

Cloud Compliance for Azure NetApp Files

Tom Onacki, Ben Cammett, Nadav Shofet
October 21, 2020

In partnership with



Table of Contents

Contents

- Cloud Compliance for Azure NetApp Files..... 1
 - Learn about Cloud Compliance 3
 - Get started 6
 - Gaining visibility and control of private data 15
 - Viewing compliance reports 28
 - Responding to a Data Subject AccessRequest 32
 - Disabling Cloud Compliance 34
 - Frequently asked questions about Cloud Compliance 34

Gain insight into data privacy

Learn about Cloud Compliance

Cloud Compliance is a data privacy and compliance service for Cloud Manager that scans your [Azure NetApp Files](#) volumes to identify the personal and sensitive data. Using Artificial Intelligence (AI) driven technology, Cloud Compliance helps organizations understand data context and identify sensitive data.

Features

Cloud Compliance provides several tools that can help you with your compliance efforts. You can use Cloud Compliance to:

- Identify Personal Identifiable Information (PII)
- Identify a wide scope of sensitive information as required by GDPR, CCPA, PCI, and HIPAA privacy regulations
- Respond to Data Subject Access Requests (DSAR)

Cost

- The cost to use Cloud Compliance depends on the amount of data that you're scanning. As of October 7th, 2020, the first 1 TB of data that Cloud Compliance scans in a Cloud Manager workspace is free.

A subscription to the Azure Marketplace is required to continue scanning data after that point. See [pricing](#) for details.

[Learn how to subscribe.](#)

- Installing Cloud Compliance requires deploying a cloud instance, resulting in charges from Azure.
- Cloud Compliance requires that you have deployed a Connector. In many cases you already have a Connector because of other storage and services you are using in Cloud Manager. This instance as well results in charges from Azure.

How Cloud Compliance works

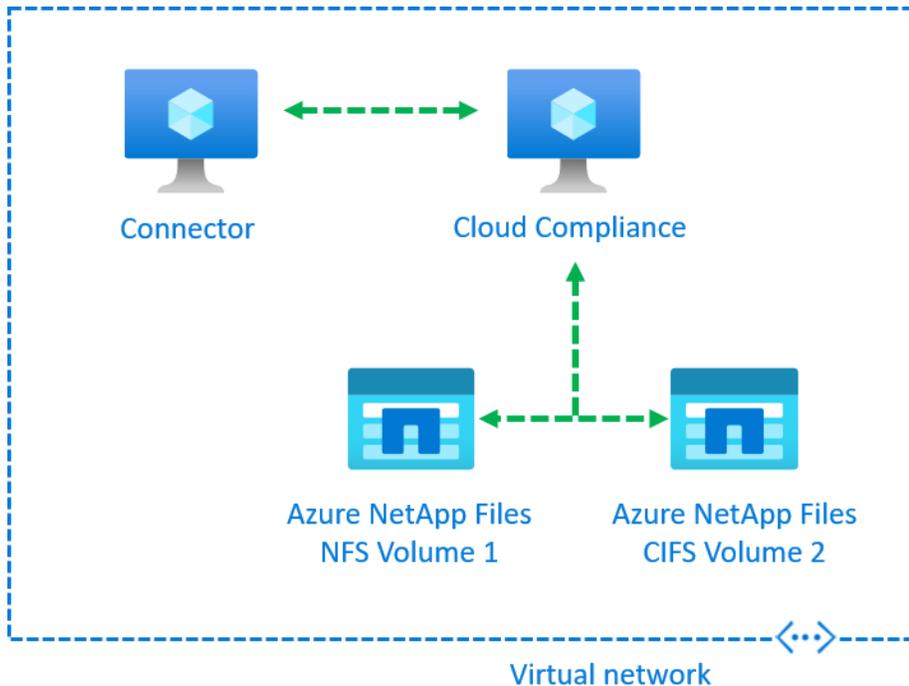
At a high-level, Cloud Compliance works like this:

1. You deploy an instance of Cloud Compliance in Cloud Manager.
2. You enable it on one or more Azure NetApp Files Volumes.
3. Cloud Compliance scans the data using an AI learning process.
4. In Cloud Manager, you click **Compliance** and use the provided dashboard and reporting tools to help in

your compliance efforts.

The Cloud Compliance instance

When you enable Cloud Compliance, Cloud Manager deploys a Cloud Compliance instance in the same subnet as the Connector. [Learn more about Connectors.](#)



Note the following about the instance:

- Cloud Compliance runs on Azure infrastructure - a Standard_D16s_v3 VM with a 512 GB disk.



Changing or resizing the instance/VM type isn't supported. You need to use the size that's provided.

- The instance is named *CloudCompliance* with a generated hash (UUID) concatenated to it. For example: *CloudCompliance-16bb6564-38ad-4080-9a92-36f5fd2f71c7*
- Only one Cloud Compliance instance is deployed per Connector.
- Upgrades of Cloud Compliance software is automated—you don't need to worry about it.



The instance should remain running at all times because Cloud Compliance continuously scans the data.

How scans work

After you enable Cloud Compliance and select the volumes you want to scan, it immediately starts scanning the

data to identify personal and sensitive data. For NFS volumes, the export policy must allow access from the Cloud Compliance instance. For CIFS volumes, Active Directory credentials are required. Cloud Compliance maps your organizational data, categorizes each file, and identifies and extracts entities and predefined patterns in the data. The result of the scan is an index of personal information, sensitive personal information, and data categories.

Cloud Compliance connects to the data like any other client by mounting NFS and CIFS volumes. NFS volumes are automatically accessed as read-only, while you need to provide Active Directory credentials to scan CIFS volumes.

After the initial scan, Cloud Compliance continuously scans each volume to detect incremental changes (this is why it's important to keep the instance running).

You can enable and disable scans at the [volume level](#).

Information that Cloud Compliance indexes

Cloud Compliance collects, indexes, and assigns categories to unstructured data (files). The data that Cloud Compliance indexes includes the following:

Standard metadata

Cloud Compliance collects standard metadata about files: the file type, its size, creation and modification dates, and so on.

Personal data

Personally identifiable information such as email addresses, identification numbers, or credit card numbers. [Learn more about personal data](#).

Sensitive personal data

Special types of sensitive information, such as health data, ethnic origin, or political opinions, as defined by GDPR and other privacy regulations. [Learn more about sensitive personal data](#).

Categories

Cloud Compliance takes the data that it scanned and divides it into different types of categories. Categories are topics based on AI analysis of the content and metadata of each file. [Learn more about categories](#).

Name entity recognition

Cloud Compliance uses AI to extract natural persons' names from documents. [Learn about responding to Data Subject Access Requests](#).

Networking overview

Cloud Manager deploys the Cloud Compliance instance with a security group that enables inbound HTTP connections from the Connector instance.

Outbound rules are completely open. Internet access is needed to install and upgrade the Cloud Compliance software and to send usage metrics.

If you have strict networking requirements, [learn about the endpoints that Cloud Compliance contacts](#).

User access to compliance information

The role each user has been assigned provides different capabilities within Cloud Manager and within Cloud Compliance:

- **Account Admins** can manage compliance settings and view compliance information for all working environments.
- **Workspace Admins** can manage compliance settings and view compliance information only for systems that they have permissions to access. If a Workspace Admin can't access a working environment in Cloud Manager, then they can't see any compliance information for the working environment in the Compliance tab.
- Users with the **Cloud Compliance Viewer** role can only view compliance information and generate reports for systems that they have permission to access. These users cannot enable/disable scanning of volumes.

[Learn more about Cloud Manager roles](#) and how to [add users with specific roles](#).

Get started

Deploy Cloud Compliance

Complete a few steps to deploy the Cloud Compliance instance in your Cloud Manager workspace.

Quick start

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



Create a Connector

If you don't already have a Connector, create a Connector in Azure see [creating a Connector in Azure](#).



Review prerequisites

Ensure that your cloud environment can meet the prerequisites, which includes 16 vCPUs for the Cloud Compliance instance, outbound internet access for the instance, connectivity between the Connector and Cloud Compliance over port 80, and more. [See the complete list](#).



Deploy Cloud Compliance

Launch the installation wizard to deploy the Cloud Compliance instance in Cloud Manager.



Subscribe to the Cloud Compliance service

The first 1 TB of data that Cloud Compliance scans in Cloud Manager is free. A subscription to the Azure Marketplace is required to continue scanning data after that point.

Creating a Connector

If you don't already have a Connector, create a Connector in Azure. See [creating a Connector in Azure](#). In most cases you will probably have a Connector set up before you attempt to activate Cloud Compliance because most [Cloud Manager features require a Connector](#), but there are cases when you need to set one up now.



You need to make sure you're deploying the Connector in the same region as the Azure NetApp Files volumes you wish to scan.

Reviewing prerequisites

Review the following prerequisites to make sure that you have a supported configuration before you deploy Cloud Compliance.

Enable outbound internet access

Cloud Compliance requires outbound internet access. If your virtual network uses a proxy server for internet access, ensure that the Cloud Compliance instance has outbound internet access to contact the following endpoints. Note that Cloud Manager deploys the Cloud Compliance instance in the same subnet as the Connector.

Endpoints	Purpose
https://cloudmanager.cloud.netapp.com	Communication with the Cloud Manager service, which includes Cloud Central accounts.
https://netapp-cloud-account.auth0.com https://auth0.com	Communication with NetApp Cloud Central for centralized user authentication.

https://hub.docker.com https://auth.docker.io https://registry-1.docker.io https://index.docker.io/ https://dseasb33srnrn.cloudfront.net/ https://production.cloudflare.docker.com/	Provides access to software images, manifests, and templates.
https://support.compliance.cloudmanager.cloud.netapp.com/	Provides access to software images, manifests, and templates. Enables NetApp to stream data from audit records, and to send logs and metrics.

Ensure that Cloud Manager has the required permissions

Ensure that Cloud Manager has permissions to deploy resources and create security groups for the Cloud Compliance instance. You can find the latest Cloud Manager permissions in [the policies provided by NetApp](#).

Check your vCPU limits

Ensure that your vCPU limit allows for the deployment of an instance with 16 cores. You'll need to verify the vCPU limit for the relevant instance family in the region where Cloud Manager is running. For more details on vCPU limits, see the following:

- [Azure documentation: Virtual machine vCPU quotas](#)

Ensure that the Connector can access Cloud Compliance

Ensure connectivity between the Connector and the Cloud Compliance instance. The security group for the Connector must allow inbound and outbound traffic over port 80 to and from the Cloud Compliance instance.

