
Amazon Macie

User Guide



Amazon Macie: User Guide

Copyright © 2022 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

- What is Amazon Macie? 1
 - Features of Amazon Macie 1
 - Accessing Amazon Macie 3
 - Pricing for Amazon Macie 4
 - Related services 4
- Getting started 5
 - Before you begin 5
 - Step 1: Enable Amazon Macie 5
 - Step 2: Configure a repository for sensitive data discovery results 6
 - Step 3: Explore sample findings 6
 - Step 4: Create a job to discover sensitive data 7
 - Step 5: Review your findings 7
- Concepts and terminology 9
 - account 9
 - administrator account 9
 - allow list 9
 - AWS Security Finding Format (ASFF) 10
 - classifiable bytes or size 10
 - classifiable object 10
 - custom data identifier 10
 - filter rule 11
 - finding 11
 - finding event 11
 - job 11
 - managed data identifier 12
 - member account 12
 - organization 12
 - policy finding 12
 - sample finding 13
 - sensitive data finding 13
 - sensitive data discovery job 13
 - sensitive data discovery result 13
 - standalone account 14
 - suppressed finding 14
 - suppression rule 14
 - unclassifiable bytes or size 14
 - unclassifiable object 14
- Monitoring Amazon S3 data 16
 - How Macie monitors Amazon S3 data 16
 - Key components 17
 - Data refreshes 18
 - Additional considerations 19
 - Assessing your Amazon S3 security posture 20
 - Displaying the dashboard 21
 - Understanding dashboard components 21
 - Understanding S3 bucket statistics on the dashboard 23
 - Analyzing your Amazon S3 security posture 25
 - Reviewing your S3 bucket inventory 26
 - Filtering your S3 bucket inventory 32
 - Allowing Macie to access S3 buckets and objects 40
- Discovering sensitive data 44
 - Using managed data identifiers 45
 - Keyword requirements 46
 - Sensitive data types: Credentials 47

Sensitive data types: Financial information	48
Sensitive data types: Personal health information (PHI)	52
Sensitive data types: Personally identifiable information (PII)	54
Building custom data identifiers	64
Defining detection criteria	65
Defining severity settings	67
Creating custom data identifiers	68
Regex support	69
Defining sensitive data exceptions with allow lists	70
Allow list options and requirements	71
Creating and managing allow lists	78
Running sensitive data discovery jobs	88
Scope options for jobs	89
Creating a job	97
Monitoring jobs	104
Reviewing job statistics and results	113
Managing jobs	116
Forecasting and monitoring job costs	122
Supported file and storage formats	124
Analyzing encrypted S3 objects	125
Encryption options for S3 objects	125
Allowing Macie to use a customer managed KMS key	127
Storing and retaining sensitive data discovery results	130
Step 1: Verify your permissions	131
Step 2: Choose an AWS KMS key	132
Step 3: Specify an S3 bucket	134
Analyzing findings	140
Types of findings	141
Types of policy findings	141
Types of sensitive data findings	142
Working with sample findings	143
Creating sample findings	143
Reviewing sample findings	144
Suppressing sample findings	146
Reviewing findings	146
Filtering findings	148
Filter fundamentals	148
Creating and applying filters	154
Creating and managing filter rules	160
Fields for filtering findings	165
Investigating sensitive data with findings	183
Locating sensitive data	184
Retrieving sensitive data samples	186
Schema for sensitive data locations	196
Suppressing findings	203
Creating suppression rules	204
Reviewing suppressed findings	206
Changing suppression rules	206
Deleting suppression rules	208
Severity scoring for findings	209
Severity scoring for policy findings	210
Severity scoring for sensitive data findings	210
Monitoring and processing findings	215
EventBridge integration	215
Using EventBridge	216
Creating EventBridge rules for finding events	216
Security Hub integration	219

How Macie publishes findings to Security Hub	220
Examples of Macie findings in Security Hub	223
Enabling and configuring Security Hub integration	226
Stopping the publication of findings to Security Hub	227
Configuring publication settings for findings	227
Choosing publication destinations	227
Determining the publication frequency	228
Changing the publication frequency	229
EventBridge event schema for findings	229
Event schema	229
Event example for a policy finding	230
Event example for a sensitive data finding	233
Managing multiple accounts	238
Administrator and member account relationships	238
Managing accounts with AWS Organizations	240
Considerations and recommendations	241
Integrating and configuring an organization	243
Reviewing organization accounts	249
Managing member accounts	252
Designating a different administrator account	257
Disabling integration with AWS Organizations	259
Managing accounts by invitation	260
Considerations and recommendations	260
Creating and managing an organization	263
Reviewing organization accounts	271
Designating a different administrator account	273
Managing your membership in an organization	274
Forecasting and monitoring costs	278
Understanding how estimated usage costs are calculated	278
Reviewing estimated usage costs	280
Reviewing estimated usage costs on the console	280
Querying estimated usage costs with the API	281
Participating in the free trial	283
Security	285
Data protection	285
Encryption at rest	286
Encryption in transit	286
Identity and access management	286
Audience	287
Authenticating with identities	287
Managing access using policies	289
How Macie works with IAM	291
Identity-based policy examples	296
Service-linked roles	302
AWS managed policies	305
Troubleshooting	308
Logging and monitoring	309
Compliance validation	309
Resilience	310
Infrastructure security	310
VPC endpoints (AWS PrivateLink)	310
Considerations for Macie VPC endpoints	311
Creating an interface VPC endpoint for Macie	311
Logging API calls	312
Macie information in CloudTrail	312
Understanding Macie log file entries	313
Tagging resources	315

Tagging fundamentals	315
Using tags in IAM policies	316
Adding tags to resources	316
Reviewing tags for resources	318
Editing tags for resources	320
Removing tags from resources	322
Creating resources with AWS CloudFormation	324
Macie and AWS CloudFormation templates	324
Learn more about AWS CloudFormation	324
Suspending or disabling Macie	326
Suspending Macie	326
Disabling Macie	327
Quotas	328
AWS glossary	330
Document history	331

What is Amazon Macie?

Amazon Macie is a fully managed data security and data privacy service that uses machine learning and pattern matching to help you discover, monitor, and protect sensitive data in your AWS environment.

Macie automates the discovery of sensitive data, such as personally identifiable information (PII) and financial data, to provide you with a better understanding of the data that your organization stores in Amazon Simple Storage Service (Amazon S3). Macie also provides you with an inventory of your S3 buckets, and it automatically evaluates and monitors those buckets for security and access control. Within minutes, Macie can identify and report overly permissive or unencrypted buckets for your organization.

If Macie detects sensitive data or potential issues with the security or privacy of your data, it creates detailed findings for you to review and remediate as necessary. You can review and analyze these findings directly in Macie, or monitor and process them by using other services, applications, and systems.

Topics

- [Features of Amazon Macie \(p. 1\)](#)
- [Accessing Amazon Macie \(p. 3\)](#)
- [Pricing for Amazon Macie \(p. 4\)](#)
- [Related services \(p. 4\)](#)

Features of Amazon Macie

Here are some of the key ways that Amazon Macie can help you discover, monitor, and protect your sensitive data in Amazon S3.

Automate the discovery of sensitive data

With Macie, you can automate discovery and reporting of sensitive data by [creating and running sensitive data discovery jobs \(p. 88\)](#). A sensitive data discovery job analyzes objects in S3 buckets to determine whether they contain sensitive data. If Macie detects sensitive data in an object, it creates a sensitive data finding for you. The finding provides a detailed report of the sensitive data that Macie found.

You can configure a job to run only once, for on-demand analysis and assessment, or on a recurring basis for periodic analysis, assessment, and monitoring. You can also choose various options to control the breadth and depth of a job's analysis—the S3 buckets to analyze, the sampling depth, and custom include and exclude criteria that derive from properties of S3 objects. With these scheduling and scope options, you can build and maintain a comprehensive view of the data that your organization stores in Amazon S3 and any security or compliance risks for that data.

Discover a variety of sensitive data types

When you create a sensitive data discovery job, you can configure the job to use built-in criteria and techniques, such as machine learning and pattern matching, to analyze objects in S3 buckets. These criteria and techniques, referred to as [managed data identifiers \(p. 45\)](#), can detect a large and growing list of sensitive data types for many countries and regions, including multiple types of personally identifiable information (PII), financial data, and credentials data.

You can also configure the job to use [custom data identifiers \(p. 64\)](#). A custom data identifier is a set of criteria that you define to detect sensitive data—a regular expression (*regex*) that defines a text pattern to match and, optionally, character sequences and a proximity rule that refine the results. With this type of identifier, you can detect sensitive data that reflects your particular scenarios, intellectual property, or proprietary data, and supplement the managed data identifiers that Macie provides.

To fine tune the analysis, a job can also use [allow lists \(p. 70\)](#). Allow lists define specific text and text patterns that you want Macie to ignore in S3 objects—for example, the names of public representatives for your organization, public phone numbers for your organization, or sample data that your organization uses for testing.

Evaluate and monitor data for security and access control

When you enable Macie, Macie immediately generates and begins maintaining a complete inventory of your S3 buckets, and it begins evaluating and monitoring the buckets for security and access control. If Macie detects a potential issue with the security or privacy of a bucket, it creates a [policy finding \(p. 141\)](#) for you.

In addition to specific findings, a [dashboard \(p. 20\)](#) gives you a snapshot of aggregated statistics for your buckets. This includes statistics that indicate how many of your buckets are publicly accessible, are shared with other AWS accounts, or don't encrypt objects by default. You can drill down on each statistic to review the supporting data.

Macie also provides detailed information and statistics for individual buckets in your inventory. This data includes breakdowns of a bucket's public access and encryption settings, and the size and number of objects that Macie can analyze to detect sensitive data in the bucket. You can [browse the inventory \(p. 25\)](#), or sort and filter the inventory by certain fields. When you choose a bucket, a panel displays the bucket's details.

Review and analyze findings

In Macie, a finding is a detailed report of sensitive data in an S3 object or a potential policy-related issue with the security or privacy of an S3 bucket. Each finding provides a severity rating, information about the affected resource, and additional details, such as when and how Macie found the issue.

To [review, analyze, and manage findings \(p. 140\)](#), you can use the **Findings** pages on the Amazon Macie console. These pages list your findings and provide the details of individual findings. They also provide multiple options for grouping, filtering, sorting, and suppressing findings. You can also use the Amazon Macie API to query, retrieve, and suppress findings. If you use the API, you can pass the data to another application, service, or system for deeper analysis, long-term storage, or reporting.

Monitor and process findings with other services and systems

To support integration with other services and systems, Macie [publishes findings to Amazon EventBridge \(p. 215\)](#) as finding events. EventBridge is a serverless event bus service that can route findings data to targets such as AWS Lambda functions and Amazon Simple Notification Service (Amazon SNS) topics. With EventBridge, you can monitor and process findings in near-real time as part of your existing security and compliance workflows.

You can configure Macie to also [publish findings to AWS Security Hub \(p. 219\)](#). Security Hub is a service that provides a comprehensive view of your security posture across your AWS environment and helps you check your environment against security industry standards and best practices. With Security Hub, you can more easily monitor and process your findings as part of a broader analysis of your organization's security posture in AWS.

Centrally manage multiple Macie accounts

If your AWS environment has multiple accounts, you can [centrally manage Macie \(p. 238\)](#) for accounts in your environment. You can do this in two ways, by integrating Macie with AWS Organizations or by sending membership invitations from Macie.

In a multiple-account configuration, a designated Macie administrator can perform certain tasks and access certain Macie settings, data, and resources for accounts that are members of the same organization. Tasks include reviewing information about S3 buckets that are owned by member accounts, reviewing policy findings for those buckets, and running sensitive data discovery jobs to detect sensitive data in the buckets. If the accounts are associated through AWS Organizations, the Macie administrator can also enable Macie for member accounts in the organization.

Develop and manage resources programmatically

In addition to the Amazon Macie console, you can interact with Macie by using the [Amazon Macie API](#). The Amazon Macie API gives you comprehensive, programmatic access to your Macie account and resources.

To develop and manage resources with the Amazon Macie API, you can send HTTPS requests directly to Macie or use a current version of an AWS command line tool or an AWS SDK. AWS provides tools and SDKs that consist of libraries and sample code for various languages and platforms, such as PowerShell, Java, Go, Python, C++, and .NET.

Accessing Amazon Macie

Amazon Macie is available in most AWS Regions. For a list of Regions where Macie is currently available, see [Amazon Macie endpoints and quotas](#) in the *Amazon Web Services General Reference*. To learn more about AWS Regions, see [Managing AWS Regions](#) in the *Amazon Web Services General Reference*.

In each Region, you can work with Macie in any of the following ways.

AWS Management Console

The AWS Management Console is a browser-based interface that you can use to create and manage AWS resources. As part of that console, the Amazon Macie console provides access to your Macie account and resources. You can perform any Macie task by using the Macie console—review statistics and other information about your S3 buckets, run sensitive data discovery jobs, review and analyze findings, and more.

AWS command line tools

With AWS command line tools, you can issue commands at your system's command line to perform Macie tasks and AWS tasks. Using the command line can be faster and more convenient than using the console. The command line tools are also useful if you want to build scripts that perform tasks.

AWS provides two sets of command line tools: the AWS Command Line Interface (AWS CLI) and the AWS Tools for PowerShell. For information about installing and using the AWS CLI, see the [AWS Command Line Interface User Guide](#). For information about installing and using the Tools for PowerShell, see the [AWS Tools for PowerShell User Guide](#).

AWS SDKs

AWS provides SDKs that consist of libraries and sample code for various programming languages and platforms—for example, Java, Go, Python, C++, and .NET. The SDKs provide convenient, programmatic access to Macie and other AWS services. They also handle tasks such as cryptographically signing requests, managing errors, and retrying requests automatically. For information about installing and using the AWS SDKs, see [Tools to Build on AWS](#).

Amazon Macie REST API

The Amazon Macie REST API gives you comprehensive, programmatic access to your Macie account and resources. With this API, you can send HTTPS requests directly to Macie. However, unlike the AWS command line tools and SDKs, use of this API requires your application to handle low-level details such as generating a hash to sign a request. For information about this API, see the [Amazon Macie API Reference](#).

Pricing for Amazon Macie

As with other AWS products, there are no contracts or minimum commitments for using Amazon Macie.

Macie pricing is based on two dimensions—evaluating and monitoring S3 buckets for security and access control, and analyzing S3 objects to discover and report sensitive data in those objects. To help you understand and forecast the cost of using Macie, Macie provides estimated usage costs for your account. You can [review these estimates \(p. 278\)](#) on the Amazon Macie console and access them with the Amazon Macie API.

Depending on how you use the service, you might incur additional costs for using other AWS services in combination with certain Macie features, such as retrieving bucket data from Amazon S3 and using customer managed AWS KMS keys to decrypt objects for analysis. For more information, see [Amazon Macie pricing](#).

When you enable Macie for the first time, your AWS account is automatically enrolled in the 30-day free trial of Macie. This includes individual accounts that are enabled as part of an organization in AWS Organizations. During the free trial, there's no charge for using Macie in the applicable AWS Region to evaluate and monitor your S3 data for security and access control. Note that the free trial doesn't include running sensitive data discovery jobs to discover and report sensitive data in S3 objects.

To help you understand and forecast the cost of using Macie after the free trial ends, Macie provides you with estimated usage costs based on your use of Macie during the trial. Your usage data also indicates the amount of time that remains before your free trial ends. You can [review this data \(p. 283\)](#) on the Amazon Macie console and access it with the Amazon Macie API.

Related services

To further secure your data, workloads, and applications in AWS, consider using the following AWS services in combination with Amazon Macie.

AWS Security Hub

AWS Security Hub gives you a comprehensive view of the security state of your AWS resources and helps you check your AWS environment against security industry standards and best practices. It does this partly by consuming, aggregating, organizing, and prioritizing your security findings from multiple AWS services (including Macie) and supported AWS Partner Network (APN) products. Security Hub helps you analyze your security trends and identify the highest priority security issues across your AWS environment.

To learn more about Security Hub, see the [AWS Security Hub User Guide](#). To learn about using Macie and Security Hub together, see [Amazon Macie integration with AWS Security Hub \(p. 219\)](#).

Amazon GuardDuty

Amazon GuardDuty is a security monitoring service that analyzes and processes certain types of AWS logs, such as AWS CloudTrail data event logs for Amazon S3 and CloudTrail management event logs. It uses threat intelligence feeds, such as lists of malicious IP addresses and domains, and machine learning to identify unexpected and potentially unauthorized and malicious activity within your AWS environment.

To learn more about GuardDuty, see the [Amazon GuardDuty User Guide](#).

To learn about additional AWS security services, see [Security, Identity, and Compliance on AWS](#).

Getting started with Amazon Macie

This tutorial provides an introduction to Amazon Macie.

Tasks

- [Before you begin](#) (p. 5)
- [Step 1: Enable Amazon Macie](#) (p. 5)
- [Step 2: Configure a repository for sensitive data discovery results](#) (p. 6)
- [Step 3: Explore sample findings](#) (p. 6)
- [Step 4: Create a job to discover sensitive data](#) (p. 7)
- [Step 5: Review your findings](#) (p. 7)

Before you begin

When you sign up for Amazon Web Services (AWS), your account is automatically signed up for all AWS services, including Amazon Macie. However, to enable and use Macie, you have to first set up permissions that allow you to access the Amazon Macie console and API operations. You can do this by using the AWS Identity and Access Management (IAM) console to attach the `AmazonMacieFullAccess` managed policy to your IAM identity. To learn more, see [AWS managed policies](#) in the *IAM User Guide*.

Step 1: Enable Amazon Macie

After you set up the required permissions, you can enable Macie. Follow these steps to enable Macie.

To enable Macie

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. By using the AWS Region selector in the upper-right corner of the page, select the Region in which you want to enable Macie.
3. Choose **Get started**.
4. (Optional) When you enable Macie, Macie creates a service-linked role that grants Macie the permissions that it requires to call other AWS services on your behalf. To learn more about this role, see [Service-linked roles for Amazon Macie](#) (p. 302).
5. Choose **Enable Macie**.

Within minutes, Macie generates an inventory of the Amazon Simple Storage Service (Amazon S3) buckets for your account in the current Region. Macie also begins monitoring the buckets for security and access control.

To review your bucket inventory, choose **S3 buckets** in the navigation pane on the console. To then display details about a bucket, choose the bucket's name in the table. The details panel displays statistics and other information that provides insight into the security and privacy of the bucket's data. To learn more about these details, see [Reviewing your S3 bucket inventory](#) (p. 26).

Step 2: Configure a repository for sensitive data discovery results

With Macie, you detect sensitive data by creating and running sensitive data discovery jobs. A sensitive data discovery job analyzes objects in S3 buckets to determine whether the objects contain sensitive data. If Macie discovers sensitive data in an object, Macie creates a *sensitive data finding*. A *sensitive data finding* is a detailed report of sensitive data that Macie found in an object.

Macie also creates a *sensitive data discovery result* for each object that you configure a job to analyze. A *sensitive data discovery result* is a record that logs details about the analysis of an object. This includes objects that don't contain sensitive data, and therefore don't produce sensitive data findings, and objects that Macie can't analyze due to issues such as permissions settings. If an object does contain sensitive data, the sensitive data discovery result includes data from the corresponding sensitive data finding. It provides additional information too.

Macie stores your sensitive data discovery results for 90 days. To access the results and enable long-term storage and retention of them, configure Macie to store the results in an S3 bucket. You must do this within 30 days of enabling Macie. After you do this, the S3 bucket can serve as a definitive, long-term repository for all of your discovery results.

To learn how to configure a repository for your discovery results, see [Storing and retaining sensitive data discovery results](#) (p. 130).

Step 3: Explore sample findings

Macie provides two categories of findings, *policy findings* and *sensitive data findings*. A *finding* is a detailed report of a potential policy violation for an S3 bucket or sensitive data in an S3 object. Macie generates a policy finding when the policies or settings for an S3 bucket are changed in a way that reduces the security or privacy of the bucket and the bucket's objects. Macie generates a sensitive data finding when it discovers sensitive data in an S3 object that you configure a sensitive data discovery job to analyze. Within each category, there are multiple types of findings.

To explore and learn about the different categories and types of findings that Macie can generate, optionally create and review sample findings. Sample findings use example data and placeholder values to demonstrate the kinds of information that Macie might include in each type of finding. Follow these steps to create and review sample findings.

To create and review sample findings

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, choose **Settings**.
3. Under **Sample findings**, choose **Generate sample findings**.

Macie generates one sample finding for each type of finding that Macie supports.

4. In the navigation pane, choose **Findings**.

The **Findings** page displays current findings for your account in the current AWS Region. This includes the sample findings that you created in the preceding step.

5. On the **Findings** page, locate findings whose type begins with **[SAMPLE]**.
6. To review the details of a specific sample finding, choose any field other than the check box for the finding. The details panel displays information for the finding.

For information about each type of finding, see [Types of findings \(p. 141\)](#). For more information about creating and reviewing sample findings, see [Working with sample findings \(p. 143\)](#).

Step 4: Create a job to discover sensitive data

In Macie, you create and run sensitive data discovery jobs to analyze S3 objects and report sensitive data in those objects. To analyze objects, a job can use built-in, managed data identifiers that Macie provides, custom data identifiers that you create, or a combination of the two. For information about the types of S3 objects that Macie can analyze, see [Discovering sensitive data \(p. 44\)](#). For information about the types of sensitive data that Macie can detect, see [Using managed data identifiers \(p. 45\)](#).

Follow these steps to create a job that runs once, immediately after you create it, and uses default settings. To learn how to create a job that runs periodically or uses custom settings, see [Creating a sensitive data discovery job \(p. 97\)](#).

To create a sensitive data discovery job

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, choose **Jobs**.
3. Choose **Create job**.
4. For the **Choose S3 buckets** step, choose **Select specific buckets**.

Macie displays a complete inventory of the S3 buckets for your account in the current AWS Region.

5. Select the check box for each bucket that you want the job to analyze.

Tip

To find specific buckets more easily, enter filter criteria in the filter bar above the table. You can also sort the inventory by choosing a column heading in the table.

6. When you finish selecting buckets, choose **Next**.
7. For the **Review S3 buckets** step, review and verify your bucket selections. Then choose **Next**.
8. For the **Refine the scope** step, choose **One-time job**, and then choose **Next**.
9. For the **Select managed data identifiers** step, choose **All**, and then choose **Next**.
10. For the **Select custom data identifiers** step, choose **Next**.
11. For the **Select allow lists** step, choose **Next**.
12. For the **Enter general settings** step, enter a name and, optionally, a description of the job. Then choose **Next**.
13. For the **Review and create** step, review the job's configuration settings and verify that they're correct.

You can also review the total estimated cost (in US Dollars) of running the job. To learn more about this estimate, see [Forecasting the cost of a sensitive data discovery job \(p. 122\)](#).

14. When you finish reviewing and verifying the job's settings, choose **Submit**.

Macie immediately starts running the job. You can then [monitor and check the status of the job \(p. 119\)](#).

Step 5: Review your findings

Macie automatically monitors your S3 buckets for security and access control, and it creates policy findings to report potential issues with the security or privacy of your buckets. If you create and run a

sensitive data discovery job, Macie also creates sensitive data findings to report sensitive data that it discovers in S3 objects. For more information about findings, see [Analyzing findings \(p. 140\)](#).

Follow these steps to review your findings.

To review your findings

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, choose **Findings**. The **Findings** page displays current findings for your account in the current AWS Region.
3. (Optional) To filter the findings by specific criteria, enter the criteria in the filter bar above the table. To learn more about filters, see [Filtering findings \(p. 148\)](#).
4. To review the details of a specific finding, choose any field other than the check box for the finding. The details panel displays information for the finding.

For more information, including how to group, filter, and download findings, see [Reviewing findings \(p. 146\)](#).

Amazon Macie concepts and terminology

In Amazon Macie, we build on [common AWS concepts and terminology](#) and use these additional terms.

account

A standard AWS account that contains your AWS resources and the identities that can access those resources.

To use Macie, you sign in to AWS with your AWS account credentials, select the AWS Region in which you want to use Macie, and then enable Macie for your AWS account in that Region. For more information, see [Getting started with Amazon Macie \(p. 5\)](#).

There are three types of accounts in Macie:

- **Administrator account** – This type of account manages Macie accounts for an organization. An *organization* is a set of Macie accounts that are associated with each other and centrally managed as a group of related accounts in a specific AWS Region.
- **Member account** – This type of account is associated with and managed by the Macie administrator account for an organization.
- **Standalone account** – This type of account is neither an administrator nor a member account. It isn't part of an organization.

You can add Macie accounts to an organization in two ways: by integrating Macie with AWS Organizations or by sending and accepting Macie membership invitations. For more information, see [Managing multiple accounts \(p. 238\)](#).

administrator account

In Macie, an account that manages Macie accounts for an organization. An *organization* is a set of Macie accounts that are associated with each other and centrally managed as a group of related accounts in a specific AWS Region.

Users of a Macie administrator account have access to Amazon Simple Storage Service (Amazon S3) inventory data, [policy findings \(p. 12\)](#), and certain Macie settings and resources for all the accounts in their organization. They can also run [sensitive data discovery jobs \(p. 13\)](#) to detect sensitive data in S3 buckets that the accounts own. Depending on how an account is designated as an administrator account, they may also be able to perform additional tasks for other accounts in their organization.

For more information, see [Managing multiple accounts \(p. 238\)](#).

allow list

In Macie, an allow list specifies text or a text pattern that you want Macie to ignore when it inspects S3 objects for sensitive data.

You can create two types of allow lists in Macie: a plaintext file that lists specific words and other kinds of character sequences to ignore, or a regular expression (*regex*) that defines a text pattern to ignore. If an object contains text that matches an entry or pattern in an allow list, Macie doesn't report the text in [sensitive data findings \(p. 13\)](#) or [sensitive data discovery results \(p. 13\)](#), even if the text matches the criteria of a [managed data identifier \(p. 12\)](#) or a [custom data identifier \(p. 10\)](#).

For more information, see [Defining sensitive data exceptions with allow lists \(p. 70\)](#).

AWS Security Finding Format (ASFF)

A standardized JSON format for the contents of [findings \(p. 11\)](#) that are published to or generated by AWS Security Hub. The ASFF includes details about the source of a security issue, the affected resources, and the status of a finding.

For information about ASFF, see [AWS Security Finding Format \(ASFF\)](#) in the *AWS Security Hub User Guide*. For information about publishing Macie findings to Security Hub, see [Amazon Macie integration with AWS Security Hub \(p. 219\)](#).

classifiable bytes or size

In the S3 bucket statistics that Macie provides, the total storage size of all the [classifiable objects \(p. 10\)](#) in an S3 bucket.

If versioning is enabled for a bucket, this value is based on the storage size of the latest version of each classifiable object in the bucket. If an object is a compressed file, this value doesn't reflect the actual size of the file's contents after the file is decompressed.

For more information, see [Reviewing your S3 bucket inventory \(p. 26\)](#) and [Assessing your Amazon S3 security posture \(p. 20\)](#).

classifiable object

An S3 object that Macie can analyze to detect sensitive data.

When calculating S3 bucket statistics, Macie determines that an object is *classifiable* based on the object's storage class and file name extension. An object is *classifiable* if it uses a supported Amazon S3 storage class (S3 Intelligent-Tiering, S3 One Zone-IA, S3 Standard, or S3 Standard-IA) and has a file name extension for a supported file or storage format.

For more information, see [Reviewing your S3 bucket inventory \(p. 26\)](#) and [Assessing your Amazon S3 security posture \(p. 20\)](#).

For [sensitive data discovery jobs \(p. 13\)](#), Macie determines that an object is *classifiable* based on the object's storage class, file name extension, and contents. An object is *classifiable* if it uses a supported Amazon S3 storage class, has a file name extension for a supported file or storage format, and Macie verified that it can extract and analyze data from the object.

For more information, see [Discovering sensitive data \(p. 44\)](#) and [Forecasting and monitoring costs \(p. 278\)](#).

custom data identifier

A set of criteria that you define to detect sensitive data.

The criteria consist of a regular expression (*regex*) that defines a text pattern to match and, optionally, character sequences and a proximity rule that refine the results. The character sequences can be:

- *Keywords*, which are words or phrases that must be in proximity of text that matches the regex, or
- *Ignore words*, which are words or phrases to exclude from the results.

In addition to detection criteria, you can define custom severity settings for the [sensitive data findings \(p. 13\)](#) that a custom data identifier produces.

For more information, see [Building custom data identifiers \(p. 64\)](#).

filter rule

A set of attribute-based filter criteria that you create and save to analyze [findings \(p. 11\)](#) on the Amazon Macie console. Filter rules can help you perform consistent analysis of findings that have specific characteristics, such as all high-severity findings that report a specific type of sensitive data.

For more information, see [Creating and managing filter rules for findings \(p. 160\)](#).

finding

A detailed report of sensitive data in an S3 object or a potential issue with the security or privacy of an S3 bucket. Each finding provides details such as a severity rating, information about the affected resource, and when Macie found the data or issue.

Macie generates two categories of findings: [sensitive data findings \(p. 13\)](#), for sensitive data that Macie detects in S3 objects, and [policy findings \(p. 12\)](#), for potential issues with the security or privacy of S3 buckets. Within each category, there are specific types of findings.

For more information, see [Types of Amazon Macie findings \(p. 141\)](#).

finding event

An Amazon EventBridge event that contains the details of a [sensitive data finding \(p. 13\)](#) or [policy finding \(p. 12\)](#).

Macie automatically publishes sensitive data findings and policy findings to Amazon EventBridge as *events*. Each event is a JSON object that conforms to the EventBridge schema for AWS events. You can use these events to monitor, process, and act upon findings by using other applications, services, and systems.

For information about the schema for these events, see [Amazon EventBridge event schema for Amazon Macie findings \(p. 229\)](#). For information about how Macie publishes findings to EventBridge, see [Amazon Macie integration with Amazon EventBridge \(p. 215\)](#).

job

See [sensitive data discovery job \(p. 13\)](#).

managed data identifier

A set of built-in criteria and techniques that are designed to detect a specific type of sensitive data. Examples of sensitive data include credit card numbers, AWS secret access keys, or passport numbers for a particular country or region. These identifiers can detect a large and growing list of sensitive data types for many countries and regions.

For more information, see [Using managed data identifiers \(p. 45\)](#).

member account

A Macie account that's managed by the designated Macie [administrator account \(p. 9\)](#) for an organization. An *organization* is a set of Macie accounts that are associated with each other and centrally managed as a group of related accounts in a specific AWS Region.

An account can become a member account in two ways: by integrating Macie with the account's organization in AWS Organizations or by accepting a Macie membership invitation.

If you have a member account, your Macie administrator has access to Amazon S3 inventory data, [policy findings \(p. 12\)](#), and certain Macie settings and resources for your account. Your administrator can also run [sensitive data discovery jobs \(p. 13\)](#) to detect sensitive data in your S3 buckets. They may also be able to perform additional tasks for your account, depending on how your account became a member account.

For more information, see [Managing multiple accounts \(p. 238\)](#).

organization

A set of Macie accounts that are associated with each other and centrally managed as a group of related accounts in a specific AWS Region.

Each organization consists of a designated Macie [administrator account \(p. 9\)](#) and one or more associated [member accounts \(p. 12\)](#). The administrator account can access certain Macie settings, data, and resources for member accounts. You can create an organization in two ways: by integrating Macie with AWS Organizations or by sending and accepting membership invitations in Macie.

For more information, see [Managing multiple accounts \(p. 238\)](#).

policy finding

A detailed report of a potential policy violation or issue with the security and access control settings for an S3 bucket. The details include a severity rating, information about the affected resource, and when Macie found the issue.

Macie generates policy findings when the policies or settings for an S3 bucket are changed in a way that reduces the security or privacy of the bucket and the bucket's objects. Macie generates these findings as part of its ongoing monitoring activities for your Amazon S3 data. Macie can generate several types of policy findings.

For more information, see [Types of Amazon Macie findings \(p. 141\)](#) and [Monitoring Amazon S3 data \(p. 16\)](#).

sample finding

A [finding](#) (p. 11) that uses example data and placeholder values to demonstrate the kinds of information that a finding might contain.

For more information, see [Working with sample findings](#) (p. 143).

sensitive data finding

A detailed report of sensitive data that Macie found in an S3 object. The details include a severity rating, information about the affected resource, the type and number of occurrences of the sensitive data that Macie found, and when Macie found the sensitive data.

Macie generates sensitive data findings when it discovers sensitive data in S3 objects that you configure a [sensitive data discovery job](#) (p. 13) to analyze. Macie can generate several types of sensitive data findings.

For more information, see [Types of Amazon Macie findings](#) (p. 141) and [Discovering sensitive data](#) (p. 44).

sensitive data discovery job

Also referred to as a *job*, a series of automated processing and analysis tasks that Macie performs to analyze S3 objects and determine whether the objects contain sensitive data. When you create a job, you specify how often you want the job to run, and you define the scope and nature of the job's analysis.

When a job runs, Macie produces records of the sensitive data that it finds ([sensitive data findings](#) (p. 13)) and the analysis that it performs ([sensitive data discovery results](#) (p. 13)). Macie also publishes logging data to Amazon CloudWatch Logs.

For more information, see [Running sensitive data discovery jobs](#) (p. 88).

sensitive data discovery result

A record that logs details about the analysis that Macie performed on an S3 object to determine whether the object contains sensitive data. Macie generates and writes these records to JSON Lines (.jsonl) files, which it encrypts and stores in an S3 bucket that you specify. The records adhere to a standardized schema.

When you run a [sensitive data discovery job](#) (p. 13), Macie creates a sensitive data discovery result for each object that you configured the job to analyze. This includes:

- Objects that contain sensitive data, and therefore also produce [sensitive data findings](#) (p. 13).
- Objects that don't contain sensitive data, and therefore don't produce sensitive data findings.
- Objects that Macie can't analyze due to issues such as permissions settings or use of an unsupported file or storage format.

For more information, see [Reviewing job statistics and results](#) (p. 113).

standalone account

A Macie account that's neither an administrator nor a member account in an [organization \(p. 12\)](#). The account isn't part of an organization.

suppressed finding

A [finding \(p. 11\)](#) that was archived automatically by a [suppression rule \(p. 14\)](#). That is to say, Macie automatically changed the status of the finding to *archived* because the finding matched the criteria of a suppression rule when Macie generated the finding.

For more information, see [Suppressing findings \(p. 203\)](#).

suppression rule

A set of attribute-based filter criteria that you create and save to archive (suppress) [findings \(p. 11\)](#) automatically. Suppression rules are helpful in situations where you've reviewed a class of findings and don't want to be notified of them again.

If you suppress findings with a suppression rule, Macie continues to generate findings that match the rule's criteria. However, Macie automatically changes the status of the findings to *archived*. This means that the findings don't appear by default on the Amazon Macie console and Macie doesn't publish them to other AWS services.

For more information, see [Suppressing findings \(p. 203\)](#).

unclassifiable bytes or size

In the S3 bucket statistics that Macie provides, the total storage size of all the [unclassifiable objects \(p. 14\)](#) in an S3 bucket.

If versioning is enabled for a bucket, this value is based on the storage size of the latest version of each unclassifiable object in the bucket. If an object is a compressed file, this value doesn't reflect the actual size of the file's contents after the file is decompressed.

For more information, see [Reviewing your S3 bucket inventory \(p. 26\)](#) and [Assessing your Amazon S3 security posture \(p. 20\)](#).

unclassifiable object

An S3 object that Macie can't analyze to detect sensitive data.

When calculating S3 bucket statistics, Macie determines that an object is *unclassifiable* based on the object's storage class and file name extension. An object is *unclassifiable* if it doesn't use a supported Amazon S3 storage class (S3 Intelligent-Tiering, S3 One Zone-IA, S3 Standard, or S3 Standard-IA) or doesn't have a file name extension for a supported file or storage format.

For more information, see [Reviewing your S3 bucket inventory \(p. 26\)](#) and [Assessing your Amazon S3 security posture \(p. 20\)](#).

For [sensitive data discovery jobs \(p. 13\)](#), Macie determines that an object is *unclassifiable* based on the object's storage class, file name extension, and contents. An object is *unclassifiable* if it doesn't use a supported Amazon S3 storage class, doesn't have a file name extension for a supported file or storage format, or Macie wasn't able to extract and analyze data from the object. For example, the object is a malformed file.

For more information, see [Discovering sensitive data \(p. 44\)](#) and [Forecasting and monitoring costs \(p. 278\)](#).

Monitoring Amazon S3 data with Amazon Macie

When you enable Amazon Macie for your AWS account, Macie automatically generates and begins maintaining a complete inventory of your Amazon Simple Storage Service (Amazon S3) buckets in the current AWS Region. Macie also begins monitoring and evaluating the buckets for security and access control. If Macie detects an event that reduces the security or privacy of an S3 bucket, Macie creates a [policy finding \(p. 141\)](#) for you to review and remediate as necessary.

To also monitor S3 buckets for the presence of sensitive data, you can create and run [sensitive data discovery jobs \(p. 88\)](#) that analyze bucket objects on a daily, weekly, or monthly basis. If you do this and Macie detects sensitive data in an object, Macie creates a [sensitive data finding \(p. 142\)](#) to notify you of the sensitive data that Macie found.

In addition to findings, Macie provides constant visibility into the security and privacy of your Amazon S3 data. To assess the security posture of your data and determine where to take action, you can use the **Summary** dashboard on the console. This dashboard provides a snapshot of aggregated statistics for your Amazon S3 data. The statistics include data for key security metrics such as the number of buckets that are publicly accessible, don't encrypt new objects by default, or are shared with other AWS accounts. The dashboard also displays groups of aggregated findings data for your account—for example, the names of 1–5 buckets that have the most findings for the preceding seven days. You can drill down on each statistic to view its supporting data. If you prefer to query the statistics programmatically, you can use the [Amazon S3 Data Source Statistics](#) resource of the Amazon Macie API.

For deeper analysis and evaluation, Macie also provides detailed information and statistics for individual buckets in your inventory. This includes breakdowns of each bucket's public access and encryption settings, and the size and number of objects that Macie can analyze to detect sensitive data in the bucket. The inventory also indicates whether any sensitive data discovery jobs are configured to analyze objects in a bucket and, if so, when one of those jobs most recently ran. You can browse, sort, and filter the inventory by using the Amazon Macie console or the [Amazon S3 Data Source](#) resource of the Amazon Macie API.

If you're the Macie administrator for an organization, you can access statistical and other data for S3 buckets that are owned by member accounts in your organization. You can also access policy findings that Macie creates for the buckets, and create sensitive data discovery jobs to detect sensitive data in the buckets. This means that you can use Macie to evaluate and monitor your organization's security posture across your Amazon S3 environment. For more information, see [Managing multiple accounts \(p. 238\)](#).

Topics

- [How Amazon Macie monitors Amazon S3 data \(p. 16\)](#)
- [Assessing your Amazon S3 security posture with Amazon Macie \(p. 20\)](#)
- [Analyzing your Amazon S3 security posture with Amazon Macie \(p. 25\)](#)
- [Allowing Amazon Macie to access S3 buckets and objects \(p. 40\)](#)

How Amazon Macie monitors Amazon S3 data

When you enable Amazon Macie for your AWS account, Macie creates an AWS Identity and Access Management (IAM) [service-linked role \(p. 302\)](#) for your account in the current AWS Region. The

permissions policy for this role allows Macie to monitor AWS resources and call other AWS services on your behalf. By using this role, Macie generates and maintains a complete inventory of your Amazon Simple Storage Service (Amazon S3) buckets in the Region, and Macie monitors and evaluates the buckets for security and access control.

If you're the Macie administrator for an organization, the inventory includes statistical and other data about S3 buckets that are owned by your account and by member accounts in your organization. With this data, you can use Macie to monitor and evaluate your organization's security posture across your Amazon S3 environment. For more information, see [Managing multiple accounts \(p. 238\)](#).

Topics

- [Key components \(p. 17\)](#)
- [Data refreshes \(p. 18\)](#)
- [Additional considerations \(p. 19\)](#)

Key components

Amazon Macie uses a combination of features and techniques to provide and maintain data about your S3 buckets, and to monitor and evaluate the buckets for security and access control.

Gathering metadata and calculating statistics

To generate and maintain metadata and statistics for your bucket inventory, Macie retrieves bucket and object metadata directly from Amazon S3. For each bucket, the metadata includes:

- General information about the bucket, such as the bucket's name, Amazon Resource Name (ARN), creation date, default encryption settings, tags, and the account ID for the AWS account that owns the bucket.
- Account-level permissions settings that apply to the bucket, such as the block public access setting for the account.
- Bucket-level permissions settings for the bucket, such as the block public access setting for the bucket and settings that derive from a bucket policy or access control list (ACL).
- Shared access and replication settings for the bucket, including whether bucket data is replicated to or shared with AWS accounts that aren't part of your organization.
- Object counts and settings for objects in the bucket, such as the number of objects in the bucket and breakdowns of object counts by encryption type, file type, and storage class.

Macie provides this information to you directly. Macie also uses the information to calculate statistics and provide assessments about the security and privacy of your bucket inventory overall and individual buckets in your inventory. For example, you can find the total storage size and number of buckets in your inventory, the total storage size and number of objects in those buckets, and the total storage size and number of objects that Macie can analyze to detect sensitive data in the buckets.

Monitoring bucket security and privacy

To help ensure the accuracy of bucket-level data in your inventory, Macie monitors and analyzes certain [AWS CloudTrail](#) events that can occur for Amazon S3 data. If a relevant event occurs, Macie updates the appropriate inventory data.

For example, if you enable default encryption for a bucket, Macie updates all data about the bucket's default encryption settings. Similarly, if you add or update the bucket policy for a bucket, Macie analyzes the policy and updates the relevant data in your inventory.

Macie monitors and analyzes data for the following CloudTrail events:

- **Account-level events** – DeletePublicAccessBlock and PutPublicAccessBlock
- **Bucket-level events** – CreateBucket, DeleteAccountPublicAccessBlock, DeleteBucket, DeleteBucketEncryption, DeleteBucketPolicy, DeleteBucketPublicAccessBlock, DeleteBucketReplication, DeleteBucketTagging, PutAccountPublicAccessBlock, PutBucketAcl, PutBucketEncryption, PutBucketPolicy, PutBucketPublicAccessBlock, PutBucketReplication, PutBucketTagging, and PutBucketVersioning

You can't enable monitoring for additional CloudTrail events or disable monitoring for any of the preceding events. For detailed information about corresponding operations for the preceding events, see the [Amazon Simple Storage Service API Reference](#).

Tip

To monitor object-level events, we recommend that you use the Amazon S3 protection feature of Amazon GuardDuty. This feature monitors object-level, Amazon S3 data events and analyzes them for malicious and suspicious activity. For more information, see [Amazon S3 protection in Amazon GuardDuty](#) in the *Amazon GuardDuty User Guide*.

Evaluating bucket security and access control

To evaluate bucket-level security and access control, Macie uses automated, logic-based reasoning to analyze resource-based policies that apply to a bucket. Macie also analyzes the account- and bucket-level permissions settings that apply to a bucket. This analysis factors bucket policies, bucket-level ACLs, and block public access settings for the account and the bucket.

For resource-based policies, Macie uses [Zelkova](#). Zelkova is an automated reasoning engine that translates IAM policies into logical statements and runs a suite of general-purpose and specialized logical solvers (*satisfiability modulo theories*) against the decision problem. Macie applies Zelkova repeatedly to a policy with increasingly specific queries to characterize the classes of behaviors that the policy allows. To learn more about the nature of the solvers that Zelkova uses, see [Satisfiability Modulo Theories](#).

Important


To perform the preceding tasks for a bucket, Macie must be allowed to access the bucket. If a bucket's permissions settings prevent Macie from retrieving metadata for the bucket or the bucket's objects, Macie can only provide a subset of information about the bucket, such as the bucket's name and creation date. Macie can't perform any additional tasks for the bucket. For more information, see [Allowing Macie to access S3 buckets and objects](#) (p. 40).

Data refreshes

When you enable Macie for your AWS account, Macie retrieves metadata for your S3 buckets and objects directly from Amazon S3. Thereafter, Macie automatically retrieves both bucket and object metadata directly from Amazon S3 on a daily basis as part of a daily refresh cycle.

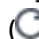
To determine when Macie most recently retrieved both bucket and object metadata for your account as part of the daily refresh cycle, you can refer to the **Last updated** field on the console. This field appears on the **Summary** dashboard and the **S3 buckets** page, and in the [bucket details panel](#) (p. 28). (If you use the Amazon Macie API to query inventory data, the `lastUpdated` field provides this information.) If you're the Macie administrator for an organization, the **Last updated** field indicates the earliest date and time when Macie retrieved the data for an account in your organization.

Macie also retrieves bucket metadata directly from Amazon S3 when any of the following occurs:

- You refresh your inventory data by choosing refresh () on the Amazon Macie console. You can refresh the data as frequently as every five minutes.
- You send a [DescribeBuckets](#) request to the Amazon Macie API programmatically, and you haven't sent a **DescribeBuckets** request within the preceding five minutes.

- Macie detects a relevant AWS CloudTrail event.

Macie can also retrieve the latest object metadata for a specific bucket if you choose to manually refresh that data. This can be helpful if you recently created a bucket or made significant changes to a bucket's objects during the past 24 hours. To manually refresh object metadata for a bucket, choose refresh

 in the **Object statistics** section of the [bucket details panel \(p. 28\)](#) on the console. This feature is available for buckets that contain 30,000 or fewer objects.

Each time Macie retrieves bucket or object metadata, Macie automatically updates all the relevant data in your inventory. If Macie detects differences that affect the security or privacy of a bucket, Macie immediately begins evaluating and analyzing the changes. When the analysis is complete, Macie updates the relevant data in your inventory. If any differences reduce the security or privacy of a bucket, Macie also creates the appropriate [policy findings \(p. 141\)](#) for you to review and remediate as necessary.


On rare occasions under certain conditions, latency and other issues might prevent Macie from retrieving bucket and object metadata. They might also delay notifications that Macie receives about changes to your bucket inventory or the permissions settings and policies for individual buckets. For example, delivery issues with CloudTrail events might cause delays. If this happens, Macie analyzes new and updated data the next time it performs the daily refresh, which is within 24 hours.

Additional considerations

As you use Macie to monitor and assess the security posture of your Amazon S3 data, keep the following in mind:

- Inventory data applies only to S3 buckets in the current AWS Region. To access the data for additional Regions, enable and use Macie in each additional Region.
- If you're the Macie administrator for an organization, you can access inventory data for a member account only if Macie is enabled for that account in the current Region.
- If a bucket's permissions settings prevent Macie from retrieving information about the bucket or the bucket's objects, Macie can't evaluate and monitor the security and privacy of the bucket's data or provide detailed information about the bucket.

To help you identify a bucket where this is the case, Macie does the following:

- In your bucket inventory, Macie displays a warning icon () for the bucket. For the bucket's details, Macie displays only a subset of fields and data: the account ID for the AWS account that owns the bucket; the bucket's name, Amazon Resource Name (ARN), creation date, and Region; and, the date and time when Macie most recently retrieved both bucket and object metadata for the bucket as part of the daily refresh cycle. If you use the Amazon Macie API to query inventory data, Macie provides an error code and message for the bucket and the value for most of the bucket's properties is null.
- On the **Summary** dashboard, the bucket has a value of *Unknown* for the **Public access**, **Encryption**, and **Sharing** statistics. (If you use the Amazon Macie API to query the statistics, the bucket has a value of `unknown` for these statistics.) In addition, Macie excludes the bucket when it calculates data for **Storage** and **Objects** statistics.

To investigate the issue, review the bucket's policy and permissions settings in Amazon S3. For example, the bucket might have a restrictive bucket policy. For more information, see [Allowing Macie to access S3 buckets and objects \(p. 40\)](#).

- Data about access and permissions is limited to account- and bucket-level settings. It doesn't reflect object-level settings that determine access to specific objects in a bucket. For example, if public access is enabled for a specific object in a bucket, Macie doesn't report that the bucket or the bucket's objects are publicly accessible.

To monitor object-level operations and identify potential security risks, we recommend that you use the Amazon S3 protection feature of Amazon GuardDuty. This feature monitors object-level, Amazon

S3 data events and analyzes them for malicious and suspicious activity. For more information, see [Amazon S3 protection in Amazon GuardDuty](#) in the *Amazon GuardDuty User Guide*.

- If you manually refresh object metadata for a specific bucket, Macie temporarily reports *Unknown* for encryption statistics that apply to the objects. The next time Macie performs the daily data refresh (within 24 hours), Macie re-evaluates the encryption metadata for the objects and reports quantitative data for the statistics again.
- In rare cases, Macie might not be able to determine whether a bucket is publicly accessible, is shared with another AWS account, or requires server-side encryption of new objects. For example, a temporary issue might prevent Macie from retrieving and analyzing the requisite data. Or Macie might not be able to fully determine whether one or more policy statements grant access to an external entity. In these cases, Macie reports *Unknown* for the relevant statistics and fields in the inventory. To investigate these cases, review the bucket's policy and permissions settings in Amazon S3.

Also note that Macie generates policy findings only if the security or privacy of a bucket is reduced after you enable Macie for your account. For example, if you disable default encryption for a bucket after you enable Macie, Macie generates a **Policy:IAMUser/S3BucketEncryptionDisabled** finding for the bucket. However, if default encryption was disabled for a bucket when you enabled Macie and default encryption continues to be disabled, Macie doesn't generate a **Policy:IAMUser/S3BucketEncryptionDisabled** finding for the bucket.

In addition, when Macie assesses the security and privacy of a bucket, it doesn't examine access logs or analyze users, roles, and other relevant configurations for accounts. Instead, Macie analyzes and reports data for key settings that indicate *potential* security risks. For example, if a policy finding indicates that a bucket is publicly accessible, it doesn't necessarily mean that an external entity accessed the bucket. Similarly, if a policy finding indicates that a bucket is shared with an AWS account outside your organization, Macie doesn't attempt to determine whether this access is intended and safe. Instead, these findings indicate that an external entity can potentially access the bucket's data, which may be an unintended security risk.

Assessing your Amazon S3 security posture with Amazon Macie

To assess the overall security posture of your Amazon Simple Storage Service (Amazon S3) data and determine where to take action, you can use the **Summary** dashboard on the Amazon Macie console.

The **Summary** dashboard provides a snapshot of aggregated statistics for your Amazon S3 data in the current AWS Region. The statistics include data for key security metrics such as the number of buckets that are publicly accessible or don't encrypt new objects by default. The dashboard also displays groups of aggregated findings data for your account—for example, the types of findings that had the highest number of occurrences during the preceding seven days. If you're the Macie administrator for an organization, the dashboard includes aggregated statistics and data for member accounts in your organization.

To perform deeper analysis, you can drill down and review the supporting data for individual items on the dashboard. You can also [review and analyze your S3 bucket inventory \(p. 25\)](#) by using the Amazon Macie console, or query and analyze inventory data by using the [Amazon S3 Data Source](#) resource of the Amazon Macie API.

Topics

- [Displaying the Summary dashboard \(p. 21\)](#)
- [Understanding components of the Summary dashboard \(p. 21\)](#)
- [Understanding S3 bucket statistics on the Summary dashboard \(p. 23\)](#)

Displaying the Summary dashboard

On the Amazon Macie console, the **Summary** dashboard provides a snapshot of aggregated statistics and findings data for your Amazon S3 data in the current AWS Region. If you prefer to query and review this data programmatically, you can use the [Amazon S3 Data Source Statistics](#) resource of the Amazon Macie API.

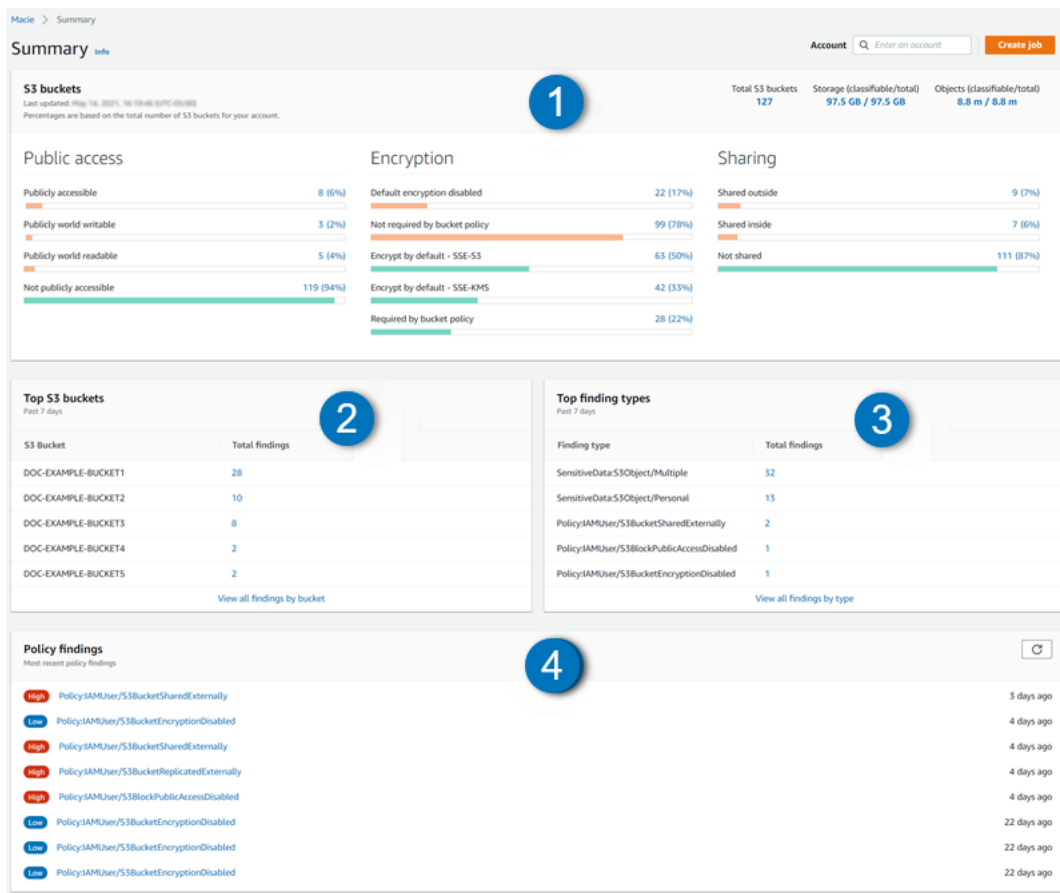
To display the Summary dashboard

1. Open the Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, choose **Summary**. Macie displays the **Summary** dashboard.
3. To determine when Macie most recently retrieved both bucket and object metadata for your account, refer to the **Last updated** field at the top of the dashboard. For more information, see [Data refreshes \(p. 18\)](#).
4. To review the supporting data for an item on the dashboard, choose the item.

If you're the Macie administrator for an organization, the dashboard displays aggregated statistics and data for your account and member accounts in your organization. To filter the dashboard and display data only for a particular account, enter the account's ID in the **Account** box above the dashboard.

Understanding components of the Summary dashboard

On the **Summary** dashboard, statistics and data are organized into four sections, as shown in the following image.



Each section provides insight into key metrics or recent findings data that can help you assess the security and privacy of your Amazon S3 data in the current AWS Region.

1. S3 buckets

This section provides statistics about the amount of data that you store in Amazon S3 and how much of that data Macie can analyze to detect sensitive data. It also provides statistics that indicate potential security and privacy risks for the data. For details about each statistic, see [Understanding S3 bucket statistics on the dashboard](#) (p. 23).

This section also indicates when Macie most recently retrieved bucket and object metadata from Amazon S3 as part of the daily refresh cycle. You can find this information in the **Last updated** field. For more information, see [Data refreshes](#) (p. 18).

2. Top S3 buckets

This section lists the S3 buckets that generated the most findings (of any type) during the preceding seven days, for as many as five buckets. It also indicates the number of findings that Macie created for each bucket.

To display and optionally drill down on all the findings for a bucket for the preceding seven days, choose the value in the **Total findings** field. To display all current findings for all of your buckets, grouped by bucket, choose **View all findings by bucket**.

This section is empty if Macie didn't create any findings during the preceding seven days.

3. Top finding types

This section lists the [types of findings](#) (p. 141) that had the highest number of occurrences during the preceding seven days, for as many as five types of findings. It also indicates the number of findings that Macie created for each type.

To display and optionally drill down on all findings of a particular type for the preceding seven days, choose the value in the **Total findings** field. To display all current findings, grouped by finding type, choose **View all findings by type**.

This section is empty if Macie didn't create any findings during the preceding seven days.

4. Policy findings

This section lists the [policy findings](#) (p. 141) that Macie created or updated most recently, for as many as ten findings. To display the details of a particular finding, choose the finding.

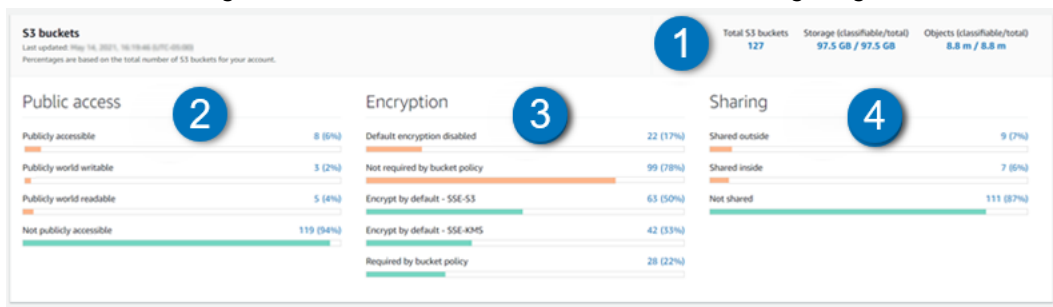
This section is empty if Macie didn't create or update any policy findings during the preceding seven days.

Note that findings data on the **Summary** dashboard doesn't include findings that were suppressed by a [suppression rule](#) (p. 203).

Understanding S3 bucket statistics on the Summary dashboard

The **S3 buckets** section of the **Summary** dashboard provides statistics about the amount of data that you store in Amazon S3 in the current AWS Region and how much of that data Macie can analyze to detect sensitive data. It also provides statistics that can help you identify and investigate potential security risks. For example, you might use this data to identify S3 buckets that are publicly accessible or don't encrypt new objects by default, and then [create a sensitive data discovery job](#) (p. 97) to determine whether the buckets also contain sensitive data.

The statistics are organized into four sections, as shown in the following image.



1. Storage and sensitive data discovery

The statistics at the top of the **S3 buckets** section indicate how much data you store in Amazon S3 and how much of that data Macie can analyze to detect sensitive data:

- **Total S3 buckets** – The total number of S3 buckets, including buckets that don't contain any objects.
- **Storage**
 - **Classifiable** – The total storage size of all the objects that Macie can analyze in the buckets.
 - **Total** – The total storage size of all the objects in the buckets, including objects that Macie can't analyze.

If any of the objects are compressed files, these values don't reflect the actual size of those files after they're decompressed. If versioning is enabled for any of the buckets, these values are based on the storage size of the latest version of each object in those buckets.

- **Objects**

- **Classifiable** – The total number of objects that Macie can analyze in the buckets.
- **Total** – The total number of objects in the buckets, including objects that Macie can't analyze.

In the preceding statistics, data and objects are *classifiable* if they use a supported Amazon S3 storage class (S3 Intelligent-Tiering, S3 One Zone-IA, S3 Standard, or S3 Standard-IA) and have a file name extension for a [supported file or storage format \(p. 124\)](#). You can detect sensitive data in these objects by creating and running sensitive data discovery jobs.

Note that the **Storage** and **Objects** statistics don't include data about objects in buckets that Macie isn't allowed to access. To identify buckets where this is the case, you can [review your bucket inventory \(p. 26\)](#). If the warning icon (⚠) appears next to a bucket's name in your inventory, Macie isn't allowed to access the bucket.

2. Public access

This section indicates how many S3 buckets are or aren't publicly accessible:

- **Publicly accessible** – The number and percentage of buckets that allow the general public to have read or write access to the bucket.
- **Publicly world writable** – The number and percentage of buckets that allow the general public to have write access to the bucket.
- **Publicly world readable** – The number and percentage of buckets that allow the general public to have read access to the bucket.
- **Not publicly accessible** – The number and percentage of buckets that don't allow the general public to have read or write access to the bucket.

To calculate each percentage, Macie divides the number of applicable buckets by the total number of buckets in your bucket inventory.

To determine the values in this section, Macie analyzes a combination of account- and bucket-level settings for each bucket: the block public access setting for the account; the block public access setting for the bucket; the bucket policy for the bucket; and, the access control list (ACL) for the bucket. For information about these settings, see [Identity and access management in Amazon S3](#) and [Blocking public access to your Amazon S3 storage](#) in the *Amazon Simple Storage Service User Guide*.

In certain cases, this section also displays values for *Unknown*. If these values appear, Macie wasn't able to evaluate the public access settings for the specified number and percentage of buckets. For example, a temporary issue or the buckets' permissions settings prevented Macie from retrieving the requisite data. Or Macie wasn't able to fully determine whether one or more policy statements allow an external entity to access the buckets.

3. Encryption

This section indicates how many S3 buckets are or aren't configured to encrypt new objects automatically, and how many S3 buckets do or don't require server-side encryption of objects when objects are uploaded to the buckets:

- **Default encryption disabled** – The number and percentage of buckets that don't encrypt new objects automatically. Default encryption is disabled for these buckets.
- **Not required by bucket policy** – The number and percentage of buckets that don't have bucket policies or have bucket policies that don't require server-side encryption of new objects. For these buckets, [PutObject](#) requests don't have to include a valid server-side encryption header.
- **Encrypt by default – SSE-S3** – The number and percentage of buckets that encrypt new objects automatically using an Amazon S3 managed key. Default encryption is enabled for these buckets.

