1 and newer, DataLock is now supported with Google Cloud Platform. This complements existing DataLock support with Amazon AWS, Microsoft Azure, and NetApp StorageGRID.

25 August 2025

This NetApp Backup and Recovery release includes the following updates.

Support for protecting VMware workloads in Preview

This release adds preview support for protecting VMware workloads. Back up VMware VMs and datastores from on-premises ONTAP systems to Amazon Web Services and StorageGRID.



Documentation about protecting VMware workloads is provided as a technology preview. With this preview offering, NetApp reserves the right to modify offering details, contents, and timeline before General Availability.

[Learn more about protecting VMware workloads with NetApp Backup and Recovery](#).

High performance indexing for AWS, Azure, and GCP is generally available

In February 2025, we announced the preview of high performance indexing (Indexed Catalog v2) for AWS, Azure, and GCP. This feature is now generally available (GA). In June 2025, we provided it to all *new* customers by default. With this release, the support is available to *all* customers. High performance indexing improves the performance of backup and restore operations for workloads that are protected to object storage.

Enabled by default:

- If you are a new customer, high performance indexing is enabled by default.
- If you are an existing customer, you can enable reindexing by going to the Restore section of the UI.

12 August 2025

This NetApp Backup and Recovery release includes the following updates.

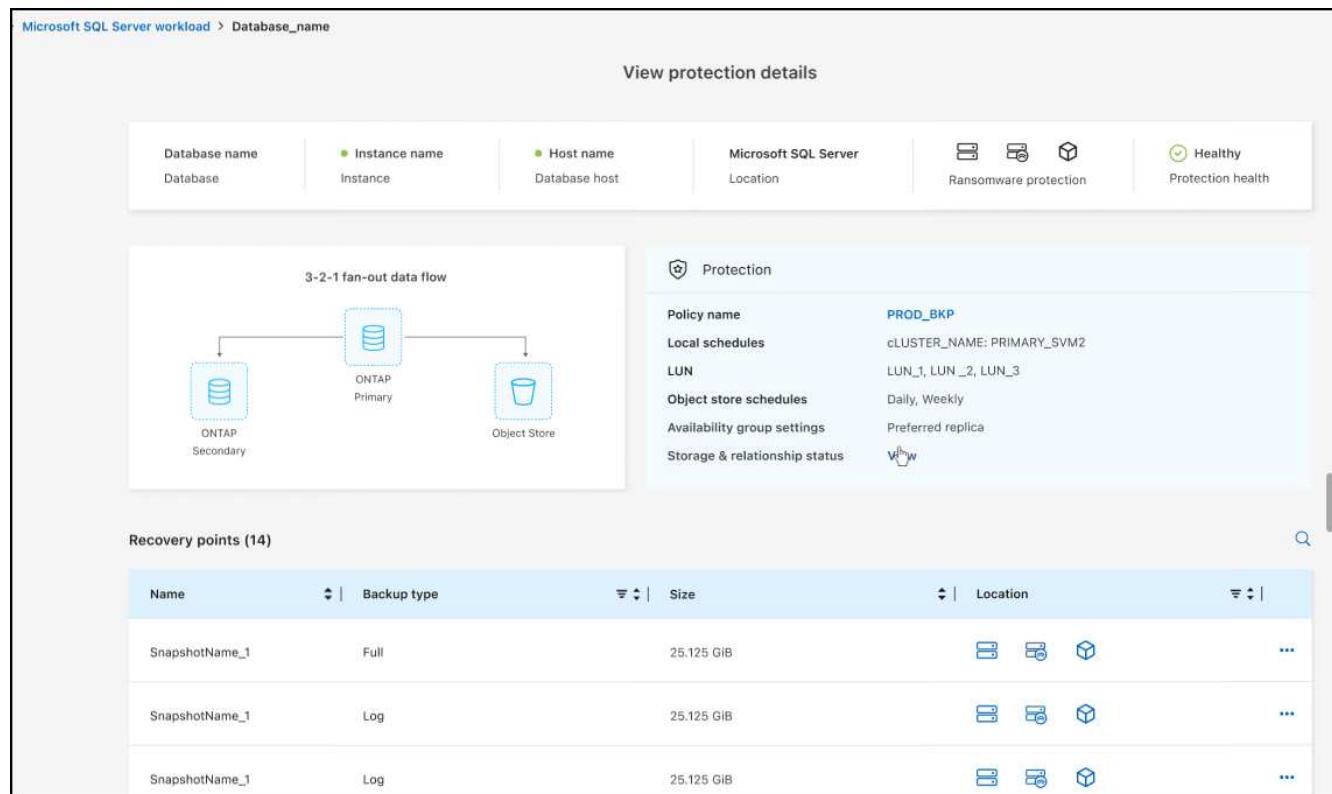
Microsoft SQL Server workload supported in General Availability (GA)

Microsoft SQL Server workload support is now generally available (GA) in NetApp Backup and Recovery. Organizations using an MSSQL environment on ONTAP, Cloud Volumes ONTAP, and Amazon FSx for NetApp ONTAP storage can now take advantage of this new backup and recovery service to protect their data.

This release includes the following enhancements to the Microsoft SQL Server workload support from the previous preview version:

- **SnapMirror active sync:** This version now supports SnapMirror active sync (also referred to as SnapMirror Business Continuity [SM-BC]), which enables business services to continue operating even through a complete site failure, supporting applications to fail over transparently using a secondary copy. NetApp Backup and Recovery now supports protection of Microsoft SQL Server databases in a SnapMirror active sync and Metrocluster configuration. The information appears in the **Storage and relationship status** section of the Protection details page. The relationship information is displayed in the updated **Secondary settings** section of the Policy page.

Refer to [Use policies to protect your workloads](#).



Microsoft SQL Server workload > Database_name

View protection details

Database name Database Instance name Instance Host name Database host Microsoft SQL Server Location Ransomware protection Protection health

3-2-1 fan-out data flow

ONTAP Primary

ONTAP Secondary

Object Store

Protection

Policy name PROD_BKP

Local schedules cLUSTER_NAME: PRIMARY_SVM2

LUN LUN_1, LUN_2, LUN_3

Object store schedules Daily, Weekly

Availability group settings Preferred replica

Storage & relationship status

Recovery points (14)

Name	Backup type	Size	Location	Actions
SnapshotName_1	Full	25.125 GiB		...   
SnapshotName_1	Log	25.125 GiB		...   
SnapshotName_1	Log	25.125 GiB		...   

- **Multi-bucket support:** You can now protect the volumes within a working environment with up to 6 buckets per working environment across different cloud providers.

- **Licensing and free trial updates** for SQL Server workloads: You can now use the existing NetApp Backup and Recovery licensing model to protect SQL Server workloads. There is no separate licensing requirement for SQL Server workloads.

For details, refer to [Set up licensing for NetApp Backup and Recovery](#).

- **Custom snapshot name**: You can now use your own snapshot name in a policy that governs the backups for Microsoft SQL Server workloads. Enter this information in the **Advanced settings** section of the Policy page.

The screenshot shows the 'Create policy' interface. At the top, it says 'Create policy' and 'Create a backup and recovery policy to protect your data.' Below this, there are several sections with expandable dropdowns:

- Details:** Workload type Microsoft SQL Server | Name Test123 | Name Test123
- Backup architecture:** Data flow 3-2-1 cascade
- Local snapshot settings:** Schedule Daily, Weekly, Monthly, Yearly | Log backup Enabled
- Secondary settings:** Backup Hourly, Daily, Weekly, Monthly, Yearly | Backup targets ONTAP targets | SVM | AGGR
- Object store settings:** Backup Weekly, Monthly | Backup target Registered object stores | Retention ...

At the bottom, the 'Advanced settings' section is expanded. It includes:

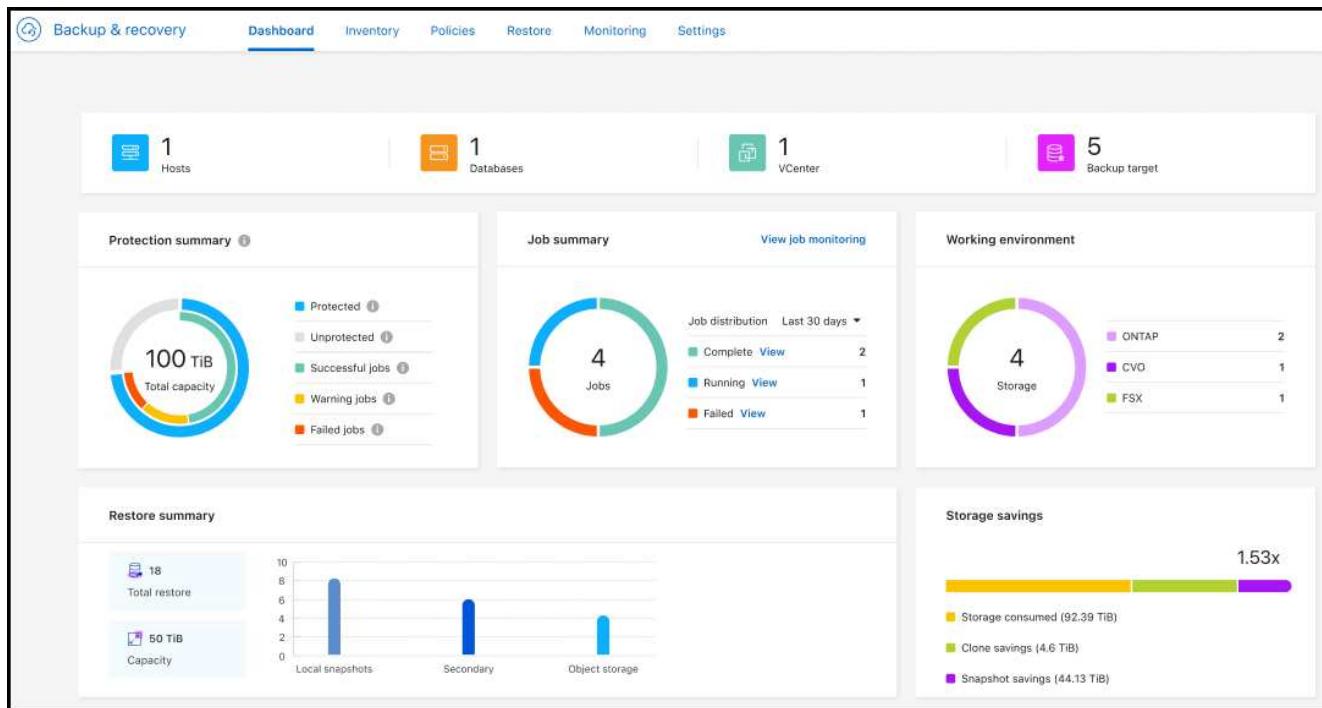
- SnapMirror volume and snapshot format:**
 - Use custom name format for snapshot copy
 - Snapshot name format: Protection group X \$Policy X +5 (with a 'Custom text' input field containing 'Test_text')
 - Provide SnapMirror volume format (ONTAP Secondary)
 - Prefix: Vol_
 - Suffix: <sourceVolumeName>_Dest

Refer to [Use policies to protect your workloads](#).

- **Secondary volume prefix and suffix**: You can enter a custom prefix and suffix in the **Advanced settings** section of the Policy page.
- **Identity and access**: You can now control users' access to features.

Refer to [Log in to NetApp Backup and Recovery](#) and [NetApp Backup and Recovery access to features](#).

- **Restore from object storage to an alternate host**: You can now restore from object storage to an alternate host even if the primary storage is down.
- **Log backup data**: The database protection details page now shows log backups. You can see the Backup type column that shows whether the backup is a full backup or a log backup.
- **Enhanced Dashboard**: The Dashboard now shows Storage and Clone savings.



ONTAP volume workload enhancements

- Multi-folder restore for ONTAP volumes:** Until now, you could restore either one folder or multiple files at a time from the Browse and restore feature. NetApp Backup and Recovery now provides the ability to select multiple folders at a time using the Browse and restore feature.
- View and manage backups of deleted volumes:** The NetApp Backup and Recovery Dashboard now gives an option to show and manage volumes that are deleted from ONTAP. With this, you can view and delete backups from volumes that no longer exist in ONTAP.
- Force delete backups:** In some extreme cases, you might want NetApp Backup and Recovery not to have access to backups any longer. This might happen for example, if the service no longer has access to the backup bucket or backups are DataLock protected but you don't want them anymore. Previously, you could not delete these yourself and needed to call NetApp Support. With this release, you can use the option to force delete backups (at volume and work environment levels).

! Use this option carefully and only in extreme cleanup needs. NetApp Backup and Recovery will not have access to these backups any longer even if they are not deleted in the object storage. You will need to go to your cloud provider and manually delete the backups.

Refer to [Protect ONTAP workloads](#).

28 July 2025

This NetApp Backup and Recovery release includes the following updates.

Kubernetes workload support as a Preview

This release of NetApp Backup and Recovery introduces support for discovering and managing Kubernetes workloads:

- Discover Red Hat OpenShift and open-source Kubernetes clusters, backed by NetApp ONTAP, without sharing kubeconfig files.

- Discover, manage, and protect applications across multiple Kubernetes clusters using a unified control plane.
- Offload data movement operations for backup and recovery of Kubernetes applications to NetApp ONTAP.
- Orchestrate local and object-storage-based application backups.
- Back up and restore entire applications and individual resources to any Kubernetes clusters.
- Work with containers and virtual machines running on Kubernetes.
- Create application-consistent backups using execution hooks and templates.

For details about protecting Kubernetes workloads, refer to [Protect Kubernetes workloads overview](#).

14 July 2025

This NetApp Backup and Recovery release includes the following updates.

Enhanced ONTAP volume Dashboard

In April 2025, we launched a preview of an enhanced ONTAP volume Dashboard that is much faster and more efficient.

This dashboard was designed to help enterprise customers with a high number of workloads. Even for customers with 20,000 volumes, the new dashboard loads in <10 seconds.

After a successful preview and great feedback from preview customers, we are now making it the default experience for all our customers. Be ready for a blazingly fast dashboard.

For details, see [View protection health in the Dashboard](#).

Microsoft SQL Server workload support as a Public Technology Preview

This release of NetApp Backup and Recovery provides an updated user interface that enables you to manage Microsoft SQL Server workloads using a 3-2-1 protection strategy, familiar in NetApp Backup and Recovery. With this new version, you can back up these workloads to primary storage, replicate them to secondary storage, and back them up to cloud object storage.

You can sign up for the preview by completing this [Preview Signup Form](#).



This documentation about protecting Microsoft SQL Server workloads is provided as a technology preview. With this preview offering, NetApp reserves the right to modify offering details, contents, and timeline before general availability.

This version of NetApp Backup and Recovery includes the following updates:

- **3-2-1 backup capability:** This version integrates SnapCenter capabilities, enabling you to manage and protect your SnapCenter resources with a 3-2-1 data protection strategy from the NetApp Backup and Recovery user interface.
- **Import from SnapCenter:** You can import SnapCenter backup data and policies into NetApp Backup and Recovery.
- **A redesigned user interface** provides a more intuitive experience for managing your backup and recovery tasks.
- **Backup targets:** You can add buckets in Amazon Web Services (AWS), Microsoft Azure Blob Storage,

StorageGRID, and ONTAP S3 environments to use as backup targets for your Microsoft SQL Server workloads.

- **Workload support:** This version enables you to back up, restore, verify, and clone Microsoft SQL Server databases and availability groups. (Support for other workloads will be added in future releases.)
- **Flexible restore options:** This version enables you to restore databases to both original and alternate locations in case of corruption or accidental data loss.
- **Instant production copies:** Generate space-efficient production copies for development, testing, or analytics in minutes instead of hours or days.
- This version includes the ability to create detailed reports.

For details about protecting Microsoft SQL Server workloads, see [Protect Microsoft SQL Server workloads overview](#).

09 June 2025

This NetApp Backup and Recovery release includes the following updates.

Indexed catalog support updates

In February 2025, we introduced the updated indexing feature (Indexed Catalog v2) that you use during the Search & Restore method of restoring data. The previous release significantly improved data indexing performance in on-premises environments. With this release, the indexing catalog is now available with Amazon Web Services, Microsoft Azure, and Google Cloud Platform (GCP) environments.

If you are a new customer, the Indexed Catalog v2 is enabled by default for all new environments. If you are an existing customer, you can re-index your environment to leverage the Indexed Catalog v2.

How do you enable indexing?

Before you can use the Search & Restore method of restoring data, you need to enable "Indexing" on each source working environment from which you're planning to restore volumes or files. Select the **Enable Indexing** option when you are performing a Search & Restore.

The Indexed Catalog can then track every volume and backup file, making your searches quick and efficient.

For more information, refer to [Enable indexing for Search & Restore](#).

Azure private link endpoints and service endpoints

Typically, NetApp Backup and Recovery establishes a private endpoint with the cloud provider to handle protection tasks. This release introduces an optional setting that lets you enable or disable NetApp Backup and Recovery from automatically creating a private endpoint. This might be useful to you if you want more control over the private endpoint creation process.

You can enable or disable this option when you enable protection or start the restore process.

If you disable this setting, you must manually create the private endpoint for NetApp Backup and Recovery to function properly. Without proper connectivity, you might not be able to perform backup and recovery tasks successfully.

Support for SnapMirror to Cloud Resync on ONTAP S3

The previous release introduced support for SnapMirror to Cloud Resync (SM-C Resync). The feature streamlines data protection during volume migration in NetApp environments. This release adds support for SM-

C Resync on ONTAP S3 as well as other S3-compatible providers such as Wasabi and MinIO.

Bring your own bucket for StorageGRID

When you create backup files in object storage for a working environment, by default, NetApp Backup and Recovery creates the container (bucket or storage account) for the backup files in the object storage account that you configured. Previously, you could override this and specify your own container for Amazon S3, Azure Blob Storage, and Google Cloud Storage. With this release, you can now bring your own StorageGRID object storage container.

See [Create your own object storage container](#).

13 May 2025

This NetApp Backup and Recovery release includes the following updates.

SnapMirror to Cloud Resync for volume migrations

The SnapMirror to Cloud Resync feature streamlines data protection and continuity during volume migrations in NetApp environments. When a volume is migrated using SnapMirror Logical Replication (LRSE) from one on-premises NetApp deployment to another, or to a cloud-based solution such as Cloud Volumes ONTAP, SnapMirror to Cloud Resync ensures that existing cloud backups remain intact and operational.

This feature eliminates the need for a time-consuming and resource-intensive re-baseline operation, enabling backup operations to continue post-migration. This feature is valuable in workload migration scenarios, supporting both FlexVols and FlexGroups, and is available starting with ONTAP version 9.16.1.

By maintaining backup continuity across environments, SnapMirror to Cloud Resync enhances operational efficiency and reduces the complexity of hybrid and multi-cloud data management.

For details on how to perform the resync operation, see [Migrate volumes using SnapMirror to Cloud Resync](#).

Support for third-party MinIO object store (Preview)

NetApp Backup and Recovery now extends its support to third-party object stores with a primary focus on MinIO. This new preview feature enables you to leverage any S3-compatible object store for your backup and recovery needs.

With this preview version, we hope to ensure robust integration with third-party object stores before the full functionality is rolled out. You are encouraged to explore this new capability and provide feedback to help enhance the service.



This feature should not be used in production.

Preview mode limitations

While this feature is in preview, there are certain limitations:

- Bring Your Own Bucket (BYOB) is not supported.
- Enabling DataLock in the policy is not supported.
- Enabling Archival mode in the policy is not supported.
- Only on-premises ONTAP environments are supported.

- MetroCluster is not supported.
- Options to enable bucket-level encryption are not supported.

Getting started

To begin using this preview feature, you must enable a flag on the Console agent. You can then enter the connection details of your MinIO third-party object store in the protection workflow by choosing **Third party Compatible** object store in the backup section.

16 April 2025

This NetApp Backup and Recovery release includes the following updates.

UI improvements

This release enhances your experience by simplifying the interface:

- The removal of the Aggregate column from the Volumes tables, along with the Snapshot Policy, Backup Policy, and Replication Policy columns from the Volume table in the V2 Dashboard, results in a more streamlined layout.
- Excluding non-activated working environments from the drop-down list makes the interface less cluttered, the navigation more efficient, and loading faster.
- While sorting on the Tags column is disabled, you can still view the tags, ensuring that important information remains easily accessible.
- The removal of labels on protection icons contributes to a cleaner look and decreases loading time.
- During the working environment activation process, a dialog box displays a loading icon to provide feedback until the discovery process is complete, enhancing transparency and confidence in the system's operations.

Enhanced Volume Dashboard (Preview)

The Volume Dashboard now loads in under 10 seconds, providing a much faster and more efficient interface. This preview version is available to select customers, offering them an early look at these improvements.

Support for third-party Wasabi object store (Preview)

NetApp Backup and Recovery now extends its support to third-party object stores with a primary focus on Wasabi. This new preview feature enables you to leverage any S3-compatible object store for your backup and recovery needs.

Getting started with Wasabi

To begin using third-party storage as an object store, you must enable a flag within the Console agent. Then, you can enter the connection details for your third-party object store and integrate it into your backup and recovery workflows.

Steps

1. SSH into your connector.
2. Go into the NetApp Backup and Recovery cbs server container:

```
docker exec -it cloudmanager_cbs sh
```

3. Open the `default.json` file inside the `config` folder via VIM or any other editor:

```
vi default.json
```

4. Modify `allow-s3-compatible: false` to `allow-s3-compatible: true`.
5. Save the changes.
6. Exit from the container.
7. Restart the NetApp Backup and Recovery cbs server container.

Result

After the container is ON again, open the NetApp Backup and Recovery UI. When you initiate a backup or edit a backup strategy, you will see the new provider "S3 Compatible" listed along with other backup providers of AWS, Microsoft Azure, Google Cloud, StorageGRID, and ONTAP S3.

Preview mode limitations

While this feature is in preview, consider the following limitations:

- Bring Your Own Bucket (BYOB) is not supported.
- Enabling DataLock in a policy is not supported.
- Enabling Archival mode in a policy is not supported.
- Only on-premises ONTAP environments are supported.
- MetroCluster is not supported.
- Options to enable bucket-level encryption are not supported.

During this preview, we encourage you to explore this new feature and provide feedback about integration with third-party object stores before the full functionality is rolled out.

17 March 2025

This NetApp Backup and Recovery release includes the following updates.

SMB snapshot browsing

This NetApp Backup and Recovery update resolved an issue that prevented customers from browsing local snapshots in an SMB environment.

AWS GovCloud environment update

This NetApp Backup and Recovery update fixed an issue that prevented the UI from connecting to an AWS GovCloud environment due to TLS certificate errors. The issue was resolved by using the Console agent host name instead of the IP address.

Backup policy retention limits

Previously, the NetApp Backup and Recovery UI limited backups to 999 copies, while the CLI allowed more. Now, you can attach up to 4,000 volumes to a backup policy and include 1,018 volumes not attached to a backup policy. This update includes additional validations that prevent exceeding these limits.

SnapMirror Cloud resync

This update ensures that SnapMirror Cloud resync cannot be started from NetApp Backup and Recovery for unsupported ONTAP versions after a SnapMirror relationship has been deleted.

21 February 2025

This NetApp Backup and Recovery release includes the following updates.

High performance indexing

NetApp Backup and Recovery introduces an updated indexing feature that makes the indexing of data on the source system more efficient. The new indexing feature includes updates to the UI, improved performance of the Search & Restore method of restoring data, upgrades to global search capabilities, and better scalability.

Here's a breakdown of the improvements:

- **Folder consolidation:** The updated version groups folders together using names that include specific identifiers, making the indexing process smoother.
- **Parquet file compaction:** The updated version reduces the number of files used for indexing each volume, simplifying the process and removing the need for an extra database.
- **Scale-out with more sessions:** The new version adds more sessions to handle indexing tasks, speeding up the process.
- **Support for multiple index containers:** The new version uses multiple containers to better manage and distribute indexing tasks.
- **Split index workflow:** The new version divides the indexing process into two parts, enhancing efficiency.
- **Improved concurrency:** The new version makes it possible to delete or move directories at the same time, speeding up the indexing process.

Who benefits from this feature?

The new indexing feature is available to all new customers.

How do you enable indexing?

Before you can use the Search & Restore method of restoring data, you need to enable "Indexing" on each source system from which you're planning to restore volumes or files. This allows the Indexed Catalog to track every volume and every backup file, making your searches quick and efficient.

Enable indexing on the source working environment by selecting the "Enable Indexing" option when you are performing a Search & Restore.

For more information, see the documentation [how to restore ONTAP data using Search & Restore](#).

Supported scale

The new indexing feature supports the following:

- Global search efficiency in less than 3 minutes

- Up to 5 billion files
- Up to 5000 volumes per cluster
- Up to 100K snapshots per volume
- Maximum time for baseline indexing is less than 7 days. The actual time will vary depending on your environment.

Global search performance improvements

This release also includes enhancements to global search performance. You will now see progress indicators and more detailed search results, including the count of files and the time taken for the search. Dedicated containers for search and indexing ensure that global searches are completed in under five minutes.

Note these considerations related to global search:

- The new index is not performed on snapshots labeled as hourly.
- The new indexing feature works only on snapshots on FlexVols, and not for snapshots on FlexGroups.

13 February 2025

This NetApp Backup and Recovery release includes the following updates.

NetApp Backup and Recovery Preview Release

This Preview release of NetApp Backup and Recovery provides an updated user interface that enables you to manage Microsoft SQL Server workloads using a 3-2-1 protection strategy, familiar in NetApp Backup and Recovery. With this new version, you can back up these workloads to primary storage, replicate them to secondary storage, and back them up to cloud object storage.



This documentation is provided as a technology preview. With this preview offering, NetApp reserves the right to modify offering details, contents, and timeline before General Availability.

This version of NetApp Backup and Recovery Preview 2025 includes the following updates.

- A redesigned user interface that provides a more intuitive experience for managing your backup and recovery tasks.
- The Preview version enables you to back up and restore Microsoft SQL Server databases. (Support for other workloads will be added in future releases.)
- This version integrates SnapCenter capabilities, enabling you to manage and protect your SnapCenter resources with a 3-2-1 data protection strategy from the NetApp Backup and Recovery user interface.
- This version enables you to import SnapCenter workloads into NetApp Backup and Recovery.

22 November 2024

This NetApp Backup and Recovery release includes the following updates.

SnapLock Compliance and SnapLock Enterprise protection modes

NetApp Backup and Recovery now can back up both FlexVol and FlexGroup on-premises volumes that are configured using either SnapLock Compliance or SnapLock Enterprise protection modes. Your clusters must be running ONTAP 9.14 or greater for this support. Backing up FlexVol volumes using SnapLock Enterprise

mode has been supported since ONTAP version 9.11.1. Earlier ONTAP releases provide no support for backing up SnapLock protection volumes.

See the complete list of supported volumes in the [Learn about NetApp Backup and Recovery](#).

Indexing for Search & Restore process on Volumes page

Before you can use Search & Restore, you need to enable "Indexing" on each source system from which you'll want to restore volume data. This enables the Indexed Catalog to track the backup files for every volume. The Volumes page now shows the indexing status:

- Indexed: Volumes have been indexed.
- In-progress
- Not Indexed
- Indexing paused
- Error
- Not Enabled

27 September 2024

This NetApp Backup and Recovery release includes the following updates.

Podman support on RHEL 8 or 9 with Browse and Restore

NetApp Backup and Recovery now supports file and folder restores on Red Hat Enterprise Linux (RHEL) versions 8 and 9 using the Podman engine. This applies to the NetApp Backup and Recovery Browse and Restore method.

Console agent version 3.9.40 supports certain versions of Red Hat Enterprise Linux versions 8 and 9 for any manual installation of the Console agent software on a RHEL 8 or 9 host, regardless of the location in addition to the operating systems mentioned in the [host requirements](#). These newer RHEL versions require the Podman engine instead of the Docker engine. Previously, NetApp Backup and Recovery had two limitations when using the Podman engine. These limitations have been removed.

[Learn more about restoring ONTAP data from backup files.](#)

Faster catalog indexing improves Search and Restore

This release includes an improved catalog index that completes the baseline indexing much faster. Faster indexing enables you to use the Search and Restore feature more quickly.

[Learn more about restoring ONTAP data from backup files.](#)

Known limitations with NetApp Backup and Recovery for ONTAP volumes

Platforms, devices, or features that do not work or do not work well with this version are listed here. Read these limitations carefully.

- NetApp Backup and Recovery can back up Cloud Volumes ONTAP to an object store in the AWS China regions (including Beijing and Ningxia); however, you might need to manually modify Identity and access

policies first.

For details about creating a Console agent in AWS, refer to [Installing a Console agent in AWS](#).

For additional details, refer to the blog post [NetApp Backup and Recovery Feature Blog May 2023](#).

- NetApp Backup and Recovery does not support Microsoft Azure China regions.

For details about creating a Console agent in Azure, refer to [Installing a Console agent in Azure](#).

- NetApp Backup and Recovery does not support backups of FlexCache volumes.

Replication limitations for ONTAP volumes

- You can select only one FlexGroup volume at a time for replication. You'll need to activate backups separately for each FlexGroup volume.

There is no limitation for FlexVol volumes - you can select all FlexVol volumes in your system and assign the same backup policies.

- The following functionality is supported in [NetApp Replication](#), but not when using the replication feature of NetApp Backup and Recovery:
 - There is no support for a cascade configuration where replication occurs from volume A to volume B and from volume B to volume C. Support includes replication from volume A to volume B.
 - There is no support for replicating data to and from FSx for ONTAP systems.
 - There is no support for creating a one-time replication of a volume.
- When creating replications from on-premises ONTAP systems, if the ONTAP version on the target Cloud Volumes ONTAP system is 9.8, 9.9, or 9.11, only mirror-vault policies are allowed.
- NetApp Backup & Recovery does not support converting a FlexVol volume with an active cloud backup relationship to a FlexGroup volume while maintaining cloud backup functionality.

Backup-to-object limitations for ONTAP volumes

- When backing up data, NetApp Backup and Recovery will not maintain NetApp Volume Encryption (NVE). This means that encrypted data on the NVE volume will be decrypted while the data is being transferred to the destination and the encryption will not be maintained.

For an explanation about these encryption types, refer to [Configure NetApp Volume Encryption overview](#).

- If long-term retention snapshots are enabled on a SnapMirror destination volume using the schedule in the SnapMirror policy, snapshots are created directly on the destination volume. In this case, you should not back up those volumes using NetApp Backup and Recovery because those snapshots will not be moved to object storage.
- When backing up data, NetApp Backup and Recovery will not maintain NetApp Volume Encryption (NVE). This means that encrypted data on the NVE volume will be decrypted while the data is being transferred to the destination and the encryption will not be maintained.

For an explanation about these encryption types, refer to [Configure NetApp Volume Encryption overview](#).

- If long-term retention snapshots are enabled on a SnapMirror destination volume using the schedule in the

SnapMirror policy, snapshots are created directly on the destination volume. In this case, you should not back up those volumes using NetApp Backup and Recovery because those snapshots will not be moved to object storage.

- When you create or edit a backup policy when no volumes are assigned to the policy, the number of retained backups can be a maximum of 1018. After you assign volumes to the policy, you can edit the policy to create up to 4000 backups.
- When backing up data protection (DP) volumes:
 - Relationships with the SnapMirror labels `app_consistent` and `all_source_snapshot` won't be backed up to cloud.
 - If you create local copies of Snapshots on the SnapMirror destination volume (irrespective of the SnapMirror labels used) these Snapshots will not be moved to the cloud as backups. At this time you'll need to create a Snapshot policy with the desired labels to the source DP volume in order for NetApp Backup and Recovery to back them up.
- FlexGroup volume backups can't be moved to archival storage.
- FlexGroup volume backups can use DataLock and Ransomware protection if the cluster is running ONTAP 9.13.1 or greater.
- SVM-DR volume backup is supported with the following restrictions:
 - Backups are supported from the ONTAP secondary only.
 - The Snapshot policy applied to the volume must be one of the policies recognized by NetApp Backup and Recovery, including daily, weekly, monthly, etc. The default "sm_created" policy (used for **Mirror All Snapshots**) is not recognized and the DP volume will not be shown in the list of volumes that can be backed up.
 - SVM-DR and volume backup and recovery work fully independently when the backup is taken from either the source or destination. The only restriction is that SVM-DR does not replicate the SnapMirror cloud relationship. In the DR scenario when the SVM goes online in the secondary location, you must manually update the SnapMirror cloud relationship.
- MetroCluster support:
 - When you use ONTAP 9.12.1 GA or greater, backup is supported when connected to the primary system. The entire backup configuration is transferred to the secondary system so that backups to the cloud continue automatically after switchover. You don't need to set up backup on the secondary system (in fact, you are restricted from doing so).
 - When you use ONTAP 9.12.0 and earlier, backup is supported only from the ONTAP secondary system.
 - Beginning with ONTAP 9.18.1, FlexGroup volume backups are supported in MetroCluster configurations.
- Ad-hoc volume backup using the **Backup Now** button isn't supported on data protection volumes.
- SM-BC configurations are not supported.
- ONTAP doesn't support fan-out of SnapMirror relationships from a single volume to multiple object stores; therefore, this configuration is not supported by NetApp Backup and Recovery.
- WORM/Compliance mode on an object store is supported on Amazon S3, Azure, and StorageGRID at this time. This is known as the DataLock feature, and it must be managed by using NetApp Backup and Recovery settings, not by using the cloud provider interface.

Restore limitations for ONTAP volumes

These limitations apply to both the Search & Restore and the Browse & Restore methods of restoring files and folders; unless called out specifically.

- Browse & Restore can restore up to 100 individual files at a time.
- Search & Restore can restore 1 file at a time.
- When using ONTAP 9.13.0 or greater, Browse & Restore and Search & Restore can restore a folder along with all files and sub-folders within it.

When using a version of ONTAP greater than 9.11.1 but before 9.13.0, the restore operation can restore only the selected folder and the files in that folder - no sub-folders, or files in sub-folders, are restored.

When using a version of ONTAP before 9.11.1, folder restore is not supported.

- Directory/folder restore is supported for data that resides in archival storage only when the cluster is running ONTAP 9.13.1 and greater.
- Directory/folder restore is supported for data that is protected using DataLock only when the cluster is running ONTAP 9.13.1 and greater.
- Directory/folder restore is not currently supported from replications and/or local snapshots.
- Restoring from FlexGroup volumes to FlexVol volumes, or FlexVol volumes to FlexGroup volumes is not supported.
- The file being restored must be using the same language as the language on the destination volume. You will receive an error message if the languages are not the same.
- The *High* restore priority is not supported when restoring data from Azure archival storage to StorageGRID systems.
- If you back up a DP volume and then decide to break the SnapMirror relationship to that volume, you cannot restore files to that volume unless you also delete the SnapMirror relationship or reverse the SnapMirror direction.
- Quick restore limitations:
 - The destination location must be a Cloud Volumes ONTAP system using ONTAP 9.13.0 or greater.
 - It is not supported with backups located in archived storage.
 - FlexGroup volumes are supported only if the source system from which the cloud backup was created was running ONTAP 9.12.1 or greater.
 - SnapLock volumes are supported only if the source system from which the cloud backup was created was running ONTAP 9.11.0 or greater.

Known limitations with NetApp Backup and Recovery for Microsoft SQL Server workloads

Platforms, devices, or features that do not work or do not work well with this version are listed here. Read these limitations carefully.

Clone lifecycle support

- Cloning from object storage is not supported.

- Bulk clone operations are not supported for on-demand clones.
- Choosing I-groups is not supported.
- Choosing QOS (maximum throughput) options is not supported.

Standard deployment mode only

This NetApp Backup and Recovery version works only in standard deployment mode, not restricted or private modes.

Windows cluster name restriction

The Windows cluster name cannot contain an underscore (_) character.

SnapCenter migration issues

The migration of resources from SnapCenter into NetApp Backup and Recovery has the following limitations.

For details about how SnapCenter policies migrate to NetApp Backup and Recovery policies, see [Policies in SnapCenter compared to those in NetApp Backup and Recovery](#).

Resource group limitations

If all the resources in a resource group are protected and one of those resources is also protected outside of the resource group, the migration from SnapCenter is blocked.

Workaround: Protect the resource either in a resource group or by itself, but not in both.

Resources with multiple policies using the same schedule tier not supported

You cannot have assign multiple policies that use the same schedule tier (for example, hourly, daily, weekly, etc.) to a resource. NetApp Backup and Recovery will not import those resources from SnapCenter.

Workaround: Attach only one policy using the same schedule tier to a resource.

Hourly policies must begin at the start of the hour

If you have a SnapCenter policy that repeats hourly but does not use intervals at the start of the hour, NetApp Backup and Recovery will not import the resource. For example, policies with schedules of 1:30, 2:30, 3:30, etc. are not supported, while policies with schedules of 1:00, 2:00, 3:00, etc. are supported.

Workaround: Use a policy that repeats in 1-hour intervals starting at the top of the hour.

Both daily and monthly policies attached to one resource not supported

If a SnapCenter policy repeats both in day and month intervals, NetApp Backup and Recovery will not import the policy.

For example, you cannot attach a daily policy (with less than or equal to 7 days or greater than 7 days) to a resource and also attach a monthly policy to the same resource.

Workaround: Use a policy that uses a daily or a monthly interval, but not both.

On demand backup policies not migrated

NetApp Backup and Recovery does not import on demand backup policies from SnapCenter.

Log-only backup policies not migrated

NetApp Backup and Recovery does not import log-only backup policies from SnapCenter. If a SnapCenter policy includes log-only backups, NetApp Backup and Recovery will not import the resource.

Workaround: Use a policy in SnapCenter that uses more than just log-only backups.

Host mapping

SnapCenter does not have map storage clusters or SVMs for the resources to hosts, but NetApp Backup and Recovery does. The on-premises ONTAP cluster or SVM will not be mapped to a host in NetApp Backup and Recovery preview versions. Additionally, NetApp Console does not support SVMs.

Workaround: Before importing resources from SnapCenter, create a system in NetApp Backup and Recovery for all the on-premises ONTAP storage systems that are registered in on-premises SnapCenter. Then, import the resources for that cluster from SnapCenter into NetApp Backup and Recovery.

Schedules not in 15-minute intervals

If you have a SnapCenter policy schedule that starts at a certain time and repeats in intervals other than 15-minute intervals, NetApp Backup and Recovery will not import the schedule.

Workaround: Use SnapCenter to adjust the policy so that it repeats in 15-minute intervals.

Limited support for virtualization management software

When you protect KVM workloads, NetApp Backup and Recovery does not support discovery of KVM workloads when virtualization management software such as Apache CloudStack or Red Hat OpenShift Virtualization is in use.

Known limitations with NetApp Backup and Recovery for VMware workloads

Platforms, devices, or features that do not work or do not work well with this version are listed here. Read these limitations carefully.

The following actions are not supported in the preview version of VMware workloads in NetApp Backup and Recovery:

- Mount
- Unmount
- Attach VMDK
- Detach VMDK
- vVol support
- NVMe support
- Email integration

- Edit policy
- Edit protection group
- Role-based access control (RBAC) support

Known limitations with NetApp Backup and Recovery for Hyper-V workloads

Platforms, devices, or features that do not work or do not work well with this version are listed here. Read these limitations carefully.

Unsupported actions

The following actions are not supported in the private preview version of Hyper-V workloads in NetApp Backup and Recovery:

- Create resource groups using VMs from multiple Hyper-V hosts
- Restore VMs to an alternate location
- Spanning disks (across multiple CIFS shares)
- Protect VMs over SAN
- You cannot restore VMs or VM data between systems with different CPU vendors (Intel to AMD or vice versa), regardless of the "Processor compatibility" setting in Hyper-V. This setting only supports compatibility between different generations of the same vendor (e.g., Intel to Intel or AMD to AMD).



In the January 19 2026 release, you cannot upgrade the NetApp plugins for Hyper-V or Windows using the **Upgrade** option in the Actions menu. Instead, remove each Hyper-V host and add it again to update the plugins.

Known limitations with NetApp Backup and Recovery for KVM workloads

Platforms, devices, or features that do not work or do not work well with this version are listed here. Read these limitations carefully.

The following actions and configurations are not supported in the private preview version of KVM workloads in NetApp Backup and Recovery:

Unsupported actions

The following actions are not supported in the private preview release:

- Clone, mount, or unmount VMs
- Restore VMs to an alternate location
- Protect VMs stored on SAN
- Protect applications
- Edit protection groups

- Create protection groups using VMs from multiple KVM hosts
- Create user-defined backups (only backups initiated from the NetApp Console are supported)

Unsupported configurations

The following configurations are not supported:

- Role-based access control (RBAC)
- Disks directly attached to the KVM host
- Disks spanned across multiple NFS mount points or shares
- RAW disk format
- Storage pool types other than NetFS (only NetFS is supported)

Troubleshooting notes

Make note of the following when using the private preview of KVM workloads with NetApp Backup and Recovery:

- To ensure that KVM workload restores complete successfully, make sure that the **Enable VM-consistent snapshot** setting is active in the protection policy you use for KVM backups.
- You cannot back up a storage pool with KVM hosts managed by Apache CloudStack unless you add all managed hosts to NetApp Backup and Recovery. As a workaround, add every CloudStack-managed KVM host to NetApp Backup and Recovery.
- You cannot back up a stopped VM that belongs to a protection group. As a workaround, remove the stopped VM from the protection group before starting the backup.

Known limitations with NetApp Backup and Recovery for Oracle Database workloads

Platforms, devices, or features that do not work or do not work well with this version are listed here. Read these limitations carefully.

The following action is not supported in the private preview version of Oracle Database workloads in NetApp Backup and Recovery:

- Offline backup

Oracle Database is supported only as a standalone deployment using NFS, SAN, or ASM SAN in the private preview version of Oracle Database workloads.

Get started

Learn about NetApp Backup and Recovery

NetApp Backup and Recovery is a data service that provides efficient, secure, and cost-effective data protection for all your ONTAP workloads, including volumes, databases, virtual machines, and Kubernetes workloads.

Support for Backup and Recovery is already built in to all ONTAP systems, so there is no need for additional hardware, software licenses or media gateways. This makes backup operations simple and cost effective. The NetApp Console simplifies the implementation of any backup strategy, including the full spectrum of 3-2-1 backup variants, without needing multiple resource managers or specialized personnel.



Documentation about protecting VMware, KVM, Hyper-V, and Kubernetes workloads is provided as a technology preview. With this preview offering, NetApp reserves the right to modify offering details, contents, and timeline before General Availability.

What you can do with NetApp Backup and Recovery

Use NetApp Backup and Recovery to accomplish the following goals:

- **ONTAP volume workloads:**

- Create local snapshots, replicate to secondary storage, and back up ONTAP volumes from on-premises ONTAP or Cloud Volumes ONTAP systems to object storage in your public or private cloud account.
- Create block-level, incremental forever backups that are stored on another ONTAP cluster and in object storage in the cloud.
- Use NetApp Backup and Recovery along with SnapCenter.
- Refer to [Protect ONTAP volumes](#).

- **Microsoft SQL Server workloads:**

- Back up Microsoft SQL Server instances and databases from on-premises ONTAP, Cloud Volumes ONTAP, or Amazon FSx for NetApp ONTAP.
- Restore Microsoft SQL Server databases.
- Clone Microsoft SQL Server databases.
- Use NetApp Backup and Recovery without SnapCenter.
- Refer to [Protect Microsoft SQL Server workloads](#).

- **VMware workloads (Preview with new UI without SnapCenter Plug-in for VMware vSphere):**

- Protect your VMware VMs and datastores with NetApp Backup and Recovery.
- Back up VMware workloads to Amazon Web Services S3 or StorageGRID (for Preview).
- Restore VMware data from the cloud back to the on-premises vCenter.
- You can restore the VM to the exact same location from where the backup was taken or to an alternate location.
- Use NetApp Backup and Recovery without SnapCenter Plug-in for VMware vSphere.

- Refer to [Protect VMware workloads](#).

- **VMware workloads (With SnapCenter Plug-in for VMware vSphere):**

- Back up VMs and datastores to Amazon Web Services S3, Microsoft Azure Blob, Google Cloud Platform, and StorageGRID and restore VMs back to the on-premises SnapCenter Plug-in for VMware vSphere host.
- Restore VM data from the cloud back to the on-premises vCenter with NetApp Backup and Recovery. You can restore the VM to the exact same location from where the backup was taken or to an alternate location.
- Use NetApp Backup and Recovery along with SnapCenter Plug-in for VMware vSphere.
- Refer to [Protect VMware workloads](#).

- **KVM workloads (Preview):**

- Back up and restore virtual machines
- Back up KVM storage pools
- Use protection groups to manage backup tasks
- Refer to [Protect KVM workloads](#).

- **Hyper-V workloads (Preview):**

- Back up and restore virtual machines
- Use protection groups to manage backup tasks
- Refer to [Protect Hyper-V workloads](#).

- **Oracle Database workloads (Preview):**

- Back up and restore databases and logs
- Use protection groups to manage backup tasks
- Create policies to manage database and log backups
- Protecting a database with a 3-2-1 backup architecture
- Configure backup retention
- Mount and unmount ARCHIVELOG backups
- Refer to [Protect Oracle Database workloads](#).

- **Kubernetes workloads (Preview):**

- Manage and protect your Kubernetes applications and resources all in one place.
- Use protection policies to structure your incremental backups.
- Restore applications and resources to the same or different clusters and namespaces.
- Use NetApp Backup and Recovery without SnapCenter.
- Refer to [Protect Kubernetes workloads](#).

Benefits of using NetApp Backup and Recovery

NetApp Backup and Recovery provides the following benefits:

- **Efficient:** NetApp Backup and Recovery performs block-level, incremental-forever replication, which significantly reduces the amount of data that's replicated and stored. This helps to minimize network traffic and storage costs.

- **Secure:** NetApp Backup and Recovery encrypts data in transit and at rest, and it uses secure communication protocols to protect your data.
- **Cost-effective:** NetApp Backup and Recovery uses the lowest-cost storage tiers available in your cloud account, which helps to reduce costs.
- **Automated:** NetApp Backup and Recovery automatically generates backups based on a predefined schedule, which helps to ensure that your data is protected.
- **Flexible:** NetApp Backup and Recovery enables you to restore data to the same or different system, which provides flexibility in data recovery.

Cost

NetApp doesn't charge you for using the trial version. However, you are responsible for the costs associated with the cloud resources that you use, such as storage and data transfer costs.

There are two types of costs associated with using the backup-to-object feature of NetApp Backup and Recovery with ONTAP systems:

- Resource charges
- Service charges

There is no charge to create snapshots or replicated volumes - other than the disk space required to store the snapshots and replicated volumes.

Resource charges

Resource charges are paid to the cloud provider for object storage capacity and for writing and reading backup files to the cloud.

- For Backup to object storage, you pay your cloud provider for object storage costs.

Because NetApp Backup and Recovery preserves the storage efficiencies of the source volume, you pay the cloud provider object storage costs for the data *after* ONTAP efficiencies (for the smaller amount of data after deduplication and compression have been applied).

- For restoring data using Search & Restore, certain resources are provisioned by your cloud provider, and there is per-TiB cost associated with the amount of data that is scanned by your search requests. (These resources are not needed for Browse & Restore.)
 - In AWS, [Amazon Athena](#) and [AWS Glue](#) resources are deployed in a new S3 bucket.
 - In Azure, an [Azure Synapse workspace](#) and [Azure Data Lake Storage](#) are provisioned in your storage account to store and analyze your data.
 - In Google, a new bucket is deployed, and the [Google Cloud BigQuery services](#) are provisioned on an account/project level.
- If you plan to restore volume data from a backup file that has been moved to archival object storage, then there's an additional per-GiB retrieval fee and per-request fee from the cloud provider.
- If you plan to scan a backup file for ransomware during the process of restoring volume data (if you enabled DataLock and Ransomware Resilience for your cloud backups), then you'll incur extra egress costs from your cloud provider as well.

Service charges

For ONTAP volume workloads, you are only charged for volumes protected to object storage. Charges are

based on the logical used capacity of the source ONTAP volumes before efficiencies are applied, also known as Front-End Terabytes (FETB).

For Kubernetes workloads, you are charged based on the combined size of all persistent volumes.

For all other workloads, you are charged for resources protected to at least one secondary or object storage target. Charges are calculated using the logical size of the source workload. For databases, this means the database size; for VMs, the VM size.

There are three ways to pay for Backup and Recovery:

- The first option is to subscribe from your cloud provider, which enables you to pay per month.
- The second option is to purchase an annual contract.
- The third option is to purchase licenses directly from NetApp. Refer to the [Licensing](#) section for details.

Licensing

NetApp Backup and Recovery offers a free trial, allowing you to use it without a license key for a limited time.

A Backup license is only required for backup and restore operations involving object storage. Creating snapshots and replicated volumes does not require a license.

You can choose from three licensing options:

- **Bring Your Own License (BYOL):**

Purchase a term-based (1, 2, or 3 years) and capacity-based (in 1-TiB increments) license from NetApp. Enter the provided serial number in the NetApp Console to activate. The license covers all source systems in your organization. Renewal is required when the term or capacity limit is reached.

- **Pay As You Go (PAYGO):**

Subscribe through your cloud provider's marketplace and pay per GiB of backed-up data, billed monthly. No upfront payment is required. A 30-day free trial is available when you first sign up. For more information, refer to [use a NetApp Backup and Recovery PAYGO subscription](#).

- **Annual Contract:**

Available through AWS and Azure marketplaces for 1, 2, or 3 years. Two annual contracts are available:

- **Cloud Backup:** Backs up Cloud Volumes ONTAP and on-premises ONTAP data.

- **CVO Professional:** Bundles Cloud Volumes ONTAP and NetApp Backup and Recovery, with unlimited backups for Cloud Volumes ONTAP volumes (backup capacity is not counted against the license).

- With the CVO Professional plan, there are two types of charges:

- **Resource charges:** Based on storage usage. For more information, refer to [licensing for Cloud Volumes ONTAP](#).

- **Service charges:** Fees for NetApp Backup and Recovery. However, if the source volume is in a storage system using the CVO Professional plan, NetApp Backup and Recovery is provided free of charge.

When you use Google Cloud Platform, request a private offer from NetApp and select your plan during activation in the Google Cloud Marketplace.

[Learn how to set up licenses.](#)

Supported workloads, systems, and backup targets

Supported workloads

NetApp Backup and Recovery protects the following types of workloads:

- ONTAP volumes
- Microsoft SQL Server instances and databases stored on physical disk and VMware Virtual Machine Disk (VMDK) over VMFS or NFS
- VMware VMs and datastores
- KVM workloads (Preview)
- Hyper-V workloads (Preview)
- Oracle Database workloads (Preview)
- Kubernetes workloads (Preview)

Supported systems

- On-premises ONTAP SAN (iSCSI protocol) and NAS (using NFS and CIFS protocols) with ONTAP version 9.8 or greater
- Cloud Volumes ONTAP 9.8 or greater for AWS (using SAN and NAS)
- Cloud Volumes ONTAP 9.8 or greater for Google Cloud Platform (using NFS and CIFS protocols)
- Cloud Volumes ONTAP 9.8 or greater for Microsoft Azure (using SAN and NAS)
- Amazon FSx for NetApp ONTAP (Microsoft SQL Server workloads only)

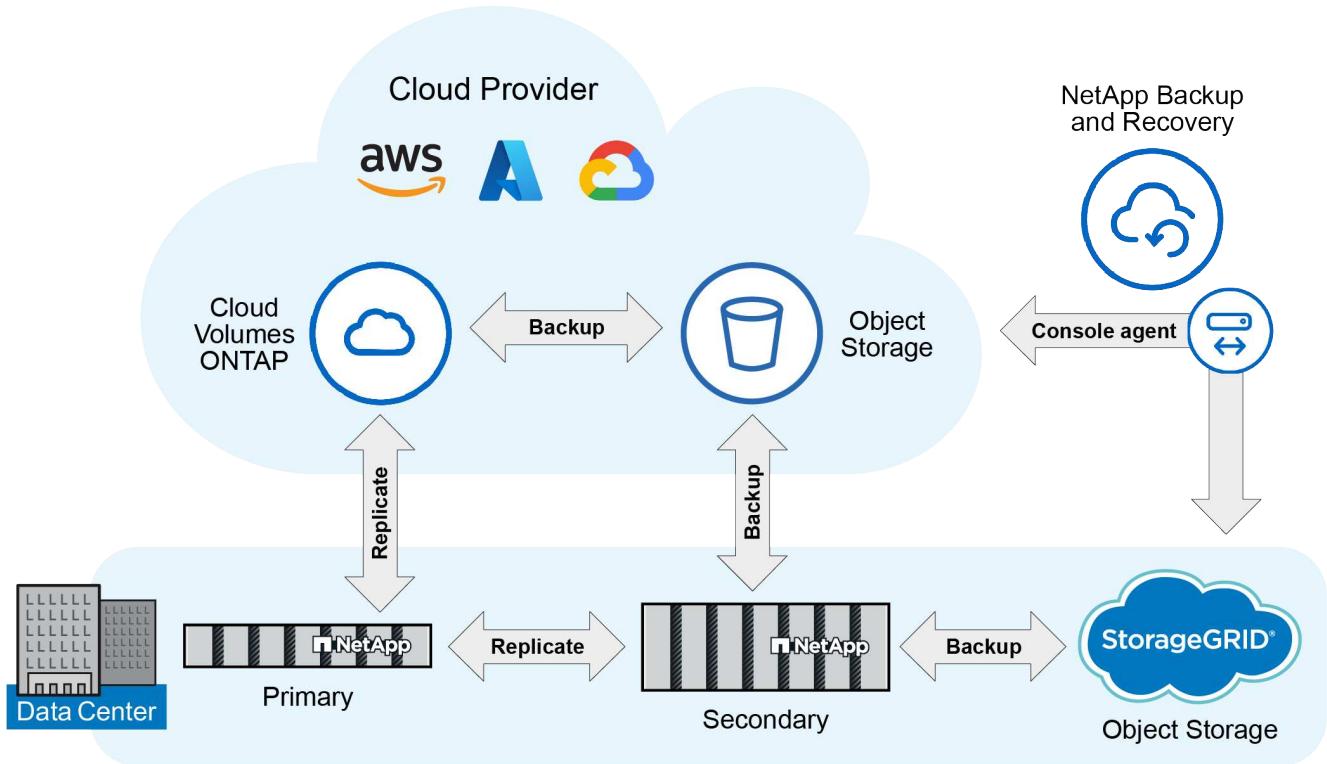
Supported backup targets

- Amazon Web Services (AWS) S3
- Google Cloud Storage
- Microsoft Azure Blob (not available for VMware workloads in Preview)
- StorageGRID
- ONTAP S3 (Not available for VMware workloads in Preview)

How NetApp Backup and Recovery works

When you enable NetApp Backup and Recovery, the service performs a full backup of your data. After the initial backup, all additional backups are incremental. This keeps network traffic to a minimum.

The following image shows the relationship among components.



Primary to object storage is also supported, not just from secondary storage to object storage.

Where backups reside in object store locations

Backup copies are stored in an object store that the NetApp Console creates in your cloud account. There's one object store per cluster or system, and the Console names the object store as follows: netapp-backup-clusteruuid. Be sure not to delete this object store.

- In AWS, the NetApp Console enables the [Amazon S3 Block Public Access feature](#) on the S3 bucket.
- In Azure, the NetApp Console uses a new or existing resource group with a storage account for the Blob container. the Console [blocks public access to your blob data](#) by default.
- In StorageGRID, the Console uses an existing storage account for the object store bucket.
- In ONTAP S3, the Console uses an existing user account for the S3 bucket.

Backup copies are associated with your NetApp Console organization

Backup copies are associated with the NetApp Console organization in which the Console agent resides. [Learn about NetApp Console Identity and access](#).

If you have multiple Console agents in the same NetApp Console organization, each Console agent displays the same list of backups.

Terms that might help you with NetApp Backup and Recovery

You might benefit by understanding some terminology related to protection.

- **Protection:** Protection in NetApp Backup and Recovery means ensuring that snapshots and immutable

backups occur on a regular basis to a different security domain using protection policies.

- **Workload:** A workload in NetApp Backup and Recovery can include ONTAP volumes, Microsoft SQL Server instances and databases; VMware VMs and datastores; or Kubernetes clusters and applications.

NetApp Backup and Recovery prerequisites

Get started with NetApp Backup and Recovery by verifying the readiness of your operational environment, NetApp Console agent, and NetApp Console account. To use NetApp Backup and Recovery, you'll need these prerequisites.

Prerequisite for ONTAP 9.8 and later

An ONTAP One license must be enabled on the on-premises ONTAP instance.

Prerequisites for backups to object storage

To use object storage as backup targets, you need an account with AWS S3, Microsoft Azure Blob, StorageGRID, or ONTAP and the appropriate access permissions configured.

- Protect your ONTAP volume data

Requirements for protecting Microsoft SQL Server workloads

To use NetApp Backup and Recovery for Microsoft SQL Server workloads, you need the following host system, space, and sizing prerequisites.

Item	Requirements
Operating systems	Microsoft Windows For the latest information about supported versions, see the NetApp Interoperability Matrix Tool .
Microsoft SQL Server versions	Version 2012 and later are supported for VMware Virtual Machine File System (VMFS) and VMware Virtual Machine Disk (VMDK) NFS.
SnapCenter Server version	SnapCenter Server version 5.0 or greater is required if you are going to import your existing data from SnapCenter into NetApp Backup and Recovery.  If you already have SnapCenter, first check to be sure you've met the prerequisites before importing from SnapCenter. See Prerequisites for importing resources from SnapCenter .
Minimum RAM for the plug-in on the SQL Server host	1 GB

Item	Requirements
Minimum install and log space for the plug-in on the SQL Server host	<p>5 GB</p> <p>Allocate sufficient disk space and monitor the storage consumption by the logs folder. The log space required varies depending on the number of backups performed and the frequency of data protection operations. If there is not sufficient space, the logs will not be created for the operations.</p>
Required software packages	<ul style="list-style-type: none"> ASP.NET Core Runtime 8.0.12 Hosting Bundle (and all subsequent 8.0.x patches) PowerShell 7.4.2 <p>For the latest information about supported versions, see the NetApp Interoperability Matrix Tool.</p>

Requirements for protecting VMware workloads

You need specific requirements to discover and protect your VMware workloads.

Software support

- NFS and VMFS datastores are supported.
- NFS versions supported: NFS 3 and NFS 4.1
- VMware ESXi Server versions supported: 7.0U1 and above
- VMware vCenter vSphere versions supported: 7.0U1 and above
- IP addresses: IPv4 and IPv6
- VMware TLS: 1.2, 1.3
- Connected storage supported: ONTAP 9.13 or later

Connection and port requirements for protecting VMware workloads

Type of port	Pre-configured port
VMware ESXi Server port	443 (HTTPS), bidirectional. The Guest File Restore feature uses this port.
Storage cluster or storage VM port	443 (HTTPS), bidirectional. 80 (HTTP), bidirectional. This port is used for communication between the virtual appliance and the storage VM or the cluster containing the storage VM.

Role-based access control (RBAC) requirements for protecting VMware workloads

The vCenter administrator account must have the required vCenter privileges.

For a list of vCenter privileges needed, see [SnapCenter Plug-in for VMware vSphere vCenter privileges needed](#).

Requirements for protecting KVM workloads

You need specific requirements to discover and protect KVM virtual machines.

- A modern Linux distribution running kernel version 5.14.0-503.22.1.el9_5.x86_64 (longterm) or later
- Your KVM hosts and VMs must be managed by a management platform. NetApp Backup and Recovery supports the following management platforms:
 - Apache CloudStack 4.22.0.0
- Ensure that inbound network traffic to port 22 is allowed from the Console agent to the KVM host
- QEMU Guest Agent version 9.0.0 or later
- libvirt version 10.5.0 or later



To ensure that KVM workload restores complete successfully, make sure that the **Enable VM-consistent snapshot** setting is active in the protection policy you use for KVM backups.

To enable protection of KVM VMs administered by non-root users, use the following steps:

1. Mount the volume as type NFS3 to avoid the use of the nobody user and group.
2. Use the following command to add a non-root user to the `qemu` group while preserving their existing groups:

```
usermod -aG qemu <non-root-user>
```



3. Use the following command to grant ownership of the mount path to the `qemu` user and group and change permissions for the mount path:

```
chown -R qemu:qemu <kvm_vm_mount_path> & chmod 771  
<kvm_vm_mount_path>
```

4. Delete the existing `NetApp_SnapCenter_Backups` directory if present.

Requirements for protecting Oracle Database workloads

Ensure your environment meets specific requirements to discover and protect Oracle resources.

- Oracle Database:
 - Oracle 19C and 21C are supported in a standalone deployment.
 - Oracle Database must be deployed in primary or secondary NetApp ONTAP storage.
 - Host OS support: Red Hat Enterprise Linux 8 and 9
- Object storage support:
 - Azure Object Storage
 - Amazon AWS
 - NetApp StorageGRID

- ONTAP S3

Requirements for protecting Kubernetes applications

You need specific requirements to discover Kubernetes resources and protect your Kubernetes applications.

For NetApp Console requirements, refer to [In NetApp Console](#).

- A primary ONTAP system (ONTAP 9.16.1 or later)
- A Kubernetes cluster - Supported Kubernetes distributions and versions include:
 - Anthos On-Prem (VMware) and Anthos on bare metal 1.16
 - Kubernetes 1.27 - 1.33
 - OpenShift 4.10 - 4.18
 - Rancher Kubernetes Engine 2 (RKE2) v1.26.7+rke2r1, v1.28.5+rke2r1
 - Suse Rancher
- NetApp Trident 24.10 or later
- NetApp Trident Protect 25.07 or later (installed during Kubernetes workload discovery)
- NetApp Trident Protect Connector 25.07 or later (installed during Kubernetes workload discovery)
 - Make sure that TCP port 443 is unfiltered in the outbound direction between the Kubernetes cluster, the Trident Protect Connector, and the Trident Protect proxy.

Requirements for protecting Hyper-V workloads

Ensure your Hyper-V instance meets specific requirements to discover and protect virtual machines.

- Software requirements for the Hyper-V Windows Server host:
 - Microsoft Hyper-V 2019, 2022 & 2025 editions
 - ASP.NET Core Runtime 8.0.12 Hosting Bundle (and all subsequent 8.0.x patches)
 - PowerShell 7.4.2 or later
 - If users that are not part of an administrator domain will be protecting Hyper-V VMs, ensure that the user has the following permissions:
 - Ensure the user is a member of the local administrators group.
 - Ensure the user is part of the "Log on as service" local security policy.
 - Ensure that two-way HTTPS traffic is allowed for the following ports in the Windows Firewall settings:
 - 8144 (NetApp Plugin for Hyper-V)
 - 8145 (NetApp Plugin for Windows)
- Hardware requirements for the Hyper-V host:
 - Standalone and FCI-clustered hosts are supported
 - 1GB RAM minimum for the NetApp Hyper-V plug-in on the Hyper-V host
 - 5GB minimum installation and log space for the plug-in on the Hyper-V Host



Ensure that you allocate enough disk space on the Hyper-V host for the logs folder and regularly monitor its usage. The required space depends on how often backups and data protection operations occur. If there isn't enough space, logs will not be generated.

- NetApp ONTAP configuration requirements:

- A primary ONTAP system (ONTAP 9.14.1 or later)
- For Hyper-V deployments using CIFS shares to store virtual machine data, ensure that the continuous availability share property is enabled on the ONTAP system. Refer to the [ONTAP documentation](#) for instructions.

In NetApp Console

Ensure NetApp Console meets the following requirements.

- A Console user should have the required role and privileges to perform operations on Microsoft SQL Server and Kubernetes workloads. To discover the resources, you must have the NetApp Backup and Recovery role of Super admin. See [NetApp Backup and Recovery role-based access to features](#) for details about the roles and permissions required to perform operations in NetApp Backup and Recovery.
- A Console organization with at least one active Console agent that connects to on-premises ONTAP clusters or Cloud Volumes ONTAP.
- At least one Console system with a NetApp on-premises ONTAP or Cloud Volumes ONTAP cluster.
- A Console agent

Refer to [Learn how to configure a Console agent](#) and [standard NetApp Console requirements](#).

- The preview version requires the Ubuntu 22.04 LTS operating system for the Console agent.

Set up NetApp Console

The next step is to set up the Console and NetApp Backup and Recovery.

Review [standard NetApp Console requirements](#).

Create a Console agent

You should reach out to your NetApp Product Team to try out Backup and Recovery. Then, when you use the Console agent, it will include the appropriate capabilities for the service.

To create a Console agent in the NetApp Console before using the service, refer to the Console documentation that describes [how to create a Console agent](#).

Where to install the Console agent

To complete a restore operation, the Console agent can be installed in the following locations:

- For Amazon S3, the Console agent can be deployed on your premises.
- For Azure Blob, the Console agent can be deployed on your premises.
- For StorageGRID, the Console agent must be deployed in your premises; with or without internet access.
- For ONTAP S3, the Console agent can be deployed in your premises (with or without internet access) or in a cloud provider environment



References to "on-premises ONTAP systems" includes FAS and AFF systems.

Set up licensing for NetApp Backup and Recovery

You can license NetApp Backup and Recovery by purchasing a pay-as-you-go (PAYGO) or annual marketplace subscription to **NetApp Intelligent Services** from your cloud provider, or by purchasing a bring-your-own-license (BYOL) from NetApp. A valid license is required to activate NetApp Backup and Recovery on a system, to create backups of your production data, and to restore backup data to a production system.

A few notes before you read any further:

- If you've already subscribed to the pay-as-you-go (PAYGO) subscription in your cloud provider's marketplace for a Cloud Volumes ONTAP system, then you're automatically subscribed to NetApp Backup and Recovery as well. You won't need to subscribe again.
- The NetApp Backup and Recovery bring-your-own-license (BYOL) is a floating license that you can use across all systems associated with your NetApp Console organization or account. So if you have sufficient backup capacity available from an existing BYOL license, you won't need to purchase another BYOL license.
- If you are using a BYOL license, it is recommended that you subscribe to a PAYGO subscription as well. If you back up more data than allowed by your BYOL license, or if the term of your license expires, then backup continues through your pay-as-you-go subscription - there is no disruption of service.
- When backing up on-prem ONTAP data to StorageGRID, you need a BYOL license, but there's no cost for cloud provider storage space.

[Learn more about the costs related to using NetApp Backup and Recovery.](#)

30-day free trial

A NetApp Backup and Recovery 30-day free trial is available if you sign up for a pay-as-you-go subscription in your cloud provider's marketplace to **NetApp Intelligent Services**. The free trial starts at the time that you subscribe to the marketplace listing. Note that if you pay for the marketplace subscription when deploying a Cloud Volumes ONTAP system, and then start your NetApp Backup and Recovery free trial 10 days later, you'll have 20 days remaining to use the free trial.

When the free trial ends, you'll be switched over automatically to the PAYGO subscription without interruption. If you decide not to continue using NetApp Backup and Recovery, just [unregister NetApp Backup and Recovery from the system](#) before the trial ends and you won't be charged.

End the free trial

If you want to continue using NetApp Backup and Recovery after the free trial ends, you must set up a paid subscription. You can do this from the NetApp Console interface by navigating to the billing section and selecting a subscription plan that fits your needs. If you don't want to continue using NetApp Backup and Recovery, you can end the free trial.

When you end the free trial without subscribing to a paid plan, your data is automatically deleted 60 days after the free trial ends. You can optionally have the system delete your data immediately.

Steps

1. From the NetApp Backup and Recovery landing page, select **View free trial**.
2. Select **End free trial**.
3. Select **Delete data immediately after ending my free trial** to delete your data immediately.
4. Type **end trial** in the box.
5. Select **End** to confirm.

Use a NetApp Backup and Recovery PAYGO subscription

For pay-as-you-go, you'll pay your cloud provider for object storage costs and for NetApp backup licensing costs on an hourly basis in a single subscription. You should subscribe to **NetApp Intelligent Services** in the Marketplace even if you have a free trial or if you bring your own license (BYOL):

- Subscribing ensures that there's no disruption of service after your free trial ends. When the trial ends, you'll be charged hourly according to the amount of data that you back up.
- If you back up more data than allowed by your BYOL license, then data backup and restore operations continue through your pay-as-you-go subscription. For example, if you have a 10 TiB BYOL license, all capacity beyond the 10 TiB is charged through the PAYGO subscription.

You won't be charged from your pay-as-you-go subscription during your free trial or if you haven't exceeded your BYOL license.

There are a few PAYGO plans for NetApp Backup and Recovery:

- A "Cloud Backup" package that enables you to back up Cloud Volumes ONTAP data and on-premises ONTAP data.
- A "CVO Professional" package that enables you to bundle Cloud Volumes ONTAP and NetApp Backup and Recovery. This includes unlimited backups for the Cloud Volumes ONTAP system using the license (backup capacity is not counted against the licensed capacity). This option doesn't enable you to back up on-premises ONTAP data.

Note that this option also requires a Backup and recovery PAYGO subscription, but no charges will be incurred for eligible Cloud Volumes ONTAP systems.

[Learn more about these capacity-based license packages](#).

Use these links to subscribe to NetApp Backup and Recovery from your cloud provider marketplace:

- AWS: [Go to the Marketplace offering for NetApp Intelligent Services for pricing details](#).
- Azure: [Go to the Marketplace offering for NetApp Intelligent Services for pricing details](#).
- Google Cloud: [Go to the Marketplace offering for NetApp Intelligent Services for pricing details](#).

Use an annual contract

Pay for NetApp Backup and Recovery annually by purchasing an annual contract. They're available in 1-, 2-, or 3-year terms.

If you have an annual contract from a marketplace, all NetApp Backup and Recovery consumption is charged against that contract. You can't mix and match an annual marketplace contract with a BYOL.

When you use AWS, there are two annual contracts available from the [AWS Marketplace page](#) for Cloud

Volumes ONTAP and on-premises ONTAP systems:

- A "Cloud Backup" plan that enables you to back up Cloud Volumes ONTAP data and on-premises ONTAP data.

If you want to use this option, set up your subscription from the Marketplace page and then [associate the subscription with your AWS credentials](#). Note that you'll also need to pay for your Cloud Volumes ONTAP systems using this annual contract subscription since you can assign only one active subscription to your AWS credentials in the Console.

- A "CVO Professional" plan that enables you to bundle Cloud Volumes ONTAP and NetApp Backup and Recovery. This includes unlimited backups for the Cloud Volumes ONTAP system using the license (backup capacity is not counted against the licensed capacity). This option doesn't enable you to back up on-premises ONTAP data.

See the [Cloud Volumes ONTAP licensing topic](#) to learn more about this licensing option.

If you want to use this option, you can set up the annual contract when you create a Cloud Volumes ONTAP system and the Console prompts you to subscribe to the AWS Marketplace.

When you use Azure, there are two annual contracts available from the [Azure Marketplace page](#) for Cloud Volumes ONTAP and on-premises ONTAP systems:

- A "Cloud Backup" plan that enables you to back up Cloud Volumes ONTAP data and on-premises ONTAP data.

If you want to use this option, set up your subscription from the Marketplace page and then [associate the subscription with your Azure credentials](#). Note that you'll also need to pay for your Cloud Volumes ONTAP systems using this annual contract subscription since you can assign only one active subscription to your Azure credentials in the Console.

- A "CVO Professional" plan that enables you to bundle Cloud Volumes ONTAP and NetApp Backup and Recovery. This includes unlimited backups for the Cloud Volumes ONTAP system using the license (backup capacity is not counted against the licensed capacity). This option doesn't enable you to back up on-premises ONTAP data.

See the [Cloud Volumes ONTAP licensing topic](#) to learn more about this licensing option.

If you want to use this option, you can set up the annual contract when you create a Cloud Volumes ONTAP system and the Console prompts you to subscribe to the Azure Marketplace.

When you use GCP, contact your NetApp sales representative to purchase an annual contract. The contract is available as a private offer in the Google Cloud Marketplace.

After NetApp shares the private offer with you, you can select the annual plan when you subscribe from the Google Cloud Marketplace during NetApp Backup and Recovery activation.

Use a NetApp Backup and Recovery BYOL license

Bring-your-own licenses from NetApp provide 1-, 2-, or 3-year terms. You pay only for the data that you protect, calculated by the logical used capacity (*before* any efficiencies) of the source ONTAP volumes which are being backed up. This capacity is also known as Front-End Terabytes (FETB).

The BYOL NetApp Backup and Recovery license is a floating license where the total capacity is shared across

all systems associated with your NetApp Console organization or account. For ONTAP systems, you can get a rough estimate of the capacity you'll need by running the CLI command `volume show -fields logical-used-by-afs` for the volumes you plan to back up.

If you don't have a NetApp Backup and Recovery BYOL license, click the chat icon in the lower-right of the Console to purchase one.

Optionally, if you have an unassigned node-based license for Cloud Volumes ONTAP that you won't be using, you can convert it to a NetApp Backup and Recovery license with the same dollar-equivalence and the same expiration date. [Go here for details.](#)

You use the NetApp Console to manage BYOL licenses. You can add new licenses, update existing licenses, and view license status from the Console.

[Learn about adding licenses.](#)

Set up security certificates for StorageGRID and ONTAP in NetApp Backup and Recovery

Create a security certificate to enable communication between NetApp Backup and Recovery and StorageGRID or ONTAP.

Create a security certificate for StorageGRID

If the communication between NetApp Backup and Recovery containers and StorageGRID should verify the StorageGRID certificate, then complete following steps.

The generated certificate should have CN and Subject Alternative Name as the name provided in NetApp Backup and Recovery when you were activating the backup.

Steps

1. Follow the steps in the StorageGRID documentation to create the StorageGRID certificate.

[StorageGRID information on configuring certificates](#)

2. Update StorageGRID with the certificate if you have not already done so.
3. Log in to the Console agent as a root user. Run:

```
sudo su
```

4. Get the NetApp Backup and Recovery (Cloud Backup Service) Docker volume. Run:

```
docker volume ls | grep cbs
```

Output example:

```
local service-manager-2_cloudmanager_cbs_volume"
```



The volume name differs among Standard, Private, and Restricted deployment modes. This example uses Standard mode. Refer to [NetApp Console deployment modes](#).

5. Find the mount point of the NetApp Backup and Recovery volume. Run:

```
docker volume inspect service-manager-2_cloudmanager_cbs_volume | grep Mountpoint
```

Output example:

```
"Mountpoint": "/var/lib/docker/volumes/service-manager-2_cloudmanager_cbs_volume/_data"
```



The mount point differs among Standard, Private, and Restricted deployment modes. This example shows a Standard cloud deployment. Refer to [NetApp Console deployment modes](#).

6. Change to the MountPoint directory. Run:

```
cd /var/lib/docker/volumes/service-manager-2_cloudmanager_cbs_volume/_data
```

7. If StorageGRID's certificate is signed by the root CA and an intermediate CA, then append the pem files of both into one file named `sgws.crt` in the current location. Do not add the leaf certificate to this file.

Steps for `cloudmanager_cbs` container

You'll need to enable the StorageGRID Server certificate verification in NetApp Backup and Recovery (Cloud Backup Service).

1. Change directories to the Docker volume obtained in earlier steps.

```
cd /var/lib/docker/volumes/service-manager-2_cloudmanager_cbs_volume/_data
```

2. Change directories to the config directory.

```
cd cbs_config
```

3. Create and save a configuration file as shown below with one of the following names based on your deployment environment:

- `production-customer.json` Used for Standard mode and Restricted mode deployments.
- `darksite-customer.json` Used for Private mode deployments.

Refer to [NetApp Console deployment modes](#).

Configuration file

```
{  
  "protocols": {  
    "sgws": {  
      "certificates": {  
        "reject-unauthorized": true,  
        "ca-bundle": "/config/sgws.crt"  
      }  
    }  
  }  
}
```

4. Exit the container. Run:

```
exit
```

5. Restart `cloudmanager_cbs`. Run:

```
docker restart cloudmanager_cbs
```

Steps for `cloudmanager_cbs_catalog` container

Next, you'll need to enable the StorageGRID Server certificate verification for the Cataloging Service.

1. Change directories to the Docker volume:

```
cd /var/lib/docker/volumes/service-manager-  
2_cloudmanager_cbs_volume/_data
```

2. Configure the catalog. Run:

```
cd cbs_catalog_config
```

3. Create a config file as shown below with one of the following names based on your deployment environment:

- `production-customer.json` Used for Standard mode and Restricted mode deployments.
- `darksite-customer.json` Used for Private mode deployments.

Refer to [NetApp Console deployment modes](#).

Catalog configuration file

```
{  
  "protocols": {  
    "sgws": {  
      "certificates": {  
        "reject-unauthorized": true,  
        "ca-bundle": "/config/sgws.crt"  
      }  
    }  
  }  
}
```

4. Restart the catalog. Run:

```
docker restart cloudmanager_cbs_catalog
```

Update the Console agent certificate with the StorageGRID certificate based on the agent operating system

Ubuntu

1. Copy the SGWS certificate to /usr/local/share/ca-certificates. Here is an example:

```
cp /config/sgws.crt /usr/local/share/ca-certificates/
```

where `sgws.crt` is the root CA certificate.

2. Update the host certificates with the StorageGRID certificate. Run

```
sudo update-ca-certificates
```

Red Hat Enterprise Linux

1. Copy the SGWS certificate to /etc/pki/ca-trust/source/anchors/.

```
cp /config/sgws.crt /etc/pki/ca-trust/source/anchors/
```

where `sgws.crt` is the root CA certificate.

2. Update the host certificates with the StorageGRID certificate.

```
update-ca-trust extract
```

3. Update the ca-bundle.crt

```
cd /etc/pki/tls/certs/  
openssl x509 -in ca-bundle.crt -text -noout
```

4. To check whether the certificates are present, run the following command:

```
openssl crl2pkcs7 -nocrl -certfile /etc/pki/tls/certs/ca-bundle.crt |  
openssl pkcs7 -print_certs | grep subject | head
```

Create a security certificate for ONTAP

If the communication between the NetApp Backup and Recovery containers and ONTAP should validate the ONTAP certificate, then complete the following steps.

NetApp Backup and Recovery uses the Cluster Management IP to connect to ONTAP. Enter the IP address of the cluster in the Subject Alternative names of the Certificate. Specify this step when you generate the CSR using the System Manager UI.

