



Backing up on-premises ONTAP data to the public cloud

Cloud Manager

Tom Onacki
June 07, 2021

Table of Contents

- Backing up on-premises ONTAP data to the public cloud 1
 - Quick start 1
 - Overview diagrams 3
 - Requirements 5
 - Enabling Cloud Backup 10

Backing up on-premises ONTAP data to the public cloud

Complete a few steps to get started backing up data from your on-premises ONTAP systems to low-cost object storage in the public cloud. This includes creating backup files on Amazon S3, Azure Blob, and Google Cloud Storage.

If you have a NetApp StorageGRID system and want to create your backups there, see how to [back up on-premises ONTAP data to the private cloud](#).

TIP

In most cases you'll use Cloud Manager for all backup and restore operations. However, starting with ONTAP 9.9.1 you can initiate volume backup operations of your on-premises ONTAP clusters using ONTAP System Manager. [See how to use System Manager to back up your volumes to the cloud using Cloud Backup](#).

A Beta feature released in January 2021 allows you to run compliance scans on the backed up volumes from your on-premises systems. Typically, compliance scans are free up to 1 TB of data, and then a cost for the service is applied for data over 1 TB. When combining Backup and Data Sense for your on-premises volumes, the cost for scans on those on-prem volumes is free. Learn more about how [Cloud Data Sense](#) can get your business applications and cloud environments privacy ready.

Quick start

Get started quickly by following these steps, or scroll down to the remaining sections for full details.



Verify support for your configuration

- You have discovered the on-premises cluster and added it to a working environment in Cloud Manager. See [Discovering ONTAP clusters](#) for details.
 - The cluster is running ONTAP 9.7P5 or later.
 - The cluster has a SnapMirror license — it is included as part of the PREM or Data Protection bundle.
- You have subscribed to the [Azure](#), the [AWS](#), or the [Google](#) Cloud Manager Marketplace Backup offering, or you have purchased [and activated](#) a Cloud Backup BYOL license from NetApp.
- You have a valid cloud provider subscription for the object storage space where your backups will be located.
- For AWS and GCP, you need to have an account that has an access key and the required permissions so the ONTAP cluster can back up the data.



Enable Cloud Backup on the system

Select the working environment and click **Enable** next to the Backup & Compliance service in the right-panel, and then follow the setup wizard.



3

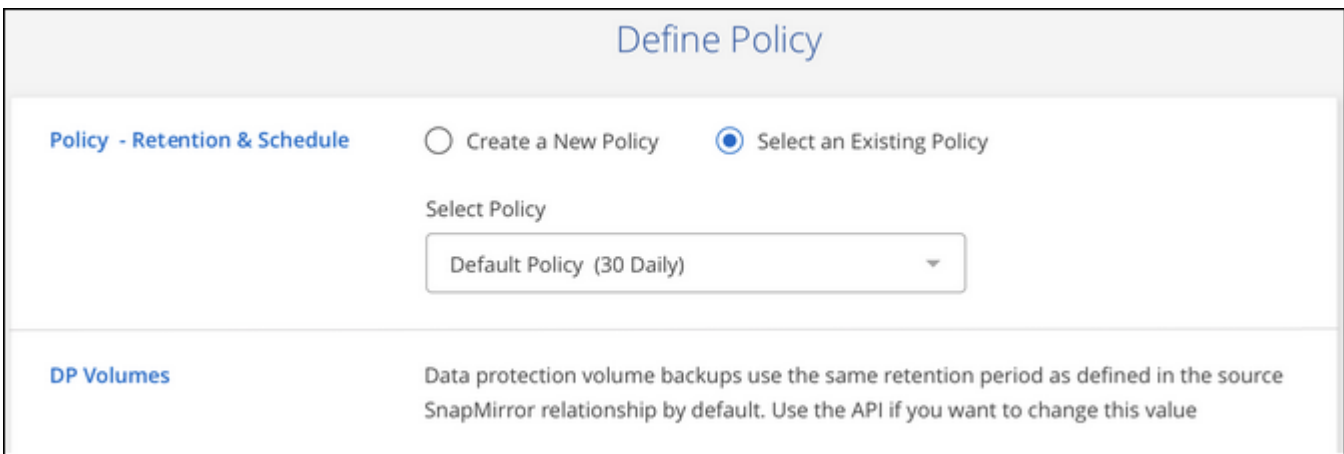
Select the cloud provider and enter the provider details

Select the provider and then enter the provider details. You also need to specify the IPspace in the ONTAP cluster where the volumes reside.

4

Define the backup policy

The default policy backs up volumes every day and retains the most recent 30 backup copies of each volume. Change to daily, weekly, or monthly backups, or select one of the system-defined policies that provide more options. You can also change the number of backup copies to retain.