- **Encrypt by default – SSE-KMS** – The number and percentage of buckets that encrypt new objects automatically using an AWS KMS key. Default encryption is enabled for these buckets.
- **Required by bucket policy** – The number and percentage of buckets whose bucket policies require server-side encryption of new objects. For these buckets, **PutObject** requests must include a valid server-side encryption header. Otherwise, Amazon S3 denies the request.

Note that the totals in this section might exceed the total number of buckets in your inventory. This is because a bucket can have a combination of encryption settings. For example, a bucket might not be configured to encrypt new objects automatically and it might not have a bucket policy that requires server-side encryption of new objects.

To calculate each percentage in this section, Macie divides the number of applicable buckets by the total number of buckets in your bucket inventory.

To determine the values in this section, Macie analyzes the default encryption settings and, if applicable, the bucket policy for each bucket. For information about default encryption settings, see [Setting default server-side encryption behavior for S3 buckets](#) in the *Amazon Simple Storage Service User Guide*. For information about using bucket policies to require server-side encryption of new objects, see [Protecting data using server-side encryption](#) in the *Amazon Simple Storage Service User Guide*.

In certain cases, this section also displays values for *Unknown*. If these values appear, Macie wasn't able to evaluate the default encryption settings or bucket policy for the specified number and percentage of buckets. For example, a temporary issue or the buckets' permissions settings prevented Macie from retrieving the requisite data. Or Macie wasn't able to fully determine whether the buckets' policies require server-side encryption of new objects.

4. Sharing

This section indicates how many S3 buckets are or aren't shared with other AWS accounts:

- **Shared outside** – The number and percentage of buckets that are shared with accounts that aren't in the same organization.
- **Shared inside** – The number and percentage of buckets that are shared with accounts in the same organization.
- **Not shared** – The number and percentage of buckets that aren't shared with other accounts.

To calculate each percentage, Macie divides the number of applicable buckets by the total number of buckets in your bucket inventory.

To determine the values in this section, Macie analyzes the bucket policy and ACL for each bucket. In addition, an *organization* is defined as a set of Macie accounts that are centrally managed as a group of related accounts through AWS Organizations or by Macie invitation.

In certain cases, this section also displays values for *Unknown*. If these values appear, Macie wasn't able to determine whether the specified number and percentage of buckets are shared with another account. For example, a temporary issue or the buckets' permissions settings prevented Macie from retrieving the requisite data. Or Macie wasn't able to fully determine whether the buckets' policies or ACLs are configured to share the buckets with another account.

Analyzing your Amazon S3 security posture with Amazon Macie

To help you perform in-depth analysis and evaluate the security posture of your Amazon Simple Storage Service (Amazon S3) data, Amazon Macie maintains a complete inventory of your S3 buckets in each AWS Region where you use Macie. To learn how Macie maintains this inventory for you, see [How Macie](#)

[monitors Amazon S3 data \(p. 16\)](#). If you're the Macie administrator for an organization, the inventory includes data for S3 buckets that are owned by member accounts in your organization.

By using this inventory, you can review your Amazon S3 data estate, and examine details and statistics for key security settings and metrics that apply to individual S3 buckets. For example, you can access breakdowns of each bucket's public access and encryption settings, and the size and number of objects that Macie can analyze to detect sensitive data in each bucket. You can also determine whether you've configured any sensitive data discovery jobs to analyze data in a bucket. If you have, the inventory indicates when one of those jobs most recently ran.

You can browse, sort, and filter inventory data by using the **S3 buckets** page on the Amazon Macie console. You can also access your inventory data programmatically by using the [Amazon S3 Data Source](#) resource of the Amazon Macie API.

Topics

- [Reviewing your S3 bucket inventory with Amazon Macie \(p. 26\)](#)
- [Filtering your S3 bucket inventory with Amazon Macie \(p. 32\)](#)

Reviewing your S3 bucket inventory with Amazon Macie

On the Amazon Macie console, the **S3 buckets** page provides detailed insight into the security and privacy of your Amazon Simple Storage Service (Amazon S3) data. With this page, you can review and analyze a complete inventory of your S3 buckets in the current AWS Region, and review detailed information and statistics for individual buckets. If you're the Macie administrator for an organization, your inventory includes details and statistics for S3 buckets that are owned by member accounts in your organization.

The **S3 buckets** page also indicates when Macie most recently retrieved both bucket and object metadata for your account as part of the daily refresh cycle. You can find this information in the **Last updated** field at the top of the page. For more information, see [Data refreshes \(p. 18\)](#).

Note that most inventory data is limited to the buckets that Macie is allowed to access for your account. If a bucket's permissions settings prevent Macie from retrieving information about the bucket or the bucket's objects, Macie can only provide a subset of information about the bucket. If this is the case for a particular bucket, Macie displays a warning icon (⚠) and message for the bucket in your bucket inventory. For the bucket's details, Macie displays only a subset of fields and data: the account ID for the AWS account that owns the bucket; the bucket's name, Amazon Resource Name (ARN), creation date, and Region; and, the date and time when Macie most recently retrieved both bucket and object metadata for the bucket as part of the daily refresh cycle. To investigate the issue, review the bucket's policy and permissions settings in Amazon S3. For example, the bucket might have a restrictive bucket policy. For more information, see [Allowing Macie to access S3 buckets and objects \(p. 40\)](#).

If you prefer to access and query your inventory data programmatically, you can use the [Amazon S3 Data Source](#) resource of the Amazon Macie API.

Topics


- [Reviewing your S3 bucket inventory \(p. 26\)](#)
- [Reviewing the details of S3 buckets \(p. 28\)](#)


Reviewing your S3 bucket inventory

The **S3 buckets** page on the Amazon Macie console provides information about your S3 buckets in the current AWS Region. On this page, a table displays summary information for each bucket in your inventory. To customize your view, you can sort and filter the table.

If you choose a bucket in the table, the details panel displays additional information about the bucket. This includes details and statistics for settings and metrics that provide insight into the security and privacy of the bucket's data.

To review your S3 bucket inventory

1. Open the Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, choose **S3 buckets**. The **S3 buckets** page opens and displays the number of buckets in your inventory and a table of the buckets.
3. At the top of the page, optionally choose refresh () to retrieve the latest bucket metadata from Amazon S3.

If the information icon () appears next to any bucket names, we recommend that you do this. This icon indicates that a bucket was created during the past 24 hours, possibly after Macie last retrieved bucket and object metadata from Amazon S3 as part of the [daily refresh cycle \(p. 18\)](#).

4. On the **S3 buckets** page, use the table to review a subset of information about each bucket in your inventory:
 - **Bucket** – The name of the bucket.
 - **Account** – The account ID for the AWS account that owns the bucket.
 - **Classifiable objects** – The total number of objects that Macie can analyze to detect sensitive data in the bucket.
 - **Classifiable size** – The total storage size of all the objects that Macie can analyze to detect sensitive data in the bucket.

Note that this value doesn't reflect the actual size of any compressed objects after they're decompressed. Also, if versioning is enabled for the bucket, this value is based on the storage size of the latest version of each object in the bucket.

- **Monitored** – Whether any sensitive data discovery jobs are configured to periodically analyze objects in the bucket on a daily, weekly, or monthly basis.



If the value for this field is *Yes*, the bucket is explicitly included in a periodic job or the bucket matched the criteria for a periodic job within the past 24 hours. In addition, the status of at least one of those jobs is not *Cancelled*. Macie updates this data on a daily basis.

- **Latest job run** – If any one-time or periodic sensitive data discovery jobs are configured to analyze objects in the bucket, the value for this field indicates the most recent time when one of those jobs started to run. Otherwise, this field is empty.

In the preceding data, objects are *classifiable* if they use a supported Amazon S3 storage class and have a file name extension for a supported file or storage format. You can detect sensitive data in these objects by creating and running a sensitive data discovery job: select the check box for each bucket that contains objects to analyze, and then choose **Create job**. For more information, see [Discovering sensitive data \(p. 44\)](#).

5. To analyze your inventory by using the table, do any of the following:
 - To sort the table by a specific field, click the column heading for the field. To change the sort order, click the column heading again.
 - To filter the table and display only those buckets that have a specific value for a field, place your cursor in the filter bar, and then add a filter condition for the field. To further refine the results, add filter conditions for additional fields. For more information, see [Filtering your S3 bucket inventory \(p. 32\)](#).
6. To review details and statistics for a particular bucket, choose the bucket's name in the table, and then refer to the details panel.

Tip


You can pivot and drill down on many of the fields in the details panel. To show buckets that have the same value for a field, choose  in the field. To show buckets that have other values for a field, choose  in the field.


Reviewing the details of S3 buckets

On the Amazon Macie console, you can use the details panel on the **S3 buckets** page to review statistics and other information about individual S3 buckets in your bucket inventory. This includes details and statistics for settings and metrics that provide insight into the security and privacy of a bucket's data.

For example, you can review breakdowns of a bucket's public access settings, and determine whether a bucket replicates objects or is shared with other AWS accounts. You can also determine whether any sensitive data discovery jobs are configured to inspect the bucket for sensitive data. If there are, you can access details about the job that ran most recently and optionally display any findings that the job produced.

To review the details of an S3 bucket



1. Open the Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, choose **S3 buckets**.
3. On the **S3 buckets** page, optionally choose refresh () to retrieve the latest bucket metadata from Amazon S3.

If the information icon () appears next to any bucket names, we recommend that you do this. This icon indicates that a bucket was created during the past 24 hours, possibly after Macie last retrieved bucket and object metadata from Amazon S3 as part of the [daily refresh cycle \(p. 18\)](#).

4. In the **S3 buckets** table, choose the name of the bucket whose details you want to review. The details panel displays statistics and other information about the bucket.

In the details panel, bucket statistics and information are organized into the following primary sections:

- [Overview \(p. 28\)](#)
- [Object statistics \(p. 29\)](#)
- [Server-side encryption \(p. 30\)](#)
- [Sensitive data discovery \(p. 31\)](#)
- [Public access \(p. 31\)](#)
- [Replication \(p. 31\)](#)
- [Tags \(p. 32\)](#)

As you review the information in each section, you can optionally pivot and drill down on certain fields. To show buckets that have the same value for a field, choose  in the field. To show buckets that have other values for a field, choose  in the field.

Overview

This section provides general information about the bucket, such as the bucket's name, when the bucket was created, and the account ID for the AWS account that owns the bucket. The **Last updated** field indicates when Macie most recently retrieved metadata from Amazon S3 for both the bucket and the bucket's objects as part of the [daily refresh cycle \(p. 18\)](#).

Of special note, the **Shared access** field indicates whether the bucket is shared with other AWS accounts and, if so, whether those accounts are internal to (part of) or external to (not part of) your organization. An *organization* is a set of Macie accounts that are centrally managed as a group of related accounts through AWS Organizations or by Macie invitation. To determine the value for this field, Macie analyzes the bucket policy and access control list (ACL) for the bucket. Note that this data is limited to bucket-level settings. It doesn't reflect any object-level settings for sharing specific objects with another account.

Object statistics

This section provides information about the objects in the bucket, starting with the total number of objects in the bucket, the total storage size of all those objects, and the total storage size of all the objects that are compressed (.gz, .gzip, or .zip) files. If versioning is enabled for the bucket, the size values are based on the size of the latest version of each object in the bucket.

If you recently created the bucket or made significant changes to the bucket's objects during the past 24 hours, optionally choose refresh (🔄) to retrieve the latest metadata for the bucket's objects. Macie displays the information icon (ℹ️) to help you determine whether this might be the case. The refresh option is available if a bucket contains 30,000 or fewer objects.

Note

If you refresh object metadata for a bucket, Macie temporarily reports *Unknown* for encryption statistics that apply to the objects. Macie will re-evaluate and update the data for these statistics when it performs the next daily refresh of bucket and object metadata, which is within 24 hours.

Additional statistics in this section can help you assess how much data Macie can analyze to detect sensitive data in the bucket.

Classifiable objects

This section indicates the total number of objects that Macie can analyze to detect sensitive data and the total storage size of those objects. These objects use a supported Amazon S3 storage class (S3 Intelligent-Tiering, S3 One Zone-IA, S3 Standard, or S3 Standard-IA) and have a file name extension for a supported file or storage format. This means that you can detect sensitive data in the objects by creating and running a sensitive data discovery job. For more information, see [Discovering sensitive data](#) (p. 44).

Note that the value in the **Total storage size** field doesn't reflect the actual size of any compressed objects after they're decompressed. Also, if versioning is enabled for the bucket, this value is based on the storage size of the latest version of each object in the bucket.

Unclassifiable objects

This section indicates the total number of objects that Macie can't analyze to detect sensitive data and the total storage size of those objects. These objects don't use a supported Amazon S3 storage class (S3 Intelligent-Tiering, S3 One Zone-IA, S3 Standard, or S3 Standard-IA) or don't have a file name extension for a [supported file or storage format](#) (p. 124).

Note that the value in the **Total storage size** field doesn't reflect the actual size of any compressed objects after they're decompressed. Also, if versioning is enabled for the bucket, this value is based on the storage size of the latest version of each object in the bucket.

Unclassifiable objects: Storage class

This section provides a breakdown of the number and storage size of the objects that Macie can't analyze because the objects don't use a supported Amazon S3 storage class.

The value in the **Storage size** field doesn't reflect the actual size of any compressed objects after they're decompressed. Also, if versioning is enabled for the bucket, this value is based on the storage size of the latest version of each applicable object in the bucket.

Unclassifiable objects: File type

This section provides a breakdown of the number and storage size of the objects that Macie can't analyze because the objects don't have a file name extension for a supported file or storage format.

The value in the **Storage size** field doesn't reflect the actual size of any compressed objects after they're decompressed. Also, if versioning is enabled for the bucket, this value is based on the storage size of the latest version of each applicable object in the bucket.

Objects by encryption type

This section provides a breakdown of the number of objects that use each type of encryption that Amazon S3 supports:

- **Customer managed** – The number of objects that are encrypted with a customer-provided key. These objects use SSE-C encryption.
- **SSE-KMS managed** – The number of objects that are encrypted with an AWS KMS key, either an AWS managed KMS key or a customer managed KMS key. These objects use SSE-KMS encryption.
- **SSE-S3 managed** – The number of objects that are encrypted with an Amazon S3 managed key. These objects use SSE-S3 encryption.
- **No encryption** – The number of objects that aren't encrypted or use client-side encryption. (If an object is encrypted using client-side encryption, Macie can't access and report encryption data for the object.)
- **Unknown** – The number of objects that Macie doesn't have current encryption metadata for. This typically occurs if you recently chose to manually refresh the metadata for the bucket's objects. Macie will update the encryption statistics when it performs the next daily refresh of bucket and object metadata, which is within 24 hours.

For information about each supported encryption type, see [Protecting data using encryption](#) in the *Amazon Simple Storage Service User Guide*.

Server-side encryption

This section provides insight into the server-side encryption settings for the bucket.

The **Encryption required by bucket policy** field indicates whether the bucket's policy requires server-side encryption of objects when objects are uploaded to the bucket:

- **No** – The bucket doesn't have a bucket policy or the bucket's policy doesn't require server-side encryption of new objects. If a bucket policy exists, it doesn't require [PutObject](#) requests to include a valid server-side encryption header.
- **Yes** – The bucket's policy requires server-side encryption of new objects. [PutObject](#) requests for the bucket must include a valid server-side encryption header. Otherwise, Amazon S3 denies the request.
- **Unknown** – Macie wasn't able to evaluate the bucket policy to determine whether it requires server-side encryption of new objects.

For this assessment, valid server-side encryption headers are: `x-amz-server-side-encryption` with a value of `AES256` or `aws:kms`, and `x-amz-server-side-encryption-customer-algorithm` with a value of `AES256`. For information about using bucket policies to require server-side encryption of new objects, see [Protecting data using server-side encryption](#) in the *Amazon Simple Storage Service User Guide*.

The **Default encryption** field indicates whether default encryption is enabled for the bucket and, if so, the type of server-side encryption that's used:

- **AES256** – New objects are encrypted automatically with an Amazon S3 managed key. Default encryption is enabled for the bucket and it uses SSE-S3 encryption.

- **aws:kms** – New objects are encrypted automatically with an AWS KMS key, either an AWS managed KMS key or a customer managed KMS key. Default encryption is enabled for the bucket and it uses SSE-KMS encryption. The **KMS key** field shows the Amazon Resource Name (ARN) or unique identifier (key ID) for the KMS key that's used.
- **None** – New objects aren't encrypted automatically. Default encryption is disabled for the bucket.

For information about configuring default encryption settings, see [Setting default server-side encryption behavior for S3 buckets](#) in the *Amazon Simple Storage Service User Guide*.

Sensitive data discovery

This section indicates whether any periodic sensitive data discovery jobs are configured to inspect the bucket for sensitive data on a daily, weekly, or monthly basis. If the value for the **Actively monitored by job** field is *Yes*, the bucket is explicitly included in a periodic job or the bucket matched the criteria for a periodic job within the past 24 hours. In addition, the status of at least one of those jobs is not *Cancelled*. Macie updates this data on a daily basis.

If any type of sensitive data discovery job (either a periodic job or a one-time job) is configured to inspect the bucket, the **Latest job** field provides the unique identifier for the job that most recently started to run. The **Latest job run** field indicates when that job started to run.

Tip

To display all the sensitive data findings that the job produced, choose the link in the **Latest job** field. In the job details panel that appears, choose **Show results** at the top of the panel, and then choose **Show findings**.

Public access

This section indicates whether the bucket is publicly accessible, and it provides a breakdown of the various account- and bucket-level settings that determine whether the bucket is publicly accessible. The **Effective permission** field indicates the cumulative result of these settings:

- **Not public** – The bucket isn't publicly accessible.
- **Public** – The bucket is publicly accessible.
- **Unknown** – Macie wasn't able to evaluate all the public access settings for the bucket.

Note that this data is limited to account- and bucket-level settings. It doesn't reflect object-level settings that enable public access to specific objects in a bucket.

To learn about Amazon S3 settings for managing public access to buckets and bucket data, see [Identity and access management in Amazon S3](#) and [Blocking public access to your Amazon S3 storage](#) in the *Amazon Simple Storage Service User Guide*.

Replication

In this section, the **Replicated** field indicates whether the bucket is configured to replicate objects to buckets that are owned by other AWS accounts. If the bucket is configured to do this, this section also lists the account IDs for those accounts.

The **Replicated externally** field indicates whether bucket objects are replicated to AWS accounts that are external to (aren't part of) your organization. An *organization* is a set of Macie accounts that are centrally managed as a group of related accounts through AWS Organizations or by Macie invitation.

To learn about Amazon S3 options and settings for replicating bucket objects, see [Replicating objects](#) in the *Amazon Simple Storage Service User Guide*.

Tags

If tags are associated with the bucket, this section appears in the panel and lists those tags. Tags are labels that you can define and assign to certain types of AWS resources, including S3 buckets. Each tag consists of a required tag key and an optional tag value.

To learn about tagging buckets, see [Using cost allocation S3 bucket tags](#) in the *Amazon Simple Storage Service User Guide*.

Filtering your S3 bucket inventory with Amazon Macie

To identify and focus on buckets that have specific characteristics, you can filter your S3 bucket inventory on the Amazon Macie console and in queries that you submit programmatically using the Amazon Macie API. When you create a filter, you use specific bucket attributes to define criteria for including or excluding buckets from a view or from query results. A *bucket attribute* is a field that stores specific metadata for a bucket.

In Macie, a filter consists of one or more conditions. Each condition, also referred to as a *criterion*, consists of three parts:

- An attribute-based field, such as **Bucket name**, **Tag key**, or **Defined in job**.
- An operator, such as *equals* or *not equals*.
- One or more values. The type and number of values depends on the field and operator that you choose.


How you define and apply filter conditions depends on whether you use the Amazon Macie console or the Amazon Macie API.

Topics

- [Filtering your inventory on the Amazon Macie console \(p. 32\)](#)
- [Filtering your inventory programmatically with the Amazon Macie API \(p. 34\)](#)

Filtering your inventory on the Amazon Macie console

If you use the Amazon Macie console to filter your bucket inventory, Macie provides options to help you choose fields, operators, and values for individual conditions. You access these options by using the filter bar on the **S3 buckets** page, as shown in the following image.




When you place your cursor in the filter bar, Macie displays a list of fields that you can use in filter conditions. The fields are organized by logical category. For example, the **Common fields** category includes fields that store general information about a bucket, and the **Public access** category includes fields that store data about the various types of public access settings that can apply to a bucket. The fields are sorted alphabetically within each category.

To add a condition, start by choosing a field from the list. To find a field, browse the complete list, or enter part of the field's name to narrow the list of fields.

Depending on the field that you choose, Macie displays different options. The options reflect the type and nature of the field that you choose. For example, if you choose the **Defined in job** field, Macie displays a list of values to choose from. If you choose the **Bucket name** field, Macie displays a text box in which you can enter a bucket name. Whichever field you choose, Macie guides you through the steps to add a condition that includes the required settings for the field.

After you add a condition, Macie applies the condition's criteria and adds the condition to a filter box in the filter bar, as shown in the following image.




In this example, the condition is configured to include all buckets that are publicly accessible, and to exclude all other buckets. It returns buckets where the value for the **Effective permission** field *equals* **Public**.

As you add more conditions, Macie applies their criteria and adds them to the filter bar. If you add multiple conditions, Macie uses AND logic to join the conditions and evaluate the filter criteria. This means that a bucket meets the filter criteria only if it matches all the conditions in the filter.

You can refer to the filter bar at any time to see which criteria you've applied.

To filter your inventory by using the console

1. Open the Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, choose **S3 buckets**. The **S3 buckets** page opens and displays the number of buckets in your inventory and a table of the buckets.
3. To retrieve the latest bucket metadata from Amazon S3, choose refresh () at the top of the page.
4. Place your cursor in the filter bar, and then choose the field to use for the condition.
5. Choose or enter the appropriate type of value for the field, keeping the following tips in mind.

Dates, times, and time ranges

For dates and times, use the **From** and **To** boxes to define an inclusive time range:

- To define a fixed time range, use the **From** and **To** boxes to specify the first date and time and the last date and time in the range, respectively.
- To define a relative time range that starts at a certain date and time and ends at the current time, enter the start date and time in the **From** boxes, and delete any text in **To** boxes.
- To define a relative time range that ends at a certain date and time, enter the end date and time in the **To** boxes, and delete any text in the **From** boxes.

Note that time values use 24-hour notation. If you use the date picker to choose dates, you can refine the values by entering text directly in the **From** and **To** boxes.

Numbers and numeric ranges

For numeric values, use the **From** and **To** boxes to enter integers that define an inclusive numeric range:

- To define a fixed numeric range, use the **From** and **To** boxes to specify the lowest and highest numbers in the range, respectively.
- To define a fixed numeric range that's limited to one specific value, enter the value in both the **From** and **To** boxes. For example, to include only those buckets that contain exactly 15 objects, enter **15** in the **From** and **To** boxes.
- To define a relative numeric range that starts at a certain number, enter the number in the **From** box, and don't enter any text in the **To** box.

- To define a relative numeric range that ends at a certain number, enter the number in the **To** box, and don't enter any text in the **From** box.

Text (string) values

For this type of value, enter a complete, valid value for the field. Values are case sensitive.

Note that you can't use a partial value or wildcard characters in this type of value. The only exception is the **Bucket name** field. For that field, you can specify a prefix instead of a complete bucket name. For example, to find all S3 buckets whose names begin with *my-S3*, enter **my-S3** as the filter value for **Bucket name** field. If you enter any other value, such as **My-s3** or **my***, Macie won't return the buckets.

6. When you finish adding a value for the field, choose **Apply**. Macie applies the filter criteria and adds the condition to a filter box in the filter bar.

Tip

For many fields, you can change a condition's operator from *equals* to *not equals* by choosing the equals icon (●) in the filter box. If you do this, Macie changes the operator to *not equals* and displays the not equals icon (⊄) in the filter box. To switch to the *equals* operator again, choose the not equals icon.

7. Repeat steps 4 through 6 for each additional condition that you want to add.
8. To remove a condition, choose the remove condition icon (⊗) in the filter box for the condition.
9. To change a condition, remove the condition by choosing the remove condition icon (⊗) in the filter box for the condition. Then repeat steps 4 through 6 to add a condition with the correct settings.

Filtering your inventory programmatically with the Amazon Macie API

To filter your bucket inventory programmatically, specify filter criteria in queries that you submit using the [DescribeBuckets](#) operation of the Amazon Macie API. This operation returns an array of objects. Each object contains statistical data and other information about a bucket that meets the filter criteria.

To specify filter criteria in a query, include a map of filter conditions in your request. For each condition, specify a field, an operator, and one or more values for the field. The type and number of values depends on the field and operator that you choose. For information about the fields, operators, and types of values that you can use in a condition, see [Amazon S3 Data Source](#) in the *Amazon Macie API Reference*.

The following examples show you how to specify filter criteria in queries that you submit using the [AWS Command Line Interface \(AWS CLI\)](#). You can also do this by sending HTTPS requests directly to Macie, or by using a current version of another AWS command line tool or an AWS SDK. For information about AWS tools and SDKs, see [Tools to Build on AWS](#).

Examples

- [Example 1: Find buckets by bucket name \(p. 37\)](#)
- [Example 2: Find buckets that are publicly accessible \(p. 38\)](#)
- [Example 3: Find buckets that contain unencrypted objects \(p. 38\)](#)
- [Example 4: Find buckets that aren't monitored by a job \(p. 39\)](#)
- [Example 5: Find buckets that replicate data to external accounts \(p. 39\)](#)
- [Example 6: Find buckets based on multiple criteria \(p. 39\)](#)

The examples use the `describe-buckets` command. If an example runs successfully, Macie returns a `buckets` array. The array contains an object for each bucket that's in the current AWS Region and meets the filter criteria. For an example of this output, expand the following section.

Example of a buckets array

In this example, the `buckets` array provides details about two buckets that met the filter criteria specified in a query.

```
{
  "buckets": [
    {
      "accountId": "123456789012",
      "bucketArn": "arn:aws:s3:::DOC-EXAMPLE-BUCKET1",
      "bucketCreatedAt": "2020-05-18T19:54:00+00:00",
      "bucketName": "DOC-EXAMPLE-BUCKET1",
      "allowsUnencryptedObjectUploads": "FALSE",
      "classifiableObjectCount": 13,
      "classifiableSizeInBytes": 1592088,
      "jobDetails": {
        "isDefinedInJob": "TRUE",
        "isMonitoredByJob": "TRUE",
        "lastJobId": "08c81dc4a2f3377fae45c9ddaexample",
        "lastJobRunTime": "2021-04-26T14:55:30.270000+00:00"
      },
      "lastUpdated": "2021-04-30T07:33:06.337000+00:00",
      "objectCount": 13,
      "objectCountByEncryptionType": {
        "customerManaged": 0,
        "kmsManaged": 2,
        "s3Managed": 7,
        "unencrypted": 4,
        "unknown": 0
      },
      "publicAccess": {
        "effectivePermission": "NOT_PUBLIC",
        "permissionConfiguration": {
          "accountLevelPermissions": {
            "blockPublicAccess": {
              "blockPublicAcls": true,
              "blockPublicPolicy": true,
              "ignorePublicAcls": true,
              "restrictPublicBuckets": true
            }
          },
          "bucketLevelPermissions": {
            "accessControlList": {
              "allowsPublicReadAccess": false,
              "allowsPublicWriteAccess": false
            },
            "blockPublicAccess": {
              "blockPublicAcls": true,
              "blockPublicPolicy": true,
              "ignorePublicAcls": true,
              "restrictPublicBuckets": true
            }
          },
          "bucketPolicy": {
            "allowsPublicReadAccess": false,
            "allowsPublicWriteAccess": false
          }
        }
      },
      "region": "us-east-1",
      "replicationDetails": {
        "replicated": false,
        "replicatedExternally": false,
        "replicationAccounts": []
      }
    },
  ]
}
```

```
"serverSideEncryption": {
  "kmsMasterKeyId": null,
  "type": "NONE"
},
"sharedAccess": "NOT_SHARED",
"sizeInBytes": 4549746,
"sizeInBytesCompressed": 0,
"tags": [
  {
    "key": "Division",
    "value": "HR"
  },
  {
    "key": "Team",
    "value": "Recruiting"
  }
],
"unclassifiableObjectCount": {
  "fileType": 0,
  "storageClass": 0,
  "total": 0
},
"unclassifiableObjectSizeInBytes": {
  "fileType": 0,
  "storageClass": 0,
  "total": 0
},
"versioning": false
},
{
  "accountId": "123456789012",
  "bucketArn": "arn:aws:s3:::DOC-EXAMPLE-BUCKET2",
  "bucketCreatedAt": "2020-11-25T18:24:38+00:00",
  "bucketName": "DOC-EXAMPLE-BUCKET2",
  "allowsUnencryptedObjectUploads": "TRUE",
  "classifiableObjectCount": 8,
  "classifiableSizeInBytes": 133810,
  "jobDetails": {
    "isDefinedInJob": "TRUE",
    "isMonitoredByJob": "FALSE",
    "lastJobId": "188d4f6044d621771ef7d65f2example",
    "lastJobRunTime": "2021-04-09T19:37:11.511000+00:00"
  },
  "lastUpdated": "2021-04-30T07:33:06.337000+00:00",
  "objectCount": 8,
  "objectCountByEncryptionType": {
    "customerManaged": 0,
    "kmsManaged": 0,
    "s3Managed": 8,
    "unencrypted": 0,
    "unknown": 0
  },
  "publicAccess": {
    "effectivePermission": "NOT_PUBLIC",
    "permissionConfiguration": {
      "accountLevelPermissions": {
        "blockPublicAccess": {
          "blockPublicAcls": true,
          "blockPublicPolicy": true,
          "ignorePublicAcls": true,
          "restrictPublicBuckets": true
        }
      },
      "bucketLevelPermissions": {
        "accessControlList": {
          "allowsPublicReadAccess": false,

```

```
        "allowsPublicWriteAccess": false
      },
      "blockPublicAccess": {
        "blockPublicAcls": true,
        "blockPublicPolicy": true,
        "ignorePublicAcls": true,
        "restrictPublicBuckets": true
      },
      "bucketPolicy": {
        "allowsPublicReadAccess": false,
        "allowsPublicWriteAccess": false
      }
    }
  },
  "region": "us-east-1",
  "replicationDetails": {
    "replicated": false,
    "replicatedExternally": false,
    "replicationAccounts": []
  },
  "serverSideEncryption": {
    "kmsMasterKeyId": null,
    "type": "AES256"
  },
  "sharedAccess": "EXTERNAL",
  "sizeInBytes": 175978,
  "sizeInBytesCompressed": 0,
  "tags": [
    {
      "key": "Division",
      "value": "HR"
    },
    {
      "key": "Team",
      "value": "Recruiting"
    }
  ],
  "unclassifiableObjectCount": {
    "fileType": 0,
    "storageClass": 0,
    "total": 0
  },
  "unclassifiableObjectSizeInBytes": {
    "fileType": 0,
    "storageClass": 0,
    "total": 0
  },
  "versioning": true
}
]
```

If no buckets meet the filter criteria, Macie returns an empty `buckets` array.

```
{
  "buckets": []
}
```

Example 1: Find buckets by bucket name

This example uses the [describe-buckets](#) command to query metadata for all buckets whose names begin with `my-S3` and are in the current AWS Region.

For Linux, macOS, or Unix:

```
$ aws macie2 describe-buckets --criteria '{"bucketName":{"prefix":"my-S3"}}'
```

For Microsoft Windows:

```
C:\> aws macie2 describe-buckets --criteria={"bucketName":{"prefix":"my-S3"}}
```

Where:

- *bucketName* specifies the JSON name of the **Bucket name** field.
- *prefix* specifies the *prefix* operator.
- *my-S3* is the value for the **Bucket name** field.

Example 2: Find buckets that are publicly accessible

This example uses the [describe-buckets](#) command to query metadata for buckets that are in the current AWS Region and, based on a combination of permissions settings, are publicly accessible.

For Linux, macOS, or Unix:

```
$ aws macie2 describe-buckets --criteria '{"publicAccess.effectivePermission":{"eq":["PUBLIC"]}]'
```

For Microsoft Windows:

```
C:\> aws macie2 describe-buckets --criteria={"publicAccess.effectivePermission":{"eq":["PUBLIC"]}}
```

Where:

- *publicAccess.effectivePermission* specifies the JSON name of the **Effective permission** field.
- *eq* specifies the *equals* operator.
- *PUBLIC* is an enumerated value for the **Effective permission** field.

Example 3: Find buckets that contain unencrypted objects

This example uses the [describe-buckets](#) command to query metadata for buckets that are in the current AWS Region and contain unencrypted objects.

For Linux, macOS, or Unix:

```
$ aws macie2 describe-buckets --criteria '{"objectCountByEncryptionType.unencrypted":{"gte":1}]'
```

For Microsoft Windows:

```
C:\> aws macie2 describe-buckets --criteria={"objectCountByEncryptionType.unencrypted":{"gte":1}}
```

Where:

- *objectCountByEncryptionType.unencrypted* specifies the JSON name of the **No encryption** field.

- *gte* specifies the *greater than or equal to* operator.
- *1* is the lowest value in an inclusive, relative numeric range for the **No encryption** field.

Example 4: Find buckets that aren't monitored by a job

This example uses the `describe-buckets` command to query metadata for buckets that are in the current AWS Region and aren't associated with any periodic sensitive data discovery jobs.

For Linux, macOS, or Unix:

```
$ aws macie2 describe-buckets --criteria '{"jobDetails.isMonitoredByJob":{"eq":["FALSE"]}}'
```

For Microsoft Windows:

```
C:\> aws macie2 describe-buckets --criteria={"jobDetails.isMonitoredByJob":{"eq":["FALSE"]}}
```

Where:

- `jobDetails.isMonitoredByJob` specifies the JSON name of the **Actively monitored by job** field.
- `eq` specifies the *equals* operator.
- `FALSE` is an enumerated value for the **Actively monitored by job** field.

Example 5: Find buckets that replicate data to external accounts

This example uses the `describe-buckets` command to query metadata for buckets that are in the current AWS Region and are configured to replicate objects to an AWS account that isn't part of your organization.

For Linux, macOS, or Unix:

```
$ aws macie2 describe-buckets --criteria '{"replicationDetails.replicatedExternally":{"eq":["true"]}}'
```

For Microsoft Windows:

```
C:\> aws macie2 describe-buckets --criteria={"replicationDetails.replicatedExternally":{"eq":["true"]}}
```

Where:

- `replicationDetails.replicatedExternally` specifies the JSON name of the **Replicated externally** field.
- `eq` specifies the *equals* operator.
- `true` specifies a Boolean value for the **Replicated externally** field.

Example 6: Find buckets based on multiple criteria

This example uses the `describe-buckets` command to query metadata for buckets that are in the current AWS Region and meet the following criteria: are publicly accessible based on a combination of permission settings; contain unencrypted objects; and aren't associated with any periodic sensitive data discovery jobs.

For Linux, macOS, or Unix, using the backslash (\) line-continuation character to improve readability:

```
$ aws macie2 describe-buckets \  
--criteria '{"publicAccess.effectivePermission":{"eq":  
["PUBLIC"]},"objectCountByEncryptionType.unencrypted":  
{"gte":1},"jobDetails.isMonitoredByJob":{"eq":["FALSE"]}}'
```

For Microsoft Windows, using the caret (^) line-continuation character to improve readability:

```
C:\> aws macie2 describe-buckets ^  
--criteria={\"publicAccess.effectivePermission\":{\"eq\":[\"PUBLIC\"]},  
\"objectCountByEncryptionType.unencrypted\":{\"gte\":1},\"jobDetails.isMonitoredByJob\":  
{\"eq\":[\"FALSE\"]}}
```

Where:

- `publicAccess.effectivePermission` specifies the JSON name of the **Effective permission** field, and:
 - `eq` specifies the *equals* operator.
 - `PUBLIC` is an enumerated value for the **Effective permission** field.
- `objectCountByEncryptionType.unencrypted` specifies the JSON name of the **No encryption** field, and:
 - `gte` specifies the *greater than or equal to* operator.
 - `1` is the lowest value in an inclusive, relative numeric range for the **No encryption** field.
- `jobDetails.isMonitoredByJob` specifies the JSON name of the **Actively monitored by job** field, and:
 - `eq` specifies the *equals* operator.
 - `FALSE` is an enumerated value for the **Actively monitored by job** field.

Allowing Amazon Macie to access S3 buckets and objects

When you enable Amazon Macie for your AWS account, Macie creates a [service-linked role \(p. 302\)](#) that grants Macie the permissions that it requires to call Amazon Simple Storage Service (Amazon S3) and other AWS services on your behalf. A service-linked role simplifies the process of setting up an AWS service because you don't have to manually add permissions for the service to complete actions on your behalf. To learn more about this type of role, see [Using service-linked roles](#) in the *AWS Identity and Access Management User Guide*.

The permissions policy for the Macie service-linked role (`AWSServiceRoleForAmazonMacie`) allows Macie to perform actions that include retrieving information about your S3 buckets and objects, and retrieving objects from your S3 buckets. If you're the Macie administrator for an organization, the policy also allows Macie to perform these actions on your behalf for member accounts in your organization.

Macie uses these permissions to perform tasks such as:

- Generate and maintain an inventory of your S3 buckets
- Provide statistical and other data about the buckets and objects in the buckets
- Monitor and evaluate the buckets for security and access control
- Analyze objects in the buckets to detect sensitive data

In most cases, Macie has the permissions that it needs to perform these tasks. However, if a bucket has a restrictive bucket policy, the policy might prevent Macie from performing some or all of these tasks.

A *bucket policy* is a resource-based AWS Identity and Access Management (IAM) policy that specifies which actions a principal (AWS account, IAM user, or IAM role) can perform on an S3 bucket, and the conditions under which a principal can perform those actions. The actions and conditions can apply to bucket-level operations, such as retrieving information about a bucket, and object-level operations, such as retrieving objects from a bucket.

Bucket policies typically grant or restrict access by using explicit `Allow` or `Deny` statements and conditions. For example, a bucket policy might contain an `Allow` or `Deny` statement that denies access to the bucket unless specific source IP addresses, Amazon Virtual Private Cloud (Amazon VPC) endpoints, or VPCs are used to access the bucket. For information about using bucket policies to grant or restrict access to buckets, see [Bucket policies and user policies](#) and [How Amazon S3 authorizes a request](#) in the *Amazon Simple Storage Service User Guide*.

If a bucket policy uses an explicit `Allow` statement, the policy doesn't prevent Macie from retrieving information about the bucket and the bucket's objects, or retrieving objects from the bucket. This is because the `Allow` statements in the permissions policy for the Macie service-linked role grant these permissions.

However, if a bucket policy uses an explicit `Deny` statement with one or more conditions, Macie might not be allowed to retrieve information about the bucket or the bucket's objects, or retrieve the bucket's objects. For example, if a bucket policy explicitly denies access from all sources except a specific IP address, Macie won't be allowed to analyze the bucket's objects when you run a sensitive data discovery job. This is because restrictive bucket policies take precedence over the `Allow` statements in the permissions policy for the Macie service-linked role.

To allow Macie to access a bucket that has a restrictive bucket policy, you can add a condition for the Macie service-linked role (`AWSServiceRoleForAmazonMacie`) to the bucket policy. The condition can exclude the Macie service-linked role from matching the `Deny` restriction in the policy. It can do this by using the `aws:PrincipalArn` [global condition context key](#) and the Amazon Resource Name (ARN) of the Macie service-linked role.

The following procedure guides you through this process and provides an example.

To add the Macie service-linked role to a bucket policy

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the navigation pane, choose **Buckets**.
3. Choose the bucket that you want to allow Macie to access.
4. On the **Permissions** tab, under **Bucket policy**, choose **Edit**.
5. In the **Bucket policy** editor, identify each `Deny` statement that restricts access and prevents Macie from accessing the bucket or the bucket's objects.
6. In each `Deny` statement, add a condition that uses the `aws:PrincipalArn` global condition context key and specifies the ARN of the Macie service-linked role for your AWS account.

The value for the condition key should be `arn:aws:iam::123456789012:role/aws-service-role/macie.amazonaws.com/AWSServiceRoleForAmazonMacie`, where `123456789012` is the account ID for your AWS account.

Where you add this to a bucket policy depends on the structure, elements, and conditions that the policy currently contains. To learn about supported structures and elements, see [Policies and permissions in Amazon S3](#) in the *Amazon Simple Storage Service User Guide*.

The following is an example of a bucket policy that uses an explicit `Deny` statement to restrict access to a bucket named `DOC-EXAMPLE-BUCKET`. With the current policy, the bucket can be accessed only from

the VPC endpoint whose ID is `vpce-1a2b3c4d`. Access from all other VPC endpoints is denied, including access from the AWS Management Console and Macie.

```
{
  "Version": "2012-10-17",
  "Id": "Policy1415115example",
  "Statement": [
    {
      "Sid": "Access from specific VPCE only",
      "Effect": "Deny",
      "Principal": "*",
      "Action": "s3:*",
      "Resource": [
        "arn:aws:s3:::DOC-EXAMPLE-BUCKET",
        "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"
      ],
      "Condition": {
        "StringNotEquals": {
          "aws:SourceVpce": "vpce-1a2b3c4d"
        }
      }
    }
  ]
}
```

To change this policy and allow Macie to access the bucket and the bucket's objects, we can add a condition that uses the `StringNotLike` [condition operator](#) and the `aws:PrincipalArn` [global condition context key](#). This additional condition excludes the Macie service-linked role from matching the Deny restriction.

```
{
  "Version": "2012-10-17",
  "Id": "Policy1415115example",
  "Statement": [
    {
      "Sid": "Access from specific VPCE and Macie only",
      "Effect": "Deny",
      "Principal": "*",
      "Action": "s3:*",
      "Resource": [
        "arn:aws:s3:::DOC-EXAMPLE-BUCKET",
        "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"
      ],
      "Condition": {
        "StringNotEquals": {
          "aws:SourceVpce": "vpce-1a2b3c4d"
        },
        "StringNotLike": {
          "aws:PrincipalArn": "arn:aws:iam::123456789012:role/aws-service-role/macie.amazonaws.com/AWSServiceRoleForAmazonMacie"
        }
      }
    }
  ]
}
```

In the preceding example, the `StringNotLike` condition operator uses the `aws:PrincipalArn` condition context key to specify the ARN of the Macie service-linked role, where:

- `123456789012` is the account ID for the AWS account that's permitted to use Macie to retrieve information about the bucket and the bucket's objects, and retrieve objects from the bucket.
- `macie.amazonaws.com` is the identifier for the Macie service principal.

- `AWSServiceRoleForAmazonMacie` is the name of the Macie service-linked role.

We used the `StringNotLike` operator because the policy already uses a `StringNotEquals` operator. A policy can use the `StringNotEquals` operator only once.

For additional policy examples and detailed information about managing access to Amazon S3 resources, see [Identity and access management in Amazon S3](#) in the *Amazon Simple Storage Service User Guide*.

Discovering sensitive data with Amazon Macie

To discover sensitive data with Amazon Macie, you create and run sensitive data discovery jobs. A sensitive data discovery job analyzes objects in Amazon Simple Storage Service (Amazon S3) buckets to determine whether the objects contain sensitive data, and it provides detailed reports of the sensitive data that it finds and the analysis that it performs. By creating and running jobs, you can automate discovery, logging, and reporting of sensitive data in S3 buckets.

A job can analyze S3 objects by using managed data identifiers that Macie provides, custom data identifiers that you define, or a combination of the two:

- A *managed data identifier* is a set of built-in criteria and techniques that detect a specific type of sensitive data—for example, credit card numbers, AWS secret access keys, or passport numbers for a particular country or region. These identifiers can detect a large and growing list of sensitive data types for many countries and regions, including multiple types of personally identifiable information (PII), credentials data, and financial data.
- A *custom data identifier* is a set of criteria that you define to detect sensitive data. The criteria consist of a regular expression (*regex*) that defines a text pattern to match and, optionally, character sequences and a proximity rule that refine the results. With custom data identifiers, you can detect sensitive data that reflects your organization's particular scenarios, intellectual property, or proprietary data—for example, employee IDs, customer account numbers, or internal data classifications.

To fine tune the analysis, a job can also use allow lists that you define. An *allow list* specifies text or a text pattern to ignore in S3 objects, typically sensitive data exceptions for your particular scenarios or environment—for example, public names or phone numbers for your organization, or sample data that your organization uses for testing. If an object contains text that matches an entry or pattern in an allow list, Macie doesn't report the text as sensitive data, even if the text matches the criteria of a managed data identifier or a custom data identifier.

When you create a job, you start by specifying which S3 buckets contain objects to analyze. Macie can analyze an S3 object if the following is true:

- The object uses a supported file or storage format. For more information, see [Supported file and storage formats \(p. 124\)](#).
- If the object is encrypted, it's encrypted with a key that Macie is allowed to use. For more information, see [Analyzing encrypted S3 objects \(p. 125\)](#).
- The object is stored directly in Amazon S3 and uses a supported storage class—S3 Intelligent-Tiering, S3 One Zone-IA, S3 Standard, or S3 Standard-IA. Macie can't analyze data that's stored in Amazon S3 Glacier or other AWS services.

Tip

Although Macie is optimized for Amazon S3, you can use it to discover sensitive data that you currently store elsewhere. You can do this by moving the data to Amazon S3 temporarily or permanently. For example, export Amazon RDS or Amazon Aurora snapshots to Amazon S3 in Apache Parquet format. Or export an Amazon DynamoDB table to Amazon S3. You can then create a job to analyze the data in Amazon S3.

- If the object is stored in a bucket that has a restrictive bucket policy, the policy allows Macie to access objects in the bucket. For more information, see [Allowing Macie to access S3 buckets and objects \(p. 40\)](#).

As you create and configure a job, you also choose options to define the schedule and scope of the job's analysis. You can run a job only once, for on-demand analysis and assessment, or on a recurring basis for periodic analysis, assessment, and monitoring. To define the breadth and depth of a job's analysis, you can also choose various scope options for the job. These options include custom criteria that derive from properties of S3 buckets and objects, such as tags. For more information, see [Scope options for jobs](#) (p. 89).

To help you meet and maintain compliance with your data security and privacy requirements, each job produces records of the sensitive data that it finds and the analysis that it performs—*sensitive data findings* and *sensitive data discovery results*. A *sensitive data finding* is a detailed report of sensitive data in an object. A *sensitive data discovery result* is a record that logs details about the analysis of an object. Each type of record adheres to a standardized schema, which can help you query, monitor, and process the records by using other applications, services, and systems as necessary. For more information, see [Reviewing job statistics and results](#) (p. 113).

Topics

- [Using managed data identifiers in Amazon Macie](#) (p. 45)
- [Building custom data identifiers in Amazon Macie](#) (p. 64)
- [Defining sensitive data exceptions with Amazon Macie allow lists](#) (p. 70)
- [Running sensitive data discovery jobs in Amazon Macie](#) (p. 88)
- [Supported file and storage formats in Amazon Macie](#) (p. 124)
- [Analyzing encrypted S3 objects with Amazon Macie](#) (p. 125)
- [Storing and retaining sensitive data discovery results with Amazon Macie](#) (p. 130)

Using managed data identifiers in Amazon Macie

Amazon Macie uses a combination of criteria and techniques, including machine learning and pattern matching, to detect sensitive data. These criteria and techniques, collectively referred to as *managed data identifiers*, can detect a large and growing list of sensitive data types for many countries and regions, including multiple types of personally identifiable information (PII), personal health information (PHI), financial data, and credentials data.

Each managed data identifier is designed to detect a specific type of sensitive data—for example, AWS secret access keys, credit card numbers, or passport numbers for a particular country or region. When you create a sensitive data discovery job, you can configure the job to use these identifiers to analyze objects in Amazon Simple Storage Service (Amazon S3) buckets that you specify.

Macie can detect the following categories of sensitive data by using managed data identifiers:

- Credentials, for credentials data such as private keys and AWS secret access keys.
- Financial information, for financial data such as credit card numbers and bank account numbers.
- Personal information, for PHI such as health insurance and medical identification numbers, and PII such as passport numbers.

Within each category, Macie can detect multiple types of sensitive data. The topics in this section list and describe each type and any relevant requirements for detecting it. For each type, they also indicate the unique identifier (ID) for the managed data identifier that's designed to detect the data. When you create a sensitive data discovery job, you can use this ID to explicitly include or exclude a managed data identifier from the job's analysis.

Topics

- [Keyword requirements \(p. 46\)](#)
- [Sensitive data types: Credentials \(p. 47\)](#)
- [Sensitive data types: Financial information \(p. 48\)](#)
- [Sensitive data types: Personal information – Personal health information \(p. 52\)](#)
- [Sensitive data types: Personal information – Personally identifiable information \(p. 54\)](#)

For information about the types of data that Macie can analyze, see [Supported file and storage formats \(p. 124\)](#).

Keyword requirements

To detect certain types of sensitive data, Macie requires a keyword to be in proximity of the data. If this is the case for a particular type of data, a subsequent topic in this section indicates specific keyword requirements for that data.

If a keyword has to be in proximity of a particular type of data, the keyword typically has to be within 30 characters (inclusively) of the data. Additional proximity requirements vary based on an S3 object's file type or storage format.

Structured, columnar data

For columnar data, a keyword has to be part of the same value or in the name of the column or field that stores a value. This is true for Microsoft Excel workbooks, CSV files, and TSV files.

For example, if the value for a field contains both *SSN* and a nine-digit number that uses the syntax of a US Social Security number (SSN), Macie can detect the *SSN* in the field. Similarly, if the name of a column contains *SSN*, Macie can detect each *SSN* in the column. Macie treats the values in that column as being in proximity of the keyword *SSN*.

Structured, record-based data

For record-based data, a keyword has to be part of the same value or in the name of an element in the path to the field or array that stores a value. This is true for Apache Avro object containers, Apache Parquet files, JSON files, and JSON Lines files.

For example, if the value for a field contains both *credentials* and a character sequence that uses the syntax of an AWS secret access key, Macie can detect the key in the field. Similarly, if the path to a field is `$.credentials.aws.key`, Macie can detect an AWS secret access key in the field. Macie treats the value in the field as being in proximity of the keyword *credentials*.

Unstructured data

For Adobe Portable Document Format files, Microsoft Word documents, and non-binary text files other than CSV, JSON, JSON Lines, and TSV files, there aren't any additional proximity requirements. A keyword typically has to be within 30 characters (inclusively) of the data. This includes any structured data, such as tables, in these types of files.

Keywords aren't case sensitive. In addition, if a keyword contains a space, Macie automatically matches keyword variations that don't contain the space or contain an underscore (`_`) or a hyphen (`-`) instead of the space. In certain cases, Macie also expands or abbreviates a keyword to address common variations of the keyword.

For a demonstration of how keywords provide context and help Macie detect specific types of sensitive data, watch the following video: [How Amazon Macie uses keywords to discover sensitive data](#).

Sensitive data types: Credentials

The following table lists and describes the types of credentials that Macie can detect using managed data identifiers.

Detection type	Managed data identifier ID	Keyword required	Additional information	Countries and regions
AWS secret access key	AWS_CREDENTIALS	Yes, including: aws_secret_access_key, credentials, secret access key, secret key, set-awscredential	–	Any
HTTP Basic Authorization header	HTTP_BASIC_AUTH_HEADER	–	Detection requires a complete header, including the field name and authentication	Any

Detection type	Managed data identifier ID	Keyword required	Additional information	Countries and regions
			scheme directive, as specified by RFC 7617 . For example: Authorization: Basic QWxhZGRpbjpvGVuIHNlc2FtZQ== and Proxy- Authorization: Basic dGVzdDoxMjPCow==	
JSON Web Token (JWT)	JSON_WEB_TOKEN	No	Macie can detect JWTs that comply with the requirements specified by RFC 7519 for JSON Web Signature (JWS) structures. The tokens can be signed or unsigned.	Any
OpenSSH private key	OPENSSSH_PRIVATE_KEY	No	–	Any
PGP private key	PGP_PRIVATE_KEY	No	–	Any
Public-Key Cryptography Standard (PKCS) private key	PKCS	No	–	Any
PuTTY private key	PUTTY_PRIVATE_KEY	No	–	Any

Sensitive data types: Financial information

The following table lists and describes the types of financial information that Macie can detect using managed data identifiers. These are in addition to certain types of data that might also qualify as personally identifiable information (PII).

Detection type	Managed data identifier ID	Keyword required	Additional information	Countries and regions
Bank account number	Depending on country or region: BANK_ACCOUNT_NUMBER (for Canadian and US bank account numbers), FRANCE_BANK_ACCOUNT_NUMBER,	Varies, see the section called "Keywords for bank account numbers" (p. 51)	This includes: <ul style="list-style-type: none"> Canadian and US bank account numbers that consist of 9–17 digit sequences and don't 	Canada, France, Germany, Italy, Spain, UK, US

Detection type	Managed data identifier ID	Keyword required	Additional information	Countries and regions
	GERMANY_BANK_ACCOUNT_NUMBER, ITALY_BANK_ACCOUNT_NUMBER, SPAIN_BANK_ACCOUNT_NUMBER, UK_BANK_ACCOUNT_NUMBER		<p>contain any spaces.</p> <ul style="list-style-type: none"> International Bank Account Numbers (IBANs) that consist of up to 34 alphanumeric characters, including elements such as country code. 	
Credit card expiration date	CREDIT_CARD_EXPIRATION	Yes. Including: exp d, exp m, exp y, expiration, expiry	Support includes most date formats, such as all digits and combinations of digits and names of months. Date components can be separated by slashes (/), hyphens (-), or applicable keywords. For example, Macie can detect dates such as 02/24, 02/2024, Feb 2024, 24-Feb, and expY=2024, expM=02.	Any
Credit card magnetic strip data	CREDIT_CARD_MAGNETIC_STRIP	Yes. Including: card data, iso7813, mag, magstripe, stripe, swipe	This includes tracks 1 and 2.	Any

Detection type	Managed data identifier ID	Keyword required	Additional information	Countries and regions
Credit card number	CREDIT_CARD_NUMBER, for credit card numbers that are in proximity of a keyword, or CREDIT_CARD_NUMBER_(NO_KEYWORD), for credit card numbers that aren't in proximity of a keyword	Varies. ¹ (p. 50)	Detection requires the data to be a 13–19 digit sequence that adheres to the Luhn check formula and uses a standard card number prefix for any of the following types of credit cards: American Express, Dankort, Diner's Club, Discover, Electron, Japanese Card Bureau (JCB), Mastercard, UnionPay, and Visa. ² (p. 50)	Any
Credit card verification code	CREDIT_CARD_SECURITY_CODE	Any, including: card id, card identification code, card identification number, card security code, card validation code, card validation number, card verification data, card verification value, cvc, cvc2, cvv, cvv2, elo verification code	–	Any

- Macie provides two managed data identifiers for credit card numbers:
 - CREDIT_CARD_NUMBER, for credit card numbers that are in proximity of a keyword.
 - CREDIT_CARD_NUMBER_(NO_KEYWORD), for credit card numbers that aren't in proximity of a keyword.

For the former, required keywords include: account number, american express, amex, bank card, card, card num, card number, cc #, ccn, check card, credit, credit card#, dankort, debit, debit card, diners club, discover, electron, japanese card bureau, jcb, mastercard, mc, pan, payment account number, payment card number, pcn, union pay, visa

- Macie doesn't report occurrences of the following sequences, which credit card issuers have reserved for public testing:

122000000000003, 2222405343248877, 2222990905257051, 2223007648726984, 2223577120017656, 30569309025904, 34343434343434, 3528000700000000,

3530111333300000, 3566002020360505, 36148900647913, 36700102000000, 371449635398431, 378282246310005, 378734493671000, 38520000023237, 4012888888881881, 4111111111111111, 42222222222222, 4444333322221111, 4462030000000000, 4484070000000000, 4911830000000, 4917300800000000, 4917610000000000, 4917610000000000003, 5019717010103742, 5105105105105100, 5111010030175156, 5185540810000019, 5200828282828210, 5204230080000017, 5204740009900014, 5420923878724339, 5454545454545454, 5455330760000018, 5506900490000436, 5506900490000444, 5506900510000234, 5506920809243667, 5506922400634930, 5506927427317625, 5553042241984105, 5555553753048194, 5555555555554444, 5610591081018250, 6011000990139424, 6011000400000000, 6011111111111117, 630490017740292441, 630495060000000000, 6331101999990016, 6759649826438453, 6799990100000000019, and 76009244561.

Keywords for bank account numbers

To detect International Bank Account Numbers (IBANs) that consist of up to 34 alphanumeric characters, including elements such as country code, Macie doesn't require a keyword to be in proximity of the numbers.

To detect Canadian and US bank account numbers that consist of 9–17 digit sequences and don't contain any spaces, Macie requires a keyword to be in proximity of the numbers. The following table lists the keywords that Macie recognizes for specific countries and regions.

Country or region	Keywords
Canada	bank account, bank acct
France	account code, account number, accountno#, accountnumber#, bban, code bancaire, compte bancaire, customer account id, customer account number, customer bank account id, iban, numéro de compte
Germany	account code, account number, accountno#, accountnumber#, bankleitzahl, bban, customer account id, customer account number, customer bank account id, geheimzahl, iban, kartenummer, kontonummer, kreditkartenummer, sepa
Italy	account code, account number, accountno#, accountnumber#, bban, codice bancario, conto bancario, customer account id, customer account number, customer bank account id, iban, numero di conto
Spain	account code, account number, accountno#, accountnumber#, bban, código cuenta, código cuenta bancaria, cuenta cliente id, customer account ID, customer account number, customer bank account id, iban, número cuenta bancaria cliente, número cuenta cliente
UK	account code, account number, accountno#, accountnumber#, bban, customer account id, customer account number, customer bank account id, iban, sepa

Country or region	Keywords
US	bank account, bank acct

Sensitive data types: Personal information – Personal health information

The following table lists and describes the types of personal health information (PHI) that Macie can detect using managed data identifiers. These are in addition to certain types of data that might also qualify as personally identifiable information (PII).

Detection type	Managed data identifier ID	Keyword required	Additional information	Countries and regions
Drug Enforcement Agency (DEA) Registration Number	US_DRUG_ENFORCEMENT_AGENCY_NUMBER	Yes, including: dea number, dea registration		US
Health Insurance Claim Number (HICN)	USA_HEALTH_INSURANCE_CLAIM_NUMBER	Yes, including: health insurance claim number, hic no, hic no., hic number, hic#, hicn, hicn#., hicno#		US
Health insurance or medical identification number	Depending on country or region: CANADA_HEALTH_NUMBER, EUROPEAN_HEALTH_INSURANCE_CARD_NUMBER, FINLAND_EUROPEAN_HEALTH_INSURANCE_NUMBER, FRANCE_HEALTH_INSURANCE_NUMBER, UK_NHS_NUMBER, USA_MEDICARE_BENEFICIARY_IDENTIFIER	Yes, see the section called "Keywords for health insurance and medical identification numbers" (p. 55)	This includes European Health Insurance Card numbers (EU, Finland), health insurance numbers (France), Medicare Beneficiary Identifiers (US), NHS numbers (UK), and Personal Health Numbers (Canada).	Canada, EU, Finland, France, UK, US
Healthcare Common Procedure Coding System (HCPCS) code	USA_HEALTHCARE_PROCEDURE_CODE	Yes, including: current procedural terminology, hcpcs, healthcare common procedure coding system	–	US
National Drug Code (NDC)	USA_NATIONAL_DRUG_CODE	Yes, including: national drug code, ndc	–	US
National Provider Identifier (NPI)	USA_NATIONAL_PROVIDER_IDENTIFIER	Yes, including: hipaa, n.p.i,	–	US

Detection type	Managed data identifier ID	Keyword required	Additional information	Countries and regions
		national provider, npi		
Unique device identifier (UDI)	MEDICAL_DEVICE_UDI	Med, including: blood, blood bag, dev id, device id, device identifier, gs1, hibcc, iccbba, med, udi, unique device id, unique device identifier	Macie can detect UDIs that comply with formats approved by the US Food and Drug Administration. This includes standard formats defined by GS1, HIBCC, and ICCBBA. ICCBA support is for the ISBT standard.	US

Keywords for health insurance and medical identification numbers

To detect various types of health insurance and medical identification numbers, Macie requires a keyword to be in proximity of the numbers. This includes European Health Insurance Card numbers (EU, Finland), health insurance numbers (France), Medicare Beneficiary Identifiers (US), National Insurance numbers (UK), NHS numbers (UK), and Personal Health Numbers (Canada).

The following table lists the keywords that Macie recognizes for specific countries and regions.

Country or region	Keywords
Canada	canada healthcare number, msp number, personal healthcare number, phn, soins de santé
EU	assicurazione sanitaria numero, carta assicurazione numero, carte d'assurance maladie, carte européenne d'assurance maladie, ceam, ehic, ehic#, finlandehicnumber#, gesundheitskarte, hälsokort, health card, health card number, health insurance card, health insurance number, insurance card number, krankenversicherungskarte, krankenversicherungsnummer, medical account number, numero conto medico, numéro d'assurance maladie, numéro de carte d'assurance, numéro de compte medical, número de cuenta médica, número de seguro de salud, número de tarjeta de seguro, sairaanhoitokortin, sairausvakuutuskortti, sairausvakuutusnumero, sjukförsäkring nummer, sjukförsäkringskort, suomi ehic-numero, tarjeta de salud, terveyskortti, tessera sanitaria assicurazione numero, versicherungsnummer

Country or region	Keywords
Finland	ehic, ehic#, finland health insurance card, finlandehicnumber#, finska sjukförsäkringskort, hälsokort, health card, health card number, health insurance card, health insurance number, sairaanhoitokortin, sairaanhoitokortin, sairausvakuutuskortti, sairausvakuutusnumero, sjukförsäkring nummer, sjukförsäkringskort, suomen sairausvakuutuskortti, suomi ehic-numero, terveyskortti
France	carte d'assuré social, carte vitale, insurance card
UK	national health service, NHS
US	mbi, medicare beneficiary

Sensitive data types: Personal information – Personally identifiable information

The following table lists and describes the types of personally identifiable information (PII) that Macie can detect using managed data identifiers.