Use the System Manager documentation to create a new CA certificate for ONTAP.

- [Manage certificates with System Manager](#)
- [How to manage ONTAP SSL certificates with System Manager](#)

Steps

1. Login to the Console agent as root. Run:

```
sudo su
```

2. Get the NetApp Backup and Recovery Docker volume. Run:

```
docker volume ls | grep cbs
```

Output example:

```
local service-manager-2_cloudmanager_cbs_volume
```



The volume name differs among Standard, Private, and Restricted deployment modes. This example shows a Standard cloud deployment. Refer to [NetApp Console deployment modes](#).

3. Obtain the mount for the volume. Run:

```
docker volume inspect service-manager-2_cloudmanager_cbs_volume | grep Mountpoint
```

Output example:

```
"Mountpoint": "/var/lib/docker/volumes/service-manager-2_cloudmanager_cbs_volume/_data"
```



The mount point differs among Standard, Private, and Restricted deployment modes. This example shows a Standard cloud deployment. Refer to [NetApp Console deployment modes](#).

4. Change to the mountpoint directory. Run:

```
cd /var/lib/docker/volumes/service-manager-2_cloudmanager_cbs_volume/_data
```

5. Complete one of the following steps:

- If the ONTAP certificate is signed by the root CA and an intermediate CA, then append the pem files of both into one file named `ontap.crt` in the current location.
- If the ONTAP certificate is signed by a single CA, then rename the pem file as `ontap.crt` and copy it in the current location. Do not add the leaf certificate to this file.

Steps for `cloudmanager_cbs` container

Next, enable the ONTAP Server certificate verification in NetApp Backup and Recovery (Cloud Backup Service).

1. Change directories to the Docker volume obtained in earlier steps.

```
cd /var/lib/docker/volumes/service-manager-2_cloudmanager_cbs_volume/_data
```

2. Change to the config directory. Run:

```
cd cbs_config
```

3. Create a configuration file as shown below with one of the following names based on your deployment environment:

- `production-customer.json` Used for Standard mode and Restricted mode deployments.
- `darksite-customer.json` Used for Private mode deployments.

Refer to [NetApp Console deployment modes](#).

Configuration file

```
{  
  "ontap": {  
    "certificates": {  
      "reject-unauthorized": true,  
      "ca-bundle": "/config/ontap.crt"  
    }  
  }  
}
```

4. Exit the container. Run:

```
exit
```

5. Restart NetApp Backup and Recovery. Run:

```
docker restart cloudmanager_cbs
```

Steps for cloudmanager_cbs_catalog container

Enable the ONTAP Server certificate verification for the Cataloging Service.

1. Change directories to the Docker volume. Run:

```
cd /var/lib/docker/volumes/service-manager-  
2_cloudmanager_cbs_volume/_data
```

2. Run:

```
cd cbs_catalog_config
```

3. Create a configuration file as shown below with one of the following names based on your deployment environment:

- `production-customer.json` Used for Standard mode and Restricted mode deployments.
- `darksite-customer.json` Used for Private mode deployments.

Refer to [NetApp Console deployment modes](#).

Configuration file

```
{  
  "ontap": {  
    "certificates": {  
      "reject-unauthorized": true,  
      "ca-bundle": "/config/ontap.crt"  
    }  
  }  
}
```

4. Restart NetApp Backup and Recovery. Run:

```
docker restart cloudmanager_cbs_catalog
```

Create a certificate for both ONTAP and StorageGRID

If you need to enable the certificate for both ONTAP and StorageGRID, then the configuration file looks like this:

Configuration file for both ONTAP and StorageGRID

```
{  
  "protocols": {  
    "sgws": {  
      "certificates": {  
        "reject-unauthorized": true,  
        "ca-bundle": "/config/sgws.crt"  
      }  
    }  
  },  
  "ontap": {  
    "certificates": {  
      "reject-unauthorized": true,  
      "ca-bundle": "/config/ontap.crt"  
    }  
  }  
}
```

Set up backup destinations before you use NetApp Backup and Recovery

Before you use NetApp Backup and Recovery, perform a few steps to set up backup destinations.

Before you begin, review [prerequisites](#) to ensure that your environment is ready.

Prepare the backup destination

Prepare one or more of the following backup destinations:

- NetApp StorageGRID.

Refer to [Discover StorageGRID](#).

Refer to [StorageGRID documentation](#) for details about StorageGRID.

- Amazon Web Services. Refer to [Amazon S3 documentation](#).

Do the following to prepare AWS as a backup destination:

- Set up an account in AWS.
- Configure S3 permissions in AWS, listed in the next section.
- For details about managing your AWS storage in the Console, refer to [Manage your Amazon S3 buckets](#).

- Microsoft Azure.
 - Refer to [Azure NetApp Files documentation](#).
 - Set up an account in Azure.
 - Configure [Azure permissions](#) in Azure.
 - For details about managing your Azure storage in the Console, refer to [Manage your Azure storage accounts](#).

After you configure options in the backup destination itself, you will later configure it as a backup destination in NetApp Backup and Recovery. For details about how to configure the backup destination in NetApp Backup and Recovery, refer to [Discover backup targets](#).

Set up S3 permissions

You'll need to configure two sets of AWS S3 permissions:

- Permissions for the Console agent to create and manage the S3 bucket.
- Permissions for the on-premises ONTAP cluster so it can read and write data to the S3 bucket.

Steps

1. Ensure that the Console agent has the required permissions. For details, see [NetApp Console policy permissions](#).



When creating backups in AWS China regions, you need to change the AWS Resource Name "arn" under all *Resource* sections in the IAM policies from "aws" to "aws-cn"; for example `arn:aws-cn:s3:::netapp-backup-*`.

2. When you activate the service, the Backup wizard will prompt you to enter an access key and secret key. These credentials are passed to the ONTAP cluster so that ONTAP can back up and restore data to the S3 bucket. For that, you'll need to create an IAM user with the following permissions.

Refer to the [AWS Documentation: Creating a Role to Delegate Permissions to an IAM User](#).

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Action": [  
        "s3:GetObject",  
        "s3:PutObject",  
        "s3:DeleteObject",  
        "s3>ListBucket",  
        "s3>ListAllMyBuckets",  
        "s3:GetBucketLocation",  
        "s3:PutEncryptionConfiguration"  
      ],  
      "Resource": "arn:aws:s3:::netapp-backup-*",  
      "Effect": "Allow",  
      "Sid": "backupPolicy"  
    },  
    {  
      "Action": [  
        "s3>ListBucket",  
        "s3:GetBucketLocation"  
      ],  
      "Resource": "arn:aws:s3:::netapp-backup*",  
      "Effect": "Allow"  
    },  
    {  
      "Action": [  
        "s3:GetObject",  
        "s3:PutObject",  
        "s3:DeleteObject",  
        "s3>ListAllMyBuckets",  
        "s3:PutObjectTagging",  
        "s3:GetObjectTagging",  
        "s3:RestoreObject",  
        "s3:GetBucketObjectLockConfiguration",  
        "s3:GetObjectRetention",  
        "s3:PutBucketObjectLockConfiguration",  
        "s3:PutObjectRetention"  
      ],  
      "Resource": "arn:aws:s3:::netapp-backup*/*",  
      "Effect": "Allow"  
    }  
  ]  
}
```

Log in to NetApp Backup and Recovery

You use the NetApp Console to log in to NetApp Backup and Recovery.

NetApp Backup and Recovery uses identity and access management to control what each user can do.

For details about the actions that each role can perform, see [NetApp Backup and Recovery user roles](#).

To log in to the NetApp Console, you can use your NetApp Support Site credentials or you can sign up for a NetApp Console login using your email and a password. [Learn more about logging in](#).

Required NetApp Console role

Backup and Recovery super admin or Backup and Recovery restore admin role. [Learn about NetApp Console access roles for all services](#).

To add a Console agent, you must have the Backup and Recovery super admin role.

Steps

1. Open a web browser and go to the [NetApp Console](#).

The NetApp Console login page appears.

2. Log in to the Console.

3. From the Console left navigation, select **Protection > Backup and Recovery**.

- If this is your first time logging in to Backup and Recovery and you haven't yet added a system to the **Systems** page, the Backup and Recovery displays the "Welcome to the new NetApp Backup and Recovery" landing page with an option to add a system. For details about adding a system to the **Systems** page, refer to [Getting started with NetApp Console standard mode](#).
- If you are logging in to Backup and Recovery for the first time and have a system in the Console but no discovered resources, the *Welcome to the new NetApp Backup and Recovery* page appears with an option to **Discover resources**.

4. If you haven't done so already, select the **Discover and manage** option.

- For Microsoft SQL Server workloads, refer to [Discover Microsoft SQL Server workloads](#).
- For VMware workloads, refer to [Discover VMware workloads](#).
- For KVM workloads, refer to [Discover KVM workloads](#).
- For Oracle Database workloads, refer to [Discover Oracle Database workloads](#).
- For Hyper-V workloads, refer to [Discover Hyper-V workloads](#).
- For Kubernetes workloads, refer to [Discover Kubernetes workloads](#).

Discover offsite backup targets in NetApp Backup and Recovery

Complete a few steps to discover or manually add offsite backup targets in NetApp Backup and Recovery.

Discover a backup target

Configure your backup targets (Amazon Web Services (AWS) S3, Microsoft Azure Blob Storage, Google Cloud Storage, or StorageGRID) before using NetApp Backup and Recovery.

You can discover these targets automatically or manually add them.

Provide credentials to access the storage account. NetApp Backup and Recovery uses these credentials to discover the workloads you want to back up.

Before you begin

You need to discover at least one workload before you can add an offsite backup target.

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select the **Offsite backup targets** tab.
3. Select **Discover backup target**.
4. Select one of the backup target types: **Amazon Web Services (AWS) S3**, **Microsoft Azure Blob Storage**, **StorageGRID** or **ONTAP S3**.
5. In the **Choose credentials location** section, choose the location where the credentials reside and then choose how to associate the credentials.
6. Select **Next**.
7. Enter the credentials information. The information differs depending on the type of backup target you selected and the credentials location that you chose.
 - For AWS:
 - **Credential name**: Enter the AWS credential name.
 - **Access key**: Enter the AWS secret.
 - **Secret key**: Enter the AWS secret key.
 - For Azure:
 - **Credential name**: Enter the Azure Blob Storage credential name.
 - **Client secret**: Enter the Azure Blob Storage client secret.
 - **Application (client) ID**: Select the Azure Blob Storage application ID.
 - **Directory tenant ID**: Enter the Azure Blob Storage tenant ID.
 - For StorageGRID:
 - **Credential name**: Enter the StorageGRID credential name.
 - **Gateway Node FQDN**: Enter a FQDN name for StorageGRID.
 - **Port**: Enter the port number for StorageGRID.
 - **Access key**: Enter the StorageGRID S3 access key.
 - **Secret key**: Enter the StorageGRID S3 secret key.
 - For ONTAP S3:
 - **Credential name**: Enter the ONTAP S3 credential name.
 - **Gateway Node FQDN**: Enter a FQDN name for ONTAP S3.

- **Port:** Enter the port number for ONTAP S3.
- **Access key:** Enter the ONTAP S3 access key.
- **Secret key:** Enter the ONTAP S3 secret key.

8. Select **Discover**.

Add a bucket for a backup target

Rather than have NetApp Backup and Recovery discover buckets automatically, you can manually add a bucket to an offsite backup target.

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select **Offsite backup targets**.
3. Select the target and on the right, select the **Actions**  icon and select **Add bucket**.
4. Enter the bucket information. The information differs depending on the type of backup target you selected.
 - For AWS:
 - **Bucket name:** Enter the name of the S3 bucket. The prefix of "netapp-backup" is a required prefix and is automatically added to the name you provide.
 - **AWS account:** Enter the AWS account name.
 - **Bucket region:** Enter the AWS region for the bucket.
 - **Enable S3 Object Lock:** Select this option to enable S3 Object Lock for the bucket. S3 Object Lock prevents objects from being deleted or overwritten for a specified retention period, providing an additional layer of data protection. You can enable this only when you are creating a bucket and you cannot turn it off later.
 - **Governance mode:** Select this option to enable governance mode for the S3 Object Lock bucket. Governance mode enables you to protect objects from being deleted or overwritten by most users, but allows certain users to alter the retention settings.
 - **Compliance mode:** Select this option to enable compliance mode for the S3 Object Lock bucket. Compliance mode prevents any user, including the root user, from altering the retention settings or deleting objects until the retention period expires.
 - **Versioning:** Select this option to enable versioning for the S3 bucket. Versioning enables you to keep multiple versions of objects in the bucket, which can be useful for backup and recovery purposes.
 - **Tags:** Select tags for the S3 bucket. Tags are key-value pairs that can be used to organize and manage your S3 resources.
 - **Encryption:** Select the type of encryption for the S3 bucket. The options are either AWS S3-managed keys or AWS Key Management Service key. If you select AWS Key Management Service keys, you must provide the key ID.
 - For Azure:
 - **Subscription:** Select the name of the Azure Blob Storage container.
 - **Resource group:** Select the name of the Azure resource group.
 - **Instance details:**
 - **Storage account name:** Enter the name of the Azure Blob Storage container.

- **Azure region:** Enter the Azure region for the container.
- **Performance type:** Select the performance type of either standard or premium for the Azure Blob Storage container indicating the level of performance required.
- **Encryption:** Select the type of encryption for the Azure Blob Storage container. The options are either Microsoft-managed keys or customer-managed keys. If you select customer-managed keys, you must provide the key vault name and key name.
- For StorageGRID:
 - **Backup target name:** Select the name of the StorageGRID bucket.
 - **Bucket name:** Enter the name of the StorageGRID bucket.
 - **Region:** Enter the StorageGRID region for the bucket.
 - **Enable versioning:** Select this option to enable versioning for the StorageGRID bucket. Versioning enables you to keep multiple versions of objects in the bucket, which can be useful for backup and recovery purposes.
 - **Object locking:** Select this option to enable object locking for the StorageGRID bucket. Object locking prevents objects from being deleted or overwritten for a specified retention period, providing an additional layer of data protection. You can enable this only when you are creating a bucket and you cannot turn it off later.
 - **Capacity:** Enter the capacity for the StorageGRID bucket. This is the maximum amount of data that can be stored in the bucket.
- For ONTAP S3:
 - **Backup target name:** Select the name of the ONTAP S3 bucket.
 - **Bucket target name:** Enter the name of the ONTAP S3 bucket.
 - **Capacity:** Enter the capacity for the ONTAP S3 bucket. This is the maximum amount of data that can be stored in the bucket.
 - **Enable versioning:** Select this option to enable versioning for the ONTAP S3 bucket. Versioning enables you to keep multiple versions of objects in the bucket, which can be useful for backup and recovery purposes.
 - **Object locking:** Select this option to enable object locking for the ONTAP S3 bucket. Object locking prevents objects from being deleted or overwritten for a specified retention period, providing an additional layer of data protection. You can enable this only when you are creating a bucket and you cannot turn it off later.

5. Select **Add**.

Change credentials for a backup target

Enter the credentials needed to access the backup target.

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select **Offsite backup targets**.
3. Select the target and on the right, select the **Actions** **...** icon and select **Change credentials**.
4. Enter the new credentials for the backup target. The information differs depending on the type of backup target you selected.
5. Select **Done**.

Switch to different NetApp Backup and Recovery workloads

You can switch among the different NetApp Backup and Recovery workloads.

Switch to a different workload

You can switch to a different workload in the NetApp Backup and Recovery UI.

Steps

1. From the Console left navigation, select **Protection > Backup and Recovery**.
2. From the top right corner of the page, select the **Switch workload** drop-down list.
3. Select the workload that you want to switch to.

The page refreshes and shows the selected workload.

Configure NetApp Backup and Recovery settings

After you set up NetApp Console, configure Backup and Recovery settings. Add credentials for host resources, import SnapCenter resources, configure log directories, and set VMware vCenter settings. Complete these steps before backing up or recovering data.

- [Add credentials for host resources](#) for any Windows, Microsoft SQL Server, Oracle Database, or Linux hosts that NetApp Backup and Recovery needs to authenticate with. This includes Windows guest OS credentials used when restoring guest files or folders.
- [Maintain VMware vCenter settings](#).
- [Import and manage SnapCenter host resources](#). (Microsoft SQL Server workloads only)
- [Add a KVM management platform](#). (KVM workloads only)
- [Configure log directories in snapshots for Windows hosts](#).
- [Create an execution hook template](#) to run scripts before and after backup jobs. (Kubernetes workloads only)

Required NetApp Console role

Backup and Recovery super admin, Backup and Recovery backup admin, Backup and Recovery restore admin. Learn about [Backup and recovery roles and privileges](#). Learn about [NetApp Console access roles for all services](#).

Add credentials for host resources

Add credentials for host resources. NetApp Backup and Recovery uses these credentials to discover workloads and apply backup policies.

If you don't have credentials, create them with permissions to access and manage host workloads.

You need to configure the following types of credentials:

- Microsoft SQL Server credentials

- SnapCenter Windows host credentials
- Windows guest OS credentials used when restoring guest files or folders
- Oracle Database credentials
- Linux host credentials

Steps

1. From the NetApp Backup and Recovery menu, select **Settings**.
2. Select the down arrow for **Credentials**.
3. Select **Add new credentials**.
4. Enter information for the credentials. Different fields appear depending on the Authentication mode you select. Hover over the Information **i** icon for more information about the fields.
 - **Credentials name:** Enter a name for the credentials.
 - **Authentication mode:** Select **Windows**, **Microsoft SQL**, **Oracle Database**, or **Linux**.



For Microsoft SQL Server workloads, you need to enter credentials for both Windows and Microsoft SQL Server, so you'll need to add two sets of credentials.

Windows

1. If you selected **Windows**:

- **Agents**: Select a Console agent from the list.
- **Domain and user name**: Enter the NetBIOS or domain FQDN and user name for the credentials.
- **Password**: Enter the password for the credentials.

Microsoft SQL Server

1. If you selected **Microsoft SQL Server**:

- **Domain and user name**: Enter the NetBIOS or domain FQDN and user name for the credentials.
- **Password**: Enter the password for the credentials.
- **Hosts**: Select a discovered SQL Server host address.
- **SQL Server instance**: Select a discovered SQL Server instance.

Oracle Database

1. If you selected **Oracle Database**:

- **Agents**: Select a Console agent from the list.
- **User name**: Enter the user name for the credentials.
- **Password**: Enter the password for the credentials.

Linux

1. If you selected **Linux**:

- **Agents**: Select a Console agent from the list.
- **User name**: Enter the user name for the credentials.
- **Password**: Enter the password for the credentials.

5. Select **Add**.

Edit credentials for host resources

You can later edit the password for any credentials that you have created.

Steps

1. From the NetApp Backup and Recovery menu, select **Settings**.
2. Select the down arrow to expand the **Credentials** section.
3. Select the Actions icon **...** > **Edit credentials**.
 - **Password**: Enter the password for the credentials.
4. Select **Save**.

Maintain VMware vCenter settings

Provide VMware vCenter credentials to discover workloads for backup. If you don't have credentials, create

them with permissions to access and manage the VMware vCenter Server workloads.

Steps

1. From the NetApp Backup and Recovery menu, select **Settings**.
2. Select the down arrow to expand the **VMware vCenter** section.
3. Select **Add vCenter**.
4. Enter the VMware vCenter Server information.
 - **vCenter FQDN or IP address**: Enter an FQDN name or the IP address for the VMware vCenter Server.
 - **Username and Password**: Enter the username and password for the VMware vCenter Server.
 - **Port**: Enter the port number for the VMware vCenter Server.
 - **Protocol**: Select **HTTP** or **HTTPS**.
5. Select **Add**.

Import and manage SnapCenter host resources

If you previously used SnapCenter to back up your resources, you can import and manage those resources in NetApp Backup and Recovery. This option enables you to import SnapCenter server information to register multiple SnapCenter servers and discover database workloads.

This is a two-part process:

- Import SnapCenter Server application and host resources
- Manage selected SnapCenter host resources

Import SnapCenter Server application and host resources

This first step imports host resources from SnapCenter and displays those resources in the NetApp Backup and Recovery Inventory page. At that point, the resources are not yet managed by NetApp Backup and Recovery.



After you import SnapCenter host resources, NetApp Backup and Recovery does not take over protection management. To do so, you must explicitly select to manage these resources in NetApp Backup and Recovery.

Steps

1. From the NetApp Backup and Recovery menu, select **Settings**.
2. Select the down arrow to expand the **Import from SnapCenter** section.
3. Select **Import from SnapCenter** to import the SnapCenter resources.
4. Enter **SnapCenter application credentials**:
 - a. **SnapCenter FQDN or IP address**: Enter the FQDN or IP address of the SnapCenter application itself.
 - b. **Port**: Enter the port number for the SnapCenter Server.
 - c. **Username and Password**: Enter the username and password for the SnapCenter Server.
 - d. **Console agent**: Select the Console agent for SnapCenter.
5. Enter **SnapCenter server host credentials**:

- a. **Existing credentials:** If you select this option, you can use the existing credentials that you have already added. Enter the credentials name.
- b. **Add new credentials:** If you don't have existing SnapCenter host credentials, you can add new credentials. Enter the credentials name, authentication mode, user name, and password.

6. Select **Import** to validate your entries and register the SnapCenter Server.



If the SnapCenter Server is already registered, you can update the existing registration details.

Result

The Inventory page shows the imported SnapCenter resources.

Manage SnapCenter host resources

After you import the SnapCenter resources, manage those host resources in NetApp Backup and Recovery. After you select to manage those imported resources, NetApp Backup and Recovery can back up and recover the resources that you are importing from SnapCenter. You no longer need to manage those resources in SnapCenter Server.

Steps

1. After you import the SnapCenter resources, on the Inventory page that appears, select the SnapCenter resources that you imported that you want to have NetApp Backup and Recovery manage from now on.
2. Select the Actions icon > **Manage** to manage the resources.
3. Select **Manage in NetApp Console**.

The Inventory page shows **Managed** under the host name to indicate that the selected host resources are now managed by NetApp Backup and Recovery.

Edit imported SnapCenter resources

You can later re-import SnapCenter resources or edit the imported SnapCenter resources to update the registration details.

You can change only the port and password details for the SnapCenter Server.

Steps

1. From the NetApp Backup and Recovery menu, select **Settings**.
2. Select the down arrow for **Import from SnapCenter**.

The Import from SnapCenter page shows all previous imports.

3. Select the Actions icon > **Edit** to update the resources.
4. Update the SnapCenter password and port details, as needed.
5. Select **Import**.

Add a KVM management platform

If you use the Apache CloudStack management platform to manage KVM resources, you need to integrate it with NetApp Backup and Recovery so that Backup and Recovery can discover and protect the managed KVM

hosts and VMs.

Steps

1. From the NetApp Backup and Recovery menu, select **Settings**.
2. Select the down arrow to expand the **Management platform** section.
3. Select **Add management platform credential**.
4. Enter the following information:
 - **Management platform IP address or FQDN**: Enter the IP address or fully qualified domain name of the management platform.
 - **API key**: Enter the API key to use to authenticate API requests.
 - **Secret Key**: Enter the secret key to use to authenticate API requests.
 - **Port**: Enter the port to use for communication between Backup and Recovery and the management platform.
 - **Agents**: Select a Console agent to use to facilitate communication between Backup and Recovery and the management platform.
5. When finished, select **Add**.

Configure log directories in snapshots for Windows hosts

Before you create policies for Windows hosts, you should configure log directories in snapshots for Windows hosts. Log directories are used to store the logs that are generated during the backup process.

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. From the Inventory page, select a workload and then select the Actions icon **...** > **View details** to display the workload details.
3. From the Inventory details page showing Microsoft SQL Server, select the Hosts tab.
4. From the Inventory details page, select a host and select the Actions icon **...** > **Configure log directory**.
5. Either browse or enter the path for the log directory.
6. Select **Save**.

Create an execution hook template

You can create a custom execution hook template that you can use to perform actions before or after a data protection operation on an application.



Templates that you create here are only usable when protecting Kubernetes workloads.

Steps

1. In the Console, go to **Protection > Backup and recovery**.
2. Select the **Settings** tab.
3. Expand the **Execution hook template** section.
4. Select **Create execution hook template**.
5. Enter a name for the execution hook.

6. Optionally, choose a type of hook. For example, a post-restore hook is run after the restore operation is complete.
7. In the **Script** text box, enter the executable shell script that you want to run as part of the execution hook template. Optionally, you can select **Upload script** to upload a script file instead.
8. Select **Create**.

After you create the template, it appears in the list of templates in the **Execution hook template** section.

Set up role-based access control in NetApp Backup and Recovery

To increase security and control resource access, configure role-based access for NetApp Backup and Recovery. The NetApp Console supports role-based access control (RBAC) for some Backup and Recovery workloads. You can assign administrative or viewer roles specific to these workloads. Other workloads that do not yet support role-based access control remain accessible to all users with Backup and Recovery roles until project-level association is supported.

Follow these steps to control access to resources in your organization. Make changes in the **Administration > Identity and access** page in the NetApp Console menu.



These steps assume that you are assigned the Organization Admin role in the Console.

Steps

1. Create the identity and access project structure.

As an Organization admin, set up the Identity and access folder and project structure where workloads will reside.

2. Assign user roles.

- a. Primary option:

Add users to each project designated for workloads and grant them the appropriate role. For example:

- **Organization admin and Backup and Recovery super admin:** A user with these roles can see all resources in all organizations, and discover Backup and Recovery workloads and assign them to projects (for example, US East or US West).
- **Folder or project admin and Backup and Recovery super admin:** A user with these roles can see only the resources in the folder or project they have permissions for, but can discover Backup and Recovery workloads and assign them to that project.

- b. Alternative option:

Instead of granting a user full Backup and Recovery admin access, you can assign yourself the Backup and Recovery super admin role and discover the workloads directly.

3. Discover workloads in Backup and Recovery.

Organization admins or Folder or project admins discover the workloads that are available and select the appropriate project (such as US East or US West). Each workload is automatically associated with the

selected project.

4. Add users to projects.

Organization admins or Folder/project admins add Console users to projects with workloads. Assign users the Organization viewer role and a Backup and Recovery role based on their access needs. Users with the right Backup and Recovery role will automatically gain access to new workloads in these projects.

Related information

- [Learn about NetApp Console identity and access management.](#)
- [NetApp Backup and Recovery roles in NetApp Console.](#)

Use NetApp Backup and Recovery

View protection health on the NetApp Backup and Recovery Dashboard

Monitoring the health of your workloads ensures that you are aware of issues with workload protection and can take steps to resolve them. View the status of your backups and restores on the NetApp Backup and Recovery Dashboard. You can review the system summary, Protection summary, Job summary, Restore summary, and more.

Required NetApp Console role

Storage viewer, Backup and Recovery super admin, Backup and Recovery backup admin, Backup and Recovery restore admin, Backup and Recovery clone admin, or Backup and Recovery viewer role. Learn about [Backup and recovery roles and privileges](#). Learn about [NetApp Console access roles for all services](#).

Steps

1. From the NetApp Console menu, select **Protection > Backup and recovery**.
2. Select a workload tile (for example, Microsoft SQL Server).
3. From the Backup and Recovery menu, select **Dashboard**.

You can review the following types of information:

- Number of hosts or VMs discovered
- Number of Kubernetes clusters discovered
- Number of backup targets on object storage
- Number of vCenters
- Number of storage clusters in ONTAP

View the Protection summary

Review the following information in the Protection summary:

- The total number of protected and unprotected databases, VMs, and datastores.



A protected database is one that has a backup policy assigned. An unprotected database is one that doesn't have a backup policy assigned to it.

- The number of backups that were successful, have a warning, or have failed.
- The total capacity discovered by the backup service and the capacity that is protected versus unprotected. Hover over the "i" icon to see the details.

View the Job summary

Review the total jobs completed, running or failed in the Job summary.

Steps

1. For each job distribution, change a filter to show the summary of failed, running and complete based on the

number of days, for example, the last 30 days, last 7 days, last 24 hours, or last 1 year.

2. View details of the failed, running and complete jobs by selecting **View job monitoring**.

View the Restore summary

Review the following information on the Restore summary:

- The total number of restore jobs performed
- The total amount of capacity that has been restored
- The number of restore jobs performed on local, secondary, and object storage. Hover over the chart to see the details.

Create and manage policies to govern backups in NetApp Backup and Recovery

In NetApp Backup and Recovery, create your own policies that govern backup frequency, the time the backup is taken, and the number of backup files that are retained.



Some of these options and configuration sections are not available for all workloads.

If you import resources from SnapCenter, you might encounter some differences with policies used in SnapCenter and those used in NetApp Backup and Recovery. See [Policy differences between SnapCenter and NetApp Backup and Recovery](#).

You can accomplish the following goals related to policies:

- Create a local snapshot policy
- Create a policy for replication to secondary storage
- Create a policy for object storage settings
- Configure advanced policy settings
- Edit policies (not available for VMware preview workloads)
- Delete policies

View policies

1. From the NetApp Backup and Recovery menu, select **Policies**.
2. Review these policy details.
 - **Workload:** Examples include Microsoft SQL Server, Volumes, VMware, KVM, Hyper-V, Oracle Database, or Kubernetes.
 - **Backup type:** Examples include full backup and log backup.
 - **Architecture:** Examples include local snapshot, fan-out, cascading, disk to disk, and disk to object store.
 - **Resources protected:** Shows how many resources out of the total resources on that workload are protected.
 - **Ransomware protection:** Shows if the policy includes snapshot locking on the local snapshot, snapshot locking on secondary storage, or DataLock locking on object storage.

Create a policy

You can create policies that govern your local snapshots, replications to secondary storage, and backups to object storage. Part of your 3-2-1 strategy involves creating a snapshot of the instances, databases, applications, or VMs on the **primary** storage system.

Required NetApp Console role

Storage viewer, Backup and Recovery super admin, Backup and Recovery backup admin. Learn about [Backup and recovery roles and privileges](#). Learn about [NetApp Console access roles for all services](#).

Before you begin

If you plan on replicating to secondary storage and want to use snapshot locking on local snapshots or on remote ONTAP secondary storage, you first need to initialize the ONTAP compliance clock on the cluster level. This is a requirement for enabling snapshot locking in the policy.

For instructions on how to do this, refer to [Initialize the compliance clock in ONTAP](#).

For information about snapshot locking in general, refer to [Snapshot locking in ONTAP](#).

Steps

1. From the NetApp Backup and Recovery menu, select **Policies**.
2. From the Policies page, select **Create new policy**.
3. In the Policies page, provide the following information.
 - **Details** section:
 - Workload type: Select the workload that will use the policy.
 - Enter a policy name.



For a list of characters to avoid, see the hover tip.

- **Select a Console agent from the Agent list.**
- **Backup architecture** section: Select the down arrow and choose the data flow for the backup, such as 3-2-1 fan-out, 3-2-1 cascade, or disk to disk.
 - **3-2-1 fanout:** Primary storage (disk) to secondary storage (disk) to cloud (object store). Creates multiple copies of data across different storage systems, such as ONTAP to ONTAP and ONTAP to object-store configurations. This can be a cloud hyperscaler object store or a private object store. These configurations help in achieving optimal data protection and disaster recovery.



This option is not available for Amazon FSx for NetApp ONTAP.

For VMware workloads, this configures the local snapshot on the datastores or VMs on the primary and replicates from primary disk storage to secondary disk storage as well as replicates from primary to cloud object storage.

- **3-2-1 cascade:** (Not available for Kubernetes workloads) Primary storage (disk) to secondary storage (disk) and primary storage (disk) to cloud storage (object store). This can be a cloud hyperscaler object store or a private object store — StorageGRID. This creates a chain of data replication across multiple systems to ensure redundancy and reliability.



This option is not available for Amazon FSx for NetApp ONTAP.

For VMware workloads, this configures the local snapshot on the datastores or VMs on the primary storage and a cascade from primary disk storage to secondary disk storage and then to cloud object storage.

- **Disk to disk:** (Not available for Kubernetes workloads) Primary storage (disk) to secondary storage (disk). The ONTAP to ONTAP data protection strategy replicates data between two ONTAP systems to ensure high availability and disaster recovery. This is typically achieved using SnapMirror, which supports both synchronous and asynchronous replication. This method ensures that your data is continuously updated and available across multiple locations, providing robust protection against data loss.

For VMware workloads, this configures the local snapshot on the datastores or VMwares on the primary storage system and then replicates the data from the primary disk storage system to the secondary disk storage system.

- **Disk-to-object store:** Primary storage (disk) to cloud (object store). This replicates data from an ONTAP system to an object storage system, such as AWS S3, Azure Blob Storage or StorageGRID. This is typically achieved using SnapMirror Cloud, which provides incremental forever backups by transferring only changed data blocks after the initial baseline transfer. This can be a cloud hyperscaler object store or a private object store — StorageGRID. This method is ideal for long-term data retention and archiving, offering a cost-effective and scalable solution for data protection.

For VMWare workloads, this configures the local snapshot on the datastores or VMs on the primary and replication from primary disk storage to cloud object storage.

- **Disk-to-disk fanout:** (Not available for Kubernetes workloads) Primary storage (disk) to secondary storage (disk) and primary storage (disk) to secondary storage (disk).



You can configure multiple secondary settings for the disk-to-disk fanout option.

For VMware workloads, this configures the primary disk storage to secondary disk storage and replicates primary disk storage to secondary disk storage.

- **Local snapshots:** Local snapshot on the selected volume (Microsoft SQL Server). Local snapshots are a key component of data protection strategies, capturing the state of your data at specific points in time. This creates read-only, point-in-time copies of production volumes where your workloads are running. The snapshot consumes minimal storage space and incurs negligible performance overhead because it records only changes to files since the last snapshot. You can use local snapshots to recover from data loss or corruption, as well as to create backups for disaster recovery purposes.

For VMware workloads, this configures the local snapshot on the datastores or VMs on the primary storage system.

Create a local snapshot policy

Provide information for the local snapshot.

- Select the **Add schedule** option to select the snapshot schedule or schedules. You can have a maximum of 5 schedules.
- **Snapshot frequency:** Select the frequency of hourly, daily, weekly, monthly, or yearly. The yearly frequency is not available for Kubernetes workloads.

- **Snapshot retention:** Enter the number of snapshots to keep.
- **Enable log backup:** (Applies to Microsoft SQL Server workloads and Oracle Database workloads only.) Enable this option to back up logs and set the frequency and retention of the log backups. To do this, you must have already configured a log backup. See [Configure log directories](#).
 - **Prune archive logs after backup:** (Oracle Database workloads only) If log backups are enabled, you can optionally enable this feature to limit how long Backup and Recovery keeps Oracle archive logs. You can choose the retention period as well as where Backup and Recovery should delete the archive logs.
- **Provider:** (Kubernetes workloads only) Select the storage provider that hosts the Kubernetes application resources.

Create a policy for secondary settings (replication to secondary storage)

Provide information for the replication to secondary storage. Schedule information from the local snapshot settings appears for you in the secondary settings. These settings are not available for Kubernetes workloads.

- **Backup:** Select the frequency of hourly, daily, weekly, monthly, or yearly.
- **Backup target:** Select the target system on secondary storage for the backup.
- **Retention:** Enter the number of snapshots to keep.
- **Enable snapshot locking:** Select whether you want to enable tamper-proof snapshots.
- **Snapshot locking period:** Enter the number of days, months, or years that you want to lock the snapshot.
- **Transfer to secondary:**
 - The **ONTAP transfer schedule - Inline** option is selected by default and that indicates that snapshots are transferred to the secondary storage system immediately. You don't need to schedule the backup.
 - Other options: If you choose a deferred transfer, the transfers are not immediate and you can set a schedule.
- **SnapMirror and SnapVault SMAS secondary relationship:** Use SnapMirror and SnapVault SMAS secondary relationships for SQL Server workloads.

Create a policy for object storage settings

Provide information for the backup to object storage. These settings are called "Backup settings" for Kubernetes workloads.



The fields that appear differ depending on the provider and architecture selected.

Create a policy for AWS object storage

Enter information in these fields:

- **Provider:** Select **AWS**.
- **AWS account:** Select the AWS account.
- **Backup target:** Select a registered S3 object storage target. Ensure that the target is accessible within your backup environment.
- **IPspace:** Select the IPspace to use for the backup operations. This is useful if you have multiple IPspaces and want to control which one is used for backups.
- **Schedule settings:** Select the schedule that was set for the local snapshots. You can remove a schedule,

but you cannot add one because the schedules are set according to the local snapshot schedules.

- **Retention copies:** Enter the number of snapshots to keep.
- **Run at:** Choose the ONTAP transfer schedule to back up data to object storage.
- **Tier your backups from object store to archival storage:** If you choose to tier backups to archive storage (for example, AWS Glacier), select the tier option and the number of days to archive.
- **Enable integrity scan:** (Not available for Kubernetes workloads) Select whether you want to enable integrity scans (snapshot locking) on the object storage. This ensures that the backups are valid and can be restored successfully. The integrity scan frequency is set to 7 days by default. To protect your backups from being modified or deleted, select the **Integrity scan** option. The scan occurs only on the latest snapshot. You can enable or disable integrity scans on the latest snapshot.

Create a policy for Microsoft Azure object storage

Enter information in these fields:

- **Provider:** Select **Azure**.
- **Azure subscription:** Select the Azure subscription from those discovered.
- **Azure resource group:** Select the Azure resource group from those discovered.
- **Backup target:** Select a registered object storage target. Ensure that the target is accessible within your backup environment.
- **IPspace:** Select the IPspace to use for the backup operations. This is useful if you have multiple IPspaces and want to control which one is used for backups.
- **Schedule settings:** Select the schedule that was set for the local snapshots. You can remove a schedule, but you cannot add one because the schedules are set according to the local snapshot schedules.
- **Retention copies:** Enter the number of snapshots to keep.
- **Run at:** Choose the ONTAP transfer schedule to back up data to object storage.
- **Tier your backups from object store to archival storage:** If you choose to tier backups to archive storage, select the tier option and the number of days to archive.
- **Enable integrity scan:** (Not available for Kubernetes workloads) Select whether you want to enable integrity scans (snapshot locking) on the object storage. This ensures that the backups are valid and can be restored successfully. The integrity scan frequency is set to 7 days by default. To protect your backups from being modified or deleted, select the **Integrity scan** option. The scan occurs only on the latest snapshot. You can enable or disable integrity scans on the latest snapshot.

Create a policy for StorageGRID object storage

Enter information in these fields:

- **Provider:** Select **StorageGRID**.
- **StorageGRID credentials:** Select the StorageGRID credentials from those discovered. These credentials are used to access the StorageGRID object storage system and were entered in the Settings option.
- **Backup target:** Select a registered S3 object storage target. Ensure that the target is accessible within your backup environment.
- **IPspace:** Select the IPspace to use for the backup operations. This is useful if you have multiple IPspaces and want to control which one is used for backups.
- **Schedule settings:** Select the schedule that was set for the local snapshots. You can remove a schedule, but you cannot add one because the schedules are set according to the local snapshot schedules.

- **Retention copies:** Enter the number of snapshots to keep for each frequency.
- **Transfer schedule for object storage:** (Not available for Kubernetes workloads) Choose the ONTAP transfer schedule to back up data to object storage.
- **Enable integrity scan:** (Not available for Kubernetes workloads) Select whether you want to enable integrity scans (snapshot locking) on the object storage. This ensures that the backups are valid and can be restored successfully. The integrity scan frequency is set to 7 days by default. To protect your backups from being modified or deleted, select the **Integrity scan** option. The scan occurs only on the latest snapshot. You can enable or disable integrity scans on the latest snapshot.
- **Tier your backups from object store to archival storage:** (Not available for Kubernetes workloads) If you choose to tier backups to archive storage, select the tier option and the number of days to archive.

Configure advanced settings in the policy

Optionally, you can configure advanced settings in the policy. These settings are available for all backup architectures, including local snapshots, replication to secondary storage, and backups to object storage. These settings are not available for Kubernetes workloads. The available advanced settings will differ depending on the workload you selected at the top of the page, so the advanced settings described here might not apply to all workloads. Advanced settings are not available when configuring a policy for Kubernetes workloads.

Steps

1. From the NetApp Backup and Recovery menu, select **Policies**.
2. From the Policies page, select **Create new policy**.
3. In the **Policy > Advanced** settings section, select the **Select advance action** menu to choose from a list of advanced settings.
4. Enable any of the settings you want to view or change, and then select **Accept**.
5. Provide the following information:
 - **Copy only backup:** (Applies to Microsoft SQL Server workloads only) Choose copy-only backup (a type of Microsoft SQL Server backup) if you need to back up your resources by using another backup application.
 - **Availability group settings:** (Applies to Microsoft SQL Server workloads only) Select preferred backup replicas or specify a particular replica. This setting is useful if you have a SQL Server availability group and want to control which replica is used for backups.
 - **Maximum transfer rate:** To not set a limit on bandwidth usage, select **Unlimited**. If you want to limit the transfer rate, select **Limited** and select the network bandwidth between 1 and 1,000 Mbps allocated to upload backups to object storage. By default, ONTAP can use an unlimited amount of bandwidth to transfer the backup data from volumes in the system to object storage. If you notice backup traffic is affecting normal user workloads, consider decreasing the amount of network bandwidth that is used during the transfer.
 - **Backup retries:** (Not applicable to VMware workloads) To retry the job in case of a failure or interruption, select **Enable job retries during failure**. Enter the maximum number of snapshot and backup job retries and the retry time interval. The recount must be less than 10. This setting is useful if you want to ensure that the backup job is retried in case of a failure or interruption.



If the snapshot frequency is set to 1 hour, the maximum delay along with the retry count shouldn't exceed 45 minutes.

- **Enable VM-consistent snapshot:** Select whether you want to enable VM-consistent snapshots. This ensures that the newly created snapshots are consistent with the state of the virtual machine at the

time of the snapshot. This is useful for ensuring that the backups can be restored successfully and that the data is in a consistent state. This does not apply to existing snapshots.

- **Ransomware scan:** Select whether you want to enable ransomware scanning on each bucket. This requires DataLock locking on object storage. Enter the frequency of the scan in days. This option applies to AWS and Microsoft Azure object storage. Note that this option might incur additional charges, depending on the cloud provider.
- **Backup verification:** (Not applicable to VMware workloads) Select whether you want to enable backup verification and whether you want it immediately or later. This feature ensures that the backups are valid and can be restored successfully. We recommend that you enable this option to ensure the integrity of your backups. By default, backup verification runs from secondary storage if secondary storage is configured. If secondary storage isn't configured, backup verification runs from primary storage.

Additionally, configure the following options:

- **Daily, Weekly, Monthly, or Yearly verification:** If you chose **Later** as the backup verification, select the frequency of backup verification. This ensures that backups are regularly checked for integrity and can be restored successfully.
- **Backup labels:** Enter a label for the backup. This label is used to identify the backup in the system and can be useful for tracking and managing backups.
- **Database consistency check:** (Not applicable to VMware workloads) Select whether you want to enable database consistency checks. This option ensures that the databases are in a consistent state before the backup is taken, which is crucial for ensuring data integrity.
- **Verify log backups:** (Not applicable to VMware workloads) Select whether you want to verify log backups. Select the verification server. If you chose disk-to-disk or 3-2-1, also select the verification storage location. This option ensures that the log backups are valid and can be restored successfully, which is important for maintaining the integrity of your databases.
- **Networking:** Select the network interface to use for the backup operations. This is useful if you have multiple network interfaces and want to control which one is used for backups.
 - **IPspace:** Select the IPspace to use for the backup operations. This is useful if you have multiple IPspaces and want to control which one is used for backups.
 - **Private endpoint configuration:** If you are using a private endpoint for your object storage, select the private endpoint configuration to use for the backup operations. This is useful if you want to ensure that the backups are transferred securely over a private network connection.
- **Notification:** Select whether you want to enable email notifications for backup operations. This is useful if you want to be notified when a backup operation starts, completes, or fails.
- **Independent disks:** (Applies to VMware workloads only) Check this to include in the backup any datastores with independent disks that contain temporary data. An independent disk is a VM disk that is not included in VMware snapshots.
- **SnapMirror volume and snapshot format:** Optionally, enter your own snapshot name in a policy that governs the backups for Microsoft SQL Server workloads. Enter the format and custom text. If you chose to backup to secondary storage, you can also add a SnapMirror volume prefix and suffix.

Edit a policy

You can edit backup architecture, backup frequency, retention policy, and other settings for a policy.

You can add another protection level when you edit a policy, but you cannot remove a protection level. For example, if the policy is only protecting local snapshots, you can add replication to secondary storage or

backups to object storage. If you have local snapshots and replication, you can add object storage. However, if you have local snapshots, replication, and object storage, you cannot remove one of these levels.

If you are editing a policy that backs up to object storage, you can enable archival.

If you imported resources from SnapCenter, you might encounter some differences policies used in SnapCenter and those used in NetApp Backup and Recovery. See [Policy differences between SnapCenter and NetApp Backup and Recovery](#).

Required NetApp Console role

Backup and Recovery super admin. [Learn about NetApp Console access roles for all services](#).

Steps

1. In the NetApp Console, go to **Protection > Backup and Recovery**.
2. Select the **Policies** option.
3. Select the policy that you want to edit.
4. Select the **Actions** **...** icon, and select **Edit**.

Delete a policy

You can delete a policy if you no longer need it.



You cannot delete a policy that is associated with a workload.

Steps

1. In the Console, go to **Protection > Backup and Recovery**.
2. Select the **Policies** option.
3. Select the policy that you want to delete.
4. Select the **Actions** **...** icon, and select **Delete**.
5. Confirm the action, and select **Delete**.

Protect ONTAP volume workloads

Protect your ONTAP volume data using NetApp Backup and Recovery

NetApp Backup and Recovery provides backup and restore capabilities for protection and long-term archive of your ONTAP volume data. You can implement a 3-2-1 strategy where you have 3 copies of your source data on 2 different storage systems along with 1 copy in the cloud.



To switch to and from NetApp Backup and Recovery workloads, refer to [Switch to different NetApp Backup and Recovery workloads](#).

After activation, backup and recovery creates block-level, incremental forever backups that are stored on another ONTAP cluster and in object storage in the cloud. In addition to your source volume, you'll have a:

- Snapshot of the volume on the source system

- Replicated volume on a different storage system
- Backup of the volume in object storage

NetApp Backup and Recovery leverages NetApp's SnapMirror data replication technology to ensure that all the backups are fully synchronized by creating snapshots and transferring them to the backup locations.

The benefits of the 3-2-1 approach include:

- Multiple data copies protect against internal and external cybersecurity threats.
- Using different types of media helps you recover if one type fails.
- You can quickly restore from the onsite copy, and use the offsite copies if the onsite copy is compromised.

When necessary, you can restore an entire *volume*, a *folder*, or one or more *files*, from any of the backup copies to the same or different system.

Features

Replication features:

- Replicate data between ONTAP storage systems to support backup and disaster recovery.
- Ensure the reliability of your DR environment with high availability.
- Native ONTAP in-flight encryption set up via Pre-Shared Key (PSK) between the two systems.
- Copied data is immutable until you make it writable and ready to use.
- Replication is self-healing in the event of a transfer failure.
- When compared to [NetApp Replication](#), the replication in NetApp Backup and Recovery includes the following features:
 - Replicate multiple FlexVol volumes at a time to a secondary system.
 - Restore a replicated volume to the source system or to a different system using the UI.

See [Replication limitations for ONTAP volumes](#) for a list of replication features that are unavailable with NetApp Backup and Recovery for ONTAP volumes.

Backup-to-object features:

- Back up independent copies of your data volumes to low-cost object storage.
- Apply a single backup policy to all volumes in a cluster, or assign different backup policies to volumes that have unique recovery point objectives.
- Create a backup policy to be applied to all future volumes created in the cluster.
- Make immutable backup files so they are locked and protected for the retention period.
- Scan backup files for possible ransomware attack - and remove/replace infected backups automatically.
- Tier older backup files to archival storage to save costs.
- Delete the backup relationship so you can archive unneeded source volumes while retaining volume backups.
- Back up from cloud to cloud, and from on-premises systems to public or private cloud.

- Backup data is secured with AES-256 bit encryption at-rest and TLS 1.2 HTTPS connections in-flight.
- Use your own customer-managed keys for data encryption instead of using the default encryption keys from your cloud provider.
- Support for up to 4,000 backups of a single volume.

Restore features:

- Restore data from a specific point in time from local snapshots, replicated volumes, or backed up volumes in object storage.
- Restore a volume, a folder, or individual files, to the source system or to a different system.
- Restore data to a system using a different subscription/account or that is in a different region.
- Perform a *quick restore* of a volume from cloud storage to a Cloud Volumes ONTAP system or to an on-premises system; perfect for disaster recovery situations where you need to provide access to a volume as soon as possible.
- Restore data on a block level, placing the data directly in the location you specify, all while preserving the original ACLs.
- Browse and search file catalogs for easy selection of individual folders and files for single file restore.

Supported systems for backup and restore operations

NetApp Backup and Recovery supports ONTAP systems and public and private cloud providers.

Supported regions

NetApp Backup and Recovery is supported with Cloud Volumes ONTAP in many Amazon Web Services, Microsoft Azure, and Google Cloud regions.

[Learn more using the Global Regions Map](#)

Supported backup destinations

NetApp Backup and Recovery enables you to back up ONTAP volumes from the following source systems to the following secondary systems and object storage in public and private cloud providers. snapshots reside on the source system.

Source system	Secondary system (Replication)	Destination Object Store (Backup)
Cloud Volumes ONTAP in AWS	Cloud Volumes ONTAP in AWS On-premises ONTAP system	Amazon S3
Cloud Volumes ONTAP in Azure	Cloud Volumes ONTAP in Azure On-premises ONTAP system	Azure Blob
Cloud Volumes ONTAP in Google	Cloud Volumes ONTAP in Google On-premises ONTAP system	Google Cloud Storage

Source system	Secondary system (Replication)	Destination Object Store (Backup)
On-premises ONTAP system	Cloud Volumes ONTAP On-premises ONTAP system	Amazon S3 Azure Blob Google Cloud Storage NetApp StorageGRID ONTAP S3

Supported restore destinations

You can restore ONTAP data from a backup file that resides in a secondary system (a replicated volume) or in object storage (a backup file) to the following systems. Snapshots reside on the source system and can be restored only to that same system.

Backup File Location	Destination system
Object Store (Backup)	Secondary System (Replication)
Amazon S3	Cloud Volumes ONTAP in AWS On-premises ONTAP system
Azure Blob	Cloud Volumes ONTAP in Azure On-premises ONTAP system
Google Cloud Storage	Cloud Volumes ONTAP in Google On-premises ONTAP system
NetApp StorageGRID	On-premises ONTAP system Cloud Volumes ONTAP
ONTAP S3	On-premises ONTAP system Cloud Volumes ONTAP

Note that references to "on-premises ONTAP systems" includes FAS, AFF, and ONTAP Select systems.

Supported volumes

NetApp Backup and Recovery supports the following types of volumes:

- FlexVol read-write volumes
- FlexGroup volumes (requires ONTAP 9.12.1 or later)
- SnapLock Enterprise volumes (requires ONTAP 9.11.1 or later)
- SnapLock Compliance for on-premises volumes (requires ONTAP 9.14 or later)
- SnapMirror data protection (DP) destination volumes



NetApp Backup and Recovery does not support backups of FlexCache volumes.

See the sections on [Backup and restore limitations for ONTAP volumes](#) for additional requirements and limitations.

Cost

There are two types of costs associated with using NetApp Backup and Recovery with ONTAP systems: resource charges and service charges. Both of these charges are for the backup to object portion of the service.

There is no charge to create snapshots or replicated volumes - other than the disk space required to store the snapshots and replicated volumes.

Resource charges

Resource charges are paid to the cloud provider for object storage capacity and for writing and reading backup files to the cloud.

- For Backup to object storage, you pay your cloud provider for object storage costs.

Since NetApp Backup and Recovery preserves the storage efficiencies of the source volume, you pay the cloud provider object storage costs for the data *after* ONTAP efficiencies (for the smaller amount of data after deduplication and compression have been applied).

- For restoring data using Search & Restore, certain resources are provisioned by your cloud provider, and there is per-TiB cost associated with the amount of data that is scanned by your search requests. (These resources are not needed for Browse & Restore.)
 - In AWS, [Amazon Athena](#) and [AWS Glue](#) resources are deployed in a new S3 bucket.
 - In Azure, an [Azure Synapse workspace](#) and [Azure Data Lake Storage](#) are provisioned in your storage account to store and analyze your data.
 - In Google, a new bucket is deployed, and the [Google Cloud BigQuery services](#) are provisioned on an account/project level.
- If you plan to restore volume data from a backup file that has been moved to archival object storage, then there's an additional per-GiB retrieval fee and per-request fee from the cloud provider.
- If you plan to scan a backup file for ransomware during the process of restoring volume data (if you have enabled DataLock and Ransomware Resilience for your cloud backups), then you'll incur extra egress costs from your cloud provider as well.

Service charges

Service charges are paid to NetApp and cover both the cost to *create* backups to object storage and to *restore* volumes, or files, from those backups. You pay only for the data that you protect in object storage, calculated by the source logical used capacity (*before* ONTAP efficiencies) of ONTAP volumes which are backed up to object storage. This capacity is also known as Front-End Terabytes (FETB).

There are three ways to pay for the Backup service. The first option is to subscribe from your cloud provider, which enables you to pay per month. The second option is to get an annual contract. The third option is to purchase licenses directly from NetApp.

Licensing

NetApp Backup and Recovery is available with the following consumption models:

- **BYOL:** A license purchased from NetApp that can be used with any cloud provider.
- **PAYGO:** An hourly subscription from your cloud provider's marketplace.
- **Annual:** An annual contract from your cloud provider's marketplace.

A Backup license is required only for backup and restore from object storage. Creating snapshots and replicated volumes do not require a license.

Bring your own license

BYOL is term-based (1, 2, or 3 years) *and* capacity-based in 1 TiB increments. You pay NetApp to use the service for a period of time, say 1 year, and for a maximum amount capacity, say 10 TiB.

You'll receive a serial number that you enter in the NetApp Console to enable the service. When either limit is reached, you'll need to renew the license. The Backup BYOL license applies to all source systems associated with your NetApp Console organization or account.

[Learn how to manage your BYOL licenses.](#)

Pay-as-you-go subscription

NetApp Backup and Recovery offers consumption-based licensing in a pay-as-you-go model. After subscribing through your cloud provider's marketplace, you pay per GiB for data that's backed up — there's no up-front payment. You are billed by your cloud provider through your monthly bill.

[Learn how to set up a pay-as-you-go subscription.](#)

Note that a 30-day free trial is available when you initially sign up with a PAYGO subscription.

Annual contract

When you use AWS, two annual contracts are available for 1, 2, or 3 year terms:

- A "Cloud Backup" plan that enables you to back up Cloud Volumes ONTAP data and on-premises ONTAP data.
- A "CVO Professional" plan that enables you to bundle Cloud Volumes ONTAP and NetApp Backup and Recovery. This includes unlimited backups for Cloud Volumes ONTAP volumes charged against this license (backup capacity is not counted against the license).

When you use Azure, two annual contracts are available for 1, 2, or 3 year terms:

- A "Cloud Backup" plan that enables you to back up Cloud Volumes ONTAP data and on-premises ONTAP data.
- A "CVO Professional" plan that enables you to bundle Cloud Volumes ONTAP and NetApp Backup and Recovery. This includes unlimited backups for Cloud Volumes ONTAP volumes charged against this license (backup capacity is not counted against the license).

When you use GCP, you can request a private offer from NetApp, and then select the plan when you subscribe from the Google Cloud Marketplace during NetApp Backup and Recovery activation.

[Learn how to set up annual contracts.](#)

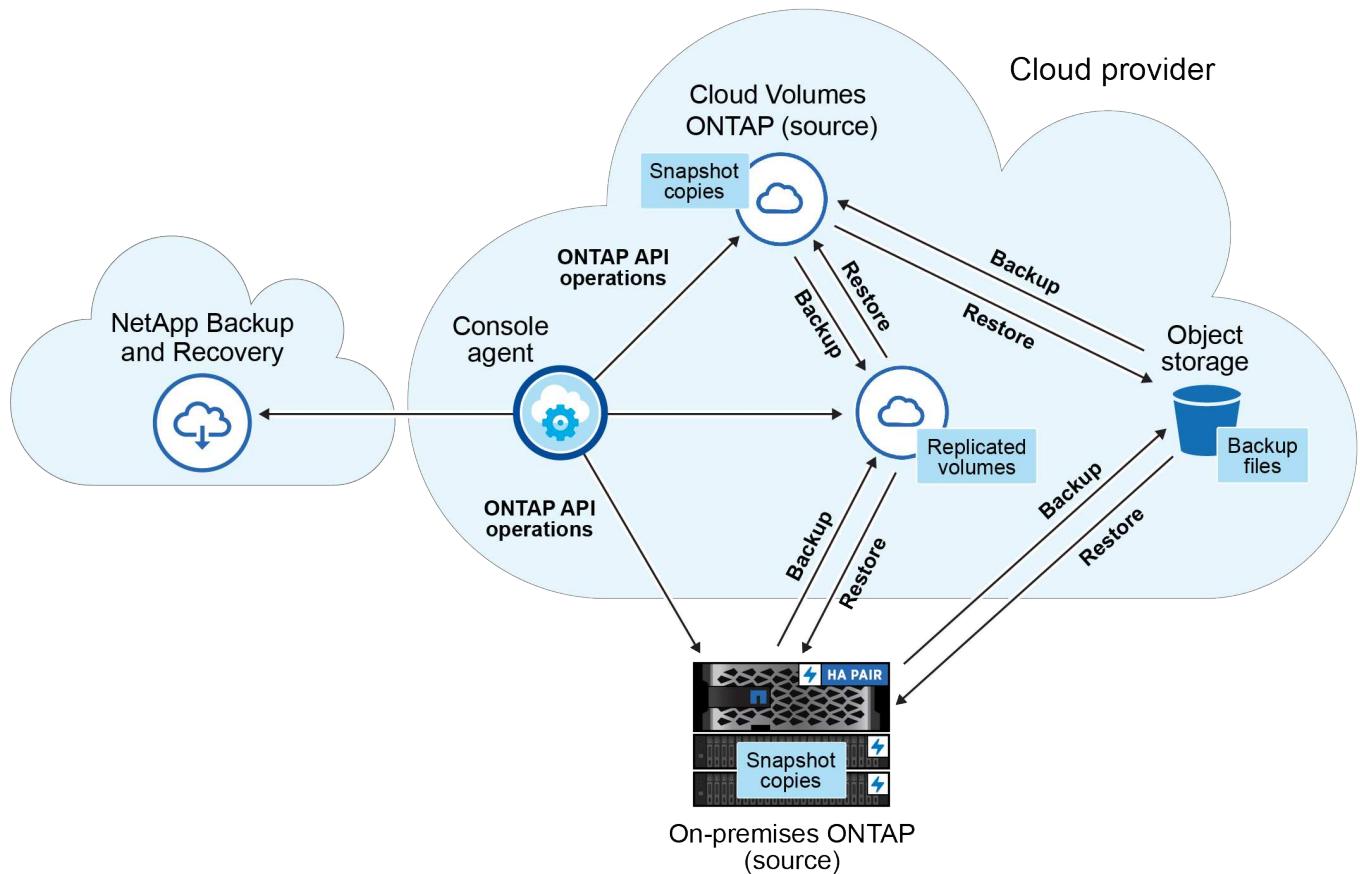
How NetApp Backup and Recovery works

When you enable NetApp Backup and Recovery on a Cloud Volumes ONTAP or on-premises ONTAP system, the service performs a full backup of your data. After the initial backup, all additional backups are incremental, which means that only changed blocks and new blocks are backed up. This keeps network traffic to a minimum. Backup to object storage is built on top of the [NetApp SnapMirror Cloud technology](#).



Any actions taken directly from your cloud provider environment to manage or change cloud backup files may corrupt the files and will result in an unsupported configuration.

The following image shows the relationship between each component:



This diagram shows volumes being replicated to a Cloud Volumes ONTAP system, but volumes could be replicated to an on-premises ONTAP system as well.

Where backups reside

Backups reside in different locations based on the type of backup:

- *Snapshots* reside on the source volume in the source system.
- *Replicated volumes* reside on the secondary storage system - a Cloud Volumes ONTAP or on-premises ONTAP system.
- *Backup copies* are stored in an object store that the Console creates in your cloud account. There's one object store per cluster/system, and the Console names the object store as follows: "netapp-backup-clusteruuid". Be sure not to delete this object store.
 - In AWS, the Console enables the [Amazon S3 Block Public Access feature](#) on the S3 bucket.
 - In Azure, the Console uses a new or existing resource group with a storage account for the Blob container. The Console [blocks public access to your blob data](#) by default.
 - In GCP, the Console uses a new or existing project with a storage account for the Google Cloud Storage bucket.
 - In StorageGRID, the Console uses an existing tenant account for the S3 bucket.

- In ONTAP S3, the Console uses an existing user account for the S3 bucket.

If you want to change the destination object store for a cluster in the future, you'll need to [unregister NetApp Backup and Recovery for the system](#), and then enable NetApp Backup and Recovery using the new cloud provider information.

Customizable backup schedule and retention settings

When you enable NetApp Backup and Recovery for a system, all the volumes you initially select are backed up using the policies that you select. You can select separate policies for snapshots, replicated volumes, and backup files. If you want to assign different backup policies to certain volumes that have different recovery point objectives (RPO), you can create additional policies for that cluster and assign those policies to the other volumes after NetApp Backup and Recovery is activated.

You can choose a combination of hourly, daily, weekly, monthly, and yearly backups of all volumes. For backup to object you can also select one of the system-defined policies that provide backups and retention for 3 months, 1 year, and 7 years. Backup protection policies that you have created on the cluster using ONTAP System Manager or the ONTAP CLI will also appear as selections. This includes policies created using custom SnapMirror labels.

 The Snapshot policy applied to the volume must have one of the labels that you're using in your replication policy and backup to object policy. If matching labels are not found, no backup files will be created. For example, if you want to create "weekly" replicated volumes and backup files, you must use a Snapshot policy that creates "weekly" snapshots.

Once you reach the maximum number of backups for a category, or interval, older backups are removed so you always have the most current backups (and so obsolete backups don't continue to take up space).

 The retention period for backups of data protection volumes is the same as defined in the source SnapMirror relationship. You can change this if you want by using the API.

Backup file protection settings

If your cluster is using ONTAP 9.11.1 or greater, you can protect your backups in object storage from deletion and ransomware attacks. Each backup policy provides a section for *DataLock* and *Ransomware Resilience* that can be applied to your backup files for a specific period of time - the *retention period*.

- *DataLock* protects your backup files from being modified or deleted.
- *Ransomware protection* scans your backup files to look for evidence of a ransomware attack when a backup file is created, and when data from a backup file is being restored.

Scheduled ransomware protection scans are enabled by default. The default setting for the scan frequency is for 7 days. The scan occurs only on the latest snapshot. The scheduled scans can be disabled to reduce your costs. You can enable or disable scheduled ransomware scans on the latest snapshot by using the option on the Advanced Settings page. If you enable it, scans are performed weekly by default. You can change that schedule to days or weeks or disable it, saving costs.

The backup retention period is the same as the backup schedule retention period, plus a maximum 31-day buffer. For example, *weekly* backups with 5 copies retained will lock each backup file for 5 weeks. *Monthly* backups with 6 copies retained will lock each backup file for 6 months.

Support is currently available when your backup destination is Amazon S3, Azure Blob, or NetApp StorageGRID. Other storage provider destinations will be added in future releases.

For more details, refer to this information:

- [How DataLock and Ransomware protection work.](#)
- [How to update Ransomware protection options in the Advanced Settings page.](#)



DataLock can't be enabled if you are tiering backups to archival storage.

Archival storage for older backup files

When using certain cloud storage you can move older backup files to a less expensive storage class/access tier after a certain number of days. You can also choose to send your backup files to archival storage immediately without being written to standard cloud storage. Note that archival storage can't be used if you have enabled DataLock.

- In AWS, backups start in the *Standard* storage class and transition to the *Standard-Infrequent Access* storage class after 30 days.

If your cluster is using ONTAP 9.10.1 or greater, you can choose to tier older backups to either *S3 Glacier* or *S3 Glacier Deep Archive* storage in the NetApp Backup and Recovery UI after a certain number of days for further cost optimization. [Learn more about AWS archival storage.](#)

- In Azure, backups are associated with the *Cool* access tier.

If your cluster is using ONTAP 9.10.1 or greater, you can choose to tier older backups to *Azure Archive* storage in the NetApp Backup and Recovery UI after a certain number of days for further cost optimization. [Learn more about Azure archival storage.](#)

- In GCP, backups are associated with the *Standard* storage class.

If your 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 for further cost optimization. [Learn more about Google archival storage.](#)

- In StorageGRID, backups are associated with the *Standard* storage class.

If your on-prem cluster is using ONTAP 9.12.1 or greater, and your StorageGRID system is using 11.4 or greater, you can archive older backup files to public cloud archival storage after a certain number of days. Current support is for AWS S3 Glacier/S3 Glacier Deep Archive or Azure Archive storage tiers. [Learn more about archiving backup files from StorageGRID.](#)

See [xref:/prev-ontap-policy-object-options.html](#) for details about archiving older backup files.

FabricPool tiering policy considerations

There are certain things you need to be aware of when the volume you are backing up resides on a FabricPool aggregate and it has an assigned tiering policy other than `none`:

- The first backup of a FabricPool-tiered volume requires reading all local and all tiered data (from the object store). A backup operation does not "reheat" the cold data tiered in object storage.

This operation could cause a one-time increase in cost to read the data from your cloud provider.

- Subsequent backups are incremental and do not have this effect.

- If the tiering policy is assigned to the volume when it is initially created you will not see this issue.
- Consider the impact of backups before assigning the `all` tiering policy to volumes. Because data is tiered immediately, NetApp Backup and Recovery will read data from the cloud tier rather than from the local tier. Because concurrent backup operations share the network link to the cloud object store, performance degradation might occur if network resources become saturated. In this case, you may want to proactively configure multiple network interfaces (LIFs) to decrease this type of network saturation.

Plan your protection journey with NetApp Backup and Recovery

NetApp Backup and Recovery enables you to create up to three copies of your source volumes to protect your data. There are many options that you can select when enabling Backup and Recovery on your volumes, so you should review your choices so you're prepared.



To switch to and from NetApp Backup and Recovery workloads, refer to [Switch to different NetApp Backup and Recovery workloads](#).

We'll go over the following options:

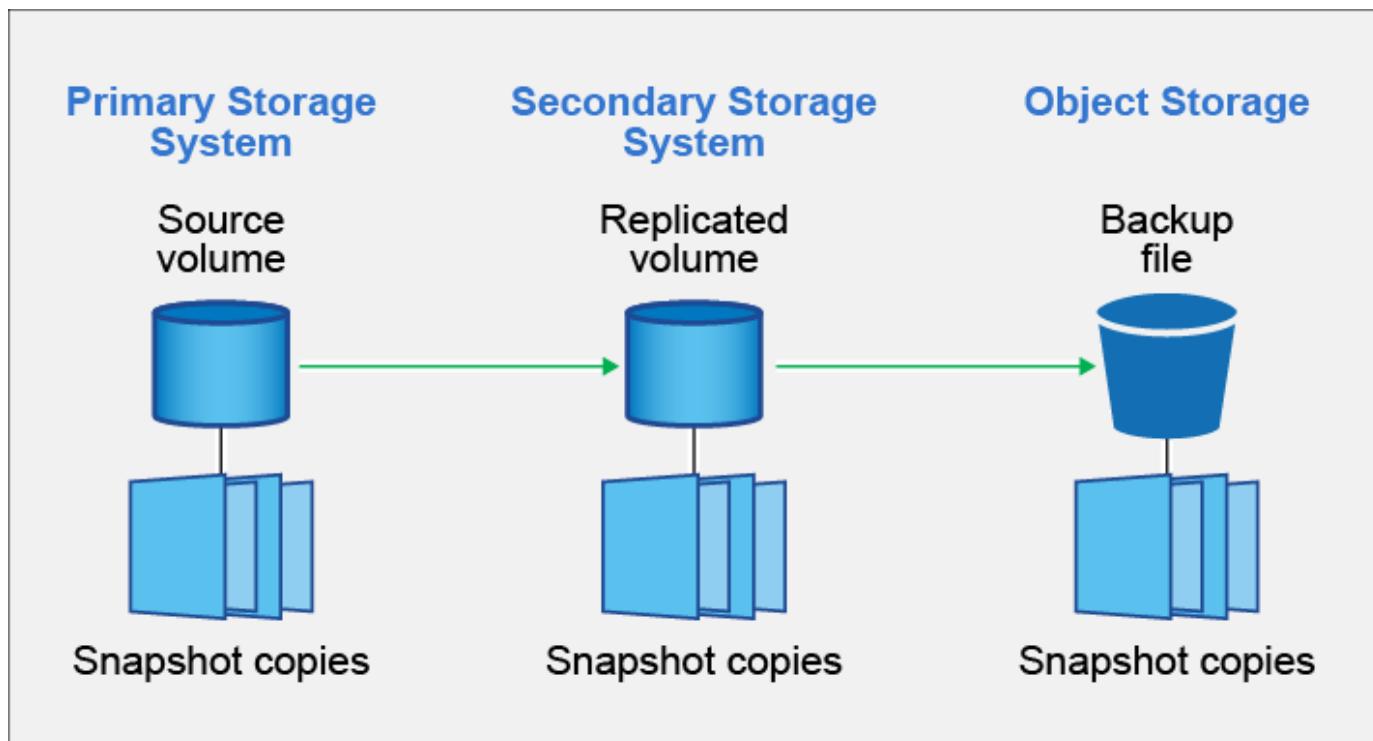
- Which protection features will you use: snapshots, replicated volumes, and/or backup to cloud
- Which backup architecture will you use: a cascade or fan-out backup of your volumes
- Will you use the default backup policies, or do you need to create custom policies
- Do you want the service to create the cloud buckets for you, or do you want to make your object storage containers before you begin
- Which Console agent deployment mode are you using (standard, restricted, or private mode)

Which protection features will you use

Before you select the features you'll use, here's a quick explanation of what each features does, and what type of protection it provides.

Backup type	Description
Snapshot	Creates a read-only, point-in-time image of a volume within the source volume as a snapshot. You can use the snapshot to recover individual files, or to restore the entire contents of a volume.
Replication	Creates a secondary copy of your data on another ONTAP storage system and continually updates the secondary data. Your data is kept current and remains available whenever you need it.
Cloud backup	Creates backups of your data to the cloud for protection and for long-term archival purposes. If necessary, you can restore a volume, folder, or individual files from the backup to the same, or different, system.

Snapshots are the basis of all the backup methods, and they are required to use the backup and recovery service. A snapshot is a read-only, point-in-time image of a volume. The image consumes minimal storage space and incurs negligible performance overhead because it records only changes to files since the last snapshot was made. The snapshot that is created on your volume is used to keep the replicated volume and backup file synchronized with changes made to the source volume - as shown in the figure.



You can choose to create both replicated volumes on another ONTAP storage system and backup files in the cloud. Or you can choose just to create replicated volumes or backup files - it's your choice.

To summarize, these are the valid protection flows you can create for volumes in your ONTAP system:

- Source volume → Snapshot → Replicated volume → Backup file
- Source volume → Snapshot → Backup file
- Source volume → Snapshot → Replicated volume



The initial creation of a replicated volume or backup file includes a full copy of the source data — this is called a *baseline transfer*. Subsequent transfers contain only differential copies of the source data (the snapshot).

Comparison of the different backup methods

The following table shows a generalized comparison of the three backup methods. While object storage space is typically less expensive than your on-premises disk storage, if you think you might restore data from the cloud frequently, then the egress fees from cloud providers can reduce some of your savings. You'll need to identify how often you need to restore data from the backup files in the cloud.

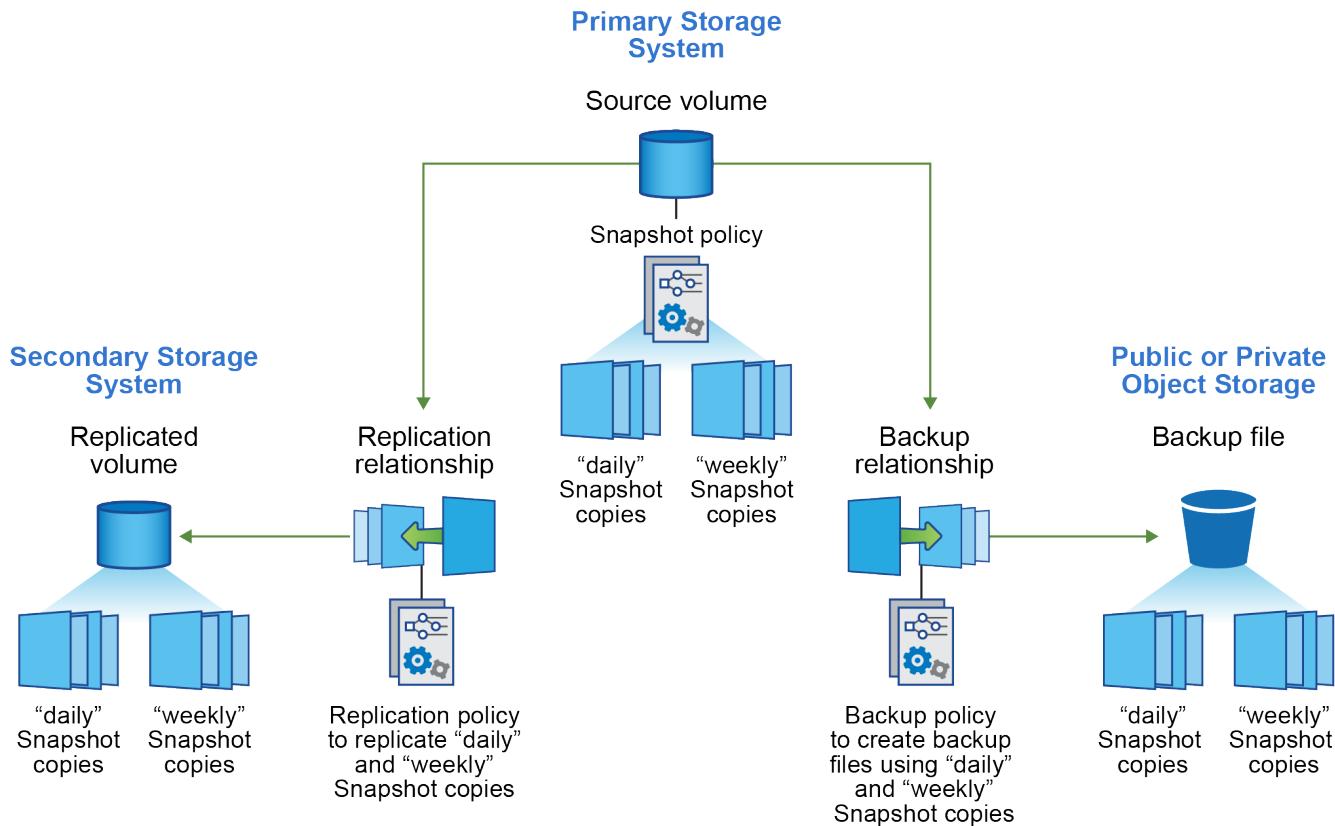
In addition to this criteria, cloud storage offers additional security options if you use the DataLock and Ransomware Resilience feature, and additional cost savings by selecting archival storage classes for older backup files. [Learn more about DataLock and Ransomware protection and archival storage settings](#).

Backup type	Backup speed	Backup cost	Restore speed	Restore cost
Snapshot	High	Low (disk space)	High	Low
Replication	Medium	Medium (disk space)	Medium	Medium (network)
Cloud backup	Low	Low (object space)	Low	High (provider fees)

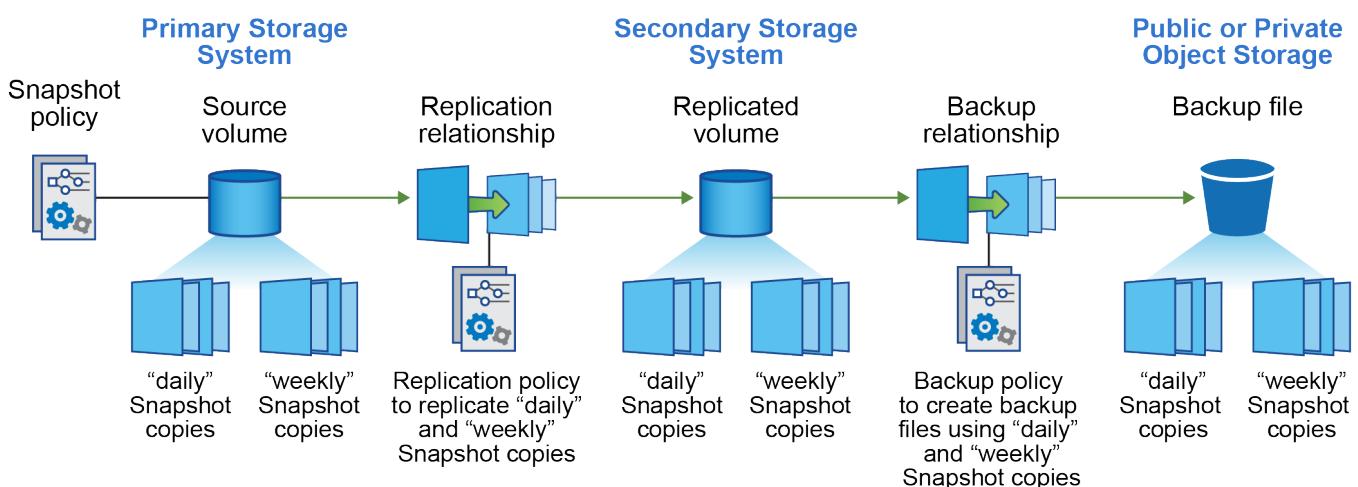
Which backup architecture will you use

When creating both replicated volumes and backup files, you can choose a fan-out or cascade architecture to back up your volumes.

A **fan-out** architecture transfers the snapshot independently to both the destination storage system and the backup object in the cloud.



A **cascade** architecture transfers the snapshot to the destination storage system first, and then that system transfers the copy to the backup object in the cloud.



Comparison of the different architecture choices

This table provides a comparison of the fan-out and cascade architectures.

Fan-out	Cascade
Small performance impact on the source system because it is sending snapshots to 2 distinct systems	Less effect on the performance of the source storage system because it sends the snapshot only once
Easier to set up because all policies, networking, and ONTAP configurations are done on the source system	Requires some networking and ONTAP configuration to be done from the secondary system as well.

Will you use the default policies for snapshots, replications, and backups

You can use the default policies provided by NetApp to create your backups, or you can create custom policies. When you use the activation wizard to enable the backup and recovery service for your volumes, you can select from the default policies and any other policies that already exist in the system (Cloud Volumes ONTAP or on-premises ONTAP system). If you want to use a policy different than those existing policies, you can create the policy before starting or while using the activation wizard.

- The default snapshot policy creates hourly, daily, and weekly snapshots, retaining 6 hourly, 2 daily, and 2 weekly snapshots.
- The default replication policy replicates daily and weekly snapshots, retaining 7 daily and 52 weekly snapshots.
- The default backup policy replicates daily and weekly snapshots, retaining 7 daily and 52 weekly snapshots.

If you create custom policies for replication or backup, the policy labels (for example, "daily" or "weekly") must match the labels that exist in your snapshot policies or replicated volumes and backup files won't be created.

You can create snapshot, replication, and backup to object storage policies in the NetApp Backup and Recovery UI. See the section for [adding a new backup policy](#) for details.

In addition to using NetApp Backup and Recovery to create custom policies, you can use System Manager or the ONTAP Command Line Interface (CLI):

- [Create a snapshot policy using System Manager or the ONTAP CLI](#)
- [Create a replication policy using System Manager or the ONTAP CLI](#)

Note: When using System Manager, select **Asynchronous** as the policy type for replication policies, and select **Asynchronous** and **Back up to cloud** for backup to object policies.

Here are a few sample ONTAP CLI commands that might be helpful if you are creating custom policies. Note that you must use the *admin* vserver (storage VM) as the *<vserver_name>* in these commands.

Policy Description	Command
Simple snapshot policy	<pre>snapshot policy create -policy WeeklySnapshotPolicy -enabled true -schedule1 weekly -count1 10 -vserver ClusterA -snapmirror-label1 weekly</pre>
Simple backup to cloud	<pre>snapmirror policy create -policy <policy_name> -transfer -priority normal -vserver <vserver_name> -create -snapshot-on-source false -type vault snapmirror policy add-rule -policy <policy_name> -vserver <vserver_name> -snapmirror-label <snapmirror_label> -keep</pre>

Policy Description	Command
Backup to cloud with DataLock and Ransomware protection	<pre>snapmirror policy create -policy CloudBackupService-Enterprise -snapshot-lock-mode enterprise -vserver <vserver_name> snapmirror policy add-rule -policy CloudBackupService-Enterprise -retention-period 30days</pre>
Backup to cloud with archival storage class	<pre>snapmirror policy create -vserver <vserver_name> -policy <policy_name> -archive-after-days <days> -create -snapshot-on-source false -type vault snapmirror policy add-rule -policy <policy_name> -vserver <vserver_name> -snapmirror-label <snapmirror_label> -keep</pre>
Simple replication to another storage system	<pre>snapmirror policy create -policy <policy_name> -type async-mirror -vserver <vserver_name> snapmirror policy add-rule -policy <policy_name> -vserver <vserver_name> -snapmirror-label <snapmirror_label> -keep</pre>



Only vault policies can be used for backup to cloud relationships.

Where do my policies reside?

Backup policies reside in different locations depending on the backup architecture you plan to use: Fan-out or Cascading. Replication policies and Backup policies are not designed the same way because replications pair two ONTAP storage systems and backup to object uses a storage provider as the destination.

- Snapshot policies always reside on the primary storage system.
- Replication policies always reside on the secondary storage system.
- Backup-to-object policies are created on the system where the source volume resides - this is the primary cluster for fan-out configurations, and the secondary cluster for cascading configurations.

These differences are shown in the table.

Architecture	Snapshot policy	Replication policy	Backup policy
Fan-out	Primary	Secondary	Primary
Cascade	Primary	Secondary	Secondary

So if you're planning to create custom policies when using the cascading architecture, you'll need to create the replication and backup to object policies on the secondary system where the replicated volumes will be created. If you're planning to create custom policies when using the fan-out architecture, you'll need to create the replication policies on the secondary system where the replicated volumes will be created and backup to object policies on the primary system.

If you're using the default policies that exist on all ONTAP systems, then you're all set.

Do you want to create your own object storage container

When you create backup files in object storage for a system, by default, the backup and recovery service creates the container (bucket or storage account) for the backup files in the object storage account that you

have configured. The AWS or GCP bucket is named "netapp-backup-<uuid>" by default. The Azure Blob storage account is named "netappbackup<uuid>".

You can create the container yourself in the object provider account if you want to use a certain prefix or assign special properties. If you want to create your own container, you must create it before starting the activation wizard. NetApp Backup and Recovery can use any bucket and share buckets. The backup activation wizard will automatically discover your provisioned containers for the selected Account and credentials so that you can select the one you want to use.

You can create the bucket from the Console, or from your cloud provider.

- [Create Amazon S3 buckets from the Console](#)
- [Create Azure Blob storage accounts from the Console](#)
- [Create Google Cloud Storage buckets from the Console](#)

If you plan to use a different bucket prefix than "netapp-backup-xxxxxx", then you'll need to modify the S3 permissions for the Console agent IAM Role.

Advanced bucket settings

If you plan to move older backup files to archival storage, or if you plan to enable DataLock and Ransomware protection to lock your backup files and scan them for possible ransomware, you'll need to create the container with certain configuration settings:

- Archival storage on your own buckets is supported in AWS S3 storage at this time when using ONTAP 9.10.1 or greater software on your clusters. By default, backups start in the S3 *Standard* storage class. Ensure that you create the bucket with the appropriate lifecycle rules:
 - Move the objects in the entire scope of the bucket to S3 *Standard-IA* after 30 days.
 - Move the objects with the tag "smc_push_to_archive: true" to *Glacier Flexible Retrieval* (formerly S3 *Glacier*)
- DataLock and Ransomware protection are supported in AWS storage when using ONTAP 9.11.1 or greater software on your clusters, and Azure storage when using ONTAP 9.12.1 or greater software.
 - For AWS, you must enable Object Locking on the bucket using a 30-day retention period.
 - For Azure, you need to create the Storage Class with version-level immutability support.

Which Console agent deployment mode are you using

If you're already using the Console to manage your storage, then a Console agent has already been installed. If you plan to use the same Console agent with NetApp Backup and Recovery, then you're all set. If you need to use a different Console agent, you'll need to install it before starting your backup and recovery implementation.

NetApp Console offers multiple deployment modes that enable you to use the Console in a way that meets your business and security requirements. *Standard mode* leverages the Console SaaS layer to provide full functionality, while *restricted mode* and *private mode* are available for organizations that have connectivity restrictions.

[Learn more about NetApp Console deployment modes.](#)

Support for sites with full internet connectivity

When NetApp Backup and Recovery is used in a site with full internet connectivity (also known as *standard*

mode or SaaS mode), you can create replicated volumes on any on-premises ONTAP or Cloud Volumes ONTAP systems managed by the Console, and you can create backup files on object storage in any of the supported cloud providers. [See the full list of supported backup destinations](#).

For a list of valid Console agent locations, refer to one of the following backup procedures for the cloud provider where you plan to create backup files. There are some restrictions where the Console agent must be installed manually on a Linux machine or deployed in a specific cloud provider.

- [Back up Cloud Volumes ONTAP data to Amazon S3](#)
- [Back up Cloud Volumes ONTAP data to Azure Blob](#)
- [Back up Cloud Volumes ONTAP data to Google Cloud](#)
- [Back up on-premises ONTAP data to Amazon S3](#)
- [Back up on-premises ONTAP data to Azure Blob](#)
- [Back up on-premises ONTAP data to Google Cloud](#)
- [Back up on-premises ONTAP data to StorageGRID](#)
- [Back up on-premises ONTAP to ONTAP S3](#)

Support for sites with limited internet connectivity

NetApp Backup and Recovery can be used in a site with limited internet connectivity (also known as *restricted mode*) to back up volume data. In this case, you'll need to deploy the Console agent in the destination cloud region.

- You can back up data from on-premises ONTAP systems or Cloud Volumes ONTAP systems installed in AWS commercial regions to Amazon S3. [Back up Cloud Volumes ONTAP data to Amazon S3](#).
- You can back up data from on-premises ONTAP systems or Cloud Volumes ONTAP systems installed in Azure commercial regions to Azure Blob. [Back up Cloud Volumes ONTAP data to Azure Blob](#).

Support for sites with no internet connectivity

NetApp Backup and Recovery can be used in a site with no internet connectivity (also known as *private mode* or *dark sites*) to back up volume data. In this case, you'll need to deploy the Console agent on a Linux host in the same site.

 BlueXP private mode (legacy BlueXP interface) is typically used with on-premises environments that have no internet connection and with secure cloud regions, which includes AWS Secret Cloud, AWS Top Secret Cloud, and Azure IL6. NetApp continues to support these environments with the legacy BlueXP interface. For private mode documentation in the legacy BlueXP interface, refer to the [PDF documentation for BlueXP private mode](#).

- You can back up data from local on-premises ONTAP systems to local NetApp StorageGRID systems. [Back up on-premises ONTAP data to StorageGRID](#).
- You can back up data from local on-premises ONTAP systems to local on-premises ONTAP systems or Cloud Volumes ONTAP systems configured for S3 object storage. [Back up on-premises ONTAP data to ONTAP S3](#).

Manage backup policies for ONTAP volumes with NetApp Backup and Recovery

With NetApp Backup and Recovery, use the default backup policies provided by NetApp to create your backups, or create custom policies. Policies govern the backup frequency,

the time the backup is taken, and the number of backup files that are retained.



To switch to and from NetApp Backup and Recovery workloads, refer to [Switch to different NetApp Backup and Recovery workloads](#).

When you use the activation wizard to enable the backup and recovery service for your volumes, you can select from the default policies and any other policies that already exist in the system (Cloud Volumes ONTAP or on-premises ONTAP system). If you want to use a policy different than those existing policies, you can create the policy before or while you use the activation wizard.

To learn about the default backup policies provided, refer to [Plan your protection journey](#).

NetApp Backup and Recovery provides three types of backups of ONTAP data: Snapshots, replications, and backups to object storage. Their policies reside in different locations based on the architecture that you use and the type of backup:

Architecture	Snapshot policy storage location	Replication policy storage location	Backup to object policy storage location
Fan-out	Primary	Secondary	Primary
Cascade	Primary	Secondary	Secondary

Create backup policies using the following tools depending on your environment, your preferences, and the protection type:

- NetApp Console UI
- System Manager UI
- ONTAP CLI



When using System Manager, select **Asynchronous** as the policy type for replication policies, and select **Asynchronous** and **Back up to cloud** for backup to object policies.

View policies for a system

1. In the Console UI, select **Volumes > Backup settings**.
2. From the Backup Settings page, select the system, select the **Actions** icon, and select **Policies management**.

The Policies management page appears. Snapshot policies are displayed by default.

3. To view other policies that exist in the system, select either **Replication Policies** or **Backup Policies**. If the existing policies can be used for your backup plans, you're all set. If you need to have a policy with different characteristics, you can create new policies from this page.

Create policies

You can create policies that govern your snapshots, replications and backups to object storage:

- [Create a snapshot policy before initiating the snapshot](#)
- [Create a replication policy before initiating the replication](#)

- [Create a backup-to-object-storage policy before initiating the backup](#)

Create a snapshot policy before initiating the snapshot

Part of your 3-2-1 strategy involves creating a snapshot of the volume on the **primary** storage system.

Part of the policy creation process involves identifying snapshot and SnapMirror labels that denote the schedule and retention. You can use predefined labels or create your own.

Steps

1. In the Console UI, select **Volumes > Backup settings**.
2. From the Backup Settings page, select the system, select the **Actions ...** icon, and select **Policies management**.

The Policies management page appears.

3. In the Policies page, select **Create policy > Create Snapshot policy**.
4. Specify the policy name.
5. Select the snapshot schedule or schedules. You can have a maximum of 5 labels. Or, create a schedule.
6. If you choose to create a schedule:
 - a. Select the frequency of hourly, daily, weekly, monthly, or yearly.
 - b. Specify the snapshot labels denoting the schedule and retention.
 - c. Enter when and how often the snapshot will be taken.
 - d. Retention: Enter the number of snapshots to keep.
7. Select **Create**.

Snapshot policy example using cascading architecture

This example creates a snapshot policy with two clusters:

1. Cluster 1:
 - a. Select Cluster 1 on the policy page.
 - b. Ignore the Replication and Backup to Object policy sections.
 - c. Create the snapshot policy.
2. Cluster 2:
 - a. Select Cluster 2 on the Policy page.
 - b. Ignore the snapshot policy section.
 - c. Configure the Replication and Backup to object policies.

Create a replication policy before initiating the replication

Your 3-2-1 strategy might include replicating a volume on a different storage system. The replication policy resides on the **secondary** storage system.

Steps

1. In the Policies page, select **Create policy > Create replication policy**.

2. In the Policy Details section, specify the policy name.
3. Specify the SnapMirror labels (maximum of 5) denoting the retention for each label.
4. Specify the transfer schedule.
5. Select **Create**.

Create a backup-to-object-storage policy before initiating the backup

Your 3-2-1 strategy might include backing up a volume to object storage.

This storage policy resides in different storage system locations depending on the backup architecture:

- Fan-out: Primary storage system
- Cascading: Secondary storage system

Steps

1. In the Policy management page, select **Create policy > Create backup policy**.
2. In the Policy Details section, specify the policy name.
3. Specify the SnapMirror labels (maximum of 5) denoting the retention for each label.
4. Specify the settings, including the transfer schedule and when to archive backups.
5. (Optional) To move older backup files to a less expensive storage class or access tier after a certain number of days, select the **Archive** option and indicate the number of days that should elapse before the data is archived. Enter **0** as the "Archive After Days" to send your backup file directly to archival storage.

[Learn more about archival storage settings.](#)

6. (Optional) To protect your backups from being modified or deleted, select the **DataLock & Ransomware protection** option.

If your cluster is using ONTAP 9.11.1 or greater, you can choose to protect your backups from deletion by configuring *DataLock* and *Ransomware protection*.

[Learn more about the available DataLock settings.](#)

7. Select **Create**.

Edit a policy

You can edit a custom snapshot, replication, or backup policy.

Changing the backup policy affects all volumes that are using that policy.

Steps

1. In the Policies management page, select the policy, select the **Actions**  icon, and select **Edit policy**.



The process is the same for replication and backup policies.

2. In the Edit Policy page, make the changes.
3. Select **Save**.

Delete a policy

You can delete policies that are not associated with any volumes.

If a policy is associated with a volume and you want to delete the policy, you must remove the policy from the volume first.

Steps

1. In the Policies management page, select the policy, select the **Actions**  icon, and select **Delete Snapshot policy**.
2. Select **Delete**.

Find more information

For instructions on creating policies using System Manager or ONTAP CLI, see the following:

[Create a Snapshot policy using System Manager](#)

[Create a Snapshot policy using the ONTAP CLI](#)

[Create a replication policy using System Manager](#)

[Create a replication policy using the ONTAP CLI](#)

[Create a backup to object storage policy using System Manager](#)

[Create a backup to object storage policy using the ONTAP CLI](#)

Backup-to-object policy options in NetApp Backup and Recovery

NetApp Backup and Recovery enables you to create backup policies with a variety of settings for your on-premises ONTAP and Cloud Volumes ONTAP systems.



These policy settings are relevant for backup-to-object storage only. None of these settings affect your snapshot or replication policies.



To switch to and from NetApp Backup and Recovery workloads, refer to [Switch to different NetApp Backup and Recovery workloads](#).

Backup schedule options

NetApp Backup and Recovery enables you to create multiple backup policies with unique schedules for each system (cluster). You can assign different backup policies to volumes that have different recovery point objectives (RPO).

Each backup policy provides a section for *Labels & Retention* that you can apply to your backup files. Note that the Snapshot policy applied to the volume must be one of the policies recognized by NetApp Backup and Recovery or backup files will not be created.

There are two parts of the schedule; the Label and the Retention value:

- The **label** defines how often a backup file is created (or updated) from the volume. You can select among the following types of labels:
 - You can choose one, or a combination of, **hourly**, **daily**, **weekly**, **monthly**, and **yearly** timeframes.
 - You can select one of the system-defined policies that provide backup and retention for 3 months, 1 year, or 7 years.

- If you have created custom backup protection policies on the cluster using ONTAP System Manager or the ONTAP CLI, you can select one of those policies.
- The **retention** value defines how many backup files for each label (timeframe) are retained. Once the maximum number of backups have been reached in a category, or interval, older backups are removed so you always have the most current backups. This also saves you storage costs because obsolete backups don't continue to take up space in the cloud.

For example, say you create a backup policy that creates 7 **weekly** and 12 **monthly** backups:

- each week and each month a backup file is created for the volume
- at the 8th week, the first weekly backup is removed, and the new weekly backup for the 8th week is added (keeping a maximum of 7 weekly backups)
- at the 13th month, the first monthly backup is removed, and the new monthly backup for the 13th month is added (keeping a maximum of 12 monthly backups)

Yearly backups are deleted automatically from the source system after being transferred to object storage. This default behavior can be changed in the Advanced Settings page for the system.

DataLock and Ransomware protection options

NetApp Backup and Recovery provides support for DataLock and Ransomware protection for your volume backups. These features enable you to lock your backup files and scan them to detect possible ransomware on the backup files. This is an optional setting that you can define in your backup policies when you want extra protection for your volume backups for a cluster.

Both of these features protect your backup files so that you'll always have a valid backup file to recover data from in case of a ransomware attack attempt on your backups. It's also helpful to meet certain regulatory requirements where backups need to be locked and retained for a certain period of time. When the DataLock and Ransomware Resilience option is enabled, the cloud bucket that is provisioned as a part of NetApp Backup and Recovery activation will have object locking and object versioning enabled.

This feature does not provide protection for your source volumes; only for the backups of those source volumes. Use some of the [anti-ransomware protections provided from ONTAP](#) to protect your source volumes.

- If you plan to use DataLock and Ransomware protection, you can enable it when creating your first backup policy and activating NetApp Backup and Recovery for that cluster. You can later enable or disable ransomware scanning using NetApp Backup and Recovery Advanced Settings.
- When the Console scans a backup file for ransomware when restoring volume data, you'll incur extra egress costs from your cloud provider to access the contents of the backup file.

What is DataLock

With this feature, you can lock the cloud snapshots replicated via SnapMirror to Cloud and also enable the feature to detect a ransomware attack and recover a consistent copy of the snapshot on the object store. This feature is supported on AWS, Azure, Google Cloud Platform, and StorageGRID.

DataLock protects your backup files from being modified or deleted for a certain period of time - also called *immutable storage*. This functionality uses technology from the object storage provider for "object locking."

Cloud providers use a Retention Until Date (RUD), which is calculated based on the Snapshot Retention Period. The Snapshot Retention Period is calculated based on the label and the retention count defined in the backup policy.

The minimum Snapshot Retention Period is 30 days. Let's look at some examples of how this works:

- If you choose the **Daily** label with Retention Count 20, the Snapshot Retention Period is 20 days, which defaults to the minimum 30 days.
- If you choose the **Weekly** label with Retention Count 4, the Snapshot Retention Period is 28 days, which defaults to the minimum of 30 days.
- If you choose the **Monthly** label with Retention Count 3, the Snapshot Retention Period is 90 days.
- If you choose the **Yearly** label with Retention Count 1, the Snapshot Retention Period is 365 days.

What is Retention Until Date (RUD) and how is it calculated?

The Retention Until Date (RUD) is determined based on the Snapshot Retention Period. The Retention Until Date is calculated by summing the Snapshot Retention Period and a Buffer.

- Buffer is the Buffer for Transfer Time (3 days) + Buffer for Cost Optimization (28 days), which totals as 31 days.
- The minimum Retention Until Date is 30 days + 31 days buffer = 61 days.

Here are some examples:

- If you create a Monthly backup schedule with 12 retentions, your backups are locked for 12 months (plus 31 days) before they are deleted (replaced by the next backup file).
- If you create a backup policy that creates 30 daily, 7 weekly, 12 monthly backups, there are three locked retention periods:
 - The "30 daily" backups are retained for 61 days (30 days plus 31 days buffer),
 - The "7 weekly" backups are retained for 11 weeks (7 weeks plus 31 days), and
 - The "12 monthly" backups are retained for 12 months (plus 31 days).
- If you create an Hourly backup schedule with 24 retentions, you might think that backups are locked for 24 hours. However, since that is less than the minimum of 30 days, each backup will be locked and retained for 61 days (30 days plus 31 days buffer).



Old backups are deleted after the DataLock Retention Period expires, not after the backup policy retention period.

The DataLock retention setting overrides the policy retention setting from your backup policy. This could affect your storage costs as your backup files will be saved in the object store for a longer period of time.

Enable DataLock and Ransomware protection

You can enable DataLock and Ransomware protection when you create a policy. You cannot enable, modify, or disable this after the policy is created.

1. When you create a policy, expand the **DataLock and Ransomware Resilience** section.
2. Choose one of the following:
 - **None**: DataLock protection and Ransomware Resilience are disabled.
 - **Unlocked**: DataLock protection and Ransomware Resilience are enabled. Users with specific permissions can overwrite or delete protected backup files during the retention period.
 - **Locked**: DataLock protection and Ransomware Resilience are enabled. No users can overwrite or

delete protected backup files during the retention period. This satisfies full regulatory compliance.

Refer to [How to update Ransomware protection options in the Advanced Settings page](#).

What is Ransomware protection

Ransomware protection scans your backup files to look for evidence of a ransomware attack. The detection of ransomware attacks is performed using a checksum comparison. If potential ransomware is identified in a new backup file versus the previous backup file, that newer backup file is replaced by the most recent backup file that does not show any signs of a ransomware attack. (The file that was identified as having a ransomware attack is deleted 1 day after it has been replaced.)

Scans occur in these situations:

- Scans on cloud backup objects are initiated soon after they are transferred to the cloud object storage. The scan is not performed on the backup file when it is first written to cloud storage, but when the next backup file is written.
- Ransomware scans can be initiated when the backup is selected for the restore process.
- Scans can be performed on-demand at any time.

How does the recovery process work?

When a ransomware attack is detected, the service uses the Active Data Console agent Integrity Checker REST API to start the recovery process. The oldest version of the data objects is the source of truth and is made into the current version as part of the recovery process.

Let's see how this works:

- In the event of a ransomware attack, the service tries to overwrite or delete the object in the bucket.
- Because the cloud storage is versioning-enabled, it automatically creates a new version of the backup object. If an object is deleted with versioning turned on, it is marked as deleted but is still retrievable. If an object is overwritten, previous versions are stored and marked.
- When a ransomware scan is initiated, the checksums are validated for both object versions and compared. If the checksums are inconsistent, potential ransomware has been detected.
- The recovery process involves reverting to the last known good copy.

Supported systems and object storage providers

You can enable DataLock and Ransomware protection on ONTAP volumes from the following systems when using object storage in the following public and private cloud providers.

Source system	Backup File Destination
Cloud Volumes ONTAP in AWS	Amazon S3
Cloud Volumes ONTAP in Azure	Azure Blob
Cloud Volumes ONTAP in Google Cloud	Google Cloud
On-premises ONTAP system	Amazon S3 Azure Blob Google Cloud NetApp StorageGRID

Requirements

- For AWS:
 - Your clusters must be running ONTAP 9.11.1 or greater
 - The Console agent can be deployed in the cloud or on your premises
 - The following S3 permissions must be part of the IAM role that provides the Console agent with permissions. They reside in the "backupS3Policy" section for the resource "arn:aws:s3:::netapp-backup-*":

AWS S3 permissions

- s3:GetObjectVersionTagging
- s3:GetBucketObjectLockConfiguration
- s3:GetObjectVersionAcl
- s3:PutObjectTagging
- s3:DeleteObject
- s3:DeleteObjectTagging
- s3:GetObjectRetention
- s3:DeleteObjectVersionTagging
- s3:PutObject
- s3:GetObject
- s3:PutBucketObjectLockConfiguration
- s3:GetLifecycleConfiguration
- s3:GetBucketTagging
- s3:DeleteObjectVersion
- s3>ListBucketVersions
- s3>ListBucket
- s3:PutBucketTagging
- s3:GetObjectTagging
- s3:PutBucketVersioning
- s3:PutObjectVersionTagging
- s3:GetBucketVersioning
- s3:GetBucketAcl
- s3:BypassGovernanceRetention
- s3:PutObjectRetention
- s3:GetBucketLocation
- s3:GetObjectVersion

[View the full JSON format for the policy where you can copy and paste required permissions.](#)

- For Azure:

- Your clusters must be running ONTAP 9.12.1 or greater
- The Console agent can be deployed in the cloud or on your premises
- For Google Cloud:
 - Your clusters must be running ONTAP 9.17.1 or greater
 - The Console agent can be deployed in the cloud or on your premises
- For StorageGRID:
 - Your clusters must be running ONTAP 9.11.1 or greater
 - Your StorageGRID systems must be running 11.6.0.3 or greater
 - The Console agent must be deployed on your premises (it can be installed in a site with or without internet access)
 - The following S3 permissions must be part of the IAM role that provides the Console agent with permissions:

StorageGRID S3 permissions

- s3:GetObjectVersionTagging
- s3:GetBucketObjectLockConfiguration
- s3:GetObjectVersionAcl
- s3:PutObjectTagging
- s3:DeleteObject
- s3:DeleteObjectTagging
- s3:GetObjectRetention
- s3:DeleteObjectVersionTagging
- s3:PutObject
- s3:GetObject
- s3:PutBucketObjectLockConfiguration
- s3:GetLifecycleConfiguration
- s3:GetBucketTagging
- s3:DeleteObjectVersion
- s3>ListBucketVersions
- s3>ListBucket
- s3:PutBucketTagging
- s3:GetObjectTagging
- s3:PutBucketVersioning
- s3:PutObjectVersionTagging
- s3:GetBucketVersioning
- s3:GetBucketAcl
- s3:PutObjectRetention
- s3:GetBucketLocation
- s3:GetObjectVersion

Restrictions

- The DataLock and Ransomware protection feature is not available if you have configured archival storage in the backup policy.
- The DataLock option you select when activating NetApp Backup and Recovery must be used for all backup policies for that cluster.
- You cannot use multiple DataLock modes on a single cluster.
- If you enable DataLock, all volume backups will be locked. You can't mix locked and non-locked volume backups for a single cluster.
- DataLock and Ransomware protection is applicable for new volume backups using a backup policy with DataLock and Ransomware protection enabled. You can later enable or disable these features using the Advanced Settings option.

- FlexGroup volumes can use DataLock and Ransomware protection only when using ONTAP 9.13.1 or greater.

Tips on how to mitigate DataLock costs

You can enable or disable the Ransomware Scan feature while keeping the DataLock feature active. To avoid extra charges, you can disable scheduled ransomware scans. This lets you customize your security settings and avoid incurring costs from the cloud provider.

Even if scheduled ransomware scans are disabled, you can still perform on-demand scans when needed.

You can choose different levels of protection:

- **DataLock without ransomware scans:** Provides protection for backup data in the destination storage that can be either in Governance or Compliance mode.
 - **Governance mode:** Offers flexibility to administrators to overwrite or delete protected data.
 - **Compliance mode:** Provides complete indelibility until the retention period expires. This helps meet the most stringent data security requirements of highly regulated environments. The data cannot be overwritten or modified during its lifecycle, providing the strongest level of protection for your backup copies.



Microsoft Azure uses a Lock and Unlock mode instead.

- **DataLock with ransomware scans:** Provides an additional layer of security for your data. This feature helps detect any attempts to change backup copies. If any attempt is made, a new version of the data is created discreetly. The scan frequency can be changed to 1, 2, 3, 4, 5, 6, or 7 days. If scans are set to every 7 days, the costs decrease significantly.

For more tips to mitigate DataLock costs, refer to

<https://community.netapp.com/t5/Tech-ONTAP-Blogs/Understanding-NetApp-Backup-and-Recovery-DataLock-and-Ransomware-Feature-TCO/ba-p/453475>

Additionally, you can get estimates for the cost associated with DataLock by visiting the [NetApp Backup and Recovery Total Cost of Ownership \(TCO\) calculator](#).

Archival storage options

When using AWS, Azure, or Google cloud storage, you can move older backup files to a less expensive archival storage class or access tier after a certain number of days. You can also choose to send your backup files to archival storage immediately without being written to standard cloud storage. Just enter **0** as the "Archive After Days" to send your backup file directly to archival storage. This can be especially helpful for users who rarely need to access data from cloud backups or users who are replacing a backup to tape solution.

Data in archival tiers can't be accessed immediately when needed, and will require a higher retrieval cost, so you'll need to consider how often you may need to restore data from backup files before deciding to archive your backup files.

A blue circular icon with a white letter 'i' inside, representing an information or note.

- Even if you select "0" to send all data blocks to archival cloud storage, metadata blocks are always written to standard cloud storage.
- Archival storage can't be used if you have enabled DataLock.
- You can't change the archival policy after selecting **0** days (archive immediately).

Each backup policy provides a section for *Archival Policy* that you can apply to your backup files.

- In AWS, backups start in the *Standard* storage class and transition to the *Standard-Infrequent Access* storage class after 30 days.

If your cluster is using ONTAP 9.10.1 or greater, you can tier older backups to either S3 *Glacier* or S3 *Glacier Deep Archive* storage. [Learn more about AWS 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.

- In Azure, backups are associated with the *Cool* access tier.

If your cluster is using ONTAP 9.10.1 or greater, you can tier older backups to *Azure Archive* storage. [Learn more about Azure archival storage](#).

- In GCP, backups are associated with the *Standard* storage class.

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 for further cost optimization. [Learn more about Google archival storage](#).

- In StorageGRID, backups are associated with the *Standard* storage class.

If your on-prem cluster is using ONTAP 9.12.1 or greater, and your StorageGRID system is using 11.4 or greater, you can archive older backup files to public cloud archival storage.

- For AWS, you can tier backups to AWS S3 *Glacier* or S3 *Glacier Deep Archive* storage. [Learn more about AWS archival storage](#).
- For Azure, you can tier older backups to *Azure Archive* storage. [Learn more about Azure archival storage](#).

Manage backup-to-object storage options in NetApp Backup and Recovery Advanced Settings

You can change cluster-level, backup-to-object storage settings that you set when activating NetApp Backup and Recovery for each ONTAP system by using the Advanced Settings page. You can also modify some settings that are applied as "default" backup settings. This includes changing the transfer rate of backups to object storage, whether historical snapshots are exported as backup files, and enabling or disabling ransomware scans for a system.



These settings are available for backup-to-object storage only. None of these settings affect your Snapshot or replication settings.



To switch to and from NetApp Backup and Recovery workloads, refer to [Switch to different NetApp Backup and Recovery workloads](#).

You can change the following options in the Advanced Settings page:

- Changing the storage keys that give your ONTAP system permission to access object storage
- Changing the ONTAP IPspace that is connected to object storage
- Changing the network bandwidth allocated to upload backups to object storage using the Max Transfer Rate option
- Changing whether historical snapshots are exported as backup files and included in your initial baseline backup files for future volumes
- Changing whether "yearly" snapshots are removed from the source system
- Enabling or disabling ransomware scans for a system, including scheduled scans

View cluster-level backup settings

You can view the cluster-level system settings and provider settings for each system.

Steps

1. From the Console menu, select **Protection > Backup and recovery**.
2. From the **Volumes** tab, select **Backup Settings**.
3. From the *Backup Settings page*, select the **...** for the system and select **Configure advanced settings > System settings** to view system settings and **Configure advanced settings > Provider settings** to view provider settings.

The resulting page displays the current settings for that system. When viewing provider settings, the provider settings shown are relevant for the bucket that you select at the top of the page.

Note that some options are unavailable based on the version of ONTAP on the source cluster and the cloud provider destination where the backups reside.

Change the network bandwidth available to upload backups to object storage

When you activate NetApp Backup and Recovery for a system, by default, ONTAP can use an unlimited amount of bandwidth to transfer the backup data from volumes in the system to object storage. If you notice that backup traffic is affecting normal user workloads, you can throttle the amount of network bandwidth that is used during the transfer using the Max Transfer Rate option in the Advanced Settings page.

Steps

1. From the **Volumes** tab, select **Backup Settings**.
2. From the *Backup Settings page*, click **...** for the system and select **Configure advanced settings > System settings**.
3. In the Advanced Settings page, expand the **Max Transfer Rate** section.
4. Choose a value between 1 and 1,000 Mbps as the maximum transfer rate.
5. Select the **Limited** radio button and enter the maximum bandwidth that can be used, or select **Unlimited** to indicate that there is no limit.
6. Select **Apply**.

This setting does not affect the bandwidth allocated to any other replication relationships that may be configured for volumes in the system.

Change whether historical snapshots are exported as backup files

If there are any local snapshots for volumes that match the backup schedule label you're using in this system (for example, daily, weekly, etc.), you can export those historic snapshots to object storage as backup files. This enables you to initialize your backups in the cloud by moving older snapshots into the baseline backup copy.

Note that this option only applies to new backup files for new read/write volumes, and it is not supported with data protection (DP) volumes.

Steps

1. From the **Volumes** tab, select **Backup Settings**.
2. From the *Backup Settings page*, click **...** for the system and select **Configure advanced settings > System settings**.
3. In the Advanced Settings page, expand the **Export existing snapshot copies** section.
4. Select whether you want existing snapshots to be exported.
5. Select **Apply**.

Change whether "yearly" snapshots are removed from the source system

When you select the "yearly" backup label for a backup policy for any of your volumes, the snapshot that is created is very large. By default, these yearly snapshots are deleted automatically from the source system after being transferred to object storage. You can change this default behavior from the Yearly Snapshot Deletion section.

Steps

1. From the **Volumes** tab, select **Backup Settings**.
2. From the *Backup Settings page*, click **...** for the system and select **Configure advanced settings > System settings**.
3. In the Advanced Settings page, expand the **Yearly Snapshot Deletion** section.
4. Select **Disabled** to retain the yearly snapshots on the source system.
5. Select **Apply**.

Enable or disable ransomware scans

Ransomware protection scans are enabled by default. The default setting for the scan frequency is for 7 days. The scan occurs only on the latest snapshot.

For details about DataLock and Ransomware Resilience options, refer to [DataLock and Ransomware Resilience options](#).

You can change that schedule to days or weeks or disable it, saving costs.



Enabling ransomware scans will incur extra charges depending on the cloud provider.

If the scheduled ransomware scans are disabled, you can still perform on-demand scans and the scan during a restore operation will still occur.

Refer to [Manage policies](#) for details about managing policies that implement ransomware detection.

Enable or disable ransomware scans for a system

You can enable or disable ransomware scans for a cluster.

Steps

1. From the **Volumes** tab, select **Backup Settings**.
2. From the *Backup Settings page*, click **...** for the system and select **Configure advanced settings > System settings**.
3. In the resulting page, expand the **Ransomware scan** section.
4. Enable or disable **Ransomware scan**.
5. Select **Scheduled ransomware scan**.
6. Optionally, change the every week default scan to days or weeks.
7. Set the how often in days or weeks that the scan should run.
8. Select **Apply**.

Enable or disable ransomware scans for a provider

You can enable or disable ransomware scans at the provider level by using the provider settings page. The settings on the page are relevant to the bucket that you select at the top of the page.

Steps

1. From the **Volumes** tab, select **Backup Settings**.
2. From the *Backup Settings page*, click **...** for the system and select **Configure advanced settings > Provider settings**.
3. At the top of the resulting page, select the bucket for which you need to change settings.
4. Expand the **Ransomware scan** section.
5. Enable or disable **Ransomware scan**.
6. Select **Scheduled ransomware scan**.
7. Optionally, change the every week default scan to days or weeks.
8. Set the how often in days or weeks that the scan should run.
9. Select **Apply**.

Back up Cloud Volumes ONTAP data to Amazon S3 with NetApp Backup and Recovery

Complete a few steps in NetApp Backup and Recovery to get started backing up volume data from your Cloud Volumes ONTAP systems to Amazon S3.



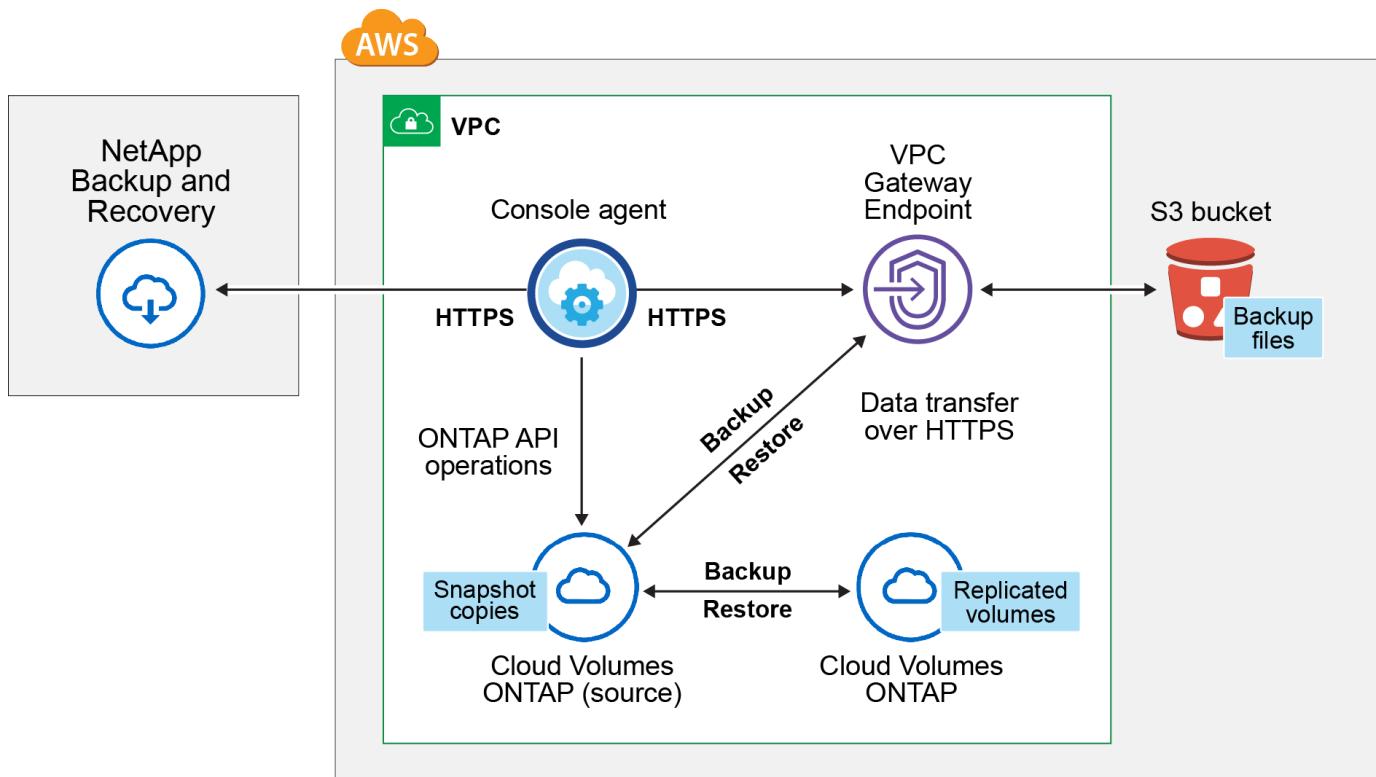
To switch to and from NetApp Backup and Recovery workloads, refer to [Switch to different NetApp Backup and Recovery workloads](#).

Verify support for your configuration

Read the following requirements to make sure that you have a supported configuration before you start backing up volumes to S3.

The following image shows each component and the connections that you need to prepare between them.

Optionally, you can connect to a secondary ONTAP system for replicated volumes using the public or private connection as well.



The VPC gateway endpoint must exist in your VPC already. [Learn more about gateway endpoints](#).

Supported ONTAP versions

Minimum of ONTAP 9.8; ONTAP 9.8P13 and later is recommended.

Required information for using customer-managed keys for data encryption

You can choose your own customer-managed keys for data encryption in the activation wizard instead of using the default Amazon S3 encryption keys. In this case you'll need to have the encryption managed keys already set up. [See how to use your own keys](#).

Verify license requirements

For NetApp Backup and Recovery PAYGO licensing, a Console subscription is available in the AWS Marketplace that enables deployments of Cloud Volumes ONTAP and NetApp Backup and Recovery. You need to [subscribe to this NetApp Console subscription](#) before you enable NetApp Backup and Recovery. Billing for NetApp Backup and Recovery is done through this subscription.

For an annual contract that enables you to back up both Cloud Volumes ONTAP data and on-premises ONTAP data, you need to subscribe from the [AWS Marketplace page](#) and then [associate the subscription with your AWS credentials](#).

For an annual contract that enables you to bundle Cloud Volumes ONTAP and NetApp Backup and Recovery, you must set up the annual contract when you create a Cloud Volumes ONTAP system. This option doesn't enable you to back up on-prem data.

For NetApp Backup and Recovery BYOL licensing, you need the serial number from NetApp that enables you to use the service for the duration and capacity of the license. [Learn how to manage your BYOL licenses](#). You must use a BYOL license when the Console agent and Cloud Volumes ONTAP system are deployed in a dark site.

And you need to have an AWS account for the storage space where your backups will be located.

Prepare your Console agent

The Console agent must be installed in an AWS region with full or limited internet access ("standard" or "restricted" mode). [See NetApp Console deployment modes for details](#).

- [Learn about Console agents](#)
- [Deploy a Console agent in AWS in standard mode \(full internet access\)](#)
- [Install the Console agent in restricted mode \(limited outbound access\)](#)

Verify or add permissions to the Console agent

The IAM role that provides the Console with permissions must include S3 permissions from the latest [Console policy](#). If the policy does not contain all of these permissions, see the [AWS Documentation: Editing IAM policies](#).

Here are the specific permissions from the policy:

