



Get started

NetApp Backup and Recovery

NetApp
March 09, 2026

Table of Contents

- Get started 1
 - Learn about NetApp Backup and Recovery 1
 - What you can do with NetApp Backup and Recovery 1
 - Benefits of using NetApp Backup and Recovery 2
 - Cost 2
 - Licensing 4
 - Supported workloads, systems, and backup targets 4
 - How NetApp Backup and Recovery works 5
 - Terms that might help you with NetApp Backup and Recovery 6
 - NetApp Backup and Recovery prerequisites 6
 - Prerequisite for ONTAP 9.8 and later 6
 - Prerequisites for backups to object storage 6
 - Requirements for protecting Microsoft SQL Server workloads 6
 - Requirements for protecting VMware workloads 7
 - Requirements for protecting KVM workloads 8
 - Requirements for protecting Oracle Database workloads 9
 - Requirements for protecting Kubernetes applications 9
 - Requirements for protecting Hyper-V workloads 10
 - In NetApp Console 10
 - Set up licensing for NetApp Backup and Recovery 11
 - 30-day free trial 12
 - Use a NetApp Backup and Recovery PAYGO subscription 12
 - Use an annual contract 13
 - Use a NetApp Backup and Recovery BYOL license 14
 - Exceeding license capacity 14
 - Set up security certificates for StorageGRID and ONTAP in NetApp Backup and Recovery 15
 - Create a security certificate for StorageGRID 15
 - Create a security certificate for ONTAP 19
 - Create a certificate for both ONTAP and StorageGRID 22
 - Set up backup destinations before you use NetApp Backup and Recovery 22
 - Prepare the backup destination 23
 - Set up S3 permissions 23
 - Log in to NetApp Backup and Recovery 26
 - Discover offsite backup targets in NetApp Backup and Recovery 26
 - Discover a backup target 27
 - Add a bucket for a backup target 28
 - Change credentials for a backup target 29
 - Switch to different NetApp Backup and Recovery workloads 30
 - Switch to a different workload 30
 - Configure NetApp Backup and Recovery settings 30
 - Add credentials for host resources 30
 - Maintain VMware vCenter settings 32
 - Import and manage SnapCenter host resources 33

Add a KVM management platform	34
Configure log directories in snapshots for Windows hosts	35
Create an execution hook template	35
Set up role-based access control in NetApp Backup and Recovery	36
Related information	37

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.

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 (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.
- 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](#).

- **KVM workloads:**

- 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:**
 - 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:**
 - 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
- Hyper-V workloads
- Oracle Database workloads (Preview)
- Kubernetes workloads

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)

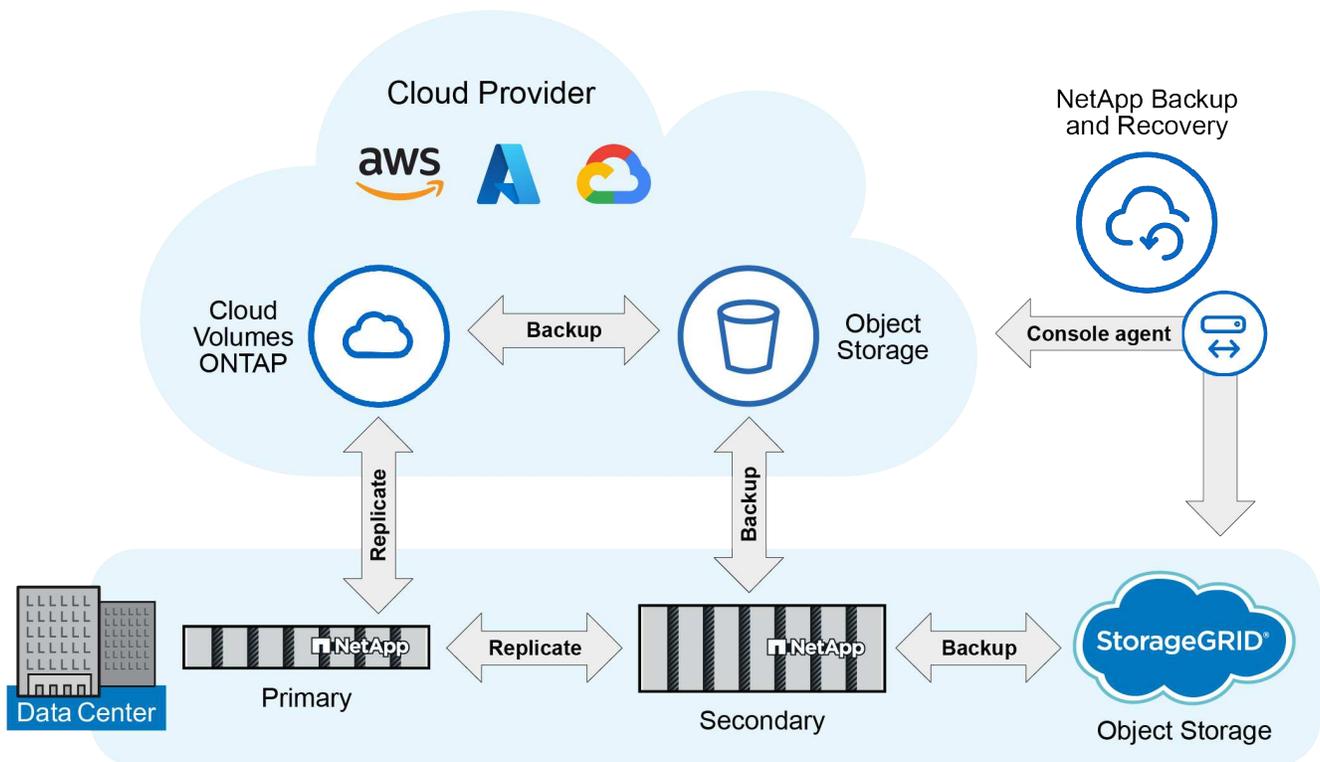
Supported backup targets

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

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	<p>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.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;">  <p>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.</p> </div>
Minimum RAM for the plug-in on the SQL Server host	1 GB
Minimum install and log space for the plug-in on the SQL Server host	5 GB 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.
Required software packages	<ul style="list-style-type: none"> • ASP.NET Core Runtime 8.0.23 Hosting Bundle (and all subsequent 8.0.x patches) • PowerShell 7.4.13 <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

- NFSv3, NFSv4.1, and VMFS datastores are supported.
- 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 9.13.1 or later
- Amazon FSx for NetApp ONTAP storage systems running version 9.13.1 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 8.0.23 Hosting Bundle (and all subsequent 8.0.x patches)
 - PowerShell 7.4.13 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-aifs` 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](#).

Exceeding license capacity

Exceeding your licensed capacity triggers PAYGO rates; without a PAYGO subscription, you cannot create new backups, though existing backups remain restorable without a service guarantee. Be sure to renew your license before it expires; an expired license prevents new backups and disrupts service.

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.](#)

Copyright information

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

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

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

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

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

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

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

Trademark information

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