Detection type	Managed data identifier ID	Keyword required	Additional information	Countries and regions
Birth date	DATE_OF_BIRTH	Yes, including: bday, b-day, birth date, birthday, date of birth, dob	Support includes most date formats, such as all digits and combinations of digits and names of months. Date components can be separated by spaces, slashes (/), or hyphens (-).	Any
Driver's license identification number	Depending on country or region: AUSTRALIA_DRIVERS_LICENSE, AUSTRIA_DRIVERS_LICENSE, BELGIUM_DRIVERS_LICENSE, BULGARIA_DRIVERS_LICENSE, CANADA_DRIVERS_LICENSE, CROATIA_DRIVERS_LICENSE, CYPRUS_DRIVERS_LICENSE, CZECHIA_DRIVERS_LICENSE, DENMARK_DRIVERS_LICENSE, DRIVERS_LICENSE (for the US), ESTONIA_DRIVERS_LICENSE, FINLAND_DRIVERS_LICENSE,	Yes, see the section called "Keywords for driver's license identification numbers" (p. 60)	–	Australia, Austria, Belgium, Bulgaria, Canada, Croatia, Cyprus, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Latvia, Lithuania, Luxembourg, Malta, Netherlands, Poland, Portugal, Romania, Slovakia,

Amazon Macie User Guide
 Sensitive data types: Personally
 identifiable information (PII)

Detection type	Managed data identifier ID	Keyword required	Additional information	Countries and regions
	FRANCE_DRIVERS_LICENSE, GERMANY_DRIVERS_LICENSE, GREECE_DRIVERS_LICENSE, HUNGARY_DRIVERS_LICENSE, IRELAND_DRIVERS_LICENSE, ITALY_DRIVERS_LICENSE, LATVIA_DRIVERS_LICENSE, LITHUANIA_DRIVERS_LICENSE, LUXEMBOURG_DRIVERS_LICENSE, MALTA_DRIVERS_LICENSE, NETHERLANDS_DRIVERS_LICENSE, POLAND_DRIVERS_LICENSE, PORTUGAL_DRIVERS_LICENSE, ROMANIA_DRIVERS_LICENSE, SLOVAKIA_DRIVERS_LICENSE, SLOVENIA_DRIVERS_LICENSE, SPAIN_DRIVERS_LICENSE, SWEDEN_DRIVERS_LICENSE, UK_DRIVERS_LICENSE			Slovenia, Spain, Sweden, UK, US
Electoral roll number	UK_ELECTORAL_ROLL_NUMBER	Yes eng: electoral #, electoral number, electoral roll #, electoral roll no., electoral roll number, electoralrollno	–	UK
Full name	NAME	No	Macie can detect full names only. Support is limited to Latin character sets.	Any

Amazon Macie User Guide
Sensitive data types: Personally
identifiable information (PII)

Detection type	Managed data identifier ID	Keyword required	Additional information	Countries and regions
Global Positioning System (GPS) coordinates	LATITUDE_LONGITUDE	Yes, including: coordinate, coordinates, lat long, latitude longitude, position	<p>Macie can detect GPS coordinates if the latitude and longitude coordinates are stored as a pair and they're in Decimal Degrees (DD) format, for example: 41.948614, -87.655311.</p> <p>Support doesn't include coordinates in Degrees Decimal Minutes (DDM) format, for example 41°56.9168'N 87°39.3187'W, or Degrees, Minutes, Seconds (DMS) format, for example 41°56'55.0104"N 87°39'19.1196"W.</p>	Any, if the coordinates are in proximity of a keyword in English
HTTP cookie	HTTP_COOKIE	No	Detection requires a complete Cookie or Set-Cookie header. The header can include one or more name-value pairs, for example: Set-Cookie: id=TWlrZQ and Cookie: session=3948; lang=en.	Any
Mailing address	ADDRESS or BRAZIL_CEP_CODE (for Brazil's Código de Endereçamento Postal)	No	Although a keyword isn't required, detection requires the address to include the name of a city or place and a ZIP or Postal Code.	Australia, Canada, France, Germany, Italy, Spain, UK, US

Amazon Macie User Guide
Sensitive data types: Personally
identifiable information (PII)

Detection type	Managed data identifier ID	Keyword required	Additional information	Countries and regions
National identification number	Depending on country or region: BRAZIL_RG_NUMBER FRANCE_NATIONAL_IDENTIFICATION_NUMBER GERMANY_NATIONAL_IDENTIFICATION_NUMBER ITALY_NATIONAL_IDENTIFICATION_NUMBER SPAIN_DNI_NUMBER	Yes, see the section called "Keywords for national identification numbers" (p. 62)	This includes Documento Nacional de Identidad (DNI) identifiers (Spain), French National Institute for Statistics and Economic Studies (INSEE) codes, German National Identity Card numbers, and Registro Geral (RG) numbers (Brazil).	Brazil, France, Germany, Italy, Spain
National Insurance Number (NINO)	UK_NATIONAL_INSURANCE_NUMBER	UK_NATIONAL_INSURANCE_NUMBER, insurance no., insurance number, insurance#, national insurance number, nationalinsurance#, nationalinsurancenum, nin, nino	–	UK
Passport number	Depending on country or region: CANADA_PASSPORT_NUMBER FRANCE_PASSPORT_NUMBER GERMANY_PASSPORT_NUMBER ITALY_PASSPORT_NUMBER SPAIN_PASSPORT_NUMBER UK_PASSPORT_NUMBER USA_PASSPORT_NUMBER	Yes, see the section called "Keywords for passport numbers" (p. 63)	–	Canada, France, Germany, Italy, Spain, UK, US

Amazon Macie User Guide
Sensitive data types: Personally
identifiable information (PII)

Detection type	Managed data identifier ID	Keyword required	Additional information	Countries and regions
Permanent residence number	CANADA_NATIONAL_IDENTIFICATION_NUMBER	Yes, including: carte résident permanent, numéro carte résident permanent, numéro résident permanent, permanent resident card, permanent resident card number, permanent resident no., permanent resident no., permanent resident number, pr no, pr no., pr non, pr number, résident permanent no., résident permanent non		Canada
Phone number	Depending on country or region: BRAZIL_PHONE_NUMBER FRANCE_PHONE_NUMBER GERMANY_PHONE_NUMBER ITALY_PHONE_NUMBER PHONE_NUMBER (for Canada and the US), SPAIN_PHONE_NUMBER UK_PHONE_NUMBER	Yes, including: cell, contact, fax, fax number, mobile, phone, phone number, tel, telephone, telephone number For Brazil, keywords also include: cel, celular, fone, móvel, número residencial, numero residencial, telefone	This includes toll-free numbers in the US and fax numbers. If a keyword is in proximity of the data, the number doesn't have to include a country code. If a keyword isn't in proximity of the data, the number has to include a country code.	Brazil, Canada, France, Germany, Italy, Spain, UK, US

Amazon Macie User Guide
Sensitive data types: Personally
identifiable information (PII)

Detection type	Managed data identifier ID	Keyword required	Additional information	Countries and regions
Social Insurance Number (SIN)	CANADA_SOCIAL_INSURANCE_NUMBER	Yes, including: canadian id, numéro d'assurance sociale, sin, social insurance number Also refer to the section called "Keywords for health insurance and medical identification numbers" (p. 53)	–	Canada
Social Security number (SSN)	Depending on country or region: SPAIN_SOCIAL_SECURITY_NUMBER or USA_SOCIAL_SECURITY_NUMBER	Yes, including: Spain: número de la seguridad social, social security no., social security number, socialsecurityno#, ssn, ssn# • US – social security, ss#, ssn	–	Spain, US
Taxpayer identification or reference number	Depending on country or region: AUSTRALIA_TAX_FILE_NUMBER BRAZIL_CNPJ_NUMBER BRAZIL_CPF_NUMBER FRANCE_TAX_IDENTIFICATION_NUMBER GERMANY_TAX_IDENTIFICATION_NUMBER SPAIN_NIE_NUMBER, SPAIN_NIF_NUMBER, SPAIN_TAX_IDENTIFICATION_NUMBER, UK_TAX_IDENTIFICATION_NUMBER, USA_INDIVIDUAL_TAX_IDENTIFICATION_NUMBER	Yes, see the section called "Keywords for taxpayer identification and reference numbers" (p. 64)	This includes: CIF, NIE, and NIF (Spain); CNPJ and CPF (Brazil); Codice Fiscale (Italy); ITIN (US); Steueridentifikationsnummer (Germany); TFN (Australia); TIN (France); and, TRN, UTR (UK).	Australia, Brazil, France, Germany, Italy, Spain, UK, US

Detection type	Managed data identifier ID	Keyword required	Additional information	Countries and regions
Vehicle identification number (VIN)	VEHICLE_IDENTIFICATION_NUMBER	Yes, NUMBER, Fahrgestellnummer, niv, numarul de identificare, numarul seriei de sasiu, numer VIN, Número de Identificação do Veículo, Número de Identificación de Automóviles, numéro d'identification du véhicule, vehicle identification number, vin, VIN numeris	Macie can detect VINs that consist of a 17-character sequence and adhere to the ISO 3779 and 3780 standards. These standards were designed for worldwide use.	Any, if the VIN is in proximity of a keyword in one of the following languages: English, French, German, Lithuanian, Polish, Portuguese, Romanian, or Spanish

Keywords for driver's license identification numbers

To detect various types of driver's license identification numbers, Macie requires a keyword to be in proximity of the numbers. The following table lists the keywords that Macie recognizes for specific countries and regions.

Country or region	Keywords
Australia	dl#, dl:, dlno#, driver licence, driver license, driver permit, drivers lic., drivers licence, driver's licence, drivers license, driver's license, drivers permit, driver's permit, drivers permit number, driving licence, driving license, driving permit
Austria	führerschein, fuhrerschein, fuhrerschein republik österreich, fuhrerschein republik osterreich
Belgium	fuehrerschein, fuehrerschein- nr, fuehrerscheinnummer, fuhrerschein, fuhrerschein, fuhrerschein- nr, fuhrerscheinnr, fuhrerscheinnummer, fuhrerscheinnummer, numéro permis conduire, permis de conduire, rijbewijs, rijbewijsnummer
Bulgaria	превозно средство, свидетелство за управление на моторно, свидетелство за управление на мпс, сумпс, шофьорска книжка
Canada	dl#, dl:, dlno#, driver licence, driver licences, driver license, driver licenses, driver permit, drivers lic., drivers licence, driver's licence, drivers licences, driver's licences, drivers license, driver's license, drivers licenses, driver's licenses, drivers permit, driver's permit, drivers permit number, driving

Amazon Macie User Guide
Sensitive data types: Personally
identifiable information (PII)

Country or region	Keywords
	licence, driving license, driving permit, permis de conduire
Croatia	vozačka dozvola
Cyprus	άρδια οδήγησης
Czech Republic	číslo licence, číslo licence řidiče, číslo řidičského průkazu, ovladače lic., povolení k jízdě, povolení řidiče, řidiči povolení, řidičský průkaz, řidičský průkaz
Denmark	kørekort, kørekortnummer
Estonia	juhi litsentsi number, juhiloa number, juhiluba, juhiluba number
Finland	ajokortin numero, ajokortti, förare lic., körkort, körkort nummer, kuljettaja lic., permis de conduire
France	permis de conduire
Germany	fuehrerschein, fuehrerschein- nr, fuehrerscheinnummer, fuhrerschein, fuhrerschein- nr, fuhrerschein- nr, fuhrerscheinnummer, fuhrerscheinnummer
Greece	δεια οδήγησης, adeia odigisis
Hungary	illesztőprogramok lic, jogosítvány, jogsí, licencszám, vezető engedély, vezetői engedély
Ireland	ceadúnas tiomána
Italy	patente di guida, patente di guida numero, patente guida, patente guida numero
Latvia	autovadītāja apliecība, licences numurs, vadītāja apliecība, vadītāja apliecības numurs, vadītāja atļauja, vadītāja licences numurs, vadītāji lic.
Lithuania	vairuotojo pažymėjimas
Luxembourg	fahrerlaubnis, fuhrerschäin
Malta	licenzja tas-sewqan
Netherlands	permis de conduire, rijbewijs, rijbewijsnummer
Poland	numer licencyjny, prawo jazdy, zezwolenie na prowadzenie
Portugal	carta de condução, carteira de habilitação, carteira de motorist, carteira habilitação, carteira motorist, licença condução, licença de condução, número de licença, número licença, permissão condução, permissão de condução

Country or region	Keywords
Romania	numărul permisului de conducere, permis de conducere
Slovakia	číslo licencie, číslo vodičského preukazu, ovládače lic., povolenia vodičov, povolenie jazdu, povolenie na jazdu, povolenie vodiča, vodičský preukaz
Slovenia	vozniško dovoljenje
Spain	carnet conductor, el carnet de conductor, licencia conductor, licencia de manejo, número carnet conductor, número de carnet de conductor, número de permiso conductor, número de permiso de conductor, número licencia conductor, número permiso conductor, permiso conducción, permiso conductor, permiso de conducción
Sweden	ajokortin numero, dlno# ajokortti, drivere lic., förare lic., körkort, körkort nummer, körkortsnummer, kuljettajat lic.
UK	dl#, dl:, dlno#, driver licence, driver licences, driver license, driver licenses, driver permit, drivers lic., drivers licence, driver's licence, drivers licences, driver's licences, drivers license, driver's license, drivers licenses, driver's licenses, drivers permit, driver's permit, drivers permit number, driving licence, driving license, driving permit
US	dl#, dl:, dlno#, driver licence, driver licences, driver license, driver licenses, driver permit, drivers lic., drivers licence, driver's licence, drivers licences, driver's licences, drivers license, driver's license, drivers licenses, driver's licenses, drivers permit, driver's permit, drivers permit number, driving licence, driving license, driving permit

Keywords for national identification numbers

To detect various types of national identification numbers, Macie requires a keyword to be in proximity of the numbers. This includes Documento Nacional de Identidad (DNI) identifiers (Spain), French National Institute for Statistics and Economic Studies (INSEE) codes, German National Identity Card numbers, and Registro Geral (RG) numbers (Brazil).

The following table lists the keywords that Macie recognizes for specific countries and regions.

Country or region	Keywords
Brazil	registro geral, rg
France	assurance sociale, carte nationale d'identité, cni, code sécurité sociale, French social security number, fssn#, insee, insurance number, national id number, nationalid#, numéro d'assurance, sécurité sociale, sécurité sociale non., sécurité

Country or region	Keywords
	sociale numéro, social, social security, social security number, socialecuritynumber, ss#, ssn, ssn#
Germany	ausweisnummer, id number, identification number, identity number, insurance number, personal id, personalausweis
Italy	codice fiscal, dati anagrafici, ehic, health card, health insurance card, p. iva, partita i.v.a., personal data, tax code, tessera sanitaria
Spain	dni, dni#, dninúmero#, documento nacional de identidad, identidad único, identidadúnico#, insurance number, national identification number, national identity, nationalid#, nationalidno#, número nacional identidad, personal identification number, personal identity no, unique identity number, uniqueid#

Keywords for passport numbers

To detect various types of passport numbers, Macie requires a keyword to be in proximity of the numbers. The following table lists the keywords that Macie recognizes for specific countries and regions.

Country or region	Keywords
Canada	passepport, passeport#, passport, passport#, passportno, passportno#
France	numéro de passeport, passeport #, passeport n °, passeport non
Germany	ausstellungsdatum, ausstellungsort, geburtsdatum, passport, passports, reiseepass, reiseepass-nr, reiseepassnummer
Italy	italian passport number, numéro passeport, numéro passeport italien, passaporto, passaporto italiana, passaporto numero, passport number, repubblica italiana passaporto
Spain	españa pasaporte, libreta pasaporte, número pasaporte, pasaporte, passport, passport book, passport no, passport number, spain passport
UK	passepport #, passeport n °, passeport non, passeportn °, passport #, passport no, passport number, passport#, passportid
US	passport, travel document

Keywords for taxpayer identification and reference numbers

To detect various types of taxpayer identification and reference numbers, Macie requires a keyword to be in proximity of the numbers. The following table lists the keywords that Macie recognizes for specific countries and regions.

Country or region	Keywords
Australia	tax file number, tfn
Brazil	cadastro de pessoa física, cadastro de pessoa física, cadastro de pessoas físicas, cadastro de pessoas físicas, cadastro nacional da pessoa jurídica, cadastro nacional da pessoa jurídica, cnpj, cpf
France	numéro d'identification fiscal, tax id, tax identification number, tax number, tin, tin#
Germany	identifikationsnummer, steuer id, steueridentifikationsnummer, steuernummer, tax id, tax identification number, tax number
Italy	codice fiscale, dati anagrafici, ehic, health card, health insurance card, p. iva, partita i.v.a., personal data, tax code, tessera sanitaria
Spain	cif, cif número, cifnúmero#, nie, nif, número de contribuyente, número de identidad de extranjero, número de identificación fiscal, número de impuesto corporativo, personal tax number, tax id, tax identification number, tax number, tin, tin#
UK	paye, tax id, tax id no., tax id number, tax identification, tax identification#, tax no., tax number, tax reference, tax#, taxid#, temporary reference number, tin, trn, unique tax reference, unique taxpayer reference, utr
US	i.t.i.n., individual taxpayer identification number, itin

Building custom data identifiers in Amazon Macie

A *custom data identifier* is a set of criteria that you define to detect sensitive data. The criteria consist of a regular expression (*regex*) that defines a text pattern to match and, optionally, character sequences and a proximity rule that refine the results.

With custom data identifiers, you can define detection criteria that reflects your organization's particular scenarios, intellectual property, or proprietary data—for example, employee IDs, customer account numbers, or internal data classifications. When you use these identifiers in sensitive data discovery jobs, you perform targeted analysis of your organization's Amazon Simple Storage Service (Amazon S3) data in a way that supplements the [managed data identifiers \(p. 45\)](#) that Amazon Macie provides.

In addition to detection criteria, you can define custom severity settings for sensitive data findings that a custom data identifier produces. By default, Macie assigns the *Medium* severity to all findings that a

custom data identifier produces—severity doesn't change based on the number of occurrences of text that matches a custom data identifier's detection criteria. By defining custom severity settings, you can specify which severity to assign based on the number of occurrences of text that matches the criteria.

Topics

- [Defining detection criteria for custom data identifiers \(p. 65\)](#)
- [Defining finding severity settings for custom data identifiers \(p. 67\)](#)
- [Creating custom data identifiers \(p. 68\)](#)
- [Regex support in custom data identifiers \(p. 69\)](#)

Defining detection criteria for custom data identifiers

When you create a custom data identifier, you specify a regular expression (*regex*) that defines a text pattern to match in S3 objects. Macie supports a subset of the regex pattern syntax provided by the [Perl Compatible Regular Expressions \(PCRE\) library](#). For more information, see [Regex support \(p. 69\)](#) later in this section.

You can also specify character sequences, such as words and phrases, and a proximity rule to refine the results.

Keywords

These are specific character sequences that must be in proximity of text that matches the regex pattern. The proximity requirements vary based on an S3 object's storage format or file type:

- For structured, columnar data, Macie includes a result if the text matches the regex pattern and a keyword is in the name of the field or column that stores the text, or the text is preceded by and within the maximum match distance of a keyword in the same field or cell value. This is true for Microsoft Excel workbooks, CSV files, and TSV files.
- For structured, record-based data, Macie includes a result if the text matches the regex pattern and the text is within the maximum match distance of a keyword. The keyword can be in the name of an element in the path to the field or array that stores the text, or it can precede and be part of the same value in the field or array that stores the text. This is true for Apache Avro object containers, Apache Parquet files, JSON files, and JSON Lines files.
- For unstructured data, Macie includes a result if the text matches the regex pattern and the text is preceded by and within the maximum match distance of a keyword. This is true for Adobe Portable Document Format files, Microsoft Word documents, and non-binary text files other than CSV, JSON, JSON Lines, and TSV files. This includes any structured data, such as tables, in these types of files.

You can specify as many as 50 keywords. Each keyword can contain 3–90 UTF-8 characters. Keywords aren't case sensitive.

Maximum match distance

This is a character-based proximity rule for keywords. Macie uses this setting to determine whether a keyword precedes text that matches the regex pattern. The setting defines the maximum number of characters that can exist between the end of a complete keyword and the end of text that matches the regex pattern. If text matches the regex pattern, occurs after at least one complete keyword, and occurs within the specified distance of the keyword, Macie includes it in the results. Otherwise, Macie excludes it from the results.

You can specify a distance of 1–300 characters. The default distance is 50 characters. For best results, this distance should be greater than the minimum number of characters of text that the regex is designed to detect. If only part of the text is within the maximum match distance of a keyword, Macie doesn't include it in the results.

Ignore words

These are specific character sequences to exclude from the results. If text matches the regex pattern but it contains an ignore word, Macie doesn't include it in the results.

You can specify as many as 10 ignore words. Each ignore word can contain 4–90 UTF-8 characters. Ignore words are case sensitive.

For example, many companies have a specific syntax for employee IDs. One such syntax might be: a capital letter that indicates whether the employee is a full-time (*F*) or part-time (*P*) employee, followed by a hyphen (-), followed by an eight-digit sequence that identifies the employee. Examples are: *F-12345678*, for a full-time employee, and *P-87654321*, for a part-time employee.

If you create a custom data identifier to detect employee IDs that use this syntax, you might use the following regex: `[A-Z]-\d{8}`. To refine the analysis and avoid false positives, you might also configure the custom data identifier to use the keywords *employee* and *employee ID* and a maximum match distance of 20 characters. With these criteria, the results include text that matches the regex only if the text occurs after the keyword *employee* or *employee ID* and all the text occurs within 20 characters of one of those keywords.

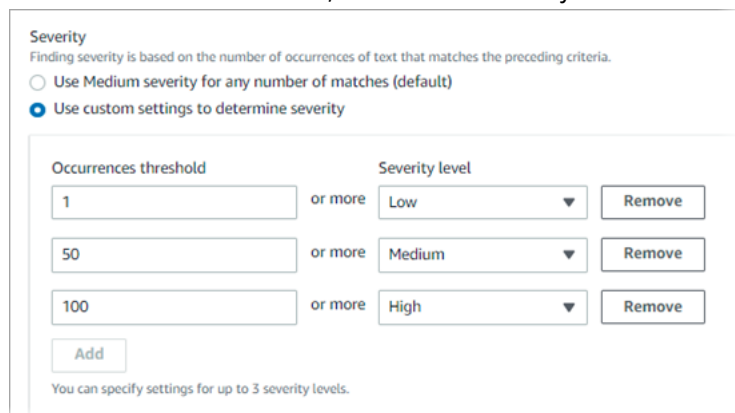
For a demonstration of how keywords can help you find sensitive data and avoid false positives, watch the following video: [How Amazon Macie uses keywords to discover sensitive data](#).

Defining finding severity settings for custom data identifiers

When you create a custom data identifier, you can also define custom severity settings for sensitive data findings that the identifier produces. By default, Macie assigns the *Medium* severity to all findings that a custom data identifier produces—if an S3 object contains at least one occurrence of text that matches a custom data identifier's detection criteria, Macie automatically assigns the *Medium* severity to the resulting finding.

With custom severity settings, you can specify which severity to assign based on the number of occurrences of text that matches the custom data identifier's detection criteria. To do this, you define *occurrences thresholds* for as many as three severity levels: *Low* (least severe), *Medium*, and *High* (most severe). An *occurrences threshold* is the minimum number of matches that must exist in an S3 object to produce a finding with the specified severity. If you specify more than one threshold, the thresholds must be in ascending order by severity, moving from *Low* to *High*.

For example, the following image shows the severity settings for a custom data identifier that specifies three occurrences thresholds, one for each severity level that Macie supports.



The following table indicates the severity of the findings that the custom data identifier produces.

Occurrences threshold	Severity level	Result
1	Low	If an S3 object contains 1–49 occurrences of text that matches the detection criteria, the severity of the resulting finding is <i>Low</i> .
50	Medium	If an S3 object contains 50–99 occurrences of text that matches the detection criteria, the severity of the resulting finding is <i>Medium</i> .
100	High	If an S3 object contains 100 or more occurrences of text that matches the detection criteria, the severity of the resulting finding is <i>High</i> .

You can also use severity settings to specify whether to create a finding at all. If an S3 object contains fewer occurrences than the lowest occurrences threshold, Macie doesn't create a finding.

Creating custom data identifiers

Follow these steps to create a custom data identifier by using the Amazon Macie console. To create a custom data identifier programmatically, you can use the [CreateCustomDataIdentifier](#) operation of the Amazon Macie API.

To create a custom data identifier

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.

2. In the navigation pane, under **Settings**, choose **Custom data identifiers**.

3. Choose **Create**.

4. For **Name**, enter a name for the custom data identifier. The name can contain as many as 128 characters.

Avoid including any sensitive data in the name. Other users of your account might be able to see the name, depending on the actions that they're allowed to perform in Macie.

5. (Optional) For **Description**, enter a brief description of the custom data identifier. The description can contain as many as 512 characters.

Avoid including any sensitive data in the description. Other users of your account might be able to see the description, depending on the actions that they're allowed to perform in Macie.

6. For **Regular expression**, enter the regular expression (*regex*) that defines the text pattern to match. The regex can contain as many as 512 characters. To learn about supported syntax and constraints, see [Regex support \(p. 69\)](#) later in this section.

7. (Optional) For **Keywords**, enter as many as 50 character sequences (separated by commas) to define specific text that must be in proximity of text that matches the regex pattern. Each keyword can contain 3–90 UTF-8 characters. Keywords aren't case sensitive.

Macie includes an occurrence in the results only if the text matches the regex pattern and the text is within the maximum match distance of one of these keywords, as explained in the [preceding topic \(p. 65\)](#).

8. (Optional) For **Ignore words**, enter as many as 10 character sequences (separated by commas) that define specific text to exclude from the results. Each ignore word can contain 4–90 UTF-8 characters. Ignore words are case sensitive.

Macie excludes an occurrence from the results if the text matches the regex pattern but it contains one of these ignore words.

9. (Optional) For **Maximum match distance**, enter the maximum number of characters that can exist between the end of a keyword and the end of text that matches the regex pattern. The distance can be 1–300 characters. The default distance is 50 characters.

Macie includes an occurrence in the results only if the text matches the regex pattern and the text is within this distance of a complete keyword, as explained in the [preceding topic \(p. 65\)](#).

10. For **Severity**, choose how you want Macie to assign severity to sensitive data findings that the custom data identifier produces:

- To automatically assign the *Medium* severity to all findings, choose **Use Medium severity for any number of matches (default)**. With this option, Macie automatically assigns the *Medium* severity to a finding if the affected S3 object contains one or more occurrences of text that matches the detection criteria.
- To assign severity based on occurrences thresholds that you specify, choose **Use custom settings to determine severity**. Then use the **Occurrences threshold** and **Severity level** options to specify

the minimum number of matches that must exist in an S3 object to produce a finding with a selected severity.

For example, to assign the *High* severity to a finding that reports 100 or more occurrences of text that matches the detection criteria, enter **100** in the **Occurrences threshold** box and then choose **High** from the **Severity level** list.

You can specify as many as three occurrences thresholds, one for each severity level that Macie supports: *Low* (for least severe), *Medium*, or *High* (for most severe). If you specify more than one, the thresholds must be in ascending order by severity, moving from *Low* to *High*. If an S3 object contains fewer occurrences than the lowest specified threshold, Macie doesn't create a finding.

11. (Optional) For **Tags**, choose **Add tag**, and then enter as many as 50 tags to assign to the custom data identifier.

A *tag* is a label that you define and assign to certain types of AWS resources. Each tag consists of a required tag key and an optional tag value. Tags can help you identify, categorize, and manage resources in different ways, such as by purpose, owner, environment, or other criteria. To learn more, see [Tagging Amazon Macie resources \(p. 315\)](#).

12. (Optional) For **Evaluate**, enter up to 1,000 characters in the **Sample data** box, and then choose **Test** to test the detection criteria. Macie evaluates the sample data and reports the number of occurrences of text that matches the criteria. You can repeat this step as many times as you like to refine and optimize the criteria.

Note

We strongly recommend that you test and refine the detection criteria before you save the custom data identifier. Because custom data identifiers are used by sensitive data discovery jobs, you can't edit a custom data identifier after you save it. This helps ensure that you have an immutable history of sensitive data findings and discovery results for data privacy and protection audits or investigations that you perform.

13. When you finish, choose **Submit**.

Macie tests the settings and verifies that it can compile the regex. If there's an issue with any of the settings, an error occurs. After you address any errors, you can save the custom data identifier. You can then [create and configure sensitive data discovery jobs \(p. 97\)](#) to use the identifier.

Regex support in custom data identifiers

Macie supports a subset of the regex pattern syntax provided by the [Perl Compatible Regular Expressions \(PCRE\) library](#). Of the constructs provided by the PCRE library, Macie doesn't support the following pattern elements:

- Backreferences
- Capturing groups
- Conditional patterns
- Embedded code
- Global pattern flags, such as */i*, */m*, and */x*
- Recursive patterns
- Positive and negative look-behind and look-ahead zero-width assertions, such as *?=*, *?!*, *?<=*, and *?<!*

To create effective regex patterns for custom data identifiers, also note the following tips and recommendations:

- **Anchors** – Use anchors (*^* or *\$*) only if you expect the pattern to appear at the beginning or end of a file, not the beginning or end of a line.

- **Bounded repeats** – For performance reasons, Macie limits the size of bounded repeat groups. For example, `\d{100,1000}` won't compile in Macie. To approximate this functionality, you can use an open-ended repeat such as `\d{100,}`.
- **Case insensitivity** – To make parts of a pattern case insensitive, you can use the `(?i)` construct instead of the `/i` flag.
- **Performance** – There's no need to optimize prefixes or alternations manually. For example, changing `/hello|hi|hey/` to `/h(?:ello|i|ey)/` won't improve performance.
- **Wildcards** – For performance reasons, Macie limits the number of repeated wildcards. For example, `a*b*a*` won't compile in Macie.

To protect against malformed or long-running expressions, Macie automatically tests regex patterns against a collection of sample text.

Defining sensitive data exceptions with Amazon Macie allow lists

With allow lists in Amazon Macie, you can define specific text and text patterns that you want Macie to ignore when it inspects Amazon Simple Storage Service (Amazon S3) objects for sensitive data. These are typically sensitive data exceptions for your particular scenarios or environment. If data matches text or a text pattern in an allow list, Macie doesn't report the data in sensitive data findings or sensitive data discovery results, even if the data matches the criteria of a [managed data identifier \(p. 45\)](#) or a [custom data identifier \(p. 64\)](#). By using allow lists, you can refine your analysis of Amazon S3 data and reduce noise.

You can create and use two types of allow lists in Macie:

- **Predefined text** – For this type of list, you specify certain character sequences to ignore—for example, the names of public representatives for your organization, specific phone numbers, or specific sample data that your organization uses for testing. If you use this type of list, Macie ignores text that exactly matches an entry in the list.

This type of allow list is helpful if you want to specify words, phrases, and other kinds of character sequences that aren't sensitive, aren't likely to change, and don't necessarily adhere to a common pattern.

- **Regular expression** – For this type of list, you specify a regular expression (*regex*) that defines a text pattern to ignore—for example, public phone numbers for your organization, email addresses for your organization's domain, or patterned sample data that your organization uses for testing. If you use this type of list, Macie ignores text that completely matches the pattern defined by the list.

This type of allow list is helpful if you want to specify text that isn't sensitive but varies or is likely to change while also adhering to a common pattern.

After you create an allow list, you can create and configure sensitive data discovery jobs to use it. When those jobs run, Macie uses the list as part of its data analysis. If Macie finds text that matches an entry or pattern in an allow list, Macie doesn't report that occurrence of text in sensitive data findings or sensitive data discovery results that the job produces. You can configure a job to use as many as 10 allow lists.

You can create and use allow lists in all the AWS Regions where Macie is currently available except the Asia Pacific (Osaka) Region.

Topics

- [Allow list options and requirements in Amazon Macie \(p. 71\)](#)

- [Creating and managing allow lists in Amazon Macie \(p. 78\)](#)

Allow list options and requirements in Amazon Macie

In Amazon Macie, you can use allow lists to specify text or text patterns that you want Macie to ignore when it inspects Amazon Simple Storage Service (Amazon S3) objects for sensitive data. Macie provides options for two types of allow lists, predefined text and regular expressions.

A list of predefined text is helpful if you want Macie to ignore specific words, phrases, and other kinds of character sequences that you don't consider sensitive. Examples are the names of public representatives for your organization, specific phone numbers, or specific sample data that your organization uses for testing. If Macie finds text that matches the criteria of a managed data identifier or a custom data identifier and the text also matches an entry in an allow list, Macie doesn't report that occurrence of text in sensitive data findings or sensitive data discovery results.

A regular expression (*regex*) is helpful if you want Macie to ignore text that varies or is likely to change while also adhering to a common pattern. The regex specifies a text pattern to ignore. Examples are public phone numbers for your organization, email addresses for your organization's domain, or patterned sample data that your organization uses for testing. If Macie finds text that matches the criteria of a managed data identifier or a custom data identifier and the text also matches a regex pattern in an allow list, Macie doesn't report that occurrence of text in sensitive data findings or sensitive data discovery results.

You can create and use both types of allow lists in all the AWS Regions where Macie is currently available except the Asia Pacific (Osaka) Region. As you create and manage allow lists, keep the following options and requirements in mind.

Topics

- [Options and requirements for lists of predefined text \(p. 71\)](#)
 - [Syntax requirements \(p. 72\)](#)
 - [Storage requirements \(p. 72\)](#)
 - [Encryption/Decryption requirements \(p. 73\)](#)
 - [Design considerations and recommendations \(p. 73\)](#)
- [Options and requirements for regular expressions in allow lists \(p. 75\)](#)
 - [Syntax support and recommendations \(p. 75\)](#)
 - [Examples \(p. 76\)](#)

Options and requirements for lists of predefined text

For this type of allow list, you provide a line-delimited plaintext file that lists specific character sequences to ignore. The list entries are typically words, phrases, and other kinds of character sequences that you don't consider sensitive, aren't likely to change, and don't necessarily adhere to a specific pattern. If you use this type of list, Amazon Macie doesn't report occurrences of text that exactly match an entry in the list. Macie treats each list entry as a string literal value.

To use this type of allow list, start by creating the list in a text editor and saving it as a plaintext file. Then upload the list to an S3 bucket and ensure that the storage and encryption settings for the bucket and the object allow Macie to retrieve and decrypt the list. Then [create and configure settings for the list \(p. 79\)](#) in Macie.

After you configure the settings in Macie, we recommend that you test the allow list with a small, representative set of data for your account or organization. To test a list, you can [create a one-time job \(p. 97\)](#) and configure the job to use the list in addition to the managed data identifiers and custom data identifiers that you typically use to analyze data. You can then review the job's results—sensitive

data findings, sensitive data discovery results, or both. If the job's results differ from what you expect, you can change and test the list until the results are what you expect.

After you finish configuring and testing an allow list, you can create and configure additional jobs to use it. When each job starts to run, Macie retrieves the latest version of the list from Amazon S3 and stores it in temporary memory. Macie then uses this temporary copy of the list when it inspects S3 objects that you configured the job to analyze. When the job finishes running, Macie permanently deletes its copy of the list from memory. (The list doesn't persist in Macie. Only the list's settings persist in Macie.) If you configure a job to run more than once, Macie performs all these tasks each time the job runs.

Important

Because lists of predefined text don't persist in Macie, it's important to [check the status of your allow lists \(p. 83\)](#). If you configure a job to use a list and Macie can't retrieve or parse the list when the job starts to run, the job continues to run but Macie doesn't use the list. This means that the job might produce unexpected results, such as sensitive data findings for text that you specified in the list.

Topics

- [Syntax requirements \(p. 72\)](#)
- [Storage requirements \(p. 72\)](#)
- [Encryption/Decryption requirements \(p. 73\)](#)
- [Design considerations and recommendations \(p. 73\)](#)

Syntax requirements

When you create this type of allow list, note the following requirements for the list's file:

- The list must be stored as a plaintext file that uses line breaks to separate individual entries. For example:

```
Akua Mansa
John Doe
Martha Rivera
425-555-0100
425-555-0101
425-555-0102
```

Macie treats each line as a single, distinct entry in the list. The file can also contain blank lines to improve readability. Macie skips blank lines when it parses the file.

- Each entry can contain 1–90 UTF–8 characters.
- Each entry must be a complete, exact match for the text to ignore. Macie doesn't support use of wildcard characters or partial values for entries. Macie treats each entry as a string literal value. Matches aren't case sensitive.
- The file can contain 1–100,000 entries.
- The total storage size of the file can't exceed 35 MB.

Storage requirements

As you add and manage allow lists in Amazon S3, note the following storage requirements and recommendations:

- **Regional support** – An allow list must be stored in an S3 bucket that's in the same AWS Region as your Macie account. Macie can't access an allow list if it's stored in a different Region.
- **Bucket ownership** – An allow list must be stored in an S3 bucket that's owned by your AWS account. If you want other accounts to use the same allow list, consider creating an Amazon S3 replication rule

to replicate the list to buckets that are owned by those accounts. For information about replicating S3 objects, see [Replicating objects](#) in the *Amazon Simple Storage Service User Guide*.

In addition, your AWS Identity and Access Management (IAM) identity must have read access to the S3 bucket and object that store the list. Otherwise, you won't be allowed to create or update the list's settings or check the list's status by using Macie.

- **Bucket policies** – If you store an allow list in an S3 bucket that has a restrictive bucket policy, ensure that the policy allows Macie to retrieve the list. To do this, you can add a condition for the Macie service-linked role to the bucket policy. For more information, see [Allowing Macie to access S3 buckets and objects \(p. 40\)](#).

Also ensure that the policy allows your IAM identity to have read access to the bucket. Otherwise, you won't be allowed to create or update the list's settings or check the list's status by using Macie.

- **Object paths** – If you store more than one allow list in Amazon S3, the object path for each list must be unique. In other words, each allow list must be stored separately as its own S3 object.
- **Storage classes** – An allow list must be stored directly in Amazon S3 using one of the following storage classes: S3 Intelligent-Tiering, S3 One Zone-IA, S3 Standard, or S3 Standard-IA.
- **Versioning** – When you add an allow list to an S3 bucket, we recommend that you also enable versioning for the bucket. You can then use date and time values to correlate versions of the list with runs of sensitive data discovery jobs that use the list. This can help with data privacy and protection audits or investigations that you perform.
- **Object Lock** – To prevent an allow list from being deleted or overwritten for a certain amount of time or indefinitely, you can enable Object Lock for the S3 bucket that stores the list. Enabling this setting doesn't prevent Macie from accessing the list. For information about this setting, see [Using S3 Object Lock](#) in the *Amazon Simple Storage Service User Guide*.

Encryption/Decryption requirements

If you encrypt an allow list in Amazon S3, the permissions policy for the [Macie service-linked role \(p. 302\)](#) typically grants Macie the permissions that it needs to decrypt the list. However, this depends on the type of encryption that's used:

- If a list is encrypted using server-side encryption with an Amazon S3 managed key (SSE-S3) or an AWS managed AWS KMS key (AWS managed SSE-KMS), Macie can decrypt the list. The service-linked role for your Macie account grants Macie the permissions that it needs.
- If a list is encrypted using server-side encryption with a customer managed AWS KMS key (customer managed SSE-KMS), Macie can decrypt the list only if you allow Macie to use the key. To learn how to do this, see [Allowing Macie to use a customer managed KMS key \(p. 127\)](#).
- If a list is encrypted using server-side encryption with a customer-provided key (SSE-C) or client-side encryption, Macie can't decrypt the list. Consider using SSE-S3 or SSE-KMS encryption instead.

If a list is encrypted with an AWS managed KMS key or a customer managed KMS key, your AWS Identity and Access Management (IAM) identity must also be allowed to use the key. Otherwise, you won't be allowed to create or update the list's settings or check the list's status by using Macie. To learn how to check or change the permissions for a KMS key, see [Key policies in AWS KMS](#) in the *AWS Key Management Service Developer Guide*.

For detailed information about encryption options for Amazon S3 data, see [Protecting data using encryption](#) in the *Amazon Simple Storage Service User Guide*.

Design considerations and recommendations

In general, Macie treats each entry in an allow list as a string literal value. That is to say, Macie ignores each occurrence of text that exactly matches a complete entry in an allow list. Matches aren't case sensitive.

However, Macie uses the entries as part of a larger data extraction and analysis framework. The framework includes machine learning and pattern matching functions that factor dimensions such as grammatical and syntactical variations and, in many cases, keyword proximity. The framework also factors an S3 object's file type or storage format. Therefore, keep the following considerations and recommendations in mind as you add and manage the entries in an allow list.

Prepare for different file types and storage formats

For unstructured data, such as text in an Adobe Portable Document Format (.pdf) file, Macie ignores text that exactly matches a complete entry in an allow list, including text that spans multiple lines or pages.

For structured data, such as columnar data in a CSV file or record-based data in a JSON file, Macie ignores text that exactly matches a complete entry in an allow list if all the text is stored in a single field, cell, or array. This requirement doesn't apply to structured data that's stored in an otherwise unstructured file, such as a table in a .pdf file.

For example, consider the following content in a CSV file:

```
Name,Account ID
Akua Mansa,111111111111
John Doe,222222222222
```

If Akua Mansa and John Doe are entries in an allow list, Macie ignores those names in the CSV file. The complete text of each list entry is stored in a single Name field.

Conversely, consider a CSV file that contains the following columns and fields:

```
First Name,Last Name,Account ID
Akua,Mansa,111111111111
John,Doe,222222222222
```

If Akua Mansa and John Doe are entries in an allow list, Macie doesn't ignore those names in the CSV file. None of the fields in the CSV file contain the complete text of an entry in the allow list.

Include common variations

Add entries for common variations of numeric data, proper nouns, terms, and alphanumeric character sequences. For example, if you add names or phrases that contain only one space between words, also add variations that include two spaces between words. Similarly, add words and phrases that do and don't contain special characters, and consider including common syntactical and semantic variations.

For the US phone number *425-555-0100*, for example, you might add these entries to an allow list:

```
425-555-0100
425.555.0100
(425) 555-0100
+1-425-555-0100
```

For the date *February 1, 2022* in a multinational context, you might add entries that include common syntactical variations for English and French, including variations that do and don't include special characters:

```
February 1, 2022
1 février 2022
1 fevrier 2022
Feb 01, 2022
1 fév 2022
1 fev 2022
```



```
02/01/2022  
01/02/2022
```

For names of people, include entries for various forms of a name that you don't consider sensitive. For example, include: the first name followed by the last name; the last name followed by the first name, the first and last name separated by one space; the first and last name separated by two spaces; and nicknames.

For the name *Martha Rivera*, for example, you might add:

```
Martha Rivera  
Martha Rivera  
Rivera, Martha  
Rivera, Martha  
Rivera Martha  
Rivera Martha
```

If you want to ignore variations of a specific name that contains many parts, create an allow list that uses a regular expression instead. For example, for the name *Dr. Martha Lyda Rivera, PhD*, you might use the following regular expression: `^(Dr.)?Martha\s(Lyda|L\.)?\s?Rivera,?(PhD)?$`.

Options and requirements for regular expressions in allow lists

For this type of allow list, you specify a regular expression (*regex*) that defines a text pattern to ignore—for example, public phone numbers for your organization, email addresses for your organization's domain, or patterned sample data that your organization uses for testing. The regex defines a common pattern for a specific kind of data that you don't consider sensitive. If you use this type of allow list, Amazon Macie doesn't report occurrences of text that completely match the specified pattern. Unlike an allow list that specifies predefined text to ignore, you create and store the regex and all other list settings in Macie.

When you create or update this type of allow list, you can test the list's regex with sample data before you save the list. We recommend that you do this with multiple sets of sample data. If you create a regex that's too general, Macie might ignore occurrences of text that you consider sensitive. If a regex is too specific, Macie might not ignore occurrences of text that you don't consider sensitive. To protect against malformed or long-running expressions, Macie also compiles and tests the regex against a collection of sample text automatically, and notifies you of issues to address.

For additional testing, we recommend that you also test the list's regex with a small, representative set of data for your account or organization. To do this, you can [create a one-time job \(p. 97\)](#) and configure the job to use the list in addition to the managed data identifiers and custom data identifiers that you typically use to analyze data. You can then review the job's results—sensitive data findings, sensitive data discovery results, or both. If the job's results differ from what you expect, you can change and test the regex until the results are what you expect.

After you configure and test an allow list, you can create and configure additional jobs to use the list. When those jobs run, Macie uses the latest version of the list's regex to analyze data. If you configure a job to run more than once, Macie does this each time the job runs.

Topics

- [Syntax support and recommendations \(p. 75\)](#)
- [Examples \(p. 76\)](#)

Syntax support and recommendations

An allow list can specify a regular expression (*regex*) that contains as many as 512 characters. Macie supports a subset of the regex pattern syntax provided by the [Perl Compatible Regular Expressions](#)

(PCRE) library. Of the constructs provided by the PCRE library, Macie doesn't support the following pattern elements:

- Backreferences
- Capturing groups
- Conditional patterns
- Embedded code
- Global pattern flags, such as `/i`, `/m`, and `/x`
- Recursive patterns
- Positive and negative look-behind and look-ahead zero-width assertions, such as `?=`, `?!`, `?<=`, and `?<!`

To create effective regex patterns for allow lists, also note the following tips and recommendations:

- **Anchors** – Use anchors (`^` or `$`) only if you expect the pattern to appear at the beginning or end of a file, not the beginning or end of a line.
- **Bounded repeats** – For performance reasons, Macie limits the size of bounded repeat groups. For example, `\d{100,1000}` won't compile in Macie. To approximate this functionality, you can use an open-ended repeat such as `\d{100,}`.
- **Case insensitivity** – To make parts of a pattern case insensitive, you can use the `(?i)` construct instead of the `/i` flag.
- **Performance** – There's no need to optimize prefixes or alternations manually. For example, changing `/hello|hi|hey/` to `/h(?:ello|i|ey)/` won't improve performance.
- **Wildcards** – For performance reasons, Macie limits the number of repeated wildcards. For example, `a*b*a*` won't compile in Macie.
- **Alternation** – To specify more than one pattern in a single allow list, you can use the alternation operator (`|`) to concatenate the patterns. If you do this, Macie uses OR logic to combine the patterns and form a new pattern. For example, if you specify `(apple|orange)`, Macie recognizes both *apple* and *orange* as a match and ignores occurrences of both words. If you concatenate patterns, be sure to limit the overall length of the concatenated expression to 512 or fewer characters.

Finally, when you develop the regex, design it to accommodate different file types and storage formats. Macie uses the regex as part of a larger data extraction and analysis framework. The framework factors an S3 object's file type or storage format. For structured data, such as columnar data in a CSV file or record-based data in a JSON file, Macie ignores text that completely matches the pattern only if all the text is stored in a single field, cell, or array. This requirement doesn't apply to structured data that's stored in an otherwise unstructured file, such as a table in an Adobe Portable Document Format (.pdf) file. For unstructured data, such as text in a .pdf file, Macie ignores text that completely matches the pattern, including text that spans multiple lines or pages.

Examples

The following examples demonstrate valid regex patterns for some common scenarios.

Email addresses

If you use a custom data identifier to detect email addresses, you can ignore email addresses that you don't consider sensitive, such as email addresses for your organization.

To ignore email addresses for a particular second-level and top-level domain, you can use this pattern:

```
[a-zA-Z0-9_+\\-]+@example\\.com
```

Where *example* is the name of the second-level domain and *com* is the top-level domain. In this case, Macie matches and ignores addresses such as *johndoe@example.com* and *john.doe@example.com*.

To ignore email addresses for a particular domain in any generic top-level domain (gTLD), such as *.com* or *.gov*, you can use this pattern:

```
[a-zA-Z0-9_+\-\-]+@example\.[a-zA-Z]{2,}
```

Where *example* is the name of the domain. In this case, Macie matches and ignores addresses such as *johndoe@example.com*, *john.doe@example.gov*, and *johndoe@example.edu*.

To ignore email addresses for a particular domain in any one country code top-level domain (ccTLD), such as *.ca* for Canada or *.au* for Australia, you can use this pattern:

```
[a-zA-Z0-9_+\-\-]+@example\.(ca|au)
```

Where *example* is the name of the domain and *ca* and *au* are specific ccTLDs to ignore. In this case, Macie matches and ignores addresses such as *johndoe@example.ca* and *john.doe@example.au*.

To ignore email addresses that are for a particular domain and gTLD and include third- and fourth-level domains, you can use this pattern:

```
[a-zA-Z0-9_+\-\-]+@([a-zA-Z0-9-]+\.)?[a-zA-Z0-9-]+\..example\.com
```

Where *example* is the name of the domain and *com* is the gTLD. In this case, Macie matches and ignores addresses such as *johndoe@www.example.com* and *john.doe@www.team.example.com*.

Phone numbers

Macie provides managed data identifiers that can detect phone numbers for several countries and regions. To ignore certain phone numbers, such as toll-free numbers or public phone numbers for your organization, you can use patterns such as the following.

To ignore toll-free, US phone numbers that use the 800 area code and are formatted as (800) ###-####:

```
^\(800\)?[ -]?\d{3}[ -]?\d{4}$
```

To ignore toll-free, US phone numbers that use the 888 area code and are formatted as (888) ###-####:

```
^\(888\)?[ -]?\d{3}[ -]?\d{4}$
```

To ignore 10-digit, French phone numbers that include the 33 country code and are formatted as +33 ## ## ## ## ##:

```
^\+33 \d( \d\d){4}$
```

To ignore US and Canadian phone numbers that use particular area and exchange codes, don't include a country code, and are formatted as (###) ###-####:

```
^\(123\)?[ -]?555[ -]?\d{4}$
```

Where *123* is the area code and *555* is the exchange code.

To ignore US and Canadian phone numbers that use particular area and exchange codes, include a country code, and are formatted as +1 (###) ###-####:

```
^\+1\(\(123\)?[ -]?555[ -]?\d{4}$
```

Where **123** is the area code and **555** is the exchange code.

Creating and managing allow lists in Amazon Macie

In Amazon Macie, an allow list defines specific text or a text pattern that you want Macie to ignore when it inspects Amazon Simple Storage Service (Amazon S3) objects for sensitive data. If text matches an entry or pattern in an allow list, Macie doesn't report the text in sensitive data findings or sensitive data discovery results, even if the text matches the criteria of a [managed data identifier \(p. 45\)](#) or a [custom data identifier \(p. 64\)](#).

You can create and manage the following types of allow lists in Macie.

Predefined text

Use this type of list to specify words, phrases, and other kinds of character sequences that aren't sensitive, aren't likely to change, and don't necessarily adhere to a common pattern. Examples are the names of public representatives for your organization, specific phone numbers, and specific sample data that your organization uses for testing. If you use this type of list, Macie ignores text that exactly matches an entry in the list.

For this type of list, you create a line-delimited plaintext file that lists specific text to ignore. You then store the file in an S3 bucket and configure settings for Macie to access the list in the bucket. You can then create and configure sensitive data discovery jobs to use the list. When those jobs run, Macie retrieves the latest version of the list from Amazon S3 and uses that version of the list as part of its analysis. If Macie finds text that exactly matches an entry in the list, Macie doesn't report that occurrence of text as sensitive data.

Regular expression

Use this type of list to specify a regular expression (*regex*) that defines a text pattern to ignore. Examples are public phone numbers for your organization, email addresses for your organization's domain, and patterned sample data that your organization uses for testing. If you use this type of list, Macie ignores text that completely matches the regex pattern defined by the list.

For this type of list, you create a regex that defines a common pattern for text that isn't sensitive but varies or is likely to change. Unlike a list of predefined text, you create and store the regex and all other list settings in Macie. You can then create and configure sensitive data discovery jobs to use the list. When those jobs run, Macie uses the regex as part of its analysis. If Macie finds text that completely matches the pattern defined by the list, Macie doesn't report that occurrence of text as sensitive data.

For detailed requirements, recommendations, and examples of each type of list, see [Allow list options and requirements \(p. 71\)](#). You can create as many as 10 allow lists for your account in each supported AWS Region, up to five allow lists that specify predefined text and up to five allow lists that specify regular expressions. You can create and use allow lists in all the AWS Regions where Macie is currently available except the Asia Pacific (Osaka) Region.

To create and manage allow lists, you can use the Amazon Macie console or the Amazon Macie API. The following topics explain how. For the API, the topics include examples of how to perform these tasks using the [AWS Command Line Interface \(AWS CLI\)](#). You can also perform these tasks by using a current version of another AWS command line tool or an AWS SDK, or by sending HTTPS requests directly to Macie. For information about AWS tools and SDKs, see [Tools to Build on AWS](#).

Topics

- [Creating allow lists \(p. 79\)](#)
- [Checking the status of allow lists \(p. 83\)](#)

- [Changing allow lists \(p. 85\)](#)
- [Deleting allow lists \(p. 87\)](#)

Creating allow lists

How you create an allow list in Amazon Macie depends on the type of list that you want to create. An allow list can be a file that lists predefined text to ignore, or it can be a regular expression (*regex*) that defines a text pattern to ignore. Choose the section for the type of list that you want to create.

Predefined text

Before you create this type of allow list in Macie, take the following steps:

1. By using a text editor, create a line-delimited plaintext file that lists specific text to ignore. For more information, see [Syntax requirements for lists of predefined text \(p. 72\)](#).
2. Upload the file to an S3 bucket and note the name of the bucket and the object. You'll need to enter these names when you configure the settings in Macie.
3. Ensure that the settings for the S3 bucket and object allow you and Macie to retrieve the list from the bucket. For more information, see [Storage requirements for lists of predefined text \(p. 72\)](#).
4. If you encrypted the S3 object, ensure that it's encrypted with a key that you and Macie are allowed to use. For more information, see [Encryption/Decryption requirements for lists of predefined text \(p. 73\)](#).

After you take these steps, you're ready to configure the list's settings in Macie. You can configure the settings by using the Amazon Macie console or the Amazon Macie API.

Console

Follow these steps to configure the settings for an allow list by using the Amazon Macie console.

To configure allow list settings in Macie

1. Open the Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, under **Settings**, choose **Allow lists**.
3. On the **Allow lists** page, choose **Create**.
4. Under **Select a list type**, choose **Predefined text**.
5. Under **List settings**, use the following options to enter additional settings for the allow list:
 - For **Name**, enter a name for the list. The name can contain as many as 128 characters.
 - For **Description**, optionally enter a brief description of the list. The description can contain as many as 512 characters.
 - For **S3 bucket name**, enter the full name of the bucket that stores the list.

In Amazon S3, you can find this value in the **Name** field of the bucket's properties. This value is case sensitive. In addition, don't use wildcard characters or partial values when you enter the name.
 - For **S3 object name**, enter the full name of the S3 object that stores the list.

In Amazon S3, you can find this value in the **Key** field of the object's properties. If the name includes a path, be sure to include the complete path when you enter the name, for example **allowlists/macie/mylist.txt**. This value is case sensitive. In addition, don't use wildcard characters or partial values when you enter the name.
6. (Optional) Under **Tags**, choose **Add tag**, and then enter as many as 50 tags to assign to the allow list.

A *tag* is a label that you define and assign to certain types of AWS resources. Each tag consists of a required tag key and an optional tag value. Tags can help you identify, categorize, and manage resources in different ways, such as by purpose, owner, environment, or other criteria. To learn more, see [Tagging Amazon Macie resources \(p. 315\)](#).

7. When you finish, choose **Create**.

Macie tests the list's settings. Macie also verifies that it can retrieve the list from Amazon S3 and parse the list's content. If an error occurs, Macie displays a message that describes the error. For detailed information that can help you troubleshoot the error, see [Options and requirements for lists of predefined text \(p. 71\)](#). After you address any errors, you can save the list's settings.

API

To configure allow list settings programmatically, use the [CreateAllowList](#) operation of the Amazon Macie API and specify the appropriate values for the required parameters.

For the `criteria` parameter, use an `s3WordsList` object to specify the name of the S3 bucket (`bucketName`) and the name of the S3 object (`objectKey`) that stores the list. To determine the bucket name, refer to the `Name` field in Amazon S3. To determine the object name, refer to the `Key` field in Amazon S3. Note that these values are case sensitive. In addition, don't use wildcard characters or partial values when you specify these names.

To configure the settings by using the AWS CLI, run the `create-allow-list` command and specify the appropriate values for the required parameters. The following examples show how to configure the settings for an allow list that's stored in an S3 bucket named `DOC-EXAMPLE-BUCKET`. The name of the S3 object that contains the list is `allowlists/macie/mylist.txt`.

This example is formatted for Linux, macOS, or Unix, and it uses the backslash (`\`) line-continuation character to improve readability.

```
$ aws macie2 create-allow-list \
--criteria '{"s3WordsList":{"bucketName":"DOC-EXAMPLE-BUCKET","objectKey":"allowlists/
macie/mylist.txt"}}' \
--name my_allow_list \
--description "Lists public phone numbers and names for Example Corp."
```

This example is formatted for Microsoft Windows and it uses the caret (`^`) line-continuation character to improve readability.

```
C:\> aws macie2 create-allow-list ^
--criteria="{\"s3WordsList\":{\"bucketName\": \"DOC-EXAMPLE-BUCKET\", \"objectKey\":
\"allowlists/macie/mylist.txt\"}} ^
--name my_allow_list ^
--description "Lists public phone numbers and names for Example Corp."
```

When you submit your request, Macie tests the list's settings. Macie also verifies that it can retrieve the list from Amazon S3 and parse the list's content. If an error occurs, your request fails and Macie returns a message that describes the error. For detailed information that can help you troubleshoot the error, see [Options and requirements for lists of predefined text \(p. 71\)](#).

If Macie can retrieve and parse the list, your request succeeds and you receive output similar to the following.

```
{
  "arn": "arn:aws:macie2:us-west-2:123456789012:allow-list/nkr81bmtu2542yyexample",
  "id": "nkr81bmtu2542yyexample"
}
```

Where `arn` is the Amazon Resource Name (ARN) of the allow list that was created, and `id` is the unique identifier for the list.

After you save the list's settings, you can [create and configure sensitive data discovery jobs \(p. 97\)](#) to use the list. Each time those jobs start to run, Macie retrieves the latest version of the list from Amazon S3. Macie then uses that version of the list when it analyzes data.

Regular expression

When you create an allow list that specifies a regular expression (*regex*), you define the regex and all other list settings directly in Macie. Macie supports a subset of the regex pattern syntax provided by the [Perl Compatible Regular Expressions \(PCRE\) library](#). For more information, see [Syntax support and recommendations \(p. 75\)](#).

You can create this type of list by using the Amazon Macie console or the Amazon Macie API.

Console

Follow these steps to create an allow list by using the Amazon Macie console.

To create an allow list

1. Open the Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, under **Settings**, choose **Allow lists**.
3. On the **Allow lists** page, choose **Create**.
4. Under **Select a list type**, choose **Regular expression**.
5. Under **List settings**, use the following options to enter additional settings for the allow list:
 - For **Name**, enter a name for the list. The name can contain as many as 128 characters.
 - For **Description**, optionally enter a brief description of the list. The description can contain as many as 512 characters.
 - For **Regular expression**, enter the regex that defines the text pattern to ignore. The regex can contain as many as 512 characters.
6. (Optional) For **Evaluate**, enter up to 1,000 characters in the **Sample data** box, and then choose **Test** to test the regex. Macie evaluates the sample data and reports the number of occurrences of text that matches the regex. You can repeat this step as many times as you like to refine and optimize the regex.

Note

We recommend that you test and refine the regex with multiple sets of sample data. If you create a regex that's too general, Macie might ignore occurrences of text that you consider sensitive. If a regex is too specific, Macie might not ignore occurrences of text that you don't consider sensitive.

7. (Optional) Under **Tags**, choose **Add tag**, and then enter as many as 50 tags to assign to the allow list.

A *tag* is a label that you define and assign to certain types of AWS resources. Each tag consists of a required tag key and an optional tag value. Tags can help you identify, categorize, and manage resources in different ways, such as by purpose, owner, environment, or other criteria. To learn more, see [Tagging Amazon Macie resources \(p. 315\)](#).

8. When you finish, choose **Create**.

Macie tests the list's settings. Macie also tests the regex to verify that it can compile the expression. If an error occurs, Macie displays a message that describes the error. For detailed information that can help you troubleshoot the error, see [Options and requirements for regular expressions in allow lists \(p. 75\)](#). After you address any errors, you can save the allow list.

API

Before you create this type of allow list in Macie, we recommend that you test and refine the regular expression with multiple sets of sample data. If you create a regex that's too general, Macie might ignore occurrences of text that you consider sensitive. If a regex is too specific, Macie might not ignore occurrences of text that you don't consider sensitive.

To test an expression with Macie, you can use the [TestCustomDataIdentifier](#) operation of the Amazon Macie API or, for the AWS CLI, run the [test-custom-data-identifier](#) command. Macie uses the same underlying code to compile expressions for allow lists and custom data identifiers. If you test an expression in this way, be sure to specify values only for the `regex` and `sampleText` parameters. Otherwise, you'll receive inaccurate results.

When you're ready to create this type of allow list, use the [CreateAllowList](#) operation of the Amazon Macie API and specify the appropriate values for the required parameters. For the `criteria` parameter, use the `regex` field to specify the regular expression that defines the text pattern to ignore. The expression can contain as many as 512 characters.

To create this type of list by using the AWS CLI, run the [create-allow-list](#) command and specify the appropriate values for the required parameters. The following examples create an allow list named `my_allow_list`. The regex is designed to ignore all email addresses that a custom data identifier might otherwise detect for the `example.com` domain.