```
{  
  "Sid": "backupPolicy",  
  "Effect": "Allow",  
  "Action": [  
    "s3:DeleteBucket",  
    "s3:GetLifecycleConfiguration",  
    "s3:PutLifecycleConfiguration",  
    "s3:PutBucketTagging",  
    "s3>ListBucketVersions",  
    "s3:GetObject",  
    "s3:DeleteObject",  
    "s3:PutObject",  
    "s3>ListBucket",  
    "s3>ListAllMyBuckets",  
    "s3:GetBucketTagging",  
    "s3:GetBucketLocation",  
    "s3:GetBucketPolicyStatus",  
    "s3:GetBucketPublicAccessBlock",  
    "s3:GetBucketAcl",  
    "s3:GetBucketPolicy",  
    "s3:PutBucketPolicy",  
    "s3:PutBucketOwnershipControls",  
    "s3:PutBucketPublicAccessBlock",  
    "s3:PutEncryptionConfiguration",  
    "s3:GetObjectVersionTagging",  
    "s3:GetBucketObjectLockConfiguration",  
    "s3:GetObjectVersionAcl",  
    "s3:PutObjectTagging",  
    "s3:DeleteObjectTagging",  
    "s3:GetObjectRetention",  
    "s3:DeleteObjectVersionTagging",  
    "s3:PutBucketObjectLockConfiguration",  
    "s3:DeleteObjectVersion",  
    "s3:GetObjectTagging",  
    "s3:PutBucketVersioning",  
    "s3:PutObjectVersionTagging",  
    "s3:GetBucketVersioning",  
    "s3:BypassGovernanceRetention",  
    "s3:PutObjectRetention",  
    "s3:GetObjectVersion",  
    "athena:StartQueryExecution",  
    "athena:GetQueryResults",  
    "athena:GetQueryExecution",  
    "glue:GetDatabase",  
    "glue:GetTable",  
  ]  
}
```

```

    "glue:CreateTable",
    "glue:CreateDatabase",
    "glue:GetPartitions",
    "glue:BatchCreatePartition",
    "glue:BatchDeletePartition"
],
"Resource": [
    "arn:aws:s3:::netapp-backup-*"
]
}

```

 When creating backups in AWS China regions, you need to change the AWS Resource Name "arn" under all *Resource* sections in the IAM policies from "aws" to "aws-cn"; for example `arn:aws-cn:s3:::netapp-backup-*`.

Required AWS Cloud Volumes ONTAP permissions

When your Cloud Volumes ONTAP system is running ONTAP 9.12.1 or greater software, the IAM role that provides that system with permissions must include a new set of S3 permissions specifically for NetApp Backup and Recovery from the latest [Cloud Volumes ONTAP policy](#).

If you created the Cloud Volumes ONTAP system using Console version 3.9.23 or greater, these permissions should be part of the IAM role already. Otherwise you'll need to add the missing permissions.

Supported AWS regions

NetApp Backup and Recovery is supported in all AWS regions, including AWS GovCloud regions.

Required setup for creating backups in a different AWS account

By default, backups are created using the same account as the one used for your Cloud Volumes ONTAP system. If you want to use a different AWS account for your backups, you must:

- Verify that the permissions "s3:PutBucketPolicy" and "s3:PutBucketOwnershipControls" are part of the IAM role that provides the Console agent with permissions.
- Add the destination AWS account credentials in the Console. [See how to do this](#).
- Add the following permissions in the user credentials in the second account:

```

"athena:StartQueryExecution",
"athena:GetQueryResults",
"athena:GetQueryExecution",
"glue:GetDatabase",
"glue:GetTable",
"glue:CreateTable",
"glue:CreateDatabase",
"glue:GetPartitions",
"glue:BatchCreatePartition",
"glue:BatchDeletePartition"

```

Create your own buckets

By default, the service creates buckets for you. If you want to use your own buckets, you can create them before you start the backup activation wizard and then select those buckets in the wizard.

[Learn more about creating your own buckets.](#)

Verify ONTAP networking requirements for replicating volumes

Unresolved directive in prev-ontap-backup-cvo-aws.adoc - include:::_include/backup-onprem-prepare-source-destination-systems-for-repl.adoc[]

- To replicate data between two Cloud Volumes ONTAP systems in different subnets, the subnets must be routed together (this is the default setting).

Enable NetApp Backup and Recovery on Cloud Volumes ONTAP

Enabling NetApp Backup and Recovery is easy. The steps differ slightly depending on whether you have an existing Cloud Volumes ONTAP system or a new one.

Enable NetApp Backup and Recovery on a new system

NetApp Backup and Recovery is enabled by default in the system wizard. Be sure to keep the option enabled.

See [Launching Cloud Volumes ONTAP in AWS](#) for requirements and details for creating your Cloud Volumes ONTAP system.

Steps

1. From the Console **Systems** page, select **Add system**, choose the cloud provider, and select **Add New**. Select **Create Cloud Volumes ONTAP**.
2. Select **Amazon Web Services** as the cloud provider and then choose a single node or HA system.
3. Fill out the Details & Credentials page.
4. On the Services page, leave the service enabled and select **Continue**.
5. Complete the pages in the wizard to deploy the system.

Result

NetApp Backup and Recovery is enabled on the system. After you've created volumes on these Cloud Volumes ONTAP systems, launch NetApp Backup and Recovery and [activate backup on each volume that you want to protect](#).

Enable NetApp Backup and Recovery on an existing system

Enable NetApp Backup and Recovery on an existing system at any time directly from the Console.

Steps

1. From the Console **Systems** page, select the cluster and select **Enable** next to Backup and Recovery in the right-panel.

If the Amazon S3 destination for your backups exists as a cluster on the **Systems** page, you can drag the cluster onto the Amazon S3 system to initiate the setup wizard.

Activate backups on your ONTAP volumes

Activate backups at any time directly from your on-premises system.

A wizard takes you through the following major steps:

- [Select the volumes that you want to back up](#)
- [Define the backup strategy](#)
- [Review your selections](#)

You can also [Show the API commands](#) at the review step, so you can copy the code to automate backup activation for future systems.

Start the wizard

Steps

1. Access the Activate backup and recovery wizard using one of the following ways:

- From the **Console Systems** page, select the system and select **Enable > Backup Volumes** next to **Backup and Recovery** in the right-panel.

If the AWS destination for your backups exists as a system on the **Console Systems** page, you can drag the ONTAP cluster onto the AWS object storage.

- Select **Volumes** in the **Backup and Recovery** bar. From the **Volumes** tab, select the **Actions** **...** icon option and select **Activate 3-2-1 Protection** for a single volume (that does not already have replication or backup to object storage already enabled).

The Introduction page of the wizard shows the protection options including local Snapshots, replication, and backups. If you did the second option in this step, the Define Backup Strategy page appears with one volume selected.

2. Continue with the following options:

- If you already have a **Console** agent, you're all set. Just select **Next**.
- If you don't already have a **Console** agent, the **Add a Console agent** option appears. Refer to [Prepare your Console agent](#).

Select the volumes that you want to back up

Choose the volumes you want to protect. A protected volume is one that has one or more of the following: Snapshot policy, replication policy, backup to object policy.

You can choose to protect FlexVol or FlexGroup volumes; however, you cannot select a mix of these volumes when activating backup for a system. See how to [activate backup for additional volumes in the system](#) (FlexVol or FlexGroup) after you have configured backup for the initial volumes.

- You can activate a backup only on a single FlexGroup volume at a time.
- The volumes you select must have the same SnapLock setting. All volumes must have SnapLock Enterprise enabled or have SnapLock disabled.

Steps

If the volumes you choose already have snapshot or replication policies applied, then the policies you select later will overwrite these existing policies.

1. In the Select Volumes page, select the volume or volumes you want to protect.
 - Optionally, filter the rows to show only volumes with certain volume types, styles, and more to make the selection easier.
 - After you select the first volume, then you can select all FlexVol volumes (FlexGroup volumes can be selected one at a time only). To back up all existing FlexVol volumes, check one volume first and then check the box in the title row.
 - To back up individual volumes, check the box for each volume.

2. Select **Next**.

Define the backup strategy

Defining the backup strategy involves setting the following options:

- Whether you want one or all of the backup options: local Snapshots, replication, and backup to object storage
- Architecture
- Local snapshot policy
- Replication target and policy



If the volumes you choose have different snapshot and replication policies than the policies you select in this step, the existing policies will be overwritten.

- Backup to object storage information (provider, encryption, networking, backup policy, and export options).

Steps

1. In the Define backup strategy page, choose one or all of the following. All three are selected by default:
 - **Local Snapshots:** If you are performing replication or back up to object storage, local snapshots must be created.
 - **Replication:** Creates replicated volumes on another ONTAP storage system.
 - **Backup:** Backs up volumes to object storage. When selecting existing buckets or configuring new buckets, you can back up volumes to up to six buckets per cluster.
2. **Architecture:** If you chose replication and backup, choose one of the following flows of information:
 - **Cascading:** Information flows from the primary storage system to the secondary, and from secondary to object storage.
 - **Fan out:** Information flows from the primary storage system to the secondary *and* from the primary to object storage.

For details about these architectures, refer to [Plan your protection journey](#).

3. **Local Snapshot:** Choose an existing snapshot policy or create a new one.



To create a custom policy before activating the snapshot, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- a. Enter the name of the policy.
- b. Select up to five schedules, typically of different frequencies.

c. Select **Create**.

4. **Replication**: Set the following options:

- **Replication target**: Select the destination system and storage VM. Optionally, select the destination aggregate or aggregates and prefix or suffix that will be added to the replicated volume name.
- **Replication policy**: Choose an existing replication policy or create one.



To create a custom policy, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to five schedules, typically of different frequencies.
- Select **Create**.

5. **Backup**: Set the following options:

- **Provider**: Select **Amazon Web Services**.

- **Provider settings**: Enter the provider details and region where the backups will be stored.

Enter the AWS Account used to store the backups. This can be a different account than where the Cloud Volumes ONTAP system resides.

If you want to use a different AWS account for your backups, you must add the destination AWS account credentials in the Console, and add the permissions "s3:PutBucketPolicy" and "s3:PutBucketOwnershipControls" to the IAM role that provides the Console with permissions.

Select the region where the backups will be stored. This can be a different region than where the Cloud Volumes ONTAP system resides.

Either create a new bucket or select an existing one.

- **Encryption**: If you created a new bucket, enter encryption key information given to you from the provider. Choose whether you'll use the default AWS encryption keys, or choose your own customer-managed keys from your AWS account, to manage encryption of your data. ([See how to use your own encryption keys](#)).

If you choose to use your own customer-managed keys, enter the key vault and key information.



If you chose an existing bucket, encryption information is already available, so you don't need to enter it now.

- **Networking**: Configure networking options for this provider.
- **Backup policy**: Select an existing backup-to-object storage policy or create one.



To create a custom policy before activating the backup, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to five schedules, typically of different frequencies.

- c. For backup-to-object policies, set the DataLock and Ransomware Resilience settings. For details on DataLock and Ransomware Resilience, refer to [Backup-to-object policy settings](#).
- d. Select **Create**.
- **Export existing snapshot:** If there are any local snapshots for volumes in this system that match the backup schedule label you just selected for this system (for example, daily, weekly, etc.), this additional prompt is displayed. Check this box to have all historic snapshots copied to object storage as backup files to ensure the most complete protection for your volumes.

6. Select **Next**.

Review your selections

This is the chance to review your selections and make adjustments, if necessary.

Steps

1. In the Review page, review your selections.
2. Optionally check the box to **Automatically fix mismatched labels on local Snapshot, replication and backup**. This creates snapshots with a label that matches the labels in the snapshot, replication and backup policies.
3. Select **Activate Backup**.

Result

NetApp Backup and Recovery starts taking the initial backups of your volumes. The baseline transfer of the replicated volume and the backup file includes a full copy of the primary storage system data. Subsequent transfers contain differential copies of the primary storage system data contained in snapshots.

A replicated volume is created in the destination cluster that will be synchronized with the primary storage volume.

An S3 bucket is created in the service account indicated by the S3 access key and secret key you entered, and the backup files are stored there.

The Volume Backup Dashboard is displayed so you can monitor the state of the backups.

You can also monitor the status of backup and restore jobs using the [Job Monitoring page](#).

Show the API commands

You might want to display and optionally copy the API commands used in the Activate backup and recovery wizard. You might want to do this to automate backup activation in future systems.

Steps

1. From the Activate backup and recovery wizard, select **View API request**.
2. To copy the commands to the clipboard, select the **Copy** icon.

Back up Cloud Volumes ONTAP data to Azure Blob storage with NetApp Backup and Recovery

Complete a few steps in NetApp Backup and Recovery to get started backing up volume data from your Cloud Volumes ONTAP systems to Azure Blob storage.



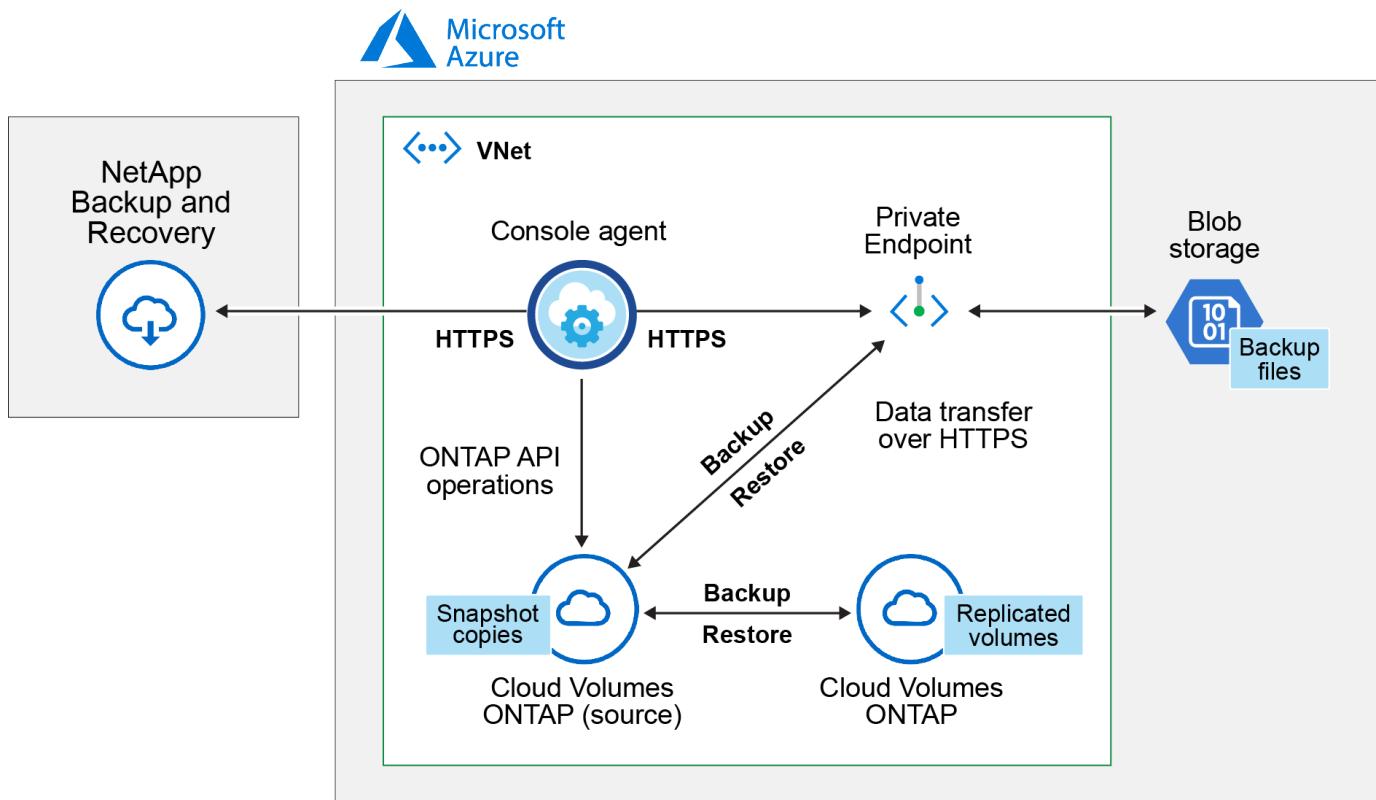
To switch to and from NetApp Backup and Recovery workloads, refer to [Switch to different NetApp Backup and Recovery workloads](#).

Verify support for your configuration

Read the following requirements to make sure that you have a supported configuration before you start backing up volumes to Azure Blob storage.

The following image shows each component and the connections that you need to prepare between them.

Optionally, you can connect to a secondary ONTAP system for replicated volumes using the public or private connection as well.



Supported ONTAP versions

Minimum of ONTAP 9.8; ONTAP 9.8P13 and later is recommended.

Supported Azure regions

NetApp Backup and Recovery is supported in all Azure regions, including Azure Government regions.

By default, NetApp Backup and Recovery provisions the Blob container with Local redundancy (LRS) for cost optimization. You can change this setting to Zone redundancy (ZRS) after NetApp Backup and Recovery has been activated if you want to make sure your data is replicated between different zones. See the Microsoft instructions for [changing how your storage account is replicated](#).

Required setup for creating backups in a different Azure subscription

By default, backups are created using the same subscription as the one used for your Cloud Volumes ONTAP system.

Verify license requirements

For NetApp Backup and Recovery PAYGO licensing, a subscription through the Azure Marketplace is required before you enable NetApp Backup and Recovery. Billing for NetApp Backup and Recovery is done through this subscription. [You can subscribe from the Details & Credentials page of the system wizard](#).

For NetApp Backup and Recovery BYOL licensing, you need the serial number from NetApp that enables you to use the service for the duration and capacity of the license. [Learn how to manage your BYOL licenses](#). You must use a BYOL license when the Console agent and Cloud Volumes ONTAP system are deployed in a dark site ("private mode").

And you need to have a Microsoft Azure subscription for the storage space where your backups will be located.

Prepare your Console agent

The Console agent can be installed in an Azure region with full or limited internet access ("standard" or "restricted" mode). [See NetApp Console deployment modes for details](#).

- [Learn about Console agents](#)
- [Deploy a Console agent in Azure in standard mode \(full internet access\)](#)
- [Install the Console agent in restricted mode \(limited outbound access\)](#)

Verify or add permissions to the Console agent

To use the NetApp Backup and Recovery Search & Restore functionality, you need to have specific permissions in the role for the Console agent so that it can access the Azure Synapse Workspace and Data Lake Storage Account. See the permissions below, and follow the steps if you need to modify the policy.

Before you start

- You must register the Azure Synapse Analytics Resource Provider (called "Microsoft.Synapse") with your Subscription. [See how to register this resource provider for your subscription](#). You must be the Subscription **Owner** or **Contributor** to register the resource provider.
- Port 1433 must be open for communication between the Console agent and the Azure Synapse SQL services.

Steps

1. Identify the role assigned to the Console agent virtual machine:
 - a. In the Azure portal, open the virtual machines service.
 - b. Select the Console agent virtual machine.
 - c. Under Settings, select **Identity**.
 - d. Select **Azure role assignments**.
 - e. Make note of the custom role assigned to the Console agent virtual machine.
2. Update the custom role:
 - a. In the Azure portal, open your Azure subscription.
 - b. Select **Access control (IAM) > Roles**.
 - c. Select the ellipsis (...) for the custom role and then select **Edit**.
 - d. Select **JSON** and add the following permissions:

```
"Microsoft.Storage/storageAccounts/listkeys/action",
"Microsoft.Storage/storageAccounts/read",
"Microsoft.Storage/storageAccounts/write",
"Microsoft.Storage/storageAccounts/blobServices/containers/read",
"Microsoft.Storage/storageAccounts/listAccountSas/action",
"Microsoft.KeyVault/vaults/read",
"Microsoft.KeyVault/vaults/accessPolicies/write",
"Microsoft.Network/networkInterfaces/read",
"Microsoft.Resources/subscriptions/locations/read",
"Microsoft.Network/virtualNetworks/read",
"Microsoft.Network/virtualNetworks/subnets/read",
"Microsoft.Resources/subscriptions/resourceGroups/read",
"Microsoft.Resources/subscriptions/resourceGroups/resources/read",
",
"Microsoft.Resources/subscriptions/resourceGroups/write",
"Microsoft.Authorization/locks/*",
"Microsoft.Network/privateEndpoints/write",
"Microsoft.Network/privateEndpoints/read",
"Microsoft.Network/privateDnsZones/virtualNetworkLinks/write",
"Microsoft.Network/virtualNetworks/join/action",
"Microsoft.Network/privateDnsZones/A/write",
"Microsoft.Network/privateDnsZones/read",
"Microsoft.Network/privateDnsZones/virtualNetworkLinks/read",
"Microsoft.Network/networkInterfaces/delete",
"Microsoft.Network/networkSecurityGroups/delete",
"Microsoft.Resources/deployments/delete",
"Microsoft.ManagedIdentity/userAssignedIdentities/assign/action",
"Microsoft.Synapse/workspaces/write",
"Microsoft.Synapse/workspaces/read",
"Microsoft.Synapse/workspaces/delete",
"Microsoft.Synapse/register/action",
"Microsoft.Synapse/checkNameAvailability/action",
"Microsoft.Synapse/workspaces/operationStatuses/read",
"Microsoft.Synapse/workspaces/firewallRules/read",
"Microsoft.Synapse/workspaces/replaceAllIpFirewallRules/action",
"Microsoft.Synapse/workspaces/operationResults/read",
"Microsoft.Synapse/workspaces/privateEndpointConnectionsApproval/
action"
```

[View the full JSON format for the policy](#)

- e. Select **Review + update** and then select **Update**.

Required information for using customer-managed keys for data encryption

You can use your own customer-managed keys for data encryption in the activation wizard instead of using the default Microsoft-managed encryption keys. In this case, you will need to have the Azure Subscription, Key Vault name, and the Key. [See how to use your own keys](#).

NetApp Backup and Recovery supports *Azure access policies*, the *Azure role-based access control* (Azure RBAC) permission model and the *Managed Hardware Security Model* (HSM) (refer to [What is Azure Key Vault Managed HSM?](#)).

Create your Azure Blob storage account

By default, the service creates storage accounts for you. If you want to use your own storage accounts, you can create them before you start the backup activation wizard and then select those storage accounts in the wizard.

[Learn more about creating your own storage accounts.](#)

Verify ONTAP networking requirements for replicating volumes

Unresolved directive in prev-ontap-backup-cvo-azure.adoc - include:.../_include/backup-onprem-prepare-source-destination-systems-for-repl.adoc[]

- To replicate data between two Cloud Volumes ONTAP systems in different subnets, the subnets must be routed together (this is the default setting).

Enable NetApp Backup and Recovery on Cloud Volumes ONTAP

Enabling NetApp Backup and Recovery is easy. The steps differ slightly depending on whether you have an existing Cloud Volumes ONTAP system or a new one.

Enable NetApp Backup and Recovery on a new system

NetApp Backup and Recovery is enabled by default in the system wizard. Be sure to keep the option enabled.

See [Launching Cloud Volumes ONTAP in Azure](#) for requirements and details for creating your Cloud Volumes ONTAP system.



If you want to pick the name of the resource group, **disable** NetApp Backup and Recovery when deploying Cloud Volumes ONTAP.

Steps

1. From the Console **Systems** page, select **Add system**, choose the cloud provider, and select **Add New**. Select **Create Cloud Volumes ONTAP**.
2. Select **Microsoft Azure** as the cloud provider and then choose a single node or HA system.
3. In the Define Azure Credentials page, enter the credentials name, client ID, client secret, and directory ID, and select **Continue**.
4. Fill out the Details & Credentials page and be sure that an Azure Marketplace subscription is in place, and select **Continue**.
5. On the Services page, leave the service enabled and select **Continue**.
6. Complete the pages in the wizard to deploy the system.

Result

NetApp Backup and Recovery is enabled on the system. After you've created volumes on these Cloud Volumes ONTAP systems, launch NetApp Backup and Recovery and [activate backup on each volume that you want to protect](#).

Enable NetApp Backup and Recovery on an existing system

Enable NetApp Backup and Recovery at any time directly from the system.

Steps

1. From the Console **Systems** page, select the system and select **Enable** next to Backup and Recovery in the right-panel.

If the Azure Blob destination for your backups exists as a system on the Console **Systems** page, you can drag the cluster onto the Azure Blob system to initiate the setup wizard.

2. Complete the pages in the wizard to deploy NetApp Backup and Recovery.
3. When you want to initiate backups, continue with [Activate backups on your ONTAP volumes](#).

Activate backups on your ONTAP volumes

Activate backups at any time directly from your on-premises system.

A wizard takes you through the following major steps:

- [Select the volumes that you want to back up](#)
- [Define the backup strategy](#)
- [Review your selections](#)

You can also [Show the API commands](#) at the review step, so you can copy the code to automate backup activation for future systems.

Start the wizard

Steps

1. Access the Activate backup and recovery wizard using one of the following ways:

- From the Console **Systems** page, select the system and select **Enable > Backup Volumes** next to Backup and Recovery in the right-panel.

If the Azure destination for your backups exists as a system on the **Systems** page, you can drag the ONTAP cluster onto the Azure Blob object storage.

- Select **Volumes** in the Backup and Recovery bar. From the Volumes tab, select the **Actions ...** icon and select **Activate Backup** for a single volume (that does not already have replication or backup to object storage already enabled).

The Introduction page of the wizard shows the protection options including local Snapshots, replication, and backups. If you did the second option in this step, the Define Backup Strategy page appears with one volume selected.

2. Continue with the following options:

- If you already have a Console agent, you're all set. Just select **Next**.

- If you don't already have a Console agent, the **Add a Console agent** option appears. Refer to [Prepare your Console agent](#).

Select the volumes that you want to back up

Choose the volumes you want to protect. A protected volume is one that has one or more of the following: Snapshot policy, replication policy, backup-to-object policy.

You can choose to protect FlexVol or FlexGroup volumes; however, you cannot select a mix of these volumes when activating backup for a system. See how to [activate backup for additional volumes in the system](#) (FlexVol or FlexGroup) after you have configured backup for the initial volumes.



- You can activate a backup only on a single FlexGroup volume at a time.
- The volumes you select must have the same SnapLock setting. All volumes must have SnapLock Enterprise enabled or have SnapLock disabled.

Steps

If the volumes you choose already have snapshot or replication policies applied, then the policies you select later will overwrite these existing policies.

1. In the Select Volumes page, select the volume or volumes you want to protect.
 - Optionally, filter the rows to show only volumes with certain volume types, styles, and more to make the selection easier.
 - After you select the first volume, then you can select all FlexVol volumes. (FlexGroup volumes can be selected one at a time only.) To back up all existing FlexVol volumes, check one volume first and then check the box in the title row.
 - To back up individual volumes, check the box for each volume.
2. Select **Next**.

Define the backup strategy

Defining the backup strategy involves setting the following options:

- Whether you want one or all of the backup options: local snapshots, replication, and backup to object storage
- Architecture
- Local snapshot policy
- Replication target and policy



If the volumes you choose have different snapshot and replication policies than the policies you select in this step, the existing policies will be overwritten.

- Backup to object storage information (provider, encryption, networking, backup policy, and export options).

Steps

1. In the Define backup strategy page, choose one or all of the following. All three are selected by default:
 - **Local Snapshots:** If you are performing replication or back up to object storage, local snapshots must be created.

- **Replication:** Creates replicated volumes on another ONTAP storage system.
- **Backup:** Backs up volumes to object storage.

2. **Architecture:** If you chose replication and backup, choose one of the following flows of information:

- **Cascading:** Information flows from the primary storage system to the secondary, and from secondary to object storage.
- **Fan out:** Information flows from the primary storage system to the secondary *and* from the primary to object storage.

For details about these architectures, refer to [Plan your protection journey](#).

3. **Local Snapshot:** Choose an existing snapshot policy or create one.



To create a custom policy before activating the snapshot, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to five schedules, typically of different frequencies.
- Select **Create**.

4. **Replication:** Set the following options:

- **Replication target:** Select the destination system and SVM. Optionally, select the destination aggregate or aggregates and prefix or suffix that will be added to the replicated volume name.
- **Replication policy:** Choose an existing replication policy or create one.



To create a custom policy before activating the replication, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to five schedules, typically of different frequencies.
- Select **Create**.

5. **Back up to Object:** If you selected **Backup**, set the following options:

- **Provider:** Select **Microsoft Azure**.
- **Provider settings:** Enter the provider details.

Enter the region where the backups will be stored. This can be a different region than where the Cloud Volumes ONTAP system resides.

Either create a new storage account or select an existing one.

Enter the Azure subscription used to store the backups. This can be a different subscription than where the Cloud Volumes ONTAP system resides.

Either create your own resource group that manages the Blob container or select the resource group type and group.



If you want to protect your backup files from being modified or deleted, ensure that the storage account was created with immutable storage enabled using a 30-day retention period.

- **Encryption key:** If you created a new Azure storage account, enter encryption key information given to you from the provider. Choose whether you'll use the default Azure encryption keys, or choose your own customer-managed keys from your Azure account, to manage encryption of your data.

If you choose to use your own customer-managed keys, enter the key vault and key information. [Learn how to use your own keys](#).



If you chose an existing Microsoft storage account, encryption information is already available, so you don't need to enter it now.

- **Networking:** Choose the IPspace, and whether you'll be using a Private Endpoint. Private Endpoint is disabled by default.
 - a. The IPspace in the ONTAP cluster where the volumes you want to back up reside. The intercluster LIFs for this IPspace must have outbound internet access.
 - b. Optionally, choose whether you'll use an Azure private endpoint that you have previously configured. [Learn about using an Azure private endpoint](#).
- **Backup policy:** Select an existing backup-to-object storage policy.



To create a custom policy before activating the backup, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- For backup-to-object policies, set the DataLock and Ransomware Resilience settings. For details on DataLock and Ransomware Resilience, refer to [Backup-to-object policy settings](#).
- Select up to five schedules, typically of different frequencies.
- Select **Create**.

◦ **Export existing snapshots to object storage as backup copies:** If there are any local snapshots for volumes in this system that match the backup schedule label you just selected for this system (for example, daily, weekly, etc.), this additional prompt is displayed. Check this box to have all historic Snapshots copied to object storage as backup files to ensure the most complete protection for your volumes.

6. Select **Next**.

Review your selections

This is the chance to review your selections and make adjustments, if necessary.

Steps

1. In the Review page, review your selections.
2. Optionally check the box to **Automatically synchronize the Snapshot policy labels with the replication and backup policy labels**. This creates Snapshots with a label that matches the labels in the replication and backup policies.
3. Select **Activate Backup**.

Result

NetApp Backup and Recovery starts taking the initial backups of your volumes. The baseline transfer of the replicated volume and the backup file includes a full copy of the primary storage system data. Subsequent transfers contain differential copies of the primary storage data contained in snapshots.

A replicated volume is created in the destination cluster that will be synchronized with the primary volume.

A Blob storage container is created in the resource group you entered, and the backup files are stored there.

By default, NetApp Backup and Recovery provisions the Blob container with Local redundancy (LRS) for cost optimization. You can change this setting to Zone redundancy (ZRS) if you want to make sure your data is replicated between different zones. See the Microsoft instructions for [changing how your storage account is replicated](#).

The Volume Backup Dashboard is displayed so you can monitor the state of the backups.

You can also monitor the status of backup and restore jobs using the [Job Monitoring page](#).

Show the API commands

You might want to display and optionally copy the API commands used in the Activate backup and recovery wizard. You might want to do this to automate backup activation in future systems.

Steps

1. From the Activate backup and recovery wizard, select **View API request**.
2. To copy the commands to the clipboard, select the **Copy** icon.

What's next?

- You can [manage your backup files and backup policies](#). This includes starting and stopping backups, deleting backups, adding and changing the backup schedule, and more.
- You can [manage cluster-level backup settings](#). This includes changing the storage keys ONTAP uses to access cloud storage, changing the network bandwidth available to upload backups to object storage, changing the automatic backup setting for future volumes, and more.
- You can also [restore volumes, folders, or individual files from a backup file](#) to a Cloud Volumes ONTAP system in AWS, or to an on-premises ONTAP system.

Back up Cloud Volumes ONTAP data to Google Cloud Storage with NetApp Backup and Recovery

Complete a few steps in NetApp Backup and Recovery to get started backing up volume data from your Cloud Volumes ONTAP systems to Google Cloud Storage.



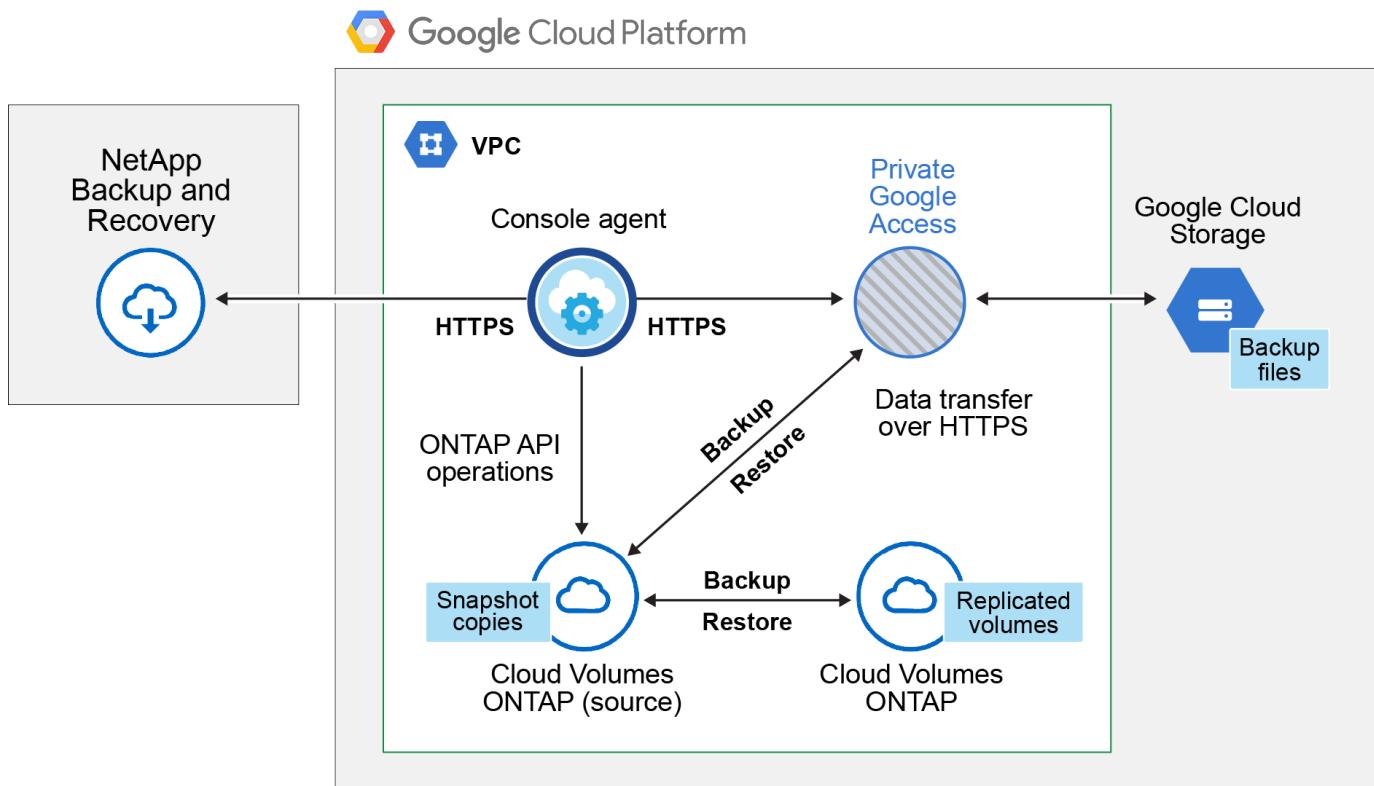
To switch to and from NetApp Backup and Recovery workloads, refer to [Switch to different NetApp Backup and Recovery workloads](#).

Verify support for your configuration

Read the following requirements to make sure that you have a supported configuration before you start backing up volumes to Google Cloud Storage.

The following image shows each component and the connections that you need to prepare between them.

Optionally, you can connect to a secondary ONTAP system for replicated volumes using the public or private connection as well.



Supported ONTAP versions

Minimum of ONTAP 9.8; ONTAP 9.8P13 and later is recommended.

Supported GCP regions

NetApp Backup and Recovery is supported in all GCP regions.

GCP Service Account

You need to have a service account in your Google Cloud Project that has the custom role. [Learn how to create a service account](#).



The Storage Admin role is no longer required for the service account that enables NetApp Backup and Recovery to access Google Cloud Storage buckets.

Verify license requirements

For NetApp Backup and Recovery PAYGO licensing, a Console subscription is available in the Google Marketplace that enables deployments of Cloud Volumes ONTAP and NetApp Backup and Recovery. You need to [subscribe to this Console subscription](#) before you enable NetApp Backup and Recovery. Billing for NetApp Backup and Recovery is done through this subscription. [You can subscribe from the Details & Credentials page of the system wizard](#).

For NetApp Backup and Recovery BYOL licensing, you need the serial number from NetApp that enables you to use the service for the duration and capacity of the license. [Learn how to manage your BYOL licenses](#).

And you need to have a Google subscription for the storage space where your backups will be located.

Prepare your Console agent

The Console agent must be installed in a Google region with internet access.

- [Learn about Console agents](#)
- [Deploy a Console agent in Google Cloud](#)

Verify or add permissions to the Console agent

To use the NetApp Backup and Recovery "Search & Restore" functionality, you need to have specific permissions in the role for the Console agent so that it can access the Google Cloud BigQuery service. See the permissions below, and follow the steps if you need to modify the policy.

Steps

1. In the [Google Cloud Console](#), go to the **Roles** page.
2. Using the drop-down list at the top of the page, select the project or organization that contains the role that you want to edit.
3. Select a custom role.
4. Select **Edit Role** to update the role's permissions.
5. Select **Add Permissions** to add the following new permissions to the role.