This connection enables deployment of the Cloud Compliance instance and enables you to view information in the Compliance tab.

Ensure that you can keep Cloud Compliance running

The Cloud Compliance instance needs to stay on to continuously scan your data.

Ensure web browser connectivity to Cloud Compliance

After Cloud Compliance is enabled, ensure that users access the Cloud Manager interface from a host that has a connection to the Cloud Compliance instance.

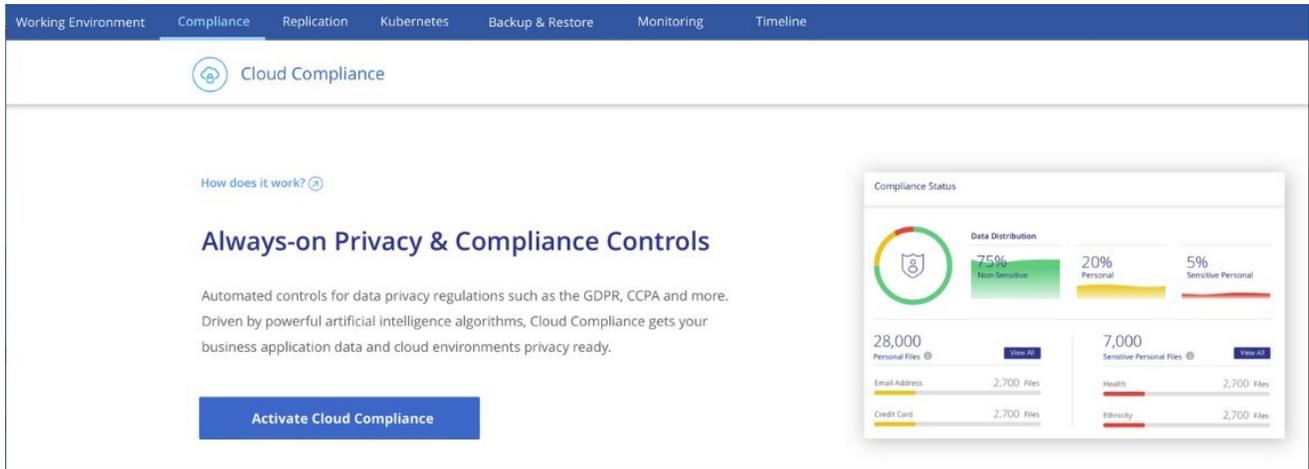
The Cloud Compliance instance uses a private IP address to ensure that the indexed data isn't accessible to the internet. As a result, the web browser that you use to access Cloud Manager must have a connection to that private IP address. That connection can come from a direct connection to Azure (for example, a VPN), or from a host that's inside the same network as the Cloud Compliance instance.

Deploying the Cloud Compliance instance

You deploy an instance of Cloud Compliance for each Cloud Manager instance.

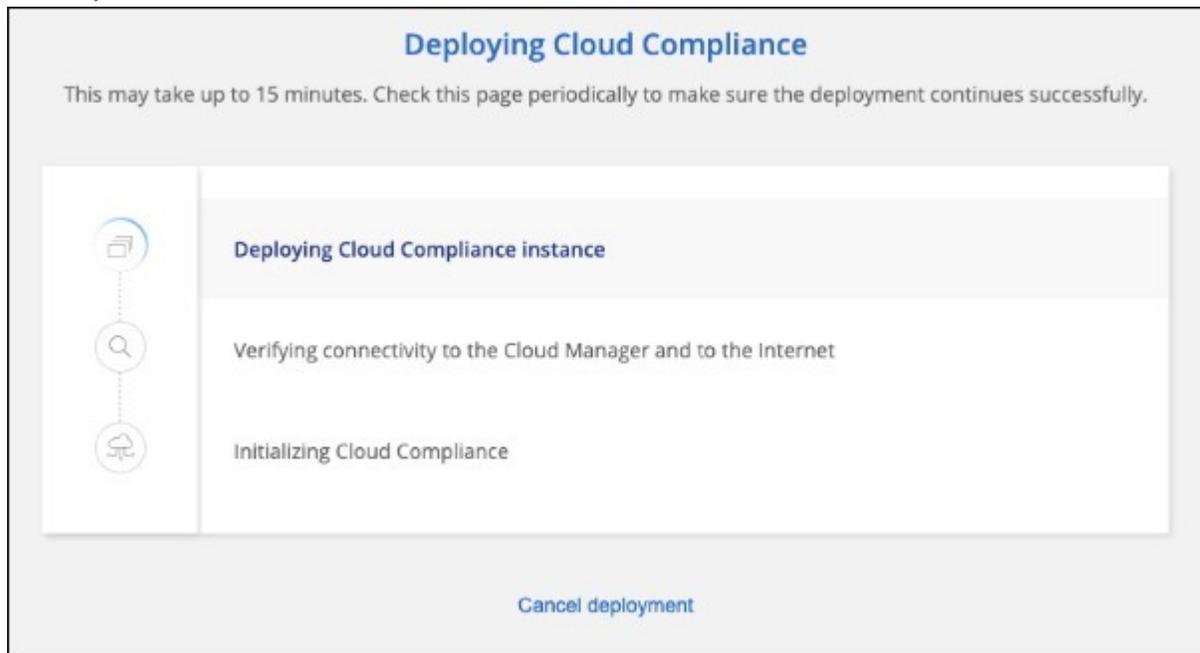
Steps

1. In Cloud Manager, click **Cloud Compliance**.
2. Click **Activate Cloud Compliance** to start the deployment wizard.



The screenshot shows the Cloud Compliance dashboard in Cloud Manager. The navigation bar includes Working Environment, Compliance (selected), Replication, Kubernetes, Backup & Restore, Monitoring, and Timeline. The main content area features a 'Cloud Compliance' header, a 'How does it work?' link, and a section titled 'Always-on Privacy & Compliance Controls'. Below this, there is a description of automated controls for data privacy regulations and a prominent 'Activate Cloud Compliance' button. On the right, a 'Compliance Status' widget displays a 'Data Distribution' chart with 75% Non-Sensitive, 20% Personal, and 5% Sensitive Personal data. It also shows 28,000 Personal Files and 7,000 Sensitive Personal Files, with sub-sections for Email Address, Credit Card, Health, and Ethnicity, each with 2,700 files.

3. The wizard displays progress as it goes through the deployment steps. It will stop and ask for input if it runs into any issues.



The screenshot shows the 'Deploying Cloud Compliance' deployment wizard. The title is 'Deploying Cloud Compliance' and the message states 'This may take up to 15 minutes. Check this page periodically to make sure the deployment continues successfully.' The progress bar shows three steps: 'Deploying Cloud Compliance instance' (completed), 'Verifying connectivity to the Cloud Manager and to the Internet' (in progress), and 'Initializing Cloud Compliance' (pending). A 'Cancel deployment' button is located at the bottom.

4. When the instance is deployed, click **Continue to configuration** to go to the *Scan Configuration* page.

Result

Cloud Manager deploys the Cloud Compliance instance in Azure.

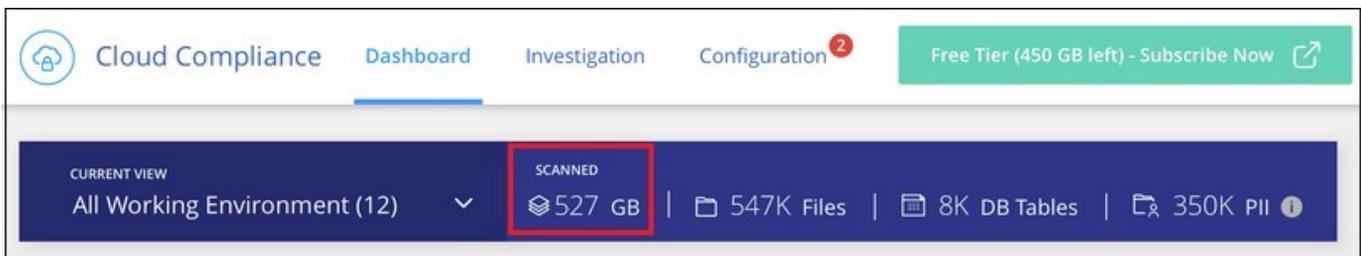
What's Next

From the Scan Configuration page you can select volumes that you want to scan for compliance.

Subscribing to the Cloud Compliance service

The first 1 TB of data that Cloud Compliance scans in a Cloud Manager workspace is free. A subscription to the Azure Marketplace is required to continue scanning data after that point.

You can subscribe at any time and you will not be charged until the amount of data exceeds 1 TB. You can always see the total amount of data that is being scanned from the Cloud Compliance Dashboard. And the *Subscribe Now* button makes it easy to subscribe when you are ready.



Note: If you are prompted by Cloud Compliance to subscribe, but you already have an Azure subscription, you're probably using the old **Cloud Manager** subscription and you need to change to the new **NetApp Cloud Manager** subscription. See

[Changing to the new NetApp Cloud Manager plan in Azure](#) for details.

Steps

These steps must be completed by a user who has the *Account Admin* role.

1. In the upper right of the Cloud Manager console, click the Settings icon, and select **Credentials**.



2. Find the credentials for the Azure Managed Service Identity.

The subscription must be added to the Instance Profile or Managed Service Identity. Charging won't work otherwise.

If you already have a subscription, then you're all set—there's nothing else that you need to do.

3. If you don't have a subscription yet, hover over the credentials and click the action menu.

4. Click **Add Subscription**.
5. Click **Add Subscription**, click **Continue**, and follow the steps.

The following video shows how to associate a Marketplace subscription to an Azure subscription:

https://docs.netapp.com/us-en/occm/media/video_subscribing_azure.mp4 (video)

Changing to the new Cloud Manager plan in Azure

Cloud Compliance was added to the Azure Marketplace named [NetApp Cloud Manager](#) as of October 7, 2020. If you already have the original Azure **Cloud Manager** subscription it will not allow you to use Cloud Compliance.

You need to follow these steps and select the new **NetApp Cloud Manager** and then remove the old **Cloud Manager** subscription.



If your existing Subscription was issued with a special private offer, you need to contact NetApp so that we can issue a new special private offer with Compliance included.

Steps

These steps are similar to adding a new subscription as described above, but vary in a few places.