This example is formatted for Linux, macOS, or Unix, and it uses the backslash (`\`) line-continuation character to improve readability.

```
$ aws macie2 create-allow-list \
--criteria '{"regex":"[a-z]@example.com"}' \
--name my_allow_list \
--description "Ignores all email addresses for Example Corp."
```

This example is formatted for Microsoft Windows and it uses the caret (`^`) line-continuation character to improve readability.

```
C:\> aws macie2 create-allow-list ^
--criteria={"regex":"[a-z]@example.com"} ^
--name my_allow_list ^
--description "Ignores all email addresses for Example Corp."
```

When you submit your request, Macie tests the list's settings. Macie also tests the regex to verify that it can compile the expression. If an error occurs, the request fails and Macie returns a message that describes the error. For detailed information that can help you troubleshoot the error, see [Options and requirements for regular expressions in allow lists \(p. 75\)](#).

If Macie can compile the expression, the request succeeds and you receive output similar to the following:

```
{
  "arn": "arn:aws:macie2:us-west-2:123456789012:allow-list/km2d4y22hp6rv05example",
  "id": "km2d4y22hp6rv05example"
}
```

Where `arn` is the Amazon Resource Name (ARN) of the allow list that was created, and `id` is the unique identifier for the list.

After you save the list, you can [create and configure sensitive data discovery jobs \(p. 97\)](#) to use it. Each time those jobs start to run, Macie retrieves the latest regex for the list. Macie then uses that regex when it analyzes data.

Checking the status of allow lists

If you configure sensitive data discovery jobs to use allow lists, it's important to check the status of your lists periodically. Otherwise, errors might cause your jobs to produce unexpected results, such as sensitive data findings for text that you specified in a list. If a job starts to run and Amazon Macie can't access or use an allow list for the job, the job continues to run. However, Macie doesn't use the list when it inspects S3 objects that you configured the job to analyze.

Errors are unlikely to occur for an allow list that specifies a regular expression (*regex*). This is partly because Macie automatically tests the regex when you create or update the list's settings. In addition, you store the regex and all other list settings in Macie.

However, errors can occur for an allow list that specifies predefined text, partly because you store the list in Amazon S3, not Macie. Common causes of errors are:

- The S3 bucket or object is deleted.
- The S3 bucket or object is renamed and the list's settings in Macie don't specify the new name.
- The bucket's permissions settings are changed and Macie loses access to the bucket and the object.
- The encryption settings for the bucket are changed and Macie can't decrypt the object.
- The policy for the encryption key is changed and Macie loses access to the key. Macie can't decrypt the object.

Important


Because these errors affect your job results, we recommend that you check the status of your allow lists periodically. We recommend that you also do this if you change the permissions or encryption settings for an S3 bucket that stores an allow list, or you change the policy for an AWS Key Management Service (AWS KMS) key that's used to encrypt a list.

You can check the status of your allow lists by using the Amazon Macie console or the Amazon Macie API. For detailed information that can help you troubleshoot errors that occur, see [Options and requirements for lists of predefined text \(p. 71\)](#).

Console

Follow these steps to check the status of your allow lists by using the Amazon Macie console.

To check the status of your allow lists

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, under **Settings**, choose **Allow lists**.
3. On the **Allow lists** page, choose refresh (). Macie tests the settings for all of your allow lists and updates the **Status** field to indicate the current status of each list.

If a list specifies a regular expression, its status is typically **OK**. This means that Macie can compile the expression. If a list specifies predefined text, its status can be any of the following values.

OK

Macie can retrieve and parse the contents of the list.


Access denied

Macie isn't allowed to access the S3 object that contains the list. Amazon S3 denied the request to retrieve the object. A list can also have this status if the object is encrypted with a customer managed AWS KMS key that Macie isn't allowed to use.

To address this error, review the bucket policy and other permissions settings for the bucket and the object. Ensure that Macie is allowed to access and retrieve the object. If the object is encrypted with a customer managed AWS KMS key, also review the key policy and ensure that Macie is allowed to use the key.

Error

A transient or internal error occurred when Macie attempted to retrieve or parse the contents of the list. An allow list can also have this status if it's encrypted with an encryption key that Amazon S3 and Macie can't access or use.

To address this error, wait a few minutes and then choose refresh () again. If the status continues to be **Error**, check the encryption settings for the S3 object. Ensure that the object is encrypted with a key that Amazon S3 and Macie can access and use.

Object is empty

Macie can retrieve the list from Amazon S3 but the list doesn't contain any content.

To address this error, download the object from Amazon S3 and ensure that it contains the correct entries. If the entries are correct, review the list's settings in Macie. Ensure that the specified bucket and object names are correct.

Object not found

The list doesn't exist in Amazon S3.

To address this error, review the list's settings in Macie. Ensure that the specified bucket and object names are correct.


Quota exceeded

Macie can access the list in Amazon S3. However, the number of entries in the list or the storage size of the list exceeds the quota for an allow list.

To address this error, break the list into multiple files. Ensure that each file contains fewer than 100,000 entries. Also ensure that the size of each file is less than 35 MB. Then, upload each file to Amazon S3. When you finish, configure allow list settings in Macie for each file. You can have as many as five lists of predefined text in each supported AWS Region.

Throttled

Amazon S3 throttled the request to retrieve the list.

To address this error, wait a few minutes and then choose refresh () again.

User access denied

Amazon S3 denied the request to retrieve the object. If the specified object exists, you're not allowed to access it or it's encrypted with an AWS KMS key that you're not allowed to use.

To address this error, work with your AWS administrator to ensure that the list's settings specify the correct bucket and object names, and you have read access to the bucket and the object. If the object is encrypted, also ensure that it's encrypted with a key that you're allowed to use.

4. To review the settings and status of a specific list, choose the list's name.

API

To check the status of an allow list programmatically, use the [GetAllowList](#) operation of the Amazon Macie API or, for the AWS CLI, run the [get-allow-list](#) command.

For the `id` parameter, specify the unique identifier for the allow list whose status you want to check. To get this identifier, you can use the [ListAllowLists](#) operation. The **ListAllowLists** operation retrieves information about all the allow lists for your account. If you're using the AWS CLI, you can run the [list-allow-lists](#) command to retrieve this information.

When you submit a **GetAllowList** request, Macie tests all the settings for the allow list. If the settings specify a regular expression (*regex*), Macie verifies that it can compile the expression. If the settings specify a list of predefined text, Macie verifies that it can retrieve and parse the list.

Macie then returns a `GetAllowListResponse` object that provides the details of the allow list. In the `GetAllowListResponse` object, the `status` object indicates the current status of the list: a status code (`code`) and, depending on the status code, a brief description of the list's status (`description`).

If the allow list specifies a *regex*, the status code is typically `OK` and there isn't an associated description. This means that Macie compiled the expression successfully.

If the allow list specifies predefined text, the status code varies depending on the test results:

- If Macie retrieved and parsed the list successfully, the status code is `OK` and there isn't an associated description.
- If an error prevented Macie from retrieving or parsing the list, the status code and description indicate the nature of the error that occurred.

For a list of possible status codes and a description of each one, see the [AllowListStatus](#) table in the *Amazon Macie API Reference*.

Changing allow lists

After you create an allow list, you can change most of the list's settings in Amazon Macie. For example, you can change the list's name and description, and you can add and edit the list's tags. The only setting that you can't change is a list's type. For example, if an existing allow list specifies a regular expression, you can't change its type to predefined text.

If an allow list specifies predefined text, you can also change the entries in the list. To do this, update the file that contains the entries, and then upload the new version of the file to Amazon S3. The next time Macie prepares to use the list, Macie retrieves the latest version of the file from Amazon S3. When you upload the new file, ensure that you store it in the same S3 bucket and object. Or, if you change the name of the bucket or object, ensure that you update the list's settings in Macie.

You can change an allow list's settings by using the Amazon Macie console or the Amazon Macie API.

Console

Follow these steps to change the settings for an allow list by using the Amazon Macie console.

To change an allow list

1. Open the Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, under **Settings**, choose **Allow lists**.
3. On the **Allow lists** page, choose the name of the allow list that you want to change. The allow list page opens and displays the current settings for the list.

- To assign or edit tags for the allow list, choose **Manage tags** in the **Tags** section. Then change the tags as necessary. When you finish, choose **Save**.
- To change other settings for the allow list, choose **Edit** in the **List settings** section. Then change the settings that you want:
 - Name** – Enter a new name for the list. The name can contain as many as 128 characters.
 - Description** – Enter a new description of the list. The description can contain as many as 512 characters.
 - If the allow list specifies predefined text:
 - S3 bucket name** – Enter the full name of the bucket that currently stores the list.
In Amazon S3, you can find this value in the **Name** field of the bucket's properties. This value is case sensitive. In addition, don't use wildcard characters or partial values when you enter the name.
 - S3 object name** – Enter the full name of the S3 object that currently stores the list.
In Amazon S3, you can find this value in the **Key** field of the object's properties. If the name includes a path, be sure to include the complete path when you enter the name, for example `allowlists/macie/mylist.txt`. This value is case sensitive. In addition, don't use wildcard characters or partial values when you enter the name.
 - If the allow list specifies a regular expression (*regex*), enter a new regex in the **Regular expression** box. The regex can contain as many as 512 characters.
After you enter the new regex, optionally test it. To do this, enter up to 1,000 characters in the **Sample data** box, and then choose **Test**. Macie evaluates the sample data and reports the number of occurrences of text that matches the regex. You can repeat this step as many times as you like to refine and optimize the regex before you save your changes.

When you finish changing the settings, choose **Save**.

Macie tests the list's settings. For a list of predefined text, Macie also verifies that it can retrieve the list from Amazon S3 and parse the list's content. For a regex, Macie also verifies that it can compile the expression. If an error occurs, Macie displays a message that describes the error. For detailed information that can help you troubleshoot the error, see [Allow list options and requirements](#) (p. 71). After you address any errors, you can save your changes.

API

To change an allow list programmatically, use the [UpdateAllowList](#) operation of the Amazon Macie API or, for the AWS CLI, run the `update-allow-list` command. In your request, use the supported parameters to specify a new value for each setting that you want to change. Note that the `criteria`, `id`, and `name` parameters are required. If you don't want to change the value for a required parameter, specify the current value for the parameter.

For example, the following command changes the name and description of an existing allow list. The example is formatted for Microsoft Windows and it uses the caret (^) line-continuation character to improve readability.

```
C:\> aws macie2 update-allow-list ^
--id km2d4y22hp6rv05example ^
--name my_allow_list-email ^
--criteria={"regex\":\"[a-z]@example.com\"} ^
--description "Ignores all email addresses for the example.com domain"
```

Where:

- `km2d4y22hp6rv05example` is the unique identifier for the list.

- `my_allow_list-email` is the new name for the list.
- `[a-z]@example.com` is the list's criteria, a regular expression.
- `Ignores all email addresses for the example.com domain` is the new description for the list.

When you submit your request, Macie tests the list's settings. If the list specifies predefined text, this includes verifying that Macie can retrieve the list from Amazon S3 and parse the list's content. If the list specifies a regex, this includes verifying that Macie can compile the expression.

If an error occurs when Macie tests the settings, your request fails and Macie returns a message that describes the error. For detailed information that can help you troubleshoot the error, see [Allow list options and requirements \(p. 71\)](#). If the request fails for another reason, Macie returns an HTTP 4xx or 500 response that indicates why the operation failed.

If your request succeeds, Macie updates the list's settings and you receive output similar to the following.

```
{
  "arn": "arn:aws:macie2:us-west-2:123456789012:allow-list/km2d4y22hp6rv05example",
  "id": "km2d4y22hp6rv05example"
}
```

Where `arn` is the Amazon Resource Name (ARN) of the allow list that was updated, and `id` is the unique identifier for the list.

Deleting allow lists

When you delete an allow list in Amazon Macie, you permanently delete all the list's settings. These settings can't be recovered after they're deleted. If the settings specify a list of predefined text that you store in Amazon S3, Macie doesn't delete the S3 object that contains the list. Only the settings in Macie are deleted.

If you configure sensitive data discovery jobs to use an allow list and you subsequently delete the list, the jobs will run as scheduled. However, your job results, both sensitive data findings and sensitive data discovery results, might report text that you previously specified in an allow list.

Before you delete an allow list, we recommend that you [review your job inventory \(p. 116\)](#) to identify jobs that use the list and are scheduled to run in the future. In the inventory, the details panel indicates whether a job is configured to use any allow lists and, if so, which ones. You might determine that it's best to change a list instead of deleting it.

As an additional safeguard, Macie checks the settings for all of your jobs when you try to delete an allow list. If you configured jobs to use the list and any of those jobs have a status other than **Complete** or **Cancelled**, Macie doesn't delete the list unless you provide additional confirmation.

You can delete an allow list by using the Amazon Macie console or the Amazon Macie API.

Console

Follow these steps to delete an allow list by using the Amazon Macie console.

To delete an allow list

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, under **Settings**, choose **Allow lists**.
3. On the **Allow lists** page, select the check box for the allow list that you want to delete.

4. On the **Actions** menu, choose **Delete**.
5. When prompted for confirmation, enter **delete**, and then choose **Delete**.

API

To delete an allow list programmatically, use the [DeleteAllowList](#) operation of the Amazon Macie API. For the `id` parameter, specify the unique identifier for the allow list to delete. You can get this identifier by using the [ListAllowLists](#) operation. The [ListAllowLists](#) operation retrieves information about all the allow lists for your account. If you're using the AWS CLI, you can run the [list-allow-lists](#) command to retrieve this information.

For the `ignoreJobChecks` parameter, specify whether to force deletion of the list, even if sensitive data discovery jobs are configured to use the list:

- If you specify `false`, Macie checks the settings for all of your jobs that have a status other than `COMPLETE` or `CANCELLED`. If none of those jobs are configured to use the list, Macie deletes the list permanently. If any of those jobs are configured to use the list, Macie rejects your request and returns an HTTP 400 (`ValidationException`) error. The error message indicates the number of applicable jobs for up to 200 jobs.
- If you specify `true`, Macie deletes the list permanently without checking the settings for any of your jobs.

To delete an allow list by using the AWS CLI, run the [delete-allow-list](#) command. For example:

```
C:\> aws macie2 delete-allow-list --id nkr81bmtu2542yyexample --ignore-job-checks false
```

Where `nkr81bmtu2542yyexample` is the unique identifier for the allow list to delete.

If your request succeeds, Macie returns an empty HTTP 200 response. Otherwise, Macie returns an HTTP 4xx or 500 response that indicates why the operation failed.

If the allow list specified predefined text, you can optionally delete the S3 object that contains the list. However, keeping this object can help ensure that you have an immutable history of sensitive data findings and sensitive data discovery results for data privacy and protection audits or investigations.

Running sensitive data discovery jobs in Amazon Macie

With Amazon Macie, you create and run sensitive data discovery jobs to automate discovery, logging, and reporting of sensitive data in Amazon Simple Storage Service (Amazon S3) buckets. A *sensitive data discovery job* is a series of automated processing and analysis tasks that Macie performs to analyze objects in S3 buckets and determine whether the objects contain sensitive data. Each job provides detailed reports of the sensitive data that Macie finds and the analysis that Macie performs.

To help you meet and maintain compliance with your data security and privacy requirements, Macie provides several options for scheduling and defining the scope of each job. With these options, you can build and maintain a comprehensive view of the data that your organization stores in Amazon S3 and any security or compliance risks for that data.

You can configure a job to run only once for on-demand analysis and assessment, or on a recurring basis for periodic analysis, assessment, and monitoring. You also define the breadth and depth of each job's

analysis. When you create a job, you start by specifying which S3 buckets contain objects that you want the job to analyze—specific buckets that you select or buckets that match specific criteria. You can then refine the scope of that analysis by choosing additional options. The options include custom include and exclude criteria that derive from properties of S3 objects, such as tags, prefixes, and the date when an object was last modified.

You also specify the types of sensitive data that you want to detect. You can configure a job to use [managed data identifiers \(p. 45\)](#) that Macie provides, [custom data identifiers \(p. 64\)](#) that you define, or a combination of the two. By selecting specific managed and custom data identifiers for a job, you can tailor the analysis to focus on specific types of sensitive data. To fine tune the analysis, you can also configure a job to use [allow lists \(p. 70\)](#) that you define. Allow lists specify text and text patterns that you want Macie to ignore, typically sensitive data exceptions for your organization's particular scenarios or environment.

Each job produces records of the sensitive data that Macie finds and the analysis that Macie performs—*sensitive data findings* and *sensitive data discovery results*. A *sensitive data finding* is a detailed report of sensitive data that Macie found in an object. A *sensitive data discovery result* is a record that logs details about the analysis of an object. Macie creates a sensitive data discovery result for each object that you configure a job to analyze. This includes objects that don't contain sensitive data and therefore don't produce sensitive data findings. Each type of record adheres to a standardized schema, which can help you query, monitor, and process the records to meet your security and compliance requirements.

Topics

- [Scope options for sensitive data discovery jobs \(p. 89\)](#)
- [Creating a sensitive data discovery job \(p. 97\)](#)
- [Monitoring sensitive data discovery jobs with Amazon CloudWatch Logs \(p. 104\)](#)
- [Reviewing statistics and results for sensitive data discovery jobs \(p. 113\)](#)
- [Managing sensitive data discovery jobs \(p. 116\)](#)
- [Forecasting and monitoring costs for sensitive data discovery jobs \(p. 122\)](#)

Scope options for sensitive data discovery jobs

In Amazon Macie, you define the scope of the data that a sensitive data discovery job analyzes. To help you do this, Macie provides several job-specific options that you can choose when you create and configure a job.

Scope options

- [S3 buckets \(p. 89\)](#)
- [Include existing S3 objects \(p. 94\)](#)
- [Sampling depth \(p. 94\)](#)
- [S3 object criteria \(p. 95\)](#)

S3 buckets

The first step in creating a sensitive data discovery job is to specify which Amazon Simple Storage Service (Amazon S3) buckets contain objects that you want the job to analyze. You can do this in either of two ways, by selecting specific S3 buckets from your bucket inventory or by specifying custom criteria that derive from properties of S3 buckets.

Selecting specific buckets

With this option, you explicitly select each S3 bucket that you want the job to analyze. Then, when the job runs, it analyzes objects in the selected buckets. If you also configure the job to run

periodically on a daily, weekly, or monthly basis, the job analyzes objects in those same buckets each time it runs.

This configuration is helpful for cases where you prefer to perform targeted analysis of a specific set of data. It gives you precise, predictable control over which buckets a job analyzes.

Specifying bucket criteria

With this option, you define runtime criteria that determine which S3 buckets the job analyzes. The criteria consist of one or more conditions that derive from bucket properties, such as public access settings and tags. When the job runs, it identifies buckets that match your criteria and then analyzes objects in those buckets. If you also configure the job to run periodically, the job does this each time it runs. Consequently, the job might analyze objects in different buckets each time it runs, depending on changes to your bucket inventory and the criteria that you define.

This configuration is helpful for cases where you want the scope of the job's analysis to dynamically adapt to changes to your bucket inventory. For example, if you configure a job to use bucket criteria and run periodically, the job can automatically identify new buckets that match the criteria and inspect those buckets for sensitive data.

The topics in this section provide additional details about each option.

Topics

- [Selecting S3 buckets \(p. 90\)](#)
- [Specifying S3 bucket criteria \(p. 92\)](#)

Selecting S3 buckets

If you choose to explicitly select each S3 bucket that you want a job to analyze, Macie provides you with a complete inventory of your buckets in the current AWS Region. You can then review your inventory and select the buckets that you want. To learn how Macie generates and maintains this inventory for you, see [How Macie monitors Amazon S3 data \(p. 16\)](#).

If you're the Macie administrator for an organization, the inventory includes buckets that are owned by member accounts in your organization. You can select as many as 1,000 of these buckets, spanning as many as 1,000 accounts.

To help you make your bucket selections, the inventory provides details and statistics for each bucket. This includes the amount of data that a job can analyze in each bucket—*classifiable objects* are objects that use a supported Amazon S3 storage class (S3 Intelligent-Tiering, S3 One Zone-IA, S3 Standard, or S3 Standard-IA) and have a file name extension for a [supported file or storage format \(p. 124\)](#). The inventory also indicates whether any existing jobs are configured to analyze objects in a bucket. These details can help you estimate the breadth of a job and refine your bucket selections.

In the inventory table:



- **Classifiable objects** – This field indicates the total number of objects that the job can analyze in a bucket.
- **Classifiable size** – This field indicates the total storage size of all the objects that the job can analyze in a bucket.


If a bucket contains compressed objects, this value doesn't reflect the actual size of those objects after they're decompressed. If versioning is enabled for a bucket, this value is based on the storage size of the latest version of each object in the bucket.

- **Monitored** – This field indicates whether any existing jobs are configured to periodically analyze objects in a bucket on a daily, weekly, or monthly basis.

If the value for this field is **Yes**, the bucket is explicitly included in a periodic job or the bucket matched the criteria for a periodic job within the past 24 hours. In addition, the status of at least one of those jobs is not *Cancelled*. Macie updates this data on a daily basis.

- **Latest job run** – If any existing periodic or one-time jobs are configured to analyze objects in a bucket, this field indicates the most recent time when one of those jobs started to run. Otherwise, this field is empty.

If the information icon () appears next to any bucket names in the table, we recommend that you retrieve the latest bucket metadata from Amazon S3. To do this, choose refresh () above the table. The information icon indicates that a bucket was created during the past 24 hours, possibly after Macie last retrieved bucket and object metadata from Amazon S3 as part of the daily refresh cycle. For more information, see [Data refreshes \(p. 18\)](#).




If the warning icon () appears next to a bucket's name in the table, Macie isn't allowed to access the bucket or the bucket's objects. (Macie can only provide a subset of information about the bucket, such as the bucket's name.) This means that the job won't be able to analyze objects in the bucket. To investigate the issue, review the bucket's policy and permissions settings in Amazon S3. For example, the bucket might have a restrictive bucket policy. For more information, see [Allowing Macie to access S3 buckets and objects \(p. 40\)](#).

To customize your view of the inventory and find specific buckets more easily, you can filter the table by entering filter criteria in the filter bar. The following table provides some examples.

To show all buckets that...	Apply this filter...
Are owned by a specific account	Account ID = <i>the 12-digit ID for the account</i>
Are publicly accessible	Effective permission = Public
Aren't included in any periodic jobs	Actively monitored by job = False
Aren't included in any periodic or one-time jobs	Defined in job = False
Have a specific tag key*	Tag key = <i>the tag key</i>
Have a specific tag value*	Tag value = <i>the tag value</i>
Contain unencrypted objects (or use client-side encryption)	Object count by encryption is No encryption and From = 1

* Tag keys and values are case sensitive. Also, you have to specify a complete, valid value for these fields in a filter. You can't specify partial values or use wildcard characters.

To display the details of a bucket, choose the bucket's name and refer to the details panel. From there, you can also:

- Pivot and drill down on certain fields by choosing a magnifying glass for the field. Choose  to show buckets with the same value, or choose  to show buckets with other values.
- Retrieve the latest metadata for objects in the bucket. This can be helpful if you recently created a bucket or made significant changes to the bucket's objects during the past 24 hours. To retrieve the data, choose refresh () in the **Object statistics** section of the panel. This option is available for buckets that contain 30,000 or fewer objects.

Specifying S3 bucket criteria

If you choose to specify bucket criteria for a job, Macie provides options for defining and testing the criteria. These are runtime criteria that determine which S3 buckets contain objects for the job to analyze. Each time the job runs, it identifies buckets that match your criteria and then analyzes objects in the appropriate buckets. If you're the Macie administrator for an organization, this includes buckets that are owned by member accounts in your organization.

Defining bucket criteria

Bucket criteria consist of one or more conditions that derive from properties of S3 buckets. Each condition, also referred to as a *criterion*, consists of the following parts:

- A property-based field, such as **Account ID** or **Effective permission**.
- An operator, either *equals* (eq) or *not equals* (neq).
- One or more values.
- An include or exclude statement that indicates whether you want the job to analyze (*include*) or skip (*exclude*) buckets that match the condition.

If you specify more than one value for a field, Macie uses OR logic to join the values. If you specify more than one condition for the criteria, Macie uses AND logic to join the conditions. In addition, exclude conditions take precedence over include conditions. For example, if you include buckets that are publicly accessible and exclude buckets that have specific tags, the job analyzes objects in any bucket that's publicly accessible unless the bucket has one of the specified tags.

You can define conditions that derive from any of the following property-based fields for S3 buckets.

Account ID

The unique identifier (ID) for the AWS account that owns a bucket. To specify multiple values for this field, enter the ID for each account and separate each entry with a comma.

Note that Macie doesn't support use of wildcard characters or partial values for this field.

Bucket name

The name of a bucket. This field correlates to the **Name** field, not the **Amazon Resource Name (ARN)** field, in Amazon S3. To specify multiple values for this field, enter the name of each bucket and separate each entry with a comma.

Note that values are case sensitive. In addition, Macie doesn't support use of wildcard characters or partial values for this field.

Effective permission

Specifies whether a bucket is publicly accessible. You can choose one or more of the following values for this field:

- **Not public** – The general public doesn't have read or write access to the bucket.
- **Public** – The general public has read or write access to the bucket.
- **Unknown** – Macie wasn't able to evaluate the public access settings for the bucket.

To determine this value for a bucket, Macie analyzes a combination of account- and bucket-level settings for the bucket: the block public access setting for the account; the block public access setting for the bucket; the bucket policy for the bucket; and, the access control list (ACL) for the bucket.

Shared access

Specifies whether a bucket is shared with other AWS accounts. You can choose one or more of the following values for this field:

- **External** – The bucket is shared with accounts that aren't in the same organization.
- **Internal** – The bucket is shared with accounts in the same organization.
- **Not shared** – The bucket isn't shared with other accounts.
- **Unknown** – Macie wasn't able to evaluate the shared access settings for the bucket.

To determine this value for a bucket, Macie analyzes the bucket policy and ACL for the bucket. In addition, an *organization* is defined as a set of Macie accounts that are centrally managed as a group of related accounts through AWS Organizations or by Macie invitation.

Tags

The tags that are associated with a bucket. Tags are labels that you can define and assign to certain types of AWS resources, including S3 buckets. Each tag consists of a required tag key and an optional tag value. For information about tagging S3 buckets, see [Using cost allocation S3 bucket tags](#) in the *Amazon Simple Storage Service User Guide*.

For a sensitive data discovery job, you can use this type of condition to include or exclude buckets that have a specific tag key, a specific tag value, or a specific tag key and tag value (as a pair). For example:

- If you specify **Project** as a tag key and don't specify any tag values for a condition, any bucket that has the *Project* tag key matches the condition's criteria, regardless of the tag values that are associated with that tag key.
- If you specify **Development** and **Test** as tag values and don't specify any tag keys for a condition, any bucket that has the **Development** or **Test** tag value matches the condition's criteria, regardless of the tag keys that are associated with those tag values.

To specify multiple tag keys in a condition, enter each tag key in the **Key** field and separate each entry with a comma. To specify multiple tag values in a condition, enter each tag value in the **Value** field and separate each entry with a comma.

Note that tag keys and values are case sensitive. In addition, Macie doesn't support use of wildcard characters or partial values in tag conditions.

Testing bucket criteria

While you define your bucket criteria, you can test and refine the criteria by previewing the results. To do this, expand the **Preview the criteria results** section that appears below the criteria on the console. This section displays a table of all the buckets that currently match the criteria.

The table also provides insight into the amount of data that the job can analyze in each bucket—*classifiable objects* are objects that use a supported Amazon S3 storage class (S3 Intelligent-Tiering, S3 One Zone-IA, S3 Standard, or S3 Standard-IA) and have a file name extension for a [supported file or storage format](#) (p. 124). The table also indicates whether any existing jobs are configured to periodically analyze objects in a bucket.

In the table:

- **Classifiable objects** – This field indicates the total number of objects that the job can analyze in a bucket.
- **Classifiable size** – This field indicates the total storage size of all the objects that the job can analyze in a bucket.

If a bucket contains compressed objects, this value doesn't reflect the actual size of those objects after they're decompressed. If versioning is enabled for a bucket, this value is based on the storage size of the latest version of each object in the bucket.

- **Monitored** – This field indicates whether any existing jobs are configured to periodically analyze objects in a bucket on a daily, weekly, or monthly basis.

If the value for this field is **Yes**, the bucket is explicitly included in a periodic job or the bucket matched the criteria for a periodic job within the past 24 hours. In addition, the status of at least one of those jobs is not *Cancelled*. Macie updates this data on a daily basis.

If the warning icon (⚠) appears next to a bucket's name, Macie isn't allowed to access the bucket or the bucket's objects. (Macie can only provide a subset of information about the bucket, such as the bucket's name.) This means that the job won't be able to analyze objects in the bucket. To investigate the issue, review the bucket's policy and permissions settings in Amazon S3. For example, the bucket might have a restrictive bucket policy. For more information, see [Allowing Macie to access S3 buckets and objects](#) (p. 40).

To refine the bucket criteria for the job, use the filter settings to add, change, or remove conditions from the criteria. Macie then updates the table to reflect your changes.

Include existing S3 objects

You can use sensitive data discovery jobs to perform ongoing, incremental analysis of objects in S3 buckets. If you configure a job to run periodically, Macie does this for you automatically—each run analyzes only those objects that are created or changed after the preceding run. With the **Include existing objects** option, you choose the starting point for the first increment:

- To analyze all existing objects immediately after you finish creating the job, select the check box for this option.
- To wait and analyze only those objects that are created or changed after you create the job and before the first run, clear the check box for this option.

Clearing this check box is helpful for cases where you've already analyzed the data and want to continue to analyze it periodically. For example, if you previously used another service or application to classify data and you recently started using Macie, you might use this option to ensure continued discovery and classification of your data without incurring unnecessary costs or duplicating classification data.

Each subsequent run of a periodic job automatically analyzes only those objects that are created or changed after the preceding run.

For both periodic and one-time jobs, you can also configure a job to analyze only those objects that are created or changed before or after a certain time or during a certain time range. To do this, add [object criteria](#) (p. 95) that use the last modified date for objects.

Sampling depth

With this option, you specify the percentage of eligible S3 objects that you want a sensitive data discovery job to analyze. If this value is less than 100%, Macie selects eligible objects to analyze at random, up to the specified percentage, and analyzes all the data in those objects. For example, if you configure a job to analyze 10,000 objects and you specify a sampling depth of 20%, the job analyzes approximately 2,000 randomly selected, eligible objects.

Reducing the sampling depth of a job can lower the cost and reduce the duration of a job. It's helpful for cases where the data in objects is highly consistent and you want to determine whether an S3 bucket, rather than each object, contains sensitive data.

Note that this option controls the percentage of *objects* that are analyzed, not the percentage of *bytes* that are analyzed. If you enter a sampling depth that's less than 100%, Macie analyzes all the data in each selected object, not that percentage of the data in each selected object.

S3 object criteria

To fine tune the scope of a sensitive data discovery job, you can also define custom criteria that determine which S3 objects are included or excluded from a job's analysis. These criteria consist of one or more conditions that derive from properties of S3 objects. The conditions apply to objects in all the S3 buckets that a job is configured to analyze. If a bucket contains multiple versions of an object, the conditions apply to the latest version of the object.

If you define multiple conditions as object criteria, Macie uses AND logic to join the conditions. In addition, exclude conditions take precedence over include conditions. For example, if you include objects that have the .pdf file name extension and exclude objects that are larger than 5 MB, the job analyzes any object that has the .pdf file name extension, unless the object is larger than 5 MB.

You can define conditions that derive from any of the following properties of S3 objects.

File name extension

This correlates to the file name extension of an S3 object. You can use this type of condition to include or exclude objects based on file type. To do this for multiple types of files, enter the file name extension for each type and separate each entry with a comma—for example: **docx, pdf, xlsx**. If you enter multiple file name extensions as values for a condition, Macie uses OR logic to join the values.

Note that values are case sensitive. In addition, Macie doesn't support the use of partial values or wildcard characters in this type of condition.

For information about the types of files that Macie can analyze, see [Supported file and storage formats \(p. 124\)](#).

Last modified

This correlates to the **Last modified** field in Amazon S3. In Amazon S3, this field stores the date and time when an S3 object was created or last changed, whichever is latest.

For a sensitive data discovery job, this condition can be a specific date, a specific date and time, or an exclusive time range:

- To analyze objects that were last modified after a certain date or date and time, enter the values in the **From** fields.
- To analyze objects that were last modified before a certain date or date and time, enter the values in the **To** fields.
- To analyze objects that were last modified during a certain time range, use the **From** fields to enter the values for the first date or date and time in the time range. Use the **To** fields to enter the values for the last date or date and time in the time range.
- To analyze objects that were last modified at any time during a certain single day, enter the date in the **From** date field. Enter the date for the next day in the **To** date field. Then verify that both time fields are blank. (Macie treats a blank time field as 00:00:00.) For example, to analyze objects that changed on August 9, 2020, enter **2020/08/09** in the **From** date field, enter **2020/08/10** in the **To** date field, and don't enter a value in either time field.

Enter any time values in Coordinated Universal Time (UTC) and use 24-hour notation.

Prefix

This correlates to the **Key** field in Amazon S3. In Amazon S3, this field stores the name of an S3 object, including the object's prefix. A *prefix* is similar to a directory path within a bucket. It enables you to group similar objects together in a bucket, much like you might store similar files together in a folder on a file system. For information about object prefixes and folders in Amazon S3, see

[Organizing objects in the Amazon S3 console using folders](#) in the *Amazon Simple Storage Service User Guide*.

You can use this type of condition to include or exclude objects whose keys (names) begin with a certain value. For example, to exclude all objects whose key begins with *AWSLogs*, enter **AWSLogs** as the value for a **Prefix** condition, and then choose **Exclude**.

If you enter multiple prefixes as values for a condition, Macie uses OR logic to join the values. For example, if you enter **AWSLogs1** and **AWSLogs2** as values for a condition, any object whose key begins with *AWSLogs1* or *AWSLogs2* matches the condition's criteria.

When you enter a value for a **Prefix** condition, keep the following in mind:

- Values are case sensitive.
- Macie doesn't support the use of wildcard characters in these values.
- In Amazon S3, an object's key doesn't include the name of the bucket that contains the object. For this reason, don't specify bucket names in these values.
- If a prefix includes a delimiter, include the delimiter in the value. For example, enter **AWSLogs/eventlogs** to define a condition for all objects whose key begins with *AWSLogs/eventlogs*. Macie supports the default Amazon S3 delimiter, which is a slash (/), and custom delimiters.

Also note that an object matches a condition's criteria only if the object's key exactly matches the value that you enter, starting with the first character in the object's key. In addition, Macie applies a condition to the complete **Key** value for an object, including the object's file name.

For example, if an object's key is *AWSLogs/eventlogs/testlog.csv* and you enter any of the following values for a condition, the object matches the condition's criteria:

- **AWSLogs**
- **AWSLogs/event**
- **AWSLogs/eventlogs/**
- **AWSLogs/eventlogs/testlog**
- **AWSLogs/eventlogs/testlog.csv**

However, if you enter **eventlogs**, the object doesn't match the criteria—the condition's value doesn't include the first part of the key, *AWSLogs/*. Similarly, if you enter **awslogs**, the object doesn't match the criteria due to differences in capitalization.

Storage size

This correlates to the **Size** field in Amazon S3. In Amazon S3, this field indicates the total storage size of an S3 object. If an object is a compressed file, this value doesn't reflect the actual size of the file after it's decompressed.

You can use this type of condition to include or exclude objects that are smaller than a certain size, larger than a certain size, or fall within a certain size range. Macie applies this type of condition to all types of objects, including compressed or archive files and the files that they contain. For information about size-based restrictions for each supported format, see [Amazon Macie quotas \(p. 328\)](#).

Tags

The tags that are associated with an S3 object. Tags are labels that you can define and assign to certain types of AWS resources, including S3 objects. Each tag consists of a required tag key and an optional tag value. For information about tagging S3 objects, see [Categorizing your storage using tags](#) in the *Amazon Simple Storage Service User Guide*.

For a sensitive data discovery job, you can use this type of condition to include or exclude objects that have a specific tag. This can be a specific tag key or a specific tag key and tag value (as a pair). If you specify multiple tags as values for a condition, Macie uses OR logic to join the values. For

example, if you specify **Project1** and **Project2** as tag keys for a condition, any object that has the *Project1* or *Project2* tag key matches the condition's criteria.

Note that tag keys and values are case sensitive. In addition, Macie doesn't support use of partial values or wildcard characters in this type of condition.

Creating a sensitive data discovery job

With Amazon Macie, you create and run sensitive data discovery jobs to automate discovery, logging, and reporting of sensitive data in Amazon Simple Storage Service (Amazon S3) buckets. A *sensitive data discovery job* is a series of automated processing and analysis tasks that Macie performs to analyze objects in S3 buckets and determine whether the objects contain sensitive data.

When you create a job, you start by specifying which S3 buckets contain objects that you want to analyze—specific buckets that you select or buckets that match specific criteria. Then you specify how often to run the job—once, or periodically on a daily, weekly, or monthly basis. You can also choose options to refine the scope of the job's analysis. The options include custom criteria that derive from properties of S3 objects, such as tags.

After you define the schedule and scope of the job, you specify which managed data identifiers and custom data identifiers you want the job to use:

- A *managed data identifier* is a set of built-in criteria and techniques that are designed to detect a specific type of sensitive data—for example, credit card numbers, AWS secret access keys, or passport numbers for a particular country or region. These identifiers can detect a large and growing list of sensitive data types for many countries and regions, including multiple types of financial data, credentials data, and personally identifiable information (PII). For more information, see [Using managed data identifiers \(p. 45\)](#).
- A *custom data identifier* is a set of criteria that you define to detect sensitive data. With custom data identifiers, you can detect sensitive data that reflects your organization's particular scenarios, intellectual property, or proprietary data—for example, employee IDs, customer account numbers, or internal data classifications. You can supplement the managed data identifiers that Macie provides. For more information, see [Building custom data identifiers \(p. 64\)](#).

You then optionally select allow lists that you want the job to use. An *allow list* specifies text or a text pattern that you want Macie to ignore, typically sensitive data exceptions for your particular scenarios or environment—for example, public names or phone numbers for your organization, or sample data that your organization uses for testing. For more information, see [Defining sensitive data exceptions with allow lists \(p. 70\)](#).

When you finish choosing these options, you're ready to enter general settings for the job, such as the job's name and description. You can then review and save the job.

Tasks

- [Before you begin \(p. 98\)](#)
- [Step 1: Choose S3 buckets \(p. 98\)](#)
- [Step 2: Review your S3 bucket selections or criteria \(p. 100\)](#)
- [Step 3: Define the schedule and refine the scope \(p. 100\)](#)
- [Step 4: Select managed data identifiers \(p. 101\)](#)
- [Step 5: Select custom data identifiers \(p. 102\)](#)
- [Step 6: Select allow lists \(p. 103\)](#)
- [Step 7: Enter general settings \(p. 103\)](#)
- [Step 8: Review and create \(p. 103\)](#)

Before you begin

Before you create a job, it's a good idea to take the following steps:

- Verify that you configured Macie to store your sensitive data discovery results in an S3 bucket. To do this, choose **Discovery results** in the navigation pane on the Amazon Macie console, and then verify that you entered the settings. To learn about these settings, see [Storing and retaining sensitive data discovery results \(p. 130\)](#).
- Create any custom data identifiers that you want the job to use. To learn how, see [Building custom data identifiers \(p. 64\)](#).
- Create any allow lists that you want the job to use. To learn how, see [Creating and managing allow lists \(p. 78\)](#).
- If you want to analyze objects that are encrypted with a customer managed AWS KMS key, ensure that Macie has permission to use the key. For more information, see [Analyzing encrypted S3 objects \(p. 125\)](#).
- If you want to analyze objects in a bucket that has a restrictive bucket policy, ensure that Macie is allowed to access objects in the bucket. For more information, see [Allowing Macie to access S3 buckets and objects \(p. 40\)](#).

If you do these things before you create a job, you streamline creation of the job and help ensure that the job can analyze the data that you want.

Step 1: Choose S3 buckets

The first step in creating a job is to specify which S3 buckets you want the job to analyze. For this step, you have two options:

- **Select specific buckets** – With this option, you explicitly select each S3 bucket that you want the job to analyze. Then, when the job runs, it analyzes objects only in the buckets that you select.
- **Specify bucket criteria** – With this option, you define runtime criteria that determine which S3 buckets the job analyzes. The criteria consist of one or more conditions that derive from bucket properties. Then, when the job runs, it identifies buckets that match your criteria and analyzes objects in those buckets.

For detailed information about these options, see [Scope options for sensitive data discovery jobs \(p. 89\)](#).

The following sections provide step-by-step instructions for choosing and configuring each option. Choose the section for the option that you want.


Select specific buckets


If you choose to explicitly select each S3 bucket that you want the job to analyze, Macie provides you with a complete inventory of your buckets in the current AWS Region. You can then use this inventory to select one or more buckets for the job to analyze. To learn about this inventory, see [Selecting S3 buckets \(p. 90\)](#).

If you're the Macie administrator for an organization, the inventory includes buckets that are owned by member accounts in your organization. You can configure the job to analyze objects in as many as 1,000 of these buckets, spanning as many as 1,000 accounts.

To select specific buckets for the job

1. Open the Macie console at <https://console.aws.amazon.com/macie/>.

2. In the navigation pane, choose **Jobs**.
3. Choose **Create job**.
4. On the **Choose S3 buckets** page, choose **Select specific buckets**. Macie displays a table of all the buckets for your account in the current Region.
5. Under **Select S3 buckets**, optionally choose refresh () to retrieve the latest bucket metadata from Amazon S3.

If the information icon () appears next to any bucket names, we recommend that you do this. This icon indicates that a bucket was created during the past 24 hours, possibly after Macie last retrieved bucket and object metadata from Amazon S3 as part of the [daily refresh cycle](#) (p. 18).

6. In the table, select the check box for each bucket that you want the job to analyze.

Tip

- To find specific buckets more easily, enter filter criteria in the filter bar above the table. You can also sort the table by choosing a column heading.
 - To quickly determine whether you already configured a job to periodically analyze objects in a bucket, refer to the **Monitored** column. If **Yes** appears in the column, the bucket is explicitly included in a periodic job or the bucket matched the criteria for a periodic job within the past 24 hours. In addition, the status of at least one of those jobs is not *Cancelled*. Macie updates this data on a daily basis.
 - To quickly determine when you most recently ran a periodic or one-time job to analyze objects in a bucket, refer to the **Latest job run** column. For additional information about that job, refer to the bucket's details.
 - To display a bucket's details, choose the bucket's name. In addition to job-related information, the details panel provides statistics and other information about the bucket, such as the bucket's public access settings. To learn more about this data, see [Reviewing your S3 bucket inventory](#) (p. 26).
7. When you finish selecting buckets, choose **Next**.

In the next step, you'll review and verify your selections.

Specify bucket criteria

If you choose to specify runtime criteria that determine which S3 buckets the job analyzes, Macie provides options to help you choose fields, operators, and values for individual conditions in the criteria. To learn more about these options, see [Specifying S3 bucket criteria](#) (p. 92).

To specify bucket criteria for the job

1. Open the Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, choose **Jobs**.
3. Choose **Create job**.
4. On the **Choose S3 buckets** page, choose **Specify bucket criteria**.
5. Under **Specify bucket criteria**, do the following to add a condition to the criteria:
 - a. Place your cursor in the filter bar, and then choose the bucket property to use for the condition.
 - b. In the first box, choose an operator for the condition, **Equals** or **Not equals**.
 - c. In the next box, enter one or more values for the property.

Depending on the type and nature of the bucket property, Macie displays different options for entering values. For example, if you choose the **Effective permission** property, Macie displays a list of values to choose from. If you choose the **Account ID** property, Macie displays a text box in

which you can enter one or more AWS account IDs. To enter multiple values in a text box, enter each value and separate each entry with a comma.

- d. Choose **Apply**. Macie adds the condition to a filter box below the filter bar.

By default, Macie adds the condition with an include statement. This means that the job is configured to analyze (*include*) objects in buckets that match the condition. To skip (*exclude*) buckets that match the condition, choose **Include** in the filter box, and then choose **Exclude**.

- e. Repeat the preceding steps for each additional condition that you want to add to the criteria.
6. To test your criteria, expand the **Preview the criteria results** section. This section displays a table of all the buckets that currently match the criteria.
7. To refine your criteria, do any of the following:
 - To remove a condition, choose **X** in the filter box for the condition.
 - To change a condition, remove the condition by choosing **X** in the filter box for the condition. Then add a condition that has the correct settings.
 - To remove all conditions, choose **Clear filters**.

Macie updates the table of criteria results to reflect your changes.

8. When you finish specifying bucket criteria, choose **Next**.

In the next step, you'll review and verify your criteria.

Step 2: Review your S3 bucket selections or criteria

For this step, verify that you chose the correct settings in the preceding step.

Review your bucket selections

If you selected specific S3 buckets for the job, review the table of buckets and change your bucket selections as necessary. The table provides insight into the projected scope and cost of the job's analysis. The data is based on the size and types of objects that are currently stored in a bucket.

The **Estimated cost** field indicates the total estimated cost (in US Dollars) of analyzing objects in a bucket. Each estimate reflects the projected amount of uncompressed data that the job will analyze in a bucket. If any objects are compressed or archive files, the estimate assumes that the files use a 3:1 compression ratio and the job can analyze all extracted files. For more information, see [Forecasting and monitoring costs for sensitive data discovery jobs \(p. 122\)](#).

Review your bucket criteria

If you specified bucket criteria for the job, review each condition in the criteria. To change the criteria, choose **Previous**, and then use the filter settings in the preceding step to enter the correct criteria. When you finish, choose **Next**.

When you finish reviewing and verifying the settings, choose **Next**.

Step 3: Define the schedule and refine the scope

For this step, specify how often you want the job to run—once, or periodically on a daily, weekly, or monthly basis. Also choose various options to refine the scope of the job's analysis. To learn about these options, see [Scope options for sensitive data discovery jobs \(p. 89\)](#).

To define the schedule and refine the scope of the job

1. On the **Refine the scope** page, choose how often you want the job to run:

- To run the job only once, immediately after you finish creating it, choose **One-time job**.
- To run the job periodically on a recurring basis, choose **Scheduled job**. For **Update frequency**, choose whether to run the job daily, weekly, or monthly. Then use the **Include existing objects** option to define the scope of the job's first run:
 - Select this check box to analyze all existing objects immediately after you finish creating the job. Each subsequent run analyzes only those objects that are created or changed after the preceding run.
 - Clear this check box to skip analysis of all existing objects. The job's first run analyzes only those objects that are created or changed after you finish creating the job and before the first run starts. Each subsequent run analyzes only those objects that are created or changed after the preceding run.

Clearing this check box is helpful for cases where you've already analyzed the data and want to continue to analyze it periodically. For example, if you previously used another service or application to classify data and you recently started using Macie, you might use this option to ensure continued discovery and classification of your data without incurring unnecessary costs or duplicating classification data.

2. (Optional) To specify the percentage of objects that you want the job to analyze, enter the percentage in the **Sampling depth** box. If this value is less than 100%, Macie selects the objects to analyze at random, up to the specified percentage, and analyzes all the data in those objects. The default value is 100%.
3. (Optional) To add specific criteria that determine which S3 objects are included or excluded from the job's analysis, expand the **Additional settings** section, and then enter the criteria. These criteria consist of individual conditions that derive from properties of objects.
 - To analyze (*include*) objects that meet a specific condition, enter the condition type and value, and then choose **Include**.
 - To skip (*exclude*) objects that meet a specific condition, enter the condition type and value, and then choose **Exclude**.

Repeat this step for each include or exclude condition that you want.

In Macie, exclude conditions take precedence over include conditions. For example, if you include objects that have the .pdf file name extension and exclude objects that are larger than 5 MB, the job analyzes any object that has the .pdf file name extension, unless the object is larger than 5 MB.

4. When you finish, choose **Next**.

Step 4: Select managed data identifiers

For this step, specify which managed data identifiers you want the job to use when it analyzes S3 objects. You can configure the job to use all, some, or none of the managed data identifiers that Macie provides. To review a detailed list of the managed data identifiers that are currently available, see [Using managed data identifiers \(p. 45\)](#). We update that list each time we release a new managed data identifier.

If you choose to use only some managed data identifiers, Macie displays a table of the managed data identifiers that are currently available. You can use the table to select each managed data identifier that you want the job to use (*include*) or not use (*exclude*), depending on the selection type that you choose for the job. In the table, each managed data identifier's ID describes the type of sensitive data that the managed data identifier detects, for example: **USA_PASSPORT_NUMBER** for US passport numbers, **CREDIT_CARD_SECURITY_CODE** for credit card verification codes, and **PGP_PRIVATE_KEY** for PGP private keys. To find specific identifiers more quickly, you can sort and filter the table by sensitive data category and type.

To select managed data identifiers for the job

1. On the **Select managed data identifiers** page, under **Selection type**, do one of the following to specify which managed data identifiers you want the job to use:

- To use all managed data identifiers, choose **All**.

If you choose this option and you configured the job to run more than once, each run will automatically use new managed data identifiers that we release, in addition to all the managed data identifiers that are currently available.

- To exclude specific managed data identifiers, choose **Exclude**. Then, in the table that appears, select the check box for each managed data identifier that you don't want the job to use.

For example, if you don't want the job to detect and report occurrences of mailing addresses, select the **ADDRESS** check box. If you do this, the job will use all managed data identifiers except the one that detects mailing addresses.

If you choose the **Exclude** option and you configured the job to run more than once, each run will automatically use new managed data identifiers that we release, in addition to all the managed data identifiers that are currently available and you didn't explicitly exclude from the job.

- To include only specific managed data identifiers, choose **Include**. Then, in the table that appears, select the check box for each managed data identifier that you want the job to use.

For example, if you want the job to only detect and report occurrences of US passport numbers, select the **USA_PASSPORT_NUMBER** check box. If you do this, the job won't use any managed data identifiers except the one that detects US passport numbers.

- To exclude all managed data identifiers, choose **None**.

If you choose this option, the job won't use any managed data identifiers. In the [next step \(p. 102\)](#), configure the job to instead use one or more custom data identifiers that you specify.

2. When you finish, choose **Next**.


Step 5: Select custom data identifiers

For this step, optionally select one or more [custom data identifiers \(p. 64\)](#) that you want the job to use when it analyzes S3 objects. The job will use the selected identifiers in addition to any managed data identifiers that you configured the job to use.

To select custom data identifiers for the job

1. On the **Select custom data identifiers** page, select the check box for each custom data identifier that you want the job to use. You can select as many as 30 custom data identifiers.

Tip

To test or review the settings for a custom data identifier before you select it, choose the link icon () next to the identifier's name. Macie opens a page that displays the identifier's settings. You can also use this page to test the identifier with sample data. To do this, enter up to 1,000 characters of text in the **Sample data** box, and then choose **Test**. Macie evaluates the sample data by using the identifier, and then reports the number of matches.

2. When you finish selecting custom data identifiers, choose **Next**.


Step 6: Select allow lists

For this step, optionally select one or more [allow lists \(p. 70\)](#) that you want the job to use when it analyzes S3 objects.

To select allow lists for the job

1. On the **Select allow lists** page, select the check box for each allow list that you want the job to use. You can select as many as 10 lists.

Tip

To review the settings for an allow list before you select it, choose the link icon () next to the list's name. Macie opens a page that displays the list's settings. If the list specifies a regular expression (*regex*), you can also use this page to test the regex with sample data. To do this, enter up to 1,000 characters of text in the **Sample data** box, and then choose **Test**. Macie evaluates the sample data by using the regex, and then reports the number of matches.

2. When you finish selecting allow lists, choose **Next**.

Step 7: Enter general settings

For this step, specify a name and, optionally, a description of the job.

You can also assign tags to the job. A *tag* is a label that you define and assign to certain types of AWS resources. Each tag consists of a required tag key and an optional tag value. Tags can help you identify, categorize, and manage resources in different ways, such as by purpose, owner, environment, or other criteria. To learn more, see [Tagging Amazon Macie resources \(p. 315\)](#).

To enter general settings for the job

1. On the **Enter general settings** page, enter a name for the job in the **Job name** box. The name can contain as many as 500 characters.
2. (Optional) For **Job description**, enter a brief description of the job. The description can contain as many as 200 characters.
3. (Optional) For **Tags**, choose **Add tag**, and then enter as many as 50 tags to assign to the job.
4. When you finish, choose **Next**.

Step 8: Review and create

For this final step, review the configuration settings for the job and verify that they're correct. This is an important step. After you create a job, you can't change any of these settings. This helps ensure that you have an immutable history of sensitive data findings and discovery results for data privacy and protection audits or investigations that you perform.

Depending on the job's settings, you can also review the total estimated cost (in US Dollars) of running the job once. If you selected specific S3 buckets for the job, the estimate is based on the size and types of objects in the buckets that you selected, and how much of that data the job can analyze. If you specified bucket criteria for the job, the estimate is based on the size and types of objects in as many as 500 buckets that currently match the criteria, and how much of that data the job can analyze. To learn about this estimate, see [Forecasting and monitoring costs for sensitive data discovery jobs \(p. 122\)](#).

To review and create the job


1. On the **Review and create** page, review each setting and verify that it's correct.

To change a setting, choose **Edit** in the section that contains the setting, and then enter the correct setting. You can also use the navigation tabs to go to the page that contains a setting.

2. When you finish verifying the settings, choose **Submit** to create and save the job. Macie checks the settings and notifies you of any issues to address.

Note

If you haven't configured a repository for your sensitive data discovery results, Macie displays a warning and doesn't save the job. To address this issue, choose **Configure** in the **Repository for sensitive data discovery results** section. Then enter the configuration settings for the repository. To learn how, see [Storing and retaining sensitive data discovery results \(p. 130\)](#). After you enter the settings, return to the **Review and create** page and

then choose refresh () in the **Repository for sensitive data discovery results** section of the page.

Although we don't recommend it, you can temporarily override the repository requirement and save the job. If you do this, you risk losing discovery results from the job—Macie will retain the results for only 90 days. To temporarily override the requirement, select the check box for the override option.

3. If Macie notifies you of issues to address, address the issues, and then choose **Submit** again to create and save the job.

If you configured the job to run once, on a daily basis, or on the current day of the week or month, Macie starts running the job immediately after you save it. Otherwise, Macie prepares to run the job on the specified day of the week or month. To monitor the job, you can [check the status of the job \(p. 119\)](#).

Monitoring sensitive data discovery jobs with Amazon CloudWatch Logs

In addition to [monitoring the overall status \(p. 119\)](#) of a sensitive data discovery job, you can monitor and analyze specific events that occur as a job progresses. You can do this by using near-real-time logging data that Amazon Macie automatically publishes to Amazon CloudWatch Logs. The data in these logs provides a record of changes to a job's progress or status, such as the exact date and time when a job started to run, was paused, or finished running.

The log data also provides details about any account- or bucket-level errors that occur while a job runs. For example, if the permissions settings for an S3 bucket prevent a job from analyzing objects in the bucket, Macie logs an event. The event indicates when the error occurred, and it identifies both the affected bucket and the account that owns the bucket. The data for these types of events can help you identify, investigate, and address errors that prevent Macie from analyzing the data that you want.

With Amazon CloudWatch Logs, you can monitor, store, and access log files from multiple systems, applications, and AWS services, including Macie. You can also query and analyze log data, and configure CloudWatch Logs to notify you when certain events occur or thresholds are met. CloudWatch Logs also provides features for archiving log data and exporting the data to Amazon S3. To learn more about CloudWatch Logs, see the [Amazon CloudWatch Logs User Guide](#).

Topics

- [How logging works for sensitive data discovery jobs \(p. 105\)](#)
- [Reviewing logs for sensitive data discovery jobs \(p. 105\)](#)
- [Log event schema for sensitive data discovery jobs \(p. 106\)](#)
- [Types of log events for sensitive data discovery jobs \(p. 107\)](#)

How logging works for sensitive data discovery jobs

When you start running sensitive data discovery jobs, Macie automatically creates and configures the appropriate resources in Amazon CloudWatch Logs to log events for all of your jobs in the current AWS Region. Macie then publishes event data to those resources automatically when your jobs run. The permissions policy for the Macie [service-linked role](#) (p. 302) for your account allows Macie to perform these tasks on your behalf. You don't need to take any steps to create or configure resources in CloudWatch Logs, or to log event data for your jobs.

In CloudWatch Logs, logs are organized into *log groups*. Each log group contains *log streams*. Each log stream contains *log events*. The general purpose of each of these resources is as follows:

- A *log group* is a collection of log streams that share the same retention, monitoring, and access control settings—for example, the collection of logs for all of your sensitive data discovery jobs.
- A *log stream* is a sequence of log events that share the same source—for example, an individual sensitive data discovery job.
- A *log event* is a record of an activity that was recorded by an application or resource—for example, an individual event that Macie recorded and published for a particular sensitive data discovery job.

Macie publishes events for all of your sensitive data discovery jobs to one log group, and each job has a unique log stream in that log group. The log group has the following prefix and name:

```
/aws/macie/classificationjobs
```

If this log group already exists, Macie uses it to store log events for your jobs. This can be helpful if your organization uses automated configuration, such as [AWS CloudFormation](#), to create log groups with predefined log retention periods, encryption settings, tags, metric filters, and so on for job events.

If this log group doesn't exist, Macie creates it with the default settings that CloudWatch Logs uses for new log groups. The settings include a log retention period of **Never Expire**, which means that CloudWatch Logs stores the logs indefinitely. To change the retention period for the log group, you can use the Amazon CloudWatch console or the Amazon CloudWatch Logs API. To learn how, see [Working with log groups and log streams](#) in the *Amazon CloudWatch Logs User Guide*.

Within this log group, Macie creates a unique log stream for each job that you run, the first time that the job runs. The name of the log stream is the unique identifier for the job, such as `85a55dc0fa6ed0be5939d0408example`, in the following format.

```
/aws/macie/classificationjobs/85a55dc0fa6ed0be5939d0408example
```

Each log stream contains all the log events that Macie recorded and published for the corresponding job. For periodic jobs, this includes events for all of the job's runs. If you delete the log stream for a periodic job, Macie creates the stream again the next time that the job runs. If you delete the log stream for a one-time job, you can't restore it.

Note that logging is enabled by default for all of your jobs. You can't disable it or otherwise prevent Macie from publishing job events to CloudWatch Logs. If you don't want to store the logs, you can reduce the retention period for the log group to as little as one day. At the end of the retention period, CloudWatch Logs automatically deletes expired event data from the log group.

Reviewing logs for sensitive data discovery jobs

You can review the logs for your sensitive data discovery jobs by using the Amazon CloudWatch console or the Amazon CloudWatch Logs API. Both the console and the API provide features that are designed to help you review and analyze log data. You can use these features to work with log streams and events for your jobs as you would work with any other type of log data in CloudWatch Logs.

For example, you can search and filter aggregate data to identify specific types of events that occurred for all of your jobs during a specific time range. Or you can perform a targeted review of all the events that occurred for a particular job. CloudWatch Logs also provides options for monitoring log data, defining metric filters, and creating custom alarms.

Tip

To navigate to the log events for a particular job by using the Amazon Macie console, do the following: On the **Jobs** page, choose the name of the job. At the top of the details panel, choose **Show results**, and then choose **Show CloudWatch logs**. Macie opens the Amazon CloudWatch console and displays a table of log events for the job.

To review the logs for your jobs (Amazon CloudWatch console)

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. By using the AWS Region selector in the upper-right corner of the page, select the Region in which you ran jobs that you want to review logs for.
3. In the navigation pane, choose **Logs**, and then choose **Log groups**.
4. On the **Log groups** page, choose the **/aws/macie/classificationjobs** log group. CloudWatch Logs displays a table of log streams for the jobs that you've run. There is one unique stream for each job. The name of each stream correlates to the unique identifier for a job.
5. Under **Log streams**, do one of the following:
 - To review the log events for a particular job, choose the log stream for the job. To find the stream more easily, enter the job's unique identifier in the filter bar above the table. After you choose the log stream, CloudWatch Logs displays a table of log events for the job.
 - To review log events for all of your jobs, choose **Search all**. CloudWatch Logs displays a table of log events for all of your jobs.
6. (Optional) In the filter bar above the table, enter terms, phrases, or values that specify characteristics of specific events to review. For more information, see [Search log data using filter patterns](#) in the *Amazon CloudWatch Logs User Guide*.
7. To review the details of a specific log event, choose the right arrow (▶) in the row for the event. CloudWatch Logs displays the event's details in JSON format.

As you familiarize yourself with the data in the log events, you can also perform tasks such as [creating metrics filters](#) that turn log data into numerical CloudWatch metrics, and [creating custom alarms](#) that make it easier for you to identify and respond to specific log events. For more information, see the [Amazon CloudWatch Logs User Guide](#).

Log event schema for sensitive data discovery jobs

Each log event for a sensitive data discovery job is a JSON object that conforms to the Amazon CloudWatch Logs event schema and contains a standard set of fields. Some types of events have additional fields that provide information that's particularly useful for that type of event. For example, events for account-level errors include the account ID for the affected AWS account. Events for bucket-level errors include the name of the affected S3 bucket. For a detailed list of job events that Macie publishes to CloudWatch Logs, see [Types of log events for jobs \(p. 107\)](#).

The following example shows the log event schema for sensitive data discovery jobs. In this example, the event reports that Macie wasn't able to analyze any objects in an S3 bucket because Amazon S3 denied access to the bucket.

```
{
  "adminAccountId": "123456789012",
  "jobId": "85a55dc0fa6ed0be5939d0408example",
  "eventType": "BUCKET_ACCESS_DENIED",
```



```
"occurredAt": "2021-04-14T17:11:30.574809Z",
"description": "Macie doesn't have permission to access the affected S3 bucket.",
"jobName": "My_Macie_Job",
"operation": "ListObjectsV2",
"runDate": "2021-04-14T17:08:30.345809Z",
"affectedAccount": "111122223333",
"affectedResource": {
  "type": "S3_BUCKET_NAME",
  "value": "DOC-EXAMPLE-BUCKET"
}
}
```

In the preceding example, Macie attempted to list the objects in the bucket by using the [ListObjectsV2](#) operation of the Amazon S3 API. When Macie sent the request to Amazon S3, Amazon S3 denied access to the bucket.