```
bigquery.jobs.get
bigquery.jobs.list
bigquery.jobs.listAll
bigquery.datasets.create
bigquery.datasets.get
bigquery.jobs.create
bigquery.tables.get
bigquery.tables.getData
bigquery.tables.list
bigquery.tables.create
```

6. Select **Update** to save the edited role.

Required information for using customer-managed encryption keys (CMEK)

You can use your own customer-managed keys for data encryption instead of using the default Google-managed encryption keys. Both cross-region and cross-project keys are supported, so you can choose a project for a bucket that is different than the project of the CMEK key. If you're planning to use your own customer-managed keys:

- You'll need to have the Key Ring and the Key Name so you can add this information in the activation wizard. [Learn more about customer-managed encryption keys](#).
- You'll need to verify that these required permissions are included in the role for the Console agent:

```
cloudkms.cryptoKeys.get
cloudkms.cryptoKeys.getIamPolicy
cloudkms.cryptoKeys.list
cloudkms.cryptoKeys.setIamPolicy
cloudkms.keyRings.get
cloudkms.keyRings.getIamPolicy
cloudkms.keyRings.list
cloudkms.keyRings.setIamPolicy
```

- You'll need to verify that the Google "Cloud Key Management Service (KMS)" API is enabled in your project. See the [Google Cloud documentation: Enabling APIs](#) for details.

CMEK considerations:

- Both HSM (hardware-backed) and software-generated keys are supported.
- Both newly created or imported Cloud KMS keys are supported.
- Only regional keys are supported; global keys are not supported.
- Currently, only the "Symmetric encrypt/decrypt" purpose is supported.
- The service agent associated with the Storage Account is assigned the "CryptoKey Encrypter/Decrypter (roles/cloudkms.cryptoKeyEncrypterDecrypter)" IAM role by NetApp Backup and Recovery.

Create your own buckets

By default, the service creates buckets for you. If you want to use your own buckets, you can create them before you start the backup activation wizard and then select those buckets in the wizard.

[Learn more about creating your own buckets.](#)

Verify ONTAP networking requirements for replicating volumes

Unresolved directive in prev-ontap-backup-cvo-gcp.adoc - include:::_include/backup-onprem-prepare-source-destination-systems-for-repl.adoc[]

- To replicate data between two Cloud Volumes ONTAP systems in different subnets, the subnets must be routed together (this is the default setting).

Enable NetApp Backup and Recovery on Cloud Volumes ONTAP

Enabling NetApp Backup and Recovery steps differ slightly depending on whether you have an existing Cloud Volumes ONTAP system or a new one.

Enable NetApp Backup and Recovery on a new system

NetApp Backup and Recovery can be enabled when you complete the system wizard to create a new Cloud Volumes ONTAP system.

You must have a Service Account already configured. If you don't select a service account when you create the Cloud Volumes ONTAP system, then you'll need to turn off the system and add the service account to Cloud Volumes ONTAP from the GCP console.

See [Launching Cloud Volumes ONTAP in GCP](#) for requirements and details for creating your Cloud Volumes ONTAP system.

Steps

1. From the Console **Systems** page, select **Add system**, choose the cloud provider, and select **Add New**. Select **Create Cloud Volumes ONTAP**.
2. **Choose a Location:** Select **Google Cloud Platform**.
3. **Choose Type:** Select **Cloud Volumes ONTAP** (either single-node or high-availability).
4. **Details & Credentials:** Enter the following information:
 - a. Click **Edit Project** and select a new project if the one you want to use is different than the default Project (where the Console agent resides).
 - b. Specify the cluster name.
 - c. Enable the **Service Account** switch and select the Service Account that has the predefined Storage Admin role. This is required to enable backups and tiering.
 - d. Specify the credentials.

Make sure that a GCP Marketplace subscription is in place.

5. **Services:** Leave NetApp Backup and Recovery enabled and click **Continue**.
6. Complete the pages in the wizard to deploy the system as described in [Launching Cloud Volumes ONTAP in GCP](#).

Result

NetApp Backup and Recovery is enabled on the system. After you've created volumes on these Cloud Volumes ONTAP systems, launch NetApp Backup and Recovery and [activate backup on each volume that you want to protect](#).

Enable NetApp Backup and Recovery on an existing system

You can enable NetApp Backup and Recovery at any time directly from the system.

Steps

1. From the Console **Systems** page, select the system and select **Enable** next to Backup and Recovery in the right-panel.

If the Google Cloud Storage destination for your backups exists as a system on the Console **Systems** page, you can drag the cluster onto the Google Cloud Storage system to initiate the setup wizard.

Prepare Google Cloud Storage as your backup target

Preparing Google Cloud Storage as your backup target involves the following steps:

- Set up permissions.
- (Optional) Create your own buckets. (The service will create buckets for you if you want.)
- (Optional) Set up customer-managed keys for data encryption

Set up permissions

You need to provide storage access keys for a service account that has specific permissions using a custom

role. A service account enables NetApp Backup and Recovery to authenticate and access Cloud Storage buckets used to store backups. The keys are required so that Google Cloud Storage knows who is making the request.

Steps

1. In the [Google Cloud Console](#), go to the **Roles** page.
2. [Create a new role](#) with the following permissions:

```
storage.buckets.create
storage.buckets.delete
storage.buckets.get
storage.buckets.list
storage.buckets.update
storage.buckets.getIamPolicy
storage.multipartUploads.create
storage.objects.create
storage.objects.delete
storage.objects.get
storage.objects.list
storage.objects.update
```

3. In the Google Cloud console, [go to the Service accounts page](#).
4. Select your Cloud project.
5. Select **Create service account** and provide the required information:
 - Service account details:** Enter a name and description.
 - Grant this service account access to project:** Select the custom role that you just created.
 - Select **Done**.
6. Go to [GCP Storage Settings](#) and create access keys for the service account:
 - Select a project, and select **Interoperability**. If you haven't already done so, select **Enable interoperability access**.
 - Under **Access keys for service accounts**, select **Create a key for a service account**, select the service account that you just created, and click **Create Key**.

You'll need to enter the keys in NetApp Backup and Recovery later when you configure the backup service.

Create your own buckets

By default, the service creates buckets for you. Or, if you want to use your own buckets, you can create them before you start the backup activation wizard and then select those buckets in the wizard.

[Learn more about creating your own buckets](#).

Set up customer-managed encryption keys (CMEK) for data encryption

You can use your own customer-managed keys for data encryption instead of using the default Google-

managed encryption keys. Both cross-region and cross-project keys are supported, so you can choose a project for a bucket that is different than the project of the CMEK key.

If you're planning to use your own customer-managed keys:

- You'll need to have the Key Ring and the Key Name so you can add this information in the activation wizard. [Learn more about customer-managed encryption keys](#).
- You'll need to verify that these required permissions are included in the role for the Console agent:

```
cloudkms.cryptoKeys.get
cloudkms.cryptoKeys.getIamPolicy
cloudkms.cryptoKeys.list
cloudkms.cryptoKeys.setIamPolicy
cloudkms.keyRings.get
cloudkms.keyRings.getIamPolicy
cloudkms.keyRings.list
cloudkms.keyRings.setIamPolicy
```

- You'll need to verify that the Google "Cloud Key Management Service (KMS)" API is enabled in your project. See the [Google Cloud documentation: Enabling APIs](#) for details.

CMEK considerations:

- Both HSM (Hardware-backed) and Software-generated keys are supported.
- Both newly created or imported Cloud KMS keys are supported.
- Only regional keys are supported, global keys are not supported.
- Currently, only the "Symmetric encrypt/decrypt" purpose is supported.
- The service agent associated with the Storage Account is assigned the "CryptoKey Encrypter/Decrypter (roles/cloudkms.cryptoKeyEncrypterDecrypter)" IAM role by NetApp Backup and Recovery.

Activate backups on your ONTAP volumes

Activate backups at any time directly from your on-premises system.

A wizard takes you through the following major steps:

- [Select the volumes that you want to back up](#)
- [Define the backup strategy](#)
- [Review your selections](#)

You can also [Show the API commands](#) at the review step, so you can copy the code to automate backup activation for future systems.

Start the wizard

Steps

1. Access the Activate backup and recovery wizard using one of the following ways:

- From the **Console Systems** page*, select the system and select **Enable > Backup Volumes** next to Backup and Recovery in the right-panel.

If the GCP destination for your backups exists as a system on the **Console Systems** page, you can drag the ONTAP cluster onto the GCP object storage.

- Select **Volumes** in the Backup and Recovery bar. From the Volumes tab, select the **Actions**  icon and select **Activate Backup** for a single volume (that does not already have replication or backup to object storage already enabled).

The Introduction page of the wizard shows the protection options including local Snapshots, replication, and backups. If you did the second option in this step, the Define Backup Strategy page appears with one volume selected.

2. Continue with the following options:

- If you already have a Console agent, you're all set. Just select **Next**.
- If you don't already have a Console agent, the **Add a Console agent** option appears. Refer to [Prepare your Console agent](#).

Select the volumes that you want to back up

Choose the volumes you want to protect. A protected volume is one that has one or more of the following: Snapshot policy, replication policy, backup to object policy.

You can choose to protect FlexVol or FlexGroup volumes; however, you cannot select a mix of these volumes when activating backup for a system. See how to [activate backup for additional volumes in the system](#) (FlexVol or FlexGroup) after you have configured backup for the initial volumes.



- You can activate a backup only on a single FlexGroup volume at a time.
- The volumes you select must have the same SnapLock setting. All volumes must have SnapLock Enterprise enabled or have SnapLock disabled.

Steps

Note that if the volumes you choose already have Snapshot or replication policies applied, then the policies you select later will overwrite these existing policies.

1. In the Select Volumes page, select the volume or volumes you want to protect.
 - Optionally, filter the rows to show only volumes with certain volume types, styles, and more to make the selection easier.
 - After you select the first volume, then you can select all FlexVol volumes (FlexGroup volumes can be selected one at a time only). To back up all existing FlexVol volumes, check one volume first and then check the box in the title row.
 - To back up individual volumes, check the box for each volume.
2. Select **Next**.

Define the backup strategy

Defining the backup strategy involves setting the following options:

- Whether you want one or all of the backup options: local Snapshots, replication, and backup to object storage

- Architecture
- Local snapshot policy
- Replication target and policy



If the volumes you choose have different Snapshot and replication policies than the policies you select in this step, the existing policies will be overwritten.

- Backup to object storage information (provider, encryption, networking, backup policy, and export options).

Steps

1. In the Define backup strategy page, choose one or all of the following. All three are selected by default:
 - **Local Snapshots:** If you are performing replication or back up to object storage, local snapshots must be created.
 - **Replication:** Creates replicated volumes on another ONTAP storage system.
 - **Backup:** Backs up volumes to object storage.
2. **Architecture:** If you chose replication and backup, choose one of the following flows of information:
 - **Cascading:** Information flows from the primary storage system to the secondary, and from secondary to object storage.
 - **Fan out:** Information flows from the primary storage system to the secondary *and* from the primary to object storage.

For details about these architectures, refer to [Plan your protection journey](#).

3. **Local Snapshot:** Choose an existing snapshot policy or create one.



To create a custom policy before activating the backup, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to five schedules, typically of different frequencies.
- For backup-to-object policies, configure Datalock and Ransomware Resilience. For details on Datalock and Ransomware Resilience, refer to [Backup-to-object policy settings](#).
- Select **Create**.

4. **Replication:** Set the following options:

- **Replication target:** Select the destination system and SVM. Optionally, select the destination aggregate or aggregates and prefix or suffix that will be added to the replicated volume name.
- **Replication policy:** Choose an existing replication policy or create one.



To create a custom policy before activating the replication, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to five schedules, typically of different frequencies.

- Select **Create**.

5. Back up to Object:

If you selected **Backup**, set the following options:

- **Provider:** Select **Google Cloud**.
- **Provider settings:** Enter the provider details and region where the backups will be stored.

Either create a new bucket or select an existing one.

- **Encryption key:** If you created a new Google bucket, enter encryption key information given to you from the provider. Choose whether you'll use the default Google Cloud encryption keys, or choose your own customer-managed keys from your Google account, to manage encryption of your data.

If you choose to use your own customer-managed keys, enter the key vault and key information.



If you chose an existing Google Cloud bucket, encryption information is already available, so you don't need to enter it now.

- **Backup policy:** Select an existing backup-to-object storage policy or create one.



To create a custom policy before activating the backup, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to five schedules, typically of different frequencies.
- Select **Create**.
- **Export existing snapshots to object storage as backup copies:** If there are any local snapshots for volumes in this system that match the backup schedule label you just selected for this system (for example, daily, weekly, etc.), this additional prompt is displayed. Check this box to have all historic Snapshots copied to object storage as backup files to ensure the most complete protection for your volumes.

6. Select **Next**.

Review your selections

This is the chance to review your selections and make adjustments, if necessary.

Steps

1. In the Review page, review your selections.
2. Optionally check the box to **Automatically synchronize the Snapshot policy labels with the replication and backup policy labels**. This creates Snapshots with a label that matches the labels in the replication and backup policies.
3. Select **Activate Backup**.

Result

NetApp Backup and Recovery starts taking the initial backups of your volumes. The baseline transfer of the replicated volume and the backup file includes a full copy of the primary storage system data. Subsequent transfers contain differential copies of the primary storage system data contained in snapshots.

A replicated volume is created in the destination cluster that will be synchronized with the primary storage

system volume.

A Google Cloud Storage bucket is created in the service account indicated by the Google access key and secret key you entered, and the backup files are stored there.

Backups are associated with the *Standard* storage class by default. You can use the lower cost *Nearline*, *Coldline*, or *Archive* storage classes. However, you configure the storage class through Google, not through the NetApp Backup and Recovery UI. See the Google topic [Changing the default storage class of a bucket](#) for details.

The Volume Backup Dashboard is displayed so you can monitor the state of the backups.

You can also monitor the status of backup and restore jobs using the [Job Monitoring page](#).

Show the API commands

You might want to display and optionally copy the API commands used in the Activate backup and recovery wizard. You might want to do this to automate backup activation in future systems.

Steps

1. From the Activate backup and recovery wizard, select **View API request**.
2. To copy the commands to the clipboard, select the **Copy** icon.

What's next?

- You can [manage your backup files and backup policies](#). This includes starting and stopping backups, deleting backups, adding and changing the backup schedule, and more.
- You can [manage cluster-level backup settings](#). This includes changing the storage keys ONTAP uses to access cloud storage, changing the network bandwidth available to upload backups to object storage, changing the automatic backup setting for future volumes, and more.
- You can also [restore volumes, folders, or individual files from a backup file](#) to a Cloud Volumes ONTAP system in AWS, or to an on-premises ONTAP system.

Back up on-premises ONTAP data to Amazon S3 with NetApp Backup and Recovery

Complete a few steps in NetApp Backup and Recovery to get started backing up volume data from your on-premises ONTAP systems to a secondary storage system and to Amazon S3 cloud storage.



"On-premises ONTAP systems" include FAS, AFF, and ONTAP Select systems.



To switch to and from NetApp Backup and Recovery workloads, refer to [Switch to different NetApp Backup and Recovery workloads](#).

Identify the connection method

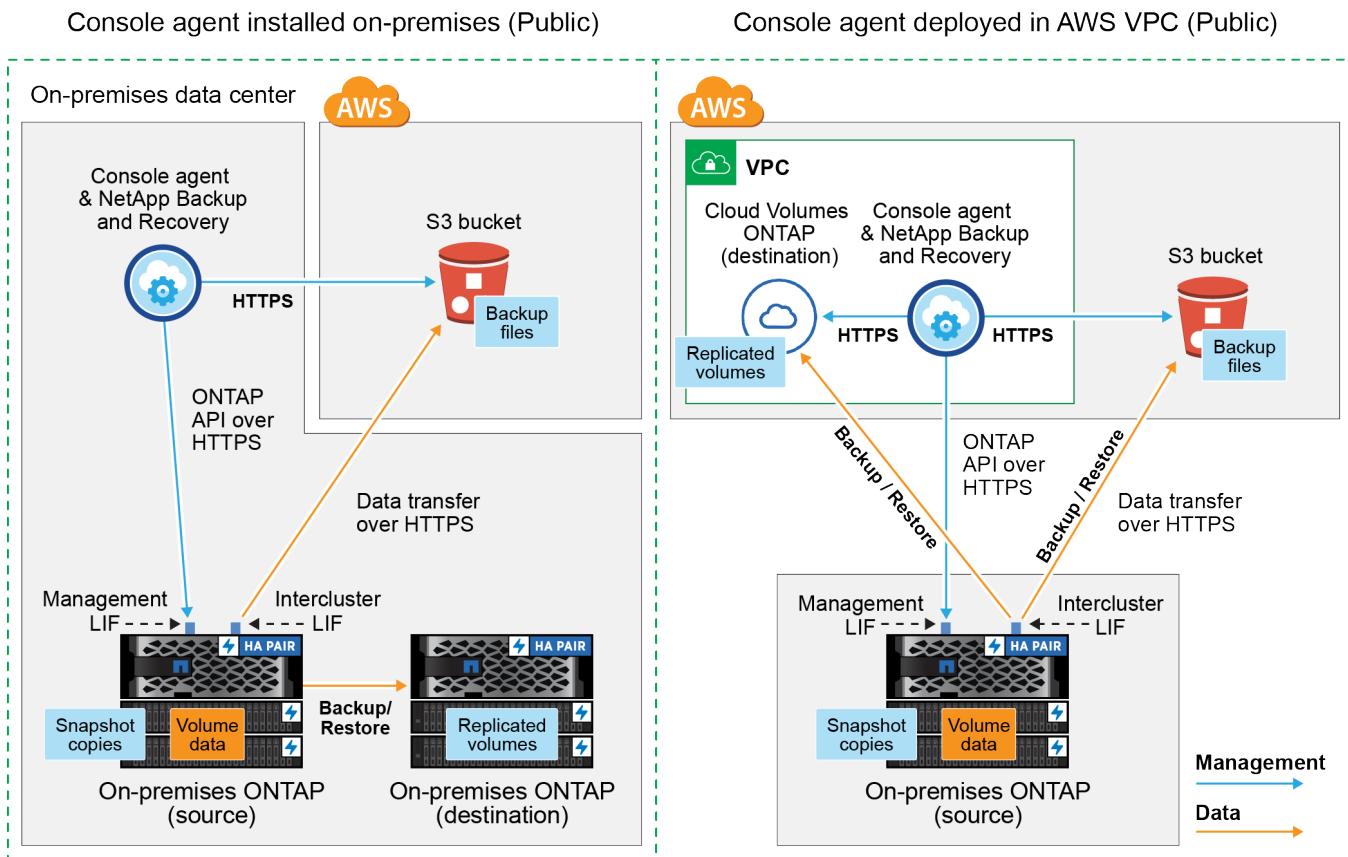
Choose which of the two connection methods you will use when configuring backups from on-premises ONTAP systems to AWS S3.

- **Public connection** - Directly connect the ONTAP system to AWS S3 using a public S3 endpoint.

- **Private connection** - Use a VPN or AWS Direct Connect and route traffic through a VPC Endpoint interface that uses a private IP address.

Optionally, you can connect to a secondary ONTAP system for replicated volumes using the public or private connection as well.

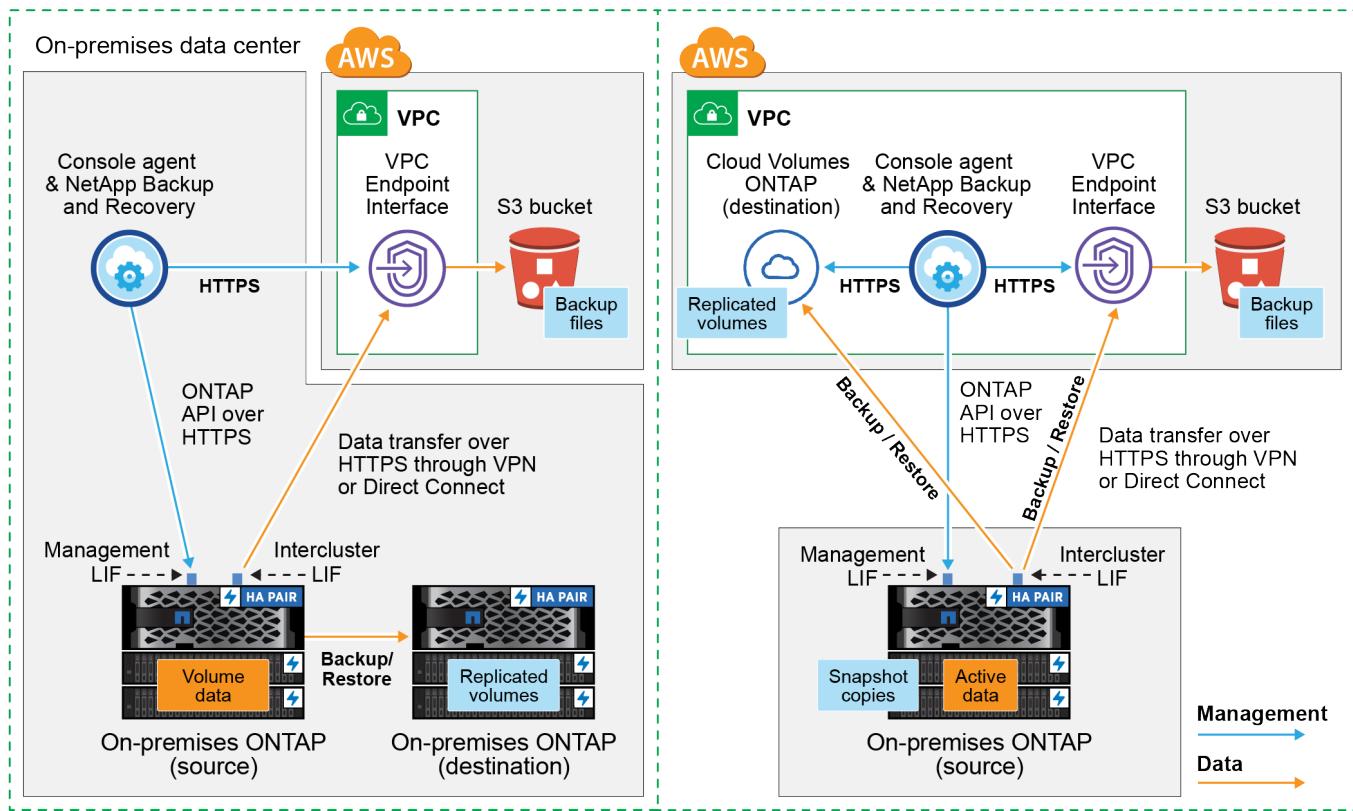
The following diagram shows the **public connection** method and the connections that you need to prepare between the components. You can use a Console agent that you've installed on your premises, or a Console agent that you've deployed in the AWS VPC.



The following diagram shows the **private connection** method and the connections that you need to prepare between the components. You can use a Console agent that you've installed on your premises, or a Console agent that you've deployed in the AWS VPC.

Console agent installed on-premises (Private)

Console agent deployed in AWS VPC (Private)



Prepare your Console agent

The Console agent is the main software for NetApp Console functionality. A Console agent is required to back up and restore your ONTAP data.

Create or switch Console agents

If you already have a Console agent deployed in your AWS VPC or on your premises, then you're all set.

If not, then you'll need to create a Console agent in one of those locations to back up ONTAP data to AWS S3 storage. You can't use a Console agent that's deployed in another cloud provider.

- [Learn about Console agents](#)
- [Install a Console agent in AWS](#)
- [Install a Console agent in your premises](#)
- [Install a Console agent in an AWS GovCloud region](#)

NetApp Backup and Recovery is supported in GovCloud regions when the Console agent is deployed in the cloud - not when it's installed in your premises. Additionally, you must deploy the Console agent from the AWS Marketplace. You can't deploy the Console agent in a Government region from the NetApp Console SaaS website.

Prepare Console agent networking requirements

Ensure that the following networking requirements are met:

- Ensure that the network where the Console agent is installed enables the following connections:
 - An HTTPS connection over port 443 to NetApp Backup and Recovery and to your S3 object storage ([see the list of endpoints](#))
 - An HTTPS connection over port 443 to your ONTAP cluster management LIF
 - Additional inbound and outbound security group rules are required for AWS and AWS GovCloud deployments. See [Rules for the Console agent in AWS](#) for details.
- If you have a Direct Connect or VPN connection from your ONTAP cluster to the VPC, and you want communication between the Console agent and S3 to stay in your AWS internal network (a **private** connection), you'll need to enable a VPC Endpoint interface to S3. [Configure your system for a private connection using a VPC endpoint interface](#).

Verify license requirements

You'll need to verify license requirements for both AWS and the NetApp Console:

- Before you can activate NetApp Backup and Recovery for your cluster, you'll need to either subscribe to a pay-as-you-go (PAYGO) NetApp Console Marketplace offering from AWS, or purchase and activate a NetApp Backup and Recovery BYOL license from NetApp. These licenses are for your account and can be used across multiple systems.
 - For NetApp Backup and Recovery PAYGO licensing, you'll need a subscription to the [NetApp Console offering from the AWS Marketplace](#). Billing for NetApp Backup and Recovery is done through this subscription.
 - For NetApp Backup and Recovery BYOL licensing, you'll need the serial number from NetApp that enables you to use the service for the duration and capacity of the license.
- You need to have an AWS subscription for the object storage space where your backups will be located.

Supported regions

You can create backups from on-premises systems to Amazon S3 in all regions, including AWS GovCloud regions. You specify the region where backups will be stored when you set up the service.

Prepare your ONTAP clusters

Unresolved directive in prev-ontap-backup-onprem-aws.adoc - include:::_include/backup-onprem-prepare-onprem-ONTAP-cluster.adoc[]

Verify ONTAP networking requirements for backing up data to object storage

You must configure the following requirements on the system that connects to object storage.

- For a fan-out backup architecture, configure the following settings on the *primary* system.
- For a cascaded backup architecture, configure the following settings on the *secondary* system.

The following ONTAP cluster networking requirements are needed:

- The cluster requires an inbound HTTPS connection from the Console agent to the cluster management LIF.
- An intercluster LIF is required on each ONTAP node that hosts the volumes you want to back up. These intercluster LIFs must be able to access the object store.

The cluster initiates an outbound HTTPS connection over port 443 from the intercluster LIFs to Amazon S3 storage for backup and restore operations. ONTAP reads and writes data to and from object storage — the

object storage never initiates, it just responds.

- The intercluster LIFs must be associated with the *IPspace* that ONTAP should use to connect to object storage. [Learn more about IPspaces](#).

When you set up NetApp Backup and Recovery, you are prompted for the *IPspace* to use. You should choose the *IPspace* that these LIFs are associated with. That might be the "Default" *IPspace* or a custom *IPspace* that you created.

If you use are using a different *IPspace* than "Default", then you might need to create a static route to get access to the object storage.

All intercluster LIFs within the *IPspace* must have access to the object store. If you can't configure this for the current *IPspace*, then you'll need to create a dedicated *IPspace* where all intercluster LIFs have access to the object store.

- DNS servers must have been configured for the storage VM where the volumes are located. See how to [configure DNS services for the SVM](#).
- Update firewall rules, if necessary, to allow NetApp Backup and Recovery connections from ONTAP to object storage through port 443 and name resolution traffic from the storage VM to the DNS server over port 53 (TCP/UDP).
- If you are using a Private VPC Interface Endpoint in AWS for the S3 connection, then in order for HTTPS/443 to be used, you'll need to load the S3 endpoint certificate into the ONTAP cluster. [Configure your system for a private connection using a VPC endpoint interface](#).
- Ensure that your ONTAP cluster has permissions to access the S3 bucket.

Verify ONTAP networking requirements for replicating volumes

Unresolved directive in prev-ontap-backup-onprem-aws.adoc - include:::_include/backup-onprem-prepare-source-destination-systems-for-repl.adoc[]

Prepare Amazon S3 as your backup target

Preparing Amazon S3 as your backup target involves the following steps:

- Set up S3 permissions.
- (Optional) Create your own S3 buckets. (The service will create buckets for you if you want.)
- (Optional) Set up customer-managed AWS keys for data encryption.
- (Optional) Configure your system for a private connection using a VPC endpoint interface.

Set up S3 permissions

You'll need to configure two sets of permissions:

- Permissions for the Console agent to create and manage the S3 bucket.
- Permissions for the on-premises ONTAP cluster so it can read and write data to the S3 bucket.

Steps

1. Ensure that the Console agent has the required permissions. For details, see [NetApp Console policy permissions](#).



When creating backups in AWS China regions, you need to change the AWS Resource Name "arn" under all *Resource* sections in the IAM policies from "aws" to "aws-cn"; for example `arn:aws-cn:s3:::netapp-backup-*`.

2. When you activate the service, the Backup wizard will prompt you to enter an access key and secret key. These credentials are passed to the ONTAP cluster so that ONTAP can back up and restore data to the S3 bucket. For that, you'll need to create an IAM user with the following permissions.

Refer to the [AWS Documentation: Creating a Role to Delegate Permissions to an IAM User](#).

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Action": [  
        "s3:GetObject",  
        "s3:PutObject",  
        "s3:DeleteObject",  
        "s3>ListBucket",  
        "s3>ListAllMyBuckets",  
        "s3:GetBucketLocation",  
        "s3:PutEncryptionConfiguration"  
      ],  
      "Resource": "arn:aws:s3:::netapp-backup-*",  
      "Effect": "Allow",  
      "Sid": "backupPolicy"  
    },  
    {  
      "Action": [  
        "s3>ListBucket",  
        "s3:GetBucketLocation"  
      ],  
      "Resource": "arn:aws:s3:::netapp-backup*",  
      "Effect": "Allow"  
    },  
    {  
      "Action": [  
        "s3:GetObject",  
        "s3:PutObject",  
        "s3:DeleteObject",  
        "s3>ListAllMyBuckets",  
        "s3:PutObjectTagging",  
        "s3:GetObjectTagging",  
        "s3:RestoreObject",  
        "s3:GetBucketObjectLockConfiguration",  
        "s3:GetObjectRetention",  
        "s3:PutBucketObjectLockConfiguration",  
        "s3:PutObjectRetention"  
      ],  
      "Resource": "arn:aws:s3:::netapp-backup*/*",  
      "Effect": "Allow"  
    }  
  ]  
}
```

Create your own buckets

By default, the service creates buckets for you. Or, if you want to use your own buckets, you can create them before you start the backup activation wizard and then select those buckets in the wizard.

[Learn more about creating your own buckets.](#)

If you create your own buckets, you should use a bucket name of "netapp-backup". If you need to use a custom name, edit the `ontapcloud-instance-policy-netapp-backup` IAMRole for the existing CVOs and add the following JSON block to the S3 permissions Statement array. You need to include `"Resource": "arn:aws:s3:::*` and assign all the necessary permissions that need to be associated with the bucket.

```
[  
  {  
    "Effect": "Allow",  
    "Action": [  
      "s3>ListBucket",  
      "s3:GetBucketLocation"  
    ],  
    "Resource": "arn:aws:s3:::*"  
  },  
  {  
    "Effect": "Allow",  
    "Action": [  
      "s3GetObject",  
      "s3PutObject",  
      "s3DeleteObject",  
      "s3ListAllMyBuckets",  
      "s3PutObjectTagging",  
      "s3GetObjectTagging",  
      "s3RestoreObject",  
      "s3GetBucketObjectLockConfiguration",  
      "s3GetObjectRetention",  
      "s3PutBucketObjectLockConfiguration",  
      "s3PutObjectRetention"  
    ],  
    "Resource": "arn:aws:s3:::*"  
  }  
]
```

Set up customer-managed AWS keys for data encryption

If you want to use the default Amazon S3 encryption keys to encrypt the data passed between your on-prem cluster and the S3 bucket, then you are all set because the default installation uses that type of encryption.

If instead you want to use your own customer-managed keys for data encryption rather than using the default

keys, then you'll need to have the encryption managed keys already set up before you start the NetApp Backup and Recovery wizard.

[Refer to how to use your own Amazon encryption keys with Cloud Volumes ONTAP.](#)

[Refer to how to use your own Amazon encryption keys with NetApp Backup and Recovery.](#)

Configure your system for a private connection using a VPC endpoint interface

If you want to use a standard public internet connection, then all the permissions are set by the Console agent and there is nothing else you need to do.

If you want to have a more secure connection over the internet from your on-prem data center to the VPC, there's an option to select an AWS PrivateLink connection in the Backup activation wizard. It's required if you plan to use a VPN or AWS Direct Connect to connect your on-premises system through a VPC Endpoint interface that uses a private IP address.

Steps

1. Create an Interface endpoint configuration using the Amazon VPC console or the command line. [Refer to details about using AWS PrivateLink for Amazon S3.](#)
2. Modify the security group configuration that's associated with the Console agent. You must change the policy to "Custom" (from "Full Access"), and you must [add the S3 permissions from the backup policy](#) as shown earlier.

If you're using port 80 (HTTP) for communication to the private endpoint, you're all set. You can enable NetApp Backup and Recovery on the cluster now.

If you're using port 443 (HTTPS) for communication to the private endpoint, you must copy the certificate from the VPC S3 endpoint and add it to your ONTAP cluster, as shown in the next 4 steps.

3. Obtain the DNS name of the endpoint from the AWS Console.
4. Obtain the certificate from the VPC S3 endpoint. You do this by [logging into the VM that hosts the Console agent](#) and running the following command. When entering the DNS name of the endpoint, add "bucket" to the beginning, replacing the "*":

```
openssl s_client -connect bucket.vpce-0ff5c15df7e00fbab-yxs71t8v.s3.us-west-2.amazonaws.com:443 -showcerts
```

5. From the output of this command, copy the data for the S3 certificate (all data between, and including, the BEGIN / END CERTIFICATE tags):

```

Certificate chain
0 s:/CN=s3.us-west-2.amazonaws.com`  

i:/C=US/O=Amazon/OU=Server CA 1B/CN=Amazon  

-----BEGIN CERTIFICATE-----  

MIIM6zCCC9OgAwIBAgIQA7MGJ4FaDBR8uL0KR3o1tTANBgkqhkiG9w0BAQsFADBG  

...  

...  

GqvbOz/oO2NWLLFCqI+xmkLcMiPrZy+/6Af+HH2mLCM4EsI2b+IpBmPkriWnnxo=  

-----END CERTIFICATE-----

```

6. Log into the ONTAP cluster CLI and apply the certificate you copied using the following command (substitute your own storage VM name):

```

cluster1::> security certificate install -vserver cluster1 -type server-
ca
Please enter Certificate: Press <Enter> when done

```

Activate backups on your ONTAP volumes

Activate backups at any time directly from your on-premises system.

A wizard takes you through the following major steps:

- [Select the volumes that you want to back up](#)
- [Define the backup strategy](#)
- [Review your selections](#)

You can also [Show the API commands](#) at the review step, so you can copy the code to automate backup activation for future systems.

Start the wizard

Steps

1. Access the Activate backup and recovery wizard using one of the following ways:

- From the Console **Systems** page, select the system and select **Enable > Backup Volumes** next to Backup and Recovery in the right-panel.

If the Amazon S3 destination for your backups exists as a system on the Console **Systems** page, you can drag the ONTAP cluster onto the Amazon S3 object storage.

- Select **Volumes** in the Backup and recovery bar. From the Volumes tab, select the **Actions ...** icon and select **Activate Backup** for a single volume (that does not already have replication or backup to object storage already enabled).

The Introduction page of the wizard shows the protection options including local Snapshots, replication, and backups. If you did the second option in this step, the Define Backup Strategy page appears with one volume selected.

2. Continue with the following options:

- If you already have a Console agent, you're all set. Just select **Next**.
- If you don't already have a Console agent, the **Add a Console agent** option appears. Refer to [Prepare your Console agent](#).

Select the volumes that you want to back up

Choose the volumes you want to protect. A protected volume is one that has one or more of the following: Snapshot policy, replication policy, backup to object policy.

You can choose to protect FlexVol or FlexGroup volumes; however, you cannot select a mix of these volumes when activating backup for a system. See how to [activate backup for additional volumes in the system](#) (FlexVol or FlexGroup) after you have configured backup for the initial volumes.



- You can activate a backup only on a single FlexGroup volume at a time.
- The volumes you select must have the same SnapLock setting. All volumes must have SnapLock Enterprise enabled or have SnapLock disabled.

Steps

If the volumes you choose already have snapshot or replication policies applied, then the policies you select later will overwrite these existing policies.

1. In the Select Volumes page, select the volume or volumes you want to protect.

- Optionally, filter the rows to show only volumes with certain volume types, styles, and more to make the selection easier.
- After you select the first volume, then you can select all FlexVol volumes (FlexGroup volumes can be selected one at a time only). To back up all existing FlexVol volumes, check one volume first and then check the box in the title row.
- To back up individual volumes, check the box for each volume.

2. Select **Next**.

Define the backup strategy

Defining the backup strategy involves setting the following options:

- Whether you want one or all of the backup options: local snapshots, replication, and backup to object storage
- Architecture
- Local snapshot policy
- Replication target and policy



If the volumes you choose have different snapshot and replication policies than the policies you select in this step, the existing policies will be overwritten.

- Backup to object storage information (provider, encryption, networking, backup policy, and export options).

Steps

1. In the Define backup strategy page, choose one or all of the following. All three are selected by default:

- **Local Snapshots:** If you are performing replication or back up to object storage, local snapshots must be created.
- **Replication:** Creates replicated volumes on another ONTAP storage system.
- **Backup:** Backs up volumes to object storage.

2. **Architecture:** If you chose replication and backup, choose one of the following flows of information:

- **Cascading:** Information flows from the primary to the secondary to object storage and from the secondary to object storage.
- **Fan out:** Information flows from the primary to the secondary *and* from the primary to object storage.

For details about these architectures, refer to [Plan your protection journey](#).

3. **Local Snapshot:** Choose an existing snapshot policy or create a policy.



To create a custom policy before activating the snapshot, refer to [Create a policy](#).

4. To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to five schedules, typically of different frequencies.
 - For backup-to-object policies, set the DataLock and Ransomware Resilience settings. For details on DataLock and Ransomware Resilience, refer to [Backup-to-object policy settings](#).
- Select **Create**.

5. **Replication:** Set the following options:

- **Replication target:** Select the destination system and SVM. Optionally, select the destination aggregate or aggregates and prefix or suffix that will be added to the replicated volume name.
- **Replication policy:** Choose an existing replication policy or create a policy.



To create a custom policy before activating the replication, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to five schedules, typically of different frequencies.
- Select **Create**.

6. **Back up to Object:** If you selected **Backup**, set the following options:

- **Provider:** Select **Amazon Web Services**.
- **Provider settings:** Enter the provider details and AWS region where the backups will be stored.

The access key and secret key are for the IAM user you created to give the ONTAP cluster access to the S3 bucket.

- **Bucket:** Either choose an existing S3 bucket or create a new one. Refer to [Add S3 buckets](#).
- **Encryption key:** If you created a new S3 bucket, enter encryption key information given to you from the provider. Choose whether you'll use the default Amazon S3 encryption keys, or choose your own customer-managed keys from your AWS account, to manage encryption of your data.



If you chose an existing bucket, encryption information is already available, so you don't need to enter it now.

- **Networking:** Choose the IPspace, and whether you'll be using a Private Endpoint. Private Endpoint is disabled by default.
 - a. The IPspace in the ONTAP cluster where the volumes you want to back up reside. The intercluster LIFs for this IPspace must have outbound internet access.
 - b. Optionally, choose whether you'll use an AWS PrivateLink that you have previously configured. [See details about using AWS PrivateLink for Amazon S3.](#)
- **Backup policy:** Select an existing backup policy or create a policy.



To create a custom policy before activating the backup, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to five schedules, typically of different frequencies.
- Select **Create**.

- **Export existing snapshots to object storage as backup copies:** If there are any local snapshots for volumes in this system that match the backup schedule label you just selected for this system (for example, daily, weekly, etc.), this additional prompt is displayed. Check this box to have all historic snapshots copied to object storage as backup files to ensure the most complete protection for your volumes.

7. Select **Next**.

Review your selections

This is the chance to review your selections and make adjustments, if necessary.

Steps

1. In the Review page, review your selections.
2. Optionally check the box to **Automatically synchronize the Snapshot policy labels with the replication and backup policy labels**. This creates Snapshots with a label that matches the labels in the replication and backup policies.
3. Select **Activate Backup**.

Result

NetApp Backup and Recovery starts taking the initial backups of your volumes. The baseline transfer of the replicated volume and the backup file includes a full copy of the primary storage system data. Subsequent transfers contain differential copies of the primary data contained in snapshots.

A replicated volume is created in the destination cluster that will be synchronized with the primary storage volume.

The S3 bucket is created in the service account indicated by the S3 access key and secret key you entered, and the backup files are stored there. The Volume Backup Dashboard is displayed so you can monitor the state of the backups.

You can also monitor the status of backup and restore jobs using the [Job Monitoring page](#).

Show the API commands

You might want to display and optionally copy the API commands used in the Activate backup and recovery wizard. You might want to do this to automate backup activation in future systems.

Steps

1. From the Activate backup and recovery wizard, select **View API request**.
2. To copy the commands to the clipboard, select the **Copy** icon.

Back up on-premises ONTAP data to Azure Blob storage with NetApp Backup and Recovery

Complete a few steps in NetApp Backup and Recovery to get started backing up volume data from your on-premises ONTAP systems to a secondary storage system and to Azure Blob storage.



"On-premises ONTAP systems" include FAS, AFF, and ONTAP Select systems.



To switch to and from NetApp Backup and Recovery workloads, refer to [Switch to different NetApp Backup and Recovery workloads](#).

Identify the connection method

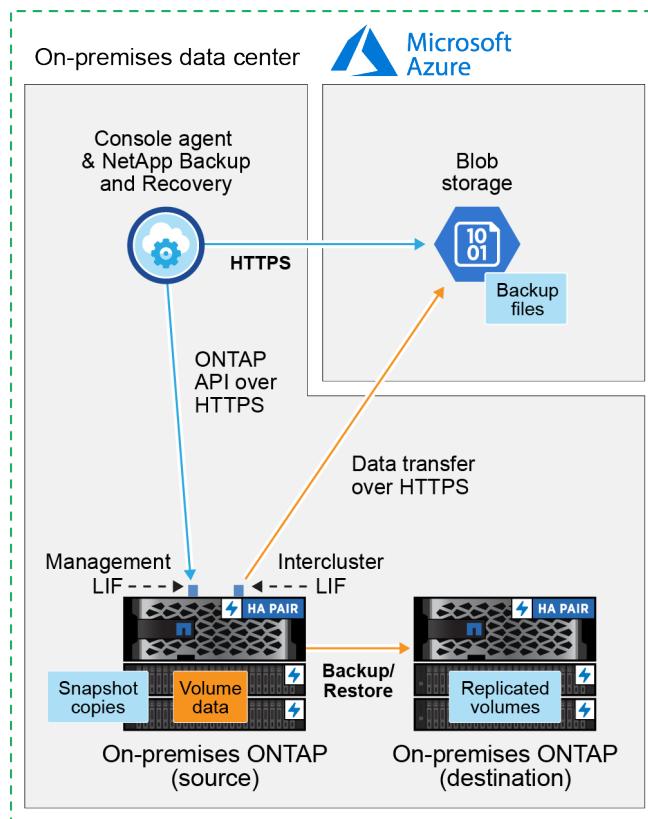
Choose which of the two connection methods you will use when configuring backups from on-premises ONTAP systems to Azure Blob.

- **Public connection** - Directly connect the ONTAP system to Azure Blob storage using a public Azure endpoint.
- **Private connection** - Use a VPN or ExpressRoute and route traffic through a VNet Private Endpoint that uses a private IP address.

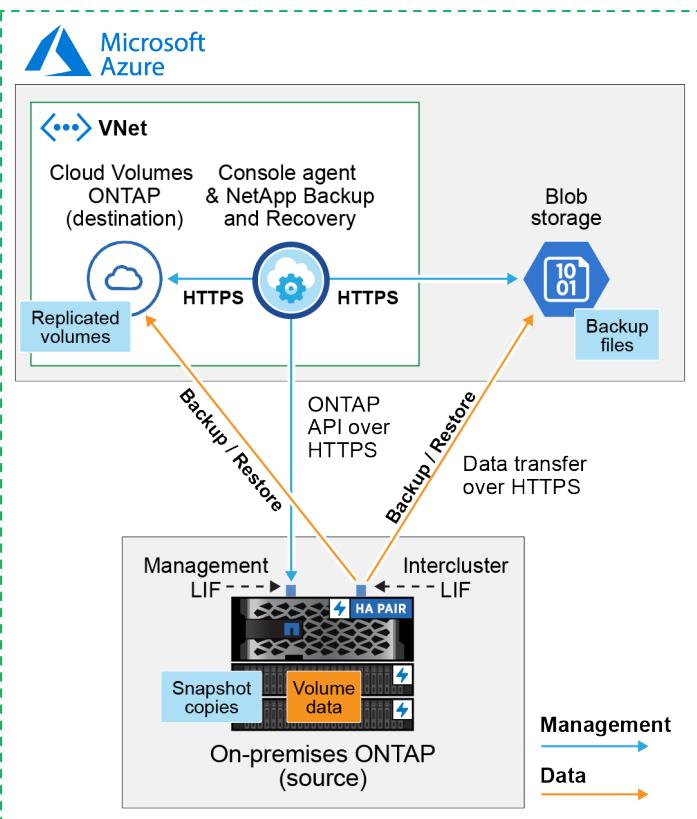
Optionally, you can connect to a secondary ONTAP system for replicated volumes using the public or private connection as well.

The following diagram shows the **public connection** method and the connections that you need to prepare between the components. You can use a Console agent that you've installed on your premises, or a Console agent that you've deployed in the Azure VNet.

Console agent installed on-premises (Public)

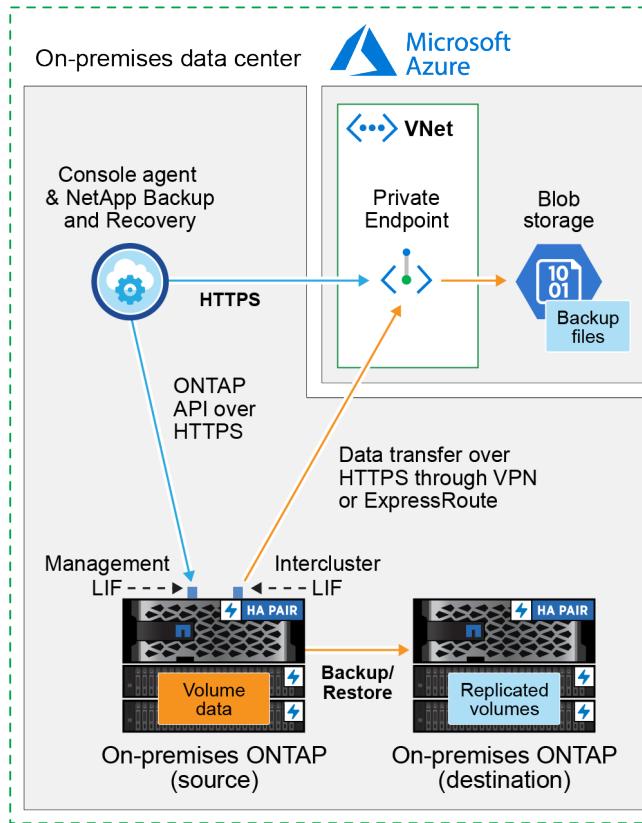


Console agent deployed in Azure VNet (Public)

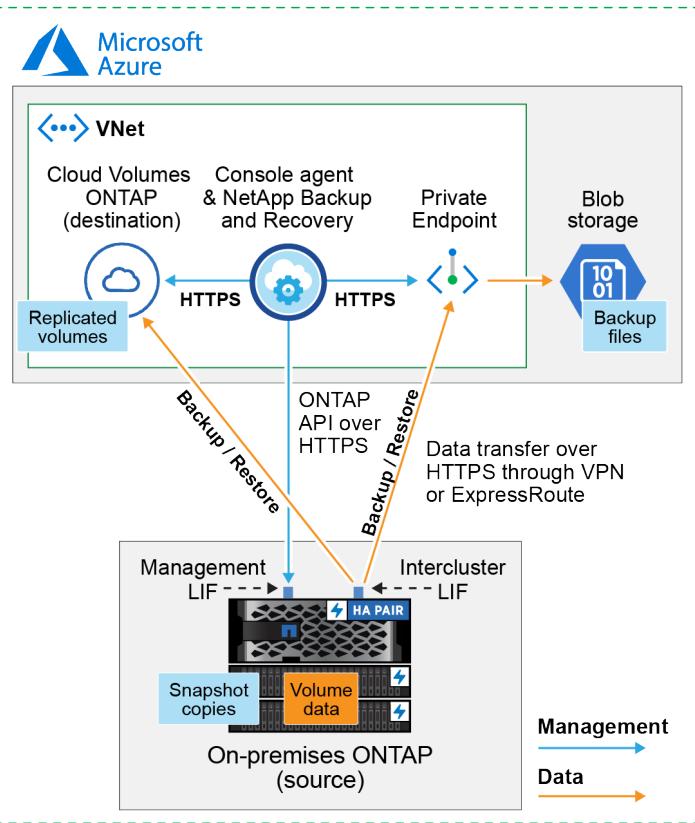


The following diagram shows the **private connection** method and the connections that you need to prepare between the components. You can use a Console agent that you've installed on your premises, or a Console agent that you've deployed in the Azure VNet.

Console agent installed on-premises (Private)



Console agent deployed in Azure VNet (Private)



Prepare your Console agent

The Console agent is the main software for NetApp Console functionality. A Console agent is required to back up and restore your ONTAP data.

Create or switch Console agents

If you already have a Console agent deployed in your Azure VNet or on your premises, then you're all set.

If not, then you'll need to create a Console agent in one of those locations to back up ONTAP data to Azure Blob storage. You can't use a Console agent that's deployed in another cloud provider.

- [Learn about Console agents](#)
- [Install a Console agent in Azure](#)
- [Install a Console agent in your premises](#)
- [Install a Console agent in an Azure Government region](#)

NetApp Backup and Recovery is supported in Azure Government regions when the Console agent is deployed in the cloud - not when it's installed in your premises. Additionally, you must deploy the Console agent from the Azure Marketplace. You can't deploy the Console agent in a Government region from the Console SaaS website.

Prepare networking for the Console agent

Ensure that the Console agent has the required networking connections.

Steps

1. Ensure that the network where the Console agent is installed enables the following connections:
 - An HTTPS connection over port 443 to NetApp Backup and Recovery and to your Blob object storage ([see the list of endpoints](#))
 - An HTTPS connection over port 443 to your ONTAP cluster management LIF
 - In order for the NetApp Backup and Recovery Search & Restore functionality to work, port 1433 must be open for communication between the Console agent and the Azure Synapse SQL services.
 - Additional inbound security group rules are required for Azure and Azure Government deployments. See [Rules for the Console agent in Azure](#) for details.
2. Enable a VNet Private Endpoint to Azure storage. This is needed if you have an ExpressRoute or VPN connection from your ONTAP cluster to the VNet and you want communication between the Console agent and Blob storage to stay in your virtual private network (a **private** connection).

Verify or add permissions to the Console agent

To use the NetApp Backup and Recovery Search & Restore functionality, you need to have specific permissions in the role for the Console agent so that it can access the Azure Synapse Workspace and Data Lake Storage Account. See the permissions below, and follow the steps if you need to modify the policy.

Before you start

You must register the Azure Synapse Analytics Resource Provider (called "Microsoft.Synapse") with your Subscription. [See how to register this resource provider for your subscription](#). You must be the **Subscription Owner** or **Contributor** to register the resource provider.

Steps

1. Identify the role assigned to the Console agent virtual machine:
 - a. In the Azure portal, open the Virtual machines service.
 - b. Select the Console agent virtual machine.
 - c. Under **Settings**, select **Identity**.
 - d. Select **Azure role assignments**.
 - e. Make note of the custom role assigned to the Console agent virtual machine.
2. Update the custom role:
 - a. In the Azure portal, open your Azure subscription.
 - b. Select **Access control (IAM) > Roles**.
 - c. Select the ellipsis (...) for the custom role and then select **Edit**.
 - d. Select **JSON** and add the following permissions:

```
"Microsoft.Storage/storageAccounts/listkeys/action",
"Microsoft.Storage/storageAccounts/read",
"Microsoft.Storage/storageAccounts/write",
"Microsoft.Storage/storageAccounts/blobServices/containers/read",
"Microsoft.Storage/storageAccounts/listAccountSas/action",
"Microsoft.KeyVault/vaults/read",
"Microsoft.KeyVault/vaults/accessPolicies/write",
"Microsoft.Network/networkInterfaces/read",
"Microsoft.Resources/subscriptions/locations/read",
"Microsoft.Network/virtualNetworks/read",
"Microsoft.Network/virtualNetworks/subnets/read",
"Microsoft.Resources/subscriptions/resourceGroups/read",
"Microsoft.Resources/subscriptions/resourceGroups/resources/read"
",
"Microsoft.Resources/subscriptions/resourceGroups/write",
"Microsoft.Authorization/locks/*",
"Microsoft.Network/privateEndpoints/write",
"Microsoft.Network/privateEndpoints/read",
"Microsoft.Network/privateDnsZones/virtualNetworkLinks/write",
"Microsoft.Network/virtualNetworks/join/action",
"Microsoft.Network/privateDnsZones/A/write",
"Microsoft.Network/privateDnsZones/read",
"Microsoft.Network/privateDnsZones/virtualNetworkLinks/read",
"Microsoft.Network/networkInterfaces/delete",
"Microsoft.Network/networkSecurityGroups/delete",
"Microsoft.Resources/deployments/delete",
"Microsoft.ManagedIdentity/userAssignedIdentities/assign/action",
"Microsoft.Synapse/workspaces/write",
"Microsoft.Synapse/workspaces/read",
"Microsoft.Synapse/workspaces/delete",
"Microsoft.Synapse/register/action",
"Microsoft.Synapse/checkNameAvailability/action",
"Microsoft.Synapse/workspaces/operationStatuses/read",
"Microsoft.Synapse/workspaces/firewallRules/read",
"Microsoft.Synapse/workspaces/replaceAllIpFirewallRules/action",
"Microsoft.Synapse/workspaces/operationResults/read",
"Microsoft.Synapse/workspaces/privateEndpointConnectionsApproval/
action"
```

[View the full JSON format for the policy](#)

- e. Select **Review + update** and then select **Update**.

Verify license requirements

You'll need to verify license requirements for both Azure and the Console:

- Before you can activate NetApp Backup and Recovery for your cluster, you'll need to either subscribe to a pay-as-you-go (PAYGO) Console Marketplace offering from Azure, or purchase and activate a NetApp Backup and Recovery BYOL license from NetApp. These licenses are for your account and can be used across multiple systems.
 - For NetApp Backup and Recovery PAYGO licensing, you'll need a subscription to the [NetApp Console offering from the Azure Marketplace](#). Billing for NetApp Backup and Recovery is done through this subscription.
 - For NetApp Backup and Recovery BYOL licensing, you'll need the serial number from NetApp that enables you to use the service for the duration and capacity of the license. [Learn how to manage your BYOL licenses](#).
- You need to have an Azure subscription for the object storage space where your backups will be located.

Supported regions

You can create backups from on-premises systems to Azure Blob in all regions, including Azure Government regions. You specify the region where the backups will be stored when you set up the service.

Prepare your ONTAP clusters

Unresolved directive in prev-ontap-backup-onprem-azure.adoc - include:../../_include/backup-onprem-prepare-onprem-ONTAP-cluster.adoc[]

Verify ONTAP networking requirements for backing up data to object storage

You must configure the following requirements on the system that connects to object storage.

- For a fan-out backup architecture, configure the following settings on the *primary* system.
- For a cascaded backup architecture, configure the following settings on the *secondary* system.

The following ONTAP cluster networking requirements are needed:

- The ONTAP cluster initiates an HTTPS connection over port 443 from the intercluster LIF to Azure Blob storage for backup and restore operations.
ONTAP reads and writes data to and from object storage. The object storage never initiates, it just responds.
- ONTAP requires an inbound connection from the Console agent to the cluster management LIF. The Console agent can reside in an Azure VNet.
- An intercluster LIF is required on each ONTAP node that hosts the volumes you want to back up. The LIF must be associated with the *IPspace* that ONTAP should use to connect to object storage. [Learn more about IPspaces](#).

When you set up NetApp Backup and Recovery, you are prompted for the IPspace to use. You should choose the IPspace that each LIF is associated with. That might be the "Default" IPspace or a custom IPspace that you created.

- The nodes' and intercluster LIFs are able to access the object store.

- DNS servers have been configured for the storage VM where the volumes are located. See how to [configure DNS services for the SVM](#).
- If you use are using a different IPspace than the Default, then you might need to create a static route to get access to the object storage.
- Update firewall rules, if necessary, to allow NetApp Backup and Recovery service connections from ONTAP to object storage through port 443 and name resolution traffic from the storage VM to the DNS server over port 53 (TCP/UDP).

Verify ONTAP networking requirements for replicating volumes

Unresolved directive in prev-ontap-backup-onprem-azure.adoc - include:::_include/backup-onprem-prepare-source-destination-systems-for-repl.adoc[]

Prepare Azure Blob as your backup target

1. You can use your own custom-managed keys for data encryption in the activation wizard instead of using the default Microsoft-managed encryption keys. In this case you will need to have the Azure Subscription, Key Vault name, and the Key. [Learn how to use your own keys](#).

Note that Backup and recovery supports *Azure access policies* as the permission model. The *Azure role-based access control* (Azure RBAC) permission model is not currently supported.

2. If you want to have a more secure connection over the public internet from your on-prem data center to the VNet, there is an option to configure an Azure Private Endpoint in the activation wizard. In this case you will need to know the VNet and Subnet for this connection. [Refer to details about using a Private Endpoint](#).

Create your Azure Blob storage account

By default, the service creates storage accounts for you. If you want to use your own storage accounts, you can create them before you start the backup activation wizard and then select those storage accounts in the wizard.

[Learn more about creating your own storage accounts](#).

Activate backups on your ONTAP volumes

Activate backups at any time directly from your on-premises system.

A wizard takes you through the following major steps:

- [Select the volumes that you want to back up](#)
- [Define the backup strategy](#)
- [Review your selections](#)

You can also [Show the API commands](#) at the review step, so you can copy the code to automate backup activation for future systems.

Start the wizard

Steps

1. Access the Activate backup and recovery wizard using one of the following ways:
 - From the **Console Systems** page, select the system and select **Enable > Backup Volumes** next to the

Backup and recovery service in the right-panel.

If the Azure destination for your backups exists on the Console **Systems** page, you can drag the ONTAP cluster onto the Azure Blob object storage.

- Select **Volumes** in the Backup and recovery bar. From the Volumes tab, select the **Actions** **...** icon and select **Activate Backup** for a single volume (that does not already have replication or backup to object storage already enabled).

The Introduction page of the wizard shows the protection options including local Snapshots, replication, and backups. If you did the second option in this step, the Define Backup Strategy page appears with one volume selected.

2. Continue with the following options:

- If you already have a Console agent, you're all set. Just select **Next**.
- If you don't already have a Console agent, the **Add a Console agent** option appears. Refer to [Prepare your Console agent](#).

Select the volumes that you want to back up

Choose the volumes you want to protect. A protected volume is one that has one or more of the following: Snapshot policy, replication policy, backup to object policy.

You can choose to protect FlexVol or FlexGroup volumes; however, you cannot select a mix of these volumes when activating backup for a system. See how to [activate backup for additional volumes in the system](#) (FlexVol or FlexGroup) after you have configured backup for the initial volumes.



- You can activate a backup only on a single FlexGroup volume at a time.
- The volumes you select must have the same SnapLock setting. All volumes must have SnapLock Enterprise enabled or have SnapLock disabled.

Steps

Note that if the volumes you choose already have Snapshot or replication policies applied, then the policies you select later will overwrite these existing policies.

1. In the Select Volumes page, select the volume or volumes you want to protect.
 - Optionally, filter the rows to show only volumes with certain volume types, styles, and more to make the selection easier.
 - After you select the first volume, then you can select all FlexVol volumes (FlexGroup volumes can be selected one at a time only). To back up all existing FlexVol volumes, check one volume first and then check the box in the title row.
 - To back up individual volumes, check the box for each volume.
2. Select **Next**.

Define the backup strategy

Defining the backup strategy involves setting the following options:

- Whether you want one or all of the backup options: local Snapshots, replication, and backup to object storage
- Architecture

- Local Snapshot policy
- Replication target and policy



If the volumes you choose have different Snapshot and replication policies than the policies you select in this step, the existing policies will be overwritten.

- Backup to object storage information (provider, encryption, networking, backup policy, and export options).

Steps

1. In the Define backup strategy page, choose one or all of the following. All three are selected by default:
 - **Local Snapshots:** If you are performing replication or back up to object storage, local snapshots must be created.
 - **Replication:** Creates replicated volumes on another ONTAP storage system.
 - **Backup:** Backs up volumes to object storage.
2. **Architecture:** If you chose replication and backup, choose one of the following flows of information:
 - **Cascading:** Information flows from the primary to the secondary, and from secondary to object storage.
 - **Fan out:** Information flows from the primary to the secondary *and* from the primary to object storage.
 For details about these architectures, refer to [Plan your protection journey](#).
3. **Local Snapshot:** Choose an existing snapshot policy or create a new one.



To create a custom policy before activating the snapshot, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to five schedules, typically of different frequencies.
- Select **Create**.

4. **Replication:** Set the following options:

- **Replication target:** Select the destination system and SVM. Optionally, select the destination aggregate or aggregates and prefix or suffix that will be added to the replicated volume name.
- **Replication policy:** Choose an existing replication policy or create a new one.



To create a custom policy before activating the replication, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to five schedules, typically of different frequencies.
- Select **Create**.

5. **Back up to Object:** If you selected **Backup**, set the following options:

- **Provider:** Select **Microsoft Azure**.
- **Provider settings:** Enter the provider details and region where the backups will be stored.

Either create a new storage account or select an existing one.

Either create your own resource group that manages the Blob container or select the resource group type and group.



If you want to protect your backup files from being modified or deleted, ensure that the storage account was created with immutable storage enabled using a 30-day retention period.



If you want to tier older backup files to Azure Archive Storage for further cost optimization, ensure that the storage account has the appropriate Lifecycle rule.

- **Encryption key:** If you created a new Azure storage account, enter encryption key information given to you from the provider. Choose whether you'll use the default Azure encryption keys, or choose your own customer-managed keys from your Azure account, to manage encryption of your data.

If you choose to use your own customer-managed keys, enter the key vault and key information.



If you chose an existing Microsoft storage account, encryption information is already available, so you don't need to enter it now.

- **Networking:** Choose the IPspace, and whether you'll be using a Private Endpoint. Private Endpoint is disabled by default.
 - a. The IPspace in the ONTAP cluster where the volumes you want to back up reside. The intercluster LIFs for this IPspace must have outbound internet access.
 - b. Optionally, choose whether you'll use an Azure private endpoint that you have previously configured. [Learn about using an Azure private endpoint](#).
- **Backup policy:** Select an existing Backup to object storage policy or create a new one.



To create a custom policy before activating the backup, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to five schedules, typically of different frequencies.
- For backup-to-object policies, set the DataLock and Ransomware Resilience settings. For details on DataLock and Ransomware Resilience, refer to [Backup-to-object policy settings](#).
- Select **Create**.

- **Export existing snapshots to object storage as backup copies:** If there are any local snapshots for volumes in this system that match the backup schedule label you just selected for this system (for example, daily, weekly, etc.), this additional prompt is displayed. Check this box to have all historic Snapshots copied to object storage as backup files to ensure the most complete protection for your volumes.

6. Select **Next**.

Review your selections

This is the chance to review your selections and make adjustments, if necessary.

Steps

1. In the Review page, review your selections.
2. Optionally check the box to **Automatically synchronize the Snapshot policy labels with the replication and backup policy labels**. This creates Snapshots with a label that matches the labels in the replication and backup policies.
3. Select **Activate Backup**.

Result

NetApp Backup and Recovery starts taking the initial backups of your volumes. The baseline transfer of the replicated volume and the backup file includes a full copy of the primary storage system data. Subsequent transfers contain differential copies of the primary storage system data contained in snapshots.

A replicated volume is created in the destination cluster that will be synchronized with the primary volume.

A Blob storage account is created in the resource group you entered, and the backup files are stored there. The Volume Backup Dashboard is displayed so you can monitor the state of the backups.

You can also monitor the status of backup and restore jobs using the [Job Monitoring page](#).

Show the API commands

You might want to display and optionally copy the API commands used in the Activate backup and recovery wizard. You might want to do this to automate backup activation in future systems.

Steps

1. From the Activate backup and recovery wizard, select **View API request**.
2. To copy the commands to the clipboard, select the **Copy** icon.

Back up on-premises ONTAP data to Google Cloud Storage with NetApp Backup and Recovery

Complete a few steps in NetApp Backup and Recovery to get started backing up volume data from your on-premises primary ONTAP systems to a secondary storage system and to Google Cloud Storage.



"On-premises ONTAP systems" include FAS, AFF, and ONTAP Select systems.



To switch to and from NetApp Backup and Recovery workloads, refer to [Switch to different NetApp Backup and Recovery workloads](#).

Identify the connection method

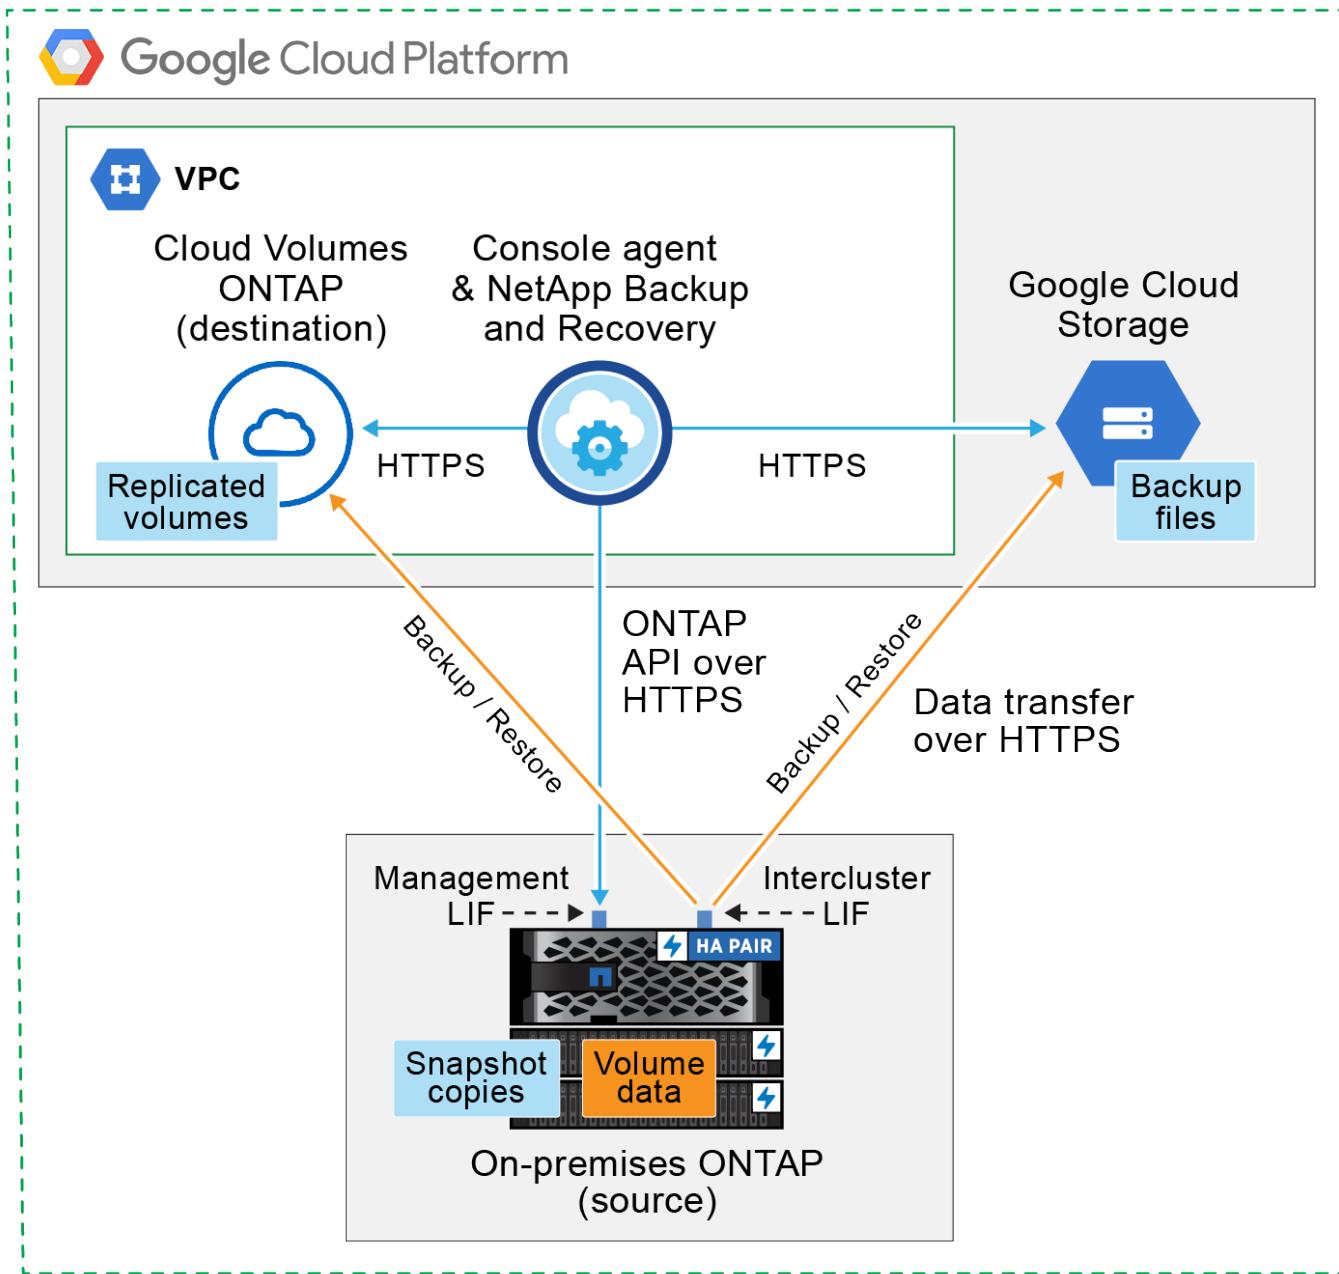
Choose which of the two connection methods you will use when configuring backups from on-premises ONTAP systems to Google Cloud Storage.

- **Public connection** - Directly connect the ONTAP system to Google Cloud Storage using a public Google endpoint.
- **Private connection** - Use a VPN or Google Cloud Interconnect and route traffic through a Private Google Access interface that uses a private IP address.

Optionally, you can connect to a secondary ONTAP system for replicated volumes using the public or private connection as well.

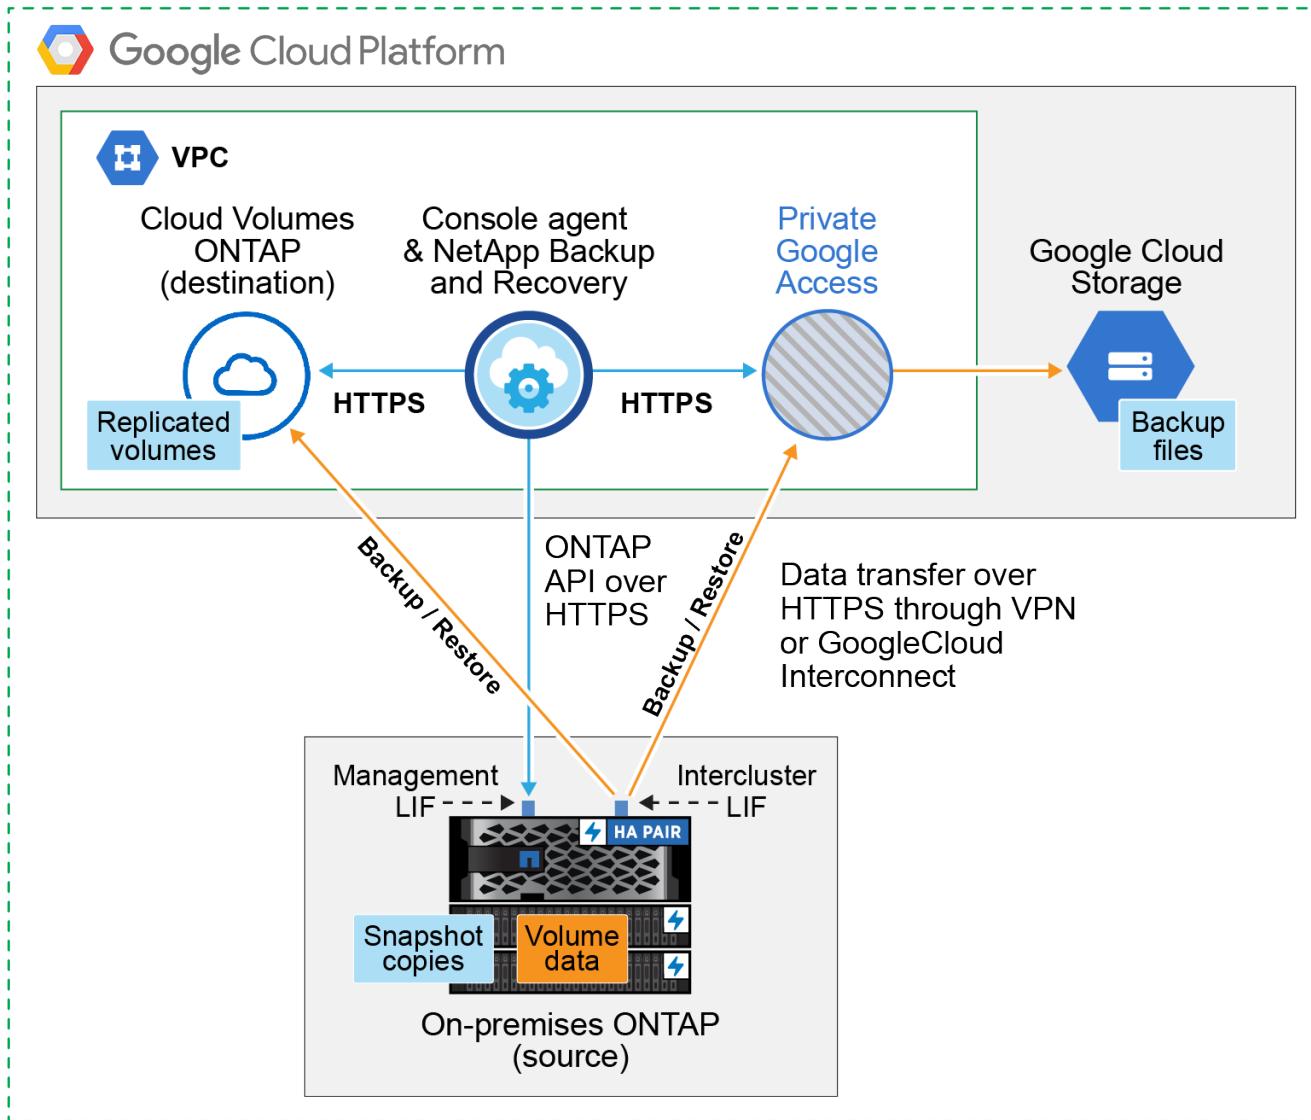
The following diagram shows the **public connection** method and the connections that you need to prepare between the components. The Console agent must be deployed in the Google Cloud Platform VPC.

Console agent deployed in Google Cloud VPC (Public)



The following diagram shows the **private connection** method and the connections that you need to prepare between the components. The Console agent must be deployed in the Google Cloud Platform VPC.

Console agent deployed in Google Cloud VPC (Private)



Prepare your Console agent

The Console agent is the main software for Console functionality. A Console agent is required to back up and restore your ONTAP data.

Create or switch Console agents

If you already have a Console agent deployed in your Google Cloud Platform VPC, then you're all set.

If not, then you'll need to create a Console agent in that location to back up ONTAP data to Google Cloud Storage. You can't use a Console agent that's deployed in another cloud provider, or on-premises.

- [Learn about Console agents](#)
- [Install a Console agent in GCP](#)

Prepare networking for the Console agent

Ensure that the Console agent has the required networking connections.