5

Select the volumes that you want to back up

Identify which volumes you want to back up from the cluster.

6

Activate Compliance scans on the backed up volumes (optional)

Choose whether you want to have Cloud Data Sense scan the volumes that are backed up in the cloud.

7

Restore your data, as needed

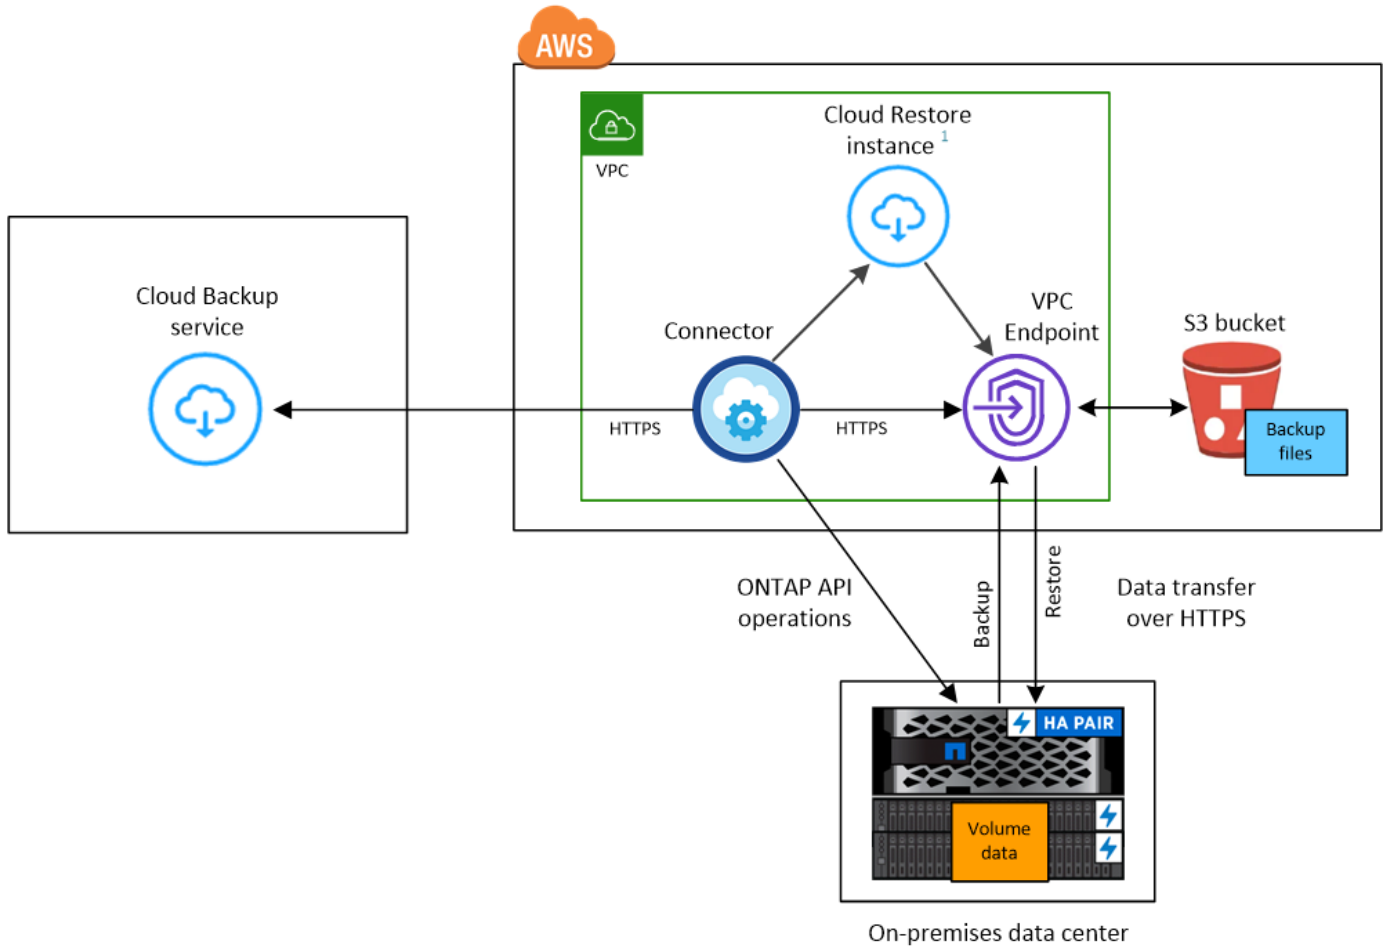
Choose to restore an entire backup to a new volume, or to restore individual files from the backup to an existing volume. You can restore data to a Cloud Volumes ONTAP system that is using the same cloud provider, or to an on-premises ONTAP system.