The following fields are common to all log events for sensitive data discovery jobs:

- **adminAccountId** – The unique identifier for the AWS account that created the job.
- **jobId** – The unique identifier for the job.
- **eventType** – The type of event that occurred. For complete lists of possible values and a description of each one, see [Types of log events for jobs \(p. 107\)](#).
- **occurredAt** – The date and time, in Coordinated Universal Time (UTC) and extended ISO 8601 format, when the event occurred.
- **description** – A brief description of the event.
- **jobName** – The custom name of the job.

Depending on the type and nature of an event, a log event can also contain the following fields:

- **affectedAccount** – The unique identifier for the AWS account that owns the affected resource.
- **affectedResource** – An object that provides details about the affected resource. In the object, the `type` field specifies a field that stores metadata about a resource. The `value` field specifies the value for the field (`type`).
- **operation** – The operation that Macie attempted to perform and caused the error.
- **runDate** – The date and time, in Coordinated Universal Time (UTC) and extended ISO 8601 format, when the applicable job or job run started.

Types of log events for sensitive data discovery jobs

Macie publishes log events for three categories of events:

- Job status events, which record changes to the status or progress of a job or a job run.
- Account-level error events, which record errors that prevented Macie from analyzing Amazon S3 data for a specific AWS account.
- Bucket-level error events, which record errors that prevented Macie from analyzing data in a specific S3 bucket.

The topics in this section list and describe the types of events that Macie publishes for each category.

Topics

- [Job status events \(p. 108\)](#)
- [Account-level error events \(p. 110\)](#)
- [Bucket-level error events \(p. 112\)](#)

Job status events

A job status event records a change to the status or progress of a job or a job run. For periodic jobs, Macie logs and publishes these events for both the overall job and individual job runs. For information about determining the overall status of a job, see [Checking the status of sensitive data discovery jobs](#) (p. 119).

The following example uses sample data to show the structure and nature of the fields in a job status event. In this example, a `SCHEDULED_RUN_COMPLETED` event indicates that a scheduled run of a periodic job finished running. The run started on April 14, 2021, at 17:09:30 UTC, as indicated by the `runDate` field. The run finished on April 14, 2021, at 17:16:30 UTC, as indicated by the `occurredAt` field.

```
{
  "adminAccountId": "123456789012",
  "jobId": "ffad0e71455f38a4c7c220f3cexample",
  "eventType": "SCHEDULED_RUN_COMPLETED",
  "occurredAt": "2021-04-14T17:16:30.574809Z",
  "description": "The scheduled job run finished running.",
  "jobName": "My_Daily_Macie_Job",
  "runDate": "2021-04-14T17:09:30.574809Z"
}
```

The following table lists and describes the types of job status events that Macie logs and publishes to CloudWatch Logs. The **Event type** column indicates the name of each event as it appears in the `eventType` field of an event. The **Description** column provides a brief description of the event as it appears in the `description` field of an event. The **Additional information** provides information about the type of job that the event applies to. The table is sorted first by the general chronological order in which events might occur, and then in ascending alphabetical order by event type.

Event type	Description	Additional information
JOB_CREATED	The job was created.	Applies to one-time and periodic jobs.
ONE_TIME_JOB_STARTED	The job started running.	Applies only to one-time jobs.
SCHEDULED_RUN_STARTED	The scheduled job run started running.	Applies only to periodic jobs. To log the start of a one-time job, Macie publishes a <code>ONE_TIME_JOB_STARTED</code> event, not this type of event.
BUCKET_MATCHED_THE_CRITERIA	The affected bucket matched the bucket criteria specified for the job.	Applies to one-time and periodic jobs that use runtime bucket criteria to determine which S3 buckets to analyze. The <code>affectedResource</code> object specifies the name of the bucket that matched the criteria and was included in the job's analysis.
NO_BUCKETS_MATCHED_THE_CRITERIA	The job started running but no buckets currently match the bucket criteria specified for the job. The job didn't analyze any data.	Applies to one-time and periodic jobs that use runtime bucket criteria to determine which S3 buckets to analyze.

Event type	Description	Additional information
SCHEDULED_RUN_COMPLETED	The scheduled job run finished running.	Applies only to periodic jobs. To log completion of a one-time job, Macie publishes a JOB_COMPLETED event, not this type of event.
JOB_PAUSED_BY_USER	The job was paused by a user.	Applies to one-time and periodic jobs that you stopped temporarily (paused).
JOB_RESUMED_BY_USER	The job was resumed by a user.	Applies to one-time and periodic jobs that you stopped temporarily (paused) and subsequently resumed.
JOB_PAUSED_BY_MACIE_SERVICE_QUOTA_MET	The job was paused by Macie. Completion of the job would exceed a monthly quota for the affected account.	Applies to one-time and periodic jobs that Macie stopped temporarily (paused). Macie automatically pauses a job when additional processing by the job or a job run would exceed the monthly sensitive data discovery quota (p. 328) for one or more accounts that the job analyzes data for. To avoid this issue, consider increasing the quota for the affected accounts.
JOB_RESUMED_BY_MACIE_SERVICE_QUOTA_LIFTED	The job was resumed by Macie. The monthly service quota was lifted for the affected account.	Applies to one-time and periodic jobs that Macie stopped temporarily (paused) and subsequently resumed. If Macie automatically paused a one-time job, Macie automatically resumes the job when the subsequent month starts or the monthly sensitive data discovery quota is increased for all the affected accounts, whichever occurs first. If Macie automatically paused a periodic job, Macie automatically resumes the job when the next run is scheduled to start or the subsequent month starts, whichever occurs first.

Event type	Description	Additional information
JOB_CANCELLED	The job was cancelled.	Applies to one-time and periodic jobs that you stopped permanently (cancelled) or, for one-time jobs, paused and didn't resume within 30 days. If you suspend or disable Macie, this type of event also applies to jobs that were active or paused when you suspended or disabled Macie. Macie automatically cancels your jobs in an AWS Region if you suspend or disable Macie in the Region.
JOB_COMPLETED	The job finished running.	Applies only to one-time jobs. To log completion of a job run for a periodic job, Macie publishes a SCHEDULED_RUN_COMPLETED event, not this type of event.

Account-level error events

An account-level error event records an error that prevented Macie from analyzing objects in S3 buckets that are owned by a specific AWS account. The `affectedAccount` field in each event specifies the account ID for that account.

The following example uses sample data to show the structure and nature of the fields in an account-level error event. In this example, an `ACCOUNT_ACCESS_DENIED` event indicates that Macie wasn't able to analyze objects in any S3 buckets that are owned by account 444455556666.

```
{
  "adminAccountId": "123456789012",
  "jobId": "85a55dc0fa6ed0be5939d0408example",
  "eventType": "ACCOUNT_ACCESS_DENIED",
  "occurredAt": "2021-04-14T17:08:30.585709Z",
  "description": "Macie doesn't have permission to access S3 bucket data for the affected account.",
  "jobName": "My_Macie_Job",
  "operation": "ListBuckets",
  "runDate": "2021-04-14T17:05:27.574809Z",
  "affectedAccount": "444455556666"
}
```

The following table lists and describes the types of account-level error events that Macie logs and publishes to CloudWatch Logs. The **Event type** column indicates the name of each event as it appears in the `eventType` field of an event. The **Description** column provides a brief description of the event as it appears in the `description` field of an event. The **Additional information** column provides any applicable tips for investigating or addressing the error that occurred. The table is sorted in ascending alphabetical order by event type.

Event type	Description	Additional information
ACCOUNT_ACCESS_DENIED	Macie doesn't have permission to access S3 bucket data for the affected account.	<p>This typically occurs because the buckets that are owned by the account have restrictive bucket policies. For information about how to address this issue, see Allowing Macie to access S3 buckets and objects (p. 40).</p> <p>The value for the <code>operation</code> field in the event can help you determine which permissions settings prevented Macie from accessing S3 data for the account. This field indicates the Amazon S3 operation that Macie attempted to perform when the error occurred.</p>
ACCOUNT_DISABLED	The job skipped resources that are owned by the affected account. Macie was disabled for the account.	To address this issue, re-enable Macie for the account in the same AWS Region.
ACCOUNT_DISASSOCIATED	The job skipped resources that are owned by the affected account. The account isn't associated with your Macie administrator account as a member account anymore.	<p>This occurs if you, as a Macie administrator for an organization, configure a job to analyze data for an associated member account and the member account is subsequently removed from your organization.</p> <p>To address this issue, re-associate the affected account with your Macie administrator account as a member account. For more information, see Managing multiple accounts (p. 238).</p>
ACCOUNT_ISOLATED	The job skipped resources that are owned by the affected account. The AWS account was isolated.	–
ACCOUNT_REGION_DISABLED	The job skipped resources that are owned by the affected account. The AWS account isn't active in the current AWS Region.	–
ACCOUNT_SUSPENDED	The job was cancelled or skipped resources that are owned by the affected account. Macie was suspended for the account.	If the specified account is your own account, Macie automatically cancelled the job when you suspended Macie in the same Region. To address

Event type	Description	Additional information
		the issue, re-enable Macie in the Region. If the specified account is a member account, re-enable Macie for that account in the same Region.
ACCOUNT_TERMINATED	The job skipped resources that are owned by the affected account. The AWS account was terminated.	–

Bucket-level error events

A bucket-level error event records an error that prevented Macie from analyzing objects in a specific S3 bucket. The `affectedAccount` field in each event specifies the account ID for the AWS account that owns the bucket. The `affectedResource` object in each event specifies the name of the bucket.

The following example uses sample data to show the structure and nature of the fields in a bucket-level error event. In this example, a `BUCKET_ACCESS_DENIED` event indicates that Macie wasn't able to analyze any objects in the S3 bucket named `DOC-EXAMPLE-BUCKET`. When Macie attempted to list the objects in the bucket by using the `ListObjectsV2` operation of the Amazon S3 API, Amazon S3 denied access to the bucket.

```
{
  "adminAccountId": "123456789012",
  "jobId": "85a55dc0fa6ed0be5939d0408example",
  "eventType": "BUCKET_ACCESS_DENIED",
  "occurredAt": "2021-04-14T17:11:30.574809Z",
  "description": "Macie doesn't have permission to access the affected S3 bucket.",
  "jobName": "My_Macie_Job",
  "operation": "ListObjectsV2",
  "runDate": "2021-04-14T17:09:30.685209Z",
  "affectedAccount": "111122223333",
  "affectedResource": {
    "type": "S3_BUCKET_NAME",
    "value": "DOC-EXAMPLE-BUCKET"
  }
}
```

The following table lists and describes the types of bucket-level error events that Macie logs and publishes to CloudWatch Logs. The **Event type** column indicates the name of each event as it appears in the `eventType` field of an event. The **Description** column provides a brief description of the event as it appears in the `description` field of an event. The **Additional information** column provides any applicable tips for investigating or addressing the error that occurred. The table is sorted in ascending alphabetical order by event type.

Event type	Description	Additional information
BUCKET_ACCESS_DENIED	Macie doesn't have permission to access the affected S3 bucket.	The value for the <code>operation</code> field in the event can help you determine which permissions settings prevented Macie from accessing the bucket. This

Event type	Description	Additional information
		field indicates the Amazon S3 operation that Macie attempted to perform when the error occurred.
BUCKET_DETAILS_UNAVAILABLE	A temporary issue prevented Macie from retrieving details about the bucket and the bucket's objects.	This occurs if a transient issue prevented Macie from retrieving the bucket and object metadata that it needs to analyze the bucket's objects. For example, an Amazon S3 exception occurred when Macie tried to verify that it's allowed to access the bucket. To address the issue for a one-time job, consider creating and running a new, one-time job to analyze objects in the bucket. For a scheduled job, Macie will try to retrieve the metadata again during the next job run.
BUCKET_DOES_NOT_EXIST	The affected S3 bucket doesn't exist anymore.	This typically occurs because a bucket was deleted.
BUCKET_IN_DIFFERENT_REGION	The affected S3 bucket was moved to a different AWS Region.	–
BUCKET_OWNER_CHANGED	The owner of the affected S3 bucket changed. Macie doesn't have permission to access the bucket anymore.	This typically occurs if ownership of a bucket was transferred to an AWS account that isn't part of your organization. The <code>affectedAccount</code> field in the event indicates the account ID for the account that previously owned the bucket.

Reviewing statistics and results for sensitive data discovery jobs

When you run a sensitive data discovery job, Amazon Macie automatically calculates and reports certain statistical data for the job. For example, Macie reports the number of times that the job has run and the approximate number of Amazon Simple Storage Service (Amazon S3) objects that the job has yet to process during its current run. Macie also produces several types of results for the job: log events, sensitive data findings, and sensitive data discovery results.

Topics

- [Types of results for sensitive data discovery jobs \(p. 114\)](#)
- [Reviewing statistics and results for a sensitive data discovery job \(p. 115\)](#)

Types of results for sensitive data discovery jobs

As a sensitive data discovery job progresses, Amazon Macie produces the following types of results for the job.

Log event

This is a record of an event that occurred while the job was running. Macie automatically logs and publishes data for certain events to Amazon CloudWatch Logs. The data in these logs provides a record of changes to the job's progress or status, such as the exact date and time when the job started or stopped running. The data also provides details about any account- or bucket-level errors that occurred while the job ran.

Log events can help you monitor a job and address any issues that prevented the job from analyzing the data that you want. If a job uses runtime criteria to determine which S3 buckets to analyze, log events can also help you determine whether and which S3 buckets matched the criteria when the job ran.

You can access log events by using the Amazon CloudWatch console or the Amazon CloudWatch Logs API. To help you navigate to the log events for a job, the Amazon Macie console provides a link to them. For more information, see [Monitoring jobs \(p. 104\)](#).

Sensitive data finding

This is a report of sensitive data that Macie found in an S3 object. Each finding provides a severity rating and details such as:

- The date and time when Macie found the sensitive data.
- The category and types of sensitive data that Macie found.
- The number of occurrences of each type of sensitive data that Macie found.
- The unique identifier for the job that produced the finding.
- The name, public access settings, encryption type, and other information about the affected S3 bucket and object.

Depending on the object's file type or storage format, the details can also include the location of as many as 15 occurrences of the sensitive data that Macie found.

A sensitive data finding doesn't include the sensitive data that Macie found. Instead, it provides information that you can use for further investigation and remediation as necessary.

Macie stores sensitive data findings for 90 days. You can access them by using the Amazon Macie console or the Amazon Macie API. You can also monitor and process them by using other applications, services, and systems. For more information, see [Analyzing findings \(p. 140\)](#).

Sensitive data discovery result

This is a record that logs details about the analysis of an S3 object. Macie creates a sensitive data discovery result for each object that you configure a job to analyze. This includes objects that don't contain sensitive data, and therefore don't produce sensitive data findings, and objects that Macie can't analyze due to issues such as permissions settings or use of an unsupported format.

If an object does contain sensitive data, the sensitive data discovery result includes data from the corresponding sensitive data finding. It provides additional information too, such as the location of as many as 1,000 occurrences of each type of sensitive data that Macie found in the object. For example:

- The column and row number for a cell or field in a Microsoft Excel workbook, CSV file, or TSV file
- The path to a field or array in a JSON or JSON Lines file

- The line number for a line in a non-binary text file other than a CSV, JSON, JSON Lines, or TSV file—for example, an HTML, TXT, or XML file
- The page number for a page in an Adobe Portable Document Format (PDF) file
- The record index and the path to a field in a record in an Apache Avro object container or Apache Parquet file

If an object is an archive file, such as a .tar or .zip file, a sensitive data discovery result also provides detailed location data for occurrences of sensitive data in individual files that Macie extracts from the archive. Macie doesn't include this information in sensitive data findings for archive files.

A sensitive data discovery result doesn't include the sensitive data that Macie found. Instead, it provides you with an analysis record that can be helpful for data privacy and protection audits or investigations.

Macie stores sensitive data discovery results for 90 days. You can't access them directly on the Amazon Macie console or through the Amazon Macie API. Instead, you configure Macie to store the results in an S3 bucket. You can then optionally access and query the results in that bucket. This configuration also ensures long-term storage and retention of the results. To learn how to configure these settings, see [Storing and retaining sensitive data discovery results \(p. 130\)](#).

After you configure Macie to store your discovery results in an S3 bucket, Macie writes the results to JSON Lines (.jsonl) files and adds those files to the bucket as GNU Zip (.gz) files. To help you navigate to the results, the Amazon Macie console provides links to them.

Sensitive data findings and sensitive data discovery results both adhere to standardized schemas. This can help you optionally query, monitor, and process them by using other applications, services, and systems.

Tip

For samples of Amazon Athena queries that you can use to analyze sensitive data discovery results, visit the [Amazon Macie Results Analytics repository](#) on GitHub. This repository also provides step-by-step instructions for configuring Athena to retrieve and decrypt your sensitive data discovery results, and scripts for creating tables for the results.

Reviewing statistics and results for a sensitive data discovery job

To review processing statistics and results for individual sensitive data discovery jobs, you can use the Amazon Macie console or the Amazon Macie API. Follow these steps to review a job's statistics and results by using the console.

To access a job's processing statistics programmatically, use the `DescribeClassificationJob` operation of the Amazon Macie API. For programmatic access to the findings that a job produced, use the `ListFindings` operation of the Amazon Macie API and specify the job's unique identifier in a filter condition for the `classificationDetails.jobId` field. To learn how, see [Creating and applying filters to findings \(p. 154\)](#). You can then use the `GetFindings` operation to retrieve the details of the findings.

To review statistics and results for a job

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, choose **Jobs**.
3. On the **Jobs** page, choose the name of the job whose statistics and results you want to review. The details panel displays statistics, settings, and other information about the job.
4. In the details panel, do any of the following:
 - To review processing statistics for the job, refer to the **Statistics** section of the panel. This section displays statistics such as the number of times that the job has run and the approximate number of objects that the job has yet to process during its current run.

- To review log events for the job, choose **Show results** at the top of the panel, and then choose **Show CloudWatch logs**. Macie opens the Amazon CloudWatch console and displays a table of the log events that Macie published for the job.
- To review all the sensitive data findings that the job produced, choose **Show results** at the top of the panel, and then choose **Show findings**. Macie opens the **Findings** page and displays all the findings from the job. To review the details of a particular finding, choose the finding in the table and refer to the details panel.

Tip

In the finding details panel, you can use the link in the **Detailed result location** field to navigate to the corresponding sensitive data discovery result in Amazon S3:

- If the finding applies to a large archive or compressed file, the link displays the folder that contains the discovery results for the file. An archive or compressed file is *large* if it generates more than 100 discovery results.
 - If the finding applies to a small archive or compressed file, the link displays the file that contains the discovery results for the file. An archive or compressed file is *small* if it generates 100 or fewer discovery results.
 - If the finding applies to another type of file, the link displays the file that contains the discovery results for the file.
- To review all the sensitive data discovery results that the job produced, choose **Show results** at the top of the panel, and then choose **Show classifications**. Macie opens the Amazon S3 console and displays the folder that contains all the discovery results for the job. This option is available only after you configure Macie to [store your sensitive data discovery results \(p. 130\)](#) in an S3 bucket.

Managing sensitive data discovery jobs

To help you manage your sensitive data discovery jobs, Amazon Macie provides a complete inventory of your jobs in each AWS Region. With this inventory, you can manage your jobs as a single collection, and access the configuration settings, status, and processing statistics for individual jobs. You can also access the [sensitive data findings and other results \(p. 113\)](#) that each job produced.

In addition to these tasks, you can create custom variations of individual jobs—copy an existing job, adjust the settings for the copy, and then save the copy as a new job. This can be helpful for cases where you want to analyze different sets of data in the same way, or the same set of data in different ways. Or you want to adjust the configuration settings for an existing job—cancel the existing job, copy it, and then adjust and save the copy as a new job.

Topics

- [Reviewing your inventory of sensitive data discovery jobs \(p. 116\)](#)
- [Reviewing configuration settings for sensitive data discovery jobs \(p. 117\)](#)
- [Checking the status of sensitive data discovery jobs \(p. 119\)](#)
- [Pausing, resuming, or cancelling sensitive data discovery jobs \(p. 120\)](#)
- [Copying sensitive data discovery jobs \(p. 121\)](#)

Reviewing your inventory of sensitive data discovery jobs

The **Jobs** page on the Amazon Macie console provides information about all the sensitive data discovery jobs for your account in the current AWS Region. For each job, the table displays summary information that includes: the current status of the job; whether the job runs on a scheduled, periodic basis; and whether the job analyzes a specific number of S3 buckets or it analyzes S3 buckets that match runtime criteria. If you choose a job in the table, the details panel displays the configuration settings and other information about the job.

To review your job inventory

1. Open the Amazon Macie console at <https://console.aws.amazon.com/maciek/>.
2. In the navigation pane, choose **Jobs**. The **Jobs** page opens and displays the number of jobs in your inventory and a table of those jobs.
3. To find a specific job more quickly, do any of the following:
 - To sort the table by a specific field, click the column heading for the field. To change the sort order, click the column heading again.
 - To show only those jobs that have a specific value for a field, place your cursor in the filter bar. In the menu that appears, choose the field to use for the filter, and enter the value for the filter. Then choose **Apply**.
 - To hide jobs that have a specific value for a field, place your cursor in the filter bar. In the menu that appears, choose the field to use for the filter, and enter the value for the filter. Then choose **Apply**. In the filter bar, choose the equals icon (●) in the filter box. This changes the filter's operator from *equals* to *not equals* (⊘).
 - To remove a filter, choose the remove filter icon (⊗) in the filter box for the filter to remove.
4. To review the configuration settings and other details for a particular job, choose the job's name in the table, and then refer to the details panel.

Reviewing configuration settings for sensitive data discovery jobs

On the Amazon Macie console, you can use the details panel on the **Jobs** page to review configuration settings and other information about individual sensitive data discovery jobs. For example, you can review a list of the S3 buckets that a job is configured to analyze and which managed data identifiers a job uses to analyze objects in those buckets.

Note

You can't change any configuration settings for an existing job. This helps ensure that you have an immutable history of sensitive data findings and discovery results for data privacy and protection audits or investigations that you perform. If you want to change an existing job, [cancel the job \(p. 120\)](#). Then [copy the job \(p. 121\)](#), configure the copy to use the settings that you want, and save the copy as a new job.

If you do this, you should also take steps to ensure that the new job doesn't analyze existing data in the same way again. To do this, note the date and time when you cancel the existing job. Then configure the scope of the new job to include only those objects that are created or changed after you cancel the original job. For example, use [object criteria \(p. 95\)](#) to add a **Last modified** exclude condition that specifies the date and time when you cancelled the original job.

To review a job's configuration settings

1. Open the Amazon Macie console at <https://console.aws.amazon.com/maciek/>.
2. In the navigation pane, choose **Jobs**.
3. On the **Jobs** page, choose the name of the job whose settings you want to review. The details panel displays the configuration settings and other information about the job. Depending on the job's settings, the panel contains the following sections.

General information

This section indicates the current status of the job and it provides general information about the job—for example, the Amazon Resource Name (ARN) of the job and the most recent date and

time when the job started to run. If you paused the job during the past 30 days, this section also indicates when you paused the job and when the job or job run will expire if you don't resume it.

Statistics

This section shows processing statistics for the job—for example, the number of times that the job has run and the approximate number of objects that the job has yet to process during its current run.

Scope

This section indicates how often the job runs. It also shows settings that refine the job's scope—for example, the sampling depth and any [object criteria \(p. 95\)](#) that include or exclude S3 objects from the job's analysis.


S3 buckets

This section appears in the panel if the job is configured to analyze buckets that you explicitly selected when you created the job. It indicates the number of AWS accounts that the job is configured to analyze data for. It also indicates the number of buckets that the job is configured to analyze and the names of those buckets (grouped by account).

To show the complete list of accounts and buckets in JSON format, choose the number in the **Total buckets** field.

S3 bucket criteria

This section appears in the panel if the job uses runtime criteria to determine which buckets to analyze. It lists the criteria that the job is configured to use.

To show the criteria in JSON format, choose **Details**, and then choose the **Criteria** tab in the window that appears. To review a table of buckets that currently match the criteria, choose **Details**, and then choose the **Matching buckets** tab in the window that appears. Optionally choose refresh () to retrieve the latest data.


Tip

If the job has already run, you can also determine whether any buckets matched the criteria when the job ran and, if so, the names of those buckets. You can do this by reviewing [log events \(p. 104\)](#) for the job: choose **Show results** at the top of the panel, and then choose **Show CloudWatch logs**. Macie opens the Amazon CloudWatch console and displays a table of log events for the job. The events include a `BUCKET_MATCHED_THE_CRITERIA` event for each bucket that matched the criteria and was included in the job's analysis.

Custom data identifiers

This section appears in the panel if the job is configured to use one or more [custom data identifiers \(p. 64\)](#). It specifies the names of those custom data identifiers.

Allow lists

This section appears in the panel if the job is configured to use one or more [allow lists \(p. 70\)](#). It specifies the names of those lists. To review the settings and status of a list, choose the link icon () next to the list's name.

Managed data identifiers

This section indicates which [managed data identifiers \(p. 45\)](#) the job is configured to use. This is determined by the managed data identifier selection type for the job:

- **Include all** – Use all the managed data identifiers that are available when the job runs.
 - **Include selected** – Use only the managed data identifiers listed in the **Selections** section.
 - **Exclude selected** – Use all the managed data identifiers that are available when the job runs, except the ones listed in the **Selections** section.
-

- **Exclude all** – Don't use any managed data identifiers.

To review these settings in JSON format, choose **Details**.

Tags

This section appears in the panel if tags are associated with the job. It lists those tags.

A *tag* is a label that you define and assign to certain types of AWS resources. Each tag consists of a required tag key and an optional tag value. Tags can help you identify, categorize, and manage resources in different ways, such as by purpose, owner, environment, or other criteria. To learn more, see [Tagging Amazon Macie resources \(p. 315\)](#).

4. To review and save the job's settings in JSON format, choose the unique identifier for the job (**Job ID**) at the top of the panel, and then choose **Download**.

Checking the status of sensitive data discovery jobs

When you create a sensitive data discovery job, its initial status is **Active (Running)** or **Active (Idle)**, depending on the job's type and schedule. The job then passes through additional states, which you can monitor as the job progresses.

Tip

In addition to monitoring the overall status of a job, you can monitor specific events that occur as a job progresses. You can do this by using logging data that Macie automatically publishes to Amazon CloudWatch Logs. The data in these logs provides a record of changes to a job's status and details about any account- or bucket-level errors that occur while a job runs. For more information, see [Monitoring jobs \(p. 104\)](#).

To check the status of a job

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, choose **Jobs**.
3. On the **Jobs** page, locate the job whose status you want to check. The **Status** field indicates the current status of the job.

Active (Idle)

For a periodic job, the previous run is complete and the next scheduled run is pending. This value doesn't apply to one-time jobs.

Active (Running)

For a one-time job, the job is currently in progress. For a periodic job, a scheduled run is in progress.

Cancelled

For any type of job, the job was stopped permanently (cancelled).

A job has this status if you explicitly cancelled it or, if it's a one-time job, you paused the job and didn't resume it within 30 days. A job can also have this status if you previously [suspended Macie \(p. 326\)](#) in the current AWS Region.

Complete

For a one-time job, the job ran successfully and is now complete. This value doesn't apply to periodic jobs. Instead, the status of a periodic job changes to **Active (Idle)** when each run completes successfully.

Paused (By Macie)

For any type of job, the job was stopped temporarily (paused) by Macie.

A job has this status if completion of the job or a job run would exceed the monthly [sensitive data discovery quota](#) (p. 328) for your account. When this happens, Macie automatically pauses the job until the next calendar month starts (and the monthly quota is reset for your account) or you increase the quota for your account.

If you're the Macie administrator for an organization and you configured the job to analyze data for member accounts, the job can also have this status if completion of the job or a job run would exceed the monthly sensitive data discovery quota for a member account.

If a job is running and the analysis of eligible objects reaches this quota for a member account, the job stops analyzing objects that are owned by the account. When the job finishes analyzing objects for all other accounts that haven't met the quota, Macie automatically pauses the job. If it's a one-time job, Macie automatically resumes the job when the next calendar month starts or the quota is increased for all the affected accounts, whichever occurs first. If it's a periodic job, Macie automatically resumes the job when the next run is scheduled to start or the next calendar month starts, whichever occurs first. If a scheduled run starts before the next calendar month starts or the quota is increased for an affected account, the job doesn't analyze objects that are owned by the account.

Paused (By user)

For any type of job, the job was stopped temporarily (paused) by you.

If you pause a one-time job and you don't resume it within 30 days, the job expires and Macie cancels it. If you pause a periodic job while it's actively running and you don't resume it within 30 days, the job's run expires and Macie cancels the run. To check the expiration date for a paused job or job run, choose the job's name in the table, and then refer to the **Expires** field in the **Status details** section of the details panel.

If a job is cancelled or paused, you can refer to the job's details to determine whether the job started to run or, for a periodic job, ran at least once before it was cancelled or paused. To do this, choose the job's name in the table, and then refer to the details panel. In the panel, the **Number of runs** field indicates the number of times that the job has run. The **Last run time** field indicates the most recent date and time when the job started to run.

Depending on the job's current status, you can optionally pause, resume, or cancel the job.

Pausing, resuming, or cancelling sensitive data discovery jobs

After you create a sensitive data discovery job, you can pause it temporarily or cancel it permanently. When you pause a job that's actively running, Macie immediately begins to pause all processing tasks for the job. When you cancel a job that's actively running, Macie immediately begins to stop all processing tasks for the job. You can't resume or restart a job after it's cancelled.

If you pause a one-time job, you can resume it within 30 days. When you resume the job, Macie immediately resumes processing from the point where you paused the job—Macie doesn't restart the job from the beginning. If you don't resume a one-time job within 30 days of pausing it, the job expires and Macie cancels it.

If you pause a periodic job, you can resume it at any time. If you resume a periodic job and the job was idle when you paused it, Macie resumes the job according to the schedule and other configuration settings that you chose when you created the job. If you resume a periodic job and the job was actively running when you paused it, how Macie resumes the job depends on when you resume the job:

- If you resume the job within 30 days of pausing it, Macie immediately resumes the latest scheduled run from the point where you paused the job—Macie doesn't restart the run from the beginning.
- If you don't resume the job within 30 days of pausing it, the latest scheduled run expires and Macie cancels all remaining processing tasks for the run. When you subsequently resume the job, Macie resumes the job according to the schedule and other configuration settings that you chose when you created the job.

To help you determine when a paused job or job run will expire, Macie adds an expiration date to the job's details while the job is paused. To check this date, choose the job's name in the table on the **Jobs** page, and then refer to the **Expires** field in the **Status details** section of the details panel. In addition, we notify you approximately seven days before the job or job run will expire. We notify you again when the job or job run expires and is cancelled. To notify you, we send email to the address that's associated with your AWS account. We also create AWS Health events and Amazon CloudWatch Events for your account.

To pause, resume, or cancel a job

1. Open the Amazon Macie console at <https://console.aws.amazon.com/maciek/>.
2. In the navigation pane, choose **Jobs**.
3. On the **Jobs** page, select the check box for the job that you want to pause, resume, or cancel, and then do one of the following on the **Actions** menu:
 - To pause the job temporarily, choose **Pause**. This option is available only if the job's current status is **Active (Idle)**, **Active (Running)**, or **Paused (By Macie)**.
 - To resume the job, choose **Resume**. This option is available only if the job's current status is **Paused (By user)**.
 - To cancel the job permanently, choose **Cancel**. If you choose this option, you can't subsequently resume or restart the job.

Copying sensitive data discovery jobs

To quickly create a new sensitive data discovery job that's similar to an existing job, you can create a copy of the job, edit the copy's settings, and then save the copy as a new job. This can be helpful for cases where you want to create a custom variation of an existing job. Or you want to adjust the configuration settings for an existing job by cancelling the job, and then copying, changing, and saving the settings as a new job.

To copy a job

1. Open the Amazon Macie console at <https://console.aws.amazon.com/maciek/>.
2. In the navigation pane, choose **Jobs**.
3. Select the check box for the job that you want to copy.
4. On the **Actions** menu, choose **Copy to new**.
5. Complete the steps on the console to review and adjust the settings for the copy of the job. On the **Scope** page, consider choosing options that prevent the job from analyzing existing data in the same way again:
 - For a one-time job, use [object criteria \(p. 95\)](#) to include only those objects that were created or changed after a certain time. For example, if you're creating a copy of a job that you cancelled, add a **Last modified** condition that specifies the date and time when you cancelled the existing job.
 - For a periodic job, clear the **Include existing objects** check box. If you do this, the first run of the job analyzes only those objects that are created or changed after you create the job and before the job's first run. You can also use [object criteria \(p. 95\)](#) to exclude objects that were last modified before a certain date and time.

6. When you finish, choose **Submit** to save the copy as a new job.

Forecasting and monitoring costs for sensitive data discovery jobs

Amazon Macie pricing is based partly on the amount of data that you analyze by running sensitive data discovery jobs. To forecast and monitor your estimated costs for running sensitive data discovery jobs, you can review cost estimates that Macie provides when you create a job and after you start running jobs.

To review and monitor your actual costs, you can use AWS Billing and Cost Management. AWS Billing and Cost Management provides features that are designed to help you track and analyze your costs for AWS services, and manage budgets for your account or organization. It also provides features that can help you forecast usage costs based on historical data. To learn more, see the [AWS Billing and Cost Management User Guide](#).

For information about Macie pricing, see [Amazon Macie pricing](#).

Topics

- [Forecasting the cost of a sensitive data discovery job \(p. 122\)](#)
- [Monitoring estimated costs for sensitive data discovery jobs \(p. 123\)](#)

Forecasting the cost of a sensitive data discovery job

When you create a sensitive data discovery job, Amazon Macie can calculate and display estimated costs during two key steps in the job creation process: when you review the table of S3 buckets that you selected for the job (step 2) and when you review all the settings for the job (step 8). These estimates can help you determine whether to adjust the job's settings before you save the job. The availability and nature of the estimates depends on the settings that you choose for the job.

Reviewing estimated costs for individual buckets (step 2)

If you explicitly select individual buckets for a job to analyze, you can review the estimated cost of analyzing objects in each of those buckets. Macie displays these estimates during step 2 of the job creation process, when you review your bucket selections. In the table for this step, the **Estimated cost** field indicates the total estimated cost (in US Dollars) of running the job once to analyze objects in a bucket.

Each estimate reflects the projected amount of uncompressed data that the job will analyze in a bucket, based on the size and types of objects that are currently stored in the bucket. The estimate also reflects Macie pricing for the current AWS Region.

Only classifiable objects are included in the cost estimate for a bucket. A *classifiable object* is an S3 object that uses a supported Amazon S3 storage class (S3 Intelligent-Tiering, S3 One Zone-IA, S3 Standard, or S3 Standard-IA) and has a file name extension for a [supported file or storage format \(p. 124\)](#). If any classifiable objects are compressed or archive files, the estimate assumes that the files use a 3:1 compression ratio and the job can analyze all extracted files.

Reviewing the total estimated cost of a job (step 8)

If you create a one-time job or you create and configure a periodic job to include existing S3 objects, Macie calculates and displays the job's total estimated cost during the final step of the job creation process. You can review this estimate while you review and verify all the settings that you selected for the job.

This estimate indicates the total projected cost (in US Dollars) of running the job once in the current Region. The estimate reflects the projected amount of uncompressed data that the job will analyze. It's based on the size and types of objects that are currently stored in buckets that you explicitly selected for the job or up to 500 buckets that currently match bucket criteria that you specified for the job, depending on the job's settings.

Note that this estimate doesn't reflect any options that you selected to refine and reduce the scope of the job—for example, a lower sampling depth, or criteria that exclude certain S3 objects from the job. It also doesn't reflect your monthly [sensitive data discovery quota \(p. 328\)](#), which might limit the scope and cost of the job's analysis, or any discounts that might apply to your account.

In addition to the total estimated cost of the job, the estimate provides aggregated data that offers insight into the projected scope and cost of the job:

- **Size** values indicate the total storage size of the objects that the job can and can't analyze.
- **Object count** values indicate the total number of objects that the job can and can't analyze.

In these values, a **Classifiable** object is an S3 object that uses a supported Amazon S3 storage class (S3 Intelligent-Tiering, S3 One Zone-IA, S3 Standard, or S3 Standard-IA) and has a file name extension for a [supported file or storage format \(p. 124\)](#). Only classifiable objects are included in the cost estimate. A **Not classifiable** object is an object that doesn't use a supported Amazon S3 storage class or doesn't have a file name extension for a supported file or storage format. These objects aren't included in the cost estimate.

The estimate provides additional aggregated data for S3 objects that are compressed or archive files. The **Compressed** value indicates the total storage size of objects that use a supported Amazon S3 storage class and have a file name extension for a supported type of compressed or archive file. The **Uncompressed** value indicates the approximate size of these objects if they're decompressed, based on a specified compression ratio. This data is relevant due to the way that Macie analyzes compressed files and archive files.

When Macie analyzes a compressed or archive file, it inspects both the full file and the contents of the file. To inspect the file's contents, Macie decompresses the file, and then inspects each extracted file that uses a supported format. The actual amount of data that a job analyzes therefore depends on:

- Whether a file uses compression and, if so, the compression ratio that it uses.
- The number, size, and format of the extracted files.

By default, Macie assumes the following when it calculates cost estimates for a job:

- All compressed and archive files use a 3:1 compression ratio.
- All the extracted files use a supported file or storage format.

These assumptions can result in a larger size estimate for the scope of the data that the job will analyze, and, consequently, a higher cost estimate for the job.

You can recalculate the job's total estimated cost based on a different compression ratio. To do this, choose the ratio from the **Choose an estimated compression ratio** list in the **Estimated cost** section. Macie then updates the estimate to match your selection.

For more information about how Macie calculates estimated costs, see [Understanding how estimated usage costs are calculated \(p. 278\)](#).

Monitoring estimated costs for sensitive data discovery jobs

If you're already running sensitive data discovery jobs, the **Usage** page on the Amazon Macie console can help you monitor the estimated cost of those jobs. The page shows your estimated costs (in US

Dollars) for using Macie in the current AWS Region during the current calendar month. For information about how Macie calculates these estimates, see [Understanding how estimated usage costs are calculated \(p. 278\)](#).

To review your estimated costs for running jobs

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. By using the AWS Region selector in the upper-right corner of the page, select the Region in which you want to review your estimated costs.
3. In the navigation pane, choose **Usage**.
4. On the **Usage** page, refer to the breakdown of estimated costs for your account. The **Data discovery jobs** item reports the total estimated cost of the jobs that you've run thus far during the current month in the current Region.

If you're the Macie administrator for an organization, the **Estimated costs** section shows estimated costs for your organization overall for the current month in the current Region. To show the total estimated cost of the jobs that were run for a specific account, choose the account in the table. The **Estimated costs** section then shows a breakdown of estimated costs for the account, including the estimated cost of the jobs that were run. To show this data for a different account, choose the account in the table. To clear your account selection, choose **X** next to the account ID.

To review and monitor your actual costs, use [AWS Billing and Cost Management](#).

Supported file and storage formats in Amazon Macie

Amazon Macie can analyze data in many different formats, including commonly used compression and archive formats. This support applies to the use of [managed data identifiers \(p. 45\)](#) and the use of [custom data identifiers \(p. 64\)](#).

When Macie analyzes data, it performs a deep inspection that factors the file or storage format for the data. For data in a compressed or archive file, Macie inspects both the full file and the contents of the file. To inspect the file's contents, Macie decompresses the file, and then inspects each extracted file that uses a supported format. Macie can do this for as many as 1,000,000 files and up to a nested depth of 10 levels.

The following table lists and describes the types of file and storage formats that Macie can analyze to detect sensitive data. For each supported type, the table also lists the applicable file name extensions.

File or storage type	Description	File name extensions
Big data	Apache Avro object containers and Apache Parquet files	.avro, .parquet
Compression or archive	GNU Zip compressed archives, TAR archives, and ZIP compressed archives	.gz, .gzip, .tar, .zip
Document	Adobe Portable Document Format files, Microsoft Excel workbooks, and Microsoft Word documents	.doc, .docx, .pdf, .xls, .xlsx

File or storage type	Description	File name extensions
Text	Non-binary text files such as comma-separated values (CSV) files, Hypertext Markup Language (HTML) files, JavaScript Object Notation (JSON) files, JSON Lines files, plaintext documents, tab-separated values (TSV) files, and Extensible Markup Language (XML) files	.csv, .htm, .html, .json, .jsonl, .tsv, .txt, .xml, and others (depending on the type of non-binary text file)

Macie doesn't analyze data in images or audio, video, and other types of multimedia content.

For information about the quotas that apply to sensitive data discovery, see [Amazon Macie quotas \(p. 328\)](#).

Analyzing encrypted S3 objects with Amazon Macie

When you enable Amazon Macie for your AWS account, Macie creates a [service-linked role \(p. 302\)](#) that grants Macie the permissions that it requires to call Amazon Simple Storage Service (Amazon S3) and other AWS services on your behalf. A service-linked role simplifies the process of setting up an AWS service because you don't have to manually add permissions for the service to complete actions on your behalf. To learn more about this type of role, see [Using service-linked roles](#) in the *AWS Identity and Access Management User Guide*.

The permissions policy for the Macie service-linked role (`AWSServiceRoleForAmazonMacie`) allows Macie to perform actions that include retrieving information about your S3 buckets and objects, and retrieving and analyzing objects in your S3 buckets. If your account is the Macie administrator account for an organization, the policy also allows Macie to perform these actions on your behalf for member accounts in your organization.

If an S3 object is encrypted, the permissions policy for the Macie service-linked role typically grants Macie the permissions that it requires to decrypt the object. However, this depends on the type of encryption that was used. It can also depend on whether Macie is allowed to use the appropriate encryption key.

Topics

- [Encryption options for S3 objects \(p. 125\)](#)
- [Allowing Amazon Macie to use a customer managed AWS KMS key \(p. 127\)](#)

Encryption options for S3 objects

Amazon S3 supports multiple encryption options for S3 objects. For most of these options, Amazon Macie can decrypt an object by using the Macie service-linked role for your account. However, this depends on the type of encryption that was used to encrypt an object.

Server-side encryption with Amazon S3 managed keys (SSE-S3)

If an object is encrypted using server-side encryption with an Amazon S3 managed key, Macie can decrypt the object.

To learn about this type of encryption, see [Protecting data using server-side encryption with Amazon S3 managed encryption keys](#) in the *Amazon Simple Storage Service User Guide*.

Server-side encryption with AWS KMS keys (SSE-KMS)

If an object is encrypted using server-side encryption with an AWS managed AWS KMS key, Macie can decrypt the object.

If an object is encrypted using server-side encryption with a customer managed AWS KMS key, Macie can decrypt the object only if you [allow Macie to use the KMS key \(p. 127\)](#). Otherwise, Macie can only store and report metadata for the object.

To learn about this type of encryption, see [Protecting data using server-side encryption with AWS Key Management Service](#) in the *Amazon Simple Storage Service User Guide*.

Server-side encryption with customer-provided keys (SSE-C)

If an object is encrypted using server-side encryption with a customer-provided key, Macie can't decrypt the object. Macie can only store and report metadata for the object.

To learn about this type of encryption, see [Protecting data using server-side encryption with customer-provided encryption keys](#) in the *Amazon Simple Storage Service User Guide*.

Client-side encryption

If an object is encrypted using client-side encryption, Macie can't decrypt the object. Macie can only store and report metadata for the object. For example, Macie can report the size of the object and the tags that are associated with the object.

To learn about this type of encryption in the context of Amazon S3, see [Protecting data using client-side encryption](#) in the *Amazon Simple Storage Service User Guide*.

You can [filter your bucket inventory \(p. 32\)](#) in Macie to determine which S3 buckets contain objects that use certain types of encryption. You can also determine which buckets use certain types of server-side encryption by default when storing new objects. The following table provides some example filters that you can apply to your bucket inventory to find this information.

To show buckets that...	Apply this filter...
Contain objects that use SSE-C encryption	Object count by encryption is Customer managed and From = 1
Contain objects that use SSE-KMS encryption	Object count by encryption is SSE-KMS managed and From = 1
Contain objects that use SSE-S3 encryption	Object count by encryption is SSE-S3 managed and From = 1
Contain objects that use client-side encryption (or aren't encrypted)	Object count by encryption is No encryption and From = 1
Encrypt new objects by default using SSE-KMS encryption	Default encryption = aws:kms
Encrypt new objects by default using SSE-S3 encryption	Default encryption = AES256

If a bucket is configured to encrypt new objects by default using SSE-KMS encryption, you can also determine which AWS KMS key is used. To do this, choose the bucket in the table on the **S3 buckets**

page. In the bucket details panel, under **Server-side encryption**, refer to the **KMS key** field. This field shows the Amazon Resource Name (ARN) or unique identifier (key ID) for the key.

Allowing Amazon Macie to use a customer managed AWS KMS key

If an S3 object is encrypted using a customer managed AWS KMS key (SSE-KMS encryption), Amazon Macie can decrypt the object only if Macie is allowed to use the KMS key. How to provide this access depends on whether the account that owns the key also owns the S3 bucket that stores the object:

- If the same account owns the KMS key and the bucket, a user of the account has to update the key's policy.
- If one account owns the KMS key and a different account owns the bucket, the account that owns the key has to allow cross-account access to the key.

This topic describes how to perform these tasks and provides examples for both scenarios.

Allowing same-account access to a customer managed key

If the same account owns both the AWS KMS key and the bucket, a user of the account has to add a statement to the policy for the KMS key. The additional statement must allow the Macie service-linked role for the account to use the key to decrypt data. For detailed information about updating a key policy, see [Changing a key policy](#) in the *AWS Key Management Service Developer Guide*.

In the statement:

- The **Principal** element must specify the Amazon Resource Name (ARN) of the Macie service-linked role for the account that owns the KMS key and the bucket.

If the account is in a manually enabled AWS Region, the ARN must also include the appropriate Region code for the Region. For example, if the account is in the Middle East (Bahrain) Region, which has the Region code *me-south-1*, the **Principal** element must specify `arn:aws:iam::123456789012:role/aws-service-role/macie.me-south-1.amazonaws.com/AWSServiceRoleForAmazonMacie`, where **123456789012** is the account ID for the account.

- The **Action** array must specify the `kms:Decrypt` action. This is the only AWS KMS action that Macie must be allowed to perform to decrypt an object that was encrypted with the key.

The following is an example of the statement to add to the policy for a KMS key.

```
{
  "Sid": "Allow the Macie service-linked role to use the key",
  "Effect": "Allow",
  "Principal": {
    "AWS": "arn:aws:iam::123456789012:role/aws-service-role/macie.amazonaws.com/AWSServiceRoleForAmazonMacie"
  },
  "Action": [
    "kms:Decrypt"
  ],
  "Resource": "*"
}
```

In the preceding example:

- The `AWS` field in the `Principal` element specifies the ARN of the Macie service-linked role (`AWSServiceRoleForAmazonMacie`) for the account. It allows the Macie service-linked role to perform the action specified by the policy statement. `123456789012` is an example account ID. Replace this ID with the account ID for the account that owns the KMS key and the bucket.
- The `Action` array specifies the action that the Macie service-linked role is allowed to perform using the KMS key—decrypt ciphertext that was encrypted with the key.

Where you add this statement to a key policy depends on the structure and elements that the policy currently contains. When you add the statement, ensure that the syntax is valid. Key policies use JSON format. This means that you have to also add a comma before or after the statement, depending on where you add the statement to the policy.

Allowing cross-account access to a customer managed key

If one account owns the AWS KMS key (*key owner*) and a different account owns the bucket (*bucket owner*), the key owner has to provide the bucket owner with cross-account access to the KMS key. To do this, the key owner first ensures that the key's policy allows the bucket owner to both use the key and create a grant for the key. The bucket owner then creates a grant for the key. The grant delegates the relevant permissions to the Macie service-linked role for the bucket owner's account.

A *grant* is a policy instrument that allows AWS principals to use KMS keys in cryptographic operations if the conditions specified by the grant are met. To learn about grants, see [Grants in AWS KMS](#) in the *AWS Key Management Service Developer Guide*.

In the key policy, the key owner should ensure that the policy includes two statements. The first statement allows the bucket owner to use the key to decrypt data. The second statement allows the bucket owner to create a grant for the Macie service-linked role for the bucket owner's account. For detailed information about updating a key policy, see [Changing a key policy](#) in the *AWS Key Management Service Developer Guide*.

In the first statement, the `Principal` element must specify the ARN of the bucket owner's account. The `Action` array must specify the `kms:Decrypt` action. This is the only AWS KMS action that Macie must be allowed to perform to decrypt an object that was encrypted with the key.

The following is an example of this statement in the policy for a KMS key.

```
{
  "Sid": "Allow account 111122223333 to use the key",
  "Effect": "Allow",
  "Principal": {
    "AWS": "arn:aws:iam::111122223333:root"
  },
  "Action": [
    "kms:Decrypt"
  ],
  "Resource": "*"
}
```

In the preceding example:

- The `AWS` field in the `Principal` element specifies the ARN of the bucket owner's account (`111122223333`). It allows the bucket owner to perform the action specified by the policy statement. `111122223333` is an example account ID. Replace this ID with the account ID for the bucket owner's account.
- The `Action` array specifies the action that the bucket owner is allowed to perform using the KMS key—decrypt ciphertext that was encrypted with the key.

The second statement in the key policy allows the bucket owner to create a grant for the Macie service-linked role for their account. In this statement, the `Principal` element must specify the ARN of the bucket's owner's account. The `Action` array must specify the `kms:CreateGrant` action. A `Condition` element can filter access to the `kms:CreateGrant` action specified in the statement.

The following is an example of this statement in the policy for a KMS key.

```
{
  "Sid": "Allow account 111122223333 to create a grant",
  "Effect": "Allow",
  "Principal": {
    "AWS": "arn:aws:iam::111122223333:root"
  },
  "Action": [
    "kms:CreateGrant"
  ],
  "Resource": "*",
  "Condition": {
    "StringEquals": {
      "kms:GranteePrincipal": "arn:aws:iam::111122223333:role/aws-service-role/macie.amazonaws.com/AWSServiceRoleForAmazonMacie"
    }
  }
}
```

In the preceding example:

- The `AWS` field in the `Principal` element specifies the ARN of the bucket owner's account (`111122223333`). It allows the bucket owner to perform the action specified by the policy statement. `111122223333` is an example account ID. Replace this ID with the account ID for the bucket owner's account.
- The `Action` array specifies the action that the bucket owner is allowed to perform on the KMS key—create a grant for the key.
- The `Condition` element uses the `StringEquals` [condition operator](#) and the `kms:GranteePrincipal` [condition key](#) to filter access to the action specified by the policy statement. In this case, the bucket owner can create a grant only for the specified `GranteePrincipal`, which is the ARN of the Macie service-linked role for the bucket owner's account. In that ARN, `111122223333` is an example account ID. Replace this ID with the account ID for the bucket owner's account.

If the bucket owner's account is in a manually enabled AWS Region, also include the appropriate Region code in the ARN of the Macie service-linked role. For example, if the account is in the Middle East (Bahrain) Region, which has the Region code `me-south-1`, replace `macie.amazonaws.com` with `macie.me-south-1.amazonaws.com` in the ARN.

Where the key owner adds these statements to the key policy depends on the structure and elements that the policy currently contains. When the key owner adds the statement, they should ensure that the syntax is valid. Key policies use JSON format. This means that the key owner has to also add a comma before or after the statement, depending on where they add the statement to the policy.

After the key owner updates the key policy as necessary, the bucket owner must create a grant for the key. The grant delegates the relevant permissions to the Macie service-linked role for their (the bucket owner's) account. Before the bucket owner creates the grant, they should verify that they're allowed to perform the `kms:CreateGrant` action for their account. This action allows them to add a grant to an existing, customer managed KMS key.

To create the grant, the bucket owner can use the [CreateGrant](#) operation of the AWS Key Management Service API. When the bucket owner creates the grant, they should specify the following values for the required parameters:

- **GranteePrincipal** – The ARN of the Macie service-linked role (`AWSServiceRoleForAmazonMacie`) for their account. This value should be `arn:aws:iam::111122223333:role/aws-service-role/macie.amazonaws.com/AWSServiceRoleForAmazonMacie`, where **111122223333** is the account ID for the bucket owner's account.

If their account is in a manually enabled Region, the ARN must include the appropriate Region code. For example, if the account is in the Middle East (Bahrain) Region, which has the Region code `me-south-1`, the ARN should be `arn:aws:iam::111122223333:role/aws-service-role/macie.me-south-1.amazonaws.com/AWSServiceRoleForAmazonMacie`, where **111122223333** is the account ID for the bucket owner's account.

- **KeyId** – The ARN of the KMS key. For cross-account access to a KMS key, this value must be an ARN. It can't be a key ID.
- **Operations** – The AWS KMS decrypt action (`Decrypt`). This is the only AWS KMS action that Macie must be allowed to perform to decrypt an object that was encrypted with the KMS key.

The following example shows how to use the [AWS Command Line Interface \(AWS CLI\)](#) to create a grant for a customer managed KMS key. The example uses the `create-grant` command of the AWS Key Management Service API. The example is formatted for Microsoft Windows and it uses the caret (^) line-continuation character to improve readability.

```
C:\> aws kms create-grant ^
--key-id arn:aws:kms:us-east-1:123456789012:key/1234abcd-12ab-34cd-56ef-1234567890ab ^
--grantee-principal arn:aws:iam::111122223333:role/aws-service-role/macie.amazonaws.com/
AWSServiceRoleForAmazonMacie ^
--operations "Decrypt"
```

Where:

- `key-id` specifies the ARN of the KMS key to apply the grant to.
- `grantee-principal` specifies the ARN of the Macie service-linked role for the account that's allowed to perform the operation specified by the grant. This value should match the ARN that's specified by the `kms:GranteePrincipal` condition of the second statement in the key policy.
- `operations` specifies the operation that the grant allows the specified principal to perform—decrypt ciphertext that was encrypted with the key.

If the command runs successfully, AWS KMS responds with output that's similar to the following.

```
{
  "GrantToken": "<grant token>",
  "GrantId": "1a2b3c4d2f5e69f440bae30eaec9570bb1fb7358824f9ddfa1aa5a0dab1a59b2"
}
```

Where `GrantToken` is a unique, non-secret, variable-length, base64-encoded string that represents the grant that was created, and `GrantId` is the unique identifier for the grant.

Storing and retaining sensitive data discovery results with Amazon Macie

When you run a sensitive data discovery job, Amazon Macie creates a record for each Amazon Simple Storage Service (Amazon S3) object that you configure the job to analyze. This includes objects that don't contain sensitive data, and therefore don't produce sensitive data findings, and objects that Macie can't analyze due to issues such as permissions settings or use of an unsupported format. If an

object does contain sensitive data, the record includes data from the corresponding finding. It provides additional information too, such as the location of as many as 1,000 occurrences of each type of sensitive data that Macie found in the object. Macie stores these records, referred to as *sensitive data discovery results*, for 90 days. To learn more about sensitive data discovery results, see [Reviewing job statistics and results \(p. 113\)](#).

To access your sensitive data discovery results and enable long-term storage and retention of them, configure Macie to store the results in an S3 bucket and encrypt them with an AWS Key Management Service (AWS KMS) key. If you do this, Macie writes your sensitive data discovery results to JSON Lines (.jsonl) files, which it adds to the S3 bucket as GNU Zip (.gz) files. The S3 bucket can then serve as a definitive, long-term repository for all of your sensitive data discovery results.

This topic guides you through the process of using the AWS Management Console to configure this type of repository for your discovery results. The configuration is a combination of an S3 bucket that stores the results, an AWS KMS key that encrypts the results, and Macie settings that indicate which bucket and key to use. If you prefer to configure the Macie settings programmatically, you can use the [PutClassificationExportConfiguration](#) operation of the Amazon Macie API.

When you configure the settings in Macie, your choices apply only to the current AWS Region. If you're the Macie administrator for an organization, your choices apply only to your account. They don't apply to any associated member accounts.

If you use Macie in multiple Regions, configure the repository settings for each Region in which you use Macie. If you prefer to store all discovery results for all Regions in one S3 bucket, you can do this by choosing the same bucket, located in one specific Region, for each Region in which you use Macie.

Tasks

- [Step 1: Verify your permissions \(p. 131\)](#)
- [Step 2: Choose an AWS KMS key and update the key policy \(p. 132\)](#)
- [Step 3: Specify the S3 bucket to use \(p. 134\)](#)

Step 1: Verify your permissions

Before you configure a repository for your sensitive data discovery results, verify that you have the permissions that you need. You can do this by using the AWS Identity and Access Management (IAM) console:

1. Sign in to the AWS Management Console and open the IAM console at <https://console.aws.amazon.com/iam/>.
2. In the navigation pane, choose **Users**.
3. Choose your user name.

The **Permissions** tab lists all the IAM policies that are attached to your user name. Choose a policy to show its details. Then compare the information in the policy to the following list of actions that you must be allowed to perform to configure the repository.

Macie

For Macie, verify that you're allowed to perform the following action:

```
macie2:PutClassificationExportConfiguration
```

This action allows you to add or change the repository settings in Macie.

Amazon S3

For Amazon S3, verify that you're allowed to perform the following actions:

- `s3:CreateBucket`
- `s3:GetBucketLocation`
- `s3:ListAllMyBuckets`
- `s3:PutBucketAcl`
- `s3:PutBucketPolicy`
- `s3:PutBucketPublicAccessBlock`
- `s3:PutObject`

These actions allow you to access and configure an S3 bucket that can serve as the repository.

AWS KMS

To use the Amazon Macie console to add or change the repository settings, also verify that you're allowed to perform the following AWS KMS actions:

- `kms:DescribeKey`
- `kms:ListAliases`

These actions allow you to retrieve information about AWS KMS keys that can encrypt data in the repository. If you plan to create a new AWS KMS key to encrypt the data, you also need to be allowed to perform the following actions: `kms:CreateKey`, `kms:GetKeyPolicy`, and `kms:PutKeyPolicy`.

If you're not allowed to perform one or more of the preceding actions, ask your AWS administrator for assistance before you proceed to the next step.

Step 2: Choose an AWS KMS key and update the key policy

After you verify your permissions, determine which AWS KMS key you want Macie to use to encrypt your sensitive data discovery results. The key must be a customer managed, symmetric encryption KMS key that's in the same AWS Region as the S3 bucket where you want to store the results.

The key can be an existing KMS key from your own account, or an existing KMS key that another account owns. If you want to use a new KMS key, create the key before proceeding. If you want to use an existing key that another account owns, obtain the Amazon Resource Name (ARN) of the key. You'll need to enter this ARN when you configure the repository settings in Macie. For information about creating and reviewing the settings for KMS keys, see [Managing keys](#) in the *AWS Key Management Service Developer Guide*.

After you determine which KMS key you want Macie to use, give Macie permission to use the key. Otherwise, Macie won't be able to encrypt or store discovery results in the repository. To give Macie permission to use the key, change the key policy for the key. For detailed information about key policies and managing access to KMS keys, see [Key policies in AWS KMS](#) in the *AWS Key Management Service Developer Guide*.

To change the key policy

1. Open the AWS KMS console at <https://console.aws.amazon.com/kms>.
2. To change the AWS Region, use the Region selector in the upper-right corner of the page.
3. Choose the key that you want to use to encrypt the results.
4. On the **Key policy** tab, choose **Edit**.
5. Copy the following statement to your clipboard and then add it to the policy:

```
{
```

```
"Sid": "Allow Macie to use the key",
"Effect": "Allow",
"Principal": {
  "Service": "macie.amazonaws.com"
},
"Action": [
  "kms:GenerateDataKey",
  "kms:Encrypt"
],
"Resource": "*",
"Condition": {
  "StringEquals": {
    "aws:SourceAccount": "111122223333"
  },
  "ArnLike": {
    "aws:SourceArn": [
      "arn:aws:macie2:Region:111122223333:export-configuration:*",
      "arn:aws:macie2:Region:111122223333:classification-job/*"
    ]
  }
}
}
```

When you add the statement, make sure that the syntax is valid. Policies use JSON format. This means that you need to also add a comma before or after the statement, depending on where you add the statement to the policy. If you add the statement as the last statement, add a comma after the closing curly brace for the preceding statement. If you add it as the first statement or between two existing statements, add a comma after the closing curly brace for the statement.

6. Update the statement with the correct values for your environment:

- In the `Condition` fields, replace the placeholder values, where:
 - `111122223333` is the account ID for your AWS account.
 - `Region` is the AWS Region in which you're using Macie and you want to allow Macie to use the key.

If you use Macie in multiple Regions and want to allow Macie to use the key in additional Regions, add `aws:SourceArn` conditions for each additional Region. For example:

```
"aws:SourceArn": [
  "arn:aws:macie2:us-east-1:111122223333:export-configuration:*",
  "arn:aws:macie2:us-east-1:111122223333:classification-job/*",
  "arn:aws:macie2:us-west-2:111122223333:export-configuration:*",
  "arn:aws:macie2:us-west-2:111122223333:classification-job/*"
]
```

Alternatively, you can allow Macie to use the key in *all* Regions. To do this, replace the placeholder value with the wildcard character (`*`). For example:

```
"aws:SourceArn": [
  "arn:aws:macie2:*:111122223333:export-configuration:*",
  "arn:aws:macie2:*:111122223333:classification-job/*"
]
```

- If you're using Macie in a manually enabled AWS Region, add the appropriate Region code to the value for the `Service` field. For example, if you're using Macie in the Middle East (Bahrain) Region, which has the Region code `me-south-1`, replace `macie.amazonaws.com` with `macie.me-south-1.amazonaws.com`.

