



Backing up on-premises ONTAP data to Amazon S3 Cloud Manager

Tom Onacki
September 15, 2021

Table of Contents

- Backing up on-premises ONTAP data to Amazon S3 1
 - Quick start 1
 - Requirements 3
 - Enabling Cloud Backup 9

Backing up on-premises ONTAP data to Amazon S3

Complete a few steps to get started backing up data from your on-premises ONTAP systems to Amazon S3 storage.

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 Premium Bundle or Data Protection Bundle.
 - The cluster must have the required network connections to S3 storage and to the Connector.
- The Connector must have the required network connections to S3 storage and to the cluster, and the required permissions.
- You have a valid AWS subscription for the object storage space where your backups will be located.
- You have an AWS Account with an access key and secret key, and the [required permissions](#) so the ONTAP cluster can back up and restore 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.



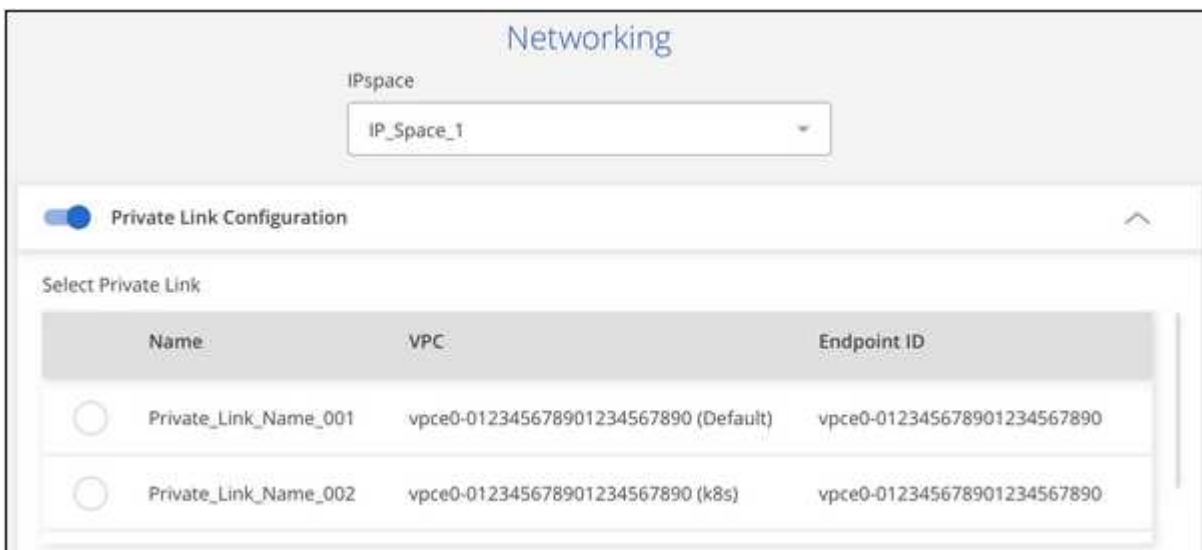
3 Select the cloud provider and enter the provider details

Select Amazon Web Services as your provider and then enter the provider details. You'll need to select the AWS Account and the region where you want to create the backups. You can also choose your own customer-managed key for data encryption instead of using the default Amazon S3 encryption key.



4 Select the cluster IPspace and optionally select an AWS PrivateLink connection

Select the IPspace in the ONTAP cluster where the volumes reside. You can also choose to use an existing AWS PrivateLink configuration for a more secure connection to the VPC from your on-prem data center.



5

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 hourly, 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.

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

6

Select the volumes that you want to back up

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

7

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.