See [Restoring volume data from backup files](#) for details.

Overview diagrams

The following diagrams show each component when backing up an on-prem ONTAP system to object storage and the connections that you need to prepare between them.

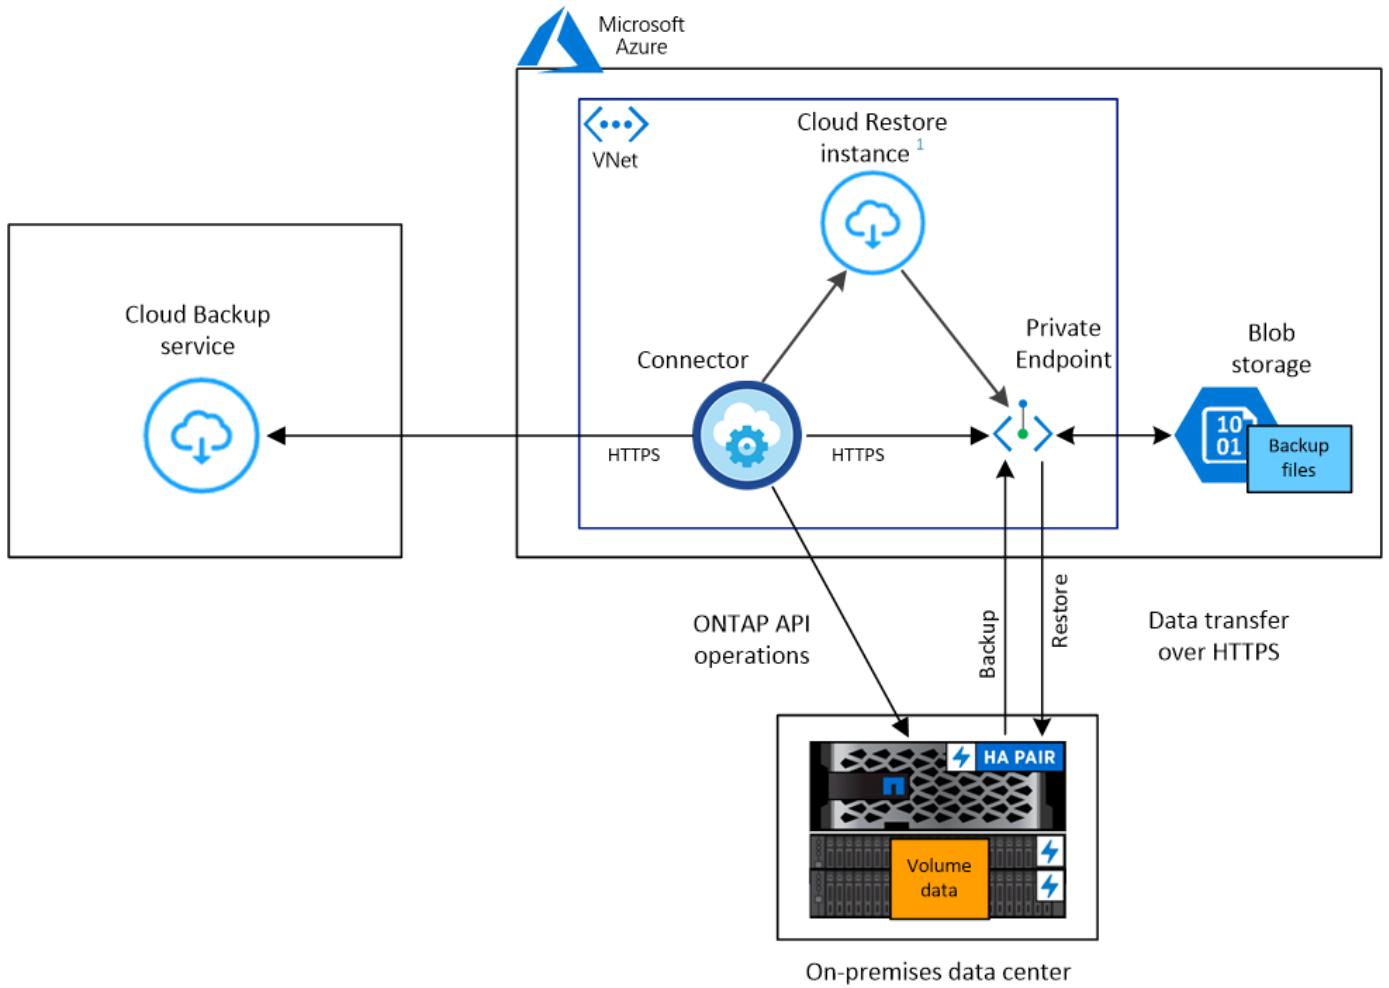
Amazon S3 configuration:



¹ Cloud Restore instance is active only during single-file restore operations.