Note that the `Condition` fields use two IAM global condition keys:

- [aws:SourceAccount](#) – This condition allows Macie to perform the specified actions only for your account. More specifically, it determines which account can perform the specified actions for the resources and actions specified by the `aws:SourceArn` condition.

To allow Macie to perform the specified actions for additional accounts, add the account ID for each additional account to this condition. For example:

```
"aws:SourceAccount": [111122223333,444455556666]
```

- [aws:SourceArn](#) – This condition prevents other AWS services from performing the specified actions. It also prevents Macie from using the key while performing other actions for your account. In other words, it allows Macie to encrypt S3 objects with the key only if the objects are sensitive data discovery results, and only if those results are for sensitive data discovery jobs that are created by the account and in the Region specified in the condition.

To allow Macie to perform the specified actions for additional accounts, add ARNs for each additional account to this condition. For example:

```
"aws:SourceArn": [  
  "arn:aws:macie2:us-east-1:111122223333:export-configuration:*",  
  "arn:aws:macie2:us-east-1:111122223333:classification-job/*",  
  "arn:aws:macie2:us-east-1:444455556666:export-configuration:*",  
  "arn:aws:macie2:us-east-1:444455556666:classification-job/*"  
]
```

The accounts specified by the `aws:SourceAccount` and `aws:SourceArn` conditions should match.

These conditions help prevent Macie from being used as a [confused deputy](#) during transactions with AWS KMS. Although we don't recommend it, you can remove these conditions from the statement.

7. When you finish adding and updating the statement, choose **Save changes**.

Step 3: Specify the S3 bucket to use

After you verify your permissions and choose the AWS KMS key to use, you're ready to specify which S3 bucket you want to use as the repository for your sensitive data discovery results. You have two options:

- **Use a new S3 bucket that Macie creates** – If you choose this option, Macie automatically creates a new S3 bucket for your discovery results. Macie also applies a bucket policy to the bucket. The policy allows Macie to add (put) objects to the bucket. To review this policy, choose **View policy** on the Amazon Macie console after you enter a name for the bucket.
- **Use an existing S3 bucket that you create** – If you prefer to store your discovery results in a particular S3 bucket that you create, create the bucket before you proceed. Then check the bucket's settings and update the bucket's policy to ensure that Macie can add (put) objects to the bucket. This topic explains which setting to check and how to update the policy. It also provides examples of the statements to add to the policy.

The following sections provide step-by-step instructions for each option. Choose the section for the option that you want.

Use a new S3 bucket that Macie creates

If you prefer to use a new S3 bucket that Macie creates for you, the final step in the process is to configure the repository settings in Macie.

To configure the repository settings in Macie

1. Open the Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, under **Settings**, choose **Discovery results**.
3. Under **Repository for sensitive data discovery results**, choose **Create bucket**.
4. In the **Create a bucket** box, enter a name for the bucket. The name must be unique across all S3 buckets. In addition, the name can consist only of lowercase letters, numbers, dots (.), and hyphens (-). For additional naming requirements, see [Bucket naming rules](#) in the *Amazon Simple Storage Service User Guide*.
5. Expand the **Advanced** section.
6. (Optional) To specify a prefix to use in the path to a location in the bucket, enter the prefix in the **Data discovery result prefix** box.

When you enter a value, Macie updates the example below the box to show the path to the bucket location where it will store your discovery results.

7. For **Block all public access**, choose **Yes** to enable all block public access settings for the bucket. For information about these settings, see [Blocking public access to your Amazon S3 storage](#) in the *Amazon Simple Storage Service User Guide*.
8. Under **Encryption settings**, specify the AWS KMS key that you want to use to encrypt the results:
 - To use a key from your own account, choose **Select a key from your account**. Then, in the **AWS KMS key** list, choose the key to use. The list displays customer managed, symmetric encryption KMS keys for your account.
 - To use a key that another account owns, choose **Enter the ARN of a key from another account**. Then, in the **AWS KMS key ARN** box, enter the Amazon Resource Name (ARN) of the key to use—for example, `arn:aws:kms:us-east-1:111122223333:key/1234abcd-12ab-34cd-56ef-1234567890ab`.
9. When you finish entering the settings, choose **Save**. Macie tests the settings to verify that they're correct. If any settings are incorrect, Macie displays an error message to help you address the issue.

After you save the repository settings, Macie adds existing discovery results for the preceding 90 days to the repository. Macie also starts adding new discovery results to the repository.

Use an existing S3 bucket that you create

If you prefer to store your sensitive data discovery results in a particular S3 bucket that you create, create and configure the bucket before you configure the repository settings in Macie.

If you enabled Object Lock for the bucket, ensure that you disable the default retention setting for that feature. Otherwise, Macie won't be able to add your discovery results to the bucket. For information about this setting, see [Using S3 Object Lock](#) in the *Amazon Simple Storage Service User Guide*.

Then add a bucket policy that allows Macie to retrieve information about the bucket and add (put) objects to the bucket. You can then configure the repository settings in Macie.

To add the bucket policy to the bucket

1. Open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. Choose the bucket that you want to store your discovery results in.
3. Choose the **Permissions** tab.
4. In the **Bucket policy** section, choose **Edit**.
5. Copy the following example policy to your clipboard:

```
{
```

Amazon Macie User Guide
Step 3: Specify an S3 bucket

```
"Version": "2012-10-17",
"Statement": [
  {
    "Sid": "Allow Macie to use the GetBucketLocation operation",
    "Effect": "Allow",
    "Principal": {
      "Service": "macie.amazonaws.com"
    },
    "Action": "s3:GetBucketLocation",
    "Resource": "arn:aws:s3:::myBucketName",
    "Condition": {
      "StringEquals": {
        "aws:SourceAccount": "111122223333"
      },
      "ArnLike": {
        "aws:SourceArn": [
          "arn:aws:macie2:Region:111122223333:export-configuration:*",
          "arn:aws:macie2:Region:111122223333:classification-job/*"
        ]
      }
    }
  },
  {
    "Sid": "Allow Macie to add objects to the bucket",
    "Effect": "Allow",
    "Principal": {
      "Service": "macie.amazonaws.com"
    },
    "Action": "s3:PutObject",
    "Resource": "arn:aws:s3:::myBucketName/[optional prefix/]*",
    "Condition": {
      "StringEquals": {
        "aws:SourceAccount": "111122223333"
      },
      "ArnLike": {
        "aws:SourceArn": [
          "arn:aws:macie2:Region:111122223333:export-configuration:*",
          "arn:aws:macie2:Region:111122223333:classification-job/*"
        ]
      }
    }
  },
  {
    "Sid": "Deny unencrypted object uploads. This is optional",
    "Effect": "Deny",
    "Principal": {
      "Service": "macie.amazonaws.com"
    },
    "Action": "s3:PutObject",
    "Resource": "arn:aws:s3:::myBucketName/[optional prefix/]*",
    "Condition": {
      "StringNotEquals": {
        "s3:x-amz-server-side-encryption": "aws:kms"
      }
    }
  },
  {
    "Sid": "Deny incorrect encryption headers. This is optional",
    "Effect": "Deny",
    "Principal": {
      "Service": "macie.amazonaws.com"
    },
    "Action": "s3:PutObject",
    "Resource": "arn:aws:s3:::myBucketName/[optional prefix/]*",
    "Condition": {
      "StringNotEquals": {
```

```
        "s3:x-amz-server-side-encryption-aws-kms-key-id":  
        "arn:aws:kms:Region:111122223333:key/KMSKeyId"  
    }  
  },  
  {  
    "Sid": "Deny non-HTTPS access",  
    "Effect": "Deny",  
    "Principal": "*",  
    "Action": "s3:*",  
    "Resource": "arn:aws:s3::myBucketName/*",  
    "Condition": {  
      "Bool": {  
        "aws:SecureTransport": "false"  
      }  
    }  
  }  
]  
}
```

6. Paste the example policy in the **Bucket policy** editor on the Amazon S3 console.
7. Update the bucket policy with the correct values for your environment:
 - In the optional statement that denies incorrect encryption headers:
 - Replace *myBucketName* with the name of the bucket.
 - In the `StringNotEquals` condition, replace the placeholder value for the specified field with the Amazon Resource Name (ARN) of the AWS KMS key to use for encryption of your discovery results.
 - In all other statements, replace the placeholder values, where:
 - *myBucketName* is the name of the bucket.
 - *Region* is the AWS Region in which you're using Macie and want to allow Macie to add discovery results to the bucket.

If you use Macie in multiple Regions and want to allow Macie to add results to the bucket for additional Regions, add `aws:SourceArn` conditions for each additional Region. For example:

```
"aws:SourceArn": [  
  "arn:aws:macie2:us-east-1:111122223333:export-configuration:*",  
  "arn:aws:macie2:us-east-1:111122223333:classification-job/*",  
  "arn:aws:macie2:us-west-2:111122223333:export-configuration:*",  
  "arn:aws:macie2:us-west-2:111122223333:classification-job/*"  
]
```

Alternatively, you can allow Macie to add results to the bucket for *all* Regions in which you use Macie. To do this, replace the placeholder value with the wildcard character (*). For example:

```
"aws:SourceArn": [  
  "arn:aws:macie2:*:111122223333:export-configuration:*",  
  "arn:aws:macie2:*:111122223333:classification-job/*"  
]
```

- *111122223333* is the account ID for your AWS account.
- If you're using Macie in a manually enabled AWS Region, add the appropriate Region code to the value for the `Service` field in each statement that specifies the Macie service principal. For example, if you're using Macie in the Middle East (Bahrain) Region, which has the Region code *me-south-1*, replace `macie.amazonaws.com` with `macie.me-south-1.amazonaws.com` in each applicable statement.

Note that the example policy includes statements that allow Macie to determine which Region the bucket resides in (`GetBucketLocation`) and to add objects to the bucket (`PutObject`). These statements define conditions that use two IAM global condition keys:

- `aws:SourceAccount` – This condition allows Macie to add sensitive data discovery results to the bucket only for your account. It prevents Macie from adding discovery results for other accounts to the bucket. More specifically, the condition specifies which account can use the bucket for the resources and actions specified by the `aws:SourceArn` condition.

To store results for additional accounts in the bucket, add the account ID for each additional account to this condition. For example:

```
"aws:SourceAccount": [111122223333,444455556666]
```

- `aws:SourceArn` – This condition restricts access to the bucket based on the source of the objects that are being added to the bucket. It prevents other AWS services from adding objects to the bucket. It also prevents Macie from adding objects to the bucket while performing other actions for your account. More specifically, the condition allows Macie to add objects to the bucket only if the objects are sensitive data discovery results, and only if those results are for sensitive data discovery jobs that are created by the account and in the Region specified in the condition.

To allow Macie to perform the specified actions for additional accounts, add ARNs for each additional account to this condition. For example:

```
"aws:SourceArn": [  
  "arn:aws:macie2:us-east-1:111122223333:export-configuration:*",  
  "arn:aws:macie2:us-east-1:111122223333:classification-job/*",  
  "arn:aws:macie2:us-east-1:444455556666:export-configuration:*",  
  "arn:aws:macie2:us-east-1:444455556666:classification-job/*"  
]
```

The accounts specified by the `aws:SourceAccount` and `aws:SourceArn` conditions should match.

Both conditions help prevent Macie from being used as a [confused deputy](#) during transactions with Amazon S3. Although we don't recommend it, you can remove these conditions from the bucket policy.

8. When you finish updating the bucket policy, choose **Save changes**.

Important

If you change the bucket path after you configure the repository settings in Macie, you have to update the bucket policy. Otherwise, Macie won't be allowed to add discovery results to the bucket.

You can now configure the repository settings in Macie.

To configure the repository settings in Macie

1. Open the Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, under **Settings**, choose **Discovery results**.
3. Under **Repository for sensitive data discovery results**, choose **Existing bucket**.
4. For **Choose a bucket**, select the bucket that you want to store your discovery results in.
5. (Optional) To specify a prefix to use in the path to a location in the bucket, expand the **Advanced** section. Then, for **Data discovery result prefix**, enter the prefix to use.

When you enter a value, Macie updates the example below the box to show the path to the bucket location where it will store your discovery results.

6. Under **Encryption settings**, specify the AWS KMS key that you want to use to encrypt the results:
 - To use a key from your own account, choose **Select a key from your account**. Then, in the **AWS KMS key** list, choose the key to use. The list displays customer managed, symmetric encryption KMS keys for your account.
 - To use a key that another account owns, choose **Enter the ARN of a key from another account**. Then, in the **AWS KMS key ARN** box, enter the ARN of the key to use—for example, **arn:aws:kms:us-east-1:111122223333:key/1234abcd-12ab-34cd-56ef-1234567890ab**.
7. When you finish entering the settings, choose **Save**. Macie tests the settings to verify that they're correct. If any settings are incorrect, Macie displays an error message to help you address the issue.

After you save the repository settings, Macie adds existing discovery results for the preceding 90 days to the repository. Macie also starts adding new discovery results to the repository.

Analyzing Amazon Macie findings

Amazon Macie generates findings when it detects potential policy violations or issues with the security or privacy of your Amazon Simple Storage Service (Amazon S3) buckets or it discovers sensitive data in S3 objects. A *finding* is a detailed report of a potential issue or sensitive data that Macie found. Each finding provides a severity rating, information about the affected resource, and additional details, such as when and how Macie found the issue or data. Macie stores your policy and sensitive data findings for 90 days.

You can review, analyze, and manage findings in the following ways.

Amazon Macie console

The **Findings** pages on the Amazon Macie console list your findings and provide detailed information for individual findings. These pages also provide options for grouping, filtering, and sorting findings, and for creating and managing [suppression rules \(p. 203\)](#). Suppression rules can help you streamline your analysis of findings.

Amazon Macie API

With the Amazon Macie API, you can query and retrieve findings data by sending HTTPS requests directly to Macie or by using the AWS Command Line Interface (AWS CLI) or another AWS tool or SDK of your choice. To query the data, you send a request to the Amazon Macie API and use supported parameters to specify which findings you want to retrieve. After you submit your query, Macie returns the results in a JSON response. You can then pass the results to another service or application for deeper analysis, long-term storage, or reporting. For more information, see the [Amazon Macie API Reference](#).

Amazon EventBridge

To further support integration with other services and systems, such as monitoring or event management systems, Macie publishes findings to Amazon EventBridge as events. EventBridge, formerly Amazon CloudWatch Events, is a serverless event bus service that can deliver a stream of real-time data from your own applications, software as a service (SaaS) applications, and AWS services such as Macie. It can route that data to targets such as AWS Lambda functions, Amazon Simple Notification Service topics, and Amazon Kinesis streams. To learn about this service, see the [Amazon EventBridge User Guide](#).

Macie automatically publishes events to EventBridge for new findings. It also publishes events automatically for subsequent occurrences of existing policy findings. Because the notifications are structured as EventBridge events, you can more easily monitor, analyze, and act upon findings by using other services and tools. For example, you might use EventBridge to automatically send specific types of new findings to an AWS Lambda function that, in turn, processes and sends the data to your security incident and event management (SIEM) system. In addition to automated processing, use of EventBridge events helps ensure longer-term retention of findings data. To learn about using EventBridge events for findings, see [Monitoring and processing findings \(p. 215\)](#).

AWS Security Hub

For additional, broader analysis of your organization's security posture, you can also review and analyze findings by using AWS Security Hub. Security Hub is a service that provides you with a comprehensive view of your security state across your AWS environment and helps you check your environment against security industry standards and best practices. To learn about this service, see the [AWS Security Hub User Guide](#). To learn about how Macie publishes findings to Security Hub, see [Monitoring and processing findings \(p. 215\)](#).

In addition to findings, Macie creates sensitive data discovery results for S3 objects that you configure sensitive data discovery jobs to analyze. A *sensitive data discovery result* is a record that logs details

about the analysis of an object. This includes objects that don't contain sensitive data, and therefore don't produce findings, and objects that Macie can't analyze due to issues such as permission settings for a bucket. To learn more about sensitive data discovery results, see [Reviewing job statistics and results \(p. 113\)](#).

You can't access sensitive data discovery results directly on the Amazon Macie console or through the Amazon Macie API. Instead, you configure Macie to store the results in an S3 bucket. You can then optionally access and query the results in that bucket. To learn how to configure Macie to store the results, see [Storing and retaining sensitive data discovery results \(p. 130\)](#). For samples of Amazon Athena queries that you can use to analyze the results, visit the [Amazon Macie Results Analytics repository](#) on GitHub. This repository also provides step-by-step instructions for configuring Athena to retrieve and decrypt sensitive data discovery results, and scripts for creating tables for the results.

Topics

- [Types of Amazon Macie findings \(p. 141\)](#)
- [Working with sample findings in Amazon Macie \(p. 143\)](#)
- [Reviewing findings on the Amazon Macie console \(p. 146\)](#)
- [Filtering Amazon Macie findings \(p. 148\)](#)
- [Investigating sensitive data with Amazon Macie findings \(p. 183\)](#)
- [Suppressing Amazon Macie findings \(p. 203\)](#)
- [Severity scoring for Amazon Macie findings \(p. 209\)](#)

Types of Amazon Macie findings

Amazon Macie generates two categories of findings: *policy findings* and *sensitive data findings*. A *policy finding* is a detailed report of a potential policy violation or issue with the security or privacy of an Amazon Simple Storage Service (Amazon S3) bucket. Macie generates these findings as part of its ongoing monitoring activities for your Amazon S3 data. A *sensitive data finding* is a detailed report of sensitive data in an S3 object. Macie generates these findings when it discovers sensitive data in S3 objects that you configure a sensitive data discovery job to analyze. Each category includes specific types of findings.

Topics

- [Types of policy findings \(p. 141\)](#)
- [Types of sensitive data findings \(p. 142\)](#)

Tip

To explore and learn about the different categories and types of findings that Macie can generate, [create sample findings \(p. 143\)](#). Sample findings use example data and placeholder values to demonstrate the kinds of information that each type of finding might contain.

Types of policy findings

Macie generates policy findings when the policies or settings for an S3 bucket are changed in a way that reduces the security or privacy of the bucket and the bucket's objects. Macie does this only if the change occurs after you enable Macie for your AWS account.

For example, if default encryption is disabled for a bucket after you enable Macie, Macie generates a **Policy:IAMUser/S3BucketEncryptionDisabled** finding for the bucket. However, if default encryption was disabled for a bucket when you enabled Macie and default encryption continues to be disabled, Macie doesn't generate a **Policy:IAMUser/S3BucketEncryptionDisabled** finding for the bucket.

Macie can generate the following types of policy findings for an S3 bucket.

Policy:IAMUser/S3BlockPublicAccessDisabled

Block public access settings were disabled for the bucket. Access to the bucket is controlled only by access control lists (ACLs) and bucket policies.

To learn about block public access settings for S3 buckets, see [Blocking public access to your Amazon S3 storage](#) in the *Amazon Simple Storage Service User Guide*.

Policy:IAMUser/S3BucketEncryptionDisabled

Default encryption was disabled for the bucket. By default, Amazon S3 won't automatically encrypt new objects when they're added to the bucket.

To learn about default encryption settings for S3 buckets, see [Setting default server-side encryption behavior for S3 buckets](#) in the *Amazon Simple Storage Service User Guide*.

Policy:IAMUser/S3BucketPublic

An ACL or bucket policy for the bucket was changed to allow access by anonymous users or by all authenticated AWS Identity and Access Management (IAM) users or roles.

To learn about ACLs and bucket policies for S3 buckets, see [Identity and access management in Amazon S3](#) in the *Amazon Simple Storage Service User Guide*.

Policy:IAMUser/S3BucketReplicatedExternally

Data replication was enabled and configured to replicate objects from the bucket to an AWS account that isn't part of your organization. An *organization* is a set of Macie accounts that are centrally managed as a group of related accounts through AWS Organizations or by Macie invitation.

To learn about replication settings for S3 buckets, see [Replicating objects](#) in the *Amazon Simple Storage Service User Guide*.

Policy:IAMUser/S3BucketSharedExternally

An ACL or bucket policy for the bucket was changed to allow the bucket to be shared with an AWS account that isn't part of your organization. An *organization* is a set of Macie accounts that are centrally managed as a group of related accounts through AWS Organizations or by Macie invitation.

To learn about ACLs and bucket policies for S3 buckets, see [Identity and access management in Amazon S3](#) in the *Amazon Simple Storage Service User Guide*.

Types of sensitive data findings

Macie generates sensitive data findings when it discovers sensitive data in S3 objects that you configure a sensitive data discovery job to analyze. Macie can generate the following types of sensitive data findings for an object.

SensitiveData:S3Object/Credentials

The object contains credentials data, such as private keys or AWS secret access keys.

SensitiveData:S3Object/CustomIdentifier

The object contains text that matches the detection criteria of one or more custom data identifiers. The object might contain more than one type of sensitive data.

SensitiveData:S3Object/Financial

The object contains financial information, such as credit card numbers or bank account numbers.

SensitiveData:S3Object/Multiple

The object contains more than one category of sensitive data—any combination of credentials data, financial information, personal information, or text that matches the detection criteria of one or more custom data identifiers.

SensitiveData:S3Object/Personal

The object contains personally identifiable information (such as full names or mailing addresses), personal health information (such as health insurance or medical identification numbers), or a combination of the two.

For detailed information about the types of sensitive data that Macie can detect, see [Using managed data identifiers \(p. 45\)](#). For information about the types of S3 objects that Macie can analyze, see [Discovering sensitive data \(p. 44\)](#).

Working with sample findings in Amazon Macie

To explore and learn about the different [types of findings \(p. 141\)](#) that Amazon Macie can generate, you can create sample findings. Sample findings use example data and placeholder values to demonstrate the kinds of information that each type of finding might contain.

For example, the **Policy:IAMUser/S3BucketPublic** sample finding contains details about a fictitious Amazon Simple Storage Service (Amazon S3) bucket. The finding's details include example data about an actor and action that changed the access control list (ACL) for the bucket and made the bucket publicly accessible. Similarly, the **SensitiveData:S3Object/Multiple** sample finding contains details about a fictitious Microsoft Excel workbook. The finding's details include example data about the types and location of sensitive data in the workbook.

In addition to familiarizing yourself with the information that different types of findings might contain, you can use sample findings to test integration with other applications, services, and systems. Depending on the [suppression rules \(p. 203\)](#) for your account, Macie can publish sample findings to Amazon EventBridge as events. By using the example data in sample findings, you can develop and test automated solutions for monitoring and processing these events. Depending on the [publication settings \(p. 227\)](#) for your account, Macie can also publish sample findings to AWS Security Hub. This means that you can also use sample findings to develop and test solutions for monitoring and processing Macie findings in Security Hub. For information about publishing findings to these services, see [Monitoring and processing findings \(p. 215\)](#).

Topics

- [Creating sample findings \(p. 143\)](#)
- [Reviewing sample findings \(p. 144\)](#)
- [Suppressing sample findings \(p. 146\)](#)

Creating sample findings

You can create sample findings by using the Amazon Macie console or the Amazon Macie API. If you use the console, Macie automatically generates one sample finding for each type of finding that Macie supports. If you use the API, you can create a sample for each type or only certain types that you specify.

Console

Follow these steps to create sample findings by using the Amazon Macie console.

To create sample findings

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, choose **Settings**.
3. Under **Sample findings**, choose **Generate sample findings**.

API

To create sample findings programmatically, use the [CreateSampleFindings](#) operation of the Amazon Macie API. When you submit your request, optionally use the `findingTypes` parameter to specify only certain types of sample findings to create. To automatically create samples of all types, don't include this parameter in your request.

To create sample findings by using the [AWS Command Line Interface \(AWS CLI\)](#), run the `create-sample-findings` command. To automatically create samples of all types of findings, don't include the `finding-types` parameter. To create samples of only certain types of findings, include this parameter and specify the types of sample findings to create. For example:

```
C:\> aws macie2 create-sample-findings --finding-types "SensitiveData:S3Object/  
Multiple" "Policy:IAMUser/S3BucketPublic"
```

Where *SensitiveData:S3Object/Multiple* is a type of sensitive data finding to create and *Policy:IAMUser/S3BucketPublic* is a type of policy finding to create.

If the command runs successfully, Macie returns an empty response.

Reviewing sample findings

To help you identify sample findings that you created, Macie sets the value for the **Sample** field of each sample finding to *True*. In addition, the name of the affected S3 bucket is the same for all sample findings: *macie-sample-finding-bucket*. If you review sample findings by using the **Findings** page on the Amazon Macie console, Macie also displays the **[SAMPLE]** prefix in the **Finding type** field for each sample finding.

Console

Follow these steps to review sample findings by using the Amazon Macie console.

To review sample findings

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, choose **Findings**.
3. On the **Findings** page, do any of the following:
 - In the **Finding type** column, locate findings whose type begins with **[SAMPLE]**, as shown in the following image.

<input type="checkbox"/>	Severity ▾	Finding type
<input type="checkbox"/>	High	[SAMPLE] Policy:IAMUser/S3BlockPublicAccessDisabled
<input type="checkbox"/>	Medium	[SAMPLE] Policy:IAMUser/S3BucketEncryptionDisabled
<input type="checkbox"/>	High	[SAMPLE] Policy:IAMUser/S3BucketPublic
<input type="checkbox"/>	High	[SAMPLE] Policy:IAMUser/S3BucketReplicatedExternally
<input type="checkbox"/>	High	[SAMPLE] Policy:IAMUser/S3BucketSharedExternally
<input type="checkbox"/>	High	[SAMPLE] SensitiveData:S3Object/Credentials
<input type="checkbox"/>	High	[SAMPLE] SensitiveData:S3Object/CustomIdentifier
<input type="checkbox"/>	High	[SAMPLE] SensitiveData:S3Object/Financial
<input type="checkbox"/>	High	[SAMPLE] SensitiveData:S3Object/Multiple
<input type="checkbox"/>	High	[SAMPLE] SensitiveData:S3Object/Multiple
<input type="checkbox"/>	Low	[SAMPLE] SensitiveData:S3Object/Personal

- By using the filter bar above the table, filter the table to display only sample findings. To do this, place your cursor in the filter bar. In the list of fields that appears, choose **Sample**. Then choose **True**, and then choose **Apply**. This adds the following filter condition to the table:



4. To review the details of a specific sample finding, choose any field other than the check box for the finding. The details panel displays information for the finding.

You can also download and save the details of one or more sample findings as a JSON file. To do this, select the check box for each sample finding that you want to download and save. Then choose **Export (JSON)** from the **Actions** menu at the top of the **Findings** page. In the window that appears, choose **Download**. For detailed descriptions of the JSON fields that a finding can include, see the [Finding table](#) in the *Amazon Macie API Reference*.

API

To review sample findings programmatically, first use the [ListFindings](#) operation of the Amazon Macie API to retrieve the unique identifier (`findingId`) for each sample finding that you created. Then use the [GetFindings](#) operation to retrieve the details of those findings.

When you submit the **ListFindings** request, you can specify filter criteria to include only sample findings in the results. To do this, add a filter condition where the value for the `sample` field is `true`. If you're using the AWS CLI, run the `list-findings` command and use the `finding-criteria` parameter to specify the filter condition. For example:

```
C:\> aws macie2 list-findings --finding-criteria={"criterion":{"sample":{"eq":{"true"}}}}
```

If your request succeeds, Macie returns a `findingIds` array. The array lists the unique identifier for each sample finding for your account in the current AWS Region.

To then retrieve the details of the sample findings, specify these unique identifiers in a **GetFindings** request or, for the AWS CLI, when you run the [get-findings](#) command.

Suppressing sample findings

Like other findings, Macie stores sample findings for 90 days. After you finish reviewing and experimenting with the samples, you can optionally archive them by [creating a suppression rule](#) (p. 203). If you do this, the sample findings stop appearing by default on the console and their status changes to *archived*.

To archive sample findings by using the Amazon Macie console, configure the rule to archive findings where the value for the **Sample** field is **True**. To archive sample findings by using the Amazon Macie API, configure the rule to archive findings where the value for the `sample` field is `true`.

Reviewing findings on the Amazon Macie console

Amazon Macie monitors your AWS environment and generates policy findings when it detects potential policy violations or issues that affect the security or privacy of your Amazon Simple Storage Service (Amazon S3) buckets. Macie generates sensitive data findings when it discovers sensitive data in S3 objects that you configure a sensitive data discovery job to analyze. Macie stores your policy and sensitive data findings for 90 days.

By using the Amazon Macie console, you can review and analyze findings, and access the details of individual findings. Each finding provides a severity rating, information about the affected resource, and additional details, such as the exact nature of the issue, and when and how Macie found the issue. To help you streamline your analysis, the console offers several options for building custom views of findings.

Use predefined groupings

Use specific pages to review findings that are grouped by criteria such as affected S3 bucket, finding type, or sensitive data discovery job. With these pages, you can review aggregated statistics for each group, such as the count of findings by severity. You can also drill down to review the details of individual findings in a group, and you can apply filters to refine your analysis.

For example, if you group all findings by S3 bucket and note that a particular S3 bucket has a policy violation, you can quickly determine whether the bucket also contains sensitive data. To do this, choose **By bucket** in the navigation pane (under **Findings**), and then choose the bucket. In the details panel that appears, the **Findings by type** section lists the types of findings that apply to the bucket. To investigate a specific type, choose the number for the type. Macie displays a table of all the findings that match the selected type and apply to the bucket. To refine the results, filter the table.

Create and apply filters


Use specific finding attributes to include or exclude certain findings from a **Findings** table. A *finding attribute* is a field that stores specific data for a finding, such as finding type, severity, or the name of the affected S3 bucket. If you filter a table, you can more easily identify findings that have specific characteristics. Then you can drill down to review the details of those findings.

For example, to review all of your policy findings, add filter criteria for the **Category** field. To refine your view and include only a specific type of policy finding, add filter criteria for the **Finding type** field. To then review the details of a particular finding, choose the finding. The details panel displays information for the finding.

You can also sort findings in ascending or descending order by certain fields. To do this, click the column heading for the field. To change the sort order, click the column heading again.

To review findings on the console

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, choose **Findings**. The **Findings** page displays findings that Macie created or updated for your account in the current AWS Region during the past 90 days. By default, this doesn't include findings that were suppressed by a [suppression rule](#) (p. 203).
3. (Optional) To pivot on and review findings by a predefined logical group, choose **By bucket**, **By type**, or **By job** in the navigation pane (under **Findings**). Then choose an item in the table. In the details panel, choose the link for the field to pivot on.
4. (Optional) To filter the findings by specific criteria, use the filter bar above the table:



- To display findings that were suppressed by a suppression rule, choose **Current** in the filter bar. Then choose **Archived** to display only suppressed findings, or choose **All** to display both current and suppressed findings.
- To display only those findings that have a specific attribute, place your cursor in the filter bar and add a filter condition for the attribute. To further refine the results, add conditions for additional attributes. For information about using filter conditions, see [Creating and applying filters to findings](#) (p. 154).
- To remove a filter condition, choose the remove condition icon (⊗) in the filter box.

To save your filter settings, choose **Save rule** in the filter bar. Then enter a name and, optionally, a description for the settings. When you finish, choose **Save**.

5. (Optional) To sort the findings by a specific field, click the column heading for the field. To change the sort order, click the column heading again.
6. To review the details of a specific finding, choose any field other than the check box for the finding. The details panel displays information for the finding.

Tip

You can use the details panel to pivot and drill down on certain fields. To show findings that have the same value for a field, choose ⊕ in the field. Or choose ⊖ to show findings that have other values for the field.

For a sensitive data finding, you can also use the details panel to investigate sensitive data that Macie found in the affected S3 object:

- To locate occurrences of sensitive data, choose a link in an **Occurrences** field. Macie displays information (in JSON format) about where Macie found the data. To learn more, see [Locating sensitive data](#) (p. 184).
- To retrieve samples of sensitive data, choose **Review** in the **Reveal samples** field. To learn more, see [Retrieving sensitive data samples](#) (p. 186).
- To navigate to the corresponding sensitive data discovery result, choose the link in the **Detailed result location** field. Macie opens the Amazon S3 console and displays the file

or folder that contains the discovery result. To learn more, see [Reviewing job statistics and results \(p. 113\)](#).

You can also download and save the details of one or more findings as a JSON file. To do this, select the check box for each finding that you want to download and save. Then choose **Export (JSON)** from the **Actions** menu at the top of the **Findings** page. In the window that appears, choose **Download**. For detailed descriptions of the JSON fields that a finding can include, see the [Finding table](#) in the *Amazon Macie API Reference*.

Filtering Amazon Macie findings

To perform targeted analysis and to analyze findings more efficiently, you can filter Amazon Macie findings. With filters, you build custom views and queries for findings, which can help you identify and focus on findings that have specific characteristics. Use the Amazon Macie console to filter findings, or submit queries programmatically using the Amazon Macie API.

When you create a filter, you use specific attributes of findings to define criteria for including or excluding findings from a view or from query results. A *finding attribute* is a field that stores specific data for a finding, such as severity, type, or the name of the S3 bucket that a finding applies to.

In Macie, a filter consists of one or more conditions. Each condition, also referred to as a *criterion*, consists of three parts:

- An attribute-based field, such as **Severity** or **Finding type**.
- An operator, such as *equals* or *not equals*.
- One or more values. The type and number of values depends on the field and operator that you choose.

If you create a filter that you want to use again, you can save it as a *filter rule*. A *filter rule* is a set of filter criteria that you create and save to reapply when you review findings on the Amazon Macie console.

You can also save a filter as a *suppression rule*. A *suppression rule* is a set of filter criteria that you create and save to automatically archive findings that match the criteria of the rule. To learn about suppression rules, see [Suppressing findings \(p. 203\)](#).

Topics

- [Fundamentals of filtering findings \(p. 148\)](#)
- [Creating and applying filters to findings \(p. 154\)](#)
- [Creating and managing filter rules for findings \(p. 160\)](#)
- [Fields for filtering findings \(p. 165\)](#)

Fundamentals of filtering findings

When you create a filter, keep the following features and guidelines in mind. Also note that filtered results are limited to the preceding 90 days and the current AWS Region. Amazon Macie stores your findings for only 90 days in each AWS Region.

Topics

- [Using multiple conditions in a filter \(p. 149\)](#)
- [Specifying values for fields \(p. 149\)](#)
- [Specifying multiple values for a field \(p. 151\)](#)

- [Using operators in conditions \(p. 151\)](#)

Using multiple conditions in a filter

A filter can include one or more conditions. Each condition, also referred to as a *criterion*, consists of three parts:

- An attribute-based field, such as **Severity** or **Finding type**. For a list of fields that you can use, see [Fields for filtering findings \(p. 165\)](#).
- An operator, such as *equals* or *not equals*. For a list of operators that you can use, see [Using operators in conditions \(p. 151\)](#).
- One or more values. The type and number of values depends on the field and operator that you choose.

If a filter contains multiple conditions, Macie uses AND logic to join the conditions and evaluate the filter criteria. This means that a finding matches the filter criteria only if it matches *all* the conditions in the filter.

For example, if you add a condition to include only high-severity findings and add another condition to include only sensitive data findings, Macie returns all high-severity, sensitive data findings. In other words, Macie excludes all policy findings and all medium-severity and low-severity sensitive data findings.

You can use a field only once in a filter. However, you can specify multiple values for many fields.

For example, if a condition uses the **Severity** field to include only high-severity findings, you can't use the **Severity** field in another condition to include medium-severity or low-severity findings. Instead, specify multiple values for the existing condition, or use a different operator for the existing condition. For example, to include all medium-severity and high-severity findings, add a **Severity equals Medium, High** condition or add a **Severity not equals Low** condition.

Specifying values for fields

When you specify a value for a field, the value has to conform to the underlying data type for the field. Depending on the field, you can specify one of the following types of values.

Array of text (strings)

Specifies a list of text (string) values for a field. Each string correlates to a predefined or existing value for a field—for example, *High* for the **Severity** field, *SensitiveData:S3Object/Financial* for the **Finding type** field, or the name of an S3 bucket for the **S3 bucket name** field.

If you use an array, note the following:

- Values are case sensitive.
- You can't specify partial values or use wildcard characters in values. You have to specify a complete, valid value for the field.

For example, to filter findings for an S3 bucket named *my-S3-bucket*, enter **my-S3-bucket** as the value for the **S3 bucket name** field. If you enter any other value, such as **my-s3-bucket** or **my-S3**, Macie won't return findings for the bucket.

For a list of valid values for each field, see [Fields for filtering findings \(p. 165\)](#).

You can specify as many as 50 values in an array. How you specify the values depends on whether you use the Amazon Macie console or the Amazon Macie API, as discussed in [Specifying multiple values for a field \(p. 151\)](#).

Boolean

Specifies one of two mutually exclusive values for a field.

If you use the Amazon Macie console to specify this type of value, the console provides a list of values to choose from. If you use the Amazon Macie API, specify `true` or `false` for the value.

Date/Time (and time ranges)

Specifies an absolute date and time for a field. If you specify this type of value, you have to specify both a date and time.

On the Amazon Macie console, date and time values are in your local time zone and use 24-hour notation. In all other contexts, these values are in Coordinated Universal Time (UTC) and extended ISO 8601 format—for example `2020-09-01T14:31:13Z` for 2:31:13 PM UTC September 1, 2020.

If a field stores a date/time value, you can use the field to define a fixed or relative time range. For example, you can include only those findings that were created between two specific dates and times, or only those findings that were created before or after a specific date and time. How you define a time range depends on whether you use the Amazon Macie console or the Amazon Macie API:

- On the console, use a date picker or enter text directly in the **From** and **To** boxes.
- With the API, define a fixed time range by adding a condition that specifies the first date and time in the range, and add another condition that specifies the last date and time in the range. If you do this, Macie uses AND logic to join the conditions. To define a relative time range, add one condition that specifies the first or last date and time in the range. Specify the values as Unix timestamps in milliseconds—for example, `1604616572653` for 22:49:32 UTC November 5, 2020.

On the console, time ranges are inclusive. With the API, time ranges can be inclusive or exclusive, depending on the operator that you choose.

Number (and numeric ranges)

Specifies a long integer for a field.

If a field stores a numeric value, you can use the field to define a fixed or relative numeric range. For example, you can include only those findings that report 50-90 occurrences of sensitive data in an S3 object. How you define a numeric range depends on whether you use the Amazon Macie console or the Amazon Macie API:

- On the console, use the **From** and **To** boxes to enter the lowest and highest numbers in the range, respectively.
- With the API, define a fixed numeric range by adding a condition that specifies the lowest number in the range, and add another condition that specifies the highest number in the range. If you do this, Macie uses AND logic to join the conditions. To define a relative numeric range, add one condition that specifies the lowest or highest number in the range.

On the console, numeric ranges are inclusive. With the API, numeric ranges can be inclusive or exclusive, depending on the operator that you choose.

Text (string)

Specifies a single text (string) value for a field. The string correlates to a predefined or existing value for a field—for example, *High* for the **Severity** field, the name of an S3 bucket for the **S3 bucket name** field, or the unique identifier for a sensitive data discovery job for the **Job ID** field.

If you specify a single text string, note the following:

- Values are case sensitive.
- You can't use partial values or use wildcard characters in values. You have to specify a complete, valid value for the field.

For example, to filter findings for an S3 bucket named *my-S3-bucket*, enter **my-S3-bucket** as the value for the **S3 bucket name** field. If you enter any other value, such as **my-s3-bucket** or **my-S3**, Macie won't return findings for the bucket.

For a list of valid values for each field, see [Fields for filtering findings \(p. 165\)](#).

Specifying multiple values for a field

With certain fields and operators, you can specify multiple values for a field. If you do this, Macie uses OR logic to join the values and evaluate the filter criteria. This means that a finding matches the criteria if it has *any* of the values for the field.

For example, if you add a condition to include findings where the value for the **Finding type** field equals *SensitiveData:S3Object/Financial*, *SensitiveData:S3Object/Personal*, Macie returns sensitive data findings for S3 objects that contain only financial data, and S3 objects that contain only personal information. In other words, Macie excludes all policy findings. Macie also excludes all sensitive data findings for objects that contain other types of sensitive data or multiple types of sensitive data.

The exception is conditions that use the *eqExactMatch* operator. For this operator, Macie uses AND logic to join the values and evaluate the filter criteria. This means that a finding matches the criteria only if it has *all* the values for the field and *only* those values for the field. To learn more about this operator, see [Using operators in conditions \(p. 151\)](#).

How you specify multiple values for a field depends on whether you use the Amazon Macie API or the Amazon Macie console. With the API, you use an array that lists the values.

On the console, you typically choose the values from a list. However, for some fields, you have to add a distinct condition for each value. For example, to include findings for data that Macie detected using certain custom data identifiers, do the following:

1. Place your cursor in the filter bar, choose the **Custom data identifier detection name** field, enter the name of a custom data identifier, and then choose **Apply**.
2. Repeat the preceding step for each additional custom data identifier that you want to specify for the filter.

For a list of fields that you need to do this for, see [Fields for filtering findings \(p. 165\)](#).

Using operators in conditions

You can use the following types of operators in individual conditions.

Equals (eq)

Matches (=) any value specified for the field. You can use the *equals* operator with the following types of values: array of text (strings), Boolean, date/time, number, and text (string).

For many fields, you can use this operator and specify as many as 50 values for the field. If you do this, Macie uses OR logic to join the values. This means that a finding matches the criteria if it has *any* of the values specified for the field.

For example:

- To include findings that report occurrences of financial information, personal information, or both financial and personal information, add a condition that uses the **Sensitive data category** field and this operator, and specify *Financial information* and *Personal information* as the values for the field.
- To include findings that report occurrences of credit card numbers, mailing addresses, or both credit card numbers and mailing addresses, add a condition for the **Sensitive data detection type**

field, use this operator, and specify `CREDIT_CARD_NUMBER` and `ADDRESS` as the values for the field.

If you use the Amazon Macie API to define a condition that uses this operator with a date/time value, specify the value as a Unix timestamp in milliseconds—for example, `1604616572653` for 22:49:32 UTC November 5, 2020.

Equals exact match (eqExactMatch)

Exclusively matches all the values specified for the field. You can use the *equals exact match* operator with a select set of fields.

If you use this operator and specify multiple values for a field, Macie uses AND logic to join the values. This means that a finding matches the criteria only if it has *all* the values specified for the field and *only* those values for the field. You can specify as many as 50 values for the field.

For example:

- To include findings that report occurrences of credit card numbers and no other type of sensitive data, add a condition for the **Sensitive data detection type** field, use this operator, and specify `CREDIT_CARD_NUMBER` as the only value for the field.
- To include findings that report occurrences of both credit card numbers and mailing addresses (and no other types of sensitive data), add a condition for the **Sensitive data detection type** field, use this operator, and specify `CREDIT_CARD_NUMBER` and `ADDRESS` as the values for the field.

Because Macie uses AND logic to join the values for a field, you can't use this operator in combination with any other operators for the same field. In other words, if you use the *equals exact match* operator with a field in one condition, you have to use it in all other conditions that use the same field.

Like other operators, you can use the *equals exact match* operator in more than one condition in a filter. If you do this, Macie uses AND logic to join the conditions and evaluate the filter. This means that a finding matches the filter criteria only if it has *all* the values specified by *all* the conditions in the filter.

For example, to include findings that were created after a certain time, report occurrences of credit card numbers, and don't report any other type of sensitive data, do the following:

1. Add a condition that uses the **Created at** field, uses the *greater than* operator, and specifies the starting date and time for the filter.
2. Add another condition that uses the **Sensitive data detection type** field, uses the *equals exact match* operator, and specifies `CREDIT_CARD_NUMBER` as the only value for the field.

You can use the *equals exact match* operator with the following fields:

- Custom data identifier detection ARN (`customDataIdentifiers.detections.arn`)
- Custom data identifier detection name (`customDataIdentifiers.detections.name`)
- S3 bucket tag key (`resourcesAffected.s3Bucket.tags.key`)
- S3 bucket tag value (`resourcesAffected.s3Bucket.tags.value`)
- S3 object tag key (`resourcesAffected.s3Object.tags.key`)
- S3 object tag value (`resourcesAffected.s3Object.tags.value`)
- Sensitive data detection type (`sensitiveData.detections.type`)
- Sensitive data category (`sensitiveData.category`)

In the preceding list, the parenthetical name uses dot notation to indicate the name of the field in JSON representations of findings and the Amazon Macie API.

Greater than (gt)

Is greater than (>) the value specified for the field. You can use the *greater than* operator with number and date/time values.

For example, to include only those findings that report more than 90 occurrences of sensitive data in an S3 object, add a condition that uses the **Sensitive data total count** field and this operator, and specify 90 as the value for the field. To do this on the Amazon Macie console, enter **91** in the **From** box, don't enter a value in the **To** box, and then choose **Apply**. Numeric and time-based comparisons are inclusive on the console.

If you use the Amazon Macie API to define a time range that uses this operator, you have to specify the date/time values as Unix timestamps in milliseconds—for example, 1604616572653 for 22:49:32 UTC November 5, 2020.

Greater than or equal to (gte)

Is greater than or equal to (\geq) the value specified for the field. You can use the *greater than or equal to* operator with number and date/time values.

For example, to include only those findings that report 90 or more occurrences of sensitive data in an S3 object, add a condition that uses the **Sensitive data total count** field and this operator, and specify 90 as the value for the field. To do this on the Amazon Macie console, enter **90** in the **From** box, don't enter a value in the **To** box, and then choose **Apply**.

If you use the Amazon Macie API to define a time range that uses this operator, you have to specify the date/time values as Unix timestamps in milliseconds—for example, 1604616572653 for 22:49:32 UTC November 5, 2020.

Less than (lt)

Is less than ($<$) the value specified for the field. You can use the *less than* operator with number and date/time values.

For example, to include only those findings that report fewer than 90 occurrences of sensitive data in an S3 object, add a condition that uses the **Sensitive data total count** field and this operator, and specify 90 as the value for the field. To do this on the Amazon Macie console, enter **89** in the **To** box, don't enter a value in the **From** box, and then choose **Apply**. Numeric and time-based comparisons are inclusive on the console.

If you use the Amazon Macie API to define a time range that uses this operator, you have to specify the date/time values as Unix timestamps in milliseconds—for example, 1604616572653 for 22:49:32 UTC November 5, 2020.

Less than or equal to (lte)

Is less than or equal to (\leq) the value specified for the field. You can use the *less than or equal to* operator with number and date/time values.

For example, to include only those findings that report 90 or fewer occurrences of sensitive data in an S3 object, add a condition that uses the **Sensitive data total count** field and this operator, and specify 90 as the value for the field. To do this on the Amazon Macie console, enter **90** in the **To** box, don't enter a value in the **From** box, and then choose **Apply**.

If you use the Amazon Macie API to define a time range that uses this operator, you have to specify the date/time values as Unix timestamps in milliseconds—for example, 1604616572653 for 22:49:32 UTC November 5, 2020.

Not equals (neq)

Doesn't match (\neq) any value specified for the field. You can use the *not equals* operator with the following types of values: array of text (strings), Boolean, date/time, number, and text (string).

For many fields, you can use this operator and specify as many as 50 values for the field. If you do this, Macie uses OR logic to join the values. This means that a finding matches the criteria if it doesn't have *any* of the values specified for the field.

For example:

- To exclude findings that report occurrences of financial information, personal information, or both financial and personal information, add a condition that uses the **Sensitive data category** field and this operator, and specify *Financial information* and *Personal information* as the values for the field.
- To exclude findings that report occurrences of credit card numbers, add a condition for the **Sensitive data detection type** field, use this operator, and specify *CREDIT_CARD_NUMBER* as the value for the field.
- To exclude findings that report occurrences of credit card numbers, mailing addresses, or both credit card numbers and mailing addresses, add a condition for the **Sensitive data detection type** field, use this operator, and specify *CREDIT_CARD_NUMBER* and *ADDRESS* as the values for the field.

If you use the Amazon Macie API to define a condition that uses this operator with a date/time value, specify the value as a Unix timestamp in milliseconds—for example, 1604616572653 for 22:49:32 UTC November 5, 2020.

Creating and applying filters to findings

To identify and focus on findings that have specific characteristics, you can filter findings on the Amazon Macie console and in queries that you submit programmatically using the Amazon Macie API. When you create a filter, you use specific attributes of findings to define criteria for including or excluding findings from a view or from query results. A *finding attribute* is a field that stores specific data for a finding, such as severity, type, or the name of the S3 bucket that a finding applies to.

In Macie, a filter consists of one or more conditions. Each condition, also referred to as a *criterion*, consists of three parts:

- An attribute-based field, such as **Severity** or **Finding type**.
- An operator, such as *equals* or *not equals*.
- One or more values. The type and number of values depends on the field and operator that you choose.

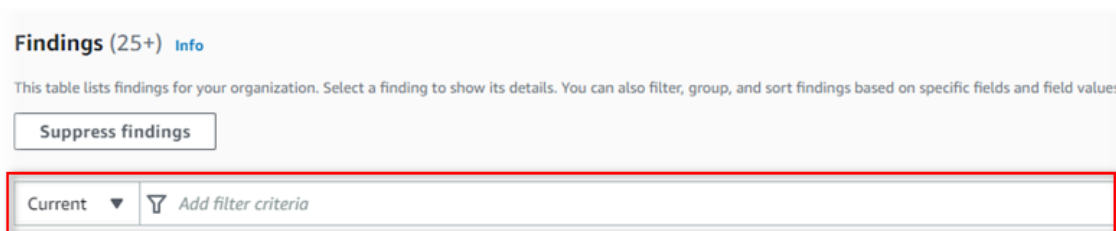
How you define and apply filter conditions depends on whether you use the Amazon Macie console or the Amazon Macie API.

Topics

- [Filtering findings on the Amazon Macie console \(p. 154\)](#)
- [Filtering findings programmatically with the Amazon Macie API \(p. 156\)](#)

Filtering findings on the Amazon Macie console

If you use the Amazon Macie console to filter findings, Macie provides options to help you choose fields, operators, and values for individual conditions. You access these options by using the filter bar on **Findings** pages, as shown in the following image.




When you place your cursor in the filter bar, Macie displays a list of fields that you can use in filter conditions. The fields are organized by logical category. For example, the **Common fields** category includes fields that apply to any type of finding, and the **Classification fields** category includes fields that apply only to sensitive data findings. The fields are sorted alphabetically within each category.

To add a condition, start by choosing a field from the list. To find a field, browse the complete list, or enter part of the field's name to narrow the list of fields.

Depending on the field that you choose, Macie displays different options. The options reflect the type and nature of the field that you choose. For example, if you choose the **Severity** field, Macie displays a list of values to choose from—**Low**, **Medium**, and **High**. If you choose the **S3 bucket name** field, Macie displays a text box in which you can enter a bucket name. Whichever field you choose, Macie guides you through the steps to add a condition that includes the required settings for the field.

After you add a condition, Macie applies the criteria for the condition and adds the condition to a filter box in the filter bar, as shown in the following image.



In this example, the condition is configured to include all medium-severity and high-severity findings, and to exclude all low-severity findings. It returns findings where the value for the **Severity** field *equals* **Medium** or **High**.

Tip

For many fields, you can change a condition's operator from *equals* to *not equals* by choosing the equals icon (●) in a filter box. If you do this, Macie changes the operator to *not equals* and displays the not equals icon (⊘) in the filter box. To switch to the *equals* operator again, choose the not equals icon.

As you add more conditions, Macie applies their criteria and adds them to the filter bar. You can refer to the filter bar at any time to see which criteria you've applied. To remove a condition, choose the remove condition icon (⊗) in the filter box for the condition.

To filter findings using the console

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, choose **Findings**.
3. (Optional) To first view and pivot on findings by a predefined logical group, choose **By bucket**, **By type**, or **By job** in the navigation pane (under **Findings**), and then choose an item in the table. In the details panel, choose the link for the field to pivot on.
4. (Optional) To display findings that were suppressed by a [suppression rule \(p. 203\)](#), choose **Current** in the filter bar. Then choose **Archived** to display only suppressed findings, or choose **All** to display both current and suppressed findings.
5. To add a filter condition:
 - a. Place your cursor in the filter bar, and then choose the field to use for the condition. For information about the fields that you can use, see [Fields for filtering findings \(p. 165\)](#).
 - b. Enter the appropriate type of value for the field. For detailed information about the different types of values, see [Specifying values for fields \(p. 149\)](#).

Array of text (strings)

For this type of value, Macie often provides a list of values to choose from. If this is the case, select each value that you want to use in the condition.

If Macie doesn't provide a list of values, enter a complete, valid value for the field. To specify additional values for the field, choose **Apply**, and then add another condition for each additional value.

Note that values are case sensitive. In addition, you can't use partial values or wildcard characters in values. For example, to filter findings for an S3 bucket named *my-S3-bucket*, enter **my-S3-bucket** as the value for the **S3 bucket name** field. If you enter any other value, such as **my-s3-bucket** or **my-S3**, Macie won't return findings for the bucket.

Boolean

For this type of value, Macie provides a list of values to choose from. Select the value that you want to use in the condition.

Date/Time (time ranges)

For this type of value, use the **From** and **To** boxes to define an inclusive time range:

- To define a fixed time range, use the **From** and **To** boxes to specify the first date and time and the last date and time in the range, respectively.
- To define a relative time range that starts at a certain date and time and ends at the current time, enter the start date and time in the **From** boxes, and delete any text in **To** boxes.
- To define a relative time range that ends at a certain date and time, enter the end date and time in the **To** boxes, and delete any text in the **From** boxes.

Note that time values use 24-hour notation. If you use the date picker to choose dates, you can refine the values by entering text directly in the **From** and **To** boxes.

Number (numeric ranges)

For this type of value, use the **From** and **To** boxes to enter one or more integers that define an inclusive, fixed or relative numeric range.

Text (string) values

For this type of value, enter a complete, valid value for the field.

Note that values are case sensitive. In addition, you can't use partial values or wildcard characters in values. For example, to filter findings for an S3 bucket named *my-S3-bucket*, enter **my-S3-bucket** as the value for the **S3 bucket name** field. If you enter any other value, such as **my-s3-bucket** or **my-S3**, Macie won't return findings for the bucket.

- c. When you finish adding values for the field, choose **Apply**. Macie applies the filter criteria and adds the condition to a filter box in the filter bar.
6. Repeat step 5 for each additional condition that you want to add.
7. To remove a condition, choose the remove condition icon (✕) in the filter box for the condition.
8. To change a condition, remove the condition by choosing the remove condition icon (✕) in the filter box for the condition. Then repeat step 5 to add a condition with the correct settings.

If you want to subsequently use this set of conditions again, you can save the filter as a filter rule. To do this, choose **Save rule** in the filter bar. Then enter a name and, optionally, a description for the rule. When you finish, choose **Save**.

Filtering findings programmatically with the Amazon Macie API

To filter findings programmatically, specify filter criteria in queries that you submit using the [ListFindings](#) or [GetFindingStatistics](#) operation of the Amazon Macie API. The **ListFindings** operation returns an array of finding IDs, one ID for each finding that matches the filter criteria. The **GetFindingStatistics** operation

returns aggregated statistical data about all the findings that match the filter criteria, grouped by a field that you specify in your request.

Note that the **ListFindings** and **GetFindingStatistics** operations are different from operations that you use to [suppress findings](#) (p. 203). Unlike suppression operations, which also specify filter criteria, the **ListFindings** and **GetFindingStatistics** operations only query findings data. They don't perform any action on findings that match filter criteria. To suppress findings, use the [Findings Filters](#) resource of the Amazon Macie API.

To specify filter criteria in a query, include a map of filter conditions in your request. For each condition, specify a field, an operator, and one or more values for the field. The type and number of values depends on the field and operator that you choose. For information about the fields, operators, and types of values that you can use in a condition, see [Fields for filtering findings](#) (p. 165), [Using operators in conditions](#) (p. 151), and [Specifying values for fields](#) (p. 149).

The following examples show you how to specify filter criteria in queries that you submit using the [AWS Command Line Interface \(AWS CLI\)](#). You can also do this by sending HTTPS requests directly to Macie, or by using a current version of another AWS command line tool or an AWS SDK. For information about AWS tools and SDKs, see [Tools to Build on AWS](#).

Examples

- [Example 1: Filter findings based on severity](#) (p. 157)
- [Example 2: Filter findings based on sensitive data category](#) (p. 158)
- [Example 3: Filter findings based on a fixed time range](#) (p. 158)
- [Example 4: Filter findings based on suppression status](#) (p. 159)
- [Example 5: Filter findings based on multiple fields and types of values](#) (p. 159)

The examples use the `list-findings` command. If an example runs successfully, Macie returns a `findingIds` array. The array lists the unique identifier for each finding that matches the filter criteria, as shown in the following example.

```
{
  "findingIds": [
    "1f1c2d74db5d8caa76859ec52example",
    "6cfa9ac820dd6d55cad30d851example",
    "702a6fd8750e567d1a3a63138example",
    "826e94e2a820312f9f964cf60example",
    "274511c3fdcd87010a19a3a42example"
  ]
}
```

If no findings match the filter criteria, Macie returns an empty `findingIds` array.

```
{
  "findingIds": []
}
```

Example 1: Filter findings based on severity

This example uses the `list-findings` command to retrieve finding IDs for all of your high-severity and medium-severity findings in the current AWS Region.

For Linux, macOS, or Unix:

```
$ aws macie2 list-findings --finding-criteria '{"criterion":{"severity.description":{"eq":["High","Medium"]}}}'
```

For Microsoft Windows:

```
C:\> aws macie2 list-findings --finding-criteria={"criterion":{"severity.description":{"eq":["High","Medium"]}}}
```

Where:

- *severity.description* specifies the JSON name of the **Severity** field.
- *eq* specifies the *equals* operator.
- *High* and *Medium* are an array of enumerated values for the **Severity** field.

Example 2: Filter findings based on sensitive data category

This example uses the `list-findings` command to retrieve finding IDs for all of your sensitive data findings that are in the current Region and report occurrences of financial data (and no other categories of sensitive data) in S3 objects.

For Linux, macOS, or Unix, using the backslash (\) line-continuation character to improve readability:

```
$ aws macie2 list-findings \
--finding-criteria '{"criterion":{"classificationDetails.result.sensitiveData.category":{"eqExactMatch":["FINANCIAL_INFORMATION"]}}}'
```

For Microsoft Windows, using the caret (^) line-continuation character to improve readability:

```
C:\> aws macie2 list-findings ^
--finding-criteria={"criterion":{"classificationDetails.result.sensitiveData.category":{"eqExactMatch":["FINANCIAL_INFORMATION"]}}}
```

Where:

- *classificationDetails.result.sensitiveData.category* specifies the JSON name of the **Sensitive data category** field.
- *eqExactMatch* specifies the *equals exact match* operator.
- *FINANCIAL_INFORMATION* is an enumerated value for the **Sensitive data category** field.

Example 3: Filter findings based on a fixed time range

This example uses the `list-findings` command to retrieve finding IDs for all of your findings that are in the current Region and were created between 07:00 UTC October 5, 2020, and 07:00 UTC November 5, 2020 (inclusively).

For Linux, macOS, or Unix:

```
$ aws macie2 list-findings --finding-criteria '{"criterion":{"createdAt":{"gte":1601881200000,"lte":1604559600000}}}'
```

For Microsoft Windows:

```
C:\> aws macie2 list-findings --finding-criteria={"criterion":{"createdAt":{"gte":1601881200000,"lte":1604559600000}}}
```

Where:

- `createdAt` specifies the JSON name of the **Created at** field.
- `gte` specifies the *greater than or equal to* operator.
- `1601881200000` is the first date and time (as a Unix timestamp in milliseconds) in the time range.
- `lte` specifies the *less than or equal to* operator.
- `1604559600000` is the last date and time (as a Unix timestamp in milliseconds) in the time range.

Example 4: Filter findings based on suppression status

This example uses the `list-findings` command to retrieve finding IDs for all of your findings that are in the current Region and were suppressed (automatically archived) by a suppression rule.

For Linux, macOS, or Unix:

```
$ aws macie2 list-findings --finding-criteria '{"criterion":{"archived":{"eq":["true"]}}}'
```

For Microsoft Windows:

```
C:\> aws macie2 list-findings --finding-criteria={"criterion":{"archived":{"eq":["true"]}}}
```

Where:

- `archived` specifies the JSON name of the **Archived** field.
- `eq` specifies the *equals* operator.
- `true` is a Boolean value for the **Archived** field.

Example 5: Filter findings based on multiple fields and types of values

This example uses the `list-findings` command to retrieve finding IDs for all of your sensitive data findings that are in the current Region and match the following criteria: were created between 07:00 UTC October 5, 2020, and 07:00 UTC November 5, 2020 (exclusively); report occurrences of financial data and no other categories of sensitive data in S3 objects; and weren't suppressed (automatically archived) by a suppression rule.

For Linux, macOS, or Unix, using the backslash (\) line-continuation character to improve readability:

```
$ aws macie2 list-findings \  
--finding-criteria '{"criterion":{"createdAt":  
{"gt":"1601881200000","lt":"1604559600000"},"classificationDetails.result.sensitiveData.category":  
{"eqExactMatch":["FINANCIAL_INFORMATION"]},"archived":{"eq":["false"]}}}'
```

For Microsoft Windows, using the caret (^) line-continuation character to improve readability:

```
C:\> aws macie2 list-findings ^  
--finding-criteria={"criterion":{"createdAt":{"gt":"1601881200000",  
"lt":"1604559600000"},"classificationDetails.result.sensitiveData.category\  
{"eqExactMatch":["FINANCIAL_INFORMATION"]},"archived":{"eq":["false"]}}}
```

Where:

- `createdAt` specifies the JSON name of the **Created at** field, and:
 - `gt` specifies the *greater than or equal to* operator.
 - `1601881200000` is the first date and time (as a Unix timestamp in milliseconds) in the time range.