1. In the upper right of the Cloud Manager console, click the Settings icon, and select **Credentials**.
2. Find the credentials for the Azure Managed Service Identity that you want to change the subscription for and hover over the credentials and click **Associate Subscription**.

The details for your current Marketplace Subscription are displayed.

3. Click **Add Subscription**, click **Continue**, and follow the steps. You are redirected to Azure portal in order to create the new subscription.
4. Make sure you select the plan **NetApp Cloud Manager** that provides access to Cloud Compliance and not **Cloud Manager**.
5. Go through the steps in the video to associate a Marketplace subscription to an Azure subscription:

https://docs.netapp.com/us-en/occm/media/video_subscribing_azure.mp4 (video)

6. Return to Cloud Manager, select the new subscription, and click **Associate**.
7. To verify your subscription has changed, hover over the “i” above subscription in the Credentials card.

Now you can unsubscribe your old subscription from the Azure portal.

In the Azure portal, go to Software as a Service (SaaS), select the subscription, and click **Unsubscribe**.

Activate scanning on Azure NetApp Files

Getting started with Cloud Compliance for Azure NetApp Files

Complete a few steps to get started with Cloud Compliance for Azure NetApp Files.

Quick start

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



Deploy the Cloud Compliance instance

[Deploy Cloud Compliance in Cloud Manager](#) if there isn't already an instance deployed.



Enable Cloud Compliance on Azure NetApp Files

Provision Azure NetApp Files Volumes, Discover them in Cloud Manager, Click **Compliance**, select the **Configuration** tab, and activate compliance scans for all Azure NetApp Files volumes or for individual volumes.



Ensure access to volumes

Now that Cloud Compliance is enabled, ensure that it can access volumes.

- The Cloud Compliance instance needs a network connection to each Azure NetApp Files subnet.
- NFS volume export policies must allow access from the Cloud Compliance instance.
- Cloud Compliance needs Active Directory credentials to scan CIFS volumes.

Click **Cloud Compliance** > **Scan Configuration** > **Edit CIFS Credentials** and provide the credentials. The credentials can be read-only, but providing admin credentials ensures that Cloud Compliance can read data that requires elevated permissions.



Configure volumes to scan

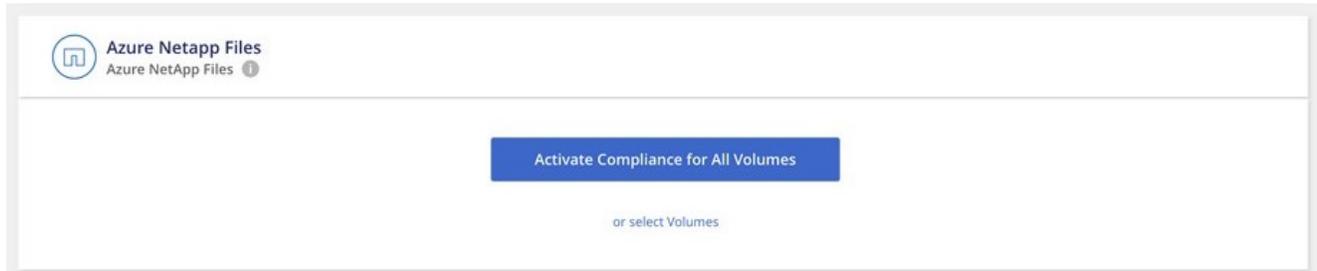
Select the volumes that you'd like to scan and Cloud Compliance will start scanning them.

Deploying the Cloud Compliance instance

[Deploy Cloud Compliance in Cloud Manager](#) if there isn't already an instance deployed.

Enabling Cloud Compliance on Azure NetApp Files

1. If you have not yet provisioned Azure NetApp Files volumes – Get started here:
 - <https://docs.microsoft.com/azure/azure-netapp-files/azure-netapp-files-introduction>
2. Set up discovery of Azure NetApp Files in Cloud Manager - Before you can scan volumes for Azure NetApp Files, Cloud Manager must be [set up to discover the configuration](#).
3. At the top of Cloud Manager, click **Cloud Compliance** and then select the **Configuration** tab.



4. To scan all volumes in Azure NetApp Files, click **Activate Compliance for All Volumes**.

To scan only certain volumes in Azure NetApp Files, click **or select Volumes** and then choose the volumes you want to scan.

See [Enabling and disabling compliance scans on volumes](#) for details.

Result

Cloud Compliance starts scanning the data on the volumes. Results will be available in the Compliance dashboard as soon as Cloud Compliance finishes the initial scans. The time that it takes depends on the amount of data—it could be a few minutes or hours.

Verifying that Cloud Compliance has access to volumes

Make sure that Cloud Compliance can access volumes by checking your networking, security groups, and export policies. You'll need to provide Cloud Compliance with CIFS credentials so it can access CIFS volumes.

Steps

1. Make sure that there's a network connection between the Cloud Compliance instance and Azure NetApp Files.



Cloud Compliance can only scan Azure NetApp Files volumes that are in the same region as Cloud Manager.

2. Ensure that NFS volume export policies include the IP address of the Cloud Compliance instance so it can access the data on each volume.
3. If you use CIFS, provide Cloud Compliance with Active Directory credentials so it can scan CIFS volumes.
 - a. At the top of Cloud Manager, click **Cloud Compliance**.

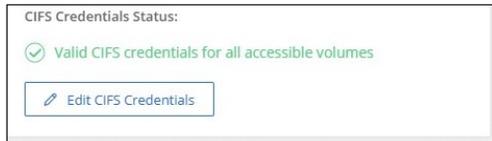
b. Click the **Configuration** tab.



c. For each working environment, click **Edit CIFS Credentials** and enter the user name and password that Cloud Compliance needs to access CIFS volumes on the system.

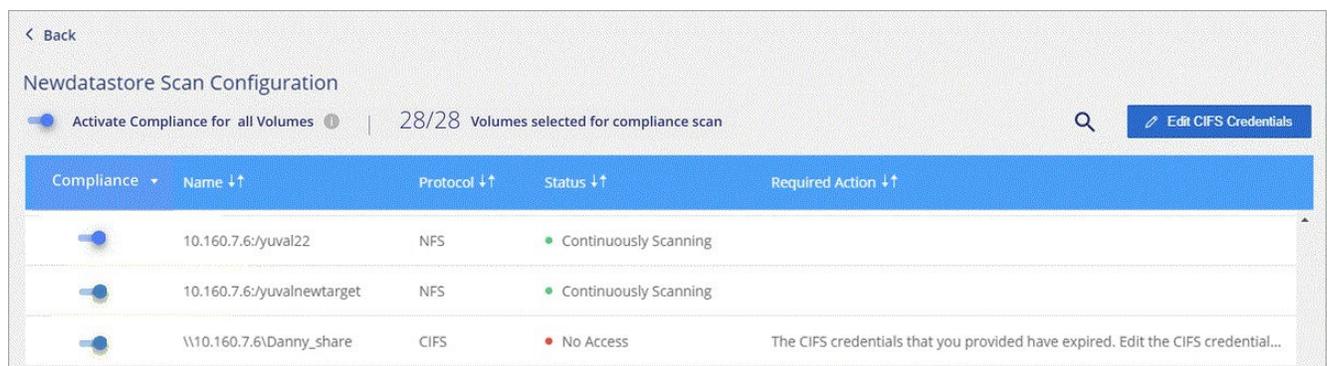
The credentials can be read-only, but providing admin credentials ensures that Cloud Compliance can read any data that requires elevated permissions. The credentials are stored on the Cloud Compliance instance.

After you enter the credentials, you should see a message that all CIFS volumes were authenticated successfully.



4. On the *Scan Configuration* page, click **View Details** to review the status for each CIFS and NFS volume and correct any errors.

For example, the following image shows three volumes; one of which Cloud Compliance can't scan due to network connectivity issues between the Cloud Compliance instance and the volume.



Enabling and disabling compliance scans on volumes

You can stop or start scanning volumes in a working environment at any time from the Scan Configuration page. We recommend that you scan all volumes.

< Back

Newdatastore Scan Configuration

Activate Compliance for all Volumes ⓘ | 27/28 Volumes selected for compliance scan

Compliance ▾	Volume Name ↓↑	Status ▾	Required Action
<input checked="" type="checkbox"/>	VolumeName1	● Not Scanning	Add CIFS Credentials ⓘ
<input checked="" type="checkbox"/>	VolumeName2	● Continuosly Scanning	
<input type="checkbox"/>	VolumeName3	● Not Scanning	
<input checked="" type="checkbox"/>	VolumeName4	● Continuosly Scanning	
<input checked="" type="checkbox"/>	VolumeName5	● Continuosly Scanning	

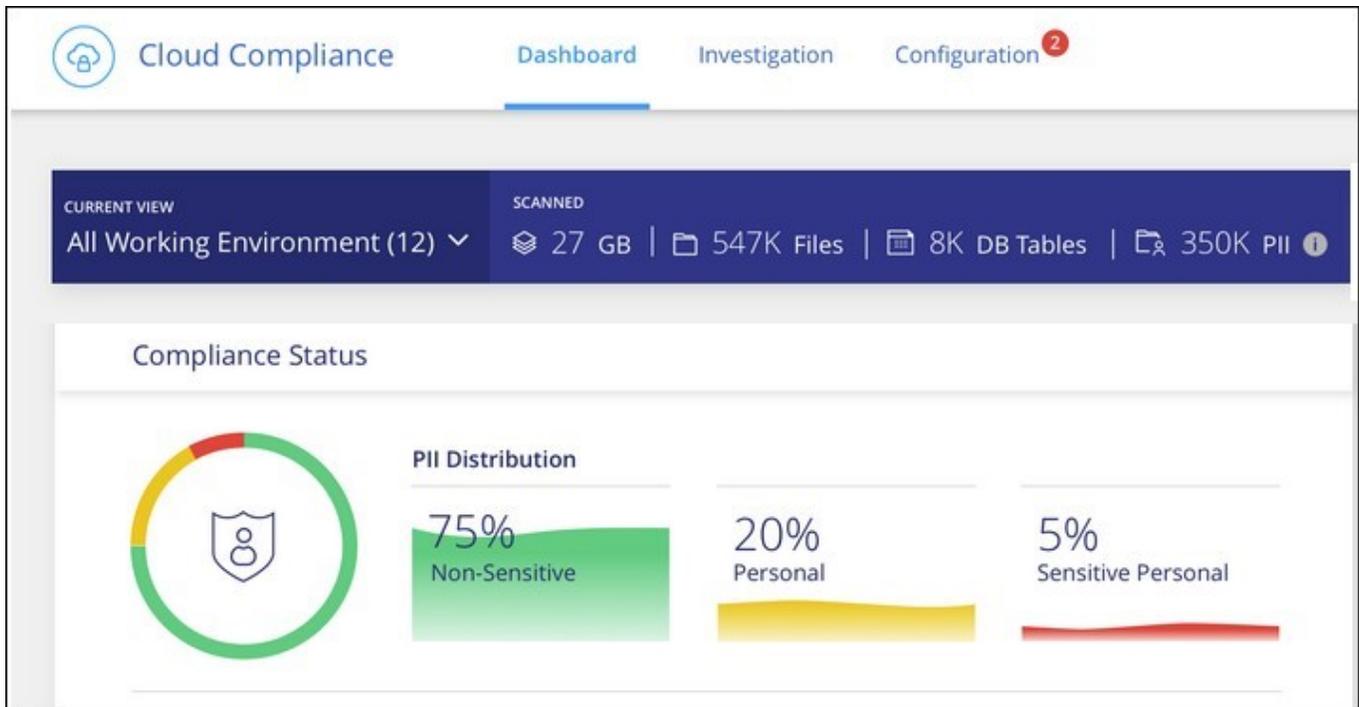
To:	Do this:
Disable scanning for a volume	Move the volume slider to the left
Disable scanning for all volumes	Move the Activate Compliance for all Volumes slider to the left
Enable scanning for a volume	Move the volume slider to the right
Enable scanning for all volumes	Move the Activate Compliance for all Volumes slider to the right