8

Restore your data, as needed

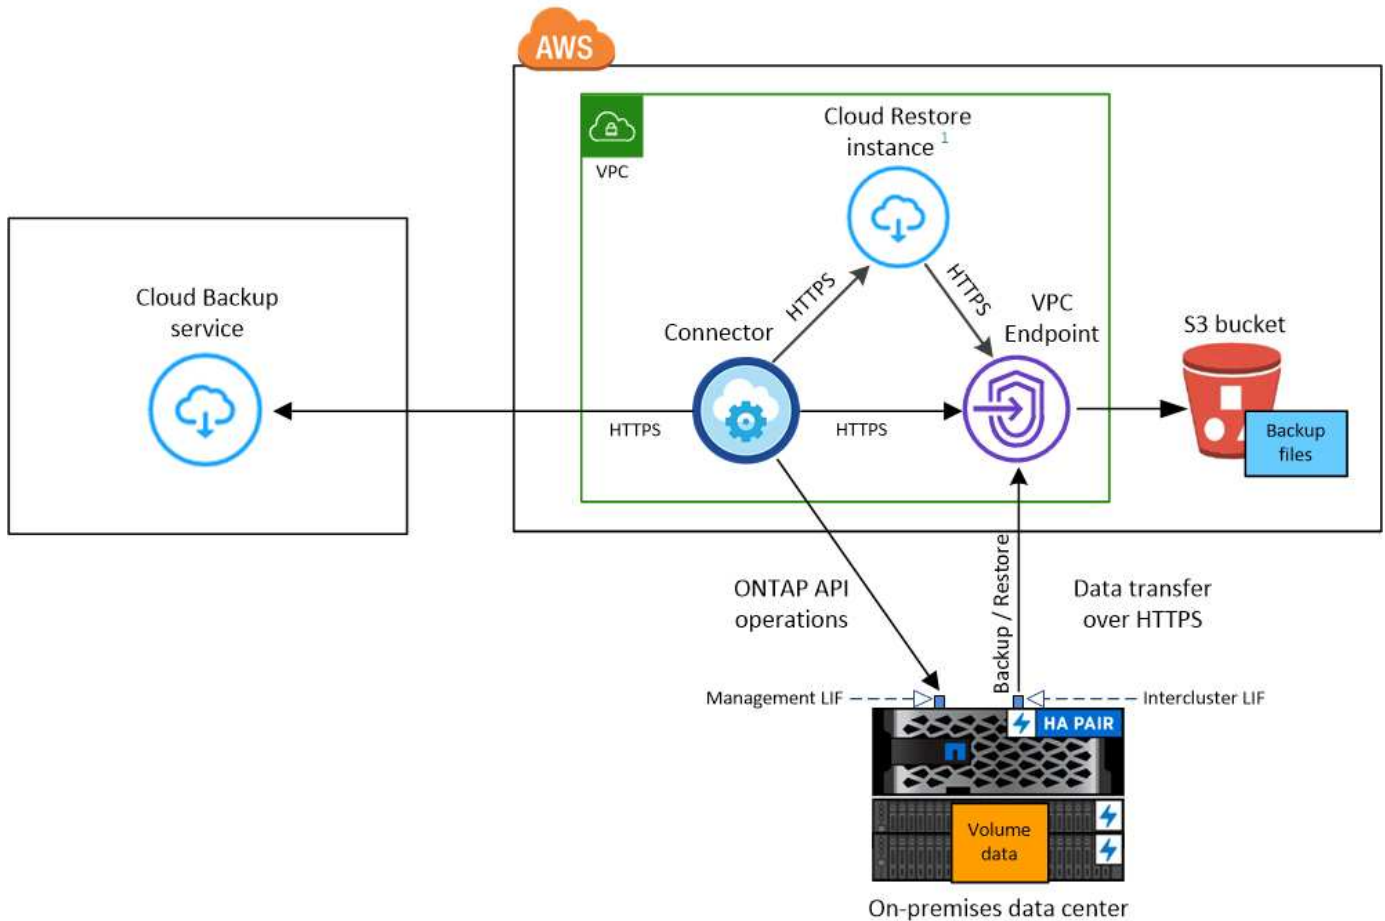
Choose to restore an entire backup to a new volume, or to restore individual files from the backup to an existing volume. You can restore data to a Cloud Volumes ONTAP system in AWS, or to an on-premises ONTAP system.