- `lt` specifies the *less than or equal to* operator.
- `1604559600000` is the last date and time (as a Unix timestamp in milliseconds) in the time range.
- `classificationDetails.result.sensitiveData.category` specifies the JSON name of the **Sensitive data category** field, and:
 - `eqExactMatch` specifies the *equals exact match* operator.
 - `FINANCIAL_INFORMATION` is an enumerated value for the field.
- `archived` specifies the JSON name of the **Archived** field, and:
 - `eq` specifies the *equals* operator.
 - `false` is a Boolean value for the field.

Creating and managing filter rules for findings

A *filter rule* is a set of filter criteria that you create and save to use again when you review findings on the Amazon Macie console. Filter rules can help you perform consistent analysis of findings that have specific characteristics. For example, you might create one filter rule for analyzing all high-severity policy findings for S3 buckets that contain unencrypted objects, and another filter rule for analyzing all high-severity sensitive data findings that report specific types of sensitive data.

Note that filter rules are different from suppression rules. A *suppression rule* is a set of filter criteria that you create and save to automatically archive findings that match the criteria of the rule. Although both types of rules store and apply filter criteria, a filter rule doesn't perform any action on findings that match the criteria of the rule. Instead, a filter rule only determines which findings appear on the console after you apply the rule. For information about suppression rules, see [Suppressing findings \(p. 203\)](#).

To create and manage filter rules, you can use the Amazon Macie console or the Amazon Macie API. The following topics explain how. For the API, the topics include examples of how to perform these tasks using the [AWS Command Line Interface \(AWS CLI\)](#). You can also perform these tasks by using a current version of another AWS command line tool or an AWS SDK, or by sending HTTPS requests directly to Macie. For information about AWS tools and SDKs, see [Tools to Build on AWS](#).

Topics

- [Creating filter rules \(p. 160\)](#)
- [Applying filter rules \(p. 162\)](#)
- [Changing filter rules \(p. 162\)](#)
- [Deleting filter rules \(p. 164\)](#)

Creating filter rules

When you create a filter rule, you specify filter criteria, a name, and, optionally, a description of the rule. You can create a filter rule by using the Amazon Macie console or the Amazon Macie API.

Console

Follow these steps to create a filter rule by using the Amazon Macie console.

To create a filter rule

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, choose **Findings**.

Tip

To use an existing filter rule as a starting point, choose the rule from the **Saved rules** list.

You can also streamline creation of a rule by first pivoting and drilling down on findings by a predefined logical group. If you do this, Macie automatically creates and applies the appropriate filter conditions, which can be a helpful starting point for creating a rule. To do this, choose **By bucket**, **By type**, or **By job** in the navigation pane (under **Findings**), and then choose an item in the table. In the details panel, choose the link for the field to pivot on.

3. In the filter bar, add conditions that define the filter criteria for the rule. To learn how, see [Creating and applying filters to findings \(p. 154\)](#).
4. When you finish defining filter criteria for the rule, choose **Save rule** in the filter bar.



5. Under **Filter rule**, enter a name and, optionally, a description of the rule.
6. Choose **Save**.

API

To create a filter rule programmatically, use the [CreateFindingsFilter](#) operation of the Amazon Macie API and specify the appropriate values for the required parameters:

- For the `action` parameter, specify `NOOP` to ensure that Macie doesn't suppress (automatically archive) findings that match the criteria of the rule.
- For the `criterion` parameter, specify a map of conditions that define the filter criteria for the rule.

In the map, each condition should specify a field, an operator, and one or more values for the field. The type and number of values depends on the field and operator that you choose. For information about the fields, operators, and types of values that you can use in a condition, see [Fields for filtering findings \(p. 165\)](#), [Using operators in conditions \(p. 151\)](#), and [Specifying values for fields \(p. 149\)](#).

To create a filter rule by using the AWS CLI, run the `create-findings-filter` command and specify the appropriate values for the required parameters. The following examples create a filter rule that returns all sensitive data findings that are in the current AWS Region and report occurrences of personal information (and no other categories of sensitive data) in S3 objects.

This example is formatted for Linux, macOS, or Unix, and it uses the backslash (\) line-continuation character to improve readability.

```
$ aws macie2 create-findings-filter \  
--action NOOP \  
--name my_filter_rule \  
--finding-criteria '{"criterion":  
{ "classificationDetails.result.sensitiveData.category": {"eqExactMatch":  
["PERSONAL_INFORMATION"]} }'
```

This example is formatted for Microsoft Windows and it uses the caret (^) line-continuation character to improve readability.

```
C:\> aws macie2 create-findings-filter ^  
--action NOOP ^  
--name my_filter_rule ^  
--finding-criteria={"criterion\  
{ "classificationDetails.result.sensitiveData.category": {"eqExactMatch\  
["PERSONAL_INFORMATION"] } }
```

Where:

- `my_filter_rule` is the custom name for the rule.
- `criterion` is a map of filter conditions for the rule:
 - `classificationDetails.result.sensitiveData.category` is the JSON name of the **Sensitive data category** field.
 - `eqExactMatch` specifies the *equals exact match* operator.
 - `PERSONAL_INFORMATION` is an enumerated value for the **Sensitive data category** field.

If the command runs successfully, you receive output similar to the following.

```
{
  "arn": "arn:aws:macie2:us-west-2:123456789012:findings-filter/9b2b4508-aa2f-4940-
b347-d1451example",
  "id": "9b2b4508-aa2f-4940-b347-d1451example"
}
```

Where `arn` is the Amazon Resource Name (ARN) of the filter rule that was created, and `id` is the unique identifier for the rule.

For additional examples of filter criteria, see [Filtering findings programmatically with the Amazon Macie API \(p. 156\)](#).

Applying filter rules

When you apply a filter rule, Macie uses the criteria of the rule to determine which findings to include or exclude from your view of findings on the console. Macie also displays the criteria in the filter bar.

Note that filter rules are designed for use with the Amazon Macie console. You can't use them directly in queries that you submit programmatically using the Amazon Macie API. However, if you're using the API to query findings, you can retrieve the filter criteria for a rule by using the [GetFindingsFilter](#) operation. You can then add the criteria to your query. For information about specifying filter criteria in a query, see [Creating and applying filters to findings \(p. 154\)](#).

Follow these steps to filter findings on the console by applying a filter rule.

To apply a filter rule

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, choose **Findings**.
3. In the **Saved rules** list, choose the filter rule that you want to apply. Macie applies the criteria of the rule and displays the criteria in the filter bar.
4. (Optional) To refine the criteria, use the filter bar to add or remove filter conditions. If you do this, your changes won't affect the settings for the rule. Macie won't save any of your changes unless you explicitly save them as a new rule.
5. To apply a different filter rule, repeat step 3.

After you apply a filter rule, you can quickly remove all of its filter criteria from your view by choosing the **X** in the filter bar.

Changing filter rules

You can change the settings for a filter rule at any time by using the Amazon Macie console or the Amazon Macie API. You can also assign and manage tags for the rule.

A *tag* is a label that you define and assign to certain types of AWS resources. Each tag consists of a required tag key and an optional tag value. Tags can help you identify, categorize, and manage resources in different ways, such as by purpose, owner, environment, or other criteria. To learn more, see [Tagging Amazon Macie resources \(p. 315\)](#).

Console

Follow these steps to change the settings for an existing filter rule by using the Amazon Macie console.

To change a filter rule

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, choose **Findings**.
3. In the **Saved rules** list, choose the edit icon (✎) next to the filter rule that you want to change.
4. Do any of the following:
 - To change the filter criteria of the rule, use the filter bar to enter conditions for the criteria that you want. To learn how, see [Creating and applying filters to findings \(p. 154\)](#).
 - To change the name of the rule, enter a new name in the **Name** box under **Filter rule**.
 - To change the description of the rule, enter a new description in the **Description** box under **Filter rule**.
 - To assign, review, or edit tags for the rule, choose **Manage tags** under **Filter rule**. Then review and change the tags as necessary. A rule can have as many as 50 tags.
5. When you finish making changes, choose **Save**.

API

To change a filter rule programmatically, use the [UpdateFindingsFilter](#) operation of the Amazon Macie API. When you submit your request, use the supported parameters to specify a new value for each setting that you want to change.

For the `id` parameter, specify the unique identifier for the rule to change. You can get this identifier by using the [ListFindingsFilter](#) operation to retrieve a list of filter and suppression rules for your account. If you're using the AWS CLI, run the `list-findings-filters` command to retrieve this list.

To change a filter rule by using the AWS CLI, run the `update-findings-filter` command and use the supported parameters to specify a new value for each setting that you want to change. For example, the following command changes the name of an existing filter rule.

```
C:\> aws macie2 update-findings-filter --id 9b2b4508-aa2f-4940-b347-d1451example --name personal_information_only
```

Where:

- `9b2b4508-aa2f-4940-b347-d1451example` is the unique identifier for the rule.
- `personal_information_only` is the new name for the rule.

If the command runs successfully, you receive output similar to the following.

```
{
  "arn": "arn:aws:macie2:us-west-2:123456789012:findings-filter/9b2b4508-aa2f-4940-b347-d1451example",
  "id": "9b2b4508-aa2f-4940-b347-d1451example"
}
```

Where `arn` is the Amazon Resource Name (ARN) of the rule that was changed, and `id` is the unique identifier for the rule.

Similarly, the following example converts a suppression rule to a filter rule by changing the value for the `action` parameter from `ARCHIVE` to `NOOP`.

```
C:\> aws macie2 update-findings-filter --id 8a1c3508-aa2f-4940-b347-d1451example --  
action NOOP
```

Where:

- `8a1c3508-aa2f-4940-b347-d1451example` is the unique identifier for the rule.
- `NOOP` is the new action for Macie to perform on findings that match the criteria of the rule—perform no action (don't suppress the findings).

If the command runs successfully, you receive output similar to the following:

```
{  
  "arn": "arn:aws:macie2:us-west-2:123456789012:findings-filter/8a1c3508-aa2f-4940-  
b347-d1451example",  
  "id": "8a1c3508-aa2f-4940-b347-d1451example"  
}
```

Where `arn` is the Amazon Resource Name (ARN) of the rule that was changed, and `id` is the unique identifier for the rule.


Deleting filter rules

You can delete a filter rule at any time by using the Amazon Macie console or the Amazon Macie API.

Console

Follow these steps to delete a filter rule by using the Amazon Macie console.

To delete a filter rule

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, choose **Findings**.
3. In the **Saved rules** list, choose the edit icon () next to the filter rule that you want to delete.
4. Under **Filter rule**, choose **Delete**.

API

To delete a filter rule programmatically, use the `DeleteFindingsFilter` operation of the Amazon Macie API. For the `id` parameter, specify the unique identifier for the filter rule to delete. You can get this identifier by using the `ListFindingsFilter` operation to retrieve a list of filter and suppression rules for your account. If you're using the AWS CLI, run the `list-findings-filters` command to retrieve this list.

To delete a filter rule by using the AWS CLI, run the `delete-findings-filter` command. For example:

```
C:\> aws macie2 delete-findings-filter --id 9b2b4508-aa2f-4940-b347-d1451example
```

Where `9b2b4508-aa2f-4940-b347-d1451example` is the unique identifier for the filter rule to delete.

If the command runs successfully, Macie returns an empty HTTP 200 response. Otherwise, Macie returns an HTTP 4xx or 500 response that indicates why the operation failed.

Fields for filtering findings

To help you analyze findings more efficiently, the Amazon Macie console and the Amazon Macie API provide access to several sets of fields for filtering findings:

- **Common fields** – These fields store data that applies to any type of finding. They correlate to common attributes of findings such as severity, finding type, and finding ID.
- **Affected resource fields** – These fields store data about the resources that a finding applies to, such as the name, public access settings, and encryption settings for an affected S3 bucket or object.
- **Policy fields** – These fields store data that's specific to policy findings, such as the action that produced a finding, and the entity that performed the action.
- **Sensitive data classification fields** – These fields store data that's specific to sensitive data findings, such as the types of sensitive data that Macie found, and the unique identifier for the sensitive data discovery job that produced a finding.

A filter can use a combination of fields from any of the preceding sets.

The topics in this section list and describe the individual fields that you can use to filter findings. For additional details about these fields, including any relationships between the fields, see [Findings](#) in the *Amazon Macie API Reference*.

Topics

- [Common fields \(p. 165\)](#)
- [Affected resource fields \(p. 167\)](#)
- [Policy fields \(p. 172\)](#)
- [Sensitive data classification fields \(p. 178\)](#)

Common fields

The following table lists and describes fields that you can use to filter findings based on common finding attributes. These fields store data that applies to any type of finding.

In the table, the **Field** column indicates the name of the field on the Amazon Macie console. The **JSON field** column uses dot notation to indicate the name of the field in JSON representations of findings and the Amazon Macie API. The **Description** column provides a brief description of the data that the field stores, and indicates any requirements for filter values. The table is sorted in ascending alphabetical order by field, and then by JSON field.

Field	JSON field	Description
Account ID*	accountId	The unique identifier for the Amazon Web Services account that the finding applies to. This is typically the account that owns the affected resource.
—	archived	A Boolean value that specifies whether the finding was archived by a suppression rule.

Field	JSON field	Description
		To add this field to a filter on the console, choose Current , Archived , or All in the filter bar.
Category	category	The category of the finding. The console provides a list of values to choose from when you add this field to a filter. In the API, valid values are: <code>CLASSIFICATION</code> , for a sensitive data finding; and, <code>POLICY</code> , for a policy finding.
—	count	The total number of occurrences of the finding. For sensitive data findings, this value is always 1. All sensitive data findings are considered unique because they derive from individual jobs. This field isn't available as a filter option on the console. With the API, you can use this field to define a numeric range for a filter.
Created at	createdAt	The date and time when Macie created the finding. You can use this field to define a time range for a filter.
Finding ID*	id	The unique identifier for the finding. This is a random string that Macie generates and assigns to a finding when it creates the finding.
Finding type*	type	The type of the finding—for example, <code>SensitiveData:S3Object/Personal</code> or <code>Policy:IAMUser/S3BucketPublic</code> . The console provides a list of values to choose from when you add this field to a filter. For a list of valid values in the API, see FindingType in the <i>Amazon Macie API Reference</i> .
Region	region	The AWS Region that Macie created the finding in—for example, <code>us-east-1</code> or <code>ca-central-1</code> .

Field	JSON field	Description
Sample	sample	<p>A Boolean value that specifies whether the finding is a sample finding. A <i>sample finding</i> is a finding that uses example data and placeholder values to demonstrate what a finding might contain.</p> <p>The console provides a list of values to choose from when you add this field to a filter.</p>
Severity	severity.description	<p>The qualitative representation of the finding's severity.</p> <p>The console provides a list of values to choose from when you add this field to a filter. In the API, valid values are: <code>Low</code>, <code>Medium</code>, and <code>High</code>.</p>
Updated at	updatedAt	<p>The date and time when the finding was last updated. For sensitive data findings, this value is the same as the value for the Created at field. All sensitive data findings are considered new because they derive from individual jobs.</p> <p>You can use this field to define a time range for a filter.</p>

* To specify multiple values for this field on the console, add a condition that uses the field and specifies a distinct value for the filter, and then repeat that step for each additional value. To do this with the API, use an array that lists the values to use for the filter.

Affected resource fields

The following topics list and describe the fields that you can use to filter findings based on the resource that a finding applies to. The topics are organized by resource type.

Topics

- [S3 bucket \(p. 167\)](#)
- [S3 object \(p. 171\)](#)

S3 bucket

The following table lists and describes fields that you can use to filter findings based on characteristics of the S3 bucket that a finding applies to.

In the table, the **Field** column indicates the name of the field on the Amazon Macie console. The **JSON field** column uses dot notation to indicate the name of the field in JSON representations of findings and the Amazon Macie API. (Longer JSON field names use the newline character sequence `(\n)` to improve

readability.) The **Description** column provides a brief description of the data that the field stores, and indicates any requirements for filter values. The table is sorted in ascending alphabetical order by field, and then by JSON field.

Field	JSON field	Description
—	resourcesAffected.s3Bucket	The date and time when the affected bucket was created. This field isn't available as a filter option on the console. With the API, you can use this field to define a time range for a filter.
S3 bucket default encryption	resourcesAffected.s3Bucket	The type of server-side encryption that's used by default to encrypt objects that are added to the affected bucket. The console provides a list of values to choose from when you add this field to a filter. For a list of valid values for the API, see EncryptionType in the <i>Amazon Macie API Reference</i> .
S3 bucket encryption KMS key id*	resourcesAffected.s3Bucket	The Amazon Resource Name (ARN) or unique identifier (key ID) for the AWS KMS key that's used by default to encrypt objects that are added to the affected bucket.
S3 bucket encryption required by bucket policy	resourcesAffected.s3Bucket	Specifies whether the bucket policy for the affected bucket requires server-side encryption of objects when objects are uploaded to the bucket. The console provides a list of values to choose from when you add this field to a filter. For a list of valid values for the API, see S3Bucket in the <i>Amazon Macie API Reference</i> .
S3 bucket name*	resourcesAffected.s3Bucket	The name of the affected bucket.
S3 bucket owner display name*	resourcesAffected.s3Bucket	The display name of the AWS user who owns the affected bucket.
S3 bucket public access permission	resourcesAffected.s3Bucket	Specifies whether the affected bucket is publicly accessible based on a combination of permissions settings that apply to the bucket.

Field	JSON field	Description
		The console provides a list of values to choose from when you add this field to a filter. For a list of valid values for the API, see BucketPublicAccess in the <i>Amazon Macie API Reference</i> .
—	<code>resourcesAffected.s3BucketPublicAccessBlockPublicAccess</code> <code>accountLevelPermissions.blockPublicAccess</code>	A Boolean value that specifies whether Amazon S3 blocks public access control lists (ACLs) for the affected bucket and objects in the bucket. This is an account-level, block public access setting for the bucket. This field isn't available as a filter option on the console.
—	<code>resourcesAffected.s3BucketPublicPolicy</code> <code>accountLevelPermissions.blockPublicPolicy</code>	A Boolean value that specifies whether Amazon S3 blocks public bucket policies for the affected bucket. This is an account-level, block public access setting for the bucket. This field isn't available as a filter option on the console.
—	<code>resourcesAffected.s3BucketPublicACLs</code> <code>accountLevelPermissions.blockPublicACLs</code>	A Boolean value that specifies whether Amazon S3 ignores public ACLs for the affected bucket and objects in the bucket. This is an account-level, block public access setting for the bucket. This field isn't available as a filter option on the console.
—	<code>resourcesAffected.s3BucketRestrictPublicBucket</code> <code>accountLevelPermissions.blockPublicBucket</code>	A Boolean value that specifies whether Amazon S3 restricts public bucket policies for the affected bucket. This is an account-level, block public access setting for the bucket. This field isn't available as a filter option on the console.
—	<code>resourcesAffected.s3BucketPublicReadAccess</code> <code>bucketLevelPermissions.allowPublicReadAccess</code>	A Boolean value that specifies whether the bucket-level ACL for the affected bucket grants the general public with read permissions for the bucket. This field isn't available as a filter option on the console.

Field	JSON field	Description
—	<code>resourcesAffected.s3Bucket</code> \n <code>bucketLevelPermissions.accessControlListWithWritePublicWriteAccessPermissions</code>	A Boolean value that specifies whether the bucket-level ACL for the affected bucket grants the specified public write access permissions for the bucket. This field isn't available as a filter option on the console.
—	<code>resourcesAffected.s3Bucket</code> \n <code>bucketLevelPermissions.blockPublicAccess</code>	A Boolean value that specifies whether Amazon S3 blocks public ACLs for the affected bucket and objects in the bucket. This is a bucket-level, block public access setting for a bucket. This field isn't available as a filter option on the console.
—	<code>resourcesAffected.s3Bucket</code> \n <code>bucketLevelPermissions.blockPublicPolicy</code>	A Boolean value that specifies whether Amazon S3 blocks public bucket policies for the affected bucket. This is a bucket-level, block public access setting for the bucket. This field isn't available as a filter option on the console.
—	<code>resourcesAffected.s3Bucket</code> \n <code>bucketLevelPermissions.blockPublicAcls</code>	A Boolean value that specifies whether Amazon S3 ignores public ACLs for the affected bucket and objects in the bucket. This is a bucket-level, block public access setting for the bucket. This field isn't available as a filter option on the console.
—	<code>resourcesAffected.s3Bucket</code> \n <code>bucketLevelPermissions.blockPublicBuckets</code>	A Boolean value that specifies whether Amazon S3 restricts public bucket policies for the affected bucket. This is a bucket-level, block public access setting for the bucket. This field isn't available as a filter option on the console.

Field	JSON field	Description
—	resourcesAffected.s3BucketLevelPermissions.bucketPublicReadAccess	A Boolean value that specifies whether the affected bucket's policy allows the general public to have read access to the bucket. This field isn't available as a filter option on the console.
—	resourcesAffected.s3BucketLevelPermissions.bucketPublicWriteAccess	A Boolean value that specifies whether the affected bucket's policy allows the general public to have write access to the bucket. This field isn't available as a filter option on the console.
S3 bucket tag key*	resourcesAffected.s3BucketTagKey	A tag key that's associated with the affected bucket.
S3 bucket tag value*	resourcesAffected.s3BucketTagValue	A tag value that's associated with the affected bucket.

* To specify multiple values for this field on the console, add a condition that uses the field and specifies a distinct value for the filter, and then repeat that step for each additional value. To do this with the API, use an array that lists the values to use for the filter.

S3 object

The following table lists and describes fields that you can use to filter findings based on characteristics of the S3 object that a finding applies to.

In the table, the **Field** column indicates the name of the field on the Amazon Macie console. The **JSON field** column uses dot notation to indicate the name of the field in JSON representations of findings and the Amazon Macie API. The **Description** column provides a brief description of the data that the field stores, and indicates any requirements for filter values. The table is sorted in ascending alphabetical order by field, and then by JSON field.

Field	JSON field	Description
S3 object encryption KMS key id*	resourcesAffected.s3ObjectKmsMasterKeyId	The Amazon Resource Name (ARN) or unique identifier (key ID) for the AWS KMS key that was used to encrypt the affected object.
S3 object encryption type	resourcesAffected.s3ObjectEncryptionType	The type of server-side encryption that was used to encrypt the affected object. The console provides a list of values to choose from when you add this field to a filter. For a list of valid values for the API, see

Field	JSON field	Description
		EncryptionType in the <i>Amazon Macie API Reference</i> .
—	<code>resourcesAffected.s3Object.fileName</code>	The file name extension of the affected object. For objects that don't have a file name extension, specify "" as the value for the filter. This field isn't available as a filter option on the console.
—	<code>resourcesAffected.s3Object.lastModifiedDate</code>	The date and time when the affected object was created or last changed, whichever is latest. This field isn't available as a filter option on the console. With the API, you can use this field to define a time range for a filter.
S3 object key*	<code>resourcesAffected.s3Object.key</code>	The full key (name) that's assigned to the affected object.
—	<code>resourcesAffected.s3Object.path</code>	The path to the affected object, including the full key (name). This field isn't available as a filter option on the console.
S3 object public access	<code>resourcesAffected.s3Object.publicAccess</code>	A Boolean value that specifies whether the affected object is publicly accessible based on a combination of permission settings that apply to the object. The console provides a list of values to choose from when you add this field to a filter.
S3 object tag key*	<code>resourcesAffected.s3Object.tagKey</code>	A tag key that's associated with the affected object.
S3 object tag value*	<code>resourcesAffected.s3Object.tagValue</code>	A tag value that's associated with the affected object.

* To specify multiple values for this field on the console, add a condition that uses the field and specifies a distinct value for the filter, and then repeat that step for each additional value. To do this with the API, use an array that lists the values to use for the filter.

Policy fields

The following table lists and describes fields that you can use to filter policy findings. These fields store data that's specific to policy findings.

In the table, the **Field** column indicates the name of the field on the Amazon Macie console. The **JSON field** column uses dot notation to indicate the name of the field in JSON representations of findings and the Amazon Macie API. (Longer JSON field names use the newline character sequence (\n) to improve readability.) The **Description** column provides a brief description of the data that the field stores, and indicates any requirements for filter values. The table is sorted in ascending alphabetical order by field, and then by JSON field.

Field	JSON field	Description
Action type	policyDetails.action.actionType	The type of action that produced the finding. The only valid value for this field is <code>AWS_API_CALL</code> .
API call name*	policyDetails.action.apiCallName	The name of the operation that was invoked most recently and produced the finding—for example, <code>DeleteBucketEncryption</code> .
API service name*	policyDetails.action.apiCallServiceName	The URI of the AWS service that provides the operation that was invoked and produced the finding—for example, <code>s3.amazonaws.com</code> .
—	policyDetails.action.apiCallTimestamp	The first date and time when any operation was invoked and produced the finding. This field isn't available as a filter option on the console. With the API, you can use this field to define a time range for a filter.
—	policyDetails.action.apiCallTimestampMs	The most recent date and time when the specified operation (API call name or <code>api</code>) was invoked and produced the finding. This field isn't available as a filter option on the console. With the API, you can use this field to define a time range for a filter.
—	policyDetails.actor.domainName	The domain name of the device that was used to perform the action. This field isn't available as a filter option on the console.
IP city*	policyDetails.actor.ipAddressCity	The name of the originating city for the IP address of the device that was used to perform the action.

Field	JSON field	Description
IP country*	<code>policyDetails.actor.ipAddress</code>	The name of the originating country for the IP address of the device that was used to perform the action—for example, United States.
—	<code>policyDetails.actor.ipAddress</code>	The Autonomous System as a Number (ASN) for the autonomous system that included the IP address of the device that was used to perform the action. This field isn't available as a filter option on the console.
IP owner ASN org*	<code>policyDetails.actor.ipAddress</code>	The organization identifier that's associated with the ASN for the autonomous system that included the IP address of the device that was used to perform the action.
IP owner ISP*	<code>policyDetails.actor.ipAddress</code>	The name of the Internet service provider (ISP) that owned the IP address of the device that was used to perform the action.
IP V4 address*	<code>policyDetails.actor.ipAddress</code>	The Internet Protocol version 4 (IPv4) address of the device that was used to perform the action.
—	<code>policyDetails.actor.userId</code>	For an action performed with temporary security credentials that were obtained using the AssumeRole operation of the AWS STS API, the AWS access key ID that identifies the credentials. This field isn't available as a filter option on the console.
User identity assumed role account id*	<code>policyDetails.actor.userId</code>	For an action performed with temporary security credentials that were obtained using the AssumeRole operation of the AWS STS API, the unique identifier for the Amazon Web Services account that owns the entity that was used to get the credentials.

Field	JSON field	Description
User identity assumed role principal id*	<code>policyDetails.actor.userId</code>	For an action performed with temporary security credentials that were obtained using the AssumeRole operation of the AWS STS API, the unique identifier for the entity that was used to get the credentials.
User identity assumed role session ARN*	<code>policyDetails.actor.userId</code>	For an action performed with temporary security credentials that were obtained using the AssumeRole operation of the AWS STS API, the Amazon Resource Name (ARN) of the source account, IAM user, or role that was used to get the credentials.
—	<code>policyDetails.actor.userId</code> <code>\n</code> <code>sessionIssuer.type</code>	For an action performed with temporary security credentials that were obtained using the AssumeRole operation of the AWS STS API, the source of the temporary security credentials—for example, Root, IAMUser, or Role. This field isn't available as a filter option on the console.
—	<code>policyDetails.actor.userId</code> <code>\n</code> <code>sessionIssuer.userName</code>	For an action performed with temporary security credentials that were obtained using the AssumeRole operation of the AWS STS API, the name or alias of the user or role that issued the session. Note that this value is null if the credentials were obtained from a root account that doesn't have an alias. This field isn't available as a filter option on the console.
User identity AWS account account id*	<code>policyDetails.actor.userId</code>	For an action performed using the credentials for another Amazon Web Services account, the unique identifier for the account.
User identity AWS account principal id*	<code>policyDetails.actor.userId</code>	For an action performed using the credentials for another Amazon Web Services account, the unique identifier for the entity that performed the action.

Field	JSON field	Description
User identity AWS service invoked by	<code>policyDetails.actor.userIdentityDetails.serviceName</code>	For any actions performed by an account that belongs to an Amazon Web Services service, the name of the service.
—	<code>policyDetails.actor.userIdentityDetails.accessKeyId</code>	For any actions performed with temporary security credentials that were obtained using the <code>GetFederationToken</code> operation of the AWS STS API, the AWS access key ID that identifies the credentials. This field isn't available as a filter option on the console.
User identity federated session ARN*	<code>policyDetails.actor.userIdentityDetails.sessionArn</code>	For any actions performed with temporary security credentials that were obtained using the <code>GetFederationToken</code> operation of the AWS STS API, the ARN of the entity that was used to get the credentials.
User identity federated user account id*	<code>policyDetails.actor.userIdentityDetails.accountId</code>	For any actions performed with temporary security credentials that were obtained using the <code>GetFederationToken</code> operation of the AWS STS API, the unique identifier for the Amazon Web Services account that owns the entity that was used to get the credentials.
User identity federated user principal id*	<code>policyDetails.actor.userIdentityDetails.principalId</code>	For any actions performed with temporary security credentials that were obtained using the <code>GetFederationToken</code> operation of the AWS STS API, the unique identifier for the entity that was used to get the credentials.
—	<code>policyDetails.actor.userIdentityDetails.sessionContext.sessionIssuer.type</code>	For any actions performed with temporary security credentials that were obtained using the <code>GetFederationToken</code> operation of the AWS STS API, the source of the temporary security credentials—for example, <code>Root</code> , <code>IAMUser</code> , or <code>Role</code> . This field isn't available as a filter option on the console.

Field	JSON field	Description
—	<code>policyDetails.actor.userIdentityContext.sessionIssuer.userName</code>	For an action performed with temporary security credentials that were obtained using the <code>GetFederationToken</code> operation of the AWS STS API, the name or alias of the user or role that issued the session. Note that this value is null if the credentials were obtained from a root account that doesn't have an alias. This field isn't available as a filter option on the console.
User identity IAM account id*	<code>policyDetails.actor.userIdentityContext.accountId</code>	For an action performed using an IAM user's credentials, the unique identifier for the Amazon Web Services account that's associated with the IAM user who performed the action.
User identity IAM principal id*	<code>policyDetails.actor.userIdentityContext.principalId</code>	For an action performed using an IAM user's credentials, the unique identifier for the IAM user who performed the action.
User identity IAM user name*	<code>policyDetails.actor.userIdentityContext.userName</code>	For an action performed using an IAM user's credentials, the user name of the IAM user who performed the action.
User identity root account id*	<code>policyDetails.actor.userIdentityContext.rootAccountId</code>	For an action performed using the credentials for your Amazon Web Services account, the unique identifier for the account.
User identity root principal id*	<code>policyDetails.actor.userIdentityContext.rootPrincipalId</code>	For an action performed using the credentials for your Amazon Web Services account, the unique identifier for the entity that performed the action.
User identity type	<code>policyDetails.actor.userIdentityContext.type</code>	The type of entity that performed the action that produced the finding. The console provides a list of values to choose from when you add this field to a filter. For a list of valid values for the API, see UserIdentityType in the <i>Amazon Macie API Reference</i> .

* To specify multiple values for this field on the console, add a condition that uses the field and specifies a distinct value for the filter, and then repeat that step for each additional value. To do this with the API, use an array that lists the values to use for the filter.

Sensitive data classification fields

The following table lists and describes fields that you can use to filter sensitive data findings. These fields store data that's specific to sensitive data findings.

In the table, the **Field** column indicates the name of the field on the Amazon Macie console. The **JSON field** column uses dot notation to indicate the name of the field in JSON representations of findings and the Amazon Macie API. The **Description** column provides a brief description of the data that the field stores, and indicates any requirements for filter values. The table is sorted in ascending alphabetical order by field, and then by JSON field.

Field	JSON field	Description
Custom data identifier detection ARN*	classificationDetails.results.detections.customDataIdentifierArn	The Amazon Resource Name (ARN) of the custom data identifier that detected the data and produced the finding.
Custom data identifier detection name*	classificationDetails.results.detections.customDataIdentifierName	The name of the custom data identifier that detected the data and produced the finding.
Custom data identifier total count	classificationDetails.results.detections.customDataIdentifierTotalCount	The total number of occurrences of data that was detected by custom data identifiers and produced the finding. You can use this field to define a numeric range for a filter.
Job ID*	classificationDetails.jobId	The unique identifier for the sensitive data discovery job that produced the finding.
Origin type	classificationDetails.originType	How Macie found the sensitive data that produced the finding. The only valid value is SENSITIVE_DATA_DISCOVERY_JOB.
—	classificationDetails.results.detections.contentType	The type of content, as a MIME type, that the finding applies to—for example, text/csv for a CSV file or application/pdf for an Adobe Portable Document Format file. This field isn't available as a filter option on the console.
—	classificationDetails.results.detections.storageSize	The total storage size, in bytes, of the S3 object that the finding applies to.

Field	JSON field	Description
		This field isn't available as a filter option on the console. With the API, you can use this field to define a numeric range for a filter.
Result status code*	<code>classificationDetails.results</code>	The status of the finding. Valid values are: <ul style="list-style-type: none"> • COMPLETE – Macie completed its analysis of the object. • PARTIAL – Macie analyzed only a subset of the data in the object. For example, the object is an archive file that contains files in an unsupported format. • SKIPPED – Macie wasn't able to analyze the object. For example, the object is a malformed file.
Sensitive data category	<code>classificationDetails.results</code>	The category of sensitive data that was detected and produced the finding. <p>The console provides a list of values to choose from when you add this field to a filter. In the API, valid values are: CREDENTIALS, FINANCIAL_INFORMATION, and PERSONAL_INFORMATION.</p>
Sensitive data detection type	<code>classificationDetails.results</code>	The type of sensitive data that was detected and produced the finding. <p>The console provides a list of values to choose from when you add this field to a filter. For a complete list of types, see Sensitive data detection types (p. 180).</p>
Sensitive data total count	<code>classificationDetails.results</code>	The total number of occurrences of the sensitive data that was detected and produced the finding. <p>You can use this field to define a numeric range for a filter.</p>

* To specify multiple values for this field on the console, add a condition that uses the field and specifies a distinct value for the filter, and then repeat that step for each additional value. To do this with the API, use an array that lists the values to use for the filter.

Sensitive data detection types

The following topics list values that you can specify for the **Sensitive data detection type** field in a filter. (The JSON name of this field is `classificationDetails.result.sensitiveData.detections.type`.) The topics are organized by the categories of sensitive data that Macie can detect using managed data identifiers.

Categories

- [Credentials \(p. 180\)](#)
- [Financial information \(p. 180\)](#)
- [Personal information \(p. 181\)](#)

To learn more about a specific detection type, see [Using managed data identifiers \(p. 45\)](#).

Credentials

You can specify the following values to filter findings that report occurrences of credentials data in S3 objects.

Detection type	Filter values
AWS secret access key	AWS_CREDENTIALS
HTTP Basic Authorization header	HTTP_BASIC_AUTH_HEADER
JSON Web Token (JWT)	JSON_WEB_TOKEN
OpenSSH private key	OPENSSSH_PRIVATE_KEY
PGP private key	PGP_PRIVATE_KEY
Public Key Cryptography Standard (PKCS) private key	PKCS
PuTTY private key	PUTTY_PRIVATE_KEY

Financial information

You can specify the following values to filter findings that report occurrences of financial information in S3 objects.

Detection type	Filter values
Bank account number	BANK_ACCOUNT_NUMBER (for Canadian and US bank account numbers), FRANCE_BANK_ACCOUNT_NUMBER, GERMANY_BANK_ACCOUNT_NUMBER, ITALY_BANK_ACCOUNT_NUMBER, SPAIN_BANK_ACCOUNT_NUMBER, UK_BANK_ACCOUNT_NUMBER
Credit card expiration date	CREDIT_CARD_EXPIRATION

Detection type	Filter values
Credit card magnetic strip data	CREDIT_CARD_MAGNETIC_STRIPE
Credit card number	CREDIT_CARD_NUMBER (for credit card numbers that are in proximity of a keyword) and CREDIT_CARD_NUMBER_(NO_KEYWORD) (for credit card numbers that aren't in proximity of a keyword)
Credit card verification code	CREDIT_CARD_SECURITY_CODE

Personal information

You can specify the following values to filter findings that report occurrences of personal health information (PHI) in S3 objects.

Detection type	Filter values
Drug Enforcement Agency (DEA) Registration Number	US_DRUG_ENFORCEMENT_AGENCY_NUMBER
Health Insurance Claim Number (HICN)	USA_HEALTH_INSURANCE_CLAIM_NUMBER
Health insurance or medical identification number	CANADA_HEALTH_NUMBER, EUROPEAN_HEALTH_INSURANCE_CARD_NUMBER, FINLAND_EUROPEAN_HEALTH_INSURANCE_NUMBER, FRANCE_HEALTH_INSURANCE_NUMBER, UK_NHS_NUMBER, USA_MEDICARE_BENEFICIARY_IDENTIFIER
Healthcare Common Procedure Coding System (HCPCS) code	USA_HEALTHCARE_PROCEDURE_CODE
National Drug Code (NDC)	USA_NATIONAL_DRUG_CODE
National Provider Identifier (NPI)	USA_NATIONAL_PROVIDER_IDENTIFIER
Unique device identifier (UDI)	MEDICAL_DEVICE_UDI

You can specify the following values to filter findings that report occurrences of personally identifiable information (PII) in S3 objects.

Detection type	Filter values
Birth date	DATE_OF_BIRTH
Driver's license identification number	AUSTRALIA_DRIVERS_LICENSE, AUSTRIA_DRIVERS_LICENSE, BELGIUM_DRIVERS_LICENSE, BULGARIA_DRIVERS_LICENSE, CANADA_DRIVERS_LICENSE, CROATIA_DRIVERS_LICENSE, CYPRUS_DRIVERS_LICENSE, CZECHIA_DRIVERS_LICENSE, DENMARK_DRIVERS_LICENSE,

Detection type	Filter values
	DRIVERS_LICENSE (for the US), ESTONIA_DRIVERS_LICENSE, FINLAND_DRIVERS_LICENSE, FRANCE_DRIVERS_LICENSE, GERMANY_DRIVERS_LICENSE, GREECE_DRIVERS_LICENSE, HUNGARY_DRIVERS_LICENSE, IRELAND_DRIVERS_LICENSE, ITALY_DRIVERS_LICENSE, LATVIA_DRIVERS_LICENSE, LITHUANIA_DRIVERS_LICENSE, LUXEMBOURG_DRIVERS_LICENSE, MALTA_DRIVERS_LICENSE, NETHERLANDS_DRIVERS_LICENSE, POLAND_DRIVERS_LICENSE, PORTUGAL_DRIVERS_LICENSE, ROMANIA_DRIVERS_LICENSE, SLOVAKIA_DRIVERS_LICENSE, SLOVENIA_DRIVERS_LICENSE, SPAIN_DRIVERS_LICENSE, SWEDEN_DRIVERS_LICENSE, UK_DRIVERS_LICENSE
Electoral roll number	UK_ELECTORAL_ROLL_NUMBER
Full name	NAME
Global Positioning System (GPS) coordinates	LATITUDE_LONGITUDE
HTTP cookie	HTTP_COOKIE
Mailing address	ADDRESS, BRAZIL_CEP_CODE
National identification number	BRAZIL_RG_NUMBER, FRANCE_NATIONAL_IDENTIFICATION_NUMBER, GERMANY_NATIONAL_IDENTIFICATION_NUMBER, ITALY_NATIONAL_IDENTIFICATION_NUMBER, SPAIN_DNI_NUMBER
National Insurance Number (NINO)	UK_NATIONAL_INSURANCE_NUMBER
Passport number	CANADA_PASSPORT_NUMBER, FRANCE_PASSPORT_NUMBER, GERMANY_PASSPORT_NUMBER, ITALY_PASSPORT_NUMBER, SPAIN_PASSPORT_NUMBER, UK_PASSPORT_NUMBER, USA_PASSPORT_NUMBER
Permanent residence number	CANADA_NATIONAL_IDENTIFICATION_NUMBER
Phone number	BRAZIL_PHONE_NUMBER, FRANCE_PHONE_NUMBER, GERMANY_PHONE_NUMBER, ITALY_PHONE_NUMBER, PHONE_NUMBER (for Canada and the US), SPAIN_PHONE_NUMBER, UK_PHONE_NUMBER

Detection type	Filter values
Social Insurance Number (SIN)	CANADA_SOCIAL_INSURANCE_NUMBER
Social Security number (SSN)	SPAIN_SOCIAL_SECURITY_NUMBER, USA_SOCIAL_SECURITY_NUMBER
Taxpayer identification or reference number	AUSTRALIA_TAX_FILE_NUMBER, BRAZIL_CNPJ_NUMBER, BRAZIL_CPF_NUMBER, FRANCE_TAX_IDENTIFICATION_NUMBER, GERMANY_TAX_IDENTIFICATION_NUMBER, SPAIN_NIE_NUMBER, SPAIN_NIF_NUMBER, SPAIN_TAX_IDENTIFICATION_NUMBER, UK_TAX_IDENTIFICATION_NUMBER, USA_INDIVIDUAL_TAX_IDENTIFICATION_NUMBER
Vehicle identification number (VIN)	VEHICLE_IDENTIFICATION_NUMBER

Investigating sensitive data with Amazon Macie findings

When you run a sensitive data discovery job, Amazon Macie captures details about the location of each occurrence of sensitive data that it finds in an Amazon Simple Storage Service (Amazon S3) object. This includes sensitive data that Macie detects using [managed data identifiers \(p. 45\)](#), and data that matches the criteria of any [custom data identifiers \(p. 64\)](#) that you configure the job to use.

With sensitive data findings, you can review these details for as many as 15 occurrences of sensitive data that Macie detects in an affected S3 object. The details provide insight into the breadth of the categories and types of sensitive data that specific S3 buckets and objects contain. They can help you determine whether to perform a deeper investigation of specific buckets and objects, and locate individual occurrences of sensitive data in S3 objects.

For additional insight, you can optionally configure and use Macie to retrieve samples of sensitive data that Macie reports in individual findings. The samples can help you verify the nature of the sensitive data that Macie found, and tailor your investigation of an affected S3 bucket and object. If you choose to retrieve sensitive data samples for a finding, Macie uses data in the finding to locate 1-10 occurrences of each type of sensitive data reported by the finding. Macie then extracts those occurrences of sensitive data from the affected S3 object and displays the data for you to review.

If an S3 object contains many occurrences of sensitive data, a finding can also help you navigate to the corresponding sensitive data discovery result. Unlike a sensitive data finding, a sensitive data discovery result provides detailed location data for as many as 1,000 occurrences of each type of sensitive data that Macie finds in an object. To learn about differences between sensitive data findings and sensitive data discovery results, see [Reviewing job statistics and results \(p. 113\)](#). Macie uses the same schema for location data in sensitive data findings and sensitive data discovery results.

The topics in this section explain how to locate and optionally retrieve occurrences of sensitive data by using sensitive data findings. They also explain the schema that Macie uses to report the location of individual occurrences of sensitive data.

Topics

- [Locating sensitive data with Amazon Macie findings \(p. 184\)](#)
- [Retrieving sensitive data samples with Amazon Macie findings \(p. 186\)](#)
- [JSON schema for sensitive data locations \(p. 196\)](#)

Locating sensitive data with Amazon Macie findings

When you run a sensitive data discovery job, Amazon Macie performs a deep inspection of the latest version of each Amazon Simple Storage Service (Amazon S3) object that you configure the job to analyze. Macie also uses a *depth-first search* algorithm to populate the job's findings with details about the location of specific occurrences of sensitive data that Macie finds. These occurrences provide insight into the categories and types of sensitive data that the affected S3 bucket and object contain. The details can help you determine whether to perform a deeper investigation of specific buckets and objects, and locate individual occurrences of sensitive data in S3 objects.

With sensitive data findings, you can determine the location of as many as 15 occurrences of sensitive data that Macie finds in an affected S3 object. This includes sensitive data that Macie detects using [managed data identifiers](#) (p. 45), and data that matches the criteria of [custom data identifiers](#) (p. 64) that you configure a job to use.

A sensitive data finding can provide details such as:

- The column and row number for a cell or field in a Microsoft Excel workbook, CSV file, or TSV file.
- The path to a field or array in a JSON or JSON Lines file.
- The line number for a line in a non-binary text file other than a CSV, JSON, JSON Lines, or TSV file—for example, an HTML, TXT, or XML file.
- The page number for a page in an Adobe Portable Document Format (PDF) file.
- The record index and the path to a field in a record in an Apache Avro object container or Apache Parquet file.

You can access these details by using the Amazon Macie console or the Amazon Macie API. You can also access these details in findings that Macie publishes to other AWS services, both Amazon EventBridge and AWS Security Hub. To learn how to access the details in findings that Macie publishes to other AWS services, see [Monitoring and processing findings](#) (p. 215). To learn about the JSON structures that Macie uses to report the location of sensitive data, see [JSON schema for sensitive data locations](#) (p. 196).

If an S3 object contains many occurrences of sensitive data, you can also use a finding to navigate to its corresponding sensitive data discovery result. Unlike a sensitive data finding, a sensitive data discovery result provides detailed location data for as many as 1,000 occurrences of each type of sensitive data that Macie finds in an object. If an S3 object is an archive file, such as a .tar or .zip file, this includes occurrences of sensitive data in individual files that Macie extracts from the archive. (Macie doesn't include this information in sensitive data findings.) For more information about sensitive data discovery results, see [Reviewing job statistics and results](#) (p. 113). Macie uses the same JSON schema for location data in sensitive data findings and sensitive data discovery results.

Locating occurrences of sensitive data

To locate occurrences of sensitive data, you can use the Amazon Macie console or the Amazon Macie API. The following steps explain how to locate sensitive data by using the console.

To locate sensitive data programmatically, use the [GetFindings](#) operation of the Amazon Macie API. If a finding includes details about the location of one or more occurrences of a specific type of sensitive data, `occurrences` objects in the finding provide these details. For more information, see [JSON schema for sensitive data locations](#) (p. 196).

To locate occurrences of sensitive data

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.

- In the navigation pane, choose **Findings**.

Tip

You can use the **Jobs** page to display all the findings from a particular job. To do this, choose **Jobs** in the navigation pane, and then choose the name of the job. At the top of the details panel, choose **Show results**, and then choose **Show findings**.

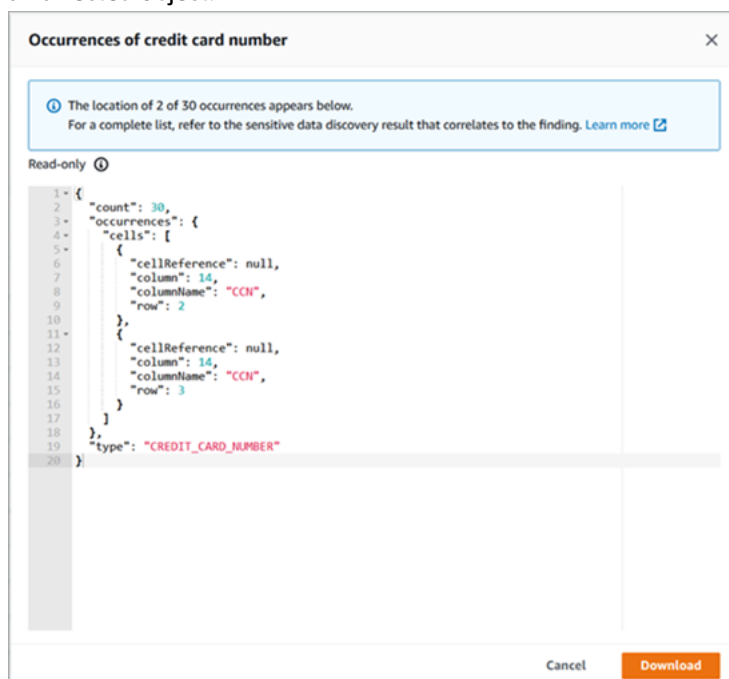
- On the **Findings** page, choose the finding for the sensitive data that you want to locate. The details panel displays information for the finding.
- In the details panel, scroll to the **Sensitive data** section. This section provides information about the categories and types of sensitive data that Macie found in the affected S3 object. It also indicates the number of occurrences of each type of sensitive data that Macie found.

For example, the following image shows some details of a finding that reports 30 occurrences of credit card numbers, 30 occurrences of names, and 30 occurrences of US Social Security numbers.

Financial information	
Credit card number	30
Personal information	
Name	30
Usa social security number	30

If the finding includes details about the location of one or more occurrences of a specific type of sensitive data, the number of occurrences is a link. Choose the link to show the details. Macie opens a new window and displays the details in JSON format.

For example, the following image shows the location of two occurrences of credit card numbers in an affected object.



To save the details as a JSON file, choose **Download**, and then specify a name and location for the file.

- (Optional) To save all the finding's details as a JSON file, choose the finding's identifier (**Finding ID**) at the top of the details panel. Macie opens a new window and displays all the details in JSON format. Choose **Download**, and then specify a name and location for the file.

To access details about the location of as many as 1,000 occurrences of each type of sensitive data in the affected object, refer to the corresponding sensitive data discovery result for the finding. To do this, scroll to the beginning of the **Details** section of the panel, and then choose the link in the **Detailed result location** field. Macie opens the Amazon S3 console and displays the file or folder that contains the discovery result. To learn more about these results, see [Reviewing job statistics and results \(p. 113\)](#).

Retrieving sensitive data samples with Amazon Macie findings

To verify the nature of sensitive data that Amazon Macie detects and reports in findings, you can optionally configure and use Macie to retrieve and reveal samples of sensitive data reported by individual findings. This includes sensitive data that Macie detects using [managed data identifiers \(p. 45\)](#), and data that matches the criteria of [custom data identifiers \(p. 64\)](#) that you configure sensitive data discovery jobs to use. The samples can help you both verify the nature of the data that Macie found, and tailor your investigation of an affected Amazon Simple Storage Service (Amazon S3) object and bucket.

Each time you retrieve and reveal sensitive data samples for a finding, Macie performs the following general tasks:

1. Verifies that the finding specifies the location of individual occurrences of sensitive data and the location of the corresponding sensitive data discovery result.
2. Evaluates the corresponding sensitive data discovery result, checking the validity of both the metadata for the affected S3 object and the location data for individual occurrences of sensitive data in the affected object.
3. By using data in the sensitive data discovery result, locates the first 1–10 occurrences of sensitive data reported by the finding, and extracts the first 1–128 characters of each occurrence from the affected S3 object. If the finding reports multiple types of sensitive data, Macie does this for up to 100 types.
4. Encrypts the extracted data with an AWS Key Management Service (AWS KMS) key that you specify.
5. Temporarily stores the encrypted data in a cache and displays the data for you to review. The data is encrypted at all times, both in transit and at rest.
6. Soon after extraction and encryption, permanently deletes the data from the cache unless additional retention is temporarily required to resolve an operational issue.

Macie doesn't use the Macie [service-linked role \(p. 302\)](#) for your account to perform these tasks. Instead, you use your AWS Identity and Access Management (IAM) identity to locate, retrieve, encrypt, and reveal the samples. You can retrieve and reveal sensitive data samples for a finding if you're allowed to access the requisite resources and data, and you're allowed to perform the requisite actions. All the requisite actions are [logged in AWS CloudTrail \(p. 312\)](#).

Important

We recommend that you restrict access to this functionality by using custom [IAM policies \(p. 286\)](#). For additional access control, we recommend that you also create a dedicated AWS KMS key for encryption of sensitive data samples that are retrieved, and restrict use of the key to only those principals who must be allowed to retrieve and reveal sensitive data samples.

The topics in this section explain how to configure and use Macie to retrieve and reveal sensitive data samples for findings. You can perform these tasks in all the AWS Regions where Macie is currently available except the Asia Pacific (Osaka) Region.

Topics

- [Configuring Amazon Macie to retrieve and reveal sensitive data samples for findings \(p. 187\)](#)
- [Retrieving and revealing sensitive data samples with Amazon Macie findings \(p. 190\)](#)

Configuring Amazon Macie to retrieve and reveal sensitive data samples for findings

You can optionally configure and use Amazon Macie to retrieve and reveal samples of sensitive data that Macie detects and reports in individual sensitive data findings. The samples can help you verify the nature of the sensitive data that Macie found, and tailor your investigation of an affected Amazon Simple Storage Service (Amazon S3) object and bucket.

When you retrieve and reveal sensitive data samples for a finding, Macie uses data in the corresponding [sensitive data discovery result \(p. 113\)](#) to locate occurrences of sensitive data in the affected S3 object. Macie then extracts samples of those occurrences from the affected object. Macie encrypts the extracted data with an AWS Key Management Service (AWS KMS) key that you specify, temporarily stores the encrypted data in a cache, and returns the data in your results for the finding. Soon after extraction and encryption, Macie permanently deletes the data from the cache unless additional retention is temporarily required to resolve an operational issue.

To retrieve and reveal sensitive data samples for findings, you first need to configure settings for your Macie account: specify the AWS KMS key that you want to use to encrypt the samples, and enable the configuration for your account. You can configure the settings in all the AWS Regions where Macie is currently available except the Asia Pacific (Osaka) Region. If you're the Macie administrator for an organization, the settings apply only to your account. They don't apply to member accounts.

This topic guides you through the process of configuring Macie to retrieve and reveal sensitive data samples. Before you start this process, verify that you [configured a repository for your sensitive data discovery results \(p. 130\)](#). Otherwise, Macie won't be able to locate sensitive data samples that you want to retrieve and reveal.

Tasks

- [Step 1: Verify your permissions \(p. 187\)](#)
- [Step 2: Choose an AWS KMS key \(p. 188\)](#)
- [Step 3: Configure Amazon Macie settings \(p. 188\)](#)

Step 1: Verify your permissions

Before you configure Macie to retrieve and reveal sensitive data samples, verify that you have the permissions that you need. You can do this by using the AWS Identity and Access Management (IAM) console:

1. Sign in to the AWS Management Console and open the IAM console at <https://console.aws.amazon.com/iam/>.
2. In the navigation pane, choose **Users**.
3. Choose your user name.

The **Permissions** tab lists all the IAM policies that are attached to your user name. Choose a policy to show its details. Then compare the information in the policy to the following list of actions that you must be allowed to perform.

Macie

For Macie, verify that you're allowed to perform the following action:

```
macie2:UpdateRevealConfiguration
```

This action allows you to change the configuration settings for retrieving and revealing sensitive data samples. This includes enabling and disabling the configuration.

Optionally verify that you're also allowed to perform the `macie2:GetRevealConfiguration` action. This action allows you to retrieve the current configuration settings and the current status of the configuration.

AWS KMS

If you plan to use the Amazon Macie console to enter the configuration settings, also verify that you're allowed to perform the following AWS KMS actions:

- `kms:DescribeKey`
- `kms:ListAliases`

These actions allow you to retrieve information about the AWS KMS keys for your account. You can then choose one of these keys when you enter the settings.

If you're not allowed to perform the requisite actions, ask your AWS administrator for assistance.

Step 2: Choose an AWS KMS key

After you verify your permissions, determine which AWS KMS key you want to use to encrypt sensitive data samples that you retrieve and reveal. The key can be an existing KMS key from your own account, or an existing KMS key that another account owns. If you want to use a key that another account owns, obtain the Amazon Resource Name (ARN) of the key. You'll need to specify this ARN when you enter the configuration settings in Macie.

The key must be a customer managed, symmetric encryption KMS key. It must also be a single-Region key that's enabled in the same AWS Region as your Macie account. In addition, the key policy for the key must allow the appropriate principals (IAM roles, IAM users, or AWS accounts) to perform the following actions:

- `kms:Decrypt`
- `kms:DescribeKey`
- `kms:GenerateDataKey`

Important

As an additional layer of access control, we recommend that you create a dedicated KMS key for encryption of sensitive data samples that are retrieved, and restrict use of the key to only those principals who must be allowed to retrieve and reveal sensitive data samples. If a user isn't allowed to perform the preceding actions for the key and the user tries to retrieve and reveal samples encrypted with the key, Macie rejects the request and doesn't return any samples for the finding.

For information about creating and reviewing the settings for KMS keys, see [Managing keys](#) in the *AWS Key Management Service Developer Guide*. For information about managing access to KMS keys, see [Key policies in AWS KMS](#) in the *AWS Key Management Service Developer Guide*.

Step 3: Configure Amazon Macie settings

After you verify your permissions and determine which AWS KMS key to use, you're ready to configure the settings for your Macie account. You can do this by using the Amazon Macie console or the Amazon Macie API.

Console

Follow these steps to configure the settings by using the Amazon Macie console.

To configure Macie settings

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. By using the AWS Region selector in the upper-right corner of the page, select the Region in which you want to configure Macie to retrieve and reveal sensitive data samples.
3. In the navigation pane, under **Settings**, choose **Reveal samples**.
4. In the **Settings** section, choose **Edit**.
5. For **Status**, choose **Enable**.
6. Under **Encryption**, specify the AWS KMS key that you want to use to encrypt sensitive data samples:
 - To use a key from your own account, choose **Select a key from your account**. Then, in the **AWS KMS key** list, choose the key to use. The list displays existing, symmetric encryption KMS keys for your account.
 - To use a key that another account owns, choose **Enter the ARN of a key from another account**. Then, in the **AWS KMS key ARN** box, enter the Amazon Resource Name (ARN) of the key to use—for example, `arn:aws:kms:us-east-1:111122223333:key/1234abcd-12ab-34cd-56ef-1234567890ab`.
7. When you finish entering the settings, choose **Save**.

To subsequently change the encryption settings, repeat the preceding steps. To subsequently disable the configuration, repeat steps 1 through 4. Then, for step 5, choose **Disable** for **Status**. When you finish, choose **Save**.

API

To configure the settings programmatically, use the [UpdateRevealConfiguration](#) operation of the Amazon Macie API and specify the appropriate values for the required parameters:

- For the `kmsKeyId` parameter, specify the AWS KMS key that you want to use to encrypt sensitive data samples.

To use a KMS key from your own account, specify the Amazon Resource Name (ARN), ID, or alias for the key. If you specify an alias, include the `alias/` prefix—for example, `alias/ExampleAlias`.

To use a KMS key that another account owns, specify the ARN of the key—for example, `arn:aws:kms:us-east-1:111122223333:key/1234abcd-12ab-34cd-56ef-1234567890ab`. Or specify the ARN of the alias for the key—for example, `arn:aws:kms:us-east-1:111122223333:alias/ExampleAlias`.

- For the `status` parameter, specify `ENABLED` to enable the configuration for your Macie account.

In your request, also ensure that you specify the AWS Region in which you want to use the specified configuration.

To configure the settings by using the [AWS Command Line Interface \(AWS CLI\)](#), run the `update-reveal-configuration` command and specify the appropriate values for the required parameters. For example, if you're running the AWS CLI on Microsoft Windows:

```
C:\> aws macie2 update-reveal-configuration ^
--region us-east-1 ^
--configuration={"kmsKeyId\":\"arn:aws:kms:us-east-1:111122223333:alias/
ExampleAlias\", \"status\":\"ENABLED\"}
```

Where:

- `us-east-1` is the Region in which to use the configuration. In this example, the US East (N. Virginia) Region.
- `arn:aws:kms:us-east-1:111122223333:alias/ExampleAlias` is the ARN of the alias for the KMS key to use. In this example, the key is owned by another account.
- `ENABLED` is the status of the configuration.

The example uses the caret (^) line-continuation character to improve readability.

When you submit your request, Macie tests the settings. If an error occurs, your request fails and Macie returns a message that describes the error.

If your request succeeds, Macie enables the configuration for your account in the specified Region and you receive output similar to the following.

```
{
  "configuration": {
    "kmsKeyId": "arn:aws:kms:us-east-1:111122223333:alias/ExampleAlias",
    "status": "ENABLED"
  }
}
```

Where `kmsKeyId` specifies the KMS key to use to encrypt sensitive data samples that are retrieved and revealed, and `status` is the status of the configuration for your Macie account.

To subsequently check the settings or status of the configuration for your account, use the [GetRevealConfiguration](#) operation or, for the AWS CLI, run the `get-reveal-configuration` command.

To subsequently change the settings or status of the configuration for your account, use the [UpdateRevealConfiguration](#) operation or, for the AWS CLI, run the `update-reveal-configuration` command. In your request, use the supported parameters to specify a new value for each setting that you want to change.

Retrieving and revealing sensitive data samples with Amazon Macie findings

By using Amazon Macie, you can retrieve and reveal samples of sensitive data that Macie detects and reports in individual sensitive data findings. This includes sensitive data that Macie detects using [managed data identifiers](#) (p. 45), and data that matches the criteria of [custom data identifiers](#) (p. 64) that you configure sensitive data discovery jobs to use. The samples can help you verify the nature of the sensitive data that Macie found, and tailor your investigation of the affected Amazon Simple Storage Service (Amazon S3) object and bucket. You can retrieve and reveal sensitive data samples in all the AWS Regions where Macie is currently available except the Asia Pacific (Osaka) Region.

When you retrieve and reveal sensitive data samples for a finding, Macie uses data in the corresponding [sensitive data discovery result](#) (p. 113) to locate the first 1–10 occurrences of sensitive data reported by the finding. Macie then extracts the first 1–128 characters of each occurrence from the affected S3 object. If a finding reports multiple types of sensitive data, Macie does this for up to 100 types of sensitive data reported by the finding.

When Macie extracts sensitive data from an affected S3 object, Macie encrypts the data with an AWS Key Management Service (AWS KMS) key that you specify, temporarily stores the encrypted data in a cache, and returns the data in your results for the finding. Soon after extraction and encryption, Macie

permanently deletes the data from the cache unless additional retention is temporarily required to resolve an operational issue.

If you choose to retrieve and reveal sensitive data samples for a finding again, Macie repeats the process for locating, extracting, encrypting, storing, and ultimately deleting the samples.