New volumes added to the working environment are automatically scanned only when the **Activate Compliance for all Volumes** setting is enabled. When this setting is disabled, you'll need to activate scanning on each new volume you create in the working environment.

Gaining visibility and control of private data

Gain control of your private data by viewing details about the personal data and sensitive personal data in your organization. You can also gain visibility by reviewing the categories and file types that Cloud Compliance found in your data.



Personal data

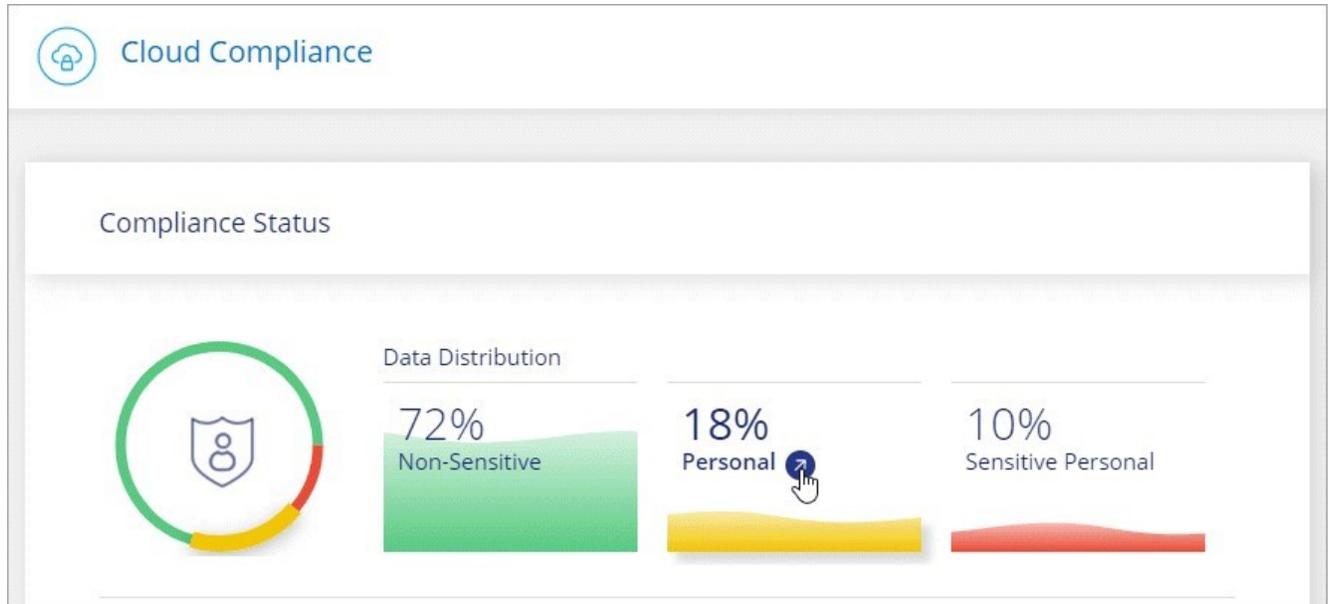
Cloud Compliance automatically identifies specific words, strings, and patterns (Regex) inside the data. For example, Personal Identification Information (PII), credit card numbers, social security numbers, bank account numbers, and more. [See the full list.](#)

For some types of personal data, Cloud Compliance uses *proximity validation* to validate its findings. The validation occurs by looking for one or more predefined keywords in proximity to the personal data that was found. For example, Cloud Compliance identifies a U.S. social security number (SSN) as a SSN if it sees a proximity word next to it—for example, *SSN* or *social security*. [The list below](#) shows when Cloud Compliance uses proximity validation.

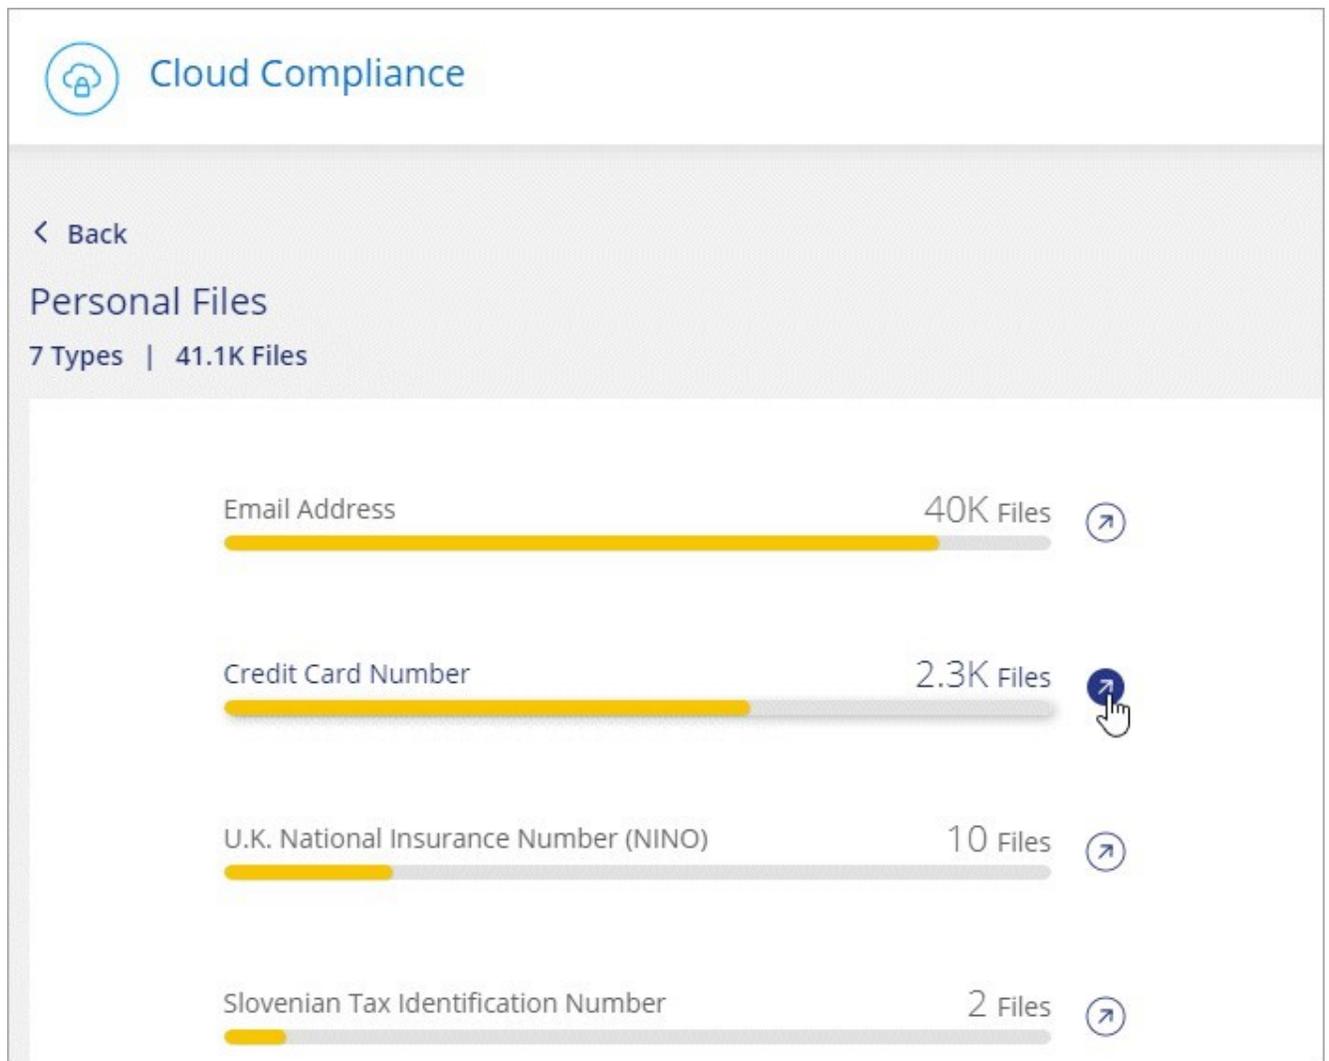
Viewing files that contain personal data

Steps

1. At the top of Cloud Manager, click **Cloud Compliance** and click the **Dashboard** tab.
2. To investigate the details for all personal data, click the icon next to the personal data percentage.



- To investigate the details for a specific type of personal data, click **View All** and then click the **Investigate Results** icon for a specific type of personal data.



Types of personal data

The personal data found in files can be general personal data or national identifiers. The third column identifies whether Cloud Compliance uses [proximity validation](#) to validate its findings for the identifier.