Note that when the Cloud Restore instance is deployed in the cloud, it is located in the same subnet as the Connector.

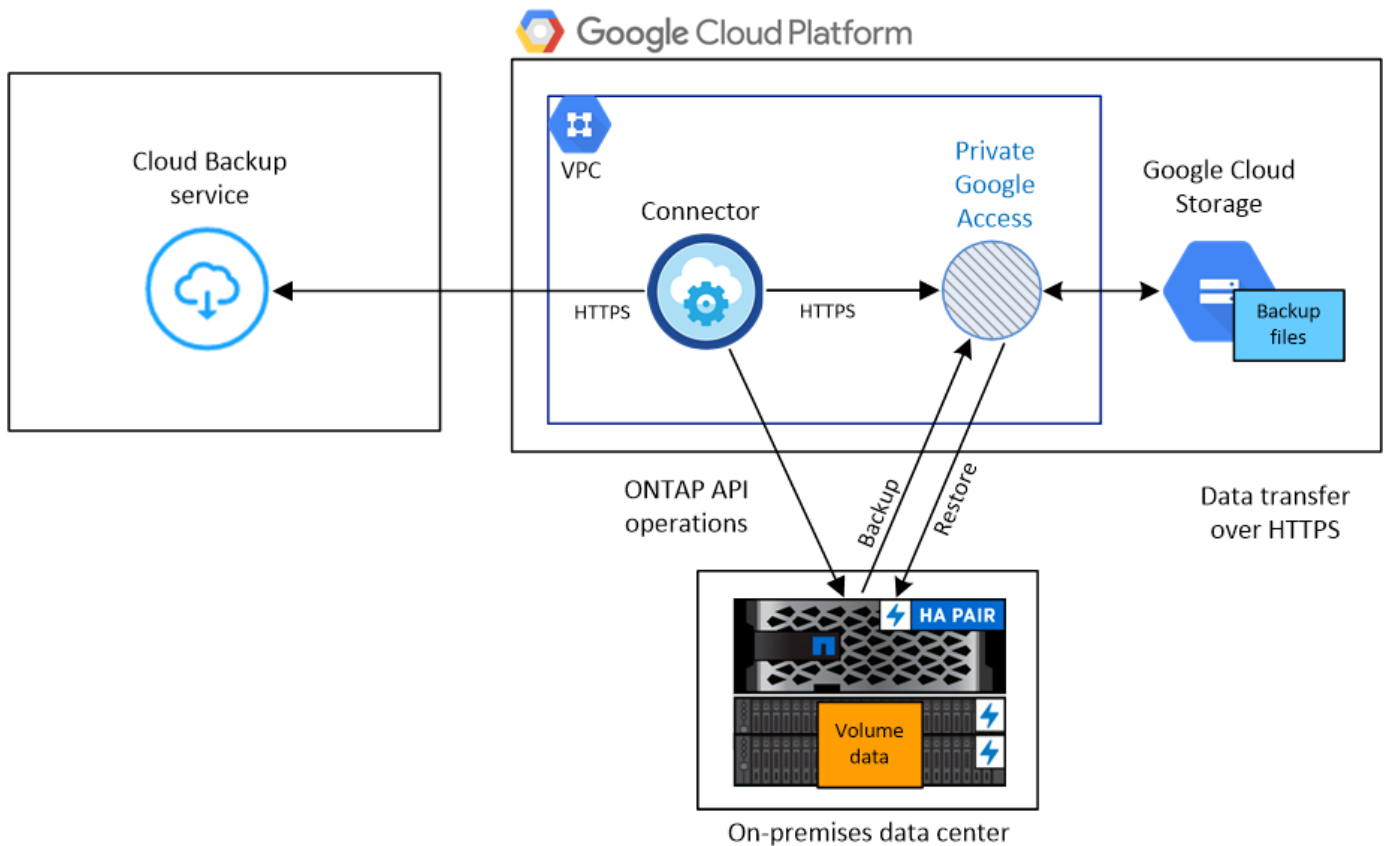
Azure Blob configuration:



¹ Cloud Restore instance is active only during single-file restore operations.

Note that when the Cloud Restore instance is deployed in the cloud, it is located in the same subnet as the Connector.

Google Cloud Storage configuration:



Note that the Cloud Restore instance is not shown in this diagram because single-file restore is not currently supported in GCP.

Requirements

Read the following requirements to make sure you have a supported configuration before you start backing up on-premises volumes to object storage.

Preparing your ONTAP clusters

You need to discover your on-premises ONTAP clusters in Cloud Manager before you can start backing up volume data.

[Learn how to discover a cluster.](#)

ONTAP requirements

- ONTAP 9.7P5 and later.
- A SnapMirror license (included as part of the PREM or Data Protection bundle).