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

Requirements

Read the following requirements to make sure you have a supported configuration before you start backing up on-premises volumes to S3 storage.

The following image shows each component and the connections that you need to prepare between them:



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.

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 Premium Bundle or Data Protection Bundle).

Note: The "Hybrid Cloud Bundle" is not required when using the Cloud Backup service.

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 from the intercluster LIF to Amazon S3 storage for backup and restore operations.

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

- ONTAP requires an inbound connection from the Connector to the cluster management LIF. The Connector can reside in an AWS VPC.
- An intercluster LIF is required on each ONTAP node that hosts the volumes you want to back up. The LIF must be associated with the *IPspace* that ONTAP should use to connect to object storage. [Learn more about IPspaces](#).

When you set up 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.

- The nodes' 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](#).
- Note that if you are using a different IPspace than the Default, then you might need to create a static route to get access to the object storage.
- Update firewall rules, if necessary, to allow 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 an AWS VPC when backing up data to AWS S3 storage. You can't use a Connector that's 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](#)
- [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 S3 object storage
 - An HTTPS connection over port 443 to your ONTAP clusters
2. 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.

Supported regions

You can create backups from on-premises systems to Amazon S3 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

Before your 30-day free trial of the Cloud Backup service expires, you need to subscribe to a pay-as-you-go (PAYGO) Cloud Manager Marketplace offering from AWS, or purchase and activate a Cloud Backup BYOL license from NetApp. These licenses are for the account and can be used across multiple systems.

- For Cloud Backup PAYGO licensing, you'll need a subscription to the [AWS](#) Cloud Manager Marketplace offering to continue using Cloud Backup. Billing for Cloud Backup is done through this subscription.
- For Cloud Backup BYOL licensing, you don't 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. [Learn how to manage your BYOL licenses](#).

You need to have an AWS subscription for the object storage space where your backups will be located.

A SnapMirror license is required on the cluster. Note that the "Hybrid Cloud Bundle" is not required when using Cloud Backup.

Preparing Amazon S3 for backups

When you are using Amazon S3, you must configure permissions for the Connector to create and manage the S3 bucket, and you must configure permissions so the on-premises ONTAP cluster can read and write to the S3 bucket.

Steps

1. Confirm that the following S3 permissions (from the latest [Cloud Manager policy](#)) are part of the IAM role that provides the Connector 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. Add the following EC2 permissions to the IAM role that provides the Connector with permissions so that it can start, stop, and terminate the Cloud Restore instance:

```

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

```

3. During the Backup wizard you will be prompted to enter an access key and secret key. For that, you will need to create an IAM user with the following permissions. Cloud Backup passes these credentials on to the ONTAP cluster so that ONTAP can backup and restore data to the S3 bucket.