Type	Identifier	Proximity validation?
General	Email address	No
	Credit card number	No
	IBAN number (International Bank Account Number)	No
National Identifiers	Belgian ID (Numero National)	Yes
	Brazilian ID (CPF)	Yes
	Bulgarian ID (UCN)	Yes
	California Driver's License	Yes
	Croatian ID (OIB)	Yes
	Cyprus Tax Identification Number (TIC)	Yes
	Czech/Slovak ID	Yes
	Danish ID (CPR)	Yes
	Dutch ID (BSN)	Yes
	Estonian ID	Yes
	Finnish ID (HETU)	Yes
	French Tax Identification Number (SPI)	Yes
	German Tax Identification Number (Steuerliche Identifikationsnummer)	Yes
	Greek ID	Yes
	Hungarian Tax Identification Number	Yes
	Irish ID (PPS)	Yes
	Israeli ID	Yes
	Italian Tax Identification Number	Yes
	Latvian ID	Yes
	Lithuanian ID	Yes
	Luxembourg ID	Yes
	Maltese ID	Yes
	Polish ID (PESEL)	Yes
	Portuguese Tax Identification Number (NIF)	Yes
	Romanian ID (CNP)	Yes
	Slovenian ID (EMSO)	Yes
	South African ID	Yes
Spanish Tax Identification Number	Yes	
Swedish ID	Yes	
U.K. ID (NINO)	Yes	
USA Social Security Number (SSN)	Yes	

Sensitive personal data

Cloud Compliance automatically identifies special types of sensitive personal information, as defined by privacy regulations such as [articles 9 and 10 of the GDPR](#). For example, information regarding a person's health, ethnic origin, or sexual orientation. [See the full list](#).

Cloud Compliance uses artificial intelligence (AI), natural language processing (NLP), machine learning (ML), and cognitive computing (CC) to understand the meaning of the content that it scans in order to extract entities and categorize it accordingly.

For example, one sensitive GDPR data category is ethnic origin. Because of its NLP abilities, Cloud Compliance can distinguish the difference between a sentence that reads "George is Mexican" (indicating sensitive data as specified in article 9 of the GDPR), versus "George is eating Mexican food."

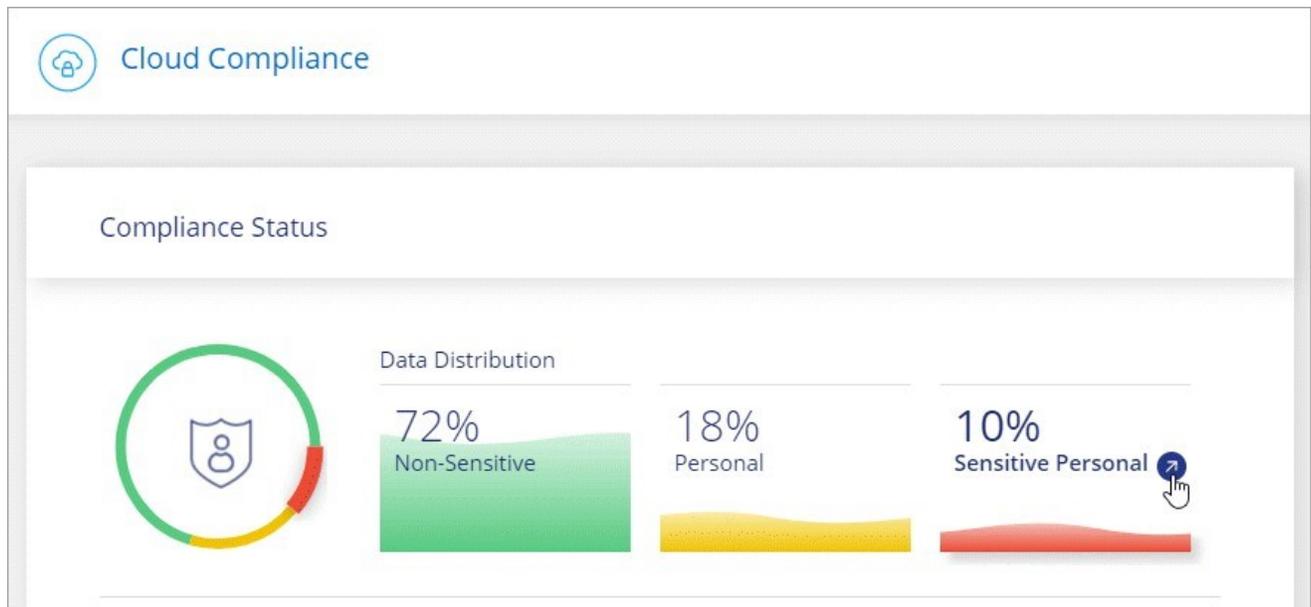


Only English is supported when scanning for sensitive personal data. Support for more languages will be added later.

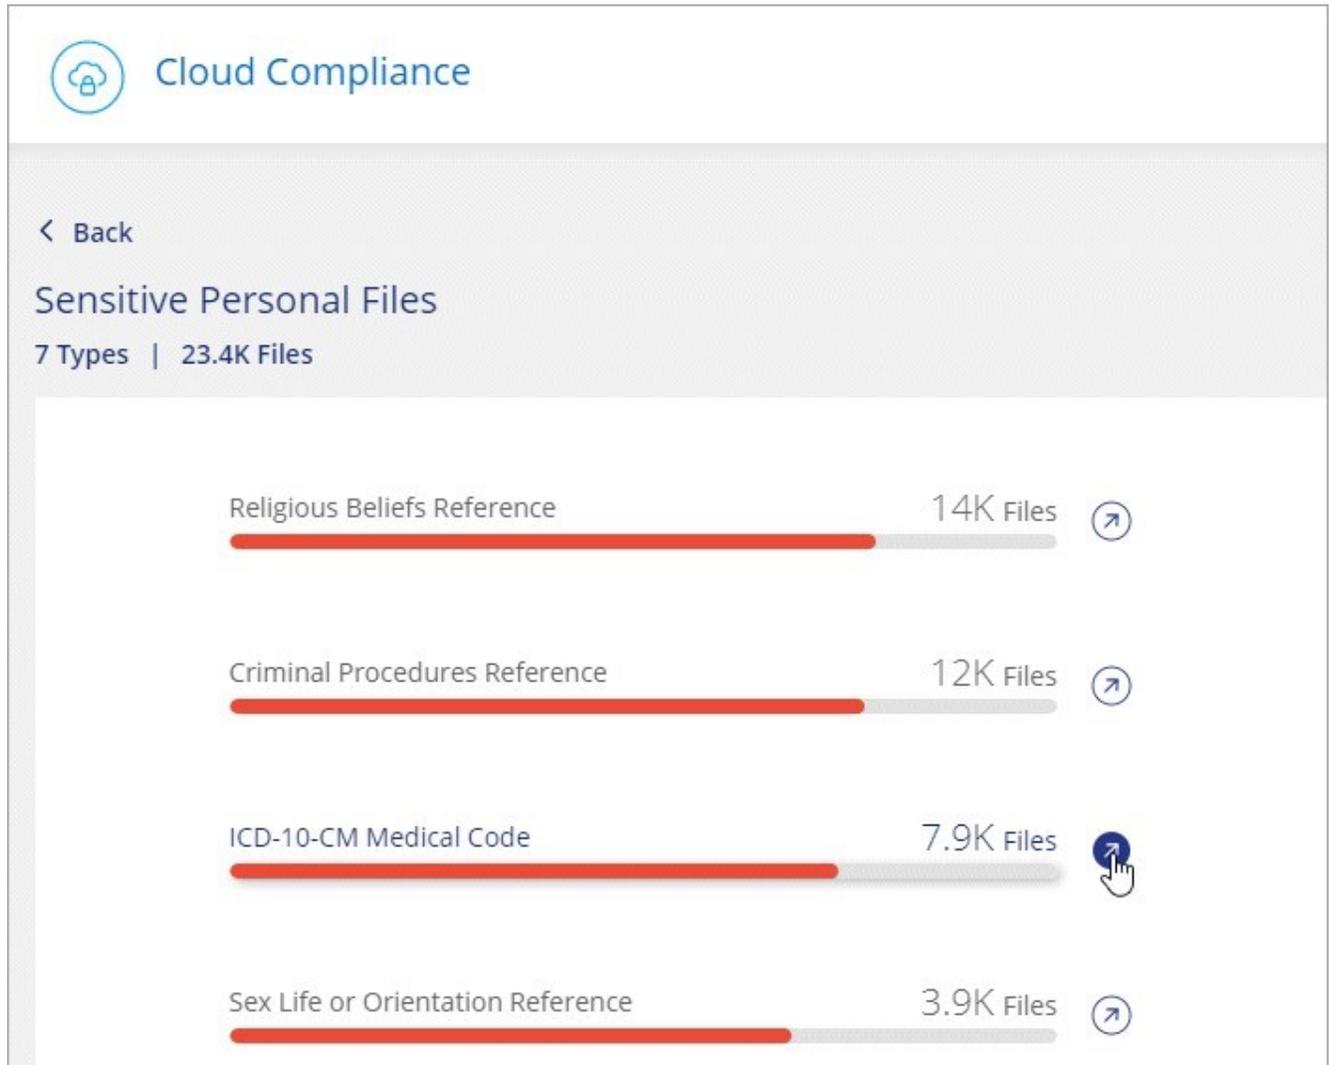
Viewing files that contain sensitive personal data

Steps

1. At the top of Cloud Manager, click **Cloud Compliance**.
2. To investigate the details for all sensitive personal data, click the icon next to the sensitive personal data percentage.



- To investigate the details for a specific type of sensitive personal data, click **View All** and then click the **Investigate Results** icon for a specific type of sensitive personal data.



- Investigate the data by searching, sorting, expanding details for a specific file, clicking **Investigate Results** to see masked information, or by downloading the file list.

Types of sensitive personal data

The sensitive personal data that Cloud Compliance can find in files includes the following:

Criminal Procedures Reference

Data concerning a natural person's criminal convictions and offenses.

Ethnicity Reference

Data concerning a natural person's racial or ethnic origin.

Health Reference

Data concerning a natural person's health.

ICD-9-CM Medical Codes

Codes used in the medical and health industry.

ICD-10-CM Medical Codes

Codes used in the medical and health industry.

Philosophical Beliefs Reference

Data concerning a natural person's philosophical beliefs.

Religious Beliefs Reference

Data concerning a natural person's religious beliefs.

Sex Life or Orientation Reference

Data concerning a natural person's sex life or sexual orientation.