See how to [manage your cluster licenses](#).

- Time and time zone are set correctly.

See how to [configure your cluster time](#).

Cluster networking requirements

- The ONTAP cluster initiates an HTTPS connection over port 443 to the cloud object storage.

ONTAP reads and writes data to and from object storage. The object storage never initiates, it just responds.

- An inbound connection is required from the Connector, which can reside in an AWS VPC, Azure VNet, or Google Cloud Platform VPC; depending on the object storage provider you are using.
- 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 Cloud Backup, 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.

- Node and intercluster LIFs are able to access the internet.
- DNS servers 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 Cloud Backup 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).

Creating or switching Connectors

A Connector is required to back up data to the cloud, and the Connector must be in the same cloud provider as the destination object storage. For example, when backing up data to AWS S3 you must use a Connector that's in an AWS VPC. You cannot use a Connector that is deployed on-premises. You'll either need to create a new Connector or make sure that the currently selected Connector resides in the correct provider.

- [Learn about Connectors](#)
- [Creating a Connector in AWS](#)
- [Creating a Connector in Azure](#)
- [Creating a Connector in GCP](#)
- [Switching between Connectors](#)

Preparing networking for the Connector

Ensure that the Connector has the required networking connections.

Steps

1. Ensure that the network where the Connector is installed enables the following connections:
 - An outbound internet connection to the Cloud Backup service over port 443 (HTTPS)
 - An HTTPS connection over port 443 to your object storage (S3, Blob, or Google)
 - An HTTPS connection over port 443 to your ONTAP clusters
2. Enable an endpoint to your object storage:
 - For AWS: Enable a VPC Endpoint to S3. This is needed if you have a Direct Connect or VPN

connection from your ONTAP cluster to the VPC and you want communication between the Connector and S3 to stay in your AWS internal network.

- For Azure: 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 Connector and Blob storage to stay in your virtual private network.
- For Google: Enable Private Google Access on the subnet where you plan to deploy the Service Connector. [Private Google Access](#) is needed if you have a direct connection from your ONTAP cluster to the VPC and you want communication between the Connector and Google Cloud Storage to stay in your virtual private network.

Note that Private Google Access works with VM instances that have only internal (private) IP addresses (no external IP addresses).

Supported regions

You can create backups from on-premises systems to the public cloud in all regions [where Cloud Volumes ONTAP is supported](#). You specify the region where the backups will be stored when you set up the service.

License requirements

For Cloud Backup PAYGO licensing, you'll need a subscription to the [Azure](#), the [AWS](#), or the [Google Cloud Manager Marketplace Backup](#) offering before you enable Cloud Backup. Billing for Cloud Backup is done through this subscription.

For Cloud Backup BYOL licensing, you do not need a subscription. You need the serial number from NetApp that enables you to use the service for the duration and capacity of the license. See [Managing your Backup BYOL license](#).

And you need to have a subscription from your cloud provider for the object storage space where your backups will be located.

Preparing Amazon S3 for backups

When you are using Amazon S3, you must configure permissions for Cloud Manager to access the S3 bucket, and you must configure permissions so the on-premises ONTAP cluster can access the S3 bucket.

Steps

1. Provide the following S3 permissions (from the latest [Cloud Manager policy](#)) to the IAM role that provides Cloud Manager with permissions:

```

{
    "Sid": "backupPolicy",
    "Effect": "Allow",
    "Action": [
        "s3:DeleteBucket",
        "s3:GetLifecycleConfiguration",
        "s3:PutLifecycleConfiguration",
        "s3:PutBucketTagging",
        "s3:ListBucketVersions",
        "s3:GetObject",
        "s3:ListBucket",
        "s3:ListAllMyBuckets",
        "s3:GetBucketTagging",
        "s3:GetBucketLocation",
        "s3:GetBucketPolicyStatus",
        "s3:GetBucketPublicAccessBlock",
        "s3:GetBucketAcl",
        "s3:GetBucketPolicy",
        "s3:PutBucketPublicAccessBlock"
    ],
    "Resource": [
        "arn:aws:s3:::netapp-backup-*"
    ]
}

```

2. Provide the following permissions to the IAM user so that the ONTAP cluster can back up data to S3.

```

"s3:ListAllMyBuckets",
"s3:ListBucket",
"s3:GetBucketLocation",
"s3:GetObject",
"s3:PutObject",
"s3:DeleteObject"

```

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

3. Provide the following permissions for the Cloud Restore instance:

```

"Action": [
    "ec2:DescribeInstanceTypeOfferings",
    "ec2:StartInstances",
    "ec2:StopInstances",
    "ec2:TerminateInstances"
]

```

- If your virtual or physical network uses a proxy server for internet access, ensure that the Cloud Restore instance has outbound internet access to contact the following endpoints.