Steps

1. Ensure that the network where the Console agent is installed enables the following connections:
 - An HTTPS connection over port 443 to NetApp Backup and Recovery and to your Google Cloud storage ([see the list of endpoints](#))
 - An HTTPS connection over port 443 to your ONTAP cluster management LIF
2. Enable Private Google Access (or Private Service Connect) on the subnet where you plan to deploy the Console agent. [Private Google Access](#) or [Private Service Connect](#) are needed if you have a direct connection from your ONTAP cluster to the VPC and you want communication between the Console agent and Google Cloud Storage to stay in your virtual private network (a **private** connection).

Follow the Google instructions for setting up these Private access options. Make sure your DNS servers have been configured to point `www.googleapis.com` and `storage.googleapis.com` to the correct internal (private) IP addresses.

Verify or add permissions to the Console agent

To use the NetApp Backup and Recovery "Search & Restore" functionality, you need to have specific permissions in the role for the Console agent so that it can access the Google Cloud BigQuery service. Review the permissions below, and follow the steps if you need to modify the policy.

Steps

1. In the [Google Cloud Console](#), go to the **Roles** page.
2. Using the drop-down list at the top of the page, select the project or organization that contains the role that you want to edit.
3. Select a custom role.
4. Select **Edit Role** to update the role's permissions.
5. Select **Add Permissions** to add the following new permissions to the role.

```
bigquery.jobs.get
bigquery.jobs.list
bigquery.jobs.listAll
bigquery.datasets.create
bigquery.datasets.get
bigquery.jobs.create
bigquery.tables.get
bigquery.tables.getData
bigquery.tables.list
bigquery.tables.create
```

6. Select **Update** to save the edited role.

Verify license requirements

- Before you can activate NetApp Backup and Recovery for your cluster, you'll need to either subscribe to a pay-as-you-go (PAYGO) Console Marketplace offering from Google, or purchase and activate a NetApp Backup and Recovery BYOL license from NetApp. These licenses are for your account and can be used across multiple systems.
 - For NetApp Backup and Recovery PAYGO licensing, you'll need a subscription to the [NetApp Console offering from the Google Marketplace](#). Billing for NetApp Backup and Recovery is done through this subscription.
 - For NetApp Backup and Recovery BYOL licensing, you'll need the serial number from NetApp that enables you to use the service for the duration and capacity of the license. [Learn how to manage your BYOL licenses](#).
- You need to have a Google subscription for the object storage space where your backups will be located.

Supported regions

You can create backups from on-premises systems to Google Cloud Storage in all regions. You specify the region where backups will be stored when you set up the service.

Prepare your ONTAP clusters

Unresolved directive in prev-ontap-backup-onprem-gcp.adoc - include:../_include/backup-onprem-prepare-onprem-ONTAP-cluster.adoc[]

Verify ONTAP networking requirements for backing up data to object storage

You must configure the following requirements on the system that connects to object storage.

- For a fan-out backup architecture, configure the following settings on the *primary* system.
- For a cascaded backup architecture, configure the following settings on the *secondary* system.

The following ONTAP cluster networking requirements are needed:

- The ONTAP cluster initiates an HTTPS connection over port 443 from the intercluster LIF to Google Cloud Storage for backup and restore operations.
ONTAP reads and writes data to and from object storage. The object storage never initiates, it just responds.
- ONTAP requires an inbound connection from the Console agent to the cluster management LIF. The Console agent can reside in a Google Cloud Platform VPC.
- An intercluster LIF is required on each ONTAP node that hosts the volumes you want to back up. The LIF must be associated with the *IPspace* that ONTAP should use to connect to object storage. [Learn more about IPspaces](#).

When you set up NetApp Backup and Recovery, you are prompted for the *IPspace* to use. You should choose the *IPspace* that each LIF is associated with. That might be the "Default" *IPspace* or a custom *IPspace* that you created.

- The nodes' intercluster LIFs are able to access the object store.
- DNS servers have been configured for the storage VM where the volumes are located. See how to [configure DNS services for the SVM](#).

If you're using Private Google Access or Private Service Connect, make sure your DNS servers have been configured to point `storage.googleapis.com` to the correct internal (private) IP address.

- Note that if you are using a different IP space than the Default, then you might need to create a static route to get access to the object storage.
- Update firewall rules, if necessary, to allow NetApp Backup and Recovery connections from ONTAP to object storage through port 443, and name resolution traffic from the storage VM to the DNS server over port 53 (TCP/UDP).

Verify ONTAP networking requirements for replicating volumes

Unresolved directive in `prev-ontap-backup-onprem-gcp.adoc` - include::`../_include/backup-onprem-prepare-source-destination-systems-for-repl.adoc`[]

Prepare Google Cloud Storage as your backup target

Preparing Google Cloud Storage as your backup target involves the following steps:

- Set up permissions.
- (Optional) Create your own buckets. (The service will create buckets for you if you want.)
- (Optional) Set up customer-managed keys for data encryption

Set up permissions

You need to provide storage access keys for a service account that has specific permissions using a custom role. A service account enables NetApp Backup and Recovery to authenticate and access Cloud Storage buckets used to store backups. The keys are required so that Google Cloud Storage knows who is making the request.

Steps

1. In the [Google Cloud Console](#), go to the **Roles** page.
2. [Create a new role](#) with the following permissions:

```
storage.buckets.create
storage.buckets.delete
storage.buckets.get
storage.buckets.list
storage.buckets.update
storage.buckets.getIamPolicy
storage.multipartUploads.create
storage.objects.create
storage.objects.delete
storage.objects.get
storage.objects.list
storage.objects.update
```

3. In the Google Cloud console, [go to the Service accounts page](#).
4. Select your Cloud project.

5. Select **Create service account** and provide the required information:
 - a. **Service account details:** Enter a name and description.
 - b. **Grant this service account access to project:** Select the custom role that you just created.
 - c. Select **Done**.
6. Go to [GCP Storage Settings](#) and create access keys for the service account:
 - a. Select a project, and select **Interoperability**. If you haven't already done so, select **Enable interoperability access**.
 - b. Under **Access keys for service accounts**, select **Create a key for a service account**, select the service account that you just created, and click **Create Key**.

You'll need to enter the keys in NetApp Backup and Recovery later when you configure the backup service.

Create your own buckets

By default, the service creates buckets for you. Or, if you want to use your own buckets, you can create them before you start the backup activation wizard and then select those buckets in the wizard.

[Learn more about creating your own buckets.](#)

Set up customer-managed encryption keys (CMEK) for data encryption

You can use your own customer-managed keys for data encryption instead of using the default Google-managed encryption keys. Both cross-region and cross-project keys are supported, so you can choose a project for a bucket that is different than the project of the CMEK key.

If you're planning to use your own customer-managed keys:

- You'll need to have the Key Ring and the Key Name so you can add this information in the activation wizard. [Learn more about customer-managed encryption keys](#).
- You'll need to verify that these required permissions are included in the role for the Console agent:

```
cloudkms.cryptoKeys.get
cloudkms.cryptoKeys.getIamPolicy
cloudkms.cryptoKeys.list
cloudkms.cryptoKeys.setIamPolicy
cloudkms.keyRings.get
cloudkms.keyRings.getIamPolicy
cloudkms.keyRings.list
cloudkms.keyRings.setIamPolicy
```

- You'll need to verify that the Google "Cloud Key Management Service (KMS)" API is enabled in your project. See the [Google Cloud documentation: Enabling APIs](#) for details.

CMEK considerations:

- Both HSM (Hardware-backed) and Software-generated keys are supported.
- Both newly created or imported Cloud KMS keys are supported.

- Only regional keys are supported, global keys are not supported.
- Currently, only the "Symmetric encrypt/decrypt" purpose is supported.
- The service agent associated with the Storage Account is assigned the "CryptoKey Encrypter/Decrypter (roles/cloudkms.cryptoKeyEncrypterDecrypter)" IAM role by NetApp Backup and Recovery.

Activate backups on your ONTAP volumes

Activate backups at any time directly from your on-premises system.

A wizard takes you through the following major steps:

- [Select the volumes that you want to back up](#)
- [Define the backup strategy](#)
- [Review your selections](#)

You can also [Show the API commands](#) at the review step, so you can copy the code to automate backup activation for future systems.

Start the wizard

Steps

1. Access the Activate backup and recovery wizard using one of the following ways:

- From the **Console Systems** page, select the system and select **Enable > Backup Volumes** next to Backup and Recovery in the right-panel.

If the Google Cloud Storage destination for your backups exists as on the **Console Systems** page, you can drag the ONTAP cluster onto the Google Cloud object storage.

- Select **Volumes** in the Backup and recovery bar. From the Volumes tab, select the **Actions ...** icon and select **Activate Backup** for a single volume (that does not already have replication or backup to object storage already enabled).

The Introduction page of the wizard shows the protection options including local Snapshots, replication, and backups. If you did the second option in this step, the Define Backup Strategy page appears with one volume selected.

2. Continue with the following options:

- If you already have a Console agent, you're all set. Just select **Next**.
- If you don't already have a Console agent, the **Add a Console agent** option appears. Refer to [Prepare your Console agent](#).

Select the volumes that you want to back up

Choose the volumes you want to protect. A protected volume is one that has one or more of the following: Snapshot policy, replication policy, backup to object policy.

You can choose to protect FlexVol or FlexGroup volumes; however, you cannot select a mix of these volumes when activating backup for a system. See how to [activate backup for additional volumes in the system](#) (FlexVol or FlexGroup) after you have configured backup for the initial volumes.

- You can activate a backup only on a single FlexGroup volume at a time.
- The volumes you select must have the same SnapLock setting. All volumes must have SnapLock Enterprise enabled or have SnapLock disabled.

Steps

If the volumes you choose already have snapshot or replication policies applied, then the policies you select later will overwrite these existing policies.

1. In the Select Volumes page, select the volume or volumes you want to protect.
 - Optionally, filter the rows to show only volumes with certain volume types, styles, and more to make the selection easier.
 - After you select the first volume, then you can select all FlexVol volumes (FlexGroup volumes can be selected one at a time only). To back up all existing FlexVol volumes, check one volume first and then check the box in the title row.
 - To back up individual volumes, check the box for each volume.
2. Select **Next**.

Define the backup strategy

Defining the backup strategy involves setting the following options:

- Whether you want one or all of the backup options: local snapshots, replication, and backup to object storage
- Architecture
- Local snapshot policy
- Replication target and policy



If the volumes you choose have different snapshot and replication policies than the policies you select in this step, the existing policies will be overwritten.

- Backup to object storage information (provider, encryption, networking, backup policy, and export options).

Steps

1. In the Define backup strategy page, choose one or all of the following. All three are selected by default:
 - **Local Snapshots:** If you are performing replication or back up to object storage, local Snapshots must be created.
 - **Replication:** Creates replicated volumes on another ONTAP storage system.
 - **Backup:** Backs up volumes to object storage.
2. **Architecture:** If you chose replication and backup, choose one of the following flows of information:
 - **Cascading:** Information flows from the primary to the secondary and from the secondary to object storage.
 - **Fan out:** Information flows from the primary to the secondary *and* from the primary to object storage.
- For details about these architectures, refer to [Plan your protection journey](#).
3. **Local Snapshot:** Choose an existing snapshot policy or create a new one.



To create a custom policy, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to five schedules, typically of different frequencies.
- Select **Create**.

4. **Replication**: Set the following options:

- **Replication target**: Select the destination system and SVM. Optionally, select the destination aggregate or aggregates and prefix or suffix that will be added to the replicated volume name.
- **Replication policy**: Choose an existing replication policy or create a new one.



To create a custom policy, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to five schedules, typically of different frequencies.
- Select **Create**.

5. **Back up to Object**: If you selected **Backup**, set the following options:

- **Provider**: Select **Google Cloud**.
- **Provider settings**: Enter the provider details and region where the backups will be stored.

Either create a new bucket or select one that you've already created.



If you want to tier older backup files to Google Cloud Archive storage for further cost optimization, ensure that the bucket has the appropriate Lifecycle rule.

Enter the Google Cloud access key and secret key.

- **Encryption key**: If you created a new Google Cloud storage account, enter encryption key information given to you from the provider. Choose whether you'll use the default Google Cloud encryption keys, or choose your own customer-managed keys from your Google Cloud account, to manage encryption of your data.



If you chose an existing Google Cloud storage account, encryption information is already available, so you don't need to enter it now.

If you choose to use your own customer-managed keys, enter the key ring and key name. [Learn more about customer-managed encryption keys](#).

- **Networking**: Choose the IPspace.

The IPspace in the ONTAP cluster where the volumes you want to back up reside. The intercluster LIFs for this IPspace must have outbound internet access.

- **Backup policy**: Select an existing Backup to object storage policy or create a new one.



To create a custom policy, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to five schedules, typically of different frequencies.
- Select **Create**.
- **Export existing snapshots to object storage as backup copies:** If there are any local snapshots for volumes in this system that match the backup schedule label you just selected for this system (for example, daily, weekly, etc.), this additional prompt is displayed. Check this box to have all historic Snapshots copied to object storage as backup files to ensure the most complete protection for your volumes.

6. Select **Next**.

Review your selections

This is the chance to review your selections and make adjustments, if necessary.

Steps

1. In the Review page, review your selections.
2. Optionally check the box to **Automatically synchronize the Snapshot policy labels with the replication and backup policy labels**. This creates snapshots with a label that matches the labels in the replication and backup policies.
3. Select **Activate Backup**.

Result

NetApp Backup and Recovery starts taking the initial backups of your volumes. The baseline transfer of the replicated volume and the backup file includes a full copy of the primary storage system data. Subsequent transfers contain differential copies of the primary storage system data contained in snapshots.

A replicated volume is created in the destination cluster that will be synchronized with the source volume.

A Google Cloud Storage bucket is created automatically in the service account indicated by the Google access key and secret key you entered, and the backup files are stored there. The Volume Backup Dashboard is displayed so you can monitor the state of the backups.

You can also monitor the status of backup and restore jobs using the [Job Monitoring page](#).

Show the API commands

You might want to display and optionally copy the API commands used in the Activate backup and recovery wizard. You might want to do this to automate backup activation in future systems.

Steps

1. From the Activate backup and recovery wizard, select **View API request**.
2. To copy the commands to the clipboard, select the **Copy** icon.

Back up on-premises ONTAP data to ONTAP S3 with NetApp Backup and Recovery

Complete a few steps in NetApp Backup and Recovery to get started backing up volume

data from your primary on-premises ONTAP systems. You can send backups to a secondary ONTAP storage system (a replicated volume) or to a bucket on an ONTAP system configured as an S3 server (a backup file), or both.

The primary on-premises ONTAP system can be a FAS, AFF, or ONTAP Select system. The secondary ONTAP system can be an on-premises ONTAP or Cloud Volumes ONTAP system. The object storage can be on an on-premises ONTAP system or a Cloud Volumes ONTAP system on which you have enabled a Simple Storage Service (S3) object storage server.



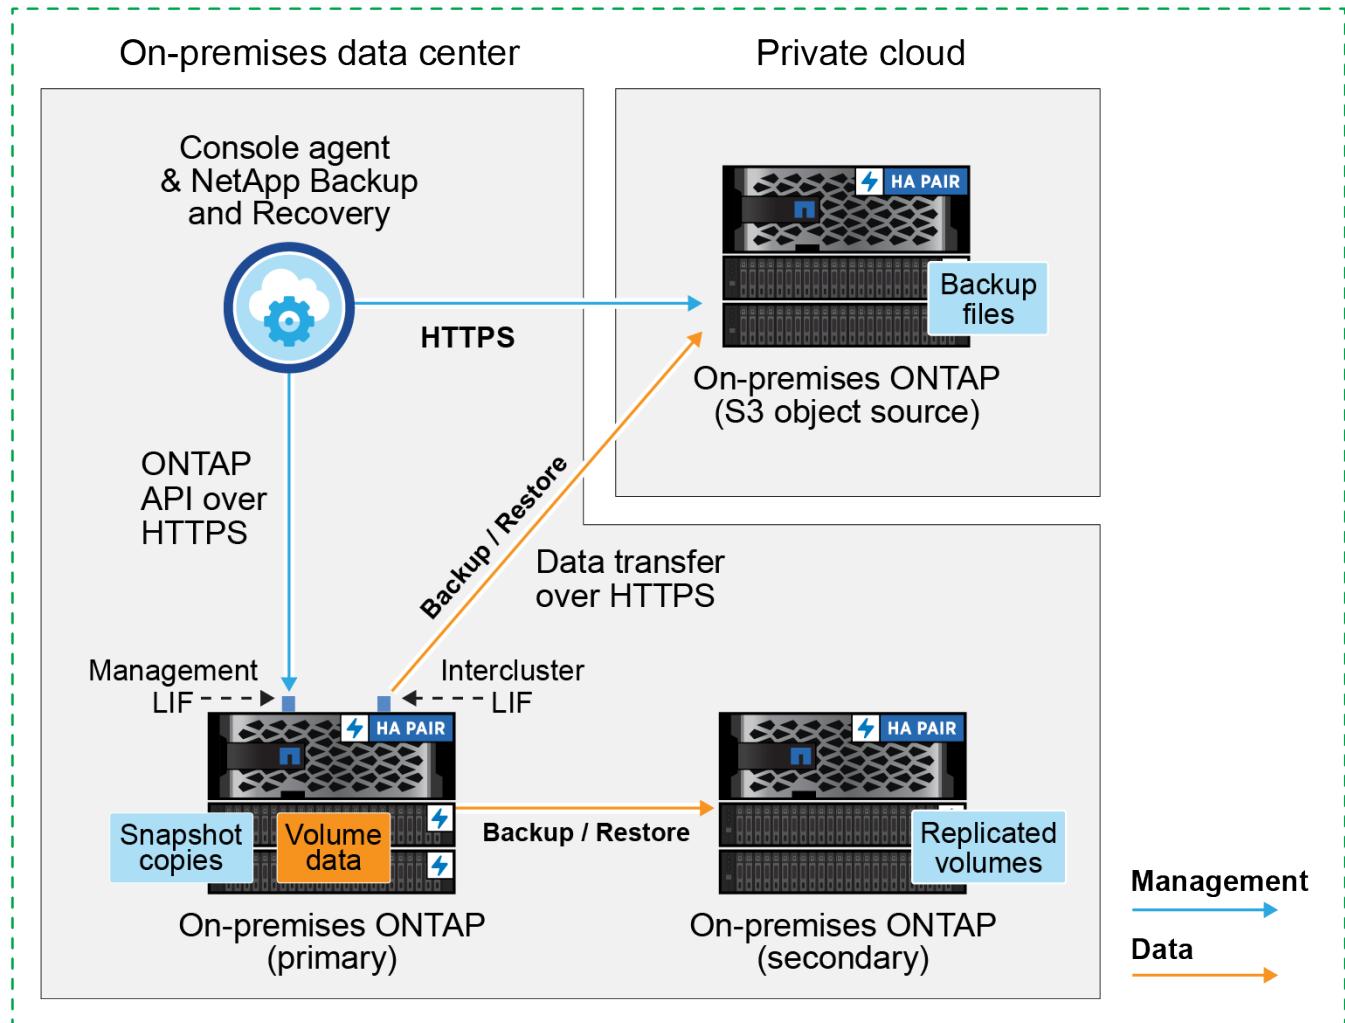
To switch to and from NetApp Backup and Recovery workloads, refer to [Switch to different NetApp Backup and Recovery workloads](#).

Identify the connection method

There are many configurations in which you can create backups to an S3 bucket on an ONTAP system. Two scenarios are shown below.

The following image shows each component when backing up a primary on-premises ONTAP system to an on-premises ONTAP system configured for S3 and the connections that you need to prepare between them. It also shows a connection to a secondary ONTAP system in the same on-premises location to replicate volumes.

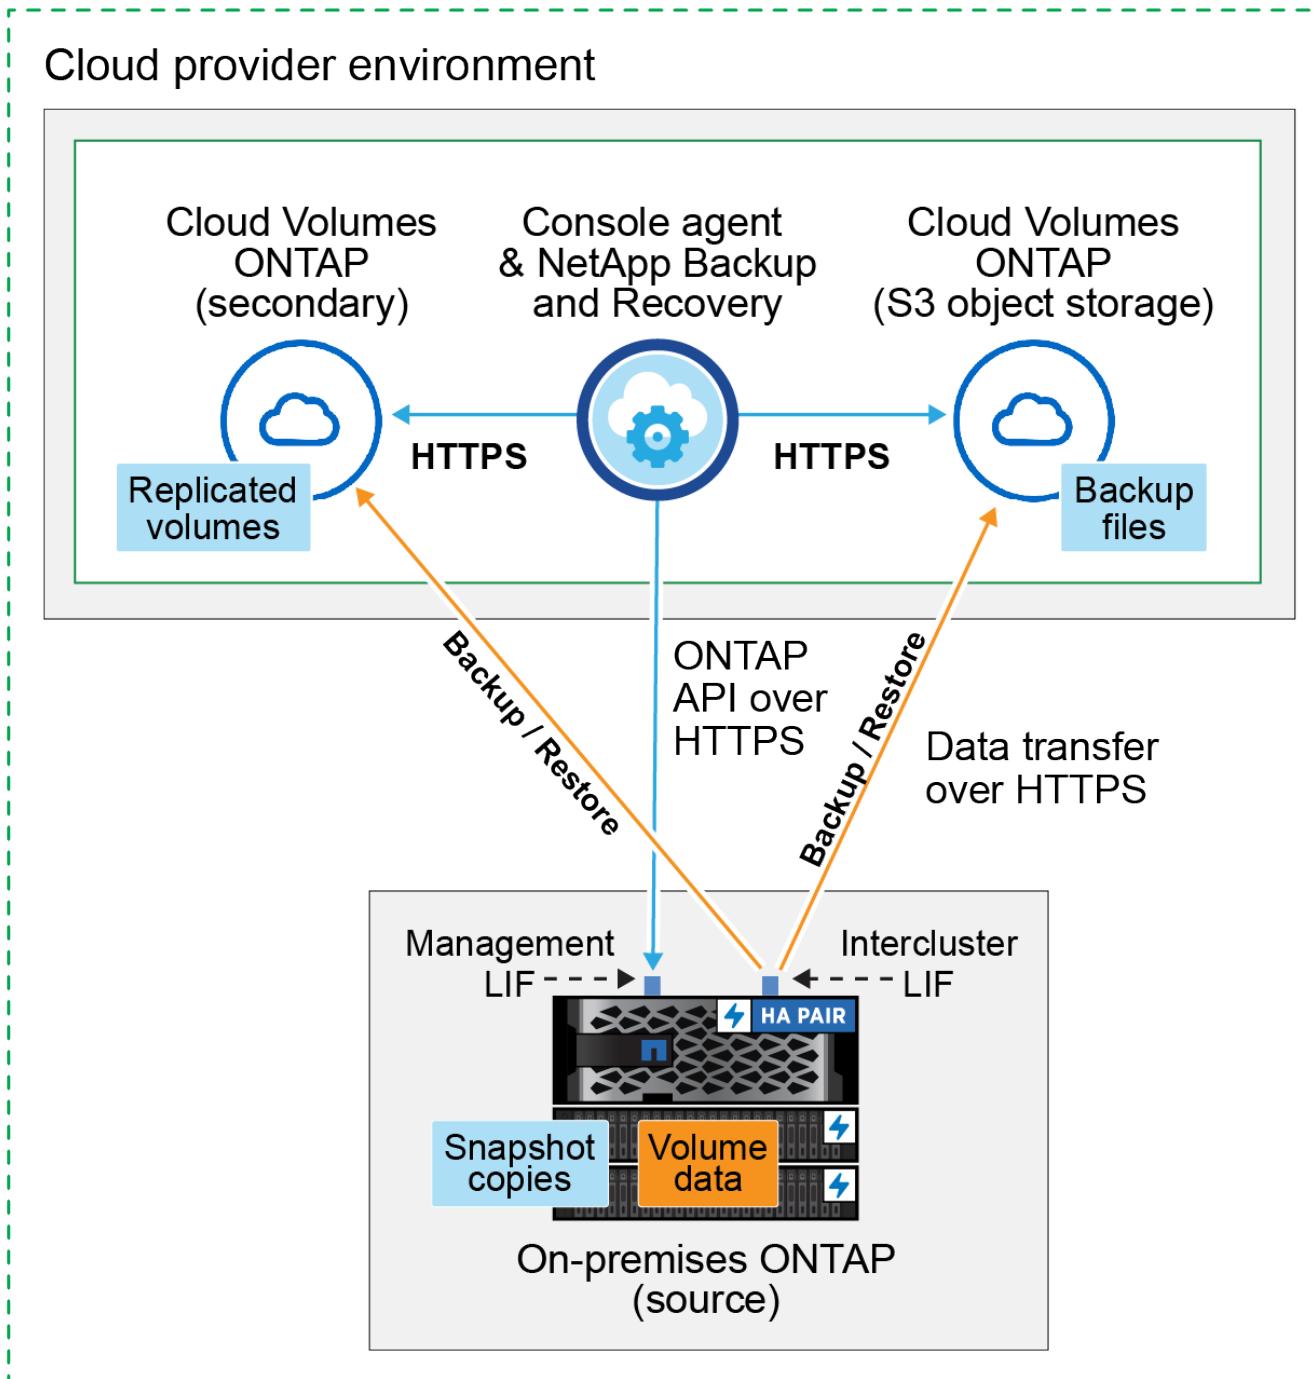
Console agent installed on premises (Public)



When the Console agent and primary on-premises ONTAP system are installed in an on-premises location without internet access (a "private" mode deployment), the ONTAP S3 system must be located in the same on-premises data center.

The following image shows each component when backing up a primary on-premises ONTAP system to a Cloud Volumes ONTAP system configured for S3 and the connections that you need to prepare between them. It also shows a connection to a secondary Cloud Volumes ONTAP system in the same cloud provider environment to replicate volumes.

Console agent deployed in cloud (Public)



In this scenario the Console agent should be deployed in the same cloud provider environment in which the Cloud Volumes ONTAP systems are deployed.

Prepare your Console agent

The Console agent is the main software for Console functionality. A Console agent is required to back up and restore your ONTAP data.

Create or switch Console agents

When you back up data to ONTAP S3, a Console agent must be available on your premises or in the cloud. You'll either need to install a new Console agent or make sure that the currently selected Console agent resides in one of these locations. The on-premises Console agent can be installed in a site with or without internet access.

- [Learn about Console agents](#)
- [Install the Console agent in your cloud environment](#)
- [Installing the Console agent on a Linux host with internet access](#)
- [Installing the Console agent on a Linux host without internet access](#)
- [Switching between Console agents](#)

Prepare Console agent networking requirements

Ensure that the network where the Console agent is installed enables the following connections:

- An HTTPS connection over port 443 to the ONTAP S3 server
- An HTTPS connection over port 443 to your source ONTAP cluster management LIF
- An outbound internet connection over port 443 to NetApp Backup and Recovery (not required when the Console agent is installed in a "dark" site)

Private mode (dark site) considerations

NetApp Backup and Recovery functionality is built into the Console agent. When it is installed in private mode, you'll need to update the Console agent software periodically to get access to new features. Check the [NetApp Backup and Recovery What's New](#) to see the new features in each NetApp Backup and Recovery release. When you want to use the new features, follow the steps to [upgrade the Console agent software](#).

When you use NetApp Backup and Recovery in a standard SaaS environment, the NetApp Backup and Recovery configuration data is backed up to the cloud. When you use NetApp Backup and Recovery in a site with no internet access, the NetApp Backup and Recovery configuration data is backed up to the ONTAP S3 bucket where your backups are being stored.

Verify license requirements

Before you can activate NetApp Backup and Recovery for your cluster, you'll need to purchase and activate a NetApp Backup and Recovery BYOL license from NetApp. The license is for backup and restore to object storage - no license is needed to create snapshots or replicated volumes. This license is for the account and can be used across multiple systems.

You'll need the serial number from NetApp that enables you to use the service for the duration and capacity of the license. [Learn how to manage your BYOL licenses](#).



PAYGO licensing is not supported when backing up files to ONTAP S3.

Prepare your ONTAP clusters

Unresolved directive in prev-ontap-backup-onprem-ontaps3.adoc - include:::_include/backup-onprem-prepare-onprem-ONTAP-cluster.adoc[]

Verify ONTAP networking requirements for backing up data to object storage

You must ensure that the following requirements are met on the system that connects to object storage.

- When you use a fan-out backup architecture, the settings must be configured on the *primary* storage system.
- When you use a cascaded backup architecture, the settings must be configured on the *secondary* storage system.

[Learn more about the types of backup architecture.](#)

The following ONTAP cluster networking requirements are needed:

- The ONTAP cluster initiates an HTTPS connection over a user-specified port from the intercluster LIF to the ONTAP S3 server for backup and restore operations. The port is configurable during backup setup. ONTAP reads and writes data to and from object storage. The object storage never initiates, it just responds.
- ONTAP requires an inbound connection from the Console agent to the cluster management LIF.
- An intercluster LIF is required on each ONTAP node that hosts the volumes you want to back up. The LIF must be associated with the *IPspace* that ONTAP should use to connect to object storage. [Learn more about IPspaces](#).

When you set up NetApp Backup and Recovery, you are prompted for the *IPspace* to use. You should choose the *IPspace* that each LIF is associated with. That might be the "Default" *IPspace* or a custom *IPspace* that you created.

- The nodes' intercluster LIFs are able to access the object store (not required when the Console agent is installed in a "dark" site).
- DNS servers have been configured for the storage VM where the volumes are located. See how to [configure DNS services for the SVM](#).
- If you use are using a different *IPspace* than Default, then you might need to create a static route to get access to the object storage.
- Update firewall rules, if necessary, to allow NetApp Backup and Recovery service connections from ONTAP to object storage through the port you specified (typically port 443) and name resolution traffic from the storage VM to the DNS server over port 53 (TCP/UDP).

Verify ONTAP networking requirements for replicating volumes

Unresolved directive in prev-ontap-backup-onprem-ontaps3.adoc - include::../_include/backup-onprem-prepare-source-destination-systems-for-repl.adoc[]

Prepare ONTAP S3 as your backup target

You must enable a Simple Storage Service (S3) object storage server in the ONTAP cluster that you plan to use for object storage backups. See the [ONTAP S3 documentation](#) for details.

Note: You can add this cluster to the Console **Systems** page, but it is not identified as being an S3 object storage server, and you can't drag and drop a source system onto this S3 system to initiate backup activation.

This ONTAP system must meet the following requirements.

Supported ONTAP versions

ONTAP 9.8 and later is required for on-premises ONTAP systems.
ONTAP 9.9.1 and later is required for Cloud Volumes ONTAP systems.

S3 credentials

You must have created an S3 user to control access to your ONTAP S3 storage. [See the ONTAP S3 docs for details.](#)

When you set up backup to ONTAP S3, the backup wizard prompts you for an S3 access key and secret key for a user account. The user account enables NetApp Backup and Recovery to authenticate and access the ONTAP S3 buckets used to store backups. The keys are required so that ONTAP S3 knows who is making the request.

These access keys must be associated with a user who has the following permissions:

```
"s3>ListAllMyBuckets",
"s3>ListBucket",
"s3GetObject",
"s3PutObject",
"s3DeleteObject",
"s3CreateBucket",
"s3GetBucketLocation"
```

Activate backups on your ONTAP volumes

Activate backups at any time directly from your on-premises system.

A wizard takes you through the following major steps:

- Select the volumes that you want to back up
- Define the backup strategy and policies
- Review your selections

You can also [Show the API commands](#) at the review step, so you can copy the code to automate backup activation for future systems.

Start the wizard

Steps

1. Access the Activate backup and recovery wizard using one of the following ways:

- From the **Console Systems** page, select the system and select **Enable > Backup Volumes** next to **Backup and Recovery** in the right-panel.
- Select **Volumes** in the **Backup and recovery** bar. From the **Volumes** tab, select the **Actions (...)** option and select **Activate Backup** for a single volume (that does not already have replication or backup to object storage enabled).

The Introduction page of the wizard shows the protection options including local snapshots, replications, and backups. If you did the second option in this step, the Define Backup Strategy page appears with one volume selected.

2. Continue with the following options:

- If you already have a Console agent, you're all set. Just select **Next**.
- If you don't have a Console agent, the **Add a Console agent** option appears. Refer to [Prepare your Console agent](#).

Select the volumes that you want to back up

Choose the volumes you want to protect. A protected volume is one that has one or more of the following: snapshot policy, replication policy, backup to object policy.

You can choose to protect FlexVol or FlexGroup volumes; however, you cannot select a mix of these volumes when activating backup for a system. See how to [activate backup for additional volumes in the system](#) (FlexVol or FlexGroup) after you have configured backup for the initial volumes.



- You can activate a backup only on a single FlexGroup volume at a time.
- The volumes you select must have the same SnapLock setting. All volumes must have SnapLock Enterprise enabled or have SnapLock disabled.

Steps

Note that if the volumes you choose already have Snapshot or replication policies applied, then the policies you select later will overwrite these existing policies.

1. In the Select Volumes page, select the volume or volumes you want to protect.

- Optionally, filter the rows to show only volumes with certain volume types, styles, and more to make the selection easier.
- After you select the first volume, then you can select all FlexVol volumes (FlexGroup volumes can be selected one at a time only). To back up all existing FlexVol volumes, check one volume first and then check the box in the title row.
- To back up individual volumes, check the box for each volume.

2. Select **Next**.

Define the backup strategy

Defining the backup strategy involves configuring the following options:

- Protection options: Whether you want to implement one or all of the backup options: local snapshots, replication, and backup to object storage
- Architecture: Whether you want to use a fan-out or cascading backup architecture
- Local snapshot policy
- Replication target and policy
- Backup to object storage information (provider, encryption, networking, backup policy, and export options).

Steps

1. In the Define Backup Strategy page, choose one or all of the following. All three are selected by default:

- **Local Snapshots:** Creates local snapshots.
- **Replication:** Creates replicated volumes on another ONTAP storage system.
- **Backup:** Backs up volumes to a bucket on an ONTAP system configured for S3.

2. **Architecture:** If you chose both replication and backup, choose one of the following flows of information:

- **Cascading:** Backup data flows from the primary to the secondary system, and then from the secondary to object storage.
- **Fan out:** Backup data flows from the primary to the secondary system *and* from the primary to object storage.

For details about these architectures, refer to [Plan your protection journey](#).

3. **Local Snapshot:** Choose an existing snapshot policy or create a new one.



If you want to create a custom policy before activating the Snapshot, you can use System Manager or the ONTAP CLI `snapmirror policy create` command. Refer to [ONTAP CLI for snapmirror policy](#).



To create a custom policy using Backup and Recovery, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to five schedules, typically of different frequencies.
- Select **Create**.

4. **Replication:** If you selected **Replication**, set the following options:

- **Replication target:** Select the destination system and SVM. Optionally, select the destination aggregate (or aggregates for FlexGroup volumes) and a prefix or suffix that will be added to the replicated volume name.
- **Replication policy:** Choose an existing replication policy or create a new one.

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to five schedules, typically of different frequencies.
- Select **Create**.

5. **Back up to Object:** If you selected **Backup**, set the following options:

- **Provider:** Select **ONTAP S3**.
- **Provider settings:** Enter the S3 server FQDN details, port, and the users' access key and secret key.

The access key and secret key are for the user you created to give the ONTAP cluster access to the S3 bucket.

- **Networking:** Choose the IPspace in the source ONTAP cluster where the volumes you want to back up reside. The intercluster LIFs for this IPspace must have outbound internet access (not required when the Console agent is installed in a "dark" site).



Selecting the correct IPspace ensures that NetApp Backup and Recovery can set up a connection from ONTAP to your ONTAP S3 object storage.

- **Backup policy:** Select an existing backup policy or create a new one.



You can create a policy with System Manager or the ONTAP CLI. To create a custom policy using the ONTAP CLI `snmpmirror policy create` command, refer to [ONTAP CLI for snapmirror policy](#).



To create a custom policy using Backup and Recovery, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to five schedules, typically of different frequencies.
- For backup-to-object policies, set the DataLock and Ransomware Resilience settings. For details on DataLock and Ransomware Resilience, refer to [Backup-to-object policy settings](#).
- Select **Create**.

◦ **Export existing snapshots to object storage as backup files:** If there are any local snapshots for volumes in this system that match the backup schedule label you just selected (for example, daily, weekly, etc.), this additional prompt is displayed. Check this box to have all historic Snapshots copied to object storage as backup files to ensure the most complete protection for your volumes.

6. Select **Next**.

Review your selections

This is the chance to review your selections and make adjustments, if necessary.

Steps

1. In the Review page, review your selections.
2. Optionally check the box to **Automatically synchronize the Snapshot policy labels with the replication and backup policy labels**. This creates Snapshots with a label that matches the labels in the replication and backup policies. If the policies don't match, backups will not be created.
3. Select **Activate Backup**.

Result

NetApp Backup and Recovery starts taking the initial backups of your volumes. The baseline transfer of the replicated volume and the backup file includes a full copy of the source data. Subsequent transfers contain differential copies of the primary storage data contained in snapshots.

A replicated volume is created in the destination cluster that will be synchronized with the primary storage volume.

An S3 bucket is created in the service account indicated by the S3 access key and secret key you entered, and the backup files are stored there.

The Volume Backup Dashboard is displayed so you can monitor the state of the backups.

You can also monitor the status of backup and restore jobs using the [Job Monitoring page](#).

Show the API commands

You might want to display and optionally copy the API commands used in the Activate backup and recovery wizard. You might want to do this to automate backup activation in future systems.

Steps

1. From the Activate backup and recovery wizard, select **View API request**.
2. To copy the commands to the clipboard, select the **Copy** icon.

Back up on-premises ONTAP data to StorageGRID with NetApp Backup and Recovery

Complete a few steps in NetApp Backup and Recovery to get started backing up volume data from your on-premises primary ONTAP systems to a secondary storage system and to object storage in your NetApp StorageGRID systems.



"On-premises ONTAP systems" include FAS, AFF, and ONTAP Select systems.

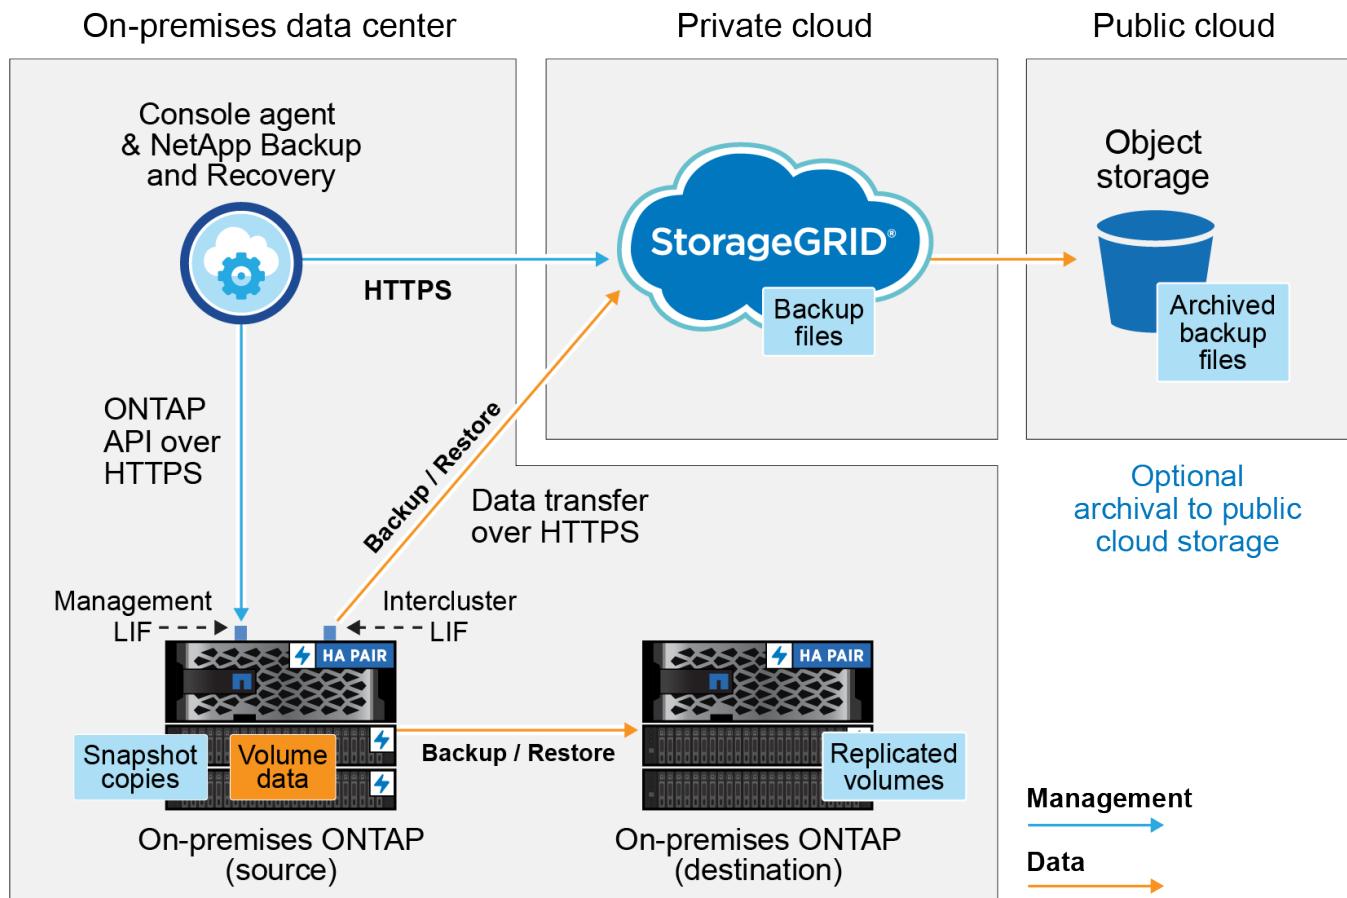


To switch to and from NetApp Backup and Recovery workloads, refer to [Switch to different NetApp Backup and Recovery workloads](#).

Identify the connection method

The following image shows each component when backing up an on-premises ONTAP system to StorageGRID and the connections that you need to prepare between them.

Optionally, you can connect to a secondary ONTAP system in the same on-premises location to replicate volumes.



When the Console agent and on-premises ONTAP system are installed in an on-premises location without internet access (a "dark site"), the StorageGRID system must be located in the same on-premises data center. Archival of older backup files to public cloud is not supported in dark site configurations.

Prepare your Console agent

The Console agent is the main software for Console functionality. A Console agent is required to back up and restore your ONTAP data.

Create or switch Console agents

When you back up data to StorageGRID, a Console agent must be available on your premises. You'll either need to install a new Console agent or make sure that the currently selected Console agent resides on-premises. The Console agent can be installed in a site with or without internet access.

- [Learn about Console agents](#)
- [Installing the Console agent on a Linux host with internet access](#)
- [Installing the Console agent on a Linux host without internet access](#)
- [Switching between Console agents](#)

Prepare Console agent networking requirements

Ensure that the network where the Console agent is installed enables the following connections:

- An HTTPS connection over port 443 to the StorageGRID Gateway Node
- An HTTPS connection over port 443 to your ONTAP cluster management LIF
- An outbound internet connection over port 443 to NetApp Backup and Recovery (not required when the Console agent is installed in a "dark" site)

Private mode (dark site) considerations

- NetApp Backup and Recovery functionality is built into the Console agent. When it is installed in private mode, you'll need to update the Console agent software periodically to get access to new features. Check the [NetApp Backup and Recovery What's New](#) to see the new features in each NetApp Backup and Recovery release. When you want to use the new features, follow the steps to [upgrade the Console agent software](#).

The new version of NetApp Backup and Recovery that includes the ability to schedule and create snapshots and replicated volumes, in addition to creating backups to object storage, requires that you are using version 3.9.31 or greater of the Console agent. So it is recommended that you get this newest release to manage all your backups.

- When you use NetApp Backup and Recovery in a SaaS environment, the NetApp Backup and Recovery configuration data is backed up to the cloud. When you use NetApp Backup and Recovery in a site with no internet access, the NetApp Backup and Recovery configuration data is backed up to the StorageGRID bucket where your backups are being stored.

Verify license requirements

Before you can activate NetApp Backup and Recovery for your cluster, you'll need to purchase and activate a NetApp Backup and Recovery BYOL license from NetApp. This license is for the account and can be used across multiple systems.

You'll need the serial number from NetApp that enables you to use the service for the duration and capacity of the license. [Learn how to manage your BYOL licenses](#).



PAYGO licensing is not supported when backing up files to StorageGRID.

Prepare your ONTAP clusters

Unresolved directive in prev-ontap-backup-onprem-storagegrid.adoc - include:../_include/backup-onprem-prepare-onprem-ONTAP-cluster.adoc[]

Verify ONTAP networking requirements for backing up data to object storage

You must configure the following requirements on the system that connects to object storage.

- When you use a fan-out backup architecture, the following settings must be configured on the *primary* storage system.
- When you use a cascaded backup architecture, the following settings must be configured on the *secondary* storage system.

The following ONTAP cluster networking requirements are needed:

- The ONTAP cluster initiates an HTTPS connection over a user-specified port from the intercluster LIF to the StorageGRID Gateway Node for backup and restore operations. The port is configurable during backup setup.
- ONTAP reads and writes data to and from object storage. The object storage never initiates, it just responds.
- ONTAP requires an inbound connection from the Console agent to the cluster management LIF. The Console agent must reside on your premises.
- An intercluster LIF is required on each ONTAP node that hosts the volumes you want to back up. The LIF must be associated with the *IPspace* that ONTAP should use to connect to object storage. [Learn more about IPspaces](#).

When you set up NetApp Backup and Recovery, you are prompted for the *IPspace* to use. You should choose the *IPspace* that each LIF is associated with. That might be the "Default" *IPspace* or a custom *IPspace* that you created.

- The nodes' intercluster LIFs are able to access the object store (not required when the Console agent is installed in a "dark" site).
- DNS servers have been configured for the storage VM where the volumes are located. See how to [configure DNS services for the SVM](#).
- If you use are using a different *IPspace* than the Default, then you might need to create a static route to get access to the object storage.
- Update firewall rules, if necessary, to allow NetApp Backup and Recovery service connections from ONTAP to object storage through the port you specified (typically port 443) and name resolution traffic from the storage VM to the DNS server over port 53 (TCP/UDP).

Verify ONTAP networking requirements for replicating volumes

Unresolved directive in prev-ontap-backup-onprem-storagegrid.adoc - include:../_include/backup-onprem-prepare-source-destination-systems-for-repl.adoc[]

Prepare StorageGRID as your backup target

StorageGRID must meet the following requirements. See the [StorageGRID documentation](#) for more information.

For details about DataLock and Ransomware Resilience requirements for StorageGRID, refer to [Backup-to-object policy options](#).

Supported StorageGRID versions

StorageGRID 10.3 and later is supported.

To use DataLock & Ransomware Resilience for your backups, your StorageGRID systems must be running version 11.6.0.3 or greater.

To tier older backups to cloud archival storage, your StorageGRID systems must be running version 11.3 or greater. Additionally, your StorageGRID systems must be discovered to the [Console Systems](#) page.

To use archival storage, admin node IP access is needed.

Gateway IP access is always needed.

S3 credentials

You must have created an S3 tenant account to control access to your StorageGRID storage. [See the StorageGRID docs for details](#).

When you set up backup to StorageGRID, the backup wizard prompts you for an S3 access key and secret key for a tenant account. The tenant account enables NetApp Backup and Recovery to authenticate and access the StorageGRID buckets used to store backups. The keys are required so that StorageGRID knows who is making the request.

These access keys must be associated with a user who has the following permissions:

```
"s3>ListAllMyBuckets",
"s3>ListBucket",
"s3.GetObject",
"s3.PutObject",
"s3>DeleteObject",
"s3>CreateBucket"
```

Object versioning

You must not enable StorageGRID object versioning manually on the object store bucket.

Prepare to archive older backup files to public cloud storage

Tiering older backup files to archival storage saves money by using a less expensive storage class for backups that you may not need. StorageGRID is an on-premises (private cloud) solution that doesn't provide archival storage, but you can move older backup files to public cloud archival storage. When used in this fashion, data that is tiered to cloud storage, or restored from cloud storage, goes between StorageGRID and the cloud storage - the Console is not involved in this data transfer.

Current support enables you to archive backups to AWS S3 *Glacier*/S3 *Glacier Deep Archive* or Azure Archive storage.

ONTAP Requirements

- Your cluster must be using ONTAP 9.12.1 or greater.

StorageGRID Requirements

- Your StorageGRID must be using 11.4 or greater.
- Your StorageGRID must be [discovered and available in the Console](#).

Amazon S3 requirements

- You'll need to sign up for an Amazon S3 account for the storage space where your archived backups will be located.
- You can choose to tier backups to AWS S3 Glacier or S3 Glacier Deep Archive storage. [Learn more about AWS archival tiers](#).
- StorageGRID should have full-control access to the bucket (s3: *); however, if this is not possible, the bucket policy must grant the following S3 permissions to StorageGRID:
 - s3:AbortMultipartUpload
 - s3:DeleteObject
 - s3:GetObject
 - s3>ListBucket
 - s3>ListBucketMultipartUploads
 - s3>ListMultipartUploadParts
 - s3:PutObject
 - s3:RestoreObject

Azure Blob requirements

- You'll need to sign up for an Azure Subscription for the storage space where your archived backups will be located.
- The activation wizard enables you to use an existing Resource Group to manage the Blob container that will store the backups, or you can create a new Resource Group.

When defining the Archival settings for the backup policy for your cluster, you'll enter your cloud provider credentials and select the storage class that you want to use. NetApp Backup and Recovery creates the cloud bucket when you activate backup for the cluster. The information required for AWS and Azure archival storage is shown below.

AWS	Azure
<input checked="" type="checkbox"/> Tier Backups to Archive Cloud Provider <div style="border: 1px solid #ccc; padding: 2px; width: 150px;">AWS</div> Account <div style="border: 1px solid #ccc; padding: 2px; width: 150px;">Select Account</div> AWS Access Key <div style="border: 1px solid #ccc; padding: 2px; width: 150px;">Enter AWS Access Key</div> Archive After (Days) <div style="border: 1px solid #ccc; padding: 2px; width: 150px;">(1-999)</div>	<input checked="" type="checkbox"/> Tier Backups to Archive Cloud Provider <div style="border: 1px solid #ccc; padding: 2px; width: 150px;">AZURE</div> Azure Subscription <div style="border: 1px solid #ccc; padding: 2px; width: 150px;">Select Account</div> Region <div style="border: 1px solid #ccc; padding: 2px; width: 150px;">Select Region</div> Resource Group Type <div style="border: 1px solid #ccc; padding: 2px; width: 150px;">Select an Existing Resource Group</div> Resource Group <div style="border: 1px solid #ccc; padding: 2px; width: 150px;">Select Resource Group</div> Archive After (Days) <div style="border: 1px solid #ccc; padding: 2px; width: 150px;">(1-999)</div> Storage Class <div style="border: 1px solid #ccc; padding: 2px; width: 150px;">S3 Glacier</div>

The archival policy settings you select will generate an information lifecycle management (ILM) policy in StorageGRID, and add the settings as "rules."

- If there is an existing active ILM policy, new rules will be added to the ILM policy to move the data to the archive tier.
- If there is an existing ILM policy in the "proposed" state, the creation and activation of a new ILM policy will not be possible. [Learn more about StorageGRID ILM policies and rules](#).

Activate backups on your ONTAP volumes

Activate backups at any time directly from your on-premises system.

A wizard takes you through the following major steps:

- [Select the volumes that you want to back up](#)
- [Define the backup strategy](#)
- [Review your selections](#)

You can also [Show the API commands](#) at the review step, so you can copy the code to automate backup activation for future systems.

Start the wizard

Steps

1. Access the Activate backup and recovery wizard using one of the following ways:

- From the **Console Systems** page, select the system and select **Enable > Backup Volumes** next to **Backup and Recovery** in the right-panel.

If the destination for your backups exists as a system on the **Console Systems** page, you can drag the ONTAP cluster onto the object storage.

- Select **Volumes** in the **Backup and recovery** bar. From the **Volumes** tab, select the **Actions (...)** option and select **Activate Backup** for a single volume (that does not already have replication or backup to object storage already enabled).

The Introduction page of the wizard shows the protection options including local Snapshots, replication, and backups. If you did the second option in this step, the Define Backup Strategy page appears with one volume selected.

2. Continue with the following options:

- If you already have a Console agent, you're all set. Just select **Next**.
- If you don't already have a Console agent, the **Add a Console agent** option appears. Refer to [Prepare your Console agent](#).

Select the volumes that you want to back up

Choose the volumes you want to protect. A protected volume is one that has one or more of the following: Snapshot policy, replication policy, backup to object policy.

You can choose to protect FlexVol or FlexGroup volumes; however, you cannot select a mix of these volumes when activating backup for a system. See how to [activate backup for additional volumes in the system](#) (FlexVol or FlexGroup) after you have configured backup for the initial volumes.



- You can activate a backup only on a single FlexGroup volume at a time.
- The volumes you select must have the same SnapLock setting. All volumes must have SnapLock Enterprise enabled or have SnapLock disabled.

Steps

If the volumes you choose already have snapshot or replication policies applied, then the policies you select later will overwrite these existing policies.

1. In the Select Volumes page, select the volume or volumes you want to protect.

- Optionally, filter the rows to show only volumes with certain volume types, styles, and more to make the selection easier.
- After you select the first volume, then you can select all FlexVol volumes (FlexGroup volumes can be selected one at a time only). To back up all existing FlexVol volumes, check one volume first and then check the box in the title row.
- To back up individual volumes, check the box for each volume.

2. Select **Next**.

Define the backup strategy

Defining the backup strategy involves setting the following options:

- Whether you want one or all of the backup options: local Snapshots, replication, and backup to object storage
- Architecture
- Local snapshot policy
- Replication target and policy



If the volumes you choose have different snapshot and replication policies than the policies you select in this step, the existing policies will be overwritten.

- Backup to object storage information (provider, encryption, networking, backup policy, and export options).

Steps

1. In the Define backup strategy page, choose one or all of the following. All three are selected by default:

- **Local Snapshots:** If you are performing replication or back up to object storage, local snapshots must be created.
- **Replication:** Creates replicated volumes on another ONTAP storage system.
- **Backup:** Backs up volumes to object storage.

2. **Architecture:** If you chose both replication and backup, choose one of the following flows of information:

- **Cascading:** Information flows from the primary to the secondary, and then from the secondary to object storage.
- **Fan out:** Information flows from the primary to the secondary *and* from the primary to object storage.

For details about these architectures, refer to [Plan your protection journey](#).

3. **Local Snapshot:** Choose an existing snapshot policy or create a new one.



To create a custom policy, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to five schedules, typically of different frequencies.
- Select **Create**.

4. **Replication:** Set the following options:

- **Replication target:** Select the destination system and SVM. Optionally, select the destination aggregate or aggregates and prefix or suffix that will be added to the replicated volume name.
- **Replication policy:** Choose an existing replication policy or create one.



To create a custom policy, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to five schedules, typically of different frequencies.
- Select **Create**.

5. **Back up to Object:** If you selected **Backup**, set the following options:

- **Provider:** Select **StorageGRID**.
- **Provider settings:** Enter the provider gateway node FQDN details, port, access key and secret key.

The access key and secret key are for the IAM user you created to give the ONTAP cluster access to the bucket.

- **Networking:** Choose the IPspace in the ONTAP cluster where the volumes you want to back up reside. The intercluster LIFs for this IPspace must have outbound internet access (not required when the Console agent is installed in a "dark" site).



Selecting the correct IPspace ensures that NetApp Backup and Recovery can set up a connection from ONTAP to your StorageGRID object storage.

- **Backup policy:** Select an existing Backup to object storage policy or create one.



To create a custom policy, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to five schedules, typically of different frequencies.
- For backup-to-object policies, set the DataLock and Ransomware Resilience settings. For details on DataLock and Ransomware Resilience, refer to [Backup-to-object policy settings](#).

If your cluster is using ONTAP 9.11.1 or greater, you can choose to protect your backups from deletion and ransomware attacks by configuring *DataLock* and *Ransomware Resilience*. *DataLock* protects your backup files from being modified or deleted, and *Ransomware Resilience* scans your backup files to look for evidence of a ransomware attack in your backup files.

- Select **Create**.

If your cluster is using ONTAP 9.12.1 or greater and your StorageGRID system is using version 11.4 or greater, you can choose to tier older backups to public cloud archive tiers after a certain number of days. Current support is for AWS S3 Glacier/S3 Glacier Deep Archive or Azure Archive storage tiers. [See how to configure your systems for this functionality](#).

- **Tier backup to public cloud:** Select the cloud provider that you want to tier backups to and enter the provider details.

Select or create a new StorageGRID cluster. For details about creating a StorageGRID cluster so the Console can discover it, refer to [StorageGRID documentation](#).

- **Export existing snapshots to object storage as backup copies:** If there are any local snapshots for volumes in this system that match the backup schedule label you just selected for this system (for example, daily, weekly, etc.), this additional prompt is displayed. Check this box to have all historic snapshots copied to object storage as backup files to ensure the most complete protection for your volumes.

6. Select **Next**.

Review your selections

This is the chance to review your selections and make adjustments, if necessary.

Steps

1. In the Review page, review your selections.
2. Optionally check the box to **Automatically synchronize the Snapshot policy labels with the replication and backup policy labels**. This creates Snapshots with a label that matches the labels in the replication and backup policies.
3. Select **Activate Backup**.

Result

NetApp Backup and Recovery starts taking the initial backups of your volumes. The baseline transfer of the replicated volume and the backup file includes a full copy of the source data. Subsequent transfers contain differential copies of the primary storage data contained in snapshots.

A replicated volume is created in the destination cluster that will be synchronized with the primary storage volume.

An S3 bucket is created in the service account indicated by the S3 access key and secret key you entered, and the backup files are stored there.

The Volume Backup Dashboard is displayed so you can monitor the state of the backups.

You can also monitor the status of backup and restore jobs using the [Job Monitoring page](#).

Show the API commands

You might want to display and optionally copy the API commands used in the Activate backup and recovery wizard. You might want to do this to automate backup activation in future systems.

Steps

1. From the Activate backup and recovery wizard, select **View API request**.
2. To copy the commands to the clipboard, select the **Copy** icon.

Migrate volumes using SnapMirror to Cloud Resync in NetApp Backup and Recovery

The SnapMirror to Cloud Resync feature in NetApp Backup and Recovery streamlines data protection and continuity during volume migrations in NetApp environments. When a volume is migrated using SnapMirror Logical Replication (LRSE) from one on-premises NetApp deployment to another, or to a cloud-based solution such as Cloud Volumes ONTAP, SnapMirror to Cloud Resync ensures that existing cloud backups remain intact and operational.

This feature removes the need for a re-baseline process and lets backups continue after migration. This feature is valuable in workload migration scenarios, supporting both FlexVols and FlexGroups, and is available starting with ONTAP version 9.16.1.



This feature is available starting with NetApp Backup and Recovery version 4.0.3 released May 2025.

SnapMirror to Cloud Resync maintains backup continuity across environments, making it easier to manage data in hybrid and multi-cloud setups.



To switch to and from NetApp Backup and Recovery workloads, refer to [Switch to different NetApp Backup and Recovery workloads](#).

Before you begin

Ensure that these prerequisites have been met:

- The destination ONTAP cluster must be running ONTAP version 9.16.1 or later.
- The old source ONTAP cluster must be protected using NetApp Backup and Recovery.
- The SnapMirror to Cloud Resync feature is available starting with NetApp Backup and Recovery version 4.0.3 released May 2025.
- Ensure that the latest backup in the object storage is the common snapshot across the old source, the new

source, and the object store. Do not use a common snapshot that is older than the latest snapshot backed up to the object store.

- Both the snapshot and SnapMirror policies used on the older ONTAP cluster must be created on the new ONTAP cluster before starting the resync operation. If you use any policy in the resync process, you must also create that policy. The Resync operation does not create policies.
- Ensure that the SnapMirror policy that is applied to the migration volume SnapMirror relationship includes the same label that the cloud relationship uses. To avoid issues, use the policy that governs an exact mirror of the volume and all snapshots.



SnapMirror to Cloud Resync after migrations using SVM-Migrate, SVM-DR, or Head Swap methods are not currently supported.

How NetApp Backup and Recovery SnapMirror to Cloud Resync works

If you complete a technical refresh or migrate volumes from one ONTAP cluster to another, it's important that your backups continue to work without interruption. NetApp Backup and Recovery SnapMirror to Cloud Resync helps with this by ensuring that your cloud backups stay consistent even after a volume migration.

Here's an example:

Imagine you have an on-premises volume called Vol1a. This volume has three snapshots: S1, S2, and S3. These snapshots are restore points. Vol1 is backed up to the cloud using SnapMirror to Cloud (SM-C), but only S1 and S2 are in the object store.

Now, you want to migrate Vol1 to another ONTAP cluster. To do this, you create a SnapMirror Logical Replication (LRSE) relationship to a new cloud volume called Vol1b. This transfers all three snapshots—S1, S2, and S3—from Vol1a to Vol1b.

After the migration is complete, you have the following setup:

- The original SM-C relationship (Vol1a → Object store) is deleted.
- The LRSE relationship (Vol1a → Vol1b) is also deleted.
- Vol1b is now your active volume.

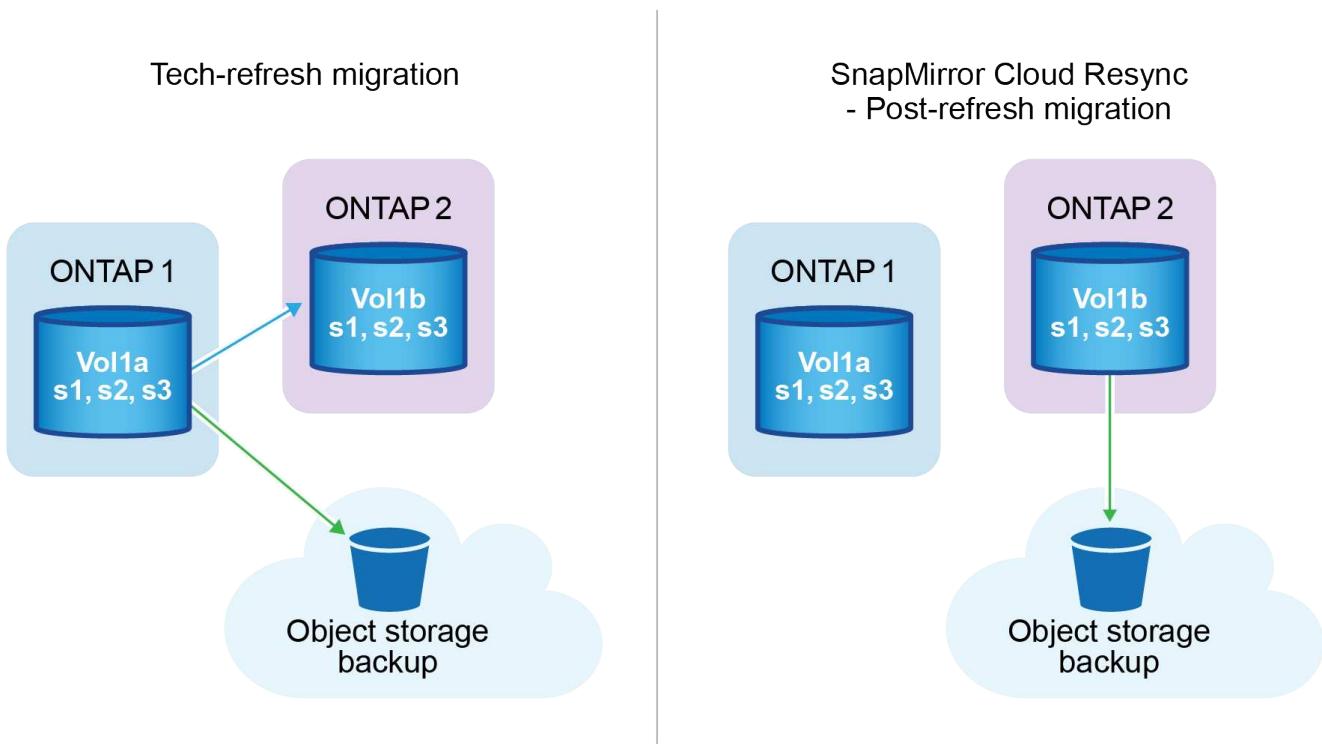
At this point, you want to continue backing up Vol1b to the same cloud endpoint. But instead of starting a full backup from scratch (which would take time and resources), you use SnapMirror to Cloud Resync.

Here's how the resync works:

- The system checks for a common snapshot between Vol1a and Object store. In this case, both have S2.
- Because of this shared snapshot, the system needs to transfer only the incremental changes between S2 and S3.

This means only the new data added after S2 is sent to object store, not the entire volume.

This process prevents duplicate backups, saves bandwidth, and keeps backups running after migration.



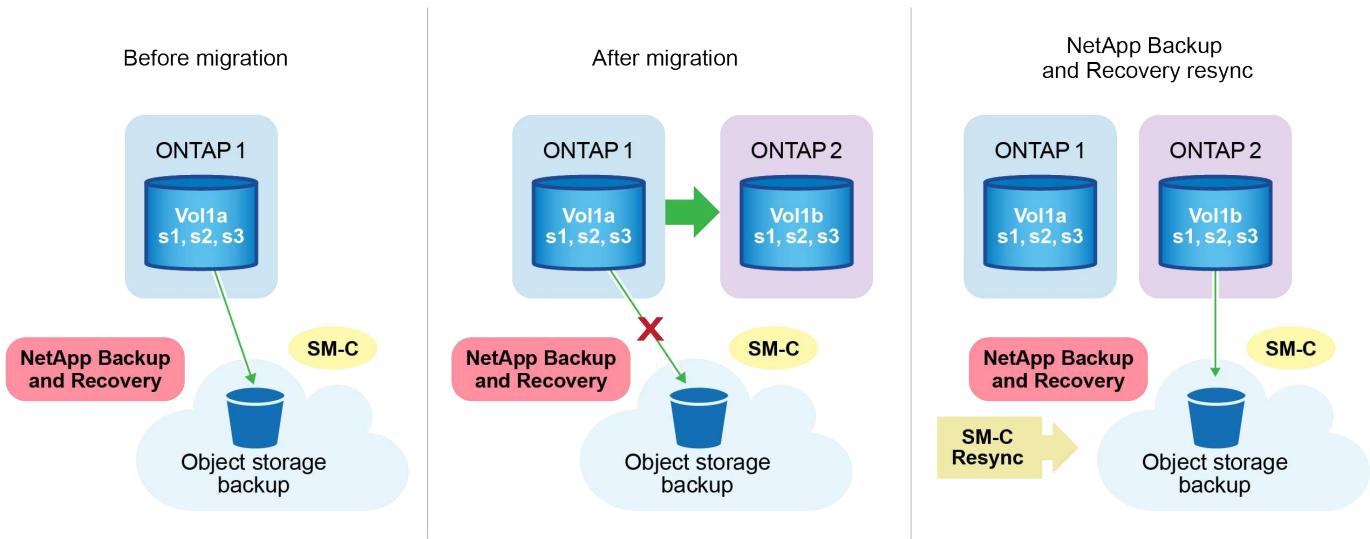
Procedure notes

- Migrations and tech refreshes are not performed using NetApp Backup and Recovery. They should be carried out by a professional services team or a qualified storage administrator.
- A NetApp migration team creates the SnapMirror relationship between the source and destination ONTAP clusters to help move volumes.
- Ensure that the migration during a tech refresh is based on SnapMirror-based migration.

How to migrate volumes using SnapMirror to Cloud Resync

Migrating volumes using SnapMirror to Cloud Resync involves the following major steps, each described in more detail below:

- **Follow a pre-migration checklist:** Before starting the migration, a NetApp Tech Refresh team ensures the following prerequisites are met to avoid data loss and ensure a smooth migration process.
- **Follow a post-migration checklist:** After the migration, a NetApp Tech Refresh team ensures the following steps are completed to establish protection and prepare for the resync.
- **Perform a SnapMirror to Cloud Resync:** After the migration, a NetApp Tech Refresh team performs a SnapMirror to Cloud Resync operation to resume cloud backups from the newly migrated volumes.



Follow a pre-migration checklist

Before migration, the NetApp Tech Refresh team checks these prerequisites to prevent data loss and ensure a smooth process.

1. Ensure all volumes that are to be migrated are protected using NetApp Backup and Recovery.
2. Record volume instance UUIDs. Write down the Instance UUIDs of all volumes before starting the migration. These identifiers are crucial for mapping and resync operations later.
3. Take a final snapshot of each volume to preserve the latest state, before deleting any SnapMirror relationships.
4. Document SnapMirror policies. Record the SnapMirror policy currently attached to each volume's relationship. This will be needed later during the SnapMirror to Cloud Resync process.
5. Delete the SnapMirror Cloud relationships with the object store.
6. Create a standard SnapMirror relationship with the new ONTAP cluster to migrate the volume to the new target ONTAP cluster.

Follow a post-migration checklist

After the migration, a NetApp Tech Refresh team ensures the following steps are completed to establish protection and prepare for the resync.

1. Record new volume instance UUIDs of all migrated volumes in the destination ONTAP cluster.
2. Confirm that all required SnapMirror policies that were available in the old ONTAP cluster are correctly configured in the new ONTAP cluster.
3. Add the new ONTAP cluster as a system in the Console **Systems** page.



The volume instance UUID should be used, not the volume ID. The volume instance UUID is a unique identifier that remains consistent across migrations, while the volume ID may change after migration.

Perform a SnapMirror to Cloud Resync

After the migration, a NetApp Tech Refresh team performs a SnapMirror to Cloud Resync operation to resume

cloud backups from the newly migrated volumes.

1. Add the new ONTAP cluster as a system in the Console **Systems** page.
2. Look at the NetApp Backup and Recovery Volumes page to ensure that the old source system details are available.
3. From the NetApp Backup and Recovery Volumes page, select **Backup Settings**.
 - Within the Backup Settings page, select **View all**.
 - From the Actions ... menu to the right of the *new* source, select **Resync backup**.
4. In the Resync system page, do the following:
 - a. **New source system**: Enter the new ONTAP cluster where the volumes have been migrated.
 - b. **Existing Target Object Store**: Select the target object store that contains the backups from the old source system.
5. Select **Download CSV Template** to download the Resync Details Excel sheet. Use this sheet to enter the details of the volumes to be migrated. In the CSV file, enter the following details:
 - The old volume instance UUID from the source cluster
 - The new volume instance UUID from the destination cluster
 - The SnapMirror policy to be applied to the new relationship.
6. Select **Upload** under the **Upload Volume Mapping Details** to upload the completed CSV sheet into the NetApp Backup and Recovery UI.



The volume instance UUID should be used, not the volume ID. The volume instance UUID is a unique identifier that remains consistent across migrations, while the volume ID may change after migration.

7. Enter provider and network configuration information required for the resync operation.
8. Select **Submit** to start the validation process.

NetApp Backup and Recovery validates that each volume selected for resync is the latest snapshot and has at least one common snapshot. This ensures that the volumes are ready for the SnapMirror to Cloud Resync operation.

9. Review validation results including the new source volume names and the resync status for each volume.
10. Check volume eligibility. The system checks if the volumes are eligible for resync. If a volume is not eligible, it means that it isn't the latest snapshot or no common snapshot was found.



To ensure that volumes remain eligible for the SnapMirror to Cloud Resync operation, take a final snapshot of each volume before deleting any SnapMirror relationships during the pre-migration phase. This preserves the latest state of the data.

11. Select **Resync** to start the resync operation. The system uses the latest and common snapshot to transfer only the incremental changes, ensuring backup continuity.
12. Monitor the resync process in the Job Monitor page.

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.](#)

Set up the StorageGRID environment details

Add the details of the StorageGRID system associated with your ONTAP systems as they were set up on the original Console agent setup using the [NetApp Console APIs](#).

The following information applies to private mode installations starting from NetApp Console 3.9.xx. For older versions, use the following procedure: [DarkSite Cloud Backup: MySQL and Indexed Catalog Backup and Restore](#).

You'll need to perform these steps for each system that is backing up data to StorageGRID.

1. Extract the authorization token using the following oauth/token API.

```
curl 'http://10.193.192.202/oauth/token' -X POST -H 'Accept: application/json' -H 'Accept-Language: en-US,en;q=0.5' -H 'Accept-Encoding: gzip, deflate' -H 'Content-Type: application/json' -d '{
  "username": "admin@netapp.com",
  "password": "Netapp@123",
  "grant_type": "password"
}'
```

While the IP address, username, and passwords are custom values, the account name is not. The account name is always "account-DARKSITE1". Also, the username must use an email-formatted name.

This API will return a response like the following. You can retrieve the authorization token as shown below.

```
{"expires_in":21600,"access_token":"eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCIsImtpZCI6IjJ1MGFjZjRiIn0eyJzdWIiOiJvY2NtYXV0aHwxIiwiYXVkJpbImh0dHBzOi8vYXBpLmNsb3VkLm51dGFwcC5jb20iXSwiaHR0cDovL2Nsb3VkLm51dGFwcC5jb20vZnVsbF9uYW11IjoiYWRtaW4iLCJodHRwOi8vY2xvdWQubmV0YXBwLmNvbS91bWFpbCI6ImFkbWluQG51dGFwcC5jb20iLCJzY29wZSI6Im9wZW5pZCBwcm9maWxlIiwiawF0IjoxNjcyNzM2MDIzLCJleHAiOjE2NzI3NTc2MjMsImlzcyI6Imh0dHA6Ly9vY2NtYXV0aDo4NDIwLyJ9CJtRpRDY23P0kyLg1if67bmgnMcYxdCvBOY-ZUYWzhrWbbY_hqUH4T-114v_pNDsPyNDyWqHaKizThdj_jHYHxm56vTz_Vdn4NqjaBDPwN9KAnC6Z88WA1cJ4WRQqj5ykODNDmrv5At_f9HHp0-xVMyHqywZ4nNFalMvAh4xEsc5jfoKOZc-IOQdWm4F4LHpMzs4qFzCYthTuSKLYtqSTUrZB81-o-ipvrOqSo1iwIeHXZJJV-USWun9daNgiYd_wX-4WWJViGENDzzwOKfUoUoe1Fg3ch--7JFkF1-rrXDOjk1sUMumN3WHV9usp1PgBE5HAcJPrEBm0ValsZcUbiA"}
```

2. Extract the system ID and the X-Agent-Id using the tenancy/external/resource API.

```
curl -X GET  
http://10.193.192.202/tenancy/external/resource?account=account-DARKSITE1 -H 'accept: application/json' -H 'authorization: Bearer eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCIsImtpZCI6IjJ1MGFjZjRiIn0eyJzdWIiOiJvY2NtYXV0aHwxIiwiYXVkJpbImh0dHBzOi8vYXBpLmNsb3VkLm51dGFwcC5jb20iXSwiaHR0cDovL2Nsb3VkLm51dGFwcC5jb20vZnVsbF9uYW11IjoiYWRtaW4iLCJodHRwOi8vY2xvdWQubmV0YXBwLmNvbS91bWFpbCI6ImFkbWluQG51dGFwcC5jb20iLCJzY29wZSI6Im9wZW5pZCBwcm9maWxlIiwiawF0IjoxNjcyNzM2MDIzLCJleHAiOjE2NzI3NDQzMMSImlzcyI6Imh0dHA6Ly9vY2NtYXV0aDo4NDIwLyJ9X_cQF8xttD0-S7sU2uph2cdw_kN-fLWpdJJX98HODwPpVUiLcxV28_sQhuopjWobozPelNISf7KvMqcoXc5kLDyX-ye0fH9gr4XgkdsjWcNvw2rRkFzjHpWrETgfqAMkZcAukV4DHuxogHWh6-DggB1NgPZT8A_szHinud5W0HJ9c4AaT0zc-sp81GaqMahPf0KcFVyjbBL4krOewgKHGfo_7ma_4mF39B1LCj7Vc2XvUd0wCaJvDMjwp19-KbZqmmBX9vDnYp7SSxC1hHJRDStcFgJLdJHtowweNH2829KsjEGBTTcBdO8SvIDtctNH_GAxwSgMT3zUfwaOimPw'
```

This API will return a response like the following. The value under the "resourceIdentifier" denotes the *WorkingEnvironment Id* and the value under "agentId" denotes *x-agent-id*.

```
[{"resourceIdentifier": "OnPremWorkingEnvironment-pMtZND0M", "resourceType": "ON_PREM", "agentId": "vB_1xShPpBtUosjD7wfB1LIhqDgIPA0wclients", "resourceClass": "ON_PREM", "name": "CBSFAS8300-01-02", "metadata": {"clusterUuid": "2cb6cb4b-dc07-11ec-9114-d039ea931e09"}, "workspaceIds": ["workspace2wKYjTy9"], "agentIds": ["vB_1xShPpBtUosjD7wfB1LIhqDgIPA0wclients"]}]
```

3. Update the NetApp Backup and Recovery database with the details of the StorageGRID system associated

with the systems. Make sure to enter the Fully Qualified Domain Name of the StorageGRID, as well as the Access-Key and Storage-Key as shown below:

```
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
eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCIsImtpZCI6IjJ1MGFizjRiIn0eyJzdWIiOiJvY
2NtYXV0aHwxIiwiYXVkJpbImh0dHBzOi8vYXBpLmNs3VkLm51dGFwcC5jb20iXswiaHR0c
DovL2Nsb3VkLm51dGFwcC5jb20vZnVsbF9uYW11IjoiYWRtaW4iLCJodHRwOi8vY2xvdWQub
mV0YXBwLmNvbS91bWFpbCI6ImFkbWluQG51dGFwcC5jb20iLCJzY29wZSI6Im9wZW5pZCBwc
m9maWx1IiwiawF0IjoxNjcyNzEzLCJleHAIoje2NzI3NDQzMTMsImlzcyI6Imh0dHA6I
y9vY2NtYXV0aDo4NDIwLyJ9X_cQF8xttD0-S7sU2uph2cdw_kN-
fLWpdJX98HODwPpVUiLcxV28_sQhuopjWobozPelNISf7KvMqcoXc5kLDyX-
yE0fH9gr4XgkdsjWcNvw2rRkFzjHpWrETgfqAMkZcAukV4DHuxogHWh6-
DggB1NgPZT8A_szHinud5W0HJ9c4AaT0zC-
sp81GaqMahPf0KcFVyjbBL4krOewgKHGfo_7ma_4mF39B1LCj7Vc2XvUd0wCaJvDMjwp19-
KbZqmmBX9vDnYp7SSxC1hHJRDStcFgJLdJHtowweNH2829KsjEGBTTcBdO8SvIDtctNH_GAx
wSgMT3zUfwaOimPw' \
> --header 'x-agent-id: vB_1xShPpBtUosjD7wfB1LIhqDgIPA0wclients' \
> -d '
> { "storage-server" : "sr630ip15.rtp.eng.netapp.com:10443", "access-
key": "2ZMYOAVAS5E70MCNH9", "secret-password":
"uk/6ikd4Lj1XQOFnzSzP/T0zR4ZQ1G0w1xgWsB" }'
```

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.