Categories

Cloud Compliance takes the data that it scanned and divides it into different types of categories. Categories are topics based on AI analysis of the content and metadata of each file. [See the list of categories.](#)

Categories can help you understand what's happening with your data by showing you the types of information that you have. For example, a category like resumes or employee contracts can include sensitive data. When you investigate the results, you might find that employee contracts are stored in an insecure location. You can then correct that issue.

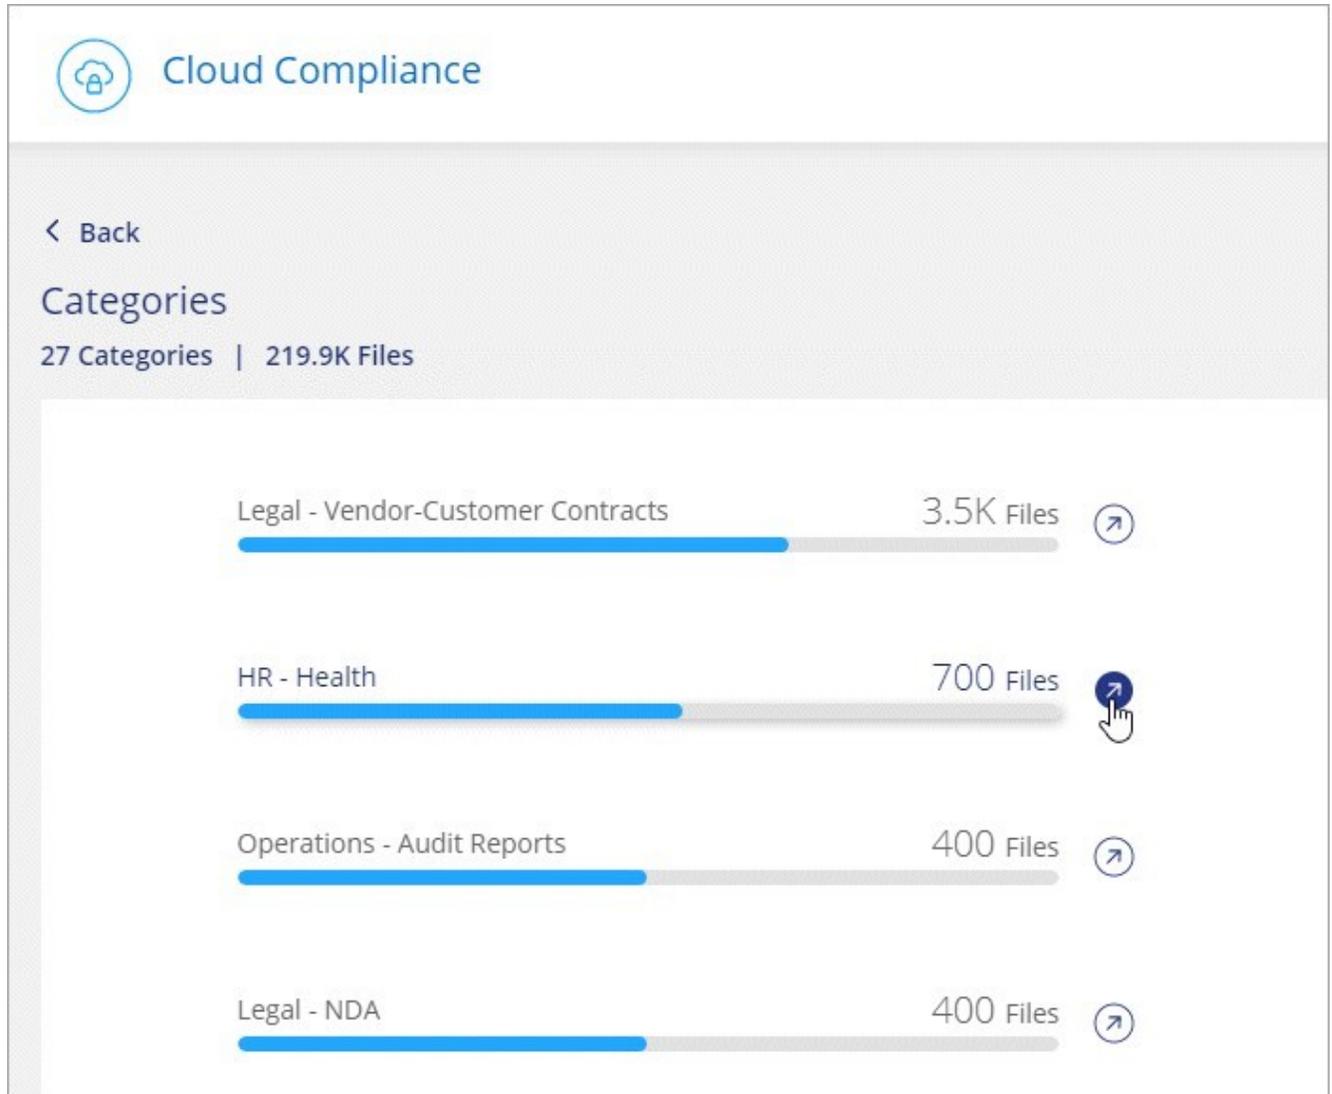


Only English is supported for categories. Support for more languages will be added later.

Viewing files by categories

Steps

1. At the top of Cloud Manager, click **Cloud Compliance**.
2. Click the **Investigate Results** icon for one of the top 4 categories directly from the main screen, or click **View All** and then click the icon for any of the categories.



3. Investigate the data by searching, sorting, expanding details for a specific file, clicking **Investigate Results** to see masked information, or by downloading the file list.

Types of categories

Cloud Compliance categorizes your data as follows:

Finance

- Balance Sheets
- Purchase Orders
- Invoices
- Quarterly Reports

HR

- Background Checks

- Compensation Plans
- Employee Contracts
- Employee Reviews
- Health
- Resumes

Legal

- NDAs
- Vendor-Customer contracts

Marketing

- Campaigns
- Conferences

Operations

- Audit Reports

Sales

- Sales Orders

Services

- RFI
- RFP
- SOW
- Training

Support

- Complaints and Tickets

Metadata categories

- Application Data
- Archive Files
- Audio
- Business Application Data
- CAD Files
- Code
- Database and index files
- Design Files

- Email Application Data
- Executables
- Financial Application Data
- Health Application Data
- Images
- Logs
- Miscellaneous Documents
- Miscellaneous Presentations
- Miscellaneous Spreadsheets
- Videos

File types

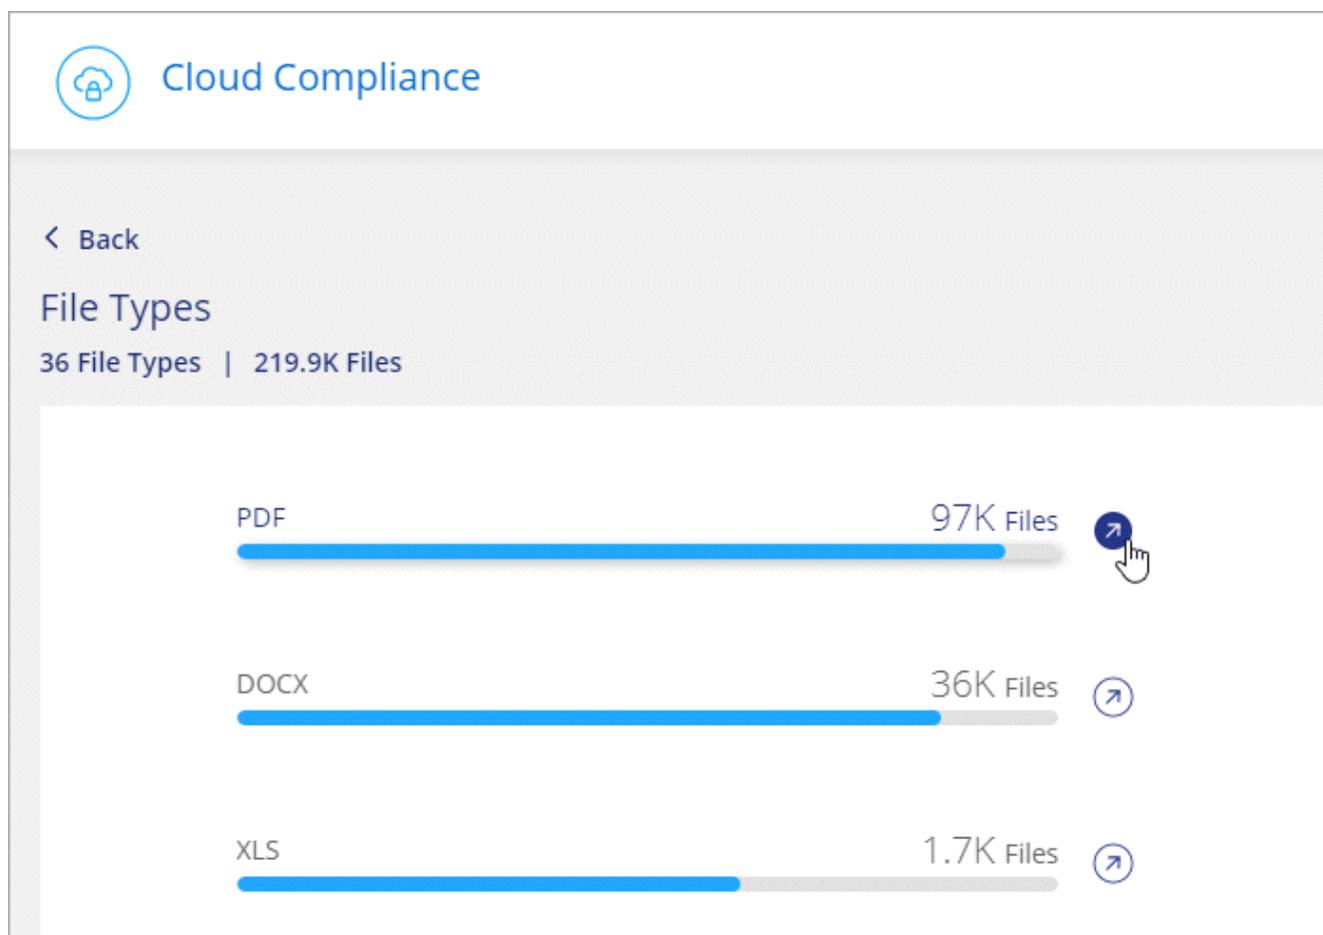
Cloud Compliance takes the data that it scanned and breaks it down by file type. Reviewing your file types can help you control your sensitive data because you might find that certain file types are not stored correctly. [See the list of file types.](#)

For example, you might be storing CAD files that include very sensitive information about your organization. If they are unsecured, you can take control of the sensitive data by restricting permissions or moving the files to another location.