Endpoints	Purpose
http://amazonlinux.us-east-1.amazonaws.com/2/extras/docker/stable/x86_64/4bf88ee77c395ffe1e0c3ca68530dfb3a683ec65a4a1ce9c0ff394be50e922b2/	CentOS package for the Cloud Restore Instance AMI.
http://cloudmanagerinfraprod.azurecr.io https://cloudmanagerinfraprod.azurecr.io	Cloud Restore Instance image repository.

- Create or locate an access key.

Cloud Backup passes the access key on to the ONTAP cluster. The credentials are not stored in the Cloud Backup service.

See the [AWS Documentation: Managing Access Keys for IAM Users](#) for details.

Preparing Azure Blob storage for backups

The Cloud Restore virtual machine requires outbound internet access. If your virtual or physical network uses a proxy server for internet access, ensure that the instance has outbound internet access to contact the following endpoints.

Endpoints	Purpose
http://olcentgbl.trafficmanager.net https://olcentgbl.trafficmanager.net	Provides CentOS packages for the Cloud Restore virtual machine.
http://cloudmanagerinfraprod.azurecr.io https://cloudmanagerinfraprod.azurecr.io	Cloud Restore virtual machine image repository.

Preparing Google Cloud Storage for backups

When you set up backup, you need to provide storage access keys for a service account that has Storage Admin permissions. A service account enables Cloud Backup 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

- [Create a service account that has the predefined Storage Admin role.](#)

2. Go to [GCP Storage Settings](#) and create access keys for the service account:
 - a. Select a project, and click **Interoperability**. If you haven't already done so, click **Enable interoperability access**.
 - b. Under **Access keys for service accounts**, click **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 Cloud Backup later when you configure the backup service.

Enabling Cloud Backup

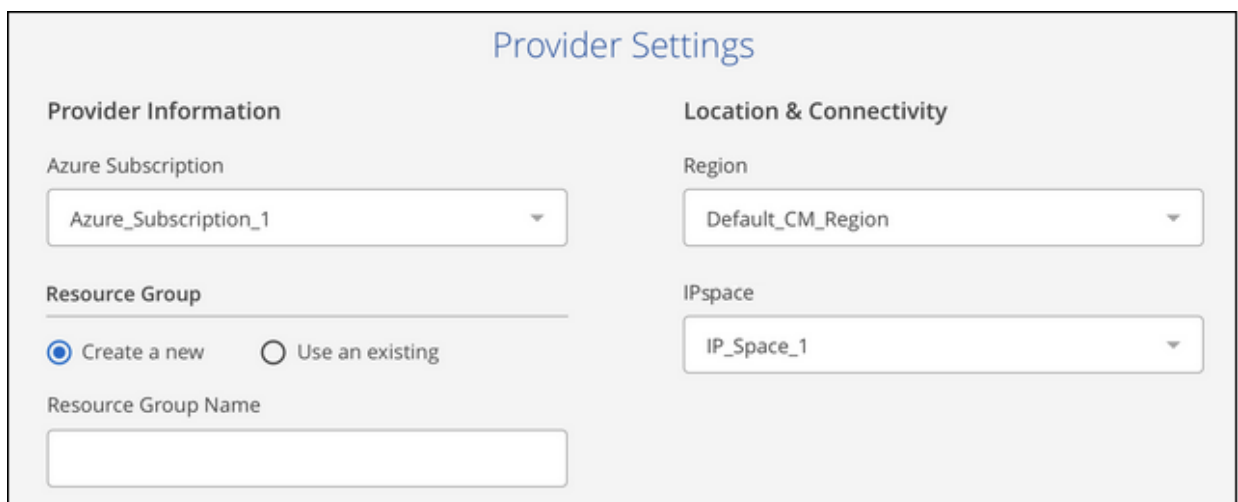
Enable Cloud Backup at any time directly from the on-premises working environment.

Steps

1. From the Canvas, select the working environment and click **Enable** next to the Backup & Compliance service in the right-panel.



2. Select the provider, click **Next**, and then enter the provider details:
 - For Azure, enter:
 - a. The Azure subscription used for backups and the Azure region where the backups will be stored.
 - b. The resource group - you can create a new resource group or select an existing resource group.
 - c. The IPspace in the ONTAP cluster where the volumes you want to back up reside.

A screenshot of a 'Provider Settings' form. The form is divided into two columns: 'Provider Information' and 'Location & Connectivity'. Under 'Provider Information', there is a dropdown for 'Azure Subscription' (value: Azure_Subscription_1), a section for 'Resource Group' with radio buttons for 'Create a new' (selected) and 'Use an existing', and a text input for 'Resource Group Name'. Under 'Location & Connectivity', there is a dropdown for 'Region' (value: Default_CM_Region) and a dropdown for 'IPspace' (value: IP_Space_1).