Manage backups for your ONTAP systems with NetApp Backup and Recovery

With NetApp Backup and Recovery, manage backups for your Cloud Volumes ONTAP and on-premises ONTAP systems by changing the backup schedule, enabling/disabling volume backups, pausing backups, deleting backups, force deleting backups, and more. This includes all types of backups, including snapshots, replicated volumes, and backup files in object storage. You can also unregister NetApp Backup and Recovery.



Do not manage or change backup files directly on your storage systems or from your cloud provider environment. This may corrupt the files and will result in an unsupported configuration.



To switch to and from NetApp Backup and Recovery workloads, refer to [Switch to different NetApp Backup and Recovery workloads](#).

View the backup status of volumes in your systems

You can view a list of all the volumes that are currently being backed up in the Volumes Backup Dashboard. This includes all types of backups, including snapshots, replicated volumes, and backup files in object storage. You can also view the volumes in those systems that are not currently being backed up.

Steps

1. From the Console menu, select **Protection > Backup and recovery**.
2. Select the **Volumes** menu to view the list of backed up volumes for your Cloud Volumes ONTAP and on-premises ONTAP systems.
3. If you are looking for specific volumes in certain systems, you can refine the list by system and volume. You can also use the search filter, or you can sort the columns based on volume style (FlexVol or FlexGroup), volume type, and more.

To show additional columns (aggregates, security style (Windows or UNIX), snapshot policy, replication policy, and backup policy), select the plus sign.

4. Review the status of the protection options in the "Existing protection" column. The 3 icons stand for "Local snapshots", "Replicated volumes", and "Backups in object storage".

Each icon is illuminated when that backup type is activated, and it's grey when the backup type is inactive. You can hover your cursor over each icon to see the backup policy that is being used, and other pertinent information for each type of backup.

Activate backup on additional volumes in a system

If you activated backup only on some of the volumes in a system when you first enabled NetApp Backup and Recovery, you can activate backups on additional volumes later.

Steps

1. From the **Volumes** tab, identify the volume on which you want to activate backups, select the Actions menu **...** at the end of the row, and select **Activate 3-2-1 Protection**.
2. In the *Define backup strategy* page, select the backup architecture, and then define the policies and other details for Local snapshots, Replicated volumes, and Backup files. See the details for backup options from the initial volumes you activated in this system. Then select **Next**.
3. Review the backup settings for this volume, and then select **Activate Backup**.

Change the backup settings assigned to existing volumes

You can change the backup policies assigned to your existing volumes that have assigned policies. You can change the policies for your local snapshots, replicated volumes, and backup files. Any new snapshot, replication, or backup policy that you want to apply to the volumes must already exist.

Edit backup settings on a single volume

Steps

1. From the **Volumes** menu, locate the volume for which you want to modify the policy settings, select the Actions menu **...** at the end of the row, and select **Edit backup strategy**.
2. In the *Edit backup strategy* page, make changes to the existing backup policies for Local snapshots, Replicated volumes, and Backup files and select **Next**.

If you enabled *DataLock and Ransomware Resilience* for cloud backups in the initial backup policy when activating NetApp Backup and Recovery for this cluster, you'll only see other policies that have been configured with DataLock. And if you did not enable *DataLock and Ransomware Resilience* when activating NetApp Backup and Recovery, you'll only see other cloud backup policies that don't have DataLock configured.

3. Review the backup settings for this volume, and then select **Activate Backup**.

Edit backup settings on multiple volumes

If you want to use the same backup settings on multiple volumes, you can activate or edit backup settings on multiple volumes at the same time. You can select volumes that have no backup settings, only snapshot settings, only backup to cloud settings, and so on, and make bulk changes across all these volumes with diverse backup settings.

When working with multiple volumes, all volumes must have these common characteristics:

- same system
- same style (FlexVol or FlexGroup volume)
- same type (Read-write or Data Protection volume)

When more than five volumes are enabled for backup, NetApp Backup and Recovery initializes only five volumes at a time. When those are finished, it continues in groups of 5 until all volumes are initialized.

Steps

1. From the **Volumes** tab, filter by the system on which the volumes reside.
2. Select all the volumes on which you want to manage backup settings.
3. Depending on the type of backup action you want to configure, click the button in the Bulk actions menu:

Backup action...	Select this button...
Manage snapshot backup settings	Manage Local Snapshots
Manage replication backup settings	Manage Replication
Manage backup to cloud backup settings	Manage Backup
Manage multiple types of backup settings. This option enables you to change the backup architecture as well.	Manage Backup and Recovery

4. In the backup page that appears, make changes to the existing backup policies for Local snapshots, Replicated volumes, or Backup files and select **Save**.

If you enabled *DataLock and Ransomware Resilience* for cloud backups in the initial backup policy when activating NetApp Backup and Recovery for this cluster, you'll only see other policies that have been

configured with DataLock. And if you did not enable *DataLock and Ransomware Resilience* when activating NetApp Backup and Recovery, you'll only see other cloud backup policies that don't have DataLock configured.

Create a manual volume backup at any time

You can create an on-demand backup at any time to capture the current state of the volume. This can be useful if very important changes have been made to a volume and you don't want to wait for the next scheduled backup to protect that data. You can also use this functionality to create a backup for a volume that is not currently being backed up and you want to capture its current state.

You can create an ad-hoc snapshot or backup to object store of a volume. You can't create an ad-hoc replicated volume.

The backup name includes the timestamp so you can identify your on-demand backup from other scheduled backups.

If you enabled *DataLock and Ransomware Resilience* when activating NetApp Backup and Recovery for this cluster, the on-demand backup also will be configured with DataLock, and the retention period will be 30 days. Ransomware scans are not supported for ad-hoc backups. [Learn more about DataLock and Ransomware protection.](#)

When you create an ad-hoc backup, a snapshot is created on the source volume. Because this snapshot is not part of a normal snapshot schedule, it will not rotate off. You may want to manually delete this snapshot from the source volume once the backup is complete. This will allow blocks related to this snapshot to be freed up. The name of the Snapshot will begin with `cbs-snapshot-adhoc-`. [See how to delete a Snapshot using the ONTAP CLI.](#)



On-demand volume backup isn't supported on data protection volumes.

Steps

1. From the **Volumes** tab, select **...** for the volume and select **Backup > Create Ad-hoc Backup**.

The Backup Status column for that volume displays "In Progress" until the backup is created.

View the list of backups for each volume

You can view the list of all backup files that exist for each volume. This page displays details about the source volume, destination location, and backup details such as last backup taken, the current backup policy, backup file size, and more.

Steps

1. From the **Volumes** tab, select **...** for the source volume and select **View volume details**.

The details for the volume and the list of snapshots are displayed.

2. Select **Snapshot**, **Replication**, or **Backup** to see the list of all backup files for each type of backup.

Run a ransomware scan on a volume backup in object storage

NetApp Backup and Recovery scans your backup files to look for evidence of a ransomware attack when a backup to object file is created, and when data from a backup file is being restored. You can also run an on-demand scan at any time to verify the usability of a specific backup file in object storage. This can be useful if you have had a ransomware issue on a particular volume and you want to verify that the backups for that

volume are not affected.

This feature is available only if the volume backup was created from a system with ONTAP 9.11.1 or greater, and if you enabled *DataLock and Ransomware Resilience* in the backup-to-object policy.

Steps

1. From the **Volumes** tab, select **...** for the source volume and select **View volume details**.

The details for the volume are displayed.

2. Select **Backup** to see the list of backup files in object storage.
3. Select **...** for the volume backup file you want to scan for ransomware and click **Scan for Ransomware**.

The Ransomware Resilience column shows that the scan is In Progress.

Manage the replication relationship with the source volume

After you set up data replication between two systems, you can manage the data replication relationship.

Steps

1. From the **Volumes** tab, select **...** for the source volume and select the **Replication** option. You can see all of the available options.
2. Select the replication action that you want to perform.

The following table describes the available actions:

Action	Description
View Replication	Shows you details about the volume relationship: transfer information, last transfer information, details about the volume, and information about the protection policy assigned to the relationship.
Update Replication	Starts an incremental transfer to update the destination volume to be synchronized with the source volume.
Pause Replication	Pause the incremental transfer of snapshots to update the destination volume. You can Resume later if you want to restart the incremental updates.
Break Replication	Breaks the relationship between the source and destination volumes, and activates the destination volume for data access - makes it read-write. This option is typically used when the source volume cannot serve data due to events such as data corruption, accidental deletion, or an offline state. Learn how to configure a destination volume for data access and reactivate a source volume in the ONTAP documentation
Abort Replication	Disables backups of this volume to the destination system, and it also disables the ability to restore a volume. Any existing backups will not be deleted. This does not delete the data protection relationship between the source and destination volumes.

Action	Description
Reverse Resync	<p>Reverses the roles of the source and destination volumes. Contents from the original source volume are overwritten by contents of the destination volume. This is helpful when you want to reactivate a source volume that went offline.</p> <p>Any data written to the original source volume between the last data replication and the time that the source volume was disabled is not preserved.</p>
Delete Relationship	<p>Deletes the data protection relationship between the source and destination volumes, which means that data replication no longer occurs between the volumes. This action does not activate the destination volume for data access - meaning it does not make it read-write. This action also deletes the cluster peer relationship and the storage VM (SVM) peer relationship, if there are no other data protection relationships between the systems.</p>

Result

After you select an action, the Console updates the relationship.

Edit an existing backup-to-cloud policy

You can change the attributes for a backup policy that is currently applied to volumes in a system. Changing the backup policy affects all existing volumes that are using the policy.

- If you enabled *DataLock and Ransomware Resilience* in the initial policy when activating NetApp Backup and Recovery for this cluster, any policies that you edit must be configured with the same DataLock setting (Governance or Compliance). And if you did not enable *DataLock and Ransomware Resilience* when activating NetApp Backup and Recovery, you can't enable DataLock now.
- When creating backups on AWS, if you chose *S3 Glacier* or *S3 Glacier Deep Archive* in your first backup policy when activating NetApp Backup and Recovery, then that tier will be the only archive tier available when editing backup policies. And if you selected no archive tier in your first backup policy, then *S3 Glacier* will be your only archive option when editing a policy.

Steps

1. From the **Volumes** tab, select **Backup Settings**.
2. From the *Backup Settings* page, select **...** for the system where you want to change the policy settings, and select **Manage Policies**.
3. From the *Manage Policies* page, select **Edit** for the backup policy you want to change in that system.
4. From the *Edit Policy* page, select the down arrow to expand the *Labels & Retention* section to change the schedule and/or backup retention, and select **Save**.

If your cluster is running ONTAP 9.10.1 or greater, you also have the option to enable or disable tiering of backups to archival storage after a certain number of days.

[Learn more about using AWS archival storage.](#)

[Learn more about using Azure archival storage.](#)

[Learn more about using Google archival storage. \(Requires ONTAP 9.12.1.\)](#)

Note that any backup files that have been tiered to archival storage are left in that tier if you stop tiering

backups to archive - they are not automatically moved back to the standard tier. Only new volume backups will reside in the standard tier.

Add a new backup-to-cloud policy

When you enable NetApp Backup and Recovery for a system, all the volumes you initially select are backed up using the default backup policy that you defined. If you want to assign different backup policies to certain volumes that have different recovery point objectives (RPO), you can create additional policies for that cluster and assign those policies to other volumes.

If you want to apply a new backup policy to certain volumes in a system, you first need to add the backup policy to the system. Then you can [apply the policy to volumes in that system](#).

- If you enabled *DataLock and Ransomware Resilience* in the initial policy when activating NetApp Backup and Recovery for this cluster, any additional policies you create must be configured with the same DataLock setting (Governance or Compliance). And if you did not enable *DataLock and Ransomware Resilience* when activating NetApp Backup and Recovery, you can't create new policies that use DataLock.
- When creating backups on AWS, if you chose *S3 Glacier* or *S3 Glacier Deep Archive* in your first backup policy when activating NetApp Backup and Recovery, then that tier will be the only archive tier available for future backup policies for that cluster. And if you selected no archive tier in your first backup policy, then *S3 Glacier* will be your only archive option for future policies.

Steps

1. From the **Volumes** tab, select **Backup Settings**.
2. From the *Backup Settings* page, select **...** for the system where you want to add the new policy, and select **Manage Policies**.
3. From the *Manage Policies* page, select **Add New Policy**.
4. From the *Add New Policy* page, select down arrow to expand the *Labels & Retention* section to define the schedule and backup retention, and select **Save**.

If your cluster is running ONTAP 9.10.1 or greater, you also have the option to enable or disable tiering of backups to archival storage after a certain number of days.

[Learn more about using AWS archival storage.](#)

[Learn more about using Azure archival storage.](#)

[Learn more about using Google archival storage. \(Requires ONTAP 9.12.1.\)](#)

Delete backups

NetApp Backup and Recovery enables you to delete a single backup file, delete all backups for a volume, or delete all backups of all volumes in a system. You might want to delete all backups if you no longer need the backups, or if you deleted the source volume and want to remove all backups.

You can't delete backup files that you have locked using DataLock and Ransomware protection. The "Delete" option will be unavailable from the UI if you selected one or more locked backup files.

 If you plan to delete a system or cluster that has backups, you must delete the backups **before** deleting the system. NetApp Backup and Recovery doesn't automatically delete backups when you delete a system, and there is no current support in the UI to delete the backups after the system has been deleted. You'll continue to be charged for object storage costs for any remaining backups.

Delete all backup files for a system

Deleting all backups on object storage for a system does not disable future backups of volumes in this system. If you want to stop creating backups of all volumes in a system, you can deactivate backups [as described here](#).

Note that this action does not affect snapshots or replicated volumes - these types of backup files are not deleted.

Steps

1. From the **Volumes** tab, select **Backup Settings**.
2. Select **...** for the system where you want to delete all backups and select **Delete All Backups**.
3. In the confirmation dialog box, enter the name of the system.
4. Select **Advanced settings**.
5. **Force delete backups**: Indicate whether or not you want to force the deletion of all backups.

In some extreme cases, you might want NetApp Backup and Recovery not to have access to backups any longer. This might happen for example, if the service no longer has access to the backup bucket or backups are DataLock protected but you don't want them anymore. Previously, you could not delete these yourself and needed to call NetApp Support. With this release, you can use the option to force delete backups (at volume and system levels).



Use this option carefully and only in extreme cleanup needs. NetApp Backup and Recovery will not have access to these backups any longer even if they are not deleted in the object storage. You will need to go to your cloud provider and manually delete the backups.

6. Select **Delete**.

Delete all backup files for a volume

Deleting all backups for a volume also disables future backups for that volume.

Steps

1. From the **Volumes** tab, click **...** for the source volume and select **Details & Backup List**.
The list of all backup files is displayed.
2. Select **Actions > Delete all Backups**.
3. Enter the volume name.
4. Select **Advanced settings**.
5. **Force delete backups**: Indicate whether or not you want to force the deletion of all backups.

In some extreme cases, you might want NetApp Backup and Recovery not to have access to backups any longer. This might happen for example, if the service no longer has access to the backup bucket or backups are DataLock protected but you don't want them anymore. Previously, you could not delete these

yourself and needed to call NetApp Support. With this release, you can use the option to force delete backups (at volume and system levels).



Use this option carefully and only in extreme cleanup needs. NetApp Backup and Recovery will not have access to these backups any longer even if they are not deleted in the object storage. You will need to go to your cloud provider and manually delete the backups.

6. Select **Delete**.

Delete a single backup file for a volume

You can delete a single backup file if you no longer need it. This includes deleting a single backup of a volume snapshot or of a backup in object storage.

You can't delete replicated volumes (data protection volumes).

Steps

1. From the **Volumes** tab, select **...** for the source volume and select **View volume details**.

The details for the volume are displayed, and you can select **Snapshot**, **Replication**, or **Backup** to see the list of all backup files for the volume. By default, the available snapshots are displayed.

2. Select **Snapshot** or **Backup** to see the type of backup files that you want to delete.
3. Select **...** for the volume backup file you want to delete and select **Delete**.
4. In the confirmation dialog box, select **Delete**.

Delete volume backup relationships

Deleting the backup relationship for a volume provides you with an archiving mechanism if you want to stop the creation of new backup files and delete the source volume, but retain all the existing backup files. This gives you the ability to restore the volume from the backup file in the future, if needed, while clearing space from your source storage system.

You don't necessarily need to delete the source volume. You can delete the backup relationship for a volume and retain the source volume. In this case you can "Activate" backup on the volume at a later time. The original baseline backup copy continues to be used in this case - a new baseline backup copy is not created and exported to the cloud. Note that if you do reactivate a backup relationship, the volume is assigned the default backup policy.

This feature is available only if your system is running ONTAP 9.12.1 or greater.

You can't delete the source volume from the NetApp Backup and Recovery user interface. However, you can open the Volume Details page on the Console **Systems** page, and [delete the volume from there](#).



You can't delete individual volume backup files once the relationship has been deleted. You can, however, you can delete all backups for the volume.

Steps

1. From the **Volumes** tab, select **...** for the source volume and select **Backup > Delete relationship**.

Deactivate NetApp Backup and Recovery for a system

Deactivating NetApp Backup and Recovery for a system disables backups of each volume on the system, and it also disables the ability to restore a volume. Any existing backups will not be deleted. This does not unregister the backup service from this system - it basically allows you to pause all backup and restore activity for a period of time.

Note that you'll continue to be charged by your cloud provider for object storage costs for the capacity that your backups use unless you [delete the backups](#).

Steps

1. From the **Volumes** tab, select **Backup Settings**.
2. From the *Backup Settings page*, select **...** for the system where you want to disable backups and select **Deactivate Backup**.
3. In the confirmation dialog box, select **Deactivate**.



An **Activate Backup** button appears for that system while backup is disabled. You can select this button when you want to re-enable backup functionality for that system.

Unregister NetApp Backup and Recovery for a system

You can unregister NetApp Backup and Recovery for a system if you no longer want to use backup functionality and you want to stop being charged for backups in that system. Typically this feature is used when you're planning to delete a system, and you want to cancel the backup service.

You can also use this feature if you want to change the destination object store where your cluster backups are being stored. After you unregister NetApp Backup and Recovery for the system, then you can enable NetApp Backup and Recovery for that cluster using the new cloud provider information.

Before you can unregister NetApp Backup and Recovery, you must perform the following steps, in this order:

- Deactivate NetApp Backup and Recovery for the system
- Delete all backups for that system

The unregister option is not available until these two actions are complete.

Steps

1. From the **Volumes** tab, select **Backup Settings**.
2. From the *Backup Settings page*, select **...** for the system where you want to unregister the backup service and select **Unregister**.
3. In the confirmation dialog box, select **Unregister**.

Restore from ONTAP backups

Restore ONTAP data from backup files with NetApp Backup and Recovery

Backups of your ONTAP volume data are stored as snapshots, on replicated volumes, or in object storage. You can restore data from any of these locations at a specific point in time. With NetApp Backup and Recovery, you can restore an entire volume, a folder, or individual files as needed.



To switch to and from NetApp Backup and Recovery workloads, refer to [Switch to different NetApp Backup and Recovery workloads](#).

- You can restore a **volume** (as a new volume) to the original system, to a different system that's using the same cloud account, or to an on-premises ONTAP system.
- You can restore a **folder** to a volume in the original system, to a volume in a different system that's using the same cloud account, or to a volume on an on-premises ONTAP system.
- You can restore **files** to a volume in the original system, to a volume in a different system that's using the same cloud account, or to a volume on an on-premises ONTAP system.

You need a valid NetApp Backup and Recovery license to restore data to a production system.

To summarize, these are the valid flows you can use to restore volume data to an ONTAP system:

- Backup file → restored volume
- Replicated volume → restored volume
- Snapshot → restored volume



If the restore operation does not complete, wait until the Job Monitor shows "Failed" before you retry the restore operation.



For limitations related to restoring ONTAP data, see [Backup and restore limitations for ONTAP volumes](#).

The Restore Dashboard

You use the Restore Dashboard to perform volume, folder, and file restore operations. To access the Restore Dashboard, select **Backup and recovery** from the Console menu, and then select the **Restore** tab. You can also select > **View Restore Dashboard** from the Backup and recovery service from the Services panel.



NetApp Backup and Recovery must already be activated for at least one system and initial backup files must exist.

The Restore Dashboard provides two different ways to restore data from backup files: **Browse & Restore** and **Search & Restore**.

Comparing Browse & Restore and Search & Restore

In broad terms, *Browse & Restore* is typically better when you need to restore a specific volume, folder, or file from the last week or month — and you know the name and location of the file, and the date when it was last in good shape. *Search & Restore* is typically better when you need to restore a volume, folder, or file, but you don't remember the exact name, or the volume in which it resides, or the date when it was last in good shape.

This table provides a feature comparison of the two methods.

Browse & Restore	Search & Restore
Browse through a folder-style structure to find the volume, folder, or file within a single backup file.	Search for a volume, folder, or file across all backup files by partial or full volume name, partial or full folder/file name, size range, and additional search filters.
Does not handle file recovery if the file has been deleted or renamed, and the user doesn't know the original file name	Handles newly created/deleted/renamed directories and newly created/deleted/renamed files
Quick restore is supported.	Quick restore is not supported.

This table provides a list of valid restore operations based on the location where your backup files reside.

Backup Type	Browse & Restore			Search & Restore		
	Restore volume	Restore files	Restore folder	Restore volume	Restore files	Restore folder
Snapshot	Yes	No	No	Yes	Yes	Yes
Replicated volume	Yes	No	No	Yes	Yes	Yes
Backup file	Yes	Yes	Yes	Yes	Yes	Yes

Before you use either restore method, configure your environment to meet the resource requirements. See the following sections for details.

See the requirements and restore steps for the type of restore operation you want to use:

- [Restore volumes using Browse & Restore](#)
- [Restore folders and files using Browse & Restore](#)
- [Restore volumes, folders, and files using Search & Restore](#)

Restore from ONTAP backups using Search & Restore

You can use Search & Restore to recover volumes, folders, or files from ONTAP backup files. Search & Restore enables you to search across all backups (including local snapshots, replicated volumes, and object storage) without needing exact system, volume, or file names.

Restoring from local snapshots or replicated volumes is typically faster and less expensive than restoring from object storage.

When restoring a full volume, NetApp Backup and Recovery creates a new volume using the backup data. You can restore to the original system, another system within the same cloud account, or an on-premises ONTAP system. Folders and files can be restored to their original location, a different volume in the same system, another system in the same cloud account, or an on-premises system.

Restore capabilities depend on your ONTAP version:

- **Folders:** Using ONTAP 9.13.0 or greater, you can restore folders with all files and sub-folders; with earlier versions, you can restore only files in the folder.

- **Archival Storage:** Restoring from archival storage (available with ONTAP 9.10.1 or greater) is slower and might incur additional costs.

- **Destination Cluster Requirements:**

- Volume restore: ONTAP 9.10.1 or greater
- File restore: ONTAP 9.11.1 or greater
- Google Archive and StorageGRID: ONTAP 9.12.1 or greater
- Folder restore: ONTAP 9.13.1 or greater

[Learn more about restoring from AWS archival storage.](#)

[Learn more about restoring from Azure archival storage.](#)

[Learn more about restoring from Google archival storage.](#)

- If the backup file in object storage has been configured with DataLock & Ransomware protection, then folder-level restore is supported only if the ONTAP version is 9.13.1 or greater. If you are using an earlier version of ONTAP, you can restore the entire volume from the backup file and then access the folder and files you need.
- If the backup file in object storage resides in archival storage, then folder-level restore is supported only if the ONTAP version is 9.13.1 or greater. If you are using an earlier version of ONTAP, you can restore the folder from a newer backup file that has not been archived, or you can restore the entire volume from the archived backup and then access the folder and files you need.
- The "High" restore priority is not supported when restoring data from Azure archival storage to StorageGRID systems.
- Restoring folders is not currently supported from volumes in ONTAP S3 object storage.



Before you start, you should have some idea of the name or location of the volume or file you want to restore.

Search & Restore supported systems and object storage providers

You can restore ONTAP data from a backup file that resides in a secondary system (a replicated volume) or in object storage (a backup file) to the following systems. Snapshots reside on the source system and can be restored only to that same system.

Note: You can restore volumes and files from any type of backup file, but you can restore a folder only from backup files in object storage at this time.

Backup File Location		Destination system
Object Store (Backup)	Secondary System (Replication)	
Amazon S3	Cloud Volumes ONTAP in AWS On-premises ONTAP system	Cloud Volumes ONTAP in AWS On-premises ONTAP system
Azure Blob	Cloud Volumes ONTAP in Azure On-premises ONTAP system	Cloud Volumes ONTAP in Azure On-premises ONTAP system
Google Cloud Storage	Cloud Volumes ONTAP in Google On-premises ONTAP system	Cloud Volumes ONTAP in Google On-premises ONTAP system

Backup File Location	Destination system
NetApp StorageGRID	On-premisesONTAP system Cloud VolumesONTAP
ONTAP S3	On-premisesONTAP system Cloud VolumesONTAP

For Search & Restore, the Console agent can be installed in the following locations:

- For Amazon S3, the Console agent can be deployed in AWS or in your premises
- For Azure Blob, the Console agent can be deployed in Azure or in your premises
- For Google Cloud Storage, the Console agent must be deployed in your Google Cloud Platform VPC
- For StorageGRID, the Console agent must be deployed in your premises; with or without internet access
- For ONTAP S3, the Console agent can be deployed in your premises (with or without internet access) or in a cloud provider environment

Note that references to "on-premises ONTAP systems" includes FAS, AFF, and ONTAP Select systems.

Search & Restore prerequisites

Ensure your environment meets these requirements before enabling Search & Restore:

- Cluster requirements:
 - The ONTAP version must be 9.8 or greater.
 - The storage VM (SVM) on which the volume resides must have a configured data LIF.
 - NFS must be enabled on the volume (both NFS and SMB/CIFS volumes are supported).
 - The SnapDiff RPC Server must be activated on the SVM. The Console does this automatically when you enable Indexing on the system. (SnapDiff is the technology that quickly identifies the file and directory differences between snapshots.)
- NetApp recommends mounting a separate volume on the Console agent to increase resiliency of Search & Restore. For instructions, refer to [mount the volume to reindex the catalog](#).

Legacy Search & Restore prerequisites (using Indexed Catalog v1)

The following are the requirements for Search & Restore when using legacy indexing:

- AWS requirements:

- Specific Amazon Athena, AWS Glue, and AWS S3 permissions must be added to the user role that provides the Console with permissions. [Make sure all the permissions are configured correctly](#).

Note that if you were already using NetApp Backup and Recovery with a Console agent you configured in the past, you'll need to add the Athena and Glue permissions to the Console user role now. They are required for Search & Restore.

- Azure requirements:

- You must register the Azure Synapse Analytics Resource Provider (called "Microsoft.Synapse") with your Subscription. [See how to register this resource provider for your subscription](#). You must be the Subscription **Owner** or **Contributor** to register the resource provider.
- Specific Azure Synapse Workspace and Data Lake Storage Account permissions must be added to the user role that provides the Console with permissions. [Make sure all the permissions are configured correctly](#).

Note that if you were already using NetApp Backup and Recovery with a Console agent you configured in the past, you'll need to add the Azure Synapse Workspace and Data Lake Storage Account permissions to the Console user role now. They are required for Search & Restore.

- The Console agent must be configured **without** a proxy server for HTTP communication to the internet. If you have configured an HTTP proxy server for your Console agent, you can't use Search & Restore functionality.

- Google Cloud requirements:

- Specific Google BigQuery permissions must be added to the user role that provides the NetApp Console with permissions. [Make sure all the permissions are configured correctly](#).

If you were already using NetApp Backup and Recovery with a Console agent you configured in the past, you'll need to add the BigQuery permissions to the Console user role now. They are required for Search & Restore.

- StorageGRID and ONTAP S3 requirements:

Depending on your configuration, there are 2 ways that Search & Restore is implemented:

- If there are no cloud provider credentials in your account, then the Indexed Catalog information is stored on the Console agent.

For information about the Indexed Catalog v2, see the section below about how to enable the Indexed Catalog.

- If you are using a Console agent in a private (dark) site, then the Indexed Catalog information is stored on the Console agent (requires Console agent version 3.9.25 or greater).
- If you have [AWS credentials](#) or [Azure credentials](#) in the account, then the Indexed Catalog is stored at the cloud provider, just like with a Console agent deployed in the cloud. (If you have both credentials, AWS is selected by default.)

Even though you are using an on-premises Console agent, the cloud provider requirements must be met for both Console agent permissions and cloud provider resources. See the AWS and

Azure requirements above when using this implementation.

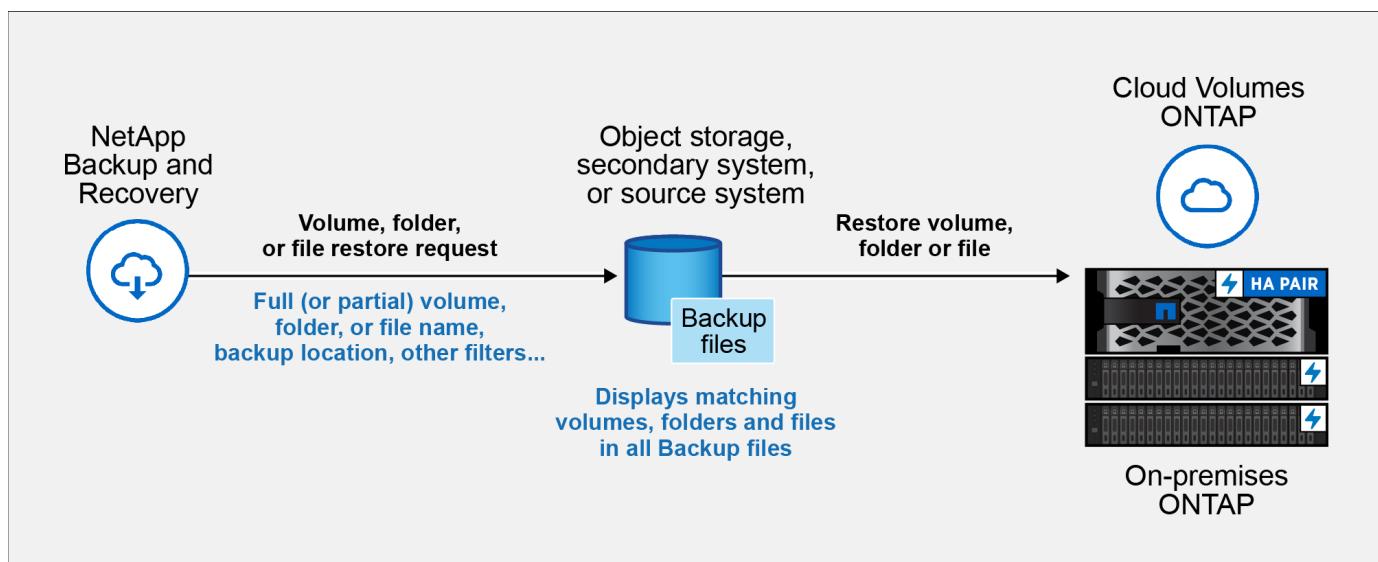
Search & Restore process

The process goes like this:

1. Before you can use Search & Restore, you need to enable "Indexing" on each source system from which you'll want to restore volume data. This allows the Indexed Catalog to track the backup files for every volume.
2. When you want to restore a volume or files from a volume backup, under **Search & Restore**, select **Search & Restore**.
3. Enter the search criteria for a volume, folder, or file by partial or full volume name, partial or full file name, backup location, size range, creation date range, other search filters, and select **Search**.

The Search Results page displays all the locations that have a file or volume that matches your search criteria.

4. Select **View All Backups** for the location you want to use to restore the volume or file, and then select **Restore** on the actual backup file you want to use.
5. Select the location where you want the volume, folder, or file(s) to be restored and select **Restore**.
6. The volume, folder, or file(s) are restored.



You only need to know a partial name and NetApp Backup and Recovery searches through all backup files that match your search.

Enable the Indexed Catalog for each system

Before you can use Search & Restore, you need to enable "Indexing" on each source system from which you're planning to restore volumes or files. This allows the Indexed Catalog to track every volume and every backup file - making your searches very quick and efficient.

The Indexed Catalog is a database that stores metadata about all the volumes and backup files in your system. It is used by the Search & Restore functionality to quickly find the backup files that contain the data you want to restore.

Indexed Catalog features

NetApp Backup and Recovery does not provision a separate bucket when you use the Indexed Catalog. Instead, for backups stored in AWS, Azure, Google Cloud Platform, StorageGRID, or ONTAP S3, the service provisions space on the Console agent or on the cloud provider environment.

The Indexed Catalog supports the following:

- Global search efficiency in less than 3 minutes
- Up to 5 billion files
- Up to 5000 volumes per cluster
- Up to 100K snapshots per volume
- Maximum time for baseline indexing is less than 7 days. The actual time will vary depending on your environment.

Steps to enable Indexing for a system:

If Indexing has already been enabled for your system, go to the next section to restore your data.

You will first need to mount a separate volume to hold catalog files. This prevents data loss if the size of the files that hold the snapshots becomes too large. This is not required on every cluster; you can mount any one volume from any of the clusters in your environment. If you don't do this, indexing might not function correctly.

For the mounted volume, use the following sizing guidance:

- Use a NetApp NFS volume
- Recommended AFF storage with 300 MB/s disk throughput. Less throughput will impact search and other operations.
- Enable NetApp snapshots to secure the catalog metadata in addition to the catalog backup zip files
- 50 GB per 1 billion files
- 20 GB for the catalog data with additional space for zip file creation and temporary files

Step to mount the volume to reindex the catalog

1. Mount the volume to /opt/application/netapp/cbs by entering the following command, where:

- volume name is the volume on the cluster where the catalog files will be stored
- /opt/application/netapp/cbs is the path where it is being mounted

```
mount <cluster IP address>:<volume name> /opt/application/netapp/cbs
```

Example:

```
mount 10.192.24.17:/CATALOG_SCALE_234 /opt/application/netapp/cbs
```

Steps to enable the index

1. Do one of the following:

- If no systems have been indexed, on the Restore Dashboard under *Search & Restore*, select **Enable**

Indexing for systems.

- If at least one system has already been indexed, on the Restore Dashboard under **Search & Restore**, select **Indexing Settings**.

2. Select **Enable Indexing** for the system.

Result

After all the services are provisioned and the Indexed Catalog has been activated, the system is shown as "Active".

Depending on the size of the volumes in the system, and the number of backup files in all 3 backup locations, the initial indexing process could take up to an hour. After that it is transparently updated hourly with incremental changes to stay current.

Restore volumes, folders, and files using Search & Restore

After you have [enabled Indexing for your system](#), you can restore volumes, folders, and files using Search & Restore. This allows you to use a broad range of filters to find the exact file or volume that you want to restore from all backup files.

Steps

1. From the Console menu, select **Protection > Backup and recovery**.
2. Select the **Restore** tab and the Restore Dashboard is displayed.
3. From the **Search & Restore** section, select **Search & Restore**.
4. From the **Search & Restore** section, select **Search & Restore**.
5. From the **Search & Restore** page:
 - a. In the **Search bar**, enter a full or partial volume name, folder name, or file name.
 - b. Select the type of resource: **Volumes**, **Files**, **Folders**, or **All**.
 - c. In the **Filter by** area, select the filter criteria. For example, you can select the system where the data resides and the file type, for example a .JPEG file. Or you can select the type of Backup Location if you want to search for results only within available snapshots or backup files in object storage.
6. Select **Search** and the **Search Results** area displays all the resources that have a file, folder, or volume that matches your search.
7. Locate the resource that has the data you want to restore and select **View All Backups** to display all the backup files that contain the matching volume, folder, or file.
8. Locate the backup file that you want to use to restore the data and select **Restore**.

Note that the results identify local volume snapshots and remote Replicated volumes that contain the file in your search. You can choose to restore from the cloud backup file, from the snapshot, or from the Replicated volume.

9. Select the destination location where you want the volume, folder, or file(s) to be restored and select **Restore**.
 - For volumes, you can select the original destination system or you can select an alternate system. When restoring a FlexGroup volume you'll need to choose multiple aggregates.
 - For folders, you can restore to the original location or you can select an alternate location; including the system, volume, and folder.
 - For files, you can restore to the original location or you can select an alternate location; including the system, volume, and folder. When selecting the original location, you can choose to overwrite the

source file(s) or to create new file(s).

If you select an on-premises ONTAP system and you haven't already configured the cluster connection to the object storage, you are prompted for additional information:

- When restoring from Amazon S3, select the IPspace in the ONTAP cluster where the destination volume will reside, enter the access key and secret key for the user you created to give the ONTAP cluster access to the S3 bucket, and optionally choose a private VPC endpoint for secure data transfer. [See details about these requirements](#).
- When restoring from Azure Blob, select the IPspace in the ONTAP cluster where the destination volume will reside, and optionally choose a private endpoint for secure data transfer by selecting the VNet and Subnet. [See details about these requirements](#).
- When restoring from Google Cloud Storage, select the IPspace in the ONTAP cluster where the destination volume will reside, and the Access Key and Secret Key to access the object storage. [See details about these requirements](#).
- When restoring from StorageGRID, enter the FQDN of the StorageGRID server and the port that ONTAP should use for HTTPS communication with StorageGRID, enter the Access Key and Secret Key needed to access the object storage, and the IPspace in the ONTAP cluster where the destination volume resides. [See details about these requirements](#).
- When restoring from ONTAP S3, enter the FQDN of the ONTAP S3 server and the port that ONTAP should use for HTTPS communication with ONTAP S3, select the Access Key and Secret Key needed to access the object storage, and the IPspace in the ONTAP cluster where the destination volume will reside. [See details about these requirements](#).

Results

The volume, folder, or file(s) are restored and you are returned to the Restore Dashboard so you can review the progress of the restore operation. You can also select the **Job Monitoring** tab to see the restore progress. See [Job monitor page](#).

Restore ONTAP data using Browse & Restore

With NetApp Backup and Recovery, restore ONTAP data using Browse & Restore. Before restoring, note the source volume name, source system and SVM, and backup file date. You can restore ONTAP data from a snapshot, a replicated volume, or from backups stored in object storage.

Restore capabilities depend on your ONTAP version:

- **Folders:** Using ONTAP 9.13.0 or greater, you can restore folders with all files and sub-folders; with earlier versions, you can restore only files in the folder.
- **Archival Storage:** Restoring from archival storage (available with ONTAP 9.10.1 or greater) is slower and might incur additional costs.
- **Destination Cluster Requirements:**
 - Volume restore: ONTAP 9.10.1 or greater
 - File restore: ONTAP 9.11.1 or greater
 - Google Archive and StorageGRID: ONTAP 9.12.1 or greater
 - Folder restore: ONTAP 9.13.1 or greater

[Learn more about restoring from AWS archival storage](#).

[Learn more about restoring from Azure archival storage.](#)
[Learn more about restoring from Google archival storage.](#)



The High priority isn't supported when restoring data from Azure archival storage to StorageGRID systems.

Browse & Restore supported systems and object storage providers

You can restore ONTAP data from a backup file that resides in a secondary system (a replicated volume) or in object storage (a backup file) to the following systems. Snapshots reside on the source system and can be restored only to that same system.

Note: You can restore a volume from any type of backup file, but you can restore a folder or individual files only from a backup file in object storage at this time.

From Object Store (Backup)	From Primary (Snapshot)	From Secondary System (Replication)	To Destination system
Amazon S3	Cloud Volumes ONTAP in AWS On-premises ONTAP system	Cloud Volumes ONTAP in AWS On-premises ONTAP system	Azure Blob
Cloud Volumes ONTAP in Azure On-premises ONTAP system	Cloud Volumes ONTAP in Azure On-premises ONTAP system	Google Cloud Storage	Cloud Volumes ONTAP in Google On-premises ONTAP system
Cloud Volumes ONTAP in Google On-premises ONTAP system	NetApp StorageGRID	On-premises ONTAP system	On-premises ONTAP system Cloud Volumes ONTAP
To on-premises ONTAP system	ONTAP S3	On-premises ONTAP system	On-premises ONTAP system Cloud Volumes ONTAP

For Browse & Restore, the Console agent can be installed in the following locations:

- For Amazon S3, the Console agent can be deployed in AWS or in your premises
- For Azure Blob, the Console agent can be deployed in Azure or in your premises
- For Google Cloud Storage, the Console agent must be deployed in your Google Cloud Platform VPC
- For StorageGRID, the Console agent must be deployed in your premises; with or without internet access
- For ONTAP S3, the Console agent can be deployed in your premises (with or without internet access) or in a cloud provider environment

Note that references to "on-premises ONTAP systems" includes FAS, AFF, and ONTAP Select systems.



If the ONTAP version on your system is less than 9.13.1, then you can't restore folders or files if the backup file has been configured with DataLock & Ransomware. In this case, you can restore the entire volume from the backup file and then access the files you need.

Restore volumes using Browse & Restore

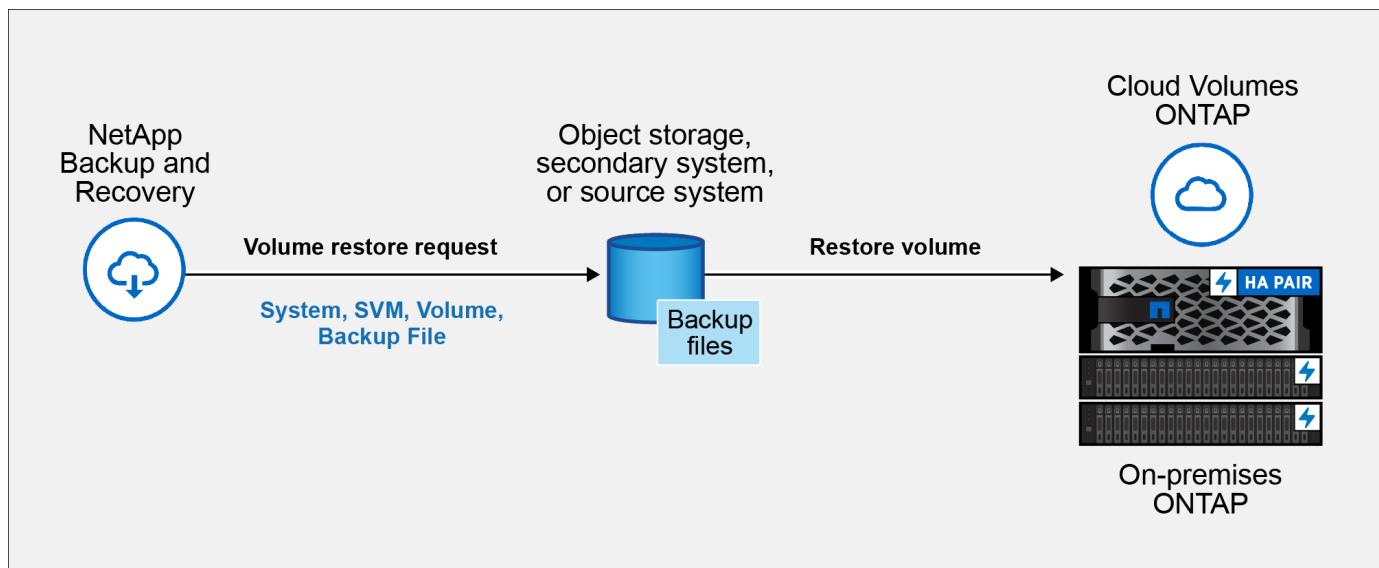
When you restore a volume from a backup file, NetApp Backup and Recovery creates a *new* volume using the data from the backup. When using a backup from object storage, you can restore the data to a volume in the original system, to a different system that's located in the same cloud account as the source system, or to an on-premises ONTAP system.

When restoring a cloud backup to a Cloud Volumes ONTAP system using ONTAP 9.13.0 or greater or to an on-premises ONTAP system running ONTAP 9.14.1, you'll have the option to perform a *quick restore* operation. The quick restore is ideal for disaster recovery situations where you need to provide access to a volume as soon as possible. A quick restore restores the metadata from the backup file to a volume instead of restoring the entire backup file. Quick restore is not recommended for performance or latency-sensitive applications, and it is not supported with backups in archived storage.



Quick restore is supported for FlexGroup volumes only if the source system from which the cloud backup was created was running ONTAP 9.12.1 or greater. And it is supported for SnapLock volumes only if the source system was running ONTAP 9.11.0 or greater.

When restoring from a replicated volume, you can restore the volume to the original system or to a Cloud Volumes ONTAP or on-premises ONTAP system.



You need the source system name, storage VM, volume name, and backup file date to restore a volume.

Steps

1. From the Console menu, select **Protection > Backup and recovery**.
2. Select the **Restore** tab and the Restore Dashboard is displayed.
3. From the *Browse & Restore* section, select **Restore Volume**.
4. In the *Select Source* page, navigate to the backup file for the volume you want to restore. Select the **system**, the **Volume**, and the **Backup** file that has the date/time stamp from which you want to restore.

The **Location** column shows whether the backup file (Snapshot) is **Local** (a snapshot on the source system), **Secondary** (a replicated volume on a secondary ONTAP system), or **Object Storage** (a backup file in object storage). Choose the file that you want to restore.

5. Select **Next**.

Note that if you select a backup file in object storage, and Ransomware Resilience is active for that backup (if you enabled DataLock and Ransomware Resilience in the backup policy), then you are prompted to run an additional ransomware scan on the backup file before restoring the data. We recommend that you scan the backup file for ransomware. (You'll incur extra egress costs from your cloud provider to access the contents of the backup file.)

6. In the *Select Destination* page, select the **system** where you want to restore the volume.
7. When restoring a backup file from object storage, if you select an on-premises ONTAP system and you haven't already configured the cluster connection to the object storage, you are prompted for additional information:
 - When restoring from Amazon S3, select the IPspace in the ONTAP cluster where the destination volume will reside, enter the access key and secret key for the user you created to give the ONTAP cluster access to the S3 bucket, and optionally choose a private VPC endpoint for secure data transfer.
 - When restoring from Azure Blob, select the IPspace in the ONTAP cluster where the destination volume will reside, select the Azure Subscription to access the object storage, and optionally choose a private endpoint for secure data transfer by selecting the VNet and Subnet.
 - When restoring from Google Cloud Storage, select the Google Cloud Project and the Access Key and Secret Key to access the object storage, the region where the backups are stored, and the IPspace in the ONTAP cluster where the destination volume will reside.
 - When restoring from StorageGRID, enter the FQDN of the StorageGRID server and the port that ONTAP should use for HTTPS communication with StorageGRID, select the Access Key and Secret Key needed to access the object storage, and the IPspace in the ONTAP cluster where the destination volume will reside.
 - When restoring from ONTAP S3, enter the FQDN of the ONTAP S3 server and the port that ONTAP should use for HTTPS communication with ONTAP S3, select the Access Key and Secret Key needed to access the object storage, and the IPspace in the ONTAP cluster where the destination volume will reside.
8. Enter the name you want to use for the restored volume, and select the Storage VM and Aggregate where the volume will reside. When restoring a FlexGroup volume you'll need to select multiple aggregates. By default, **<source_volume_name>_restore** is used as the volume name.

When restoring a backup from object storage to a Cloud Volumes ONTAP system using ONTAP 9.13.0 or greater or to an on-premises ONTAP system running ONTAP 9.14.1, you'll have the option to perform a *quick restore* operation.

And if you are restoring the volume from a backup file that resides in an archival storage tier (available starting with ONTAP 9.10.1), then you can select the Restore Priority.

[Learn more about restoring from AWS archival storage.](#)

[Learn more about restoring from Azure archival storage.](#)

[Learn more about restoring from Google archival storage.](#) Backup files in the Google Archive storage tier are restored almost immediately, and require no Restore Priority.

9. Select **Next** to choose whether you want to do a Normal restore or a Quick Restore process:
 - **Normal restore:** Use normal restore on volumes that require high performance. Volumes will not be available until the restore process is complete.
 - **Quick restore:** Restored volumes and data will be available immediately. Do not use this on volumes that require high performance because during the quick restore process, access to the data might be slower than usual.
10. Select **Restore** and you return to the Restore Dashboard so you can review the progress of the restore

operation.

Result

NetApp Backup and Recovery creates a new volume based on the backup you selected.

Note that restoring a volume from a backup file that resides in archival storage can take many minutes or hours depending on the archive tier and the restore priority. You can select the **Job Monitoring** tab to see the restore progress.

Restore folders and files using Browse & Restore

If you need to restore only a few files from an ONTAP volume backup, you can choose to restore a folder or individual files instead of restoring the entire volume. You can restore folders and files to an existing volume in the original system, or to a different system that's using the same cloud account. You can also restore folders and files to a volume on an on-premises ONTAP system.



You can restore a folder or individual files only from a backup file in object storage at this time. Restoring files and folders is not currently supported from a local snapshot or from a backup file that resides in a secondary system (a replicated volume).

If you select multiple files, they are restored to the same destination volume. To restore files to different volumes, run the process multiple times.

When using ONTAP 9.13.0 or greater, you can restore a folder along with all files and sub-folders within it. When using a version of ONTAP before 9.13.0, only files from that folder are restored - no sub-folders, or files in sub-folders, are restored.

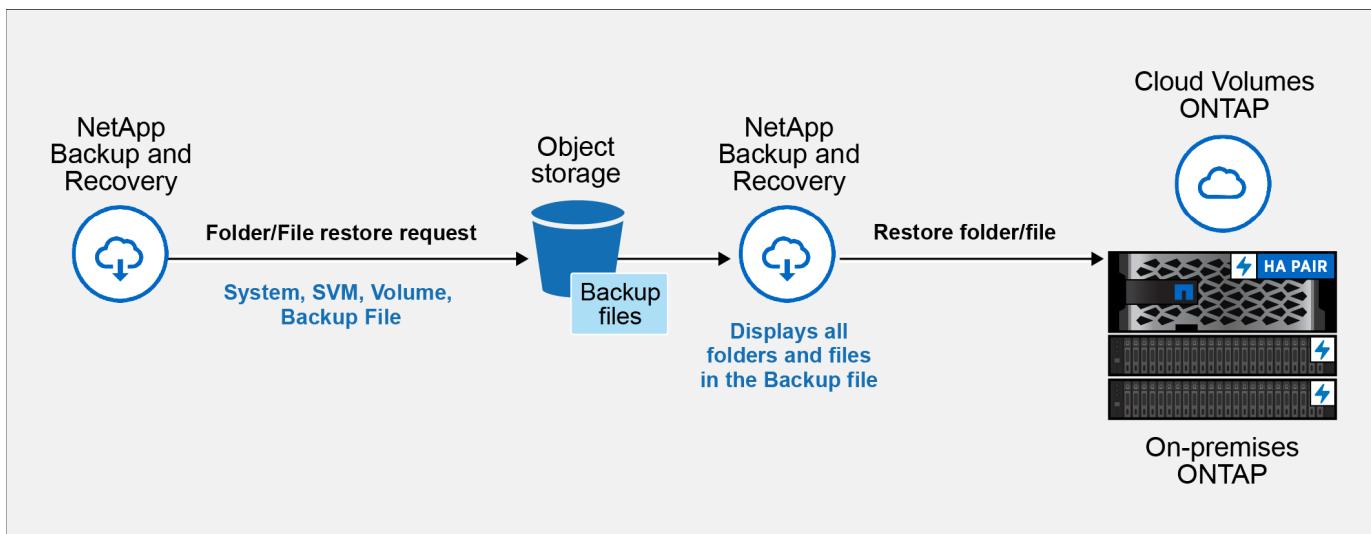


- If the backup file has been configured with DataLock & Ransomware protection, then folder-level restore is supported only if the ONTAP version is 9.13.1 or greater. If you are using an earlier version of ONTAP, you can restore the entire volume from the backup file and then access the folder and files you need.
- If the backup file resides in archival storage, then folder-level restore is supported only if the ONTAP version is 9.13.1 or greater. If you are using an earlier version of ONTAP, you can restore the folder from a newer backup file that has not been archived, or you can restore the entire volume from the archived backup and then access the folder and files you need.
- With ONTAP 9.15.1, you can restore FlexGroup folders using the "Browse and restore" option. This feature is in a Technology Preview mode.

You can test it using a special flag described in the [NetApp Backup and Recovery July 2024 Release blog](#).

Restore folders and files

Follow these steps to restore folders or files to a volume from an ONTAP volume backup. You should know the name of the volume and the date of the backup file that you want to use to restore the folder or file(s). This functionality uses Live Browsing so that you can view the list of directories and files within each backup file.



Before you begin

- The ONTAP version must be 9.6 or greater to perform *file* restore operations.
- The ONTAP version must be 9.11.1 or greater to perform *folder* restore operations. ONTAP version 9.13.1 is required if the data is in archival storage, or if the backup file is using DataLock and Ransomware protection.
- The ONTAP version must be 9.15.1 p2 or greater to restore FlexGroup directories using the Browse and restore option.

Steps

- From the Console menu, select **Protection > Backup and recovery**.
- Select the **Restore** tab and the Restore Dashboard is displayed.
- From the *Browse & Restore* section, select **Restore Files or Folder**.
- In the *Select Source* page, navigate to the backup file for the volume that contains the folder or files you want to restore. Select the **system**, the **Volume**, and the **Backup** that has the date/time stamp from which you want to restore files.
- Select **Next** and the list of folders and files from the volume backup are displayed.

If you are restoring folders or files from a backup file that resides in an archival storage tier, then you can select the Restore Priority.

[Learn more about restoring from AWS archival storage.](#)

[Learn more about restoring from Azure archival storage.](#)

[Learn more about restoring from Google archival storage.](#) Backup files in the Google Archive storage tier are restored almost immediately, and require no Restore Priority.

And if Ransomware Resilience is active for the backup file (if you enabled DataLock and Ransomware Resilience in the backup policy), then you are prompted to run an additional ransomware scan on the backup file before restoring the data. We recommend that you scan the backup file for ransomware. (You'll incur extra egress costs from your cloud provider to access the contents of the backup file.)

- In the *Select Items* page, select the folder or file(s) that you want to restore and select **Continue**. To assist you in finding the item:
 - You can select the folder or file name if you see it.
 - You can select the search icon and enter the name of the folder or file to navigate directly to the item.

- You can navigate down levels in folders using the Down arrow at the end of the row to find specific files.

As you select files they are added to the left side of the page so you can see the files that you have already chosen. You can remove a file from this list if needed by selecting the **x** next to the file name.

7. In the **Select Destination** page, select the **system** where you want to restore the items.

If you select an on-premises cluster and you haven't already configured the cluster connection to the object storage, you are prompted for additional information:

- When restoring from Amazon S3, enter the IPspace in the ONTAP cluster where the destination volume resides, and the AWS Access Key and Secret Key needed to access the object storage. You can also select a Private Link Configuration for the connection to the cluster.
- When restoring from Azure Blob, enter the IPspace in the ONTAP cluster where the destination volume resides. You can also select a Private Endpoint Configuration for the connection to the cluster.
- When restoring from Google Cloud Storage, enter the IPspace in the ONTAP cluster where the destination volumes reside, and the Access Key and Secret Key needed to access the object storage.
- When restoring from StorageGRID, enter the FQDN of the StorageGRID server and the port that ONTAP should use for HTTPS communication with StorageGRID, enter the Access Key and Secret Key needed to access the object storage, and the IPspace in the ONTAP cluster where the destination volume resides.

8. Then select the **Volume** and the **Folder** where you want to restore the folder or file(s).

You have a few options for the location when restoring folders and file(s).

- When you have chosen **Select Target Folder**, as shown above:
 - You can select any folder.
 - You can hover over a folder and click at the end of the row to drill down into subfolders, and then select a folder.
- If you have selected the same destination system and Volume as where the source folder/file was located, you can select **Maintain Source Folder Path** to restore the folder, or file(s), to the same folder where they existed in the source structure. All the same folders and sub-folders must already exist; folders are not created. When restoring files to their original location, you can choose to overwrite the source file(s) or to create new file(s).

9. Select **Restore** to return to the Restore Dashboard and review the progress of the restore operation.

Protect Microsoft SQL Server workloads

Protect Microsoft SQL workloads using NetApp Backup and Recovery overview

Back up your Microsoft SQL Server application data from on-premises ONTAP systems to AWS, Azure, or StorageGRID using NetApp Backup and Recovery. The system automatically creates and stores backups in your cloud account, following your policies. Use a 3-2-1 strategy: keep three copies of your data on two storage systems and one copy in the cloud.

The benefits of the 3-2-1 approach include:

- Multiple data copies protect against internal and external cybersecurity threats.
- Using different types of media helps you recover if one type fails.
- You can quickly restore from the onsite copy, and use the offsite copies if the onsite copy is compromised.

NetApp Backup and Recovery uses NetApp SnapMirror to synchronize backups by creating snapshots and transferring them to the backup locations.

You can do the following to protect your data:

- [Configure additional items if importing from SnapCenter](#)
- [Discover Microsoft SQL Server workloads and optionally import SnapCenter resources](#)
- [Back up workloads with local snapshots on local ONTAP primary storage](#)
- [Replicate workloads to ONTAP secondary storage](#)
- [Back up workloads to an object store location](#)
- [Back up workloads now](#)
- [Restore workloads](#)
- [Clone workloads](#)
- [Manage inventory of workloads](#)
- [Manage snapshots](#)

To back up workloads, you create policies that manage backup and restore operations. See [Create policies](#) for more information.

Supported backup destinations

NetApp Backup and Recovery enables you to back up Microsoft SQL Server instances and databases from the following source systems to the following secondary systems and object storage in public and private cloud providers. Snapshots reside on the source system.

Source system	Secondary system (Replication)	Destination Object Store (Backup)
Cloud Volumes ONTAP in AWS	Cloud Volumes ONTAP in AWS On-premises ONTAP system	Amazon S3 ONTAP S3
Cloud Volumes ONTAP in Azure	Cloud Volumes ONTAP in Azure On-premises ONTAP system	Azure Blob ONTAP S3
On-premises ONTAP system	Cloud Volumes ONTAP On-premises ONTAP system	Amazon S3 Azure Blob NetApp StorageGRID ONTAP S3
Amazon FSx for NetApp ONTAP	Amazon FSx for NetApp ONTAP	NA

Supported restore destinations

You can restore Microsoft SQL Server instances and databases from a backup that resides in primary storage or a secondary system (a replicated volume) or in object storage (a backup file) to the following systems. Snapshots reside on the source system and can be restored only to that same system.

From Backup File Location		To Destination system
Object Store (Backup)	Secondary System (Replication)	
Amazon S3	Cloud Volumes ONTAP in AWS On-premises ONTAP system	Cloud Volumes in AWS On-premises ONTAP system ONTAP S3
Azure Blob	Cloud Volumes ONTAP in Azure On-premises ONTAP system	Cloud Volumes ONTAP in Azure On-premises ONTAP system ONTAP S3
StorageGRID	Cloud Volumes ONTAP On-premises ONTAP system	On-premises ONTAP system ONTAP S3
Amazon FSx for NetApp ONTAP	Amazon FSx for NetApp ONTAP	NA



References to "on-premises ONTAP systems" include FAS and AFF systems.

Prerequisites for importing from the Plug-in service into NetApp Backup and Recovery

If you are going to import resources from the SnapCenter Plug-in service for Microsoft SQL Server into NetApp Backup and Recovery, you'll need to configure a few more items.

Create systems in NetApp Console first

If you are going to import resources from SnapCenter, you should add all on-premises SnapCenter cluster storage to the Console **Systems** page first before importing from SnapCenter. This ensures that host resources can be discovered and imported correctly.

Ensure host requirements to install the SnapCenter Plug-in

To import resources from the SnapCenter Plug-in for Microsoft SQL Server, ensure host requirements to install the SnapCenter Plug-in for Microsoft SQL Server are met.

Check specifically for the SnapCenter requirements in [NetApp Backup and Recovery prerequisites](#).

Disable User Account Control remote restrictions

Before you import resources from SnapCenter, disable User Account Control (UAC) remote restrictions on the SnapCenter Windows host. Disable UAC if you use a local administrative account to connect remotely to the SnapCenter Server host or the SQL host.

Security considerations

Consider the following issues before disabling UAC remote restrictions:

- Security risks: Disabling token filtering can expose your system to security vulnerabilities, especially if local administrative accounts are compromised by malicious actors.
- Use with caution:
 - Modify this setting only if it is essential for your administrative tasks.

- Ensure that strong passwords and other security measures are in place to protect administrative accounts.

Alternative solutions

- If remote administrative access is required, consider using domain accounts with appropriate privileges.
- Use secure remote management tools that adhere to best security practices to minimize risks.

Steps to disable User Account Control remote restrictions

1. Modify the `LocalAccountTokenFilterPolicy` registry key on the SnapCenter Windows host.

Do this by using one of the following, with instructions next:

- Method 1: Registry Editor
- Method 2: PowerShell script

Method 1: Disable User Account Control by using the Registry Editor

This is one of the methods that you can use to disable User Account Control.

Steps

1. Open the Registry Editor on the SnapCenter Windows host by doing the following:
 - a. Press `Windows+R` to open the Run dialog box.
 - b. Type `regedit` and press `Enter`.
2. Navigate to the Policy Key:
`HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System`
3. Create or modify the `DWORD` value:
 - a. Locate: `LocalAccountTokenFilterPolicy`
 - b. If it doesn't exist, create a new `DWORD` (32-bit) Value named `LocalAccountTokenFilterPolicy`.
4. The following values are supported. For this scenario, set the value to 1:
 - 0 (Default): UAC remote restrictions are enabled. Local accounts have filtered tokens when accessing remotely.
 - 1: UAC remote restrictions are disabled. Local accounts bypass token filtering and have full administrative privileges when accessing remotely.
5. Click **OK**.
6. Close the Registry Editor.
7. Restart the SnapCenter Windows host.

Example registry modification

This example sets `LocalAccountTokenFilterPolicy` to "1", disabling UAC remote restrictions.

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System]  
"LocalAccountTokenFilterPolicy"=dword:00000001
```

Method 2: Disable User Account Control by using a PowerShell script

This is another method that you can use to disable User Account Control.



Running PowerShell commands with elevated privileges can affect system settings. Ensure you understand the commands and their implications before running them.

Steps

1. Open a PowerShell window with administrative privileges on the SnapCenter Windows host:
 - a. Click on the **Start** menu.
 - b. Search for **PowerShell 7 or Windows Powershell**.
 - c. Right-click on that option and select **Run as administrator**.
2. Ensure that PowerShell is installed on your system. After installation, it should appear in the **Start** menu.



PowerShell is included by default in Windows 7 and later versions.

3. To disable UAC remote restrictions, set `LocalAccountTokenFilterPolicy` to "1" by running the following command:

```
Set-ItemProperty -Path  
"HKLM:\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System" -Name  
"LocalAccountTokenFilterPolicy" -Value 1 -Type DWord
```

4. Verify that the current value is set to "1" in `LocalAccountTokenFilterPolicy` by running:

```
Get-ItemProperty -Path  
"HKLM:\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System" -Name  
"LocalAccountTokenFilterPolicy"
```

- If the value is 1, UAC remote restrictions are disabled.
- If the value is 0, UAC remote restrictions are enabled.

5. To apply the changes, restart your computer.

Example PowerShell 7 commands to disable UAC remote restrictions:

This example with the value set to "1" indicates that UAC remote restrictions are disabled.

```

# Disable UAC remote restrictions

Set-ItemProperty -Path
"HKLM:\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System" -Name
"LocalAccountTokenFilterPolicy" -Value 1 -Type DWord

# Verify the change

Get-ItemProperty -Path
"HKLM:\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System" -Name
"LocalAccountTokenFilterPolicy"