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

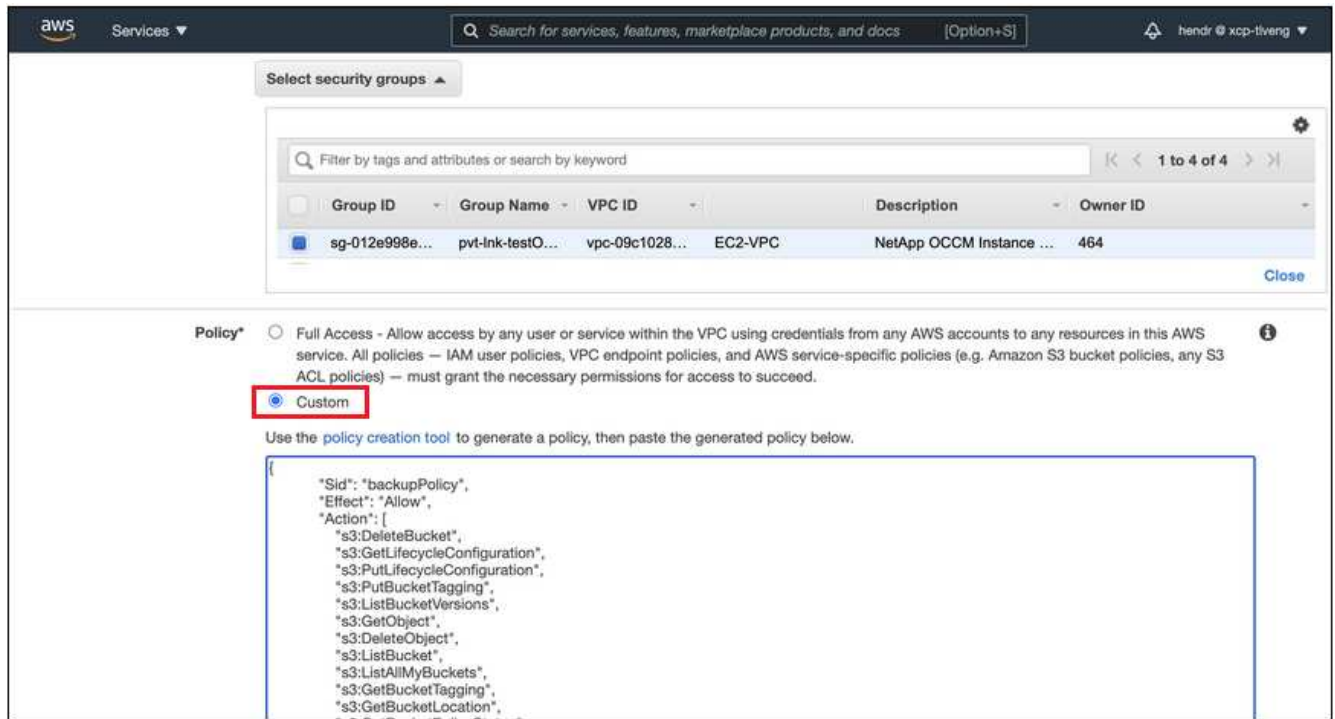
See the [AWS Documentation: Creating a Role to Delegate Permissions to an IAM User](#) for details.

4. 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.

5. You can choose your own custom-managed keys for data encryption in the activation wizard instead of using the default Amazon S3 encryption keys. In this case you'll need to have the encryption managed keys already set up. [See how to use your own keys](#).
6. If you want to have a more secure connection over the public internet from your on-prem data center to the VPC, there is an option to select an AWS PrivateLink connection in the activation wizard. It is required if you are connecting your on-premises system via VPN/DirectConnect. In this case you'll need to have created an Interface endpoint configuration using the Amazon VPC console or the command line. [See details about using AWS PrivateLink](#).

Note that you'll also need to modify the security group configuration that is associated with the Cloud Manager Connector. You must change the policy to "Custom" (from "Full Access"), and you must add the permissions from the backup policy as shown earlier (above).



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 Amazon Web Services as your provider and click **Next**.
3. Enter the provider details. Note that you can't change this information after the service has started.
 - a. The AWS Account, the AWS Access Key, and the Secret Key used to store the backups.

The access key and secret key are for the user you created to give the ONTAP cluster access to the S3 bucket.

- b. The AWS region where the backups will be stored.
- c. Whether you'll use the default Amazon S3 encryption keys or choose your own customer-managed keys from your AWS account to manage encryption of your data. ([See how to use your own keys](#)).

Provider Settings

Provider Information

AWS Account

AWS Access Key

AWS Secret Key

Location & Connectivity

Region

Encryption ⓘ

Encryption Key Type: AWS SSE-S3 [Change Key](#)

4. Click **Next** after you've entered the provider details.
5. Enter the networking details and click **Next**.
 - a. The IPspace in the ONTAP cluster where the volumes you want to back up reside. The intercluster LIFs for this IPspace must have outbound internet access.
 - b. Optionally, choose whether you'll use an AWS PrivateLink that you have previously configured. [See details about using an AWS PrivateLink.](#)

Networking

IPspace

Private Link Configuration

Select Private Link

Name	VPC	Endpoint ID
<input type="radio"/> Private_Link_Name_001	vpce0-012345678901234567890 (Default)	vpce0-012345678901234567890
<input type="radio"/> Private_Link_Name_002	vpce0-012345678901234567890 (k8s)	vpce0-012345678901234567890

6. 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](#).