Viewing file types

Steps

1. At the top of Cloud Manager, click **Cloud Compliance**.
2. Click the **Investigate Results** icon for one of the top 4 file types directly from the main screen, or click **View All** and then click the icon for any of the file types.



- Investigate the data by searching, sorting, expanding details for a specific file, clicking **Investigate Results** to see masked information, or by downloading the file list.

Types of files

Cloud Compliance scans all files for category and metadata insights and displays all file types in the file types section of the dashboard.

But when Cloud Compliance detects Personal Identifiable Information (PII), or when it performs a DSAR search, only the following file formats are supported:

.PDF, .DOCX, .DOC, .PPTX, .XLS, .XLSX, .CSV, .TXT, .RTF, and .JSON.

Accuracy of information found

NetApp can't guarantee 100% accuracy of the personal data and sensitive personal data that Cloud Compliance identifies. You should always validate the information by reviewing the data.

Based on our testing, the table below shows the accuracy of the information that Cloud Compliance finds. We break it down by *precision* and *recall*:

Precision

The probability that what Cloud Compliance finds has been identified correctly. For example, a precision rate of 90% for personal data means that 9 out of 10 files identified as containing personal information, actually contain personal information. 1 out of 10 files would be a false positive.

Recall

The probability for Cloud Compliance to find what it should. For example, a recall rate of 70% for personal data means that Cloud Compliance can identify 7 out of 10 files that actually contain personal information in your organization. Cloud Compliance would miss 30% of the data and it won't appear in the dashboard.

Cloud Compliance is in a Controlled Availability release and we are constantly improving the accuracy of our results. Those improvements will be automatically available in future Cloud Compliance releases.

Type	Precision	Recall
Personal data - General	90%-95%	60%-80%
Personal data - Country identifiers	30%-60%	40%-60%
Sensitive personal data	80%-95%	20%-30%
Categories	90%-97%	60%-80%

What's included in each file list report (CSV file)

From each Investigation page you can download file lists (in CSV format) that include details about the identified files. If there are more than 10,000 results, only the top 10,000 appear in the list.

Each file list includes the following information:

- File name
- Location type
- Working environment
- Storage repository
- Protocol
- File path
- File type
- Category
- Personal information
- Sensitive personal information
- Deletion detection date

A deletion detection date identifies the date that the file was deleted or moved. This enables you to identify when sensitive files have been moved. Deleted files aren't part of the file number count that appears in the dashboard or on the Investigation page. The files only appear in the CSV reports.

Viewing compliance reports

Cloud Compliance provides reports that you can use to better understand the status of your organization's data privacy program.



NetApp can't guarantee 100% accuracy of the personal data and sensitive personal data that Cloud Compliance identifies. You should always validate the information by reviewing the data.

Privacy Risk Assessment Report

The Privacy Risk Assessment Report provides an overview of your organization's privacy risk status, as required by privacy regulations such as GDPR and CCPA. The report includes the following information:

Compliance status

A [severity score](#) and the distribution of data, whether it's non-sensitive, personal, or sensitive personal.

Assessment overview

A breakdown of the types of personal data found, as well as the categories of data.

Data subjects in this assessment

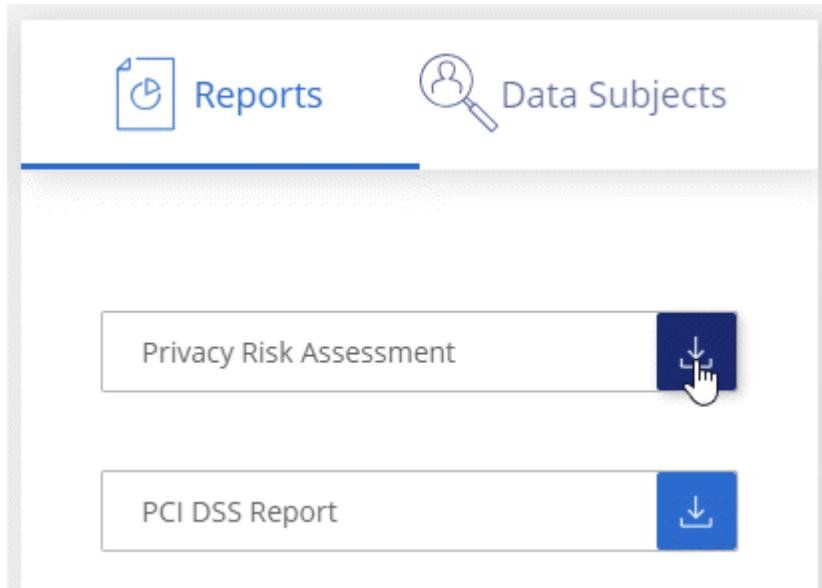
The number of people, by location, for which national identifiers were found.

Generating the Privacy Risk Assessment Report

Go to the Compliance tab to generate the report.

Steps

1. At the top of Cloud Manager, click **Cloud Compliance**.
2. Under **Reports**, click the download icon next to **Privacy Risk Assessment**.



Result

Cloud Compliance generates a PDF report that you can review and send to other groups as needed.

Severity score

Cloud Compliance calculates the severity score for the Privacy Risk Assessment Report on the basis of three variables:

- The percentage of personal data out of all data.
- The percentage of sensitive personal data out of all data.
- The percentage of files that include data subjects, determined by national identifiers such as national IDs, Social Security numbers, and tax ID numbers.

The logic used to determine the score is as follows:

Severity score	Logic
0	All three variables are exactly 0%
1	One of the variables are larger than 0%
2	One of the variables are larger than 3%
3	Two of the variables are larger than 3%
4	Three of the variables are larger than 3%
5	One of the variables are larger than 6%
6	Two of the variables are larger than 6%
7	Three of the variables are larger than 6%

8	One of the variables are larger than 15%
9	Two of the variables are larger than 15%
10	Three of the variables are larger than 15%

PCI DSS Report

The Payment Card Industry Data Security Standard (PCI DSS) Report can help you identify the distribution of credit card information across your files. The report includes the following information:

Overview

How many files contain credit card information and in which working environments.

Retention

The timeframe in which the files were last modified. This is helpful because you shouldn't keep credit card information for longer than you need to process it.

Distribution of Credit Card Information

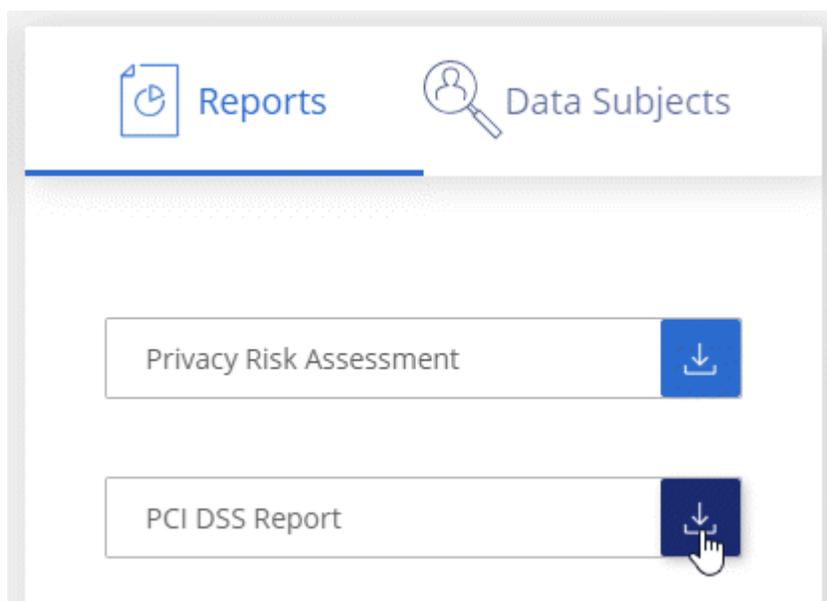
The working environments where the credit card information was found and whether encryption and ransomware protection are enabled.

Generating the PCI DSS Report

Go to the Compliance tab to generate the report.

Steps

1. At the top of Cloud Manager, click **Cloud Compliance**.
2. Under **Reports**, click the download icon next to **PCI DSS Report**.



Result

Cloud Compliance generates a PDF report that you can review and send to other groups as needed.

HIPAA Report

The Health Insurance Portability and Accountability Act (HIPAA) Report can help you identify files containing health information. It is designed to aid in your organization's requirement to comply with HIPAA data privacy laws. The information Cloud Compliance looks for includes:

- Health reference pattern
- ICD-10-CM Medical code
- ICD-9-CM Medical code
- HR – Health category
- Health Application Data category

The report includes the following information:

Overview

How many files contain health information and in which working environments.

Retention

The timeframe in which the files were last modified. This is helpful because you shouldn't keep health information for longer than you need to process it.

Distribution of Health Information

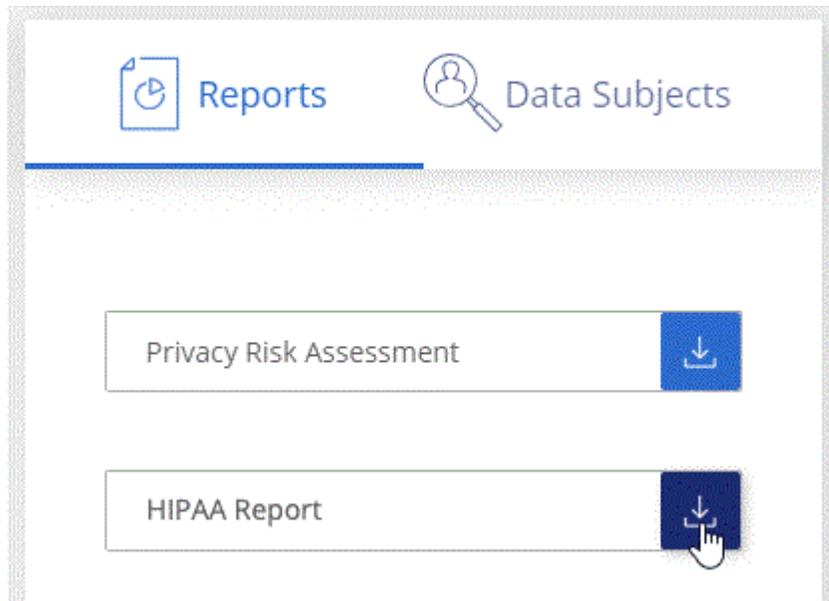
The working environments where the health information was found and whether encryption and ransomware protection are enabled.