- For AWS, enter:
 - a. The AWS Account, the AWS Access Key, and the Secret Key used to store the backups.
 - b. The AWS region where the backups will be stored.
 - c. The IPspace in the ONTAP cluster where the volumes you want to back up reside.

Provider Settings

<p>Provider Information</p> <p>AWS Account <input type="text" value="AWS_Account_1"/></p> <p>AWS Access Key <input type="text" value="Enter AWS Access Key"/></p> <p>AWS Secret Key <input type="text" value="Enter AWS Secret Key"/></p>	<p>Location & Connectivity</p> <p>Region <input type="text" value="us-east-2"/></p> <p>IPspace i <input type="text" value="IP_Space_1"/></p>
---	--

- For Google, enter:
 - a. The Google Cloud Project where you want the Google Cloud Storage bucket to be created for backups. This can be a different Project than where Cloud Manager resides. (The Project must have a Service Account that has the predefined Storage Admin role.)
 - b. The Google Access Key and Secret Key used to store the backups.
 - c. The Google region where the backups will be stored. This can be a different region than where Cloud Manager resides.
 - d. The IPspace in the ONTAP cluster where the volumes you want to back up reside.

Provider Settings

<p>Provider Information</p> <p>Google Cloud Project <input type="text" value="Cloud Manager Default Project"/></p> <p>Google Cloud Access Key <input type="text" value="Enter Google Cloud Access Key"/></p> <p>Google Cloud Secret Key <input type="text" value="Enter Google Cloud Secret Key"/></p>	<p>Location & Connectivity</p> <p>Region <input type="text" value="Cloud Manager Default Region"/></p> <p>IPspace i <input type="text" value="IP_Space_1"/></p>
--	---

- For StorageGRID, see how to [back up on-premises ONTAP data to the private cloud](#).

Note that you cannot change this information after the service has started.

3. Click **Next** after you've entered the provider details.
4. In the *Define Policy* page, select an existing backup schedule and retention value, or define a new backup policy, and click **Next**.

Define Policy

Policy - Retention & Schedule

Create a New Policy
 Select an Existing Policy

Select Policy

Default Policy (30 Daily) ▼

DP Volumes

Data protection volume backups use the same retention period as defined in the source SnapMirror relationship by default. Use the API if you want to change this value

See [the list of existing policies](#).

5. Select the volumes that you want to back up.

- To back up all volumes, check the box in the title row (Volume Name).
- To back up individual volumes, check the box for each volume (Volume_1).

Select Volumes

57 Volumes 🔍


<input checked="" type="checkbox"/>	Volume Name	Volume Type	SVM Name	Used Capacity	Allocated Capacity	Backup Status
<input checked="" type="checkbox"/>	Volume_Name_1	RW	SVM_Name_1	0.25 TB	10 TB	⊖ Not Active
<input checked="" type="checkbox"/>	Volume_Name_2	RW	SVM_Name_2	0.25 TB	10 TB	⊖ Not Active
<input checked="" type="checkbox"/>	Volume_Name_3	RW	SVM_Name_3	0.25 TB	10 TB	⊖ Not Active
<input checked="" type="checkbox"/>	Volume_Name_4	DP	SVM_Name_4	0.25 TB	10 TB	⊖ Not Active
<input checked="" type="checkbox"/>	Volume_Name_5	RW	SVM_Name_5	0.25 TB	10 TB	⊖ Not Active

6. Click **Activate Backup** and Cloud Backup starts taking the initial backups of your volumes.

When creating backup files in AWS or Azure, you are prompted whether you want to run compliance scans on the backed up volumes. Cloud Data Sense scans are free when you run them on the backed up volumes (except for the [cost of the deployed Cloud Data Sense instance](#)).