7. 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).

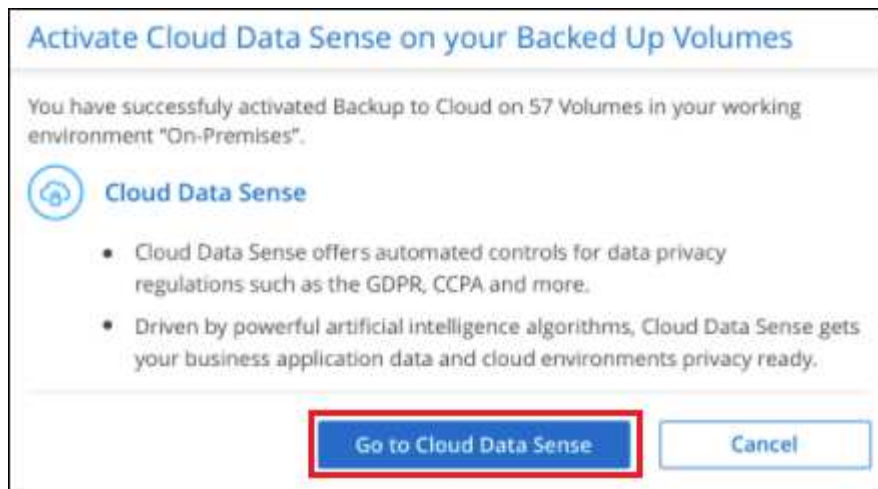
Select Volumes

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

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

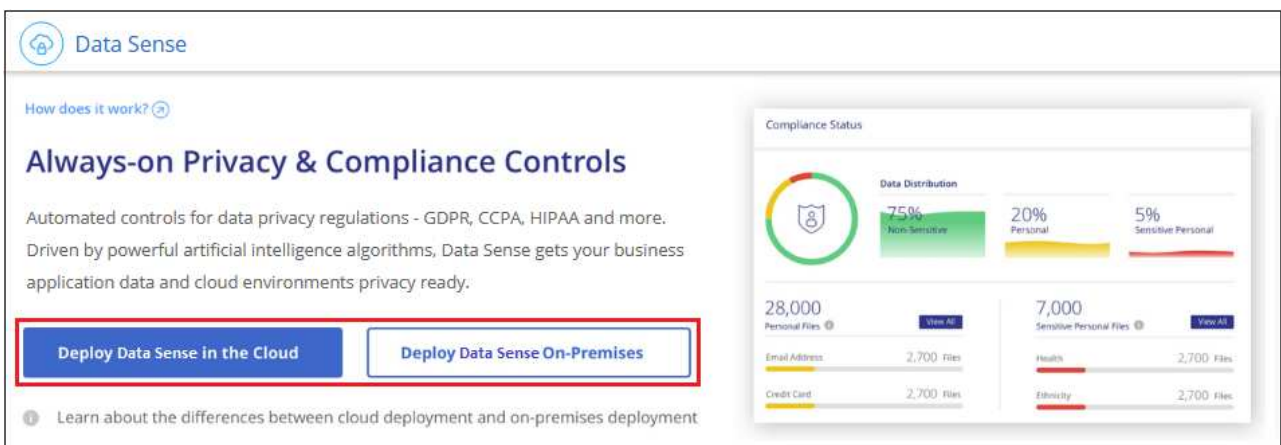
You are prompted whether you want to run compliance scans on the backed up volumes. Cloud Data Sense scans are free when you run them on the backed up volumes (except for the [cost of the deployed Cloud Data Sense instance](#)).



9. 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](#).



- 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.



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.