Generating the HIPAA Report

Go to the Compliance tab to generate the report.

Steps

1. At the top of Cloud Manager, click **Cloud Compliance**.
2. Under **Reports**, click the download icon next to **HIPAA Report**.



Result

Cloud Compliance generates a PDF report that you can review and send to other groups as needed.

Responding to a Data Subject Access Request

Respond to a Data Subject Access Request (DSAR) by searching for a subject's full name or known identifier (such as an email address) and then downloading a report. The report is designed to aid in your organization's requirement to comply with GDPR or similar data privacy laws.



NetApp can't guarantee 100% accuracy of the personal data and sensitive personal data that Cloud Compliance identifies. You should always validate the information by reviewing the data.

What is a Data Subject Access Request?

Privacy regulations such as the European GDPR grant data subjects (such as customers or employees) the right to access their personal data. When a data subject requests this information, this is known as a DSAR (data subject access request). Organizations are required to respond to these requests "without undue delay," and at the latest within one month of receipt.

How can Cloud Compliance help you respond to a DSAR?

When you perform a data subject search, Cloud Compliance finds all of the files that has that person's name or identifier in it. Cloud Compliance checks the latest pre-indexed data for the name or identifier. It doesn't initiate a new scan.

After the search is complete, you can then download the list of files for a Data Subject Access Request report. The report aggregates insights from the data and puts it into legal terms that you can send back to the person.

Searching for data subjects and downloading reports

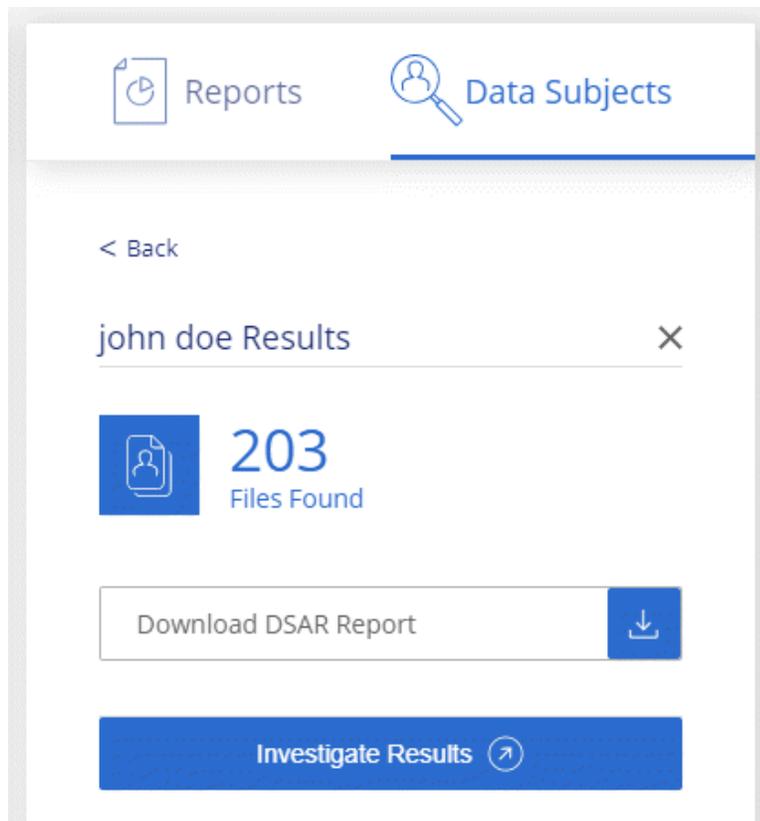
Search for the data subject's full name or known identifier and then download a file list report or DSAR report. You can search by [any personal information type](#).

Only English is supported when searching for the names of data subjects. Support for more languages will be added later.

Steps

1. At the top of Cloud Manager, click **Cloud Compliance**.
2. Click **Data Subjects**.
3. Search for the data subject's full name or known identifier.

Here's an example that shows a search for the name *john doe*:



4. Choose one of the available options:
 - **Download DSAR Report:** A formal response to the access request that you can send to the data subject. This report contains automatically-generated information based on data that Cloud Compliance found on the data subject and is designed to be used as a template. You should complete the form and

review it internally before sending it to the data subject.

- **Investigate Results:** A page that enables you to investigate the data by searching, sorting, expanding details for a specific file, and by downloading the file list.



If there are more than 10,000 results, only the top 10,000 appear in the file list.

Disabling Cloud Compliance

If you need to, you can stop Cloud Compliance from scanning one or more volumes. You can also delete the Cloud Compliance instance if you no longer want to use Cloud Compliance with your working environments.

Deactivating compliance scans for a working environment

When you deactivate scans, Cloud Compliance no longer scans the data on the system and it removes the indexed compliance insights from the Cloud Compliance instance (the data from the working environment itself isn't deleted).

Steps

From the *Scan Configuration* page, click the  button in the row for the working environment, and then click **Deactivate Compliance**.

Deleting the Cloud Compliance instance

You can delete the Cloud Compliance instance if you no longer want to use Cloud Compliance. Deleting the instance also deletes the associated disks where the indexed data resides.

Step

1. Go to your Azure portal and delete the Cloud Compliance instance.

The instance is named *CloudCompliance* with a generated hash (UUID) concatenated to it. For example:
CloudCompliance-16bb6564-38ad-4080-9a92-36f5fd2f71c7

Frequently asked questions about Cloud Compliance

This FAQ can help if you're just looking for a quick answer to a question.

What is Cloud Compliance?

Cloud Compliance is a cloud offering that uses Artificial Intelligence (AI) driven technology to help organizations understand data context and identify sensitive data across your Azure NetApp Files volumes.

Cloud Compliance provides pre-defined parameters (such as sensitive information types and categories) to address new data compliance regulations for data privacy and sensitivity, such as GDPR, CCPA, HIPAA, and more.

Why should I use Cloud Compliance?

Cloud Compliance can empower you with data to help you:

- Comply with data compliance and privacy regulations.
- Comply with data retention policies.
- Easily locate and report on specific data in response to data subjects, as required by GDPR, CCPA, HIPAA, and other data privacy regulations.

What are the common use cases for Cloud Compliance?

- Identify Personal Identifiable Information (PII).
- Identify a wide scope of sensitive information as required by GDPR and CCPA privacy regulations.
- Comply with new and upcoming data privacy regulations.

[Learn more about the use cases for Cloud Compliance.](#)

How do I access Cloud Compliance?

Cloud Compliance is operated and managed through Cloud Manager. You can access Cloud Compliance features from the **Compliance** tab in Cloud Manager.

How does Cloud Compliance work?

Cloud Compliance deploys another layer of Artificial Intelligence alongside your Cloud Manager system and storage systems. It then scans the data on Azure NetApp Files volumes indexes the data insights that are found.

[Learn more about how Cloud Compliance works.](#)

How much does Cloud Compliance cost?

The cost to use Cloud Compliance depends on the amount of data that you're scanning. The first 1 TB of data that Cloud Compliance scans in a Cloud Manager workspace is free. A subscription to the Azure Marketplace is required to continue scanning data after that point. See [pricing](#) for details.

How often does Cloud Compliance scan my data?

Data changes frequently, so Cloud Compliance scans your data continuously with no impact to your data. While the initial scan of your data might take longer, subsequent scans only scan the incremental changes, which reduces system scan times.

[Learn how scans work.](#)

Does Cloud Compliance offer reports?

Yes. The information offered by Cloud Compliance can be relevant to other stakeholders in your organizations, so we enable you to generate reports to share the insights.

The following reports are available for Cloud Compliance:

Privacy Risk Assessment report

Provides privacy insights from your data and a privacy risk score. [Learn more.](#)

Data Subject Access Request report

Enables you to extract a report of all files that contain information regarding a data subject's specific name or personal identifier. [Learn more.](#)

PCI DSS report

Helps you identify the distribution of credit card information across your files. [Learn more.](#)

HIPAA report

Helps you identify the distribution of health information across your files. [Learn more.](#)

Reports on a specific information type

Reports are available that include details about the identified files that contain personal data and sensitive personal data. You can also see files broken down by category and file type. [Learn more.](#)

What type of instance or VM is required for Cloud Compliance?

- Cloud Compliance runs on Azure infrastructure - a Standard_D16s_v3 VM with a 512 GB disk.



Changing or resizing the instance/VM type isn't supported. You need to use the default size that's provided.

[Learn more about how Cloud Compliance works.](#)

Does scan performance vary?

Scan performance can vary based on the network bandwidth and the average file size in your cloud environment.

Which file types are supported?

Cloud Compliance scans all files for category and metadata insights and displays all file types in the file types section of the dashboard.

When Cloud Compliance detects Personal Identifiable Information (PII), or when it performs a DSAR search,

only the following file formats are supported:
.PDF, .DOCX, .DOC, .PPTX, .XLS, .XLSX, .CSV, .TXT, .RTF, and .JSON.

How do I enable Cloud Compliance?

First you need to deploy an instance of Cloud Compliance in Cloud Manager. Once the instance is running, you can enable it on volumes from the **Compliance** tab or by selecting a specific working environment.

[Learn how to get started.](#)



Activating Cloud Compliance results in an immediate initial scan. Compliance results display shortly after.

How do I disable Cloud Compliance?

You can disable Cloud Compliance from the Working Environments page after you select an individual working environment.

[Learn more.](#)



To completely remove the Cloud Compliance instance, you can manually remove the Cloud Compliance instance from the Azure portal.

Can Cloud Compliance send notifications to my organization?

No, but you can download status reports that you can share internally in your organization.

Can I customize the service to my organization's need?

Cloud Compliance provides out-of-the-box insights to your data. These insights can be extracted and used for your organization's needs.

Can I limit Cloud Compliance information to specific users?

Yes, Cloud Compliance is fully integrated with Cloud Manager. Cloud Manager users can only see information for the working environments they are eligible to view according to their workspace privileges.

Additionally, if you want to allow certain users to just view Cloud Compliance scan results without having the ability to manage Cloud Compliance settings, you can assign those users the *Cloud Compliance Viewer* role.

[Learn more.](#)

Copyright Information

Copyright © 2020 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 copyrightowner.

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.