Activate Compliance on your Backed Up Volumes

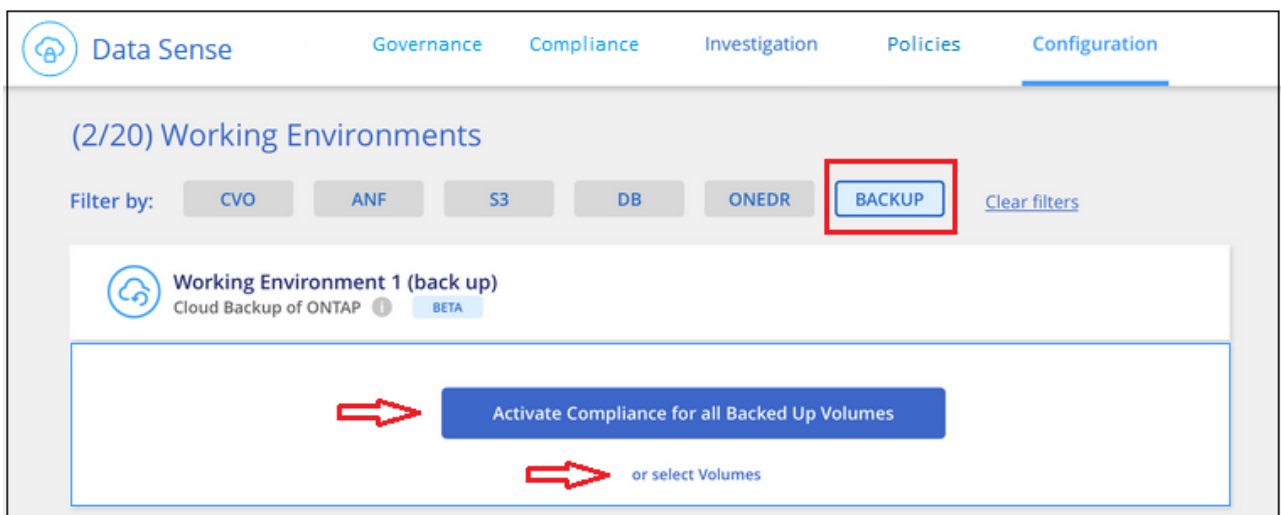
You have successfully activated Backup to Cloud on 12 Volumes in your working environment "Name 1".

 **Data Sense**

- > Cloud Compliance offer automated controls for data privacy regulations such as the GDPR, CCPA and more.
- > Driven by powerful artificial intelligence algorithms, Cloud Compliance gets your business application data and cloud environments privacy ready.

[Go to Compliance](#) [Close](#)

7. Click **Go to Compliance** to activate compliance scans on the volumes. (If you choose **Close** and not to scan these backed up volumes, you can always [enable this functionality](#) later from Cloud Data Sense.)
 - If an instance of Cloud Data Sense is already deployed in your environment, you are directed to the Configuration page to select the volumes you want to scan in each on-premises working environment that has backups. See [how to choose the volumes](#).



The screenshot shows the Data Sense Configuration page. The top navigation bar includes 'Data Sense', 'Governance', 'Compliance', 'Investigation', 'Policies', and 'Configuration'. The main content area is titled '(2/20) Working Environments'. Below this, there is a 'Filter by:' section with buttons for 'CVO', 'ANF', 'S3', 'DB', 'ONEDR', and 'BACKUP'. The 'BACKUP' button is highlighted with a red box. Below the filters, there is a section for 'Working Environment 1 (back up)' with a 'BETA' label. A large blue button labeled 'Activate Compliance for all Backed Up Volumes' is highlighted with a red arrow. Below this button, there is a link 'or select Volumes' also highlighted with a red arrow.

- If Cloud Data Sense has not been deployed, you are directed to the Compliance page where you can choose to deploy Compliance in the cloud or in your premises. We strongly recommend deploying it in the cloud. Go [here](#) for installation requirements and instructions.

The screenshot shows the Data Sense web interface. At the top left is the 'Data Sense' logo. Below it is a link 'How does it work?'. The main heading is 'Always-on Privacy & Compliance Controls'. Below this is a paragraph: 'Automated controls for data privacy regulations - GDPR, CCPA, HIPAA and more. Driven by powerful artificial intelligence algorithms, Data Sense gets your business application data and cloud environments privacy ready.' There are two buttons: 'Deploy Data Sense in the Cloud' and 'Deploy Data Sense On-Premises'. Below the buttons is a link: 'Learn about the differences between cloud deployment and on-premises deployment'. On the right side, there is a 'Compliance Status' dashboard. It features a 'Data Distribution' chart with a shield icon and a progress bar showing 75% Non-Sensitive, 20% Personal, and 5% Sensitive Personal. Below this are two sections: '28,000 Personal Files' and '7,000 Sensitive Personal Files'. Each section has a 'View All' button and a list of file types with counts: 'Email Address' (2,700 Files), 'Credit Card' (2,700 Files), 'Health' (2,700 Files), and 'Ethnicity' (2,700 Files).

After you have deployed Compliance you can choose the volumes you want to scan as described above.

Result

Cloud Backup backs up your volumes from the on-premises ONTAP system, and optionally, Cloud Data Sense runs compliance scans on the backed up volumes.

What's next?

You can [start and stop backups for volumes or change the backup schedule](#) and you can [restore entire volumes or individual files from a backup file](#).

You can also [view the results of the compliance scans](#) and review other features of Cloud Data Sense that can help you understand data context and identify sensitive data in your organization.



The scan results are not available immediately because Cloud Backup has to finish creating the backups before Cloud Data Sense can start compliance scans.

Copyright Information

Copyright © 2021 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

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.