# Output

LocalAccountTokenFilterPolicy : 1

```

Discover Microsoft SQL Server workloads and optionally import from SnapCenter in NetApp Backup and Recovery

NetApp Backup and Recovery needs to first discover Microsoft SQL Server workloads in order for you to use the service. You can optionally import backup data and policies from SnapCenter if you already have SnapCenter installed.

Required NetApp Console role

Backup and Recovery super admin. Learn about [Backup and recovery roles and privileges](#). Learn about [NetApp Console access roles for all services](#).

Discover Microsoft SQL Server workloads and optionally import SnapCenter resources

During discovery, NetApp Backup and Recovery analyzes Microsoft SQL Server instances and databases in systems within your organization.

NetApp Backup and Recovery assesses Microsoft SQL Server applications. The service assesses the existing protection level including the current backup protection policies, snapshots, and backup and recovery options.

Discovery occurs in the following ways:

- If you already have SnapCenter, import SnapCenter resources into NetApp Backup and Recovery by using the NetApp Backup and Recovery UI.



If you already have SnapCenter, first check to be sure you've met the prerequisites before importing from SnapCenter. For example, you should add on-premises SnapCenter cluster storage systems to the NetApp Console first before importing from SnapCenter. See [Prerequisites for importing resources from SnapCenter](#).

- If you don't already have SnapCenter, you can still discover workloads by adding a vCenter manually and performing discovery.

If SnapCenter is already installed, import SnapCenter resources into NetApp Backup and Recovery

If you already have SnapCenter installed, import SnapCenter resources into NetApp Backup and Recovery using these steps. NetApp Console discovers resources, hosts, credentials, and schedules from SnapCenter; you don't have to recreate all that information.

You can do this in the following ways:

- During discovery, select an option to import resources from SnapCenter.
- After discovery, from the Inventory page, select an option to import SnapCenter resources.
- After discovery, from the Settings menu, select an option to import SnapCenter resources. For details, see [Configure NetApp Backup and Recovery](#).

This is a two-part process:

- Import SnapCenter Server application and host resources
- Manage selected SnapCenter host resources

Import SnapCenter Server application and host resources

This first step imports host resources from SnapCenter and displays those resources in the NetApp Backup and Recovery Inventory page. At that point, the resources are not yet managed by NetApp Backup and Recovery.

 After you import SnapCenter host resources, NetApp Backup and Recovery does not take over protection management automatically. To do so, you must explicitly select to manage the imported resources in NetApp Backup and Recovery. This ensures that you are ready to have those resources backed up by NetApp Backup and Recovery.

Steps

1. From the NetApp Console left navigation, select **Protection > Backup and Recovery**.
2. Select **Inventory**.
3. Select **Discover resources**.
4. From the NetApp Backup and Recovery Discover workload resources page, select **Import from SnapCenter**.
5. Enter **SnapCenter application credentials**:
 - a. **SnapCenter FQDN or IP address**: Enter the FQDN or IP address of the SnapCenter application itself.
 - b. **Port**: Enter the port number for the SnapCenter Server.
 - c. **Username and Password**: Enter the username and password for the SnapCenter Server.
 - d. **Console agent**: Select the Console agent for SnapCenter.
6. Enter **SnapCenter server host credentials**:
 - a. **Existing credentials**: If you select this option, you can use the existing credentials that you have already added. Choose the credentials name.
 - b. **Add new credentials**: If you don't have existing SnapCenter host credentials, you can add new credentials. Enter the credentials name, authentication mode, user name, and password.
7. Select **Import** to validate your entries and register the SnapCenter Server.



If the SnapCenter Server is already registered, you can update the existing registration details.

Result

The Inventory page shows the imported SnapCenter resources that include MS SQL hosts, instances, and databases.

To see the details of the imported SnapCenter resources, select the **View details** option from the Actions menu.

Manage SnapCenter host resources

After you import the SnapCenter resources, manage those host resources in NetApp Backup and Recovery. After you select to manage those resources, NetApp Backup and Recovery is able to back up and recover the resources that you imported from SnapCenter. You no longer manage those resources in SnapCenter Server.

Steps

1. After you import the SnapCenter resources, from the Backup and Recovery menu, select **Inventory**.
2. From the Inventory page, select the imported SnapCenter host that you want to have NetApp Backup and Recovery to manage from now on.
3. Select the Actions icon **...** > **View details** to display the workload details.
4. From the Inventory > workload page, select the Actions icon **...** > **Manage** to display the Manage host page.
5. Select **Manage**.
6. In the Manage host page, select either to use an existing vCenter or add a new vCenter.
7. Select **Manage**.

The Inventory page shows the newly managed SnapCenter resources.

You can optionally create a report of the managed resources by selecting the **Generate reports** option from the Actions menu.

Import SnapCenter resources after discovery from the Inventory page

If you have already discovered resources, you can import SnapCenter resources from the Inventory page.

Steps

1. From the Console left navigation, select **Protection > Backup and Recovery**.
2. Select **Inventory**.
3. From the Inventory page, select **Import SnapCenter resources**.
4. Follow the steps in the **Import SnapCenter resources** section above to import SnapCenter resources.

If you don't have SnapCenter installed, add a vCenter and discover resources

If you don't already have SnapCenter installed, you can add vCenter information and have NetApp backup and recovery discover workloads. Within each Console agent, select the systems where you want to discover workloads.

This is optional if you have a VMware environment.

Steps

1. From the Console left navigation, select **Protection > Backup and Recovery**.

If you are logging in to Backup and Recovery for the first time and have a system in the Console but no discovered resources, the *Welcome to the new NetApp Backup and Recovery* page appears with an option to **Discover resources**.

2. Select **Discover resources**.

3. Enter the following information:

- a. **Workload type**: For this version, only Microsoft SQL Server is available.
- b. **vCenter settings**: Select an existing vCenter or add a new one. To add a new vCenter, enter the vCenter FQDN or IP address, user name, password, port, and protocol.
 -  If you are entering vCenter information, enter information for both vCenter settings and Host registration. If you added or entered vCenter information here, you also need to add plugin information in Advanced Settings next.
- c. **Host registration**: Select **Add credentials** and enter information about the hosts containing the workloads you want to discover.
 -  If you are adding a standalone server and not a vCenter server, enter only the host information.

4. Select **Discover**.



This process might take a few minutes.

5. Continue with Advanced Settings.

Set Advanced settings options during discovery and install the plugin

With Advanced Settings, you can manually install the plugin agent on all servers being registered. This enables you to import all SnapCenter workloads into NetApp Backup and Recovery so you can manage backups and restores there. NetApp Backup and Recovery shows the steps needed to install the plugin.

Steps

1. From the Discover resources page, continue to Advanced Settings by clicking the down arrow on the right.
2. In the Discover workload resources page, enter the following information.
 - **Enter plug-in port number**: Enter the port number that the plugin uses.
 - **Installation path**: Enter the path where the plugin will be installed.
3. If you want to install the SnapCenter agent manually, check the boxes for the following options:
 - **Use manual installation**: Check this box to install the plugin manually.
 - **Add all hosts in the cluster**: Check this box to add all hosts in the cluster to NetApp Backup and Recovery during discovery.
 - **Skip optional preinstall checks**: Check this box to skip optional preinstall checks. You might want to do this for example, if you know that memory or space considerations will be changed in the near future

and you want to install the plugin now.

4. Select **Discover**.

Continue to the NetApp Backup and Recovery Dashboard

1. From the NetApp Console menu, select **Protection > Backup and recovery**.
2. Select a workload tile (for example, Microsoft SQL Server).
3. From the Backup and Recovery menu, select **Dashboard**.
4. Review the health of data protection. The number of at risk or protected workloads increases based on the newly discovered, protected, and backed up workloads.

[Learn what the Dashboard shows you.](#)

Back up Microsoft SQL Server workloads with NetApp Backup and Recovery

Back up Microsoft SQL Server application data from on-premises ONTAP systems to Amazon Web Services, Microsoft Azure, or StorageGRID. The system automatically creates backups and stores them in an object store in your cloud account for data protection.

- To back up workloads on a schedule, create policies that manage backup and restore operations. See [Create policies](#) for instructions.
- Configure the log directory for discovered hosts before starting a backup.
- Back up workloads now (create an on-demand backup now).

View workload protection status

Before you start a backup, view the protection status of your workloads.

Required NetApp Console role

Storage viewer, Backup and Recovery super admin, Backup and Recovery backup admin, Backup and Recovery restore admin, Backup and Recovery clone admin, or Backup and Recovery viewer role. Learn about [Backup and recovery roles and privileges](#). Learn about [NetApp Console access roles for all services](#).

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select a workload to view the protection details.
3. Select the Actions icon  > **View details**.
4. Review details on the Hosts, Protection groups, Availability groups, Instances, and Databases tabs.

Configure the log directory for discovered hosts

Set the activity log path for discovered hosts to track operation status before backing up workloads.

Required NetApp Console role

Storage viewer, Backup and Recovery super admin, Backup and Recovery backup admin, or Backup and Recovery restore admin role. Learn about [NetApp Console access roles for all services](#).

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select a workload to view the protection details.
3. Select the Actions icon  > **View details**.
4. Select a host.
5. Select the Actions icon  > **Configure log directory**.
6. Enter the host path or browse through a list of hosts or nodes to find where you want to store the host log.
7. Select those on which you want to store the logs.



The fields that appear differ depending on the selected deployment model, for example, failover cluster instance or standalone.

8. Select **Save**.

Create a protection group

Create a protection group to manage backup and restore operations for multiple workloads. A protection group is a logical grouping of workloads.

Required NetApp Console role

Storage viewer, Backup and Recovery super admin, or Backup and Recovery backup admin role. [Learn about NetApp Console access roles for all services](#).

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select a workload to view the protection details.
3. Select the Actions icon  > **View details**.
4. Select the **Protection groups** tab.
5. Select **Create protection group**.
6. Provide a name for the protection group.
7. Select the instances or databases that you want to include in the protection group.
8. Select **Next**.
9. Select the **Backup policy** that you want to apply to the protection group.

If you want to create a policy, select **Create new policy** and follow the prompts to create a policy. See [Create policies](#) for more information.

10. Select **Next**.
11. Review the configuration.
12. Select **Create** to create the protection group.

Back up workloads now with an on-demand backup

Run an on-demand backup before making changes to your system to ensure your data is protected.

Required NetApp Console role

Storage viewer, Backup and Recovery super admin, or Backup and Recovery backup admin role. [Learn about NetApp Console access roles for all services](#).

[NetApp Console access roles for all services.](#)

Steps

1. From the menu, select **Inventory**.
2. Select a workload to view the protection details.
3. Select the Actions icon  > **View details**.
4. Select the **Protection Group, Instances or Databases** tab.
5. Select the instance or database you want to back up.
6. Select the Actions icon  > **Back up now**.
7. Select the policy that you want to apply to the backup.
8. Select the schedule tier.
9. Select **Back up now**.

Suspend the backup schedule

Suspend the schedule to temporarily stop backups during maintenance or troubleshooting.

Required NetApp Console role

Storage viewer, Backup and Recovery super admin, or Backup and Recovery backup admin role. [Learn about NetApp Console access roles for all services](#).

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select a workload to view the protection details.
3. Select the Actions icon  > **View details**.
4. Select the **Protection Group, Instances or Databases** tab.
5. Select the protection group, instance, or database you want to suspend.
6. Select the Actions icon  > **Suspend**.

Delete a protection group

Deleting a protection group removes it and all associated backup schedules. You might want to delete a protection group if it is no longer needed.

Required NetApp Console role

Storage viewer, Backup and Recovery super admin, or Backup and Recovery backup admin role. [Learn about NetApp Console access roles for all services](#).

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select a workload to view the protection details.
3. Select the Actions icon  > **View details**.
4. Select the **Protection groups** tab.
5. Select the Actions icon  > **Delete protection group**.

Remove protection from a workload

You can remove protection from a workload if you no longer want to back it up or if you want to stop managing it in NetApp Backup and Recovery.

Required NetApp Console role

Storage viewer, Backup and Recovery super admin, or Backup and Recovery backup admin role. [Learn about NetApp Console access roles for all services](#).

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select a workload to view the protection details.
3. Select the Actions icon  > **View details**.
4. Select the **Protection Group, Instances or Databases** tab.
5. Select the protection group, instance, or database.
6. Select the Actions icon  > **Remove protection**.
7. In the Remove protection dialog box, select whether you want to keep backups and metadata or delete them.
8. Select **Remove** to confirm the action.

Restore Microsoft SQL Server workloads with NetApp Backup and Recovery

Restore Microsoft SQL Server workloads using NetApp Backup and Recovery. Use snapshots, backups replicated to secondary storage, or backups in object storage. Restore workloads to the original system, a different system with the same cloud account, or an on-premises ONTAP system.

Restore from these locations

You can restore workloads from different starting locations:

- Restore from a primary location
- Restore from a replicated resource
- Restore from an object store backup

Restore to these points

You can restore data to the latest snapshot or to these points:

- Restore from snapshots
- Restore to a specific point in time if you know the file name, location, and last valid date
- Restore to the latest backup

Restore from object storage considerations

If you select a backup file in object storage, and Ransomware Resilience is active for that backup (if you enabled DataLock and Ransomware Resilience in the backup policy), then you are prompted to run an additional integrity check on the backup file before restoring the data. We recommend that you perform the scan.

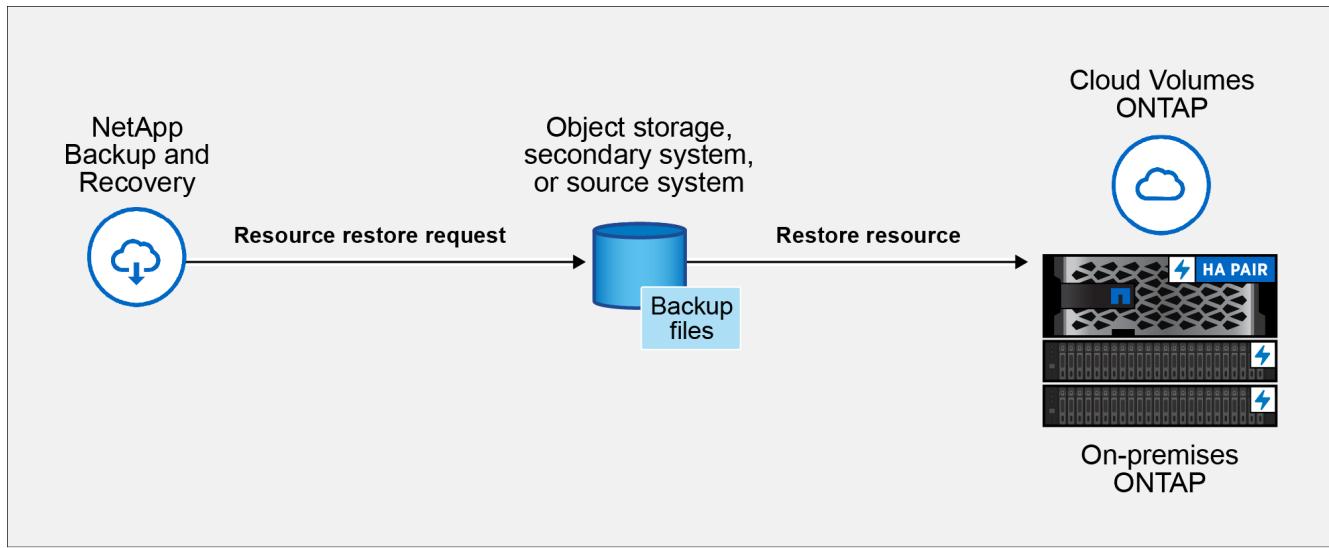


You pay extra fees to your cloud provider to access the backup file.

How restoring workloads works

When you restore workloads, the following occurs:

- When you restore a workload from a backup file, NetApp Backup and Recovery creates a *new* resource using the data from the backup.
- When you restore from a replicated workload, you can restore the workload to the original system or to an on-premises ONTAP system.



- When you restore a backup from object storage, you can restore the data to the original system or to an on-premises ONTAP system.

Restore methods

Restore workloads using one of these methods:

- From the Restore page:** Use this option to restore a resource when you do not know its name, location, or last good date. Search for the snapshot using filters.
- From the Inventory page:** Use this option to restore a specific resource when you know its name, location, and last good date. Browse the list to find the resource.

Required NetApp Console role

Storage viewer, Backup and Recovery super admin, or Backup and Recovery backup admin role. [Learn about NetApp Console access roles for all services](#).

Restore workload data from the Restore option

Restore database workloads using the Restore option.

Steps

- From the NetApp Backup and Recovery menu, select **Restore**.
- Select the database that you want to restore. Use the filters to search.
- Select the restore option:

- Restore from snapshots
- Restore to a specific point in time if you know the file name, location, and last valid date
- Restore to the latest backup

Restore workloads from snapshots

1. Continuing from the Restore options page, select **Restore from snapshots**.

A list of snapshots appears.

2. Select the snapshot you want to restore.
3. Select **Next**.

You'll see destination options next.

4. In the Destination details page, enter the following information:

- **Destination settings:** Choose whether you want to restore the data to its original location or to an alternate location. For an alternate location, select the host name and instance, enter the database name, and enter the destination path where you want to restore the snapshot.
- **Pre-restore options:**
 - **Overwrite the database with the same name during restore:** During the restore, the original database name is preserved.
 - **Retain SQL database replication settings:** Keeps the replication settings for the SQL database after the restore operation.
 - **Create transaction log backup before restore:** Creates a transaction log backup before the restore operation.* **Quit restore if transaction log backup before restore fails:** Stops the restore operation if the transaction log backup fails.
 - **Prescript:** Enter the full path for a script that should be run before the restore operation, any arguments that the script takes, and how long to wait for the script to complete.
- **Post-restore options:**
 - **Operational**, but unavailable for restoring additional transaction logs. This brings the database back online after transaction log backups are applied.
 - **Non-operational**, but available for restoring additional transaction logs. Maintains the database in a non-operational state after the restore operation while restoring transaction log backups. This option is useful for restoring additional transaction logs.
 - **Read-only mode** and available for restoring additional transaction logs. Restores the database in a read-only mode and applies transaction log backups.
 - **Postscript:** Enter the full path for a script that should be run after the restore operation and any arguments that the script takes.

5. Select **Restore**.

Restore to specific point in time

NetApp Backup and Recovery uses logs and the most recent snapshots to create a point-in-time restore of your data.

1. Continuing from the Restore options page, select **Restore to specific point in time**.

2. Select **Next**.
3. In the Restore to a specific point in time page, enter the following information:
 - **Date and time for data restoration:** Enter the exact date and time of the data that you want to restore. This date and time is from the Microsoft SQL Server Database host.
4. Select **Search**.
5. Select the snapshot that you want to restore.
6. Select **Next**.
7. In the Destination details page, enter the following information:
 - **Destination settings:** Choose whether you want to restore the data to its original location or to an alternate location. For an alternate location, select the host name and instance, enter the database name, and enter the destination path.
 - **Pre-restore options:**
 - **Preserve original database name:** During the restore, the original database name is preserved.
 - **Retain SQL database replication settings:** Keeps the replication settings for the SQL database after the restore operation.
 - **Prescript:** Enter the full path for a script that should be run before the restore operation, any arguments that the script takes, and how long to wait for the script to complete.
 - **Post-restore options:**
 - **Operational**, but unavailable for restoring additional transaction logs. This brings the database back online after transaction log backups are applied.
 - **Non-operational**, but available for restoring additional transaction logs. Maintains the database in a non-operational state after the restore operation while restoring transaction log backups. This option is useful for restoring additional transaction logs.
 - **Read-only mode** and available for restoring additional transaction logs. Restores the database in a read-only mode and applies transaction log backups.
 - **Postscript:** Enter the full path for a script that should be run after the restore operation and any arguments that the script takes.
8. Select **Restore**.

Restore to the latest backup

This option uses the latest full and log backups to restore your data to the last good state. The system scans logs from the last snapshot to the present. The process tracks changes and activities to restore the most recent and accurate version of your data.

1. Continuing from the Restore options page, select **Restore to the latest backup**.

NetApp Backup and Recovery shows you the snapshots that are available for the restore operation.
2. In the Restore to the latest state page, select the snapshot location of local, secondary storage, or object storage.
3. Select **Next**.
4. In the Destination details page, enter the following information:
 - **Destination settings:** Choose whether you want to restore the data to its original location or to an alternate location. For an alternate location, select the host name and instance, enter the database name, and enter the destination path.

name, and enter the destination path.

- **Pre-restore options:**

- **Overwrite the database with the same name during restore:** During the restore, the original database name is preserved.
- **Retain SQL database replication settings:** Keeps the replication settings for the SQL database after the restore operation.
- **Create transaction log backup before restore:** Creates a transaction log backup before the restore operation.
- **Quit restore if transaction log backup before restore fails:** Stops the restore operation if the transaction log backup fails.
- **Prescript:** Enter the full path for a script that should be run before the restore operation, any arguments that the script takes, and how long to wait for the script to complete.

- **Post-restore options:**

- **Operational**, but unavailable for restoring additional transaction logs. This brings the database back online after transaction log backups are applied.
- **Non-operational**, but available for restoring additional transaction logs. Maintains the database in a non-operational state after the restore operation while restoring transaction log backups. This option is useful for restoring additional transaction logs.
- **Read-only mode** and available for restoring additional transaction logs. Restores the database in a read-only mode and applies transaction log backups.
- **Postscript:** Enter the full path for a script that should be run after the restore operation and any arguments that the script takes.

5. Select **Restore**.

Restore workload data from the Inventory option

Restore database workloads from the Inventory page.

Using the Inventory option, you can restore only databases, not instances.

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Choose the host where the resource that you want to restore is located.
3. Select the **Actions**  icon, and select **View details**.
4. On the Microsoft SQL Server page, select the **Databases** tab.
5. In the Databases menu, select a database with "Protected" status.
6. Select the **Actions**  icon, and select **Restore**.

The same three options appear as when you restore from the Restore page:

- Restore from snapshots
- Restore to a specific point in time
- Restore to the latest backup

7. Continue with the same steps for the restore option from the Restore page

Clone Microsoft SQL Server workloads using NetApp Backup and Recovery

Clone Microsoft SQL Server application data to a VM for development, testing, or protection with NetApp Backup and Recovery. Create clones from instant or existing snapshots of your SQL Server workloads.

Choose between the following types of clones:

- **Instant snapshot and clone:** You can create a clone of your Microsoft SQL Server workloads from an instant snapshot, which is a point-in-time copy of the source data that is created from a backup. The clone is stored in an object store in your public or private cloud account. You can use the clone to restore your workloads in case of data loss or corruption.
- **Clone from an existing snapshot:** You can choose an existing snapshot from a list of snapshots that are available for the workload. This option is useful if you want to create a clone from a specific point in time. Clone to either primary or secondary storage.

You can accomplish the following protection goals:

- Create a clone
- Refresh a clone
- Split a clone
- Delete a clone

Required NetApp Console role

Storage viewer, Backup and Recovery super admin, or Backup and Recovery backup admin role. [Learn about NetApp Console access roles for all services](#).

Create a clone

You can create a clone of your Microsoft SQL Server workloads. A clone is a copy of the source data that is created from a backup. The clone is stored in an object store in your public or private cloud account. You can use the clone to restore your workloads in case of data loss or corruption.

You can create a clone from an existing snapshot or from an instant snapshot. An instant snapshot is a point-in-time copy of the source data that is created from a backup. You can use the clone to restore your workloads in case of data loss or corruption.

Steps

1. From the NetApp Backup and Recovery menu, select **Clone**.
2. Select **Create new clone**.
3. Select the clone type:
 - **Clone and database refresh from existing snapshot:** Choose a snapshot and configure clone options.
 - **Instant snapshot and clone:** Take a snapshot now of the source data and create a clone from that snapshot. This option is useful if you want to create a clone from the latest data in the source workload.
4. Complete the **Database source** section:
 - **Single clone or bulk clone:** Select whether to create a single clone or multiple clones. If you select **Bulk clone**, you can create multiple clones at once using a protection group that you already created. This option is useful if you want to create multiple clones for different workloads.

- **Source database host, instance, and name:** Select the source database host, instance, and name for the clone. The source database is the database from which the clone will be created.

5. Complete the **Database target** section:

- **Target database host, instance, and name:** Select the target database host, instance, and name for the clone. The target database is the location where the clone will be created.

Optionally, select **Suffix** from the target name drop-down list and add a suffix to the cloned database name. If you do not add a suffix, the cloned database name is the same as the source database name.

- **QoS (max throughput):** Select the quality of service (QoS) maximum throughput in MBps for the clone. The QoS defines the performance characteristics of the clone, such as the maximum throughput and IOPS.

6. Complete the **Mount** section:

- **Auto-assign mount point:** Automatically assign a mount point for the clone in the object store.
- **Define mount point path:** Enter a mount point for the clone. The mount point is the location where the clone will be mounted in the object store. Select the drive letter, enter the data file path, and enter the log file path.

7. Select **Next**.

8. Select the restore point:

- **Existing snapshots:** Select an existing snapshot from the list of snapshots that are available for the workload. This option is useful if you want to create a clone from a specific point in time.
- **Instant snapshot and clone:** Select the latest snapshot from the list of snapshots that are available for the workload. This option is useful if you want to create a clone from the latest data in the source workload.

9. If you chose to create **Instant snapshot and clone**, choose the clone storage location:

- **Local storage:** Select this option to create the clone in the local storage of the ONTAP system. The local storage is the storage that is directly attached to the ONTAP system.
- **Secondary storage:** Select this option to create the clone in the secondary storage of the ONTAP system. The secondary storage is the storage that is used for backup and recovery workloads.

10. Select the destination location for the data and logs.

11. Select **Next**.

12. Complete the **Advanced options** section.

13. If you chose **Instant snapshot and clone**, complete the following options:

- **Clone refresh schedule and expiration:** If you chose **Instant clone**, enter the date when to begin refreshing the clone. The clone schedule defines when the clone will be created.
 - **Delete clone if schedule expires:** If you want to delete the clone upon the clone expiration date.
 - **Refresh clone every:** Select how often the clone should be refreshed. You can choose to refresh the clone hourly, daily, weekly, monthly, or quarterly. This option is useful if you want to keep the clone up to date with the source workload.
- **Prescripts and postscripts:** Optionally, add scripts to run before and after the clone is created. These scripts can do extra tasks, such as setting up the clone or sending notifications.
- **Notification:** Optionally, specify email addresses to receive notifications about the clone creation status along with the Job report. You can also specify a webhook URL to receive notifications about the clone creation status. You can specify whether you want success and failure notifications or only one or the other.

- **Tags:** Select labels to help you search for resource groups later and select **Apply**. For example, if you add "HR" as a tag to multiple resource groups, you can later find all resource groups associated with the "HR" tag.

14. Select **Create**.

15. When the clone is created, you can view it in the **Inventory** page.

Refresh a clone

You can refresh a clone of your Microsoft SQL Server workloads. Refreshing a clone updates the clone with the latest data from the source workload. This is useful if you want to keep the clone up to date with the source workload.

You have the option to change the database name, use the latest instant snapshot, or refresh from an existing production snapshot.

Steps

1. From the NetApp Backup and Recovery menu, select **Clone**.
2. Select the clone you want to refresh.
3. Select the Actions icon  > **Refresh clone**.
4. Complete the **Advanced settings** section:
 - **Recovery scope:** Choose whether to recover all log backups or log backups until a specific point in time. This option is useful if you want to recover the clone to a specific point in time.
 - **Clone refresh schedule and expiration:** If you chose **Instant clone**, enter the date when to begin refreshing the clone. The clone schedule defines when the clone will be created.
 - **Delete clone if schedule expires:** If you want to delete the clone upon the clone expiration date.
 - **Refresh clone every:** Select how often the clone should be refreshed. You can choose to refresh the clone hourly, daily, weekly, monthly, or quarterly. This option is useful if you want to keep the clone up to date with the source workload.
 - **iGroup settings:** Select the iGroup for the clone. The iGroup is a logical grouping of initiators that are used to access the clone. You can select an existing iGroup or create a new one. Select the iGroup from the primary or secondary ONTAP storage system.
 - **Prescripts and postscripts:** Optionally, add scripts to run before and after the clone is created. These scripts can do extra tasks, such as setting up the clone or sending notifications.
 - **Notification:** Optionally, specify email addresses to receive notifications about the clone creation status along with the Job report. You can also specify a webhook URL to receive notifications about the clone creation status. You can specify whether you want success and failure notifications or only one or the other.
 - **Tags:** Enter one or more labels that will help you later search for the resource group. For example, if you add "HR" as a tag to multiple resource groups, you can later find all resource groups associated with the HR tag.
5. In the Refresh confirmation dialog box, to continue, select **Refresh**.

Skip a clone refresh

Skip a clone refresh to keep the clone unchanged.

Steps

1. From the NetApp Backup and Recovery menu, select **Clone**.
2. Select the clone you want to skip the refresh for.
3. Select the Actions icon **...** > **Skip refresh**.
4. In the Skip refresh confirmation dialog box, do the following:
 - a. To skip only the next refresh schedule, select **Only skip the next refresh schedule**.
 - b. To continue, select **Skip**.

Split a clone

You can split a clone of your Microsoft SQL Server workloads. Splitting a clone creates a new backup from the clone. The new backup can be used to restore the workloads.

You can choose to split a clone as independent or long-term clones. A wizard shows the list of aggregates that are part of the SVM, their sizes, and where the cloned volume resides. NetApp Backup and Recovery also indicates whether there is enough space to split the clone. After the clone is split, the clone becomes an independent database for protection.

The clone job is not be removed and it can be reused again for other clones.

Steps

1. From the NetApp Backup and Recovery menu, select **Clone**.
2. Select a clone.
3. Select the Actions icon **...** > **Split clone**.
4. Review the split clone details and select **Split**.
5. When the split clone is created, you can view it in the **Inventory** page.

Delete a clone

You can delete a clone of your Microsoft SQL Server workloads. Deleting a clone removes the clone from the object store and frees up storage space.

If a policy protects the clone, both the clone and its job are deleted.

Steps

1. From the NetApp Backup and Recovery menu, select **Clone**.
2. Select a clone.
3. Select the Actions icon **...** > **Delete clone**.
4. In the clone Delete confirmation dialog box, review the deletion details.
 - a. To delete the cloned resources from SnapCenter even if the clones or their storage is not accessible, select **Force delete**.
 - b. Select **Delete**.
5. When the clone is deleted, it is removed from the **Inventory** page.

Manage Microsoft SQL Server inventory with NetApp Backup and Recovery

NetApp Backup and Recovery helps you manage your Microsoft SQL Server hosts,

databases, and instances. You can view, change, or remove protection settings for your inventory.

You can accomplish the following tasks related to managing your inventory:

- Manage host information
 - Suspend schedules
 - Edit or delete hosts
- Manage instances information
 - Associate credentials with a resource
 - Back up now by starting an on-demand backup
 - Edit protection settings
- Manage database information
 - Protect databases
 - Restore databases
 - Edit protection settings
 - Back up now by starting an on-demand backup
- Configure the log directory (from **Inventory > Hosts**). If you want to back up logs for your database hosts in the snapshot, first configure the logs in NetApp Backup and Recovery. For details, refer to [Configure NetApp Backup and Recovery settings](#).

Manage host information

You can manage host information to ensure that the right hosts are protected. You can view, edit, and delete host information.

Required NetApp Console role

Storage viewer, Backup and Recovery super admin, Backup and Recovery backup admin, Backup and Recovery restore admin, or Backup and Recovery clone admin role. [Learn about NetApp Console access roles for all services](#).

- Configure log directory. For details, refer to [Configure NetApp Backup and Recovery settings](#).
- Suspend schedules
- Edit a host
- Delete a host

Manage hosts

You can manage the hosts that are discovered in your system. You can manage them separately or as a group.



You can manage hosts with an "Unmanaged" status in the Hosts column. NetApp Backup and Recovery already manages hosts with a "Managed" status.

After you manage the hosts in NetApp Backup and Recovery, SnapCenter no longer manages the resources on those hosts.

Required NetApp Console role

Storage viewer, or Backup and Recovery super admin. [Learn about NetApp Console access roles for all services](#).

Steps

1. From the menu, select **Inventory**.
2. Select a workload to view the protection details.
3. Select the Actions icon  > **View details**.
4. Select the **Hosts** tab.
5. Select one or more hosts. If you select multiple hosts, a Bulk actions option appears where you can select **Manage (up to 5 hosts)**.
6. Select the Actions icon  > **Manage**.
7. Review the host dependencies:
 - If the vCenter does not display, select the pencil icon to add or edit the vCenter details.
 - If you add a vCenter, you must also register the vCenter by selecting **Register vCenter**.
8. Select **Validate settings** to test your settings.
9. Select **Manage** to manage the host.

Suspend schedules

Suspend schedules to stop backup and restore operations during host maintenance.

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select the host on which you want to suspend schedules.
3. Select the Actions  icon, and select **Suspend schedules**.
4. In the confirmation dialog box, select **Suspend**.

Edit a host

You can change the vCenter server information, host registration credentials, and advanced settings options.

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select the host that you want to edit.
3. Select the Actions  icon, and select **Edit host**.
4. Edit the host information.
5. Select **Done**.

Delete a host

You can delete the host information to stop service charges.

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.

2. Select the host that you want to delete.
3. Select the **Actions**  icon, and select **Delete host**.
4. Review the confirmation information and select **Delete**.

Manage instances information

You can manage instances information to assign the appropriate credentials for resource protection and back up resources in the following ways:

- Protect instances
- Associate credentials
- Disassociate credentials
- Edit protection
- Back up now

Required NetApp Console role

Storage viewer, Backup and Recovery super admin, Backup and Recovery backup admin role. [Learn about NetApp Console access roles for all services](#).

Protect database instances

You can assign a policy to a database instance using policies that govern the schedules and retention of resource protection.

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select the workload that you want to view and select **View**.
3. Select the **Instances** tab.
4. Select the instance.
5. Select the **Actions**  icon, and select **Protect**.
6. Select a policy or create a new one.

For details about creating a policy, refer to [Create a policy](#).
7. Provide information on the scripts that you want to run before and after the backup.
 - **Pre-script:** Enter your script filename and location to run it automatically before the protect action is triggered. This is helpful for performing additional tasks or configurations that need to be executed before the protection workflow.
 - **Post-script:** Enter your script filename and location to run it automatically after the protection action is complete. This is helpful for performing additional tasks or configurations that need to be executed after the protection workflow.
8. Provide information on how you want the snapshot to be verified:
 - Storage location: Select the location where the verification snapshot will be stored.
 - Verification resource: Select whether the resource that you want to verify is on the local snapshot and on ONTAP secondary storage.
 - Verification schedule: Select the frequency of hourly, daily, weekly, monthly, or yearly.

Associate credentials with a resource

You can associate credentials with a resource so that protection can occur.

For details, see [Configure NetApp Backup and Recovery settings, including credentials](#).

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select the workload that you want to view and select **View**.
3. Select the **Instances** tab.
4. Select the instance.
5. Select the **Actions** **...** icon, and select **Associate credentials**.
6. Use existing credentials or create new ones.

Edit protection settings

You can change the policy, create a new policy, set a schedule, and set retention settings.

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select the workload that you want to view and select **View**.
3. Select the **Instances** tab.
4. Select the instance.
5. Select the **Actions** **...** icon, and select **Edit protection**.

For details about creating a policy, refer to [Create a policy](#).

Back up now

Back up your data now to protect it immediately.

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select the workload that you want to view and select **View**.
3. Select the **Instances** tab.
4. Select the instance.
5. Select the **Actions** **...** icon, and select **Back up now**.
6. Choose the backup type and set the schedule.

For details about creating an ad hoc backup, refer to [Create a policy](#).

Manage database information

You can manage database information in the following ways:

- Protect databases

- Restore databases
- View protection details
- Edit protection settings
- Back up now

Protect databases

You can change the policy, create a new policy, set a schedule, and set retention settings.

Required NetApp Console role

Storage viewer, Backup and Recovery super admin, Backup and Recovery backup admin role. [Learn about NetApp Console access roles for all services](#).

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select the workload that you want to view and select **View**.
3. Select the **Databases** tab.
4. Select the database.
5. Select the **Actions**  icon, and select **Protect**.

For details about creating a policy, refer to [Create a policy](#).

Restore databases

Restore a database to protect your data.

Required NetApp Console role

Storage viewer, Backup and Recovery super admin, Backup and Recovery backup admin role. [Learn about NetApp Console access roles for all services](#).

1. Select the **Databases** tab.
2. Select the database.
3. Select the **Actions**  icon, and select **Restore**.

For information about restoring workloads, refer to [Restore workloads](#).

Edit protection settings

You can change the policy, create a new policy, set a schedule, and set retention settings.

Required NetApp Console role

Storage viewer, Backup and Recovery super admin, Backup and Recovery backup admin role. [Learn about NetApp Console access roles for all services](#).

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select the workload that you want to view and select **View**.
3. Select the **Databases** tab.

4. Select the database.
5. Select the **Actions**  icon, and select **Edit protection**.

For details about creating a policy, refer to [Create a policy](#).

Back up now

You can back up your Microsoft SQL Server instances and databases now to protect your data immediately.

Required NetApp Console role

Storage viewer, Backup and Recovery super admin, Backup and Recovery backup admin role. [Learn about NetApp Console access roles for all services](#).

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select the workload that you want to view and select **View**.
3. Select the **Instances** or **Databases** tab.
4. Select the instance or database.
5. Select the **Actions**  icon, and select **Back up now**.

Manage Microsoft SQL Server snapshots with NetApp Backup and Recovery

You can manage Microsoft SQL Server snapshots by deleting them from NetApp Backup and Recovery.

Delete a snapshot

You can delete only local snapshots.

Required NetApp Console role

Storage viewer, Backup and Recovery super admin, Backup and Recovery backup admin role. [Learn about NetApp Console access roles for all services](#).

Steps

1. In NetApp Backup and Recovery, select **Inventory**.
2. Select the workload and select **View**.
3. Select the **Databases** tab.
4. Select the database that you want to delete a snapshot for.
5. From the Actions menu, select **View protection details**.
6. Select the local snapshot that you want to delete.



Verify that the local snapshot icon in the **Location** column on that row appears in blue.

7. Select the **Actions**  icon, and select **Delete local snapshot**.
8. In the confirmation dialog box, select **Remove**.

Create reports for Microsoft SQL Server workloads in NetApp Backup and Recovery

In NetApp Backup and Recovery, create reports for Microsoft SQL Server workloads to view backup status and details, including counts of successful and failed backups, backup types, storage systems, and timestamps.

Create a report

Required NetApp Console role

Storage viewer, Backup and Recovery super admin, Backup and Recovery backup admin, Backup and Recovery restore admin, Backup and Recovery clone admin. Learn about [Backup and recovery roles and privileges](#). Learn about [NetApp Console access roles for all services](#).

1. From the NetApp Backup and Recovery menu, select the **Reports** option.
2. Select **Create report**.
3. Enter report scope details:
 - **Report name:** Enter a unique name for the report.
 - **Report type:** Choose whether you want a report by account or by workload (Microsoft SQL Server).
 - **Select host:** If you selected by workload, select the host for which you want to generate the report.
 - **Select contents:** Choose whether you want the report to include a summary of all backups or details of each backup. (If you chose "By account")
4. Enter reporting range: Choose whether you want the report to include data from the last day, last 7 days, last 30 days, last quarter, or last year.
5. Enter report delivery details: If you want the report to be delivered by email, check **Send report using email**. Enter the email address where you want the report sent.

Configure email notifications in the Settings page. For details about configuring email notifications, see [Configure settings](#).

Protect VMware workloads

Protect VMware workloads with NetApp Backup and Recovery overview

Protect your VMware VMs and datastores with NetApp Backup and Recovery. NetApp Backup and Recovery provides fast, space-efficient, crash-consistent, and VM-consistent backup and restore operations. You can back up VMware workloads to Amazon Web Services S3 or StorageGRID and restore VMware workloads back to an on-premises VMware host.



This version of NetApp Backup and Recovery supports only VMware vCenter and does not discover vVols or VMs on vVols.

Use NetApp Backup and Recovery to implement a 3-2-1 strategy, where you have 3 copies of your source data on 2 different storage systems along with 1 copy in the cloud. The benefits of the 3-2-1 approach include:

- Multiple data copies protect against internal and external cybersecurity threats.

- Using different types of media helps you recover if one type fails.
- You can quickly restore from the onsite copy, and use the offsite copies if the onsite copy is compromised.



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

You can use NetApp Backup and Recovery to perform the following tasks related to VMware workloads:

- [Discover VMware workloads](#)
- [Create and manage protection groups for VMware workloads](#)
- [Back up VMware workloads](#)
- [Restore VMware workloads](#)

Discover VMware workloads with NetApp Backup and Recovery

The NetApp Backup and Recovery service needs to first discover VMware datastores and VMs running on ONTAP systems in order for you to use the service. You can optionally import backup data and policies from SnapCenter Plug-in for VMware vSphere if you already have it installed.

Required Console role

Backup and Recovery super admin. Learn about [Backup and recovery roles and privileges](#). Learn about [NetApp Console access roles for all services](#).

Discover VMware workloads and optionally import SnapCenter resources

During discovery, NetApp Backup and Recovery analyzes VMware workloads within your organization and assesses and imports existing protection policies, snapshots, and backup and restore options.

You can import VMware NFS and VMFS datastores and VMs from their on-premises SnapCenter Plug-in for VMware vSphere into NetApp Backup and Recovery inventory.



This version of NetApp Backup and Recovery supports only VMware vCenter and does not discover vVols or VMs on vVols.

During the import process, NetApp Backup and Recovery performs the following tasks:

- Enables secure SSH access to the vCenter server.
- Activates maintenance mode on all Resource Groups in the vCenter server.
- Prepares the metadata of the vCenter and marks it as unmanaged in the NetApp Console.
- Configures database access.
- Discovers VMware vCenter, datastores, and VMs.
- Imports existing protection policies, snapshots, and backup and restore options from SnapCenter Plug-in for VMware vSphere.
- Displays the discovered resources in the NetApp Backup and Recovery Inventory page.

Discovery occurs in the following ways:

- If you already have SnapCenter Plug-in for VMware vSphere, import SnapCenter resources into NetApp Backup and Recovery by using the NetApp Backup and Recovery UI.



If you already have SnapCenter Plug-in, ensure you've met the prerequisites before importing from SnapCenter. For example, you should create systems in NetApp Console for all on-premises SnapCenter cluster storage first before importing from SnapCenter. See [Prerequisites for importing resources from SnapCenter](#).

- If you don't already have the SnapCenter Plug-in, you can still discover workloads within your systems by adding a vCenter manually and performing discovery.

If SnapCenter Plug-in is not already installed, add a vCenter and discover resources

If you don't already have SnapCenter Plug-in for VMware installed, add vCenter information and have NetApp Backup and Recovery discover workloads. Within each Console agent, select the systems where you want to discover workloads.

Steps

1. From the NetApp Console left navigation, select **Protection > Backup and Recovery**.

If you are logging in to Backup and Recovery for the first time and have a system in the Console but no discovered resources, the *Welcome to the new NetApp Backup and Recovery* page appears with an option to **Discover resources**.

2. Select **Discover resources**.

3. Enter the following information:

- a. **Workload type:** Select **VMware**.

- b. **vCenter settings:** Add a new vCenter. To add a new vCenter, enter the vCenter FQDN or IP address, user name, password, port, and protocol.



If you are entering vCenter information, enter information for both vCenter settings and Host registration. If you added or entered vCenter information here, you also need to add plugin information in Advanced Settings next.

- c. **Host registration:** Not required for VMware.

4. Select **Discover**.



This process might take a few minutes.

5. Continue with Advanced Settings.

If SnapCenter Plug-in is already installed, import SnapCenter Plug-in for VMware resources into NetApp Backup and Recovery

If you already have SnapCenter Plug-in for VMware installed, import SnapCenter Plug-in resources into NetApp Backup and Recovery using these steps. The Console discovers ESXi hosts, datastores, and VMs in vCenters, and schedules from the Plug-in; you don't have to recreate all that information.

You can do this in the following ways:

- During discovery, select an option to import resources from SnapCenter Plug-in.

- After discovery, from the Inventory page, select an option to import SnapCenter Plug-in resources.
- After discovery, from the Settings menu, select an option to import SnapCenter Plug-in resources. For details, see [Configure NetApp Backup and Recovery](#). This is not supported for VMware.

This is a two-part process described in this section:

1. Import the vCenter metadata from SnapCenter Plug-in. The imported vCenter resources are not yet managed by NetApp Backup and Recovery.
2. Initiate management of selected vCenters, VMs, and datastores in NetApp Backup and Recovery. After you initiate management, NetApp Backup and Recovery labels the vCenter as "Managed" on the Inventory page and is able to back up and recover the resources that you imported. After you initiate management in NetApp Backup and Recovery, you no longer manage those resources in SnapCenter Plug-in.

Import vCenter metadata from SnapCenter Plug-in

This first step imports vCenter metadata from SnapCenter Plug-in. At that point, the resources are not yet managed by NetApp Backup and Recovery.

 After you import vCenter metadata from the SnapCenter Plug-in, NetApp Backup and Recovery does not take over protection management automatically. To do so, you must explicitly select to manage the imported resources in NetApp Backup and Recovery. This ensures that you are ready to have those resources backed up by NetApp Backup and Recovery.

Steps

1. From the Console left navigation, select **Protection > Backup and Recovery**.
2. Select **Inventory**.
3. From the NetApp Backup and Recovery Discover workload resources page, select **Import from SnapCenter**.
4. In the Import from field, select **SnapCenter Plug-in for VMware**.
5. Enter **VMware vCenter credentials**:
 - a. **vCenter IP/hostname**: Enter the FQDN or IP address of the vCenter you want to import into NetApp Backup and Recovery.
 - b. **vCenter port number**: Enter the port number for the vCenter.
 - c. **vCenter Username and Password**: Enter the username and password for the vCenter.
 - d. **Connector**: Select the Console agent for the vCenter.
6. Enter **SnapCenter Plug-in host credentials**:
 - a. **Existing credentials**: If you select this option, you can use the existing credentials that you have already added. Choose the credentials name.
 - b. **Add new credentials**: If you don't have existing SnapCenter Plug-in host credentials, you can add new credentials. Enter the credentials name, authentication mode, user name, and password.
7. Select **Import** to validate your entries and register the SnapCenter Plug-in.



If the SnapCenter Plug-in is already registered, you can update the existing registration details.

Result

The Inventory page shows the vCenter as unmanaged in NetApp Backup and Recovery until you explicitly select to manage it.

Manage resources imported from SnapCenter Plug-in

After you import the vCenter metadata from the SnapCenter Plug-in for VMware, manage the resources in NetApp Backup and Recovery. After you select to manage those resources, NetApp Backup and Recovery is able to back up and recover the resources that you imported. After you initiate the management in NetApp Backup and Recovery, you no longer manage those resources in SnapCenter Plug-in.

After you select to manage the resources, the resources, VMs, and policies are imported from the SnapCenter Plug-in for VMware. The resource groups, policies, and snapshots are migrated from the Plug-in and become managed in NetApp Backup and Recovery.

Steps

1. After you import the VMware resources from SnapCenter Plug-in, from the Backup and Recovery menu, select **Inventory**.
2. From the Inventory page, select the imported vCenter that you want to have NetApp Backup and Recovery manage from now on.
3. Select the Actions icon **...** > **View details** to display the workload details.
4. From the Inventory > workload page, select the Actions icon **...** > **Manage** to display the Manage vCenter page.
5. Check the box "Do you want to continue with the migration?" and select **Migrate**.

Result

The Inventory page shows the newly managed vCenter resources.

Continue to the NetApp Backup and Recovery Dashboard

1. To display the Dashboard, from the Backup and Recovery menu, select **Dashboard**.
2. Review the health of data protection. The number of at risk or protected workloads increases based on the newly discovered, protected, and backed up workloads.

[Learn what the Dashboard shows you.](#)

Create and manage protection groups for VMware workloads with NetApp Backup and Recovery

Create protection groups to manage the backup and restore operations for a set of workloads. A protection group is a logical grouping of resources such as VMs and datastores that you want to protect together.

You can perform the following tasks related to protection groups:

- Create a protection group.
- View protection details.
- Back up a protection group now. See [Back up VMware workloads now](#).
- Suspend and resume a protection group's backup schedule.

- Delete a protection group.

Create a protection group

Group workloads you want to protect into a protection group to back up and restore them together.

Required Console role

Backup and Recovery super admin or Backup and Recovery backup admin role. [Learn about NetApp Console access roles for all services](#).

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select a workload to view the protection details.
3. Select the Actions icon  > **View details**.
4. Select the **Protection groups** tab.
5. Select **Create protection group**.
6. Provide a name for the protection group.
7. Select the VMs or databases that you want to include in the protection group.
8. Select **Next**.
9. Select the **Backup policy** that you want to apply to the protection group.

If you want to create a policy, select **Create new policy** and follow the prompts to create a policy. See [Create policies](#) for more information.

10. Select **Next**.
11. Review the configuration.
12. Select **Create** to create the protection group.

Suspend a protection group's backup schedule

Suspend a protection group to pause its scheduled backups.

The protection status changes to "Under maintenance" when you suspend a protection group. You can resume the backup schedule at any time.

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select a workload to view the protection details.
3. Select the Actions icon  > **View details**.
4. Select the **Protection groups** tab.
5. Select the Actions icon  > **Suspend protection group**.
6. Review the confirmation message and select **Suspend**.

Resume a protection group's backup schedule

Resuming a suspended protection group restarts the scheduled backups for the protection group.

The protection status changes from "Under maintenance" when you suspend a protection group to "Protected" when you resume it. You can resume the backup schedule at any time.

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select a workload to view the protection details.
3. Select the Actions icon  > **View details**.
4. Select the **Protection groups** tab.
5. Select the Actions icon  > **Resume protection group**.
6. Review the confirmation message and select **Resume**.

Result

The system validates the schedules and changes the protection status to "Protected" if the schedules are valid. If the schedules are not valid, the system displays an error message and does not resume the protection group.

Delete a protection group

When you delete a protection group, you remove it and all backup schedules for the group. Delete a protection group if you do not need it anymore.

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select a workload to view the protection details.
3. Select the Actions icon  > **View details**.
4. Select the **Protection groups** tab.
5. Select the protection group that you want to delete.
6. Select the Actions icon  > **Delete**.
7. Review the confirmation message about deleting the associated backups and confirm the deletion.

Back up VMware workloads with NetApp Backup and Recovery

Back up VMware VMs and datastores from on-premises ONTAP systems to Amazon Web Services, Azure NetApp Files, or StorageGRID to ensure that your data is protected. Backups are automatically generated and stored in an object store in your public or private cloud account.

- To back up workloads on a schedule, create policies that govern the backup and restore operations. See [Create policies](#) for instructions.
- Create protection groups to manage the backup and restore operations for a set of resources. See [Create and manage protection groups for VMware workloads with NetApp Backup and Recovery](#) for more information.
- Back up workloads now (create an on-demand backup now).

Back up workloads now with an on-demand backup

Create an on-demand backup immediately. You might want to run an on-demand backup if you're about to make changes to your system and want to ensure that you have a backup before you start.

Required NetApp Console role

Storage viewer, Backup and Recovery super admin, or Backup and Recovery backup admin role. [Learn about NetApp Console access roles for all services](#).

Steps

1. From the Backup and Recovery menu, select **Inventory**.
2. Select a workload to view the protection details.
3. Select the Actions icon  > **View details**.
4. Select the **Protection Groups**, **Datastores** or **Virtual machines** tab.
5. Select the protection group, datastores, or virtual machines that you want to back up.
6. Select the Actions icon  > **Back up now**.



The policy that is applied to the backup is the same policy that is assigned to the protection group, datastore, or virtual machine.

7. Select the schedule tier.
8. Select **Back up now**.

Restore VMware workloads

Restore VMware workloads with NetApp Backup and Recovery

Restore VMware workloads from snapshots, from a workload backup replicated to secondary storage, or from backups stored in object storage using NetApp Backup and Recovery.

Restore from these locations

You can restore workloads from different starting locations:

- Restore from a primary location (local snapshot)
- Restore from a replicated resource on secondary storage
- Restore from an object storage backup

Restore to these points

You can restore data to these points:

- **Restore to the original location:** The VM is restored in the original location, to the same vCenter deployment, ESXi host, and datastore. The VM and all of its data is overwritten.
- **Restore to an alternate location:** You can choose a different vCenter, ESXi host, or datastore as a restore target for the VM. This is useful for managing different copies of the same VM in different locations and states.

Restore from object storage considerations

If Ransomware Resilience is enabled for a backup file in object storage, you are asked to run an extra check before restoring. We recommend performing the scan.

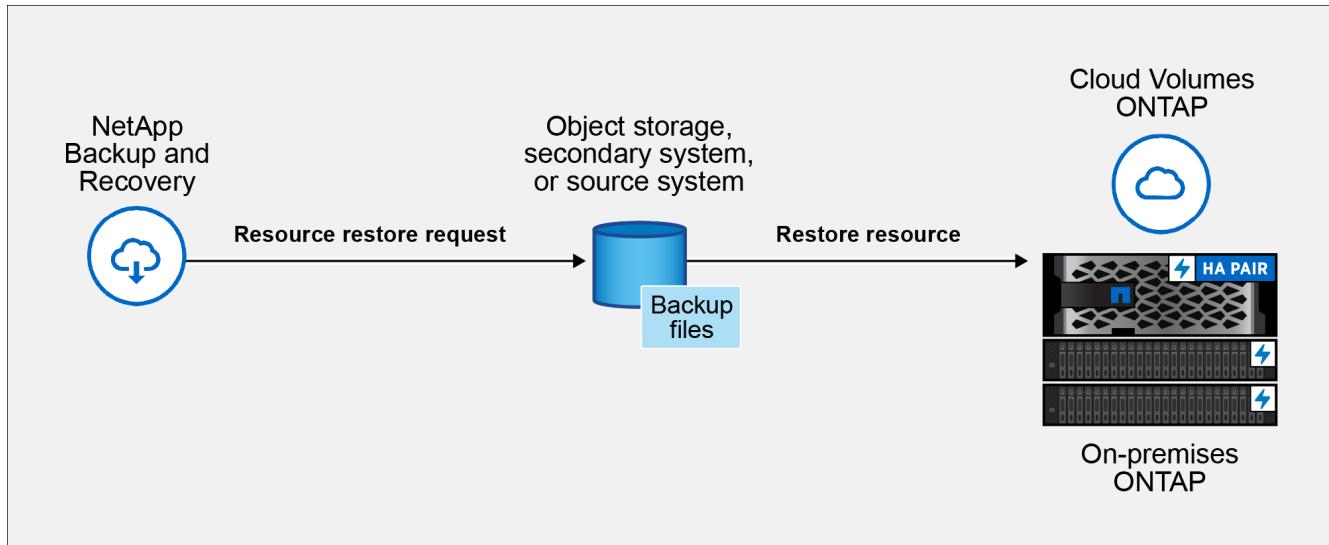


You might pay extra fees to your cloud provider to access the backup file.

How restoring workloads works

When you restore workloads, the following occurs:

- When you restore a workload from a local snapshot or remote backup, NetApp Backup and Recovery overwrites the original VM if you restore to the original location, and creates a *new* resource if you restore to an alternate location.
- When you restore from a replicated workload, you can restore the workload to the original on-premises ONTAP system or to a different on-premises ONTAP system.



- When you restore a backup from object storage, you can restore the data to the original system or to an on-premises ONTAP system.

From the Restore page (Search & Restore), you can restore a resource by searching for the snapshot with filters, even if you do not remember its exact name, location, or last known date.

Restore workload data from the Restore option (Search & Restore)

Restore VMware workloads using the Restore option. You can search for the snapshot by its name or by using filters.

Required NetApp Console role

Storage viewer, Backup and Recovery super admin, Backup and Recovery restore admin role. [Learn about NetApp Console access roles for all services](#).

Steps

1. From the NetApp Backup and Recovery menu, select **Restore**.
2. From the drop-down list to the right of the name search field, select **VMware**.
3. Enter the name of the resource you want to restore or filter for the vCenter, datacenter, or datastore where the resource that you want to restore is located.

A list of virtual machines appears that match your search criteria.

4. Find the VM that you want to restore from in the list, and select the options menu button for that VM.
5. In the resulting menu, select **Restore virtual machine**.

A list of snapshots (restore points) created on that virtual machine appears. By default, the latest snapshots are shown for the time frame that you select in the **Time frame** dropdown.

For each snapshot, any illuminated icons in the **Location** column indicate the storage locations where the snapshot is available (primary, secondary, or object storage).

6. Enable the radio button for the snapshot you want to restore.
7. Select **Next**.

Snapshot location options appear.

8. Select the restore destination for the snapshot:
 - **Local**: Restores the snapshot from the local location.
 - **Secondary**: Restores the snapshot from a remote storage location.
 - **Object store**: Restores the snapshot from object storage.

If you choose secondary storage, select the destination location from the drop-down list.

9. Select **Next** to continue.
10. Choose the restore destination and settings:

Destination selection

Restore to original location

When restoring to the original location, you cannot change the destination vCenter, ESXi host, datastore, or name of the VM. The original VM is overwritten with the restore operation.

- a. Select the **Original location** pane.
- b. Choose from the following options:
 - **Pre-restore options** section:
 - **Prescript**: Enable this option to automate additional tasks by running a custom script before the restore operation begins. Enter the full path for the script that should be run and any arguments that the script takes.
 - **Post-restore options** section:
 - **Restart virtual machine**: Enable this option to restart the virtual machine after the restore operation completes and after the post-restore script is applied.
 - **Postscript**: Enable this option to automate additional tasks by running a custom script after the restore is complete. Enter the full path for the script that should be run and any arguments that the script takes.
- c. Select **Restore**.

Restore to alternate location

When restoring to an alternate location, you can change the destination vCenter, ESXi host, datastore, and name of the VM to create a new copy of the VM in a different location or with a different name.

- a. Select the **Alternate location** pane.
- b. Enter the following information:
 - **Destination settings** section:
 - **vCenter FQDN or IP address**: Select the vCenter server where you want to restore the snapshot.
 - **ESXi host**: Select the host where you want to restore the snapshot.
 - **Network**: Select the network where you want to restore the snapshot.
 - **Datastore**: From the drop-down list, select the name of the datastore where you want to restore the snapshot.
 - **Virtual machine name**: Enter the name of the VM where you want to restore the snapshot. If the name matches a VM that already exists in the datastore, Backup and Recovery makes the name unique by appending a current timestamp.
 - **Pre-restore options** section:
 - **Prescript**: Enable this option to automate additional tasks by running a custom script before the restore operation begins. Enter the full path for the script that should be run and any arguments that the script takes.
 - **Post-restore options** section:
 - **Restart virtual machine**: Enable this option to restart the virtual machine after the restore operation completes and after the post-restore script is applied.
 - **Postscript**: Enable this option to automate additional tasks by running a custom script after the restore is complete. Enter the full path for the script that should be run and any arguments that the script takes.

c. Select **Restore**.

Restore specific virtual disks from backups

You can restore existing virtual disks (VMDKs), or deleted or detached virtual disks, from either a primary or secondary backups of traditional VMs. This enables you to restore only specific VM data or applications, so that you don't need to restore the entire VM and all of its associated virtual disks in situations where only specific data is affected. After the virtual disk is restored, it is attached to its original VM and is ready to use.

You can restore one or more virtual machine disks (VMDKs) on a VM to the same datastore or to different datastores.



For improved performance of restore operations in NFS environments, enable the VMware application vStorage API for Array Integration (VAAI).

Before you begin

- A backup must exist.
- The VM must not be in transit.

The VM that you want to restore must not be in a state of vMotion or Storage vMotion.

About this task

- If the VMDK is deleted or detached from the VM, then the restore operation attaches the VMDK to the VM.
- A restore operation might fail if the storage tier of the FabricPool where the VM is located is unavailable.
- Attach and restore operations connect VMDKs using the default SCSI controller. However, when VMDKs that are attached to a VM with a NVMe disk are backed up, the attach and restore operations use NVMe controller if available.

Steps

1. From the NetApp Backup and Recovery menu, select **Restore**.
2. From the drop-down list to the right of the name search field, select **VMware**.
3. Enter the name of the resource you want to restore or filter for the vCenter, datacenter, or datastore where the resource that you want to restore is located.

A list of virtual machines appears that match your search criteria.

4. Find the VM that you want to restore from in the list, and select the options menu button for that VM.
5. In the resulting menu, select **Restore virtual disks**.

A list of snapshots (restore points) created on that virtual machine appears. By default, the latest snapshots are shown for the time frame that you select in the **Time frame** dropdown.

For each snapshot, any illuminated icons in the **Location** column indicate the storage locations where the snapshot is available (primary, secondary, or object storage).

6. Enable the radio button for the snapshot you want to restore.

7. Select **Next**.

Snapshot location options appear.

8. Select the restore destination for the snapshot:

- **Local**: Restores the snapshot from the local location.
- **Secondary**: Restores the snapshot from a remote storage location.
- **Object store**: Restores the snapshot from object storage.

If you choose secondary storage, select the destination location from the drop-down list.

9. Select **Next** to continue.

10. Choose the restore destination and settings:

Destination selection

Restore to original location

When restoring to the original location, you cannot change the destination vCenter, ESXi host, datastore, or name of the virtual disk. The original virtual disk is overwritten.

- a. Select the **Original location** pane.
- b. In the **Destination settings** section, enable the check box for any virtual disks you want to restore.
- c. Choose from the following options:
 - **Pre-restore options** section:
 - **Prescript**: Enable this option to automate additional tasks by running a custom script before the restore operation begins. Enter the full path for the script that should be run and any arguments that the script takes.
 - **Post-restore options** section:
 - **Postscript**: Enable this option to automate additional tasks by running a custom script after the restore is complete. Enter the full path for the script that should be run and any arguments that the script takes.
 - d. Select **Restore**.

Restore to alternate location

When restoring to an alternate location, you can change the destination datastore. The virtual disk is attached to the original VM after the restore operation regardless of the datastore you choose.

- a. Select the **Alternate location** pane.
- b. In the **Destination settings** section, enable the check box for any virtual disks you want to restore.
- c. For any virtual disks you selected:
 - i. Choose **Select datastore** to choose a different datastore restore target for the virtual disk.
 - ii. Select **Select** to confirm your choice and close the selection window.
- d. Choose from the following options:
 - **Pre-restore options** section:
 - **Prescript**: Enable this option to automate additional tasks by running a custom script before the restore operation begins. Enter the full path for the script that should be run and any arguments that the script takes.
 - **Post-restore options** section:
 - **Postscript**: Enable this option to automate additional tasks by running a custom script after the restore is complete. Enter the full path for the script that should be run and any arguments that the script takes.
- e. Select **Restore**.

Restore guest files and folders

Requirements and limitations when restoring guest files and folders

You can restore files or folders from a virtual machine disk (VMDK) on a Windows guest OS.

Guest restore workflow

Guest OS restore operations include the following steps:

1. Attach

Attach a virtual disk to a guest VM and start a guest file restore session.

2. Wait

Wait for the attach operation to complete before you can browse and restore. When the attach operation finishes, a guest file restore session is automatically created.

3. Select files or folders

Browse the VMDK files and select one or more files or folders to restore.

4. Restore

Restore the selected files or folders to a specified location.

Prerequisites for restoring guest files and folders

Review all requirements before restoring files or folders from a VMDK on a Windows guest OS.

- VMware tools must be installed and running.

NetApp Backup and Recovery uses information from VMware tools to establish a connection to the VMware Guest OS.

- The Windows guest OS must be running Windows Server 2008 R2 or later.

For the latest information about supported versions, refer to [NetApp Interoperability Matrix Tool \(IMT\)](#).

- Credentials for the target VM use the built-in domain or local administrator account with the username "Administrator". Before starting the restore operation, configure the credentials for the VM where you want to attach the virtual disk. Credentials are required for both attach and restore operations. Workgroup users can use the built-in local administrator account.



If you must use an account that is not the built-in administrator account, but has administrative privileges within the VM, you must disable UAC on the guest VM.

- You must know the backup snapshot and VMDK to restore from.

NetApp Backup and Recovery does not support searching of files or folders to restore. Before you begin you must know where the files or folders are in the snapshot and the corresponding VMDK.

- Virtual disk to be attached must be in a NetApp Backup and Recovery backup.

The virtual disk that contains the file or folder you want to restore must be in a VM backup that was performed using NetApp Backup and Recovery.

- For files with non-English-alphabet names, you must restore them in a directory, not as a single file.

You can restore files with non-alphabetic names, such as Japanese Kanji, by restoring the directory in which the files are located.

Guest file restore limitations

Before you restore a file or folder from a guest OS, you should be aware of the feature limitations.

- You cannot restore dynamic disk types inside a guest OS.
- If you restore an encrypted file or folder, the encryption attribute is not retained.
- You cannot restore files or folders to an encrypted folder.
- Hidden files and folders are displayed in the file browse page, and you cannot filter them.
- You cannot restore from a Linux guest OS.

You cannot restore files and folders from a VM that is running Linux guest OS. However, you can attach a VMDK and then manually restore the files and folders. For the latest information on supported guest OS, refer to the [NetApp Interoperability Matrix Tool \(IMT\)](#).

- You cannot restore from a NTFS file system to a FAT file system.

When you try to restore from NTFS-format to FAT-format, the NTFS security descriptor is not copied because the FAT file system does not support Windows security attributes.

- You cannot restore guest files from a cloned VMDK or an uninitialized VMDK.
- You cannot restore the directory structure for a file.

When you restore a file from a nested directory, the system restores only the file, not its directory structure. To restore the entire directory tree, copy the top-level directory.

- You cannot restore guest files from a vVol VM to an alternate host.
- You cannot restore encrypted guest files.

Restore guest files and folders from VMDKs

You can restore one or more files or folders from a VMDK on a Windows guest OS.

Before you begin

You need to create credentials for the guest VM in NetApp Backup and Recovery before you can restore files and folders from it. NetApp Backup and Recovery uses these credentials to authenticate with the guest VM when attaching the virtual disk.

About this task

Guest file or folder restore performance depends upon two factors: the size of the files or folders being restored; and the number of files or folders being restored. Restoring a large number of small-sized files might take a longer time than anticipated compared to restoring a small number of large-sized files, if the data set to be restored is of same size.



Only one attach or restore operation can run at the same time on a VM. You cannot run parallel attach or restore operations on the same VM.



With the guest restore feature, you can view and restore system and hidden files and view encrypted files. Do not overwrite an existing system file or restore encrypted files to an encrypted folder. During the restore operation, the hidden, system, and encrypted attributes of guest files are not kept in the restored file. Viewing or browsing reserved partitions might cause an error.

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select the **Virtual machines** menu.
3. Choose a virtual machine from the list that contains files that you want to restore.
4. Select the Actions icon for that VM.
5. Select **Restore files and folders**.
6. Select a snapshot from which to restore and then select **Next**.
7. Choose the snapshot location to restore from. If you choose a secondary location, select the secondary snapshot from the list.
8. Select **Next**.
9. Choose virtual disk from the list to attach to the VM and then select **Next**.
10. On the *Select virtual machine credential* page, if you haven't yet stored a credential for the guest VM, select **Add credentials** and do the following:
 - a. **Credentials name:** Enter a name for the credentials.
 - b. **Authentication mode:** Select **Windows**.
 - c. **Agents:** Select a Console agent from the list that will handle communication between NetApp Backup and Recovery and this host.
 - d. **Domain and user name:** Enter the NetBIOS or domain FQDN and user name for the credentials.
 - e. **Password:** Enter a password for the credential.
 - f. Select **Add**.
11. Choose a virtual machine credential to use to authenticate with the guest VM.

NetApp Backup and Recovery attaches the virtual disk to the VM and displays all files and folders, including hidden ones. It assigns a drive letter to every partition, including system reserved partitions.

Files and folders you have selected are listed in the right pane of the screen.

12. Select **Next**.
13. Enter the UNC share path to the guest where the selected files will be restored.

- IPv4 address example: \\10.60.136.65\c\$
- IPv6 address example: \\fd20-8b1e-b255-832e-61.ipv6-literal.net\c\restore

If there are existing files with the same name, you can choose to overwrite or skip them.

14. Select **Restore**.

You can view the restore progress on the Job Monitoring page.

Guest file restore troubleshooting

When attempting to restore a guest file, you might encounter any of the following scenarios.

Guest file restore session is blank

This issue occurs if you create a guest file restore session and the guest operating system reboots during the session. VMDKs in the guest OS might stay offline, so the guest file restore session list is blank.

To correct the issue, manually put the VMDKs back online in the guest OS. When the VMDKs are online, the guest file restore session will display the correct contents.

Guest file restore attach disk operation fails

This issue occurs when you start a guest file restore operation, but the attach disk operation fails even though VMware tools is running and the Guest OS credentials are correct. If this occurs, the following error is returned:

```
Error while validating guest credentials, failed to access guest system using
specified credentials: Verify VMWare tools is running properly on system and
account used is Administrator account, Error is SystemError vix error codes =
(3016, 0).
```

To correct the issue, restart the VMware Tools Windows service on the Guest OS, and then retry the guest file restore operation.

Backups are not detached after guest file restore session is discontinued

This issue occurs when you perform a guest file restore operation from a VM-consistent backup. While the guest file restore session is active, another VM-consistent backup is performed for the same VM. When the guest file restore session is disconnected, either manually or automatically after 24 hours, the backups for the session are not detached.

To correct the issue, manually detach the VMDKs that were attached from the active guest file restore session.

Protect KVM workloads (Preview)

Protect KVM workloads overview

Protect your managed KVM VMs and storage pools with NetApp Backup and Recovery. NetApp Backup and Recovery provides fast, space-efficient, crash-consistent, and VM-consistent backup and restore operations. Your KVM hosts and VMs must be managed by a management platform such as Apache CloudStack before you can protect them using Backup and Recovery.

You can back up KVM workloads to Amazon Web Services S3, Azure NetApp Files, or StorageGRID and restore KVM workloads back to an on-premises KVM host.

Use NetApp Backup and Recovery to implement a 3-2-1 protection strategy, where you have 3 copies of your source data on 2 different storage systems along with 1 copy in the cloud. The benefits of the 3-2-1 approach include:

- Multiple data copies protect against internal and external cybersecurity threats.
- Using different types of media helps you recover if one type fails.
- You can quickly restore from the onsite copy, and use the offsite copies if the onsite copy is compromised.



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

You can use NetApp Backup and Recovery to perform the following tasks related to KVM workloads:

- [Discover KVM workloads](#)
- [Create and manage protection groups for KVM workloads](#)
- [Back up KVM workloads](#)
- [Restore KVM workloads](#)

Discover KVM workloads in NetApp Backup and Recovery

NetApp Backup and Recovery needs to discover KVM hosts and virtual machines before protecting them. Your KVM hosts and VMs must be managed by a management platform such as Apache CloudStack before you can add them to Backup and Recovery.

Required Console role

Backup and Recovery super admin. Learn about [Backup and recovery roles and privileges](#). Learn about [NetApp Console access roles for all services](#).

Add a management platform, KVM host, and discover resources

Add management platform and KVM host information and let NetApp Backup and Recovery discover workloads.

Steps

1. From the NetApp Console menu, select **Protection > Backup and Recovery**.
2. Under **Workloads**, select the **KVM** tile.

If you are logging in to Backup and Recovery for the first time and have a system in the Console but no discovered resources, the *Welcome to the new NetApp Backup and Recovery* page appears with an option to **Discover resources**.

3. Select **Discover resources**.
4. Enter the following information:
 - a. **Workload type:** Select **KVM**.
 - b. If you haven't yet integrated your management platform with Backup and Recovery, select **Add management platform**.
 - i. Enter the following information:
 - **Management platform IP address or FQDN:** Enter the IP address or fully qualified domain name of the management platform.
 - **API key:** Enter the API key to use to authenticate API requests.
 - **Secret Key:** Enter the secret key to use to authenticate API requests.

- **Port:** Enter the port to use for communication between Backup and Recovery and the management platform.
- **Agents:** Select a Console agent to use to facilitate communication between Backup and Recovery and the management platform.

ii. When finished, select **Add**.

c. **KVM settings:** Add a new KVM host by entering the following information:

- **KVM FQDN or IP address:** Enter the host's FQDN or IP address.
- **Credentials:** Enter the username and password for the KVM host.
- **Console agent:** Choose the Console agent to use for communication between Backup and Recovery and the KVM host.
- **Port number:** Enter the port to use for communication between Backup and Recovery and the KVM host.
- **Management platform:** If the KVM host is managed and you have added the management platform to Backup and Recovery, select the management platform from the list.

5. Select **Discover**.



This process might take a few minutes.

Result

The KVM workload is displayed in the list of workloads on the Inventory page.

Continue to the NetApp Backup and Recovery Dashboard

Steps

1. From the NetApp Console menu, select **Protection > Backup and recovery**.
2. Select a workload tile (for example, Microsoft SQL Server).
3. From the Backup and Recovery menu, select **Dashboard**.
4. Review the health of data protection. The number of at risk or protected workloads increases based on the newly discovered, protected, and backed up workloads.

Create and manage protection groups for KVM workloads with NetApp Backup and Recovery

Create protection groups to manage the backup operations for a set of KVM resources. A protection group is a logical grouping of resources such as VMs and storage pools that you want to protect together. You need to create a protection group to back up KVM virtual machines or storage pools.

You can perform the following tasks related to protection groups:

- Create a protection group.
- View protection details.
- Back up a protection group now. See [Back up KVM workloads now](#).
- Delete a protection group.

Create a protection group

Group VMs and storage pools that you want to protect together into a protection group.

Required Console role

Backup and Recovery super admin or Backup and Recovery backup admin role. [Learn about NetApp Console access roles for all services](#).

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select a workload to view the protection details.
3. Select the Actions icon **...** > **View details**.
4. Select the **Protection groups** tab.
5. Select **Create protection group**.
6. Provide a name for the protection group.
7. Select the VMs or storage pools that you want to include in the protection group.
8. Select **Next**.
9. Select the **Backup policy** that you want to apply to the protection group.

For more information about creating a backup policy, refer to [Create and manage policies](#).

10. Select **Next**.
11. Review the configuration.
12. Select **Create** to create the protection group.

Delete a protection group

Deleting a protection group removes it and all associated backup schedules. You might want to delete a protection group if it is no longer needed.

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select a workload to view the protection details.
3. Select the Actions icon **...** > **View details**.
4. Select the **Protection groups** tab.
5. Select the protection group that you want to delete.
6. Select the Actions icon **...** > **Delete**.
7. Review the confirmation message about deleting the associated backups and confirm the deletion.

Back up KVM workloads with NetApp Backup and Recovery

Back up KVM protection groups from on-premises ONTAP systems to Amazon Web Services, Azure NetApp Files, or StorageGRID to ensure that your data is protected. When you back up a protection group, the NetApp Console backs up the VMs and storage pools contained in the protection group. Backups are automatically generated and stored in an object store in your public or private cloud account.



To back up protection groups on a schedule, create policies that govern the backup and restore operations. See [Create policies](#) for instructions.

- Create protection groups to manage the backup and restore operations for a set of resources. See [Create and manage protection groups for KVM workloads with NetApp Backup and Recovery](#) for more information.

Back up protection groups now with an on-demand backup

You can run an on-demand backup immediately. This is helpful if you're about to make changes to your system and want to ensure that you have a backup before you start.

Required Console role

Backup and Recovery super admin or Backup and Recovery backup admin role. [Learn about NetApp Console access roles for all services](#).

Steps

1. From the NetApp Console menu, select **Protection > Backup and Recovery**.
2. In the KVM tile, select **Discover and manage**.
3. Select **Inventory**.
4. Select a workload to view the protection details.
5. Select the Actions icon > **View details**.
6. Select the **Protection Groups, Datastores or Virtual machines** tab.
7. Select the protection group that you want to back up.
8. Select the Actions icon > **Back up now**.



The policy that is applied to the backup is the same policy that is assigned to the protection group.

9. Select the schedule tier.
10. Select **Back up**.

Restore KVM virtual machines with NetApp Backup and Recovery

Restore KVM virtual machines from snapshots, from a protection group backup replicated to secondary storage, or from backups stored in object storage using NetApp Backup and Recovery.

Restore from these locations

You can restore virtual machines from different starting locations:

- Restore from a primary location (local snapshot)
- Restore from a replicated resource on secondary storage
- Restore from an object storage backup

Restore to these points

You can restore data to these points:

- Restore to the original location

Restore from object storage considerations

If you select a backup file in object storage, and ransomware protection is active for that backup (if you enabled DataLock and Ransomware Resilience in the backup policy), then you are prompted to run an additional integrity check on the backup file before restoring the data. We recommend that you perform the scan.



You'll incur extra egress costs from your cloud provider to access the contents of the backup file.

How restoring virtual machines works

When you restore virtual machines, the following occurs:

- When you restore a workload from a local backup file, NetApp Backup and Recovery creates a *new* resource using the data from the backup.
- When you restore from a replicated VM, you can restore it to the original system or to an on-premises ONTAP system.
- When you restore a backup from object storage, you can restore the data to the original system or to an on-premises ONTAP system.

From the Restore page (also known as Search & Restore), you can restore a VM, even if you don't remember the exact name, the location in which it resides, or the date when it was last in good shape. You can search for the snapshot using filters.

Restore VMs from the Restore option (Search & Restore)

Restore KVM virtual machines using the Restore option. You can search for the snapshot by its name or by using filters.

Required Console role

Backup and Recovery super admin or Backup and Recovery restore admin role. [Learn about NetApp Console access roles for all services](#).

Steps

1. From the NetApp Console menu, select **Protection > Backup and Recovery**.
2. From the NetApp Backup and Recovery menu, select **Restore**.
3. From the drop-down list to the right of the name search field, select **KVM**.
4. Enter the name of the VM you want to restore or filter for VM host or storage pool where the resource that you want to restore is located.

A list of snapshots appears that match your search criteria.

5. Select the **Restore** button for the snapshot that you want to restore.

A list of possible restore points appears.

6. Select the restore point that you want to use.

7. Select a snapshot source location.

1. Select **Next** to continue.

2. Choose the restore destination and settings:

Destination selection

Restore to original location

2. **Enable quick restore:** Select this to perform a quick restore operation. Restored volumes and data will be available immediately. Do not use this on volumes that require high performance because during the quick restore process, access to the data might be slower than usual.
3. **Pre-restore options:** Enter the full path for a script that should be run before the restore operation and any arguments that the script takes.
4. **Post-restore options:**
 - **Restart VM:** Select this to restart the VM after the restore operation completes and after the post-restore script is applied.
 - **Postscript:** Enter the full path for a script that should be run after the restore operation and any arguments that the script takes.
5. **Notification section:**
 - **Enable email notifications:** Select this to receive email notifications about the restore operation and indicate what type of notifications you want to receive.
6. Select **Restore**.

Restore to alternate location

Not available for KVM workloads preview.

Protect Hyper-V workloads

Protect Hyper-V workloads overview

Protect your Hyper-V VMs with NetApp Backup and Recovery. NetApp Backup and Recovery provides fast, space-efficient, crash-consistent, and VM-consistent backup and restore operations for both standalone and FCI cluster instances. You can also protect Hyper-V virtual machines provisioned by System Center Virtual Machine Manager (SCVMM) and hosted on a CIFS share.

You can back up Hyper-V workloads to Amazon Web Services S3 or StorageGRID and restore Hyper-V workloads back to an on-premises Hyper-V host.

Use NetApp Backup and Recovery to implement a 3-2-1 protection strategy, where you have 3 copies of your source data on 2 different storage systems along with 1 copy in the cloud. The benefits of the 3-2-1 approach include:

- Multiple data copies protect against internal and external cybersecurity threats.
- Multiple media types ensure failover viability in the case of physical or logical failure of one media type.
- The onsite copy helps you quickly restore data, and you can use the offsite copies if the onsite copy is compromised.

When you add Hyper-V hosts and discover resources, NetApp Backup and Recovery installs the NetApp

Hyper-V plug-in and the NetApp SnapCenter Windows FileSystem plug-in on the Hyper-V host to help with managing and protecting virtual machines.



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

You can use NetApp Backup and Recovery to perform the following tasks related to Hyper-V workloads:

- [Discover Hyper-V workloads](#)
- [Create and manage protection groups for Hyper-V workloads](#)
- [Back up Hyper-V workloads](#)
- [Restore Hyper-V workloads](#)

Discover Hyper-V workloads in NetApp Backup and Recovery

NetApp Backup and Recovery must discover Hyper-V virtual machines before you can protect them.

Required Console role

Backup and Recovery super admin. Learn about [Backup and recovery roles and privileges](#). Learn about [NetApp Console access roles for all services](#).

Add a Hyper-V host and discover resources

Add Hyper-V host information and let NetApp Backup and Recovery discover virtual machines. Within each Console agent, select the systems where you want to discover the resources.



When you add Hyper-V hosts and discover resources, NetApp Backup and Recovery installs the NetApp Hyper-V plug-in and the NetApp SnapCenter Windows FileSystem plug-in on the Hyper-V host to help with managing and protecting virtual machines.

Steps

1. From the NetApp Console menu, select **Protection > Backup and recovery**.

If this is your first time logging in to NetApp Backup and Recovery, you already have a system in the Console, but haven't discovered any resources, the "Welcome to the new NetApp Backup and Recovery" landing page appears and shows an option to **Discover resources**.

2. Select **Discover resources**.

3. Enter the following information:

- a. **Workload type:** Select **Hyper-V**.

- b. If you haven't yet stored credentials for this Hyper-V host, select **Add credentials**.

- i. Select the Console agent to use with this host.

- ii. Enter a name for this credential.

- iii. Enter the user name and password for the account.

- iv. Select **Done**.

- c. **Host registration:** Add a new Hyper-V host by entering the host's FQDN or IP address, credentials, Console agent, and port number. If the FQDN is not resolvable by the Console agent, use the IP

address instead. For FCI clusters, enter the FCI cluster management IP address.

4. Select **Discover**.



This process might take a few minutes.

Result

After NetApp Backup and Recovery discovers resources, the Inventory page displays the Hyper-V workload in the list of workloads.

Continue to the NetApp Backup and Recovery Dashboard

Steps

1. From the NetApp Console menu, select **Protection > Backup and recovery**.
2. Select a workload tile (for example, Microsoft SQL Server).
3. From the Backup and Recovery menu, select **Dashboard**.
4. Review the health of data protection. The number of at risk or protected workloads increases based on the newly discovered, protected, and backed up workloads.

Create and manage protection groups for Hyper-V workloads with NetApp Backup and Recovery

Create protection groups to manage the backup operations for a set of virtual machines. A protection group is a logical grouping of resources such as VMs that you want to protect together.

You can perform the following tasks related to protection groups:

- Create a protection group.
- View protection details.
- Back up a protection group now. See [Back up Hyper-V workloads now](#).
- Delete a protection group.

Create a protection group

Group workloads that you want to protect together into a protection group. Create a protection group to back up and restore workloads together.

Required Console role

Backup and Recovery super admin or Backup and Recovery backup admin role. [Learn about NetApp Console access roles for all services](#).

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select a workload to view the protection details.
3. Select the Actions icon **...** > **View details**.
4. Select the **Protection groups** menu.
5. Select **Create protection group**.

6. Provide a name for the protection group.
7. Select the VMs that you want to include in the protection group.
8. Select **Next**.
9. Select the **Backup policy** that you want to apply to the protection group.
10. Select **Next**.
11. Review the configuration.
12. Select **Create** to create the protection group.

Edit a protection group

Edit a protection group to change its name or settings. You might want to edit a protection group if the resources in the group have changed.

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select a workload to view the protection details.
3. Select the Actions icon **...** > **View details**.
4. Select the **Protection groups** tab.
5. Select the protection group that you want to edit.
6. Select the Actions icon **...** > **Edit**.
7. Change any settings for the protection group such as the name or what virtual machines are in the group.
8. Select **Next**.
9. Change the protection policy if needed. When finished, select **Next**.
10. Review the configuration and select **Submit**.

Delete a protection group

Deleting a protection group removes it and all associated backup schedules. You might want to delete a protection group if it is no longer needed.

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select a workload to view the protection details.
3. Select the Actions icon **...** > **View details**.
4. Select the **Protection groups** tab.
5. Select the protection group that you want to delete.
6. Select the Actions icon **...** > **Delete**.
7. Review the confirmation message about deleting the associated backups and confirm the deletion.

Back up Hyper-V workloads with NetApp Backup and Recovery

Back up Hyper-V VMs from on-premises ONTAP systems to Amazon Web Services, Azure NetApp Files, or StorageGRID to ensure that your data is protected. Backups are

automatically generated and stored in an object store in your public or private cloud account.

- To back up workloads on a schedule, create policies that govern the backup and restore operations. See [Create policies](#) for instructions.
- Create protection groups to manage the backup and restore operations for a set of resources. See [Create and manage protection groups for Hyper-V workloads with NetApp Backup and Recovery](#) for more information.
- Back up workloads now (create an on-demand backup now).

Back up workloads now with an on-demand backup

Use on-demand backup so that your data is protected before making system changes.

Required Console role

Backup and Recovery super admin or Backup and Recovery backup admin role. [Learn about NetApp Console access roles for all services](#).

Steps

1. From the menu, select **Inventory**.
2. Select a workload to view the protection details.
3. Select the Actions icon  > **View details**.
4. Select the **Protection Groups, Datastores or Virtual machines** tab.
5. Select the protection group or virtual machines that you want to back up.
6. Select the Actions icon  > **Back up now**.



The backup uses the same policy that you assigned to the protection group or virtual machine.

7. Select the schedule tier.
8. Select **Back up**.

Restore Hyper-V workloads with NetApp Backup and Recovery

Restore Hyper-V workloads from snapshots, from a workload backup replicated to secondary storage, or from backups stored in object storage using NetApp Backup and Recovery.

Restore from these locations

You can restore workloads from different starting locations:

- Restore from a primary location (local snapshot)
- Restore from a replicated resource on secondary storage
- Restore from an object storage backup

Restore to these points

You can restore data to these points:

- Restore to the original location
- Restore to an alternate location

Restore from object storage considerations

If you select a backup file in object storage, and ransomware protection is active for that backup (if you enabled DataLock and Ransomware Resilience in the backup policy), then you are prompted to run an additional integrity check on the backup file before restoring the data. We recommend that you perform the scan.



You'll incur extra egress costs from your cloud provider to access the contents of the backup file.

How restoring workloads works

When you restore workloads, the following occurs:

- When you restore a workload from a local backup file, NetApp Backup and Recovery creates a *new* resource using the data from the backup.
- When you restore from a replicated workload, you can restore the workload to the original system or to an on-premises ONTAP system.

From the Restore page (also known as Search & Restore), you can restore a resource, even if you don't remember the exact name, the location in which it resides, or the date when it was last in good shape. You can search for the snapshot using filters.

Restore workload data from the Restore option (Search & Restore)

Restore Hyper-V workloads using the Restore option. You can search for the snapshot by its name or by using filters.

Required Console role

Backup and Recovery super admin or Backup and Recovery restore admin role. [Learn about NetApp Console access roles for all services](#).

Steps

1. From the NetApp Backup and Recovery menu, select **Restore**.
2. From the drop-down list to the right of the name search field, select **Hyper-V**.
3. Enter the name of the resource you want to restore or filter for the VM name, VM host, or storage pool where the resource that you want to restore is located.

A list of snapshots appears that match your search criteria.

4. Select the **Restore** button for the snapshot that you want to restore.

A list of possible restore points appears.

5. Select the restore point that you want to use.
6. Select a snapshot source location.
7. Select **Next** to continue.
8. Choose the restore destination and settings:

Destination selection

Restore to original location

When you restore to the original location, you can view the destination settings by expanding the **Destination settings** section, but you cannot change them.

- a. In the **Post-restore options** section, consider the following option:
 - **Start the virtual machine**: Enable this option to boot the new virtual machine after it is restored.
- b. Select **Restore**.

Restore to alternate location

- a. In the **Destination settings**: section, enter the following information:
 - **Hyper-V FQDN or IP address**: Enter the fully qualified domain name or IP address of the destination Hyper-V host.
 - **Network**: Select the destination network where you want to restore the snapshot.
 - **Virtual machine name**: Enter the name of the VM that you want to restore.
 - **Destination location**: Enter the destination folder or CIFS share that should contain the restored data.
- b. In the **Pre restore options** section, consider the following options:
 - **Quick restore**: Enable this option to make the restored VM available immediately. Only the files needed to run the VM are restored from the object store, rather than the entire volume.
- c. In the **Post restore options** section, consider the following options:
 - **Start the virtual machine**: Enable this option to boot the new virtual machine after it is restored.
- d. Select **Restore**.

Protect Oracle Database workloads (Preview)

Protect Oracle Database workloads overview

Protect Oracle databases and logs using NetApp Backup and Recovery. Get fast, space-efficient, crash-consistent, and database-consistent backups and restores. Back up Oracle Database workloads to AWS S3, NetApp StorageGRID, Azure Blob Storage, or ONTAP S3. Restore backups to an on-premises Oracle host.

Use NetApp Backup and Recovery to implement a 3-2-1 protection strategy, where you have 3 copies of your source data on 2 different storage systems along with 1 copy in the cloud. The benefits of the 3-2-1 approach include:

- Multiple data copies protect against internal and external cybersecurity threats.
- Using different types of media helps you recover if one type fails.
- You can quickly restore from the onsite copy, and use the offsite copies if the onsite copy is compromised.



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

You can use NetApp Backup and Recovery to perform the following tasks related to Oracle Database workloads:

- [Discover Oracle Database workloads](#)
- [Create and manage protection groups for Oracle Database workloads](#)
- [Back up Oracle Database workloads](#)
- [Restore Oracle Database workloads](#)

Discover Oracle Database workloads in NetApp Backup and Recovery

NetApp Backup and Recovery needs to first discover your Oracle databases so that you can protect them.

Required Console role

Backup and Recovery super admin. Learn about [Backup and recovery roles and privileges](#). Learn about [NetApp Console access roles for all services](#).

Add an Oracle host and discover resources

Add Oracle host information and let NetApp Backup and Recovery discover workloads. Within each Console agent, select the systems where you want to discover workloads.

Steps

1. From the NetApp Console menu, select **Protection > Backup and Recovery**.

2. Under **Workloads**, select the **Oracle** tile.

If you are logging in to Backup and Recovery for the first time and have a system in the Console but no discovered resources, the *Welcome to the new NetApp Backup and Recovery* page appears with an option to **Discover resources**.

3. Select **Discover resources**.

4. Enter the following information:

a. **Workload type**: Select **Oracle**.

b. If you haven't yet stored credentials for this Oracle host, select **Add credentials**.

i. Select the Console agent to use with this host.

ii. Enter a name for this credential.

iii. Enter the user name and password for the account.

iv. Select **Done**.

c. **Host registration**: Add a new Oracle host. Enter the host's FQDN or IP address, credentials, Console agent, and port number.

5. Select **Discover**.



This process might take a few minutes.

Result

The Oracle workload is displayed in the list of workloads on the Inventory page.

Continue to the NetApp Backup and Recovery Dashboard

1. From the NetApp Console menu, select **Protection > Backup and recovery**.
2. Select a workload tile (for example, Microsoft SQL Server).
3. From the Backup and Recovery menu, select **Dashboard**.
4. Review the health of data protection. The number of at risk or protected workloads increases based on the newly discovered, protected, and backed up workloads.

Create and manage protection groups for Oracle Database workloads with NetApp Backup and Recovery

Create protection groups to manage the backup operations for a set of Oracle Database resources. A protection group is a logical grouping of resources such as databases that you want to protect together. You need to create a protection group to back up Oracle databases.

You can perform the following tasks related to protection groups:

- Create a protection group.
- View protection details.
- Back up a protection group now. See [Back up Oracle Database workloads now](#).
- Delete a protection group.

Create a protection group

Group VMs and storage pools that you want to protect together into a protection group.

Required Console role

Backup and Recovery super admin or Backup and Recovery backup admin role. [Learn about NetApp Console access roles for all services](#).

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select a workload to view the protection details.
3. Select the Actions icon  > **View details**.
4. Select the **Protection groups** tab.
5. Select **Create protection group**.
6. Provide a name for the protection group.
7. Select the VMs or storage pools that you want to include in the protection group.
8. Select **Next**.
9. Select the **Backup policy** that you want to apply to the protection group.

If you want to create a policy, select **Create new policy** and follow the prompts to create a policy. See [Create policies](#) for more information.

10. Select **Next**.

11. Review the configuration.
12. Select **Create** to create the protection group.

Delete a protection group

Deleting a protection group removes it and all associated backup schedules. You might want to delete a protection group if it is no longer needed.

Steps

1. From the NetApp Backup and Recovery menu, select **Inventory**.
2. Select a workload to view the protection details.
3. Select the Actions icon  > **View details**.
4. Select the **Protection groups** tab.
5. Select the protection group that you want to delete.
6. Select the Actions icon  > **Remove protection**.
7. Review the confirmation message about deleting the associated backups and confirm the deletion.

Back up Oracle Database workloads using NetApp Backup and Recovery

Use NetApp Backup and Recovery to back up Oracle Database protection groups or databases from on-premises ONTAP systems to cloud storage, including Amazon S3, NetApp StorageGRID, Microsoft Azure Blob Storage, or ONTAP S3. NetApp Backup and Recovery backs up databases and log data in each protection group.



To back up protection groups or single databases on a schedule, create policies that manage backup and restore operations. See [Create policies](#) for instructions.

- Create protection groups to manage the backup and restore operations for a set of resources. See [Create and manage protection groups for Oracle Database workloads with NetApp Backup and Recovery](#) for more information.
- Back up a protection group now (create an on-demand backup now).
- Back up a database now.

Back up protection groups now with an on-demand backup

Run an on-demand backup before making system changes to ensure your data is protected.

Required Console role

Backup and Recovery super admin or Backup and Recovery backup admin role. [Learn about NetApp Console access roles for all services](#).

Steps

1. From the NetApp Console menu, select **Protection > Backup and Recovery**.
2. Under **Workloads**, select the **Oracle** tile.
3. Select **Inventory**.
4. Select a workload to view the protection details.

5. Select the Actions icon  > **View details**.
6. Select the **Protection Groups, Datastores or Virtual machines** tab.
7. Select the protection group that you want to back up.
8. Select the Actions icon  > **Back up now**.



NetApp Backup and Recovery uses the same policy for both the backup and the protection group.

9. Select the schedule tier.
10. Select **Back up**.

Back up a database now with an on-demand backup

You can run an on-demand backup of a single database.

Required Console role

Backup and Recovery super admin or Backup and Recovery backup admin role. [Learn about NetApp Console access roles for all services](#).

Steps

1. From the NetApp Console menu, select **Protection > Backup and Recovery**.
2. Under **Workloads**, select the **Oracle** tile.
3. Select **Inventory**.
4. Select a workload to view the protection details.
5. Select the Actions icon  > **View details**.
6. Select the **Databases** tab.
7. Select the database that you want to back up.
8. Select the Actions icon  > **Back up now**.
9. Select the schedule tier.
10. Select **Back up**.

Restore Oracle databases with NetApp Backup and Recovery

Restore Oracle databases from snapshots, from a backup replicated to secondary storage, or from backups stored in object storage using NetApp Backup and Recovery.

Restore from these locations

You can restore databases from different starting locations:

- Restore from a primary location (local snapshot)
- Restore from a replicated resource on secondary storage
- Restore from an object storage backup

Restore to these points

You can restore data to the original location; restoring to an alternate location is not available in this private preview release.

- Restore to the original location

How restoring Oracle databases works

When you restore Oracle databases, the following occurs:

- When you restore a database from a local snapshot, NetApp Backup and Recovery creates a *new* resource using the data from the backup.
- When you restore from replicated storage, you can restore it to the original location.
- When you restore a backup from object storage, you can restore the data to the source storage or to an on-premises ONTAP system and recover the database from there.

From the Restore page (also known as Search & Restore), you can restore a database, even if you don't remember the exact name, the location in which it resides, or the date when it was last in good shape. You can search for the database using filters.

Restore an Oracle database

Depending on your needs, restore an Oracle database to a specific point in time, to a specific system change number (SCN), or to the last good state. You can also simply restore the database from snapshots and skip the automated recovery process. You might want to skip the automated recovery process if you want to perform recovery manually. You can search for the database using its name or with specific filters.

Required Console role

Backup and Recovery super admin or Backup and Recovery restore admin role. [Learn about NetApp Console access roles for all services](#).

Steps

1. From the NetApp Console menu, select **Protection > Backup and Recovery**.
2. From the NetApp Backup and Recovery menu, select **Restore**.
3. From the drop-down list to the right of the name search field, select **Oracle**.
4. Enter the name of the database you want to restore or filter for the database host where the database that you want to restore is located.

A list of snapshots appears that match your search criteria.

5. Select the **Restore** button for the database that you want to restore.
6. Choose a restore option:

Restore to specific point in time

- a. Select **Restore to specific point in time**.
- b. Select **Next**.
- c. Choose a date from the dropdown, and select **Search**.

A list of matching snapshots on the specified date are displayed.

Restore to a specific system change number (SCN)

- a. Select **Restore to a specific system change number (SCN)**.
- b. Select **Next**.
- c. Enter the SCN to use as a restore point, and select **Search**.

A list of matching snapshots for the specified SCN are displayed.

Restore to the latest backup (last good state)

- a. Select **Restore to the latest backup**.
- b. Select **Next**.

The latest full and log backups are displayed.

Restore from snapshots with no recovery

- a. Select **Restore from snapshots with no recovery**.
- b. Select **Next**.

The matching snapshots are displayed.

7. Select a snapshot source location.
8. Select **Next** to continue.
9. Choose the restore destination and settings:

Destination selection

Restore to original location

2. Destination settings:

- Choose to restore the entire database or only the tablespaces for the database.
- **Control files:** Optionally, enable this option to also restore the database control files.

3. Pre-restore options:

- Optionally, enable this option and enter the full path for a script that should be run before the restore operation and any arguments that the script takes.
- Choose a timeout value for the script. If the script fails to execute within this time period, the restore will proceed anyway.

4. Post-restore options:

- **Postscript:** Optionally, enable this option and enter the full path for a script that should be run after the restore operation and any arguments that the script takes.
- **Open the database or container database in READ-WRITE mode after recovery:** After the restore operation is complete, Backup and Recovery will enable READ-WRITE mode for the database.

5. Notification section:

- **Enable email notifications:** Select this to receive email notifications about the restore operation and indicate what type of notifications you want to receive.

6. Select **Restore**.

Restore to alternate location

Not available for Oracle Database workloads preview.

Mount and unmount Oracle database recovery points with NetApp Backup and Recovery

You might want to mount an Oracle Database recovery point if you need to access the database in a controlled state to perform recovery operations.

Mount an Oracle Database restore point

If you configure the protection policy for a database to retain archive logs, you can mount recovery points to view the database change history.

Steps

1. From the NetApp Console menu, select **Protection > Backup and Recovery**.
2. Select the Oracle tile.
3. In the Backup and Recovery menu, select **Inventory**.
4. For the Oracle Database workload in the list, select **View**.
5. Select the **Databases** menu.
6. Choose a database from the list and select the Actions icon  > **View protection details**.

A list of recovery points for that database appears.

7. Choose a recovery point from the list and select the Actions icon  > **Mount**.
8. In the dialog that appears, do the following:
 - a. Choose the host that should mount the recovery point from the list.
 - b. Select which location Backup and Recovery should use to mount the recovery point. For the preview release, mounting from the object store is not supported.

The mount path that Backup and Recovery should use is displayed.

9. Select **Mount**.

The recovery point is mounted on the Oracle host.

Unmount an Oracle database restore point

Unmount the recovery point when you no longer need to view changes made to that database.

Steps

1. From the NetApp Console menu, select **Protection > Backup and Recovery**.
2. Select the Oracle tile.
3. In the Backup and Recovery menu, select **Inventory**.
4. For the Oracle workload in the list, select **View**.
5. Select the **Databases** menu.
6. Choose a database from the list and select the Actions icon  > **View protection details**.

A list of recovery points for that database appears.

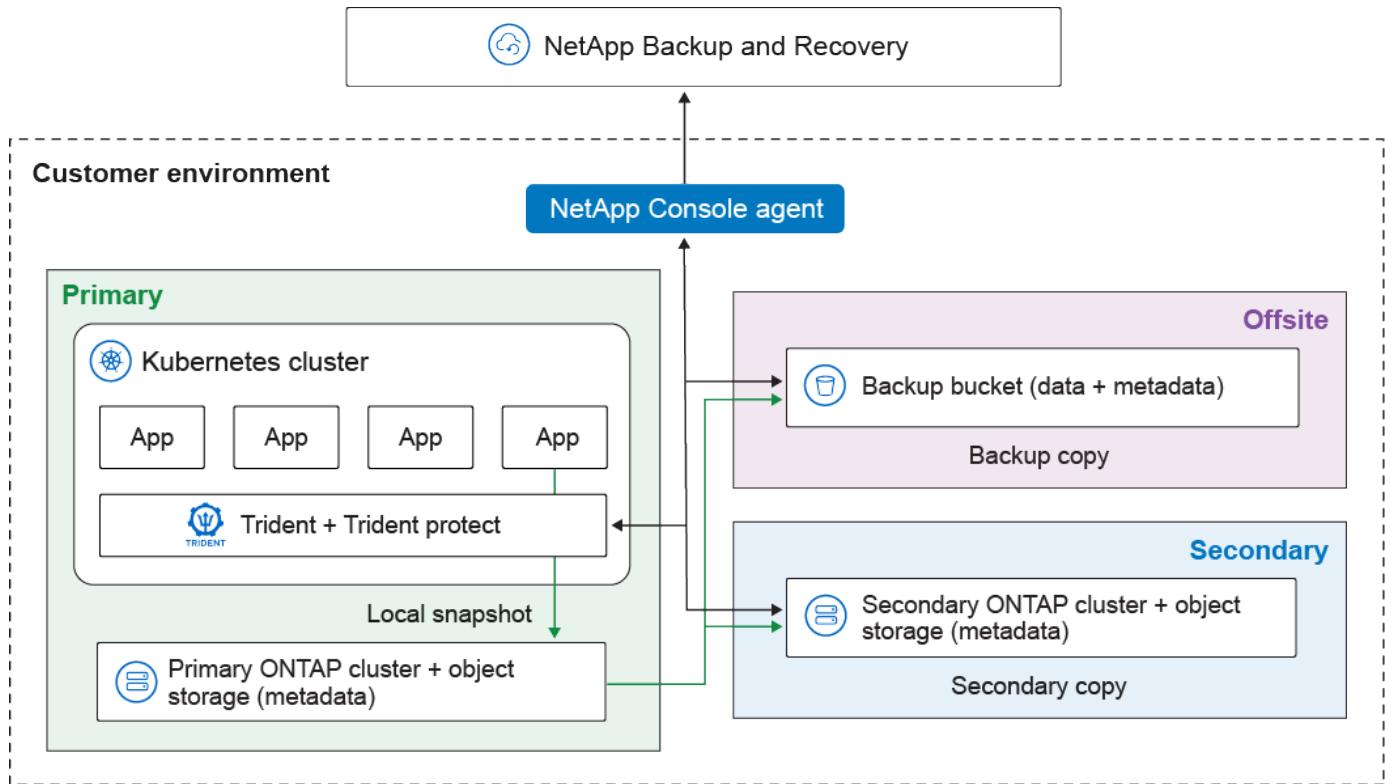
7. Choose a recovery point from the list and select the Actions icon  > **Unmount**.
8. Confirm the action by selecting **Unmount**.

Protect Kubernetes workloads (Preview)

Manage Kubernetes workloads overview

Managing Kubernetes workloads in NetApp Backup and Recovery enables you to discover, manage, and protect your Kubernetes clusters and applications all in one place. You can manage resources and applications hosted on your Kubernetes clusters. You can also create and associate protection policies with your Kubernetes workloads, all using a single interface.

The following diagram shows the components and basic architecture of backup and recovery for Kubernetes workloads and how different copies of your data can be stored in different locations:



NetApp Backup and Recovery provides the following benefits for managing Kubernetes workloads:

- A single control plane for protecting applications running across multiple Kubernetes clusters. These applications can include containers or virtual machines running on your Kubernetes clusters.
- Native integration with NetApp SnapMirror, enabling storage offloading capabilities for all backup and recovery workflows.
- Incremental forever backups for Kubernetes applications, translating to lower Recovery Point Objectives (RPOs) and Recovery Time Objectives (RTOs).



This documentation is provided as a technology preview. During the preview, Kubernetes functionality is not recommended for production workloads. With this preview offering, NetApp reserves the right to modify offering details, contents, and timeline before General Availability.

You can accomplish the following tasks related to managing Kubernetes workloads:

- [Discover Kubernetes workloads](#).
- [Manage Kubernetes clusters](#).
- [Add and protect Kubernetes applications](#).
- [Manage Kubernetes applications](#).
- [Restore Kubernetes applications](#).

Discover Kubernetes workloads in NetApp Backup and Recovery

NetApp Backup and Recovery needs to discover Kubernetes workloads before protecting them.

Required NetApp Console role

Backup and Recovery super admin. Learn about [Backup and recovery roles and privileges](#). Learn about [NetApp Console access roles for all services](#).

Discover Kubernetes workloads

In Backup and Recovery inventory, discover Kubernetes workloads in your environment. Adding a workload adds a Kubernetes cluster to NetApp Backup and Recovery. You can then add applications and protect cluster resources.



When you discover a cluster that is currently protected with Trident Protect, any backup schedules that were used with Trident Protect are disabled during the discovery process (Trident Protect backup schedules are not compatible with Backup and Recovery). To protect the cluster's applications, [create a new protection policy](#) or associate the applications with an existing policy. You can then remove the Trident Protect backup schedules if needed.

Steps

1. Do one of the following:
 - If you are discovering Kubernetes workloads for the first time, in NetApp Backup and Recovery, under **Workloads**, select the **Kubernetes** tile.
 - If you have already discovered Kubernetes workloads, in NetApp Backup and Recovery, select **Inventory > Workloads** and then select **Discover resources**.
2. Select the **Kubernetes** workload type.
3. Enter a cluster name and choose a connector to use with the cluster.
4. Follow the command line instructions that appear:
 - Create a Trident Protect namespace
 - Create a Kubernetes secret
 - Add a Helm repository
 - Install or upgrade Trident Protect and the Trident Protect connector

These steps ensure that NetApp Backup and Recovery can interact with the cluster.

5. After you complete the steps, select **Discover**.

The cluster is added to the inventory.

6. Select **View** in the associated Kubernetes workload to see the list of applications, clusters, and namespaces for that workload.

Continue to the NetApp Backup and Recovery Dashboard

Follow these steps to view the NetApp Backup and Recovery Dashboard.

1. From the NetApp Console menu, select **Protection > Backup and recovery**.
2. Select a workload tile (for example, Microsoft SQL Server).
3. From the Backup and Recovery menu, select **Dashboard**.
4. Review the health of data protection. The number of at risk or protected workloads increases based on the newly discovered, protected, and backed up workloads.

[Learn what the Dashboard shows you](#).

Add and protect Kubernetes applications

Add and protect Kubernetes applications

NetApp Backup and Recovery enables you to easily discover your Kubernetes clusters, without generating and uploading kubeconfig files. You can connect Kubernetes clusters and install the required software using simple commands copied from the NetApp Console user interface.

Required NetApp Console role

Organization admin or SnapCenter admin. [Learn about NetApp Backup and Recovery access roles](#). [Learn about NetApp Console access roles for all services](#).

Add and protect a new Kubernetes application

The first step in protecting Kubernetes applications is to create an application within NetApp Backup and Recovery. When you create an application, you make the Console aware of the running application on the Kubernetes cluster.

Before you begin

Before you can add and protect a Kubernetes application, you need to [discover Kubernetes workloads](#).

Add an application using the web UI

Steps

1. In NetApp Backup and Recovery, select **Inventory**.
2. Choose a Kubernetes instance, and select **View** to view the resources associated with that instance.
3. Select the **Applications** tab.
4. Select **Create application**.
5. Enter a name for the application.
6. Optionally, choose any of the following fields to search for the resources you want to protect:
 - Associated cluster
 - Associated namespaces
 - Resource types
 - Label selectors
7. Optionally, select **Cluster Scoped Resources** to choose any resources that are scoped at the cluster level. If you include them, they are added to the application when you create it.
8. Optionally, select **Search** to find the resources based on your search criteria.



The Console does not store the search parameters or results; the parameters are used to search the selected Kubernetes cluster for resources that can be included in the application.

9. The Console displays a list of resources that match your search criteria.
10. If the list contains the resources you want to protect, select **Next**.
11. Optionally, in the **Policy** area, choose an existing protection policy to protect the application or create a new policy. If you don't select a policy, the application is created without a protection policy. You can [add a protection policy](#) later.
12. In the **Prescripts and postscripts** area, enable and configure any prescript or postscript execution hooks that you want to run before or after backup operations. To enable prescripts or postscripts, you must have already created at least one [execution hook template](#).
13. Select **Create**.

Result

The application is created and appears in the list of applications in the **Applications** tab of the Kubernetes inventory. The NetApp Console enables protection for the application based on your settings, and you can monitor the progress in the **Monitoring** area of backup and recovery.

Add an application using a CR

Steps

1. Create the destination application CR file:
 - a. Create the custom resource (CR) file and name it (for example, `my-app-name.yaml`).
 - b. Configure the following attributes:
 - **metadata.name**: *(Required)* The name of the application custom resource. Note the name you choose because other CR files needed for protection operations refer to this value.
 - **spec.includedNamespaces**: *(Required)* Use namespace and label selector to specify the

namespaces and resources that the application uses. The application namespace must be part of this list. The label selector is optional and can be used to filter resources within each specified namespace.

- **spec.includedClusterScopedResources**: *(Optional)* Use this attribute to specify cluster-scoped resources to be included in the application definition. This attribute allows you to select these resources based on their group, version, kind, and labels.
 - **groupVersionKind**: *(Required)* Specifies the API group, version, and kind of the cluster-scoped resource.
 - **labelSelector**: *(Optional)* Filters the cluster-scoped resources based on their labels.

c. Configure the following annotations, if needed:

- **metadata.annotations.protect.trident.netapp.io/skip-vm-freeze**: *(Optional)* This annotation is only applicable to applications defined from virtual machines, such as in KubeVirt environments, where filesystem freezes occur before snapshots. Specify whether this application can write to the filesystem during a snapshot. If set to true, the application ignores the global setting and can write to the filesystem during a snapshot. If set to false, the application ignores the global setting and the filesystem is frozen during a snapshot. If specified but the application has no virtual machines in the application definition, the annotation is ignored. If not specified, the application follows the [global filesystem freeze setting](#).
- **protect.trident.netapp.io/protection-command**: *(Optional)* Use this annotation to instruct Backup and Recovery to protect or stop protecting the application. The possible values are protect or unprotect.
- **protect.trident.netapp.io/protection-policy-name**: *(Optional)* Use this annotation to specify the name of the Backup and Recovery protection policy that you want to use to protect this application. This protection policy must already exist in Backup and Recovery.

If you need to apply this annotation after an application has already been created, you can use the following command:



```
kubectl annotate application -n <application CR namespace> <application CR name> protect.trident.netapp.io/skip-vm-freeze="true"
```

Example YAML:

```

apiVersion: protect.trident.netapp.io/v1
kind: Application
metadata:
  annotations:
    protect.trident.netapp.io/skip-vm-freeze: "false"
    protect.trident.netapp.io/protection-command: "protect"
    protect.trident.netapp.io/protection-policy-name: "policy-name"
  name: my-app-name
  namespace: my-app-namespace
spec:
  includedNamespaces:
    - namespace: namespace-1
      labelSelector:
        matchLabels:
          app: example-app
    - namespace: namespace-2
      labelSelector:
        matchLabels:
          app: another-example-app
  includedClusterScopedResources:
    - groupVersionKind:
        group: rbac.authorization.k8s.io
        kind: ClusterRole
        version: v1
      labelSelector:
        matchLabels:
          mylabel: test

```

2. (Optional) Add filtering that includes or excludes resources marked with particular labels:

- **resourceFilter.resourceSelectionCriteria:** (Required for filtering) Use `Include` or `Exclude` to include or exclude a resource defined in `resourceMatchers`. Add the following `resourceMatchers` parameters to define the resources to be included or excluded:
 - **resourceFilter.resourceMatchers:** An array of `resourceMatcher` objects. If you define multiple elements in this array, they match as an OR operation, and the fields inside each element (`group`, `kind`, `version`) match as an AND operation.
 - **resourceMatchers[].group:** (Optional) Group of the resource to be filtered.
 - **resourceMatchers[].kind:** (Optional) Kind of the resource to be filtered.
 - **resourceMatchers[].version:** (Optional) Version of the resource to be filtered.
 - **resourceMatchers[].names:** (Optional) Names in the Kubernetes metadata.name field of the resource to be filtered.
 - **resourceMatchers[].namespaces:** (Optional) Namespaces in the Kubernetes metadata.name field of the resource to be filtered.
 - **resourceMatchers[].labelSelectors:** (Optional) Label selector string in the Kubernetes

metadata.name field of the resource as defined in the [Kubernetes documentation](#). For example: "trident.netapp.io/os=linux".



When both `resourceFilter` and `labelSelector` are used, `resourceFilter` runs first, and then `labelSelector` is applied to the resulting resources.

For example:

```
spec:
  resourceFilter:
    resourceSelectionCriteria: "Include"
    resourceMatchers:
      - group: my-resource-group-1
        kind: my-resource-kind-1
        version: my-resource-version-1
        names: ["my-resource-names"]
        namespaces: ["my-resource-namespaces"]
        labelSelectors: ["trident.netapp.io/os=linux"]
      - group: my-resource-group-2
        kind: my-resource-kind-2
        version: my-resource-version-2
        names: ["my-resource-names"]
        namespaces: ["my-resource-namespaces"]
        labelSelectors: ["trident.netapp.io/os=linux"]
```

3. After you create the application CR to match your environment, apply the CR. For example:

```
kubectl apply -f my-app-name.yaml
```

Back up Kubernetes applications now using the Backup and Recovery web UI

NetApp Backup and Recovery enables you to manually back up Kubernetes applications using the web interface.

Required NetApp Console role

Organization admin or SnapCenter admin. [Learn about NetApp Backup and Recovery access roles](#). [Learn about NetApp Console access roles for all services](#).

Back up a Kubernetes application now using the web UI

Manually create a backup of a Kubernetes application to establish a baseline for future backups and snapshots, or to ensure the most recent data is protected.

Steps

1. In NetApp Backup and Recovery, select **Inventory**.
2. Choose a Kubernetes instance, and select **View** to view the resources associated with that instance.
3. Select the **Applications** tab.
4. In the list of applications, choose an application you want to back up and select the associated Actions menu.
5. Select **Backup now**.
6. Ensure the correct application name is selected.
7. Select **Back up**.

Result

The Console creates a backup of the application and displays the progress in the **Monitoring** area of Backup and Recovery. The backup is created based on the protection policy associated with the application.

Back up Kubernetes applications now using custom resources in Backup and Recovery

NetApp Backup and Recovery enables you to manually back up Kubernetes applications using custom resources (CRs).

Back up a Kubernetes application now using custom resources

Manually create a backup of a Kubernetes application to establish a baseline for future backups and snapshots, or to ensure the most recent data is protected.



Cluster-scoped resources are included in a backup, snapshot, or clone if they are explicitly referenced in the application definition or if they have references to any of the application namespaces.

Unresolved directive in br-use-back-up-now-kubernetes-applications-cr.adoc - include:::_include/backup-include-sessiontoken-note.adoc[]

Create a local snapshot using a custom resource

To create a snapshot of your Kubernetes application and store it locally, use the Snapshot custom resource with specific attributes.

Steps

1. Create the custom resource (CR) file and name it `local-snapshot-cr.yaml`.
2. In the file you created, configure the following attributes:
 - **metadata.name:** *(Required)* The name of this custom resource; choose a unique and sensible name for your environment.
 - **spec.applicationRef:** The Kubernetes name of the application to snapshot.
 - **spec.appVaultRef:** *(Required)* The name of the AppVault where the snapshot contents (metadata) should be stored.
 - **spec.reclaimPolicy:** *(Optional)* Defines what happens to the AppArchive of a snapshot when the snapshot CR is deleted. This means that even when set to `Retain`, the snapshot will be deleted. Valid options:
 - `Retain` *(default)*

- Delete

```

apiVersion: protect.trident.netapp.io/v1
kind: Snapshot
metadata:
  namespace: my-app-namespace
  name: local-snapshot-cr
spec:
  applicationRef: my-application
  appVaultRef: appvault-name
  reclaimPolicy: Retain

```

3. After you populate the `local-snapshot-cr.yaml` file with the correct values, apply the CR:

```
kubectl apply -f local-snapshot-cr.yaml
```

Back up an application to an object store using a custom resource

Create a Backup CR with specific attributes to back up your application to an object store.

Steps

1. Create the custom resource (CR) file and name it `object-store-backup-cr.yaml`.
2. In the file you created, configure the following attributes:
 - **metadata.name:** *(Required)* The name of this custom resource; choose a unique and sensible name for your environment.
 - **spec.applicationRef:** *(Required)* The Kubernetes name of the application to back up.
 - **spec.appVaultRef:** *(Required, mutually exclusive with spec.appVaultTargetsRef)* If you use the same bucket to store the snapshot and backup, this is the name of the AppVault where the backup contents should be stored.
 - **spec.appVaultTargetsRef:** *(Required, mutually exclusive with spec.appVaultRef)* If you use different buckets to store the snapshot and backup, this is the name of the AppVault where the backup contents should be stored.
 - **spec.dataMover:** *(Optional)* A string indicating which backup tool to use for the backup operation. The value is case sensitive and must be CBS.
 - **spec.reclaimPolicy:** *(Optional)* Defines what happens to the backup contents (metadata/volume data) when the Backup CR is deleted. Possible values:
 - Delete
 - Retain (default)
 - **spec.cleanupSnapshot:** *(Required)* Ensures that the temporary snapshot created by the Backup CR is not deleted after the backup operation completes. Recommended value: `false`.

Example YAML when using the same bucket to store the snapshot and backup:

```

apiVersion: protect.trident.netapp.io/v1
kind: Backup
metadata:
  namespace: my-app-namespace
  name: my-cr-name
spec:
  applicationRef: my-application
  appVaultRef: appvault-name
  dataMover: CBS
  reclaimPolicy: Retain
  cleanupSnapshot: false

```

Example YAML when using different buckets to store the snapshot and backup:

```

apiVersion: protect.trident.netapp.io/v1
kind: Backup
metadata:
  namespace: my-app-namespace
  name: object-store-backup-cr
spec:
  applicationRef: my-application
  appVaultTargetsRef: appvault-targets-name
  dataMover: CBS
  reclaimPolicy: Retain
  cleanupSnapshot: false

```

3. After you populate the `object-store-backup-cr.yaml` file with the correct values, apply the CR:

```
kubectl apply -f object-store-backup-cr.yaml
```

Create a 3-2-1 fanout backup using a custom resource

Backing up using a 3-2-1 fanout architecture copies a backup to secondary storage as well as to an object store. To create a 3-2-1 fanout backup, create a Backup CR with specific attributes.

Steps

1. Create the custom resource (CR) file and name it `3-2-1-fanout-backup-cr.yaml`.
2. In the file you created, configure the following attributes:
 - **metadata.name:** *(Required)* The name of this custom resource; choose a unique and sensible name for your environment.
 - **spec.applicationRef:** *(Required)* The Kubernetes name of the application to back up.
 - **spec.appVaultTargetsRef:** *(Required)* The name of the AppVault where the backup contents should be stored.

- **spec.dataMover:** *(Optional)* A string indicating which backup tool to use for the backup operation. The value is case sensitive and must be CBS.
- **spec.reclaimPolicy:** *(Optional)* Defines what happens to the backup contents (metadata/volume data) when the Backup CR is deleted. Possible values:
 - Delete
 - Retain (default)
- **spec.cleanupSnapshot:** *(Required)* Ensures that the temporary snapshot created by the Backup CR is not deleted after the backup operation completes. Recommended value: `false`.
- **spec.replicateSnapshot:** *(Required)* Instructs Backup and Recovery to replicate the snapshot to secondary storage. Required value: `true`.
- **spec.replicateSnapshotReclaimPolicy:** *(Optional)* Defines what happens to the replicated snapshot when it is deleted. Possible values:
 - Delete
 - Retain (default)

Example YAML:

```
apiVersion: protect.trident.netapp.io/v1
kind: Backup
metadata:
  namespace: my-app-namespace
  name: 3-2-1-fanout-backup-cr
spec:
  applicationRef: my-application
  appVaultTargetsRef: appvault-targets-name
  dataMover: CBS
  reclaimPolicy: Retain
  cleanupSnapshot: false
  replicateSnapshot: true
  replicateSnapshotReclaimPolicy: Retain
```

3. After you populate the `3-2-1-fanout-backup-cr.yaml` file with the correct values, apply the CR:

```
kubectl apply -f 3-2-1-fanout-backup-cr.yaml
```

Supported backup annotations

The following table describes the annotations you can use when creating a backup CR.

Annotation	Type	Description	Default value
protect.trident.netapp.io/full-backup	string	Specifies whether a backup should be non-incremental. Set to <code>true</code> to create a non-incremental backup. It is best practice to perform a full backup periodically and then perform incremental backups in between full backups to minimize the risk associated with restores.	"false"
protect.trident.netapp.io/snapshot-completion-timeout	string	The maximum time allowed for the overall snapshot operation to complete.	"60m"
protect.trident.netapp.io/volume-snapshots-ready-to-use-timeout	string	The maximum time allowed for volume snapshots to reach the ready-to-use state.	"30m"
protect.trident.netapp.io/volume-snapshots-created-timeout	string	The maximum time allowed for volume snapshots to be created.	"5m"
protect.trident.netapp.io/pvc-bind-timeout-sec	string	Maximum time (in seconds) to wait for any newly created PersistentVolumeClaims (PVCs) to reach the <code>Bound</code> phase before the operations fails.	"1200" (20 minutes)

Restore Kubernetes applications

Restore Kubernetes applications using the web UI

NetApp Backup and Recovery enables you to restore applications that you have protected with a protection policy. To restore an application, an application needs to have at least one restore point available. A restore point consists of either the local snapshot or the backup to the object store (or both). You can restore an application using the local, secondary, or object store archive.

Before you begin

If you are restoring an application that was backed up using Trident Protect, ensure that Trident Protect is installed on both the source and destination clusters.

Required NetApp Console role

Organization admin or SnapCenter admin. [Learn about NetApp Backup and Recovery access roles](#). [Learn about NetApp Console access roles for all services](#).

Steps

1. In the NetApp Backup and Recovery menu, select **Restore**.
2. Choose a Kubernetes application from the list, and select **View and Restore** for that application.

The list of restore points appears.

3. Select the **Restore** button for the restore point you want to use.

General settings

1. Choose the source location to restore from.
2. Choose the destination cluster from the **Cluster** list.



Restoring a local snapshot created by Trident Protect to a different cluster is not supported at this time.

3. Choose to restore to the original namespaces or new namespaces.
4. If you chose to restore to new namespaces, enter the destination namespace or namespaces to use.
5. Select **Next**.

Resource selection

1. Choose whether you want to restore all resources associated with the application or use a filter to select specific resources to restore:

Restore all resources

- a. Select **Restore all resources**.
- b. Select **Next**.

Restore specific resources

- a. Select **Selective resources**.
- b. Choose the behavior of the resource filter. If you choose **Include**, the resources you select are restored. If you choose **Exclude**, the resources you select are not restored.
- c. Select **Add rules** to add rules that define filters for selecting resources. You need at least one rule to filter resources.

Each rule can filter on criteria such as the resource namespace, labels, group, version, and kind.

- d. Select **Save** to save each rule.
- e. When you have added all the rules you need, select **Search** to see the resources available in the backup archive that match your filter criteria.



The resources shown are the resources that currently exist on the cluster.

- f. When satisfied with the results, select **Next**.

Destination settings

1. Expand the **Destination settings** section and choose to restore either to the default storage class, a different storage class, or if you are restoring to a different cluster, to map the storage classes to the destination cluster.
2. If you chose to restore to a different storage class, select a destination storage class to match each source storage class.
3. Optionally, if you are restoring a backup or snapshot that was made using Trident Protect, view the details of the AppVault used as the storage bucket for the restore operation. If there is a change in your environment or the AppVault status, select **Sync App Vault** to refresh the details.



If you need to create an AppVault on a Kubernetes cluster to facilitate restoring a backup or snapshot created using Trident Protect, refer to [Use Trident Protect AppVault objects to manage buckets](#).

4. Optionally, expand the **Restore scripts** section and enable the **Postscript** option to choose an execution hook template that will run after the restore operation is complete. If needed, enter any arguments that the script needs and add label selectors to filter resources based on resource labels.
5. Select **Restore**.

Restore Kubernetes applications using a custom resource

You can use custom resources to restore your applications from a snapshot or backup. Restoring from an existing snapshot will be faster when restoring the application to the same cluster.

- When you restore an application, all execution hooks configured for the application are restored with the app. If a post-restore execution hook is present, it runs automatically as part of the restore operation.
- Restoring from a backup to a different namespace or to the original namespace is supported for qtree volumes. However, restoring from a snapshot to a different namespace or to the original namespace is not supported for qtree volumes.
- You can use advanced settings to customize restore operations. To learn more, refer to [Use advanced custom resource restore settings](#).

Restore a backup to a different namespace

When you restore a backup to a different namespace using a BackupRestore CR, Backup and Recovery restores the application in a new namespace and creates an application CR for the restored application. To protect the restored application, create on-demand backups or snapshots, or establish a protection schedule.

- Restoring a backup to a different namespace with existing resources will not alter any resources that share names with those in the backup. To restore all resources in the backup, either delete and re-create the target namespace, or restore the backup to a new namespace.
- When using a CR to restore to a new namespace, you must manually create the destination namespace before applying the CR. Backup and Recovery automatically creates namespaces only when using the CLI.

Unresolved directive in br-use-restore-kubernetes-applications-cr.adoc - include:::_include/restore-include-sessiontoken-note.adoc[]



When you restore backups using Kopia as the data mover, you can optionally specify annotations in the CR to control the behavior of the temporary storage used by Kopia. Refer to the [Kopia documentation](#) for more information about the options you can configure.

Steps

1. Create the custom resource (CR) file and name it `trident-protect-backup-restore-cr.yaml`.
2. In the file you created, configure the following attributes:

- **metadata.name:** (Required) The name of this custom resource; choose a unique and sensible name for your environment.
- **spec.appArchivePath:** The path inside AppVault where the backup contents are stored. You can use the following command to find this path:

```
kubectl get backups <BACKUP_NAME> -n my-app-namespace -o jsonpath
= '{.status.appArchivePath}'
```

- **spec.appVaultRef:** (Required) The name of the AppVault where the backup contents are stored.
- **spec.namespaceMapping:** The mapping of the source namespace of the restore operation to the destination namespace. Replace `my-source-namespace` and `my-destination-namespace` with information from your environment.

```
apiVersion: protect.trident.netapp.io/v1
kind: BackupRestore
metadata:
  name: my-cr-name
  namespace: my-destination-namespace
spec:
  appArchivePath: my-backup-path
  appVaultRef: appvault-name
  namespaceMapping: [{"source": "my-source-namespace", "destination": "my-destination-namespace"}]
```

Unresolved directive in br-use-restore-kubernetes-applications-cr.adoc - include:::_include/restore-include-selective-restore.adoc[]

3. After you populate the `trident-protect-backup-restore-cr.yaml` file with the correct values, apply the CR:

```
kubectl apply -f trident-protect-backup-restore-cr.yaml
```

Restore a backup to the original namespace

You can restore a backup to the original namespace at any time.

Unresolved directive in br-use-restore-kubernetes-applications-cr.adoc - include:::_include/restore-include-sessiontoken-note.adoc[]



When you restore backups using Kopia as the data mover, you can optionally specify annotations in the CR to control the behavior of the temporary storage used by Kopia. Refer to the [Kopia documentation](#) for more information about the options you can configure.

Steps

1. Create the custom resource (CR) file and name it `trident-protect-ipr-cr.yaml`.

2. In the file you created, configure the following attributes:

- **metadata.name:** (Required) The name of this custom resource; choose a unique and sensible name for your environment.
- **spec.appArchivePath:** The path inside AppVault where the backup contents are stored. You can use the following command to find this path:

```
kubectl get backups <BACKUP_NAME> -n my-app-namespace -o jsonpath
= '{.status.appArchivePath}'
```

- **spec.appVaultRef:** (Required) The name of the AppVault where the backup contents are stored.

For example:

```
apiVersion: protect.trident.netapp.io/v1
kind: BackupInplaceRestore
metadata:
  name: my-cr-name
  namespace: my-app-namespace
spec:
  appArchivePath: my-backup-path
  appVaultRef: appvault-name
```

Unresolved directive in br-use-restore-kubernetes-applications-cr.adoc - include:::_include/restore-include-selective-restore.adoc[]

3. After you populate the `trident-protect-backup-ipr-cr.yaml` file with the correct values, apply the CR:

```
kubectl apply -f trident-protect-backup-ipr-cr.yaml
```

Restore a backup to a different cluster

You can restore a backup to a different cluster if there is an issue with the original cluster.



- When you restore backups using Kopia as the data mover, you can optionally specify annotations in the CR to control the behavior of the temporary storage used by Kopia. Refer to the [Kopia documentation](#) for more information about the options you can configure.
- When using a CR to restore to a new namespace, you must manually create the destination namespace before applying the CR.

Before you begin

Ensure the following prerequisites are met:

- The destination cluster has Trident Protect installed.
- The destination cluster has access to the bucket path of the same AppVault as the source cluster, where

the backup is stored.

- Ensure that the AWS session token expiration is sufficient for any long-running restore operations. If the token expires during the restore operation, the operation can fail.
 - Refer to the [AWS API documentation](#) for more information about checking the current session token expiration.
 - Refer to the [AWS documentation](#) for more information about credentials with AWS resources.

Steps

1. Check the availability of the AppVault CR on the destination cluster using Trident Protect CLI plugin:

```
tridentctl-protect get appvault --context <destination_cluster_name>
```



Ensure that the namespace intended for the application restore exists on the destination cluster.

2. View the backup contents of the available AppVault from the destination cluster:

```
tridentctl-protect get appvaultcontent <appvault_name> \
--show-resources backup \
--show-paths \
--context <destination_cluster_name>
```

Running this command displays the available backups in the AppVault, including their originating clusters, corresponding application names, timestamps, and archive paths.

Example output:

CLUSTER	APP	TYPE	NAME	TIMESTAMP
PATH				
production1	wordpress	backup	wordpress-bkup-1	2024-10-30 08:37:40 (UTC)
			backuppather1	
production1	wordpress	backup	wordpress-bkup-2	2024-10-30 08:37:40 (UTC)
			backuppather2	

3. Restore the application to the destination cluster using the AppVault name and archive path:

4. Create the custom resource (CR) file and name it `trident-protect-backup-restore-cr.yaml`.

5. In the file you created, configure the following attributes:

- **metadata.name:** *(Required)* The name of this custom resource; choose a unique and sensible name for your environment.
- **spec.appVaultRef:** *(Required)* The name of the AppVault where the backup contents are stored.
- **spec.appArchivePath:** The path inside AppVault where the backup contents are stored. You can use the following command to find this path:

```
kubectl get backups <BACKUP_NAME> -n my-app-namespace -o jsonpath
= '{.status.appArchivePath}'
```



If BackupRestore CR is not available, you can use the command mentioned in step 2 to view the backup contents.

- **spec.namespaceMapping:** The mapping of the source namespace of the restore operation to the destination namespace. Replace `my-source-namespace` and `my-destination-namespace` with information from your environment.

For example:

```
apiVersion: protect.trident.netapp.io/v1
kind: BackupRestore
metadata:
  name: my-cr-name
  namespace: my-destination-namespace
spec:
  appVaultRef: appvault-name
  appArchivePath: my-backup-path
  namespaceMapping: [{"source": "my-source-namespace", "destination": "my-destination-namespace"}]
```

6. After you populate the `trident-protect-backup-restore-cr.yaml` file with the correct values, apply the CR:

```
kubectl apply -f trident-protect-backup-restore-cr.yaml
```

Restore a snapshot to a different namespace

You can restore data from a snapshot using a custom resource (CR) file either to a different namespace or the original source namespace. When you restore a snapshot to a different namespace using a SnapshotRestore CR, Backup and Recovery restores the application in a new namespace and creates an application CR for the restored application. To protect the restored application, create on-demand backups or snapshots, or establish a protection schedule.

- SnapshotRestore supports the `spec.storageClassMapping` attribute, but only when the source and destination storage classes use the same storage backend. If you attempt to restore to a StorageClass that uses a different storage backend, the restore operation will fail.
- When using a CR to restore to a new namespace, you must manually create the destination namespace before applying the CR.

Unresolved directive in br-use-restore-kubernetes-applications-cr.adoc - include:::_include/restore-include-sessiontoken-note.adoc[]

Steps

1. Create the custom resource (CR) file and name it `trident-protect-snapshot-restore-cr.yaml`.
2. In the file you created, configure the following attributes:
 - **metadata.name:** *(Required)* The name of this custom resource; choose a unique and sensible name for your environment.
 - **spec.appVaultRef:** *(Required)* The name of the AppVault where the snapshot contents are stored.
 - **spec.appArchivePath:** The path inside AppVault where the snapshot contents are stored. You can use the following command to find this path:

```
kubectl get snapshots <SNAPSHOT_NAME> -n my-app-namespace -o jsonpath
= '{.status.appArchivePath}'
```

- **spec.namespaceMapping:** The mapping of the source namespace of the restore operation to the destination namespace. Replace `my-source-namespace` and `my-destination-namespace` with information from your environment.

```
apiVersion: protect.trident.netapp.io/v1
kind: SnapshotRestore
metadata:
  name: my-cr-name
  namespace: my-app-namespace
spec:
  appVaultRef: appvault-name
  appArchivePath: my-snapshot-path
  namespaceMapping: [{"source": "my-source-namespace", "destination": "my-destination-namespace"}]
```

Unresolved directive in br-use-restore-kubernetes-applications-cr.adoc - include:::_include/restore-include-selective-restore.adoc[]

3. After you populate the `trident-protect-snapshot-restore-cr.yaml` file with the correct values, apply the CR:

```
kubectl apply -f trident-protect-snapshot-restore-cr.yaml
```

Restore a snapshot to the original namespace

You can restore a snapshot to the original namespace at any time.

Unresolved directive in br-use-restore-kubernetes-applications-cr.adoc - include:::_include/restore-include-sessiontoken-note.adoc[]

Steps

1. Create the custom resource (CR) file and name it `trident-protect-snapshot-ipr-cr.yaml`.
2. In the file you created, configure the following attributes:
 - **metadata.name:** *(Required)* The name of this custom resource; choose a unique and sensible name for your environment.
 - **spec.appVaultRef:** *(Required)* The name of the AppVault where the snapshot contents are stored.
 - **spec.appArchivePath:** The path inside AppVault where the snapshot contents are stored. You can use the following command to find this path:

```
kubectl get snapshots <SNAPSHOT_NAME> -n my-app-namespace -o
jsonpath='{.status.appArchivePath}'
```

```
apiVersion: protect.trident.netapp.io/v1
kind: SnapshotInplaceRestore
metadata:
  name: my-cr-name
  namespace: my-app-namespace
spec:
  appVaultRef: appvault-name
  appArchivePath: my-snapshot-path
```

Unresolved directive in br-use-restore-kubernetes-applications-cr.adoc - include:::_include/restore-include-selective-restore.adoc[]

3. After you populate the `trident-protect-snapshot-ipr-cr.yaml` file with the correct values, apply the CR:

```
kubectl apply -f trident-protect-snapshot-ipr-cr.yaml
```

Use advanced custom resource restore settings

You can customize restore operations using advanced settings such as annotations, namespace settings, and storage options to meet your specific requirements.

Unresolved directive in br-use-kubernetes-advanced-restore-settings.adoc - include:::_include/namespace-anno-labels.adoc[]

Supported fields

This section describes additional fields available for restore operations.

Storage class mapping

The `spec.storageClassMapping` attribute defines a mapping from a storage class present in the source application to a new storage class on the target cluster. You can use this when migrating applications between clusters with different storage classes or when changing the storage backend for `BackupRestore` operations.

Example:

```
storageClassMapping:
  - destination: "destinationStorageClass1"
    source: "sourceStorageClass1"
  - destination: "destinationStorageClass2"
    source: "sourceStorageClass2"
```

Supported annotations

This section lists the supported annotations for configuring various behaviors in the system. If an annotation is not explicitly set by the user, the system will use the default value.

Annotation	Type	Description	Default value
protect.trident.netapp.io/data-mover-timeout-sec	string	The maximum time (in seconds) allowed for data mover operation to be stalled.	"300"
protect.trident.netapp.io/kopia-content-cache-size-limit-mb	string	The maximum size limit (in megabytes) for the Kopia content cache.	"1000"
protect.trident.netapp.io/pvc-bind-timeout-sec	string	Maximum time (in seconds) to wait for any newly created PersistentVolumeClaims (PVCs) to reach the Bound phase before the operations fails. Applies to all restore CR types (<code>BackupRestore</code> , <code>BackupInplaceRestore</code> , <code>SnapshotRestore</code> , <code>SnapshotInplaceRestore</code>). Use a higher value if your storage backend or cluster often requires more time.	"1200" (20 minutes)

Manage Kubernetes clusters

NetApp Backup and Recovery enables you to discover and manage your Kubernetes clusters so that you can protect resources hosted by the clusters.

Required NetApp Console role

Organization admin or SnapCenter admin. [Learn about NetApp Backup and Recovery access roles](#). [Learn about NetApp Console access roles for all services](#).



To discover Kubernetes clusters, refer to [Discover Kubernetes workloads](#).

Edit Kubernetes cluster information

You can edit a cluster if you need to change its name.

Steps

1. In NetApp Backup and Recovery, select **Inventory > Clusters**.
2. In the list of clusters, choose a cluster you want to edit and select the associated Actions menu.
3. Select **Edit cluster**.
4. Make any required changes to the cluster name. The cluster name needs to match the name that you used with the Helm command during the discovery process.
5. Select **Done**.

Remove a Kubernetes cluster

To stop protecting a Kubernetes cluster, disable protection and delete associated applications, then remove the cluster from NetApp Backup and Recovery. NetApp Backup and Recovery does not delete the cluster or its resources; it only removes the cluster from the NetApp Console inventory.

Steps

1. In NetApp Backup and Recovery, select **Inventory > Clusters**.
2. In the list of clusters, choose a cluster you want to edit and select the associated Actions menu.
3. Select **Remove cluster**.
4. Review the information in the confirmation dialog box, and select **Remove**.

Manage Kubernetes applications

NetApp Backup and Recovery enables you to unprotect and delete your Kubernetes applications and associated resources.

Required NetApp Console role

Organization admin or SnapCenter admin. [Learn about NetApp Backup and Recovery access roles](#). [Learn about NetApp Console access roles for all services](#).

Unprotect a Kubernetes application

You can unprotect an application if you no longer want to protect it. When you unprotect an application, NetApp Backup and Recovery stops protecting the application but keeps all associated backups and snapshots.



You cannot unprotect an application while protection operations are still running for it. Either wait for the operation to finish, or as a workaround, [remove the restore point](#) the running protection operation is using. You can then unprotect the application.

Steps

1. In NetApp Backup and Recovery, select **Inventory**.
2. Choose a Kubernetes instance, and select **View** to view the resources associated with that instance.

3. Select the **Applications** tab.
4. In the list of applications, choose an application you want to unprotect and select the associated Actions menu.
5. Select **Unprotect**.
6. Read the notice, and when ready, select **Unprotect**.

Delete a Kubernetes application

Delete an application you no longer need. NetApp Backup and Recovery stops protection and removes all backups and snapshots for deleted applications.

Steps

1. In NetApp Backup and Recovery, select **Inventory**.
2. Choose a Kubernetes instance, and select **View** to view the resources associated with that instance.
3. Select the **Applications** tab.
4. In the list of applications, choose an application you want to delete and select the associated Actions menu.
5. Select **Delete**.
6. Enable **Delete snapshots and backups** to remove all snapshots and backups of the application.



You will no longer be able to restore the application using these snapshots and backups.

7. Confirm the action and select **Delete**.

Remove a restore point for a Kubernetes application

You might need to remove a restore point for an application if you need to unprotect it and protection operations are currently running.

Steps

1. In the NetApp Backup and Recovery menu, select **Restore**.
2. Choose a Kubernetes application from the list, and select **View and Restore** for that application.

The list of restore points appears.

3. Choose the recovery point you need to delete and select the Actions icon **...** > **Delete recovery point** to delete it.

Manage NetApp Backup and Recovery execution hook templates for Kubernetes workloads

An execution hook is a custom action that runs with a data protection operation in a managed Kubernetes application. For example, create application-consistent snapshots by using an execution hook to pause database transactions before a snapshot and resume them after. When you create an execution hook template, specify the hook type, the script to run, and filters for target containers. Use the template to link execution hooks to your applications.

NetApp Backup and Recovery freezes and unfreezes filesystems for applications like KubeVirt during data protection. You can disable this behavior globally or for specific applications using the Trident Protect documentation:

- To disable this behavior for all applications, refer to [Protecting data with KubeVirt VMs](#).
- To disable this behavior for a specific application, refer to [Define an application](#).

Required NetApp Console role

Organization admin or SnapCenter admin. [Learn about NetApp Backup and Recovery access roles](#). [Learn about NetApp Console access roles for all services](#).

Types of execution hooks

NetApp Backup and Recovery supports the following types of execution hooks, based on when they can be run:

- Pre-snapshot
- Post-snapshot
- Pre-backup
- Post-backup
- Post-restore

Order of execution

When a data protection operation is run, execution hook events take place in the following order:

1. Any applicable custom pre-operation execution hooks are run on the appropriate containers. You can create multiple custom pre-operation hooks, but their execution order is not guaranteed or configurable.
2. Filesystem freezes occur, if applicable.
3. The data protection operation is performed.
4. Frozen filesystems are unfrozen, if applicable.
5. NetApp Backup and Recovery runs any applicable custom post-operation execution hooks on the appropriate containers. You can create multiple custom post-operation hooks, but their execution order is not guaranteed or configurable.

If you create multiple hooks of the same type, their execution order is not guaranteed. Hooks of different types always run in the specified order. For example, the following is the order of execution of a configuration that has all of the different types of hooks:

1. Pre-snapshot hooks executed
2. Post-snapshot hooks executed
3. Pre-backup hooks executed
4. Post-backup hooks executed

 Test execution hook scripts before enabling them in production. Use 'kubectl exec' to test scripts, then verify snapshots and backups by cloning the app to a temporary namespace and restoring.



If a pre-snapshot execution hook adds, changes, or removes Kubernetes resources, those changes are included in the snapshot or backup and in any subsequent restore operation.

Important notes about custom execution hooks

Consider the following when planning execution hooks for your apps.

- An execution hook must use a script to perform actions. Many execution hooks can reference the same script.
- Execution hooks need to be written in the format of executable shell scripts.
- Script size is limited to 96KB.
- Execution hook settings and any matching criteria are used to determine which hooks are applicable to a snapshot, backup, or restore operation.



Execution hooks can reduce or disable application functionality. Make your custom hooks run as quickly as possible. If you start a backup or snapshot operation with associated execution hooks but then cancel it, the hooks are still allowed to run if the backup or snapshot operation has already begun. This means that the logic used in a post-backup execution hook cannot assume that the backup was completed.

Execution hook filters

When you add or edit an execution hook for an application, you can add filters to the execution hook to manage which containers the hook will match. Filters are useful for applications that use the same container image on all containers, but might use each image for a different purpose (such as Elasticsearch). Filters allow you to create scenarios where execution hooks run on some but not necessarily all identical containers. If you create multiple filters for a single execution hook, they are combined with a logical AND operator. You can have up to 10 active filters per execution hook.

Each filter you add to an execution hook uses a regular expression to match containers in your cluster. When a hook matches a container, the hook will run its associated script on that container. Regular expressions for filters use the Regular Expression 2 (RE2) syntax, which does not support creating a filter that excludes containers from the list of matches. For information on the syntax that NetApp Backup and Recovery supports for regular expressions in execution hook filters, see [Regular Expression 2 \(RE2\) syntax support](#).



If you add a namespace filter to an execution hook that runs after a restore or clone operation and the restore or clone source and destination are in different namespaces, the namespace filter is only applied to the destination namespace.

Execution hook examples

Visit the [NetApp Verda GitHub project](#) to download real execution hooks for popular apps such as Apache Cassandra and Elasticsearch. You can also see examples and get ideas for structuring your own custom execution hooks.

Create an execution hook template

You can create a custom execution hook template that you can use to perform actions before or after a data protection operation on an application.



Templates that you create here are only usable when protecting Kubernetes workloads.

Steps

1. In the Console, go to **Protection > Backup and recovery**.
2. Select the **Settings** tab.
3. Expand the **Execution hook template** section.
4. Select **Create execution hook template**.
5. Enter a name for the execution hook.
6. Optionally, choose a type of hook. For example, a post-restore hook is run after the restore operation is complete.
7. In the **Script** text box, enter the executable shell script that you want to run as part of the execution hook template. Optionally, you can select **Upload script** to upload a script file instead.
8. Select **Create**.

After you create the template, it appears in the list of templates in the **Execution hook template** section.

Monitor jobs in NetApp Backup and Recovery

With NetApp Backup and Recovery, monitor local snapshots, replications, and backup jobs you start. Track restore jobs you initiate. View jobs that are complete, in progress, or failed to help diagnose problems. Enable email notifications in the NetApp Console Notification Center to stay informed about system activity when not logged in. Use the Console Timeline to see details of all actions started from the UI or API.

NetApp Backup and Recovery keeps job information for 15 days, then deletes the job information and removes it from the Job Monitor.

Required NetApp Console role

Storage viewer, Backup and Recovery super admin, Backup and Recovery backup admin, Backup and Recovery restore admin, Backup and Recovery clone admin, or Backup and Recovery viewer role. Learn about [Backup and recovery roles and privileges](#). Learn about [NetApp Console access roles for all services](#).

View job status on the Job Monitor

You can view a list of all the snapshot, replication, backup to object storage, and restore operations and their current status in the **Job Monitoring** tab. This includes operations from your Cloud Volumes ONTAP, on-premises ONTAP, applications, and virtual machines. Each operation, or job, has a unique ID and a status.

The status can be:

- Success
- In Progress
- Queued
- Warning
- Failed

Snapshots, replications, backups to object storage, and restore operations that you initiated from the NetApp Backup and Recovery UI and API are available in the Job Monitoring tab.



If you have upgraded your ONTAP systems to 9.13.x and do not see ongoing scheduled backup operations in the Job Monitor, restart NetApp Backup and Recovery. [Learn how to restart NetApp Backup and Recovery](#).

Steps

1. From the NetApp Backup and Recovery menu, select **Monitoring**.
2. To show additional columns (System, SVM, User Name, Workload, Policy Name, Snapshot Label), select the plus sign.

Search and filter the list of jobs

You can filter the operations on the Job Monitoring page using several filters, such as policy, snapshot label, type of operation (protection, restore, retention, or other) and protection type (local snapshot, replication, or backup to the cloud).

By default, the Job Monitoring page shows protection and recovery jobs from the last 24 hours. You can change the timeframe using the Timeframe filter.

Steps

1. From the NetApp Backup and Recovery menu, select **Monitoring**.
2. To sort the results differently, select each column heading to sort by Status, Start Time, Resource Name, and more.
3. If you're looking for specific jobs, select the **Advanced Search & Filtering** area to open the Search panel.

Use this panel to enter a free text search for any resource; for example "volume 1" or "application 3". You can also filter the jobs list according to the items in the drop-down menus.

Most of the filters are self-explanatory. The filter for "Workload" enables you to view jobs in the following categories:

- ONTAP volumes (Cloud Volumes ONTAP and on-premises ONTAP volumes)
- Microsoft SQL Server
- Virtual Machines
- Kubernetes
 - You can search for data within a specific "SVM" only if you have first selected a System.
 - You can search using the "Protection type" filter only when you have selected the "Type" of "Protection".

4.

To update the page immediately, select the  button. Otherwise, this page refreshes every 15 minutes so that you'll always see the most recent job status results.



View job details

You can view details corresponding to a specific completed job. You can export details for a particular job in a JSON format.

You can view details such as job type (scheduled or on-demand), SnapMirror backup type (initial or periodic) start and end times, duration, amount of transferred data from system to object storage, average transfer rate,

policy name, retention lock enabled, ransomware scan performed, protection source details, and protection target details.

Restore jobs show details such as backup target provider (Amazon Web Services, Microsoft Azure, Google Cloud, on-premises), S3 bucket name, SVM name, source volume name, destination volume, snapshot label, recovered objects count, file names, file sizes, last modification date, and full file path.

Steps

1. From the NetApp Backup and Recovery menu, select **Monitoring**.
2. Select the name of the job.
3. Select the Actions menu **...** and select **View Details**.
4. Expand each section to see details.

Download Job Monitoring results as a report

You can download the contents of the main Job Monitoring page as a report after you filter or sort the results. NetApp Backup and Recovery generates and downloads a .CSV file that you can review and send to other groups as needed. The .CSV file includes up to 10,000 rows of data.

From the Job Monitoring Details information, you can download a JSON file containing details for a single job.

Steps

1. From the NetApp Backup and Recovery menu, select **Monitoring**.
2. To download a CSV file for all jobs, select the Download button and locate the file in your download directory.
3. To download a JSON file for a single job, select the Actions menu **...** for the job, select **Download JSON File**, and locate the file in your download directory.

Review retention (backup lifecycle) jobs

Monitor retention (*backup lifecycle*) flows to check backups, keep them safe, and support audits. Identify when backup copies expire to track the lifecycle.

A backup lifecycle job tracks all snapshots that are deleted or in the queue to be deleted. Beginning with ONTAP 9.13, you can view all job types called "Retention" on the Job Monitoring page.

The "Retention" job type captures all snapshot deletion jobs initiated on a volume that is protected by NetApp Backup and Recovery.

Steps

1. From the NetApp Backup and Recovery menu, select **Monitoring**.
2. Select the **Advanced Search & Filtering** area to open the Search panel.
3. Select "Retention" as the job type.

Review backup and restore alerts in the NetApp Console Notification Center

The NetApp Console Notification Center tracks the progress of backup and restore jobs that you've initiated so you can verify whether the operation was successful or not.

You can view alerts in the Notification Center and configure the Console to send email alerts for important system activity, even when you are not logged in. [Learn more about the Notification Center and how to send](#)

alert emails for backup and restore jobs.

The Notification Center displays numerous snapshot, replication, backup to cloud, and restore events, but only certain events trigger email alerts:

Operation type	Event	Alert generated	Email sent
Activation	Backup and Recovery activation failed for system	Yes	Yes
Activation	Backup and Recovery edit failed for system	Yes	Yes
Activation	Volume now associated with snapshot policy	Yes	Yes
Activation	Volume backup or state modified	Yes	Yes
Activation	Backup and Recovery activation successful for system	Yes	Yes
Activation	Ad-hoc volume backup failed	Yes	Yes
Activation	Ad-hoc volume backup successful	Yes	No
Activation	Multi-volume backup failed	Yes	Yes
Cron operations	Checking for missing snapshot labels	Yes	Yes
Cron operations	Failed to send security token to ONTAP for this system	Yes	Yes
Pub/Sub events	Connection failure	Yes	No
Pub/Sub events	Failed to delete a scheduled snapshot	Yes	No
Pub/Sub events	Scheduled backup of volume failed	Yes	No
Pub/Sub events	Restore of volume succeeded	Yes	No
Pub/Sub events	Restore of volume failed	Yes	No
Ransomware	Potential Ransomware attack identified on backup copy	Yes	Yes
Ransomware	Potential Ransomware attack identified on backup copy for this system	Yes	Yes
Local snapshot	NetApp Backup and Recovery ad-hoc snapshot creation job failure	Yes	Yes
Replication	Modification of replication relationship of volume failure	Yes	Yes
Replication	NetApp Backup and Recovery ad-hoc replication job failure	Yes	Yes
Replication	NetApp Backup and Recovery replication pause job failure	Yes	No
Replication	NetApp Backup and Recovery replication break job failure	Yes	No

Operation type	Event	Alert generated	Email sent
Replication	NetApp Backup and Recovery replication resync job failure	Yes	No
Replication	NetApp Backup and Recovery replication stop job failure	Yes	No
Replication	NetApp Backup and Recovery replication reverse resync job failure	Yes	Yes
Replication	NetApp Backup and Recovery replication delete job failure	Yes	Yes
Target operations	Restore to local or cloud destination failure	Yes	Yes
Target operations	On-demand restore failure	Yes	Yes
System operations	Ad-hoc volume snapshot creation failure	Yes	Yes



Beginning with ONTAP 9.13.0, all alerts appear for Cloud Volumes ONTAP and on-premises ONTAP systems. For systems with Cloud Volumes ONTAP 9.13.0 and on-premises ONTAP, only the alert related to "Restore job completed, but with warnings" appears.

By default, NetApp Console organization and account admins receive emails for all "Critical" and "Recommendation" alerts. By default, the system does not set up other users and recipients to receive notification emails. Configure email alerts for any Console users in your NetApp Cloud Account or to other recipients who need to know about backup and restore activity.

To receive the NetApp Backup and Recovery email alerts, you'll need to select the notification severity types "Critical", "Warning", and "Error" in the Notifications settings page.

[Learn how to send alert emails for backup and restore jobs.](#)

Steps

1. From the Console menu, select the .
2. Review the notifications.

Review operation activity in Console Timeline

You can view details of backup and restore operations for further investigation in the Console Timeline. The Console Timeline provides details of each event, whether user-initiated or system-initiated and shows actions initiated in the UI or via the API.

[Learn about the differences between the Timeline and the Notification Center.](#)

Restart NetApp Backup and Recovery

There may be situations where you'll need to restart NetApp Backup and Recovery.

The Console agent includes NetApp Backup and Recovery functionality.

Steps

1. Connect to the Linux system that the Console agent is running on.

Console agent location	Procedure
Cloud deployment	Follow the instructions for connecting to the Console agent Linux virtual machine depending on the cloud provider you're using.
Manual installation	Log in to the Linux system.

2. Enter the command to restart the service.

Console agent location	Docker command	Podman command
Cloud deployment	<code>docker restart cloudmanager_cbs</code>	<code>podman restart cloudmanager_cbs</code>
Manual installation with internet access	<code>docker restart cloudmanager_cbs</code>	<code>podman restart cloudmanager_cbs</code>
Manual installation without internet access	<code>docker restart ds_cloudmanager_cbs_1</code>	<code>podman restart ds_cloudmanager_cbs_1</code>

Automate with NetApp Backup and Recovery REST APIs

The NetApp Backup and Recovery capabilities that are available through the web UI are also available through the Backup and Recovery REST API.

There are ten categories of endpoints defined within NetApp Backup and Recovery:

- `backup` - manages backup operations of cloud and on-premises resources, and retrieves details of the backup data
- `catalog` - manages the indexed catalog search for files based on a query (Search & Restore)
- `cloud` - retrieves information about various cloud provider resources from the NetApp Console
- `job` - manages job detail entries on the NetApp Console database
- `license` - retrieves the license validity of the systems from NetApp Console
- `ransomware scan` - initiates a ransomware scan on a specific backup file
- `restore` - enables you to perform volume, file, and folder-level restore operations
- `sfr` - retrieves files from a backup file for single file-level restore operations (Browse & Restore)
- `storagegrid` - retrieves details about a StorageGRID server, and enables you to discover a StorageGRID server
- `system` - manages the backup policies, and configures the destination object store associated with a system

API reference

Documentation for each NetApp Backup and Recovery API is available from [NetApp Console automation for NetApp Backup and Recovery](#).

Getting started

To get started with the NetApp Backup and Recovery APIs, you'll need to obtain a user token, your NetApp Console account ID, and the Console agent ID.

When making API calls, you'll add the user token in the Authorization header, and the Console agent ID in the `x-agent-id` header. You should use the NetApp Console account ID in the APIs.



If you are using a service account, you should use the service access token instead of a user token. The value for "client_id" ("Mu0V1ywgYtel6w1MbD15fKfVIUrNXGWC") is a fixed value and cannot be changed. In this case, follow the instructions here: [Create a service access token](#).

Steps

1. Obtain a user token from the NetApp NetApp Console web site.

Make sure you generate the refresh token from the following xref:./ <https://services.cloud.netapp.com/refresh-token/>. The refresh token is an alpha-numeric string that you'll use to generate a user token.

```
curl --location --request POST 'https://netapp-cloud-account.auth0.com/oauth/token?=' \
--header 'Content-Type: application/json' \
-d '{
  "grant_type": "refresh_token",
  "refresh_token": "JxaVHn9cGkX92aPVCKhat3zxxxxwsC9qM1_pLhkZtsVA",
  "client_id": "Mu0V1ywgYteI6w1MbD15fKfVIUrNXGWC"
}'
```



The user token from the NetApp Console web site has an expiration date. The API response includes an "expires_in" field that states when the token expires. To refresh the token, you'll need to call this API again.

2. Obtain your NetApp Console account ID.

```
GET 'https://api.bluexp.netapp.com/tenancy/account' -H 'authority:
api.bluexp.netapp.com'
Header:
-H 'accept: application/json'
-H 'accept-language: en-GB,en;q=0.9'
-H 'authorization: Bearer eyJhbGciOiJSUzI1NiIsInR
```

This API will return a response like the following. You can retrieve the account ID by parsing the output from **[0].[accountPublicId]**.

```
{
  "accountPublicId": "account-i6vJXvZW",
  "accountName": "rashidn",
  "isSaas": true,
  "isGov": false,
  "isPrivatePreviewEnabled": false,
  "is3rdPartyServicesEnabled": false,
  "accountSerial": "96064469711530003565",
  "userRole": "Role-1"
}
```

3. Obtain the x-agent-id which contains the Console agent ID.

```
GET 'https://api.services.cloud.netapp.com/occm/list-occms/account-  
OOnAR4ZS?excludeStandalone=true&source=saas' \  
Header:  
-H 'authority: api.services.cloud.netapp.com' \  
-H 'accept: application/json' \  
-H 'accept-language: en-GB,en;q=0.9' \  
-H 'authorization: Bearer eyJhbGciOiJSUzI1NiIsInR5.....'
```

You can retrieve the agent id from the response by parsing the output from **occm.[0].[agent].[agentId]**.

Example using the APIs

The following example shows an API call to activate NetApp Backup and Recovery on a system with a new policy that has daily, hourly, and weekly labels set, archive after days set to 180 days, in East-US-2 region in Azure cloud. Note that this only enables backup on the system, but no volumes are backed up.

API Request

You'll see that we use the NetApp Console account ID account-DpTFcxN3, Console agent ID iZwFFeVCZjWnzGlw8RgD0QQNANZvpP7Iclients, and user token Bearer eyJhbGciOiJSUzI1NiIsInR5cCI6IkpxVCIsImtpZCI6Ik5rSX1PVFUzUWpZek1E...y6nyhBjwkeMwHc4V alobjUmju2x0xUH48g in this command.

```

curl --location --request POST
'https://api.bluexp.netapp.com/account/account-
DpTFCxN3/providers/cloudmanager_cbs/api/v3/backup/working-
environment/VsaWorkingEnvironment-99hPYEgk' \
--header 'x-agent-id: iZwFFeVCZjWnzGlw8RgD0QQNANZvpP7Iclients' \
--header 'Accept: application/json' \
--header 'Content-Type: application/json' \
--header 'Authorization: Bearer
eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCIsImtpZCI6Ik5rSXlPVFUzUWpZek1E...y6nyhBjwk
eMwHc4ValobjUmju2x0xUH48g' \
--data-raw '{
  "provider": "AZURE",
  "backup-policy": {
    "archive-after-days": 180,
    "rule": [
      {
        "label": "hourly",
        "retention": "2"
      },
      {
        "label": "daily",
        "retention": "30"
      },
      {
        "label": "weekly",
        "retention": "52"
      }
    ]
  },
  "ip-space": "Default",
  "region": "eastus2",
  "azure": {
    "resource-group": "rn-test-backup-rg",
    "subscription": "3beb4dd0-25d4-464f-9bb0-303d7cf5c0c2"
  }
}

```

Response is a job ID that you can then monitor:

```
{
  "job-id": "1b34b6f6-8f43-40fb-9a52-485b0dfe893a"
}
```

Monitor the response:

```
curl --location --request GET
'https://api.bluexp.netapp.com/account/account-
DpTFcxN3/providers/cloudmanager_cbs/api/v1/job/1b34b6f6-8f43-40fb-9a52-
485b0dfe893a' \
--header 'x-agent-id: iZwFFeVCZjWnzG1w8RgD0QQNANZvpP7Iclients' \
--header 'Accept: application/json' \
--header 'Content-Type: application/json' \
--header 'Authorization: Bearer
eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCIsImtpZCI6Ik5rSX1PVFUzUWpZek1E...hE9ss2Nub
K6wZRHUdSaORI7JvcOorUhJ8srqdiUiW6MvuGIFAQIh668of2M3dLbhVDBe8BBMtsa939UGnJx
7Qz6Eg'
```

Response:

```
{
  "job": [
    {
      "id": "1b34b6f6-8f43-40fb-9a52-485b0dfe893a",
      "type": "backup-working-environment",
      "status": "PENDING",
      "error": "",
      "time": 1651852160000
    }
  ]
}
```

Monitor until "status" is "COMPLETED":

```
{
  "job": [
    {
      "id": "1b34b6f6-8f43-40fb-9a52-485b0dfe893a",
      "type": "backup-working-environment",
      "status": "COMPLETED",
      "error": "",
      "time": 1651852160000
    }
  ]
}
```

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](#).

Set up the StorageGRID environment details

Add the details of the StorageGRID system associated with your ONTAP systems as they were set up on the original Console agent setup using the [NetApp Console APIs](#).

The following information applies to private mode installations starting from NetApp Console 3.9.xx. For older versions, use the following procedure: [DarkSite Cloud Backup: MySQL and Indexed Catalog Backup and Restore](#).

You'll need to perform these steps for each system that is backing up data to StorageGRID.

1. Extract the authorization token using the following oauth/token API.

```
curl 'http://10.193.192.202/oauth/token' -X POST -H 'Accept: application/json' -H 'Accept-Language: en-US,en;q=0.5' -H 'Accept-Encoding: gzip, deflate' -H 'Content-Type: application/json' -d '{"username":"admin@netapp.com", "password":"Netapp@123", "grant_type":"password"}' > '
```

While the IP address, username, and passwords are custom values, the account name is not. The account name is always "account-DARKSITE1". Also, the username must use an email-formatted name.

This API will return a response like the following. You can retrieve the authorization token as shown below.

```
{"expires_in":21600,"access_token":"eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCIsImtpZCI6IjJ1MGFiZjRiIn0eyJzdWIiOiJvY2NtYXV0aHwxIiwiYXVkIjpbImh0dHBzOi8vYXBpLmNsb3VkLm51dGFwcC5jb20iXSwiaHR0cDovL2Nsb3VkLm51dGFwcC5jb20vZnVsbF9uYW11IjoiYWRtaW4iLCJodHRwOi8vY2xvdWQubmV0YXBwLmNvbS91bWFpbCI6ImFkbWluQG51dGFwcC5jb20iLCJzY29wZSI6Im9wZW5pZCBwcm9maWx1IiwiWF0IjoxNjcyNzM2MDIzLCJleHAiOjE2NzI3NTc2MjMsImlzcyI6Imh0dHA6Ly9vY2NtYXV0aDo4NDIwLyJ9CJtRpRDY23PokYlg1if67bmgnMcYxdCvBOY-ZUYWzhrWbbY_hqUH4T-114v_pNDsPyNDyWqHaKizThdjjHYHxm56vTz_Vdn4NqjaBDPwN9KAnC6Z88WA1cJ4WRQqj5ykODNDmrv5At_f9HHp0-xVmYHqywZ4nNFa1MvAh4xESc5jfkoKOZc-IOQdWm4F4LHpMzs4qFzCYthTuSKLYtqSTUrZB81-o-ipvrOqSo1iwIeHXZJJV-USWun9daNgiYd_wX-4WWJViGENDzzwOKfuUoUoe1Fg3ch--7JFkF1-rrXDOjk1sUMumN3WHV9usp1PgBE5HAcJPrEBm0ValSzCUbA"}
```

2. Extract the system ID and the X-Agent-Id using the tenancy/external/resource API.

```
curl -X GET
http://10.193.192.202/tenancy/external/resource?account=account-
DARKSITE1 -H 'accept: application/json' -H 'authorization: Bearer
eyJhbGciOiJSUzI1NiIzInR5cCI6IkpXVCiSImtpZCI6IjJ1MGFizjRiIn0eyJzdWIiOjvY
2NtYXV0aHwxIiwiYXVkJpbImh0dHBzOi8vYXBpLmNsb3VkLm51dGFwcC5jb20ixSwiaHR0c
DovL2Nsb3VkLm51dGFwcC5jb20vZnVsbF9uYW11IjoiYWRtaW4iLCJodHRwOi8vY2xvdWQub
mV0YXBwLmNvbS91bWFpbCI6ImFkbWluQG51dGFwcC5jb20iLCJzY29wZSI6Im9wZW5pZCBwc
m9maWx1IiwiWF0IjoxNjcyNzIyNzEzLCJleHAiOjE2NzI3NDQzMMSImlzcyI6Imh0dHA6L
y9vY2NtYXV0aDo4NDIwLyJ9X_cQF8xttD0-S7sU2uph2cdw_kN-
fLWpdJJX98HODwPpVUiLcxV28_sQhuopjWobozPelNISf7KvMqcoXc5kLDyX-
yE0fH9gr4XgkdsjWcNvw2rRkFzjHpWrETgfqAMkZcAukV4DHuxogHWh6-
DggB1NgPZT8A_szHinud5W0HJ9c4AaT0zC-
sp81GaqMahPf0KcFVyjbBL4krOewgKHGfo_7ma_4mF39B1LCj7Vc2XvUd0wCaJvDMjwp19-
KbZqmmBX9vDnYp7SSxC1hHJRDStcFgJLdJHtowweNH2829KsjEGBTcBdO8SvIDtctNH_GAx
wSgMT3zUfwaOimPw'
```

This API will return a response like the following. The value under the "resourceIdentifier" denotes the *WorkingEnvironment Id* and the value under "agentId" denotes *x-agent-id*.

```
[{"resourceIdentifier": "OnPremWorkingEnvironment-
pMtZND0M", "resourceType": "ON_PREM", "agentId": "vB_1xShPpBtUosjD7wfB1LIhqD
gIPA0wclients", "resourceClass": "ON_PREM", "name": "CBSFAS8300-01-
02", "metadata": "{\"clusterUuid\": \"2cb6cb4b-dc07-11ec-9114-
d039ea931e09\"}", "workspaceIds": ["workspace2wKYjTy9"], "agentIds": ["vB_1x
ShPpBtUosjD7wfB1LIhqDgIPA0wclients"]}]
```

3. Update the NetApp Backup and Recovery database with the details of the StorageGRID system associated with the systems. Make sure to enter the Fully Qualified Domain Name of the StorageGRID, as well as the Access-Key and Storage-Key as shown below:

```

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
eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCIsImtpZCI6IjJ1MGFizjRiIn0eyJzdWIiOiJvY
2NtYXV0aHwxIiwiYXVkJpbImh0dHBzOi8vYXBpLmNsb3VkLm51dGFwcC5jb20iXswiaHR0c
DovL2Nsb3VkLm51dGFwcC5jb20vZnVsbF9uYW11ijojYWRtaW4iLCJodHRwOi8vY2xvdWQub
mV0YXBwLmNvbS91bWFpbCI6ImFkbWluQG51dGFwcC5jb20iLCJzY29wZSI6Im9wZW5pZCBwc
m9maWx1IiwigiaWF0IjoxNjcyNzIyNzEzLCJleHAiOjE2NzI3NDQzMMSImlzcyI6Imh0dHA6L
y9vY2NtYXV0aDo4NDIwLyJ9X_cQF8xttD0-S7sU2uph2cdw_kN-
fLWpdJX98HODwPpVuitLcxV28_sQhuopjWobozPelNISf7KvMqcoXc5kLDyX-
yE0fH9gr4XgkdsbjWcNvw2rRkFzjHpWrETgfqAMkZcAukV4DHuxogHWh6-
DggB1NgPZT8A_szHinud5W0HJ9c4AaT0zC-
sp81GaqMahPf0KcFVyjbBL4kr0ewgKHGfo_7ma_4mF39B1LCj7Vc2XvUd0wCaJvDMjwp19-
KbZqmmBX9vDnYp7SSxC1hHJRDStcFgJLdJHtowweNH2829KsjEGBTTcBdO8SvIDtctNH_GAx
wSgMT3zUfwaOimPw' \
> --header 'x-agent-id: vB_1xShPpBtUosjD7wfB1LIhqDgIPA0wclients' \
> -d '
> { "storage-server" : "sr630ip15.rtp.eng.netapp.com:10443", "access-
key": "2ZMYOAVAS5E70MCNH9", "secret-password":
"uk/6ikd4Lj1XQOFnzSzP/T0zR4ZQ1G0w1xgWsB" }'

```

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.

Legal notices

Legal notices provide access to copyright statements, trademarks, patents, and more.

Copyright

<https://www.netapp.com/company/legal/copyright/>

Trademarks

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<https://www.netapp.com/company/legal/trademarks/>

Patents

A current list of NetApp owned patents can be found at:

<https://www.netapp.com/pdf.html?item=/media/11887-patentspage.pdf>

Privacy policy

<https://www.netapp.com/company/legal/privacy-policy/>

Open source

Notice files provide information about third-party copyright and licenses used in NetApp software.

- [Notice for NetApp Console](#)
- [Notice for the NetApp Backup and Recovery](#)
- [Notice for Single File Restore](#)

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—with 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.