For a demonstration of how you can retrieve and reveal sensitive data samples by using the Amazon Macie console, watch the following video: [Amazon Macie One-Click Temporary Retrieval](#).

Topics

- [Before you begin \(p. 191\)](#)
- [Determining whether sensitive data samples are available for findings \(p. 192\)](#)
- [Retrieving and revealing sensitive data samples for findings \(p. 194\)](#)

Before you begin

To retrieve and reveal sensitive data samples for findings, you first need to [configure the settings for your Macie account \(p. 187\)](#). You also need to work with your AWS administrator to verify that you have the permissions that you need.

Each time you retrieve and reveal sensitive data samples, you use your AWS Identity and Access Management (IAM) identity to locate, retrieve, encrypt, and reveal the samples. Macie doesn't use the Macie [service-linked role \(p. 302\)](#) for your account to perform these tasks on your behalf.

This means that you must be allowed to access the following data and resources to retrieve and reveal samples for a finding: the finding; the corresponding sensitive data discovery result for the finding; the affected S3 bucket; and, the affected S3 object. You must also be allowed to use the AWS KMS key that was used to encrypt the affected S3 object, if applicable, and the AWS KMS key that you configured Macie to use for encryption of sensitive data samples. If any IAM policies, resource policies, or other permissions settings deny you the requisite access, an error occurs and Macie doesn't return any samples for the finding.

You must also be allowed to perform the following Macie actions:

- `macie2:GetFindings`
- `macie2:ListFindings`
- `macie2:GetSensitiveDataOccurrences`

The first two actions allow you to retrieve the details of findings. The third action allows you to retrieve sensitive data samples for findings.

To use the Amazon Macie console to retrieve and reveal sensitive data samples, you must also be allowed to perform the following action:

`macie2:GetSensitiveDataOccurrencesAvailability`

This action allows you to determine whether samples are available for individual findings. You don't need permission to perform this action to retrieve and reveal samples programmatically. However, having this permission can streamline your retrieval of samples.

If you're not allowed to perform the requisite actions or access the requisite data and resources, ask your AWS administrator for assistance.

Determining whether sensitive data samples are available for findings

To retrieve and reveal sensitive data samples for a finding, the finding needs to meet certain criteria. It has to include location data for specific occurrences of sensitive data. In addition, it has to specify the location of a valid, corresponding [sensitive data discovery result](#) (p. 113).

The affected S3 object also needs to meet certain criteria. The total storage size of the object can't exceed 10 MB. Also, the MIME type of the object must be one of the following:

- *application/avro*, for an Apache Avro object container (.avro) file
- *application/gzip*, for a GNU Zip compressed archive (.gz or .gzip) file
- *application/json*, for a JSON or JSON Lines (.json or .jsonl) file
- *application/parquet*, for an Apache Parquet (.parquet) file
- *application/vnd.openxmlformats-officedocument.spreadsheetml.sheet*, for a Microsoft Excel workbook (.xlsx) file
- *application/zip*, for a ZIP compressed archive (.zip) file
- *text/csv*, for a CSV (.csv) file
- *text/plain*, for a non-binary text file other than a CSV, JSON, or JSON Lines file

In addition, the contents of the object must be the same as when the finding was created. Macie checks the object's entity tag (ETag) to determine whether it matches the ETag specified by the finding.

If a finding and the affected S3 object meet the preceding criteria, sensitive data samples are available for the finding. You can optionally determine whether this is the case for a particular finding before you try to retrieve and reveal samples for the finding.

To determine whether sensitive data samples are available for a finding

You can use the Amazon Macie console or the Amazon Macie API to determine whether sensitive data samples are available for a finding.

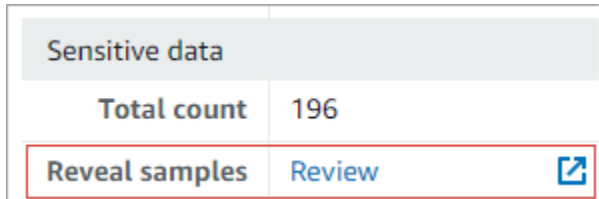
Console


Follow these steps on the Amazon Macie console to determine whether sensitive data samples are available for a finding.

To determine whether samples are available for a finding

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, choose **Findings**.
3. On the **Findings** page, choose the finding. The details panel displays information for the finding.
4. In the details panel, scroll to the **Sensitive data** section. Then refer to the **Reveal samples** field.

If sensitive data samples are available for the finding, a **Review** link appears in the field, as shown in the following image.



Sensitive data	
Total count	196
Reveal samples	Review 

If sensitive data samples aren't available for the finding, the **Reveal samples** field displays text indicating why:

- **Invalid classification result** – Macie can't verify the location of the sensitive data to retrieve. There isn't a corresponding sensitive data discovery result for the finding. Or the corresponding sensitive data discovery result isn't available, is malformed or corrupted, or uses an unsupported storage format. The information in the **Detailed result location** field of the finding can help you investigate the issue. This field specifies the original path to the discovery result in Amazon S3.
- **Object exceeds size quota** – The storage size of the object is larger than 10 MB. It exceeds the size quota for retrieving and revealing samples of sensitive data.
- **Object unavailable** – The object isn't available. It might have been renamed, moved, or deleted. Or the contents of the object changed after Macie created the finding.
- **Unsupported object type** – The MIME type of the object isn't one of the values in the [preceding list \(p. 192\)](#). The object uses a file or storage format that Macie doesn't support for retrieving and revealing samples of sensitive data.

API

To programmatically determine whether sensitive data samples are available for a finding, use the [GetSensitiveDataOccurrencesAvailability](#) operation of the Amazon Macie API. When you submit your request, use the `findingId` parameter to specify the unique identifier for the finding. To obtain this identifier, you can use the [ListFindings](#) operation.

If you're using the [AWS Command Line Interface \(AWS CLI\)](#), run the `get-sensitive-data-occurrences-availability` command and use the `finding-id` parameter to specify the unique identifier for the finding. To obtain this identifier, you can run the `list-findings` command.

If your request succeeds and samples are available for the finding, you receive output similar to the following:

```
{
  "code": "AVAILABLE",
  "reasons": []
}
```

If your request succeeds and samples aren't available for the finding, the value for the code field is UNAVAILABLE and the reasons array specifies why. For example:

```
{
  "code": "UNAVAILABLE",
  "reasons": [
    "UNSUPPORTED_OBJECT_TYPE"
  ]
}
```

Retrieving and revealing sensitive data samples for findings

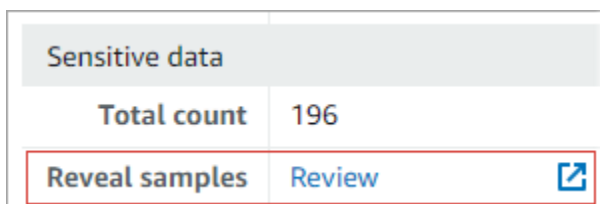
To retrieve and reveal sensitive data samples for a finding, you can use the Amazon Macie console or the Amazon Macie API.


Console

Follow these steps to retrieve and reveal sensitive data samples for a finding by using the Amazon Macie console.

To retrieve and reveal sensitive data samples for a finding

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, choose **Findings**.
3. On the **Findings** page, choose the finding. The details panel displays information for the finding.
4. In the details panel, scroll to the **Sensitive data** section. Then, in the **Reveal samples** field, choose **Review**:



Sensitive data	
Total count	196
Reveal samples	Review 

Note

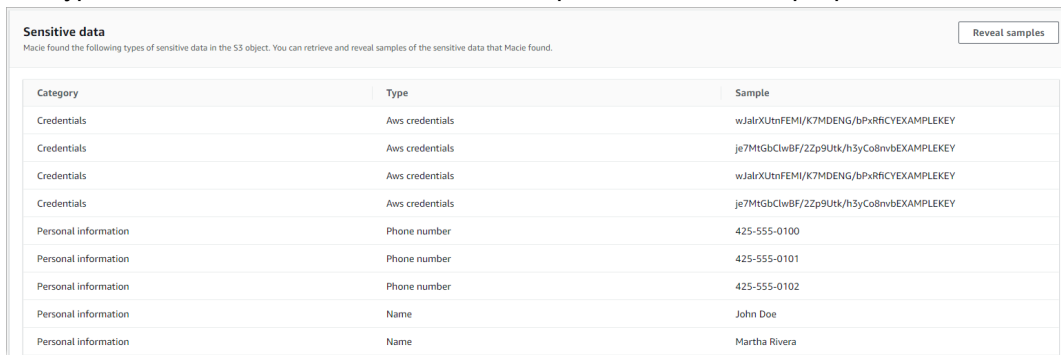
If the **Review** link doesn't appear in the **Reveal samples** field, sensitive data samples aren't available for the finding. For information about why this is the case, see the [preceding section \(p. 192\)](#).

After you choose **Review**, Macie displays a page that summarizes key details of the finding. The details include the categories, types, and number of occurrences of sensitive data that Macie found in the affected S3 object.

5. In the **Sensitive data** section of the page, choose **Reveal samples**.

Macie retrieves and reveals samples of the first 1–10 occurrences of sensitive data reported by the finding. Each sample contains the first 1–128 characters of an occurrence of sensitive data. This can take several minutes.

If the finding reports multiple types of sensitive data, Macie retrieves and reveals samples for up to 100 types. For example, the following image shows samples that span multiple categories and types of sensitive data—AWS credentials, US phone numbers, and people's names.



Category	Type	Sample
Credentials	Aws credentials	wJatRXUtnFEMI/K7MDENG/bP4rRfCYEXAMPLEKEY
Credentials	Aws credentials	je7MfGbcIwBF/2Zp9Utk/h3yCo8mnbEXAMPLEKEY
Credentials	Aws credentials	wJatRXUtnFEMI/K7MDENG/bP4rRfCYEXAMPLEKEY
Credentials	Aws credentials	je7MfGbcIwBF/2Zp9Utk/h3yCo8mnbEXAMPLEKEY
Personal information	Phone number	425-555-0100
Personal information	Phone number	425-555-0101
Personal information	Phone number	425-555-0102
Personal information	Name	John Doe
Personal information	Name	Martha Rivera

The samples are organized first by sensitive data category, and then by sensitive data type.

API

To retrieve and reveal sensitive data samples for a finding programmatically, use the [GetSensitiveDataOccurrences](#) operation of the Amazon Macie API. When you submit your request, use the `findingId` parameter to specify the unique identifier for the finding. To obtain this identifier, you can use the [ListFindings](#) operation.

To retrieve and reveal sensitive data samples by using the AWS CLI, run the [get-sensitive-data-occurrences](#) command and use the `finding-id` parameter to specify the unique identifier for the finding. For example:

```
C:\> aws macie2 get-sensitive-data-occurrences --finding-id  
"1f1c2d74db5d8caa76859ec52example"
```

Where `1f1c2d74db5d8caa76859ec52example` is the unique identifier for the finding. To obtain this identifier by using the AWS CLI, you can run the [list-findings](#) command.

If your request succeeds, Macie begins processing your request and you receive output similar to the following:

```
{  
  "status": "PROCESSING"  
}
```

It can take several minutes to process your request. Within a few minutes, submit your request again.

If Macie can locate, retrieve, and encrypt the sensitive data samples, Macie returns the samples in a `sensitiveDataOccurrences` map. The map specifies 1–100 types of sensitive data reported by the finding and, for each type, 1–10 samples. Each sample contains the first 1–128 characters of an occurrence of sensitive data reported by the finding.

In the map, each key is the ID of the [managed data identifier](#) (p. 45) that detected the sensitive data or the name and unique identifier for the custom data identifier that detected the sensitive data. The values are samples for the specified managed data identifier or custom data identifier. For example, the following response provides three samples of people's names and two samples of AWS secret access keys that were detected by managed data identifiers (`NAME` and `AWS_CREDENTIALS`, respectively).

```
{
  "sensitiveDataOccurrences": {
    "NAME": [
      {
        "value": "Akua Mansa"
      },
      {
        "value": "John Doe"
      },
      {
        "value": "Martha Rivera"
      }
    ],
    "AWS_CREDENTIALS": [
      {
        "value": "wJalrXUtnFEMI/K7MDENG/bPxrFicYEXAMPLEKEY"
      },
      {
        "value": "je7MtGbClwBF/2Zp9Utk/h3yCo8nvbEXAMPLEKEY"
      }
    ]
  },
  "status": "SUCCESS"
}
```

If your request succeeds but sensitive data samples aren't available for the finding, you receive an `UnprocessableEntityException` message that indicates why samples aren't available. For example:

```
{
  "message": "An error occurred (UnprocessableEntityException) when calling the GetSensitiveDataOccurrences operation: OBJECT_UNAVAILABLE"
}
```

In the preceding example, Macie attempted to retrieve samples from the affected S3 object but the object doesn't exist anymore. The object was renamed, moved, or deleted, or the contents of the object changed after Macie created the finding.

If your request succeeds but another type of error prevented Macie from retrieving and revealing sensitive data samples for the finding, you receive output similar to the following:

```
{
  "error": "Macie can't retrieve the samples. You're not allowed to access the affected S3 object or the object is encrypted with a key that you're not allowed to use.",
  "status": "ERROR"
}
```

The value for the `status` field is `ERROR` and the `error` field describes the error that occurred. The information in the preceding topics can help you troubleshoot the error.

JSON schema for sensitive data locations

Amazon Macie uses standardized JSON structures to store information about where it finds sensitive data in Amazon Simple Storage Service (Amazon S3) objects. The structures are used by sensitive data findings and sensitive data discovery results. For sensitive data findings, the structures are part of the JSON schema for Macie findings. To view the complete JSON schema for Macie findings, see [Findings](#) in the *Amazon Macie API Reference*.

The JSON schema for a sensitive data finding includes one `customDataIdentifiers` object and one `sensitiveData` object. The `customDataIdentifiers` object provides details about data that Macie detected using [custom data identifiers \(p. 64\)](#). The `sensitiveData` object provides details about sensitive data that Macie detected using [managed data identifiers \(p. 45\)](#).

Each `customDataIdentifiers` and `sensitiveData` object contains one or more `detections` arrays:

- In a `customDataIdentifiers` object, the `detections` array indicates which custom data identifiers detected the data and produced the finding. For each custom data identifier, the array also indicates the number of occurrences of the data that the identifier detected. It can also indicate the location of the data that the identifier detected.
- In a `sensitiveData` object, a `detections` array indicates the types of sensitive data that Macie detected using managed data identifiers. For each type of sensitive data, the array also indicates the number of occurrences of the data, and it can indicate the location of the data.

For a sensitive data finding, a `detections` array can include 1–15 `occurrences` objects. Each `occurrences` object specifies where Macie found individual occurrences of a specific type of sensitive data.

For example, the following `detections` array indicates the location of three occurrences of sensitive data (US Social Security numbers) in a CSV file.

```
"sensitiveData": [
  {
    "category": "PERSONAL_INFORMATION",
    "detections": [
      {
        "count": 30,
        "occurrences": {
          "cells": [
            {
              "cellReference": null,
              "column": 1,
              "columnName": "SSN",
              "row": 2
            },
            {
              "cellReference": null,
              "column": 1,
              "columnName": "SSN",
              "row": 3
            },
            {
              "cellReference": null,
              "column": 1,
              "columnName": "SSN",
              "row": 4
            }
          ]
        }
      },
      {
        "type": "USA_SOCIAL_SECURITY_NUMBER"
      }
    ]
  }
]
```

The location and number of `occurrences` objects in a `detections` array varies based on the categories, types, and number of occurrences of sensitive data that Macie detects when it runs a sensitive data discovery job. This variation occurs because Macie includes location data for only 1–15 occurrences of the sensitive data that it detects when it runs a job. These 1–15 occurrences are indicative of the categories and types of sensitive data that the affected S3 buckets and objects contain.

An `occurrences` object can contain any the following structures, depending on an S3 object's file type or storage format:

- `cells` array – This array applies to Microsoft Excel workbooks, CSV files, and TSV files. An object in this array specifies a cell or field that contains an occurrence of sensitive data.
- `lineRanges` array – This array applies to non-binary text files other than CSV, JSON, JSON Lines, and TSV files—for example, HTML, TXT, and XML files. An object in this array specifies a line or an inclusive range of lines that contains an occurrence of sensitive data, and the position of the data on the specified line or lines.

In certain cases, an object in a `lineRanges` array specifies the location of sensitive data in a file type or storage format that's supported by another type of array. Those cases are: sensitive data in an unstructured section of an otherwise structured file, such as a comment in a file; sensitive data in a malformed file that Macie analyzes as plaintext; and, a CSV or TSV file that has one or more column names that contain sensitive data.

- `offsetRanges` array – This array is reserved for future use. If this array is present, the value for it is always null.
- `pages` array – This array applies to Adobe Portable Document Format (PDF) files. An object in this array specifies a page that contains an occurrence of sensitive data.
- `records` array – This array applies to Apache Avro object containers, Apache Parquet files, JSON files, and JSON Lines files. For Avro object containers and Parquet files, an object in this array specifies a record index and the path to a field in a record that contains an occurrence of sensitive data. For JSON and JSON Lines files, an object in this array specifies the path to a field or array that contains an occurrence of sensitive data. For JSON Lines files, it also specifies the index of the line that contains the data.

The contents of these arrays vary based on an affected S3 object's file type or storage format and its contents. The next topic provides details and examples of each array.

JSON details and examples for sensitive data locations

Amazon Macie tailors the contents of the JSON structures that it uses to indicate the location of sensitive data in specific types of files and content. The following topics explain and provide examples of these structures.

Topics

- [Cells array \(p. 198\)](#)
- [LineRanges array \(p. 199\)](#)
- [Pages array \(p. 201\)](#)
- [Records array \(p. 201\)](#)

For a complete list of JSON structures that can be included in a sensitive data finding, see [Findings](#) in the *Amazon Macie API Reference*.

Cells array

Applies to: Microsoft Excel workbooks, CSV files, and TSV files

In a `cells` array, a `Cell` object specifies a cell or field that contains an occurrence of sensitive data. The following table describes the purpose of each field in a `Cell` object.

Field	Type	Description
<code>cellReference</code>	String	The location of the cell, as an absolute cell reference, that contains the sensitive data. This field applies only to Excel

Field	Type	Description
		workbooks. This value is null for CSV and TSV files.
column	Integer	The column number of the column that contains the sensitive data. For an Excel workbook, this value correlates to the alphabetical character(s) for a column identifier—for example, 1 for column A, 2 for column B, and so on.
columnName	String	The name of the column that contains the sensitive data, if available.
row	Integer	The row number of the row that contains the sensitive data.

The following example shows the structure of a `Cell` object that reports an occurrence of sensitive data in a CSV file.

```
"cells": [  
  {  
    "cellReference": null,  
    "column": 3,  
    "columnName": "SSN",  
    "row": 5  
  }  
]
```

In the preceding example, the finding indicates that the field in the fifth row of the third column (named *SSN*) of the file contains sensitive data.

The following example shows the structure of a `Cell` object that reports an occurrence of sensitive data in an Excel workbook.

```
"cells": [  
  {  
    "cellReference": "Sheet2!C5",  
    "column": 3,  
    "columnName": "SSN",  
    "row": 5  
  }  
]
```

In the preceding example, the finding indicates that the worksheet named *Sheet2* in the workbook contains sensitive data. In that worksheet, the sensitive data is in the cell in the fifth row of the third column (column C, named *SSN*).

LineRanges array

Applies to: Non-binary text files other than CSV, JSON, JSON Lines, and TSV files—for example, HTML, TXT, and XML files

In a `lineRanges` array, a `Range` object specifies a line or an inclusive range of lines that contains an occurrence of sensitive data, and the position of the data on the specified line or lines.

This object is often empty for file types that are supported by other types of arrays in `occurrences` objects. Exceptions are:

- Data in unstructured sections of an otherwise structured file, such as a comment in a file.
- Data in a malformed file that Macie analyzes as plaintext.
- A CSV or TSV file that has one or more column names that contain sensitive data.

The following table describes the purpose of each field in a `Range` object of a `lineRanges` array.

Field	Type	Description
<code>end</code>	Integer	The number of lines from the beginning of the file to the end of the sensitive data.
<code>start</code>	Integer	The number of lines from the beginning of the file to the beginning of the sensitive data.
<code>startColumn</code>	Integer	The number of characters, with spaces and starting from 1, from the beginning of the first line that contains the sensitive data (<code>start</code>) to the beginning of the sensitive data.

The following example shows the structure of a `Range` object that reports an occurrence of sensitive data that's stored on a single line in a TXT file.

```
"lineRanges": [  
  {  
    "end": 1,  
    "start": 1,  
    "startColumn": 119  
  }  
]
```

In the preceding example, the finding indicates that the first line of the file contains a complete occurrence of sensitive data (a mailing address). The first character in the occurrence is 119 characters (with spaces) from the beginning of that line.

The following example shows the structure of a `Range` object that reports an occurrence of sensitive data that spans multiple lines in a TXT file.

```
"lineRanges": [  
  {  
    "end": 54,  
    "start": 51,  
    "startColumn": 1  
  }  
]
```

In the preceding example, the finding indicates that lines 51 through 54 of the file contain an occurrence of sensitive data (a mailing address). The first character in the occurrence is the first character on line 51 of the file.

Pages array

Applies to: Adobe Portable Document Format (PDF) files

In a `pages` array, a `Page` object specifies a page that contains an occurrence of sensitive data. The object contains a `pageNumber` field. The `pageNumber` field stores an integer that specifies the page number of the page that contains the sensitive data.

The following example shows the structure of a `Page` object that reports an occurrence of sensitive data in a PDF file.

```
"pages": [  
  {  
    "pageNumber": 10  
  }  
]
```

In the preceding example, the finding indicates that page 10 of the file contains sensitive data.

Records array

Applies to: Apache Avro object containers, Apache Parquet files, JSON files, and JSON Lines files

For an Avro object container or a Parquet file, a `Record` object in a `records` array specifies a record index and the path to a field in a record that contains an occurrence of sensitive data. For JSON and JSON Lines files, a `Record` object specifies the path to a field or array that contains an occurrence of sensitive data. For JSON Lines files, it also specifies the index of the line that contains the data.

The following table describes the purpose of each field in a `Record` object.

Field	Type	Description
<code>jsonPath</code>	String	<p>The path, as a <code>JSONPath</code> expression, to the sensitive data.</p> <p>For an Avro object container or a Parquet file, this is the path to the field in the record (<code>recordIndex</code>) that contains the data. For a JSON or JSON Lines file, this is the path to the field or array that contains the data. If the data is a value in an array, the path also indicates which value contains the data.</p> <p>If Macie detects sensitive data in the name of any element in the path, Macie omits the <code>jsonPath</code> field from a <code>Record</code> object. If the name of a path element exceeds 20 characters, Macie truncates the name by removing characters from the beginning of the name. If the resulting full path exceeds 250 characters, Macie also truncates the path,</p>

Field	Type	Description
		starting with the first element in the path, until the path contains 250 or fewer characters.
recordIndex	Integer	For an Avro object container or a Parquet file, the record index, starting from 0, for the record that contains the sensitive data. For a JSON Lines file, the line index, starting from 0, for the line that contains the sensitive data. This value is always 0 for JSON files.

The following example shows the structure of a `Record` object that reports an occurrence of sensitive data in a Parquet file. In this example, Macie truncated the name of the field that contains the data, specified in the `jsonPath` field, to meet the character limit.

```
"records": [  
  {  
    "jsonPath": "$['...hijklmnopqrstuvwxyz']",  
    "recordIndex": 7663  
  }  
]
```

In the preceding example, the finding indicates that the record of index 7663 (record number 7664) contains sensitive data. In that record, the sensitive data is in the field whose name ends with `hijklmnopqrstuvwxyz`. The full JSON path to the field in the record is `$.abcdefghijklmnopqrstuvwxy`.

The following example also shows the structure of a `Record` object that reports an occurrence of sensitive data in a Parquet file. In this example, Macie truncated both the full path and the name of the field that contains the data.

```
"records": [  
  {  
    "jsonPath":  
    "$.usssn2.usssn3.usssn4.usssn5.usssn6.usssnfield7.usssn8.usssn9.usssn10.usssn11.usssn12.usssn13.usssn14.usssn15.usssn16.usssn17.usssn18.usssn19.usssn20.usssn21.usssn22.usssn23.usssn24.usssn25.usssn26.usssn27.usssn28.usssn29[ 'abcdefghijklmnopqrstuvwxy ' ]",  
    "recordIndex": 2335  
  }  
]
```

In the preceding example, the finding indicates that the record of index 2335 (record number 2336) contains sensitive data. In that record, the sensitive data is in the field whose name ends with `hijklmnopqrstuvwxyz`. The full JSON path to the field in the record is: `$['1234567890']usssn1.usssn2.usssn3.usssn4.usssn5.usssn6.usssnfield7.usssn8.usssn9.usssn10.usssn11.usssn12.usssn13.usssn14.usssn15.usssn16.usssn17.usssn18.usssn19.usssn20.usssn21.usssn22.usssn23.usssn24.usssn25.usssn26.usssn27.usssn28.usssn29['abcdefghijklmnopqrstuvwxy ']`

The following example shows the structure of a `Record` object that reports an occurrence of sensitive data in a JSON file. In this example, the sensitive data is a specific value in an array.

```
"records": [  
  {  
    "jsonPath": "$.access.key[2]",  
  }  
]
```



```
    "recordIndex": 0  
  }  
]
```

In the preceding example, the finding indicates that the second value in an array named `key` contains sensitive data. The array is a child of an object named `access`.

The following example shows the structure of a `Record` object that reports an occurrence of sensitive data in a JSON Lines file.

```
"records": [  
  {  
    "jsonPath": "$.access.key",  
    "recordIndex": 3  
  }  
]
```

In the preceding example, the finding indicates that the third value (line) in the file contains sensitive data. In that line, the sensitive data is in a field named `key`, which is a child of an object named `access`.

Suppressing Amazon Macie findings

To streamline your analysis of findings, you can create and use *suppression rules*. A *suppression rule* is a set of attribute-based filter criteria that defines cases where you want Amazon Macie to archive findings automatically. Suppression rules are helpful in situations where you've reviewed a class of findings and don't want to be notified of them again.

For example, you might decide to allow S3 buckets to contain mailing addresses, if the buckets don't allow public access and they encrypt new objects by default. In this case, you can create a suppression rule that specifies filter criteria for the following fields: **Sensitive data detection type**, **S3 bucket public access permission**, and **S3 bucket default encryption**. The rule suppresses future findings that match the filter criteria.

If you suppress findings with a suppression rule, Macie continues to generate findings for subsequent occurrences of sensitive data and potential policy violations that match the rule's criteria. However, Macie automatically changes the status of the findings to *archived*. This means that the findings don't appear by default on the Amazon Macie console, but they persist in Macie until they expire. Macie stores findings for 90 days.

In addition to changing the status of suppressed findings, Macie doesn't publish the findings to Amazon EventBridge as events or to AWS Security Hub. Macie does, however, continue to create and store [sensitive data discovery results \(p. 113\)](#) that correlate to sensitive data findings that you suppress. This helps ensure that you have an immutable history of sensitive data findings for data privacy and protection audits or investigations that you perform.

Note

If your account is part of an organization that centrally manages multiple Macie accounts, suppression rules might work differently for your account. This depends on the category of findings that you want to suppress, and whether you have a Macie administrator or member account:

- **Policy findings** – Only a Macie administrator can suppress policy findings for the organization's accounts.

If you have a Macie administrator account and you create a suppression rule, Macie applies the rule to policy findings for all the accounts in your organization unless you configure the rule

to exclude specific accounts. If you have a Macie member account and you want to suppress policy findings for your account, contact your Macie administrator.

- **Sensitive data findings** – A Macie administrator and individual members can suppress sensitive data findings that their sensitive data discovery jobs produce.

If you have a Macie administrator or member account and you create a suppression rule, Macie applies the rule to sensitive data findings for your account. Only the account that creates a sensitive data discovery job can suppress or otherwise access any sensitive data findings that the job produces.

For more information about the tasks that administrators and members can perform, see [Understanding the relationship between Amazon Macie administrator and member accounts \(p. 238\)](#).

To create and manage suppression rules, you can use the Amazon Macie console or the Amazon Macie API. The following topics explain how. For the API, the topics include examples of how to perform these tasks using the [AWS Command Line Interface \(AWS CLI\)](#). You can also perform these tasks by using a current version of another AWS command line tool or an AWS SDK, or by sending HTTPS requests directly to Macie. For information about AWS tools and SDKs, see [Tools to Build on AWS](#).

Topics

- [Creating suppression rules \(p. 204\)](#)
- [Reviewing suppressed findings \(p. 206\)](#)
- [Changing suppression rules \(p. 206\)](#)
- [Deleting suppression rules \(p. 208\)](#)

Creating suppression rules

Before you create a suppression rule, it's important to note that you can't restore (unarchive) findings that you suppress using a suppression rule. You can, however, [review suppressed findings \(p. 206\)](#) on the Amazon Macie console and access suppressed findings with the Amazon Macie API.

When you create a suppression rule, you specify filter criteria, a name, and, optionally, a description of the rule. You can create a suppression rule by using the Amazon Macie console or the Amazon Macie API.

Console

Follow these steps to create a suppression rule by using the Amazon Macie console.

To create a suppression rule


1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, choose **Findings**.

Tip

To use an existing suppression or filter rule as a starting point, choose the rule from the **Saved rules** list.

You can also streamline creation of a rule by first pivoting and drilling down on findings by a predefined logical group. If you do this, Macie automatically creates and applies the appropriate filter conditions, which can be a helpful starting point for creating a rule. To do this, choose **By bucket**, **By type**, or **By job** in the navigation pane (under **Findings**), and then choose an item in the table. In the details panel, choose the link for the field to pivot on.

3. In the filter bar, add filter conditions that specify attributes of the findings that you want the rule to suppress.



To learn how to add filter conditions, see [Creating and applying filters to findings \(p. 154\)](#).

4. When you finish adding filter conditions for the rule, choose **Suppress findings** above the filter bar.
5. Under **Suppression rule**, enter a name and, optionally, a description of the rule.
6. Choose **Save**.

API

To create a suppression rule programmatically, use the [CreateFindingsFilter](#) operation of the Amazon Macie API and specify the appropriate values for the required parameters:

- For the `action` parameter, specify `ARCHIVE` to ensure that Macie suppresses findings that match the criteria of the rule.
- For the `criterion` parameter, specify a map of conditions that define the filter criteria for the rule.

In the map, each condition should specify a field, an operator, and one or more values for the field. The type and number of values depends on the field and operator that you choose. For information about the fields, operators, and types of values that you can use in a condition, see [Fields for filtering findings \(p. 165\)](#), [Using operators in conditions \(p. 151\)](#), and [Specifying values for fields \(p. 149\)](#).

To create a suppression rule by using the AWS CLI, run the `create-findings-filter` command and specify the appropriate values for the required parameters. The following examples create a suppression rule that returns all sensitive data findings that are in the current AWS Region and report occurrences of mailing addresses (and no other types of sensitive data) in S3 objects.

This example is formatted for Linux, macOS, or Unix, and it uses the backslash (`\`) line-continuation character to improve readability.

```
$ aws macie2 create-findings-filter \
--action ARCHIVE \
--name my_suppression_rule \
--finding-criteria '{"criterion":
{"classificationDetails.result.sensitiveData.detections.type":{"eqExactMatch":
["ADDRESS"]}}}'
```

This example is formatted for Microsoft Windows and it uses the caret (`^`) line-continuation character to improve readability.

```
C:\> aws macie2 create-findings-filter ^
--action ARCHIVE ^
--name my_suppression_rule ^
--finding-criteria={"criterion":
{"classificationDetails.result.sensitiveData.detections.type":{"eqExactMatch":
["ADDRESS"]}}}
```

Where:

- `my_suppression_rule` is the custom name for the rule.
- `criterion` is a map of filter conditions for the rule:
 - `classificationDetails.result.sensitiveData.detections.type` is the JSON name of the **Sensitive data detection type** field.
 - `eqExactMatch` specifies the *equals exact match* operator.
 - `ADDRESS` is an enumerated value for the **Sensitive data detection type** field.

If the command runs successfully, you receive output similar to the following.

```
{
  "arn": "arn:aws:macie2:us-west-2:123456789012:findings-filter/8a3c5608-aa2f-4940-
b347-d1451example",
  "id": "8a3c5608-aa2f-4940-b347-d1451example"
}
```

Where `arn` is the Amazon Resource Name (ARN) of the suppression rule that was created, and `id` is the unique identifier for the rule.

For additional examples of filter criteria, see [Filtering findings programmatically with the Amazon Macie API \(p. 156\)](#).

Reviewing suppressed findings

By default, Macie doesn't display suppressed findings on the Amazon Macie console. However, you can review these findings on the console by changing your filter settings.

To review suppressed findings on the console

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, choose **Findings**. The **Findings** page displays findings that Macie created or updated for your account in the current AWS Region during the past 90 days. By default, this doesn't include findings that were suppressed by a suppression rule.
3. In the filter bar, do one of the following:
 - To display only suppressed findings, choose **Current**, and then choose **Archived**.
 - To display both suppressed and current findings, choose **Current**, and then choose **All**.

You can also access suppressed findings by using the Amazon Macie API. To retrieve a list of suppressed findings, use the [ListFindings](#) operation and include a filter condition that specifies `true` for the `archived` field. For an example of how to do this using the AWS CLI, see [Filtering findings programmatically \(p. 156\)](#). To then retrieve the details of one or more suppressed findings, use the [GetFindings](#) operation and specify the unique identifier for each finding to retrieve.

Changing suppression rules

You can change the settings for a suppression rule at any time by using the Amazon Macie console or the Amazon Macie API. You can also assign and manage tags for the rule.


A *tag* is a label that you define and assign to certain types of AWS resources. Each tag consists of a required tag key and an optional tag value. Tags can help you identify, categorize, and manage resources

in different ways, such as by purpose, owner, environment, or other criteria. To learn more, see [Tagging Amazon Macie resources \(p. 315\)](#).

Console

Follow these steps to change the settings for an existing suppression rule by using the Amazon Macie console.

To change a suppression rule

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, choose **Findings**.
3. In the **Saved rules** list, choose the edit icon () next to the suppression rule that you want to change.
4. Do any of the following:
 - To change the filter criteria of the rule, use the filter bar to enter conditions that specify attributes of the findings that you want the rule to suppress. To learn how, see [Creating and applying filters to findings \(p. 154\)](#).
 - To change the name of the rule, enter a new name in the **Name** box under **Suppression rule**.
 - To change the description of the rule, enter a new description in the **Description** box under **Suppression rule**.
 - To assign, review, or edit tags for the rule, choose **Manage tags** under **Suppression rule**. Then review and change the tags as necessary. A rule can have as many as 50 tags.
5. When you finish making changes, choose **Save**.

API

To change a suppression rule programmatically, use the [UpdateFindingsFilter](#) operation of the Amazon Macie API. When you submit your request, use the supported parameters to specify a new value for each setting that you want to change.

For the `id` parameter, specify the unique identifier for the rule to change. You can get this identifier by using the [ListFindingsFilter](#) operation to retrieve a list of suppression and filter rules for your account. If you're using the AWS CLI, run the `list-findings-filters` command to retrieve this list.

To change a suppression rule by using the AWS CLI, run the `update-findings-filter` command and use the supported parameters to specify a new value for each setting that you want to change. For example, the following command changes the name of an existing suppression rule.

```
C:\> aws macie2 update-findings-filter --id 8a3c5608-aa2f-4940-b347-d1451example --name mailing_addresses_only
```

Where:

- `8a3c5608-aa2f-4940-b347-d1451example` is the unique identifier for the rule.
- `mailing_addresses_only` is the new name for the rule.

If the command runs successfully, you receive output similar to the following.

```
{
  "arn": "arn:aws:macie2:us-west-2:123456789012:findings-filter/8a3c5608-aa2f-4940-b347-d1451example",
  "id": "8a3c5608-aa2f-4940-b347-d1451example"
}
```

Where `arn` is the Amazon Resource Name (ARN) of the rule that was changed, and `id` is the unique identifier for the rule.

Similarly, the following example converts a filter rule to a suppression rule by changing the value for the `action` parameter from `NOOP` to `ARCHIVE`.

```
C:\> aws macie2 update-findings-filter --id 8a1c3508-aa2f-4940-b347-d1451example --  
action ARCHIVE
```

Where:

- `8a1c3508-aa2f-4940-b347-d1451example` is the unique identifier for the rule.
- `ARCHIVE` is the new action for Macie to perform on findings that match the criteria of the rule—suppress the findings.

If the command runs successfully, you receive output similar to the following:

```
{  
  "arn": "arn:aws:macie2:us-west-2:123456789012:findings-filter/8a1c3508-aa2f-4940-  
b347-d1451example",  
  "id": "8a1c3508-aa2f-4940-b347-d1451example"  
}
```

Where `arn` is the Amazon Resource Name (ARN) of the rule that was changed, and `id` is the unique identifier for the rule.

Deleting suppression rules


You can delete a suppression rule at any time by using the Amazon Macie console or the Amazon Macie API. If you delete a suppression rule, Macie stops suppressing new and subsequent occurrences of findings that match the criteria of the rule and aren't suppressed by other rules. Note, however, that Macie might continue to suppress findings that it's currently processing and match the rule's criteria.

After you delete a suppression rule, new and subsequent occurrences of findings that match the rule's criteria have a status of *current*. This means that they appear by default on the Amazon Macie console. In addition, Macie publishes these findings to Amazon EventBridge as events. Depending on the [publication settings \(p. 227\)](#) for your account, Macie also publishes the findings to AWS Security Hub.

Console

Follow these steps to delete a suppression rule by using the Amazon Macie console.

To delete a suppression rule

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, choose **Findings**.
3. In the **Saved rules** list, choose the edit icon () next to the suppression rule that you want to delete.
4. Under **Suppression rule**, choose **Delete**.

API

To delete a suppression rule programmatically, use the [DeleteFindingsFilter](#) operation of the Amazon Macie API. For the `id` parameter, specify the unique identifier for the suppression rule

to delete. You can get this identifier by using the [ListFindingsFilter](#) operation to retrieve a list of suppression and filter rules for your account. If you're using the AWS CLI, run the [list-findings-filters](#) command to retrieve this list.

To delete a suppression rule by using the AWS CLI, run the [delete-findings-filter](#) command. For example:

```
C:\> aws macie2 delete-findings-filter --id 8a3c5608-aa2f-4940-b347-d1451example
```

Where *8a3c5608-aa2f-4940-b347-d1451example* is the unique identifier for the suppression rule to delete.

If the command runs successfully, Macie returns an empty HTTP 200 response. Otherwise, Macie returns an HTTP 4xx or 500 response that indicates why the operation failed.

Severity scoring for Amazon Macie findings

When Amazon Macie generates a policy or sensitive data finding, it automatically assigns a severity to the finding. A finding's severity reflects the principal characteristics of the finding and can help you assess and prioritize your findings. A finding's severity doesn't imply or otherwise indicate the criticality or importance that an affected resource might have for your organization.

For policy findings, severity is based on the nature of a potential issue with the security or privacy of your Amazon Simple Storage Service (Amazon S3) data. For sensitive data findings, severity is based on the nature and number of occurrences of sensitive data that Macie found in an S3 object.

In Macie, a finding's severity is represented in two ways.

Severity level

This is a qualitative representation of severity. Severity levels range from *Low*, for least severe, to *High*, for most severe.

Severity levels appear directly on the Amazon Macie console. They're also available in JSON representations of findings on the Macie console, from the Amazon Macie API, and in sensitive data discovery results that correlate to sensitive data findings. Severity levels are also included in finding events that Macie publishes to Amazon EventBridge and findings that Macie publishes to AWS Security Hub.

Severity score

This is a numerical representation of severity. Severity scores range from 1 through 3 and map directly to severity levels:

Severity score	Severity level
1	Low
2	Medium
3	High

Severity scores don't appear directly on the Amazon Macie console. However, they're available in JSON representations of findings on the Macie console, from the Amazon Macie API, and in sensitive

data discovery results that correlate to sensitive data findings. Severity scores are also included in finding events that Macie publishes to Amazon EventBridge. They aren't included in findings that Macie publishes to AWS Security Hub.

The topics in this section indicate how Macie determines the severity of policy findings and sensitive data findings.

Topics

- [Severity scoring for policy findings \(p. 210\)](#)
- [Severity scoring for sensitive data findings \(p. 210\)](#)

Severity scoring for policy findings

The severity of a policy finding is based on the nature of a potential issue with the security or privacy of an S3 bucket. The following table lists the severity levels that Macie assigns to each type of policy finding. For a description of each type, see [Types of findings \(p. 141\)](#).

Finding type	Severity level
Policy:IAMUser/S3BlockPublicAccessDisabled	High
Policy:IAMUser/S3BucketEncryptionDisabled	Low
Policy:IAMUser/S3BucketPublic	High
Policy:IAMUser/S3BucketReplicatedExternally	High
Policy:IAMUser/S3BucketSharedExternally	High

The severity of a policy finding doesn't change based on the number of occurrences of the finding.

Severity scoring for sensitive data findings

The severity of a sensitive data finding is based on the nature and number of occurrences of sensitive data that Macie found in an S3 object. The following topics indicate how Macie determines the severity of each type of sensitive data finding:

- [SensitiveData:S3Object/Credentials \(p. 210\)](#)
- [SensitiveData:S3Object/CustomIdentifier \(p. 211\)](#)
- [SensitiveData:S3Object/Financial \(p. 211\)](#)
- [SensitiveData:S3Object/Personal \(p. 212\)](#)
- [SensitiveData:S3Object/Multiple \(p. 214\)](#)

For detailed information about the types of sensitive data that Macie can detect and report in sensitive data findings, see [Using managed data identifiers \(p. 45\)](#) and [Building custom data identifiers \(p. 64\)](#).

SensitiveData:S3Object/Credentials

A **SensitiveData:S3Object/Credentials** finding indicates that an S3 object contains credentials data. For this type of finding, Macie determines severity based on the type and number of occurrences of the credentials data that Macie found in the object.

The following table indicates the severity levels that Macie assigns to findings that report occurrences of credentials data in an S3 object.

Detection type	1 occurrence	2–99 occurrences	100 or more occurrences
AWS secret access key	High	High	High
HTTP Basic Authorization header	High	High	High
JSON Web Token (JWT)	High	High	High
OpenSSH private key	High	High	High
PGP private key	High	High	High
Public-Key Cryptography Standard (PKCS) private key	High	High	High
PuTTY private key	High	High	High

SensitiveData:S3Object/CustomIdentifier

A **SensitiveData:S3Object/CustomIdentifier** finding indicates that an S3 object contains text that matches the detection criteria of one or more custom data identifiers. The object might contain more than one type of sensitive data.

By default, Macie assigns the **Medium** severity level to this type of finding—if the S3 object contains at least one occurrence of text that matches the detection criteria of at least one custom data identifier, Macie automatically assigns the **Medium** severity level to the finding. The severity of the finding doesn't change based on the number of occurrences of text that matches a custom data identifier's criteria.

However, the severity of this type of finding can vary if you defined custom severity settings for a custom data identifier that produced the finding. If this is the case, Macie determines severity as follows:

- If the S3 object contains text that matches the detection criteria of only one custom data identifier, Macie determines the finding's severity based on the severity settings for that identifier.
- If the S3 object contains text that matches the detection criteria of more than one custom data identifier, Macie determines the finding's severity by evaluating the severity settings for each custom data identifier, determining which of those settings produces the highest severity, and then assigning that highest severity to the finding.

To review the severity settings for a custom data identifier, choose **Custom data identifiers** in the navigation pane on the Amazon Macie console. Then choose the name of the custom data identifier. The **Severity** section shows the settings. For more information, see [Defining finding severity settings for custom data identifiers \(p. 67\)](#).

SensitiveData:S3Object/Financial

A **SensitiveData:S3Object/Financial** finding indicates that an S3 object contains financial information. For this type of finding, Macie determines severity based on the type and number of occurrences of the financial information that Macie found in the object.

The following table indicates the severity levels that Macie assigns to findings that report occurrences of financial information in an S3 object.

Detection type	1 occurrence	2–99 occurrences	100 or more occurrences
Bank account number	High	High	High
Credit card expiration date	Low	Medium	High
Credit card magnetic strip data	High	High	High
Credit card number*	High	High	High
Credit card verification code	Medium	High	High

* Severity levels are the same for credit card numbers that are or aren't in proximity of a keyword.

If a finding reports multiple types of financial information in an object, Macie determines the finding's severity by calculating the severity for each type of financial information that Macie found, determining which type produces the highest severity, and assigning that highest severity to the finding. For example, if Macie detects 10 credit card expiration dates (**Medium** severity level) and 10 credit card numbers (**High** severity level) in an object, Macie assigns a **High** severity level to the finding.

SensitiveData:S3Object/Personal

A **SensitiveData:S3Object/Personal** finding indicates that an S3 object contains personal information—personal health information (PHI), personally identifiable information (PII), or a combination of the two. For this type of finding, Macie determines severity based on the type and number of occurrences of the personal information that Macie found in the object.

The following table indicates the severity levels that Macie assigns to sensitive data findings that report occurrences of PHI in an S3 object.

Detection type	1 occurrence	2–99 occurrences	100 or more occurrences
Drug Enforcement Agency (DEA) Registration Number	High	High	High
Health Insurance Claim Number (HICN)	High	High	High
Health insurance or medical identification number	High	High	High
Healthcare Common Procedure Coding System (HCPCS) code	High	High	High
National Drug Code (NDC)	High	High	High
National Provider Identifier (NPI)	High	High	High

Amazon Macie User Guide
Severity scoring for sensitive data findings

Detection type	1 occurrence	2–99 occurrences	100 or more occurrences
Unique device identifier (UDI)	Low	Medium	High

The following table indicates the severity levels that Macie assigns to sensitive data findings that report occurrences of PII in an S3 object.

Detection type	1 occurrence	2–99 occurrences	100 or more occurrences
Birth date	Low	Medium	High
Driver's license identification number	Low	Medium	High
Electoral roll number	High	High	High
Full name	Low	Medium	High
Global Positioning System (GPS) coordinates	Low	Medium	Medium
HTTP cookie	Low	Medium	High
Mailing address	Low	Medium	High
National identification number	High	High	High
National Insurance Number (NINO)	High	High	High
Passport number	Medium	High	High
Permanent residence number	High	High	High
Phone number	Low	Medium	High
Social Insurance Number (SIN)	High	High	High
Social Security number (SSN)	High	High	High
Taxpayer identification or reference number	High	High	High
Vehicle identification number (VIN)	Low	Low	Medium

If a finding reports multiple types of PHI, PII, or both PHI and PII in an object, Macie determines the finding's severity by calculating the severity for each detection type, determining which detection type produces the highest severity, and assigning that highest severity to the finding.

For example, if Macie detects 10 full names (**Medium** severity level) and 5 passport numbers (**High** severity level) in an object, Macie assigns a **High** severity level to the finding. Similarly, if Macie detects 10 full names (**Medium** severity level) and 10 health insurance identification numbers (**High** severity level) in an object, Macie assigns a **High** severity level to the finding.

SensitiveData:S3Object/Multiple

A **SensitiveData:S3Object/Multiple** finding indicates that an S3 object contains data spanning multiple sensitive data categories—any combination of credentials, financial information, personal information, or text that matches the detection criteria of one or more custom data identifiers.

For this type of finding, Macie determines severity by calculating the severity for each type of sensitive data that Macie found (as indicated in the preceding topics), determining which type produces the highest severity, and assigning that highest severity to the finding.

For example, if Macie detects 10 full names (**Medium** severity level) and 10 AWS secret access keys (**High** severity level) in an object, Macie assigns a **High** severity level to the finding.

Monitoring and processing Amazon Macie findings

To support integration with other applications, services, and systems, such as monitoring or event management systems, Amazon Macie automatically publishes policy and sensitive data findings to Amazon EventBridge as events. For additional support, you can configure Macie to also publish policy and sensitive data findings to AWS Security Hub.

Amazon EventBridge, formerly Amazon CloudWatch Events, is a serverless event bus service that delivers a stream of real-time data from applications and services, and routes that data to targets such as AWS Lambda functions, Amazon Simple Notification Service topics, and Amazon Kinesis streams. With EventBridge, you can automate monitoring and processing of certain types of events, including events that Macie publishes for findings. To learn more about EventBridge, see the [Amazon EventBridge User Guide](#). To learn about using EventBridge to monitor and process findings, see [EventBridge integration \(p. 215\)](#).

AWS Security Hub is a security service that provides you with a comprehensive view of your security state in AWS and helps you check your environment against security industry standards and best practices. Security Hub collects security data from multiple AWS services and supported AWS Partner Network security solutions, and it helps you analyze your security trends and identify the highest priority security issues. With Security Hub, you can analyze Macie findings as part of a broader analysis of your organization's security posture. To learn more about Security Hub, see the [AWS Security Hub User Guide](#). To learn about using Security Hub to monitor and process findings, see [Security Hub integration \(p. 219\)](#).

When Macie creates a finding, it automatically publishes the finding to EventBridge as a new event. Depending on the publication settings that you choose for your account, Macie can also publish the finding to Security Hub. Macie publishes each new finding immediately after it finishes processing the finding. If Macie detects a subsequent occurrence of an existing policy finding, it publishes an update to the existing EventBridge event for the finding. Depending on your publication settings, Macie can also publish the update to Security Hub. Macie publishes these updates on a recurring basis, using a publication frequency that you specify in the publication settings for your account. For details about these settings and the timing with which Macie publishes findings, see [Configuring publication settings for findings \(p. 227\)](#).

Topics

- [Amazon Macie integration with Amazon EventBridge \(p. 215\)](#)
- [Amazon Macie integration with AWS Security Hub \(p. 219\)](#)
- [Configuring publication settings for Amazon Macie findings \(p. 227\)](#)
- [Amazon EventBridge event schema for Amazon Macie findings \(p. 229\)](#)

Amazon Macie integration with Amazon EventBridge

Amazon EventBridge, formerly Amazon CloudWatch Events, is a serverless event bus service. EventBridge delivers a stream of real-time data from applications and services, and routes that data to targets such as AWS Lambda functions, Amazon Simple Notification Service (Amazon SNS) topics, and Amazon Kinesis streams. To learn more about EventBridge, see the [Amazon EventBridge User Guide](#).

With EventBridge, you can automate monitoring and processing of certain types of events. This includes events that Amazon Macie publishes automatically for new policy findings and sensitive data findings. This also includes events that Macie publishes automatically for subsequent occurrences of existing policy findings. For details about how and when Macie publishes these events, see [Configuring publication settings for findings \(p. 227\)](#).

By using EventBridge and the events that Macie publishes for findings, you can monitor and process findings in near-real time. You can then act upon findings by using other applications and services. For example, you might use EventBridge to send specific types of new findings to an AWS Lambda function. The Lambda function might then process and send the data to your security incident and event management (SIEM) system.

In addition to automated monitoring and processing, use of EventBridge enables longer-term retention of your findings data. Macie stores findings for 90 days. With EventBridge, you can send findings data to your preferred storage platform and store the data for as long as you like.

Note

For long-term retention, also configure Macie to store your sensitive data discovery results in an S3 bucket. To learn more, see [Storing and retaining sensitive data discovery results \(p. 130\)](#).

Topics

- [Using Amazon EventBridge \(p. 216\)](#)
- [Creating Amazon EventBridge rules for finding events \(p. 216\)](#)

Using Amazon EventBridge

With Amazon EventBridge, you create rules to specify which events you want to monitor and which targets you want to perform automated actions for those events. A *target* is a destination that EventBridge sends events to.

To automate monitoring and processing tasks for findings, you can create an EventBridge rule that automatically detects Amazon Macie finding events and sends those events to another application or service for processing or other action. You can tailor the rule to send only those events that meet certain criteria. To do this, specify criteria that derive from the [Amazon EventBridge event schema for Amazon Macie findings \(p. 229\)](#).

For example, you can create a rule that sends specific types of new findings to an AWS Lambda function. The Lambda function can then perform tasks such as: process and send the data to your SIEM system; automatically apply encryption to an S3 object; or, restrict access to an object by changing the object's access control list (ACL). Or you can create a rule that automatically sends new high-severity findings to an Amazon SNS topic, which then notifies your incident response team of the finding.

In addition to invoking Lambda functions and notifying Amazon SNS topics, EventBridge supports other types of targets and actions, such as relaying events to Amazon Kinesis streams, activating AWS Step Functions state machines, and invoking the AWS Systems Manager run command. For information about supported targets, see [Amazon EventBridge targets](#) in the *Amazon EventBridge User Guide*.

Creating Amazon EventBridge rules for finding events

The following procedures explain how to use the Amazon EventBridge console and the [AWS Command Line Interface \(AWS CLI\)](#) to create an EventBridge rule for Amazon Macie findings. The rule detects events that use the event schema and pattern for Macie findings, and sends those events to an AWS Lambda function for processing.

AWS Lambda is a compute service that you can use to run code without provisioning or managing servers. You package your code and upload it to AWS Lambda as a *Lambda function*. AWS Lambda then runs the function when the function is invoked. A function can be invoked manually by you,

automatically in response to events, or in response to requests from applications or services. For information about creating and invoking Lambda functions, see the [AWS Lambda Developer Guide](#).

Console

This procedure explains how to use the Amazon EventBridge console to create a rule that automatically sends all Macie finding events to a Lambda function for processing. The rule uses default settings for rules that run when specific events are received. For details about rule settings or to learn how to create a rule that uses custom settings, see [Creating rules that react to events](#) in the *Amazon EventBridge User Guide*.

Tip

You can also create a rule that uses a custom pattern to detect and act upon only a subset of Macie finding events. This subset can be based on specific fields that Macie includes in a finding event. To learn about the available fields, see [EventBridge event schema for findings \(p. 229\)](#). To learn how to create this type of rule, see [Content filtering in event patterns](#) in the *Amazon EventBridge User Guide*.

Before you create this rule, create the Lambda function that you want the rule to use as a target. When you create the rule, you'll need to specify this function as the target for the rule.

To create an event rule by using the console

1. Open the Amazon EventBridge console at <https://console.aws.amazon.com/events/>.
2. In the navigation pane, under **Events**, choose **Rules**.
3. In the **Rules** section, choose **Create rule**.
4. On the **Define rule detail** page, do the following:
 - For **Name**, enter a name for the rule.
 - (Optional) For **Description**, enter a brief description of the rule.
 - For **Event bus**, ensure that **default** is selected and **Enable the rule on the selected event bus** is turned on.
 - For **Rule type**, choose **Rule with an event pattern**.
5. When you finish, choose **Next**.
6. On the **Build event pattern** page, do the following:
 - For **Event source**, choose **AWS events or EventBridge partner**.
 - (Optional) For **Sample event**, review a sample finding event for Macie to learn what an event might contain. To do this, choose **AWS events**. Then, for **Sample events**, choose **Macie Finding**.
 - For **Event pattern**, choose **Event pattern form**. Then enter the following settings:
 - For **Event source**, choose **AWS services**.
 - For **AWS service**, enter **Macie**.
 - For **Event type**, enter **Macie Finding**.
7. When you finish, choose **Next**.
8. On the **Select targets** page, do the following:
 - For **Target types**, choose **AWS service**.
 - For **Select a target**, enter **Lambda function**. Then, for **Function**, choose the Lambda function that you want to send finding events to.
 - For **Configure version/alias**, enter version and alias settings for the target Lambda function.
 - (Optional) For **Additional settings**, enter custom settings to specify which event data you want to send to the Lambda function. You can also specify how to handle events that aren't delivered to the function successfully.

9. When you finish, choose **Next**.
10. On the **Configure tags** page, optionally enter one or more tags to assign to the rule. Then choose **Next**.
11. On the **Review and create** page, review the rule's settings and verify that they're correct.

To change a setting, choose **Edit** in the section that contains the setting, and then enter the correct setting. You can also use the navigation tabs to go to the page that contains a setting.

12. When you finish verifying the settings, choose **Create rule**.

AWS CLI

This procedure explains how to use the AWS CLI to create an EventBridge rule that sends all Macie finding events to a Lambda function for processing. The rule uses default settings for rules that run when specific events are received. In the procedure, the commands are formatted for Microsoft Windows. For Linux, macOS, or Unix, replace the caret (^) line-continuation character with a backslash (\).

Before you create this rule, create the Lambda function that you want the rule to use as a target. When you create the function, note the Amazon Resource Name (ARN) of the function. You'll need to enter this ARN when you specify the target for the rule.

To create an event rule by using the AWS CLI

1. Create a rule that detects events for all the findings that Macie publishes to EventBridge. To do this, use the EventBridge [put-rule](#) command. For example:

```
C:\> aws events put-rule ^  
--name MacieFindings ^  
--event-pattern "{\"source\":[\"aws.macie\"]}"
```

Where *MacieFindings* is the name that you want for the rule.

If the command runs successfully, EventBridge responds with the ARN of the rule. Note this ARN. You'll need to enter it in step 3.

Tip

You can also create a rule that uses a custom pattern to detect and act upon only a subset of Macie finding events. This subset can be based on specific fields that Macie includes in a finding event. To learn about the available fields, see [EventBridge event schema for findings \(p. 229\)](#). To learn how to create this type of rule, see [Content filtering in event patterns](#) in the *Amazon EventBridge User Guide*.

2. Specify the Lambda function to use as a target for the rule. To do this, use the EventBridge [put-targets](#) command. For example:

```
C:\> aws events put-targets ^  
--rule MacieFindings ^  
--targets Id=1,Arn=arn:aws:lambda:regionalEndpoint:accountID:function:my-findings-  
function
```

Where *MacieFindings* is the name that you specified for the rule in step 1, and the value for the *Arn* parameter is the ARN of the function that you want the rule to use as a target.

3. Add permissions that allow the rule to invoke the target Lambda function. To do this, use the Lambda [add-permission](#) command. For example:

```
C:\> aws lambda add-permission ^  
--function-name my-findings-function ^
```



```
--statement-id Sid ^  
--action lambda:InvokeFunction ^  
--principal events.amazonaws.com ^  
--source-arn arn:aws:events:regionalEndpoint:accountId:rule:MacieFindings
```

Where:

- `my-findings-function` is the name of the Lambda function that you want the rule to use as a target.
- `Sid` is a statement identifier that you define to describe the statement in the Lambda function policy.
- `source-arn` is the ARN of the EventBridge rule.

If the command runs successfully, you receive output similar to the following:

```
{  
  "Statement": "{\"Sid\":\"sid\",  
    \"Effect\":\"Allow\",  
    \"Principal\":{\"Service\":\"events.amazonaws.com\"},  
    \"Action\":\"lambda:InvokeFunction\",  
    \"Resource\":\"arn:aws:lambda:us-east-1:111122223333:function:my-findings-  
function\",  
    \"Condition\":  
      {\"ArnLike\":  
        {\"AWS:SourceArn\":  
          \"arn:aws:events:us-east-1:111122223333:rule/MacieFindings\"}}}"  
}
```

The `Statement` value is a JSON string version of the statement that was added to the Lambda function policy.

Amazon Macie integration with AWS Security Hub

AWS Security Hub is a service that provides you with a comprehensive view of your security posture across your AWS environment and helps you check your environment against security industry standards and best practices. It does this partly by consuming, aggregating, organizing, and prioritizing findings from multiple AWS services and supported AWS Partner Network security solutions. Security Hub helps you analyze your security trends and identify the highest priority security issues. To learn more about Security Hub, see the [AWS Security Hub User Guide](#).

Amazon Macie integrates with Security Hub, which enables you to publish findings from Macie to Security Hub automatically. Security Hub can then include those findings in its analysis of your security posture. This means that you can use Security Hub to monitor and process policy and sensitive data findings as part of a larger, aggregated set of findings data for your AWS environment. In other words, you can analyze Macie findings while you perform a broader analysis of your organization's security posture and remediate findings as necessary. Security Hub reduces the complexity of addressing large volumes of findings from multiple providers.

In addition, Security Hub uses a standard format for all findings, including findings from Macie. Use of this format, the *AWS Security Finding Format (ASFF)*, eliminates the need for you to perform time-consuming data conversion efforts.

Topics

- [How Amazon Macie publishes findings to AWS Security Hub \(p. 220\)](#)
- [Examples of Amazon Macie findings in AWS Security Hub \(p. 223\)](#)

- [Enabling and configuring AWS Security Hub integration \(p. 226\)](#)
- [Stopping the publication of findings to AWS Security Hub \(p. 227\)](#)

How Amazon Macie publishes findings to AWS Security Hub

In AWS Security Hub, security issues are tracked as findings. Some findings come from issues that are detected by other AWS services or by supported AWS Partner Network security solutions. Security Hub also has a set of rules that it uses to detect security issues and generate findings.

Security Hub provides tools to manage findings from all of these sources. You can view and filter lists of findings and view the details of individual findings. To learn how, see [Viewing findings](#) in the *AWS Security Hub User Guide*. You can also track the status of an investigation into a finding. To learn how, see [Taking action on findings](#) in the *AWS Security Hub User Guide*.

All findings in Security Hub use a standard JSON format called the *AWS Security Finding Format (ASFF)*. The ASFF includes details about the source of an issue, the affected resources, and the current status of a finding. For more information, see [AWS Security Finding Format \(ASFF\)](#) in the *AWS Security Hub User Guide*.

Amazon Macie is one of the AWS services that publishes findings to Security Hub.

Types of findings that Macie publishes

Depending on the publication settings that you choose for your Macie account, Macie can publish all the findings that it creates to Security Hub, both sensitive data findings and policy findings. For information about these settings and how to change them, see [Configuring publication settings for findings \(p. 227\)](#). By default, Macie publishes only new and updated policy findings to Security Hub. Macie doesn't publish sensitive data findings to Security Hub.

Sensitive data findings

If you configure Macie to publish [sensitive data findings \(p. 142\)](#) to Security Hub, Macie automatically publishes each sensitive data finding that it creates for your account and it does so immediately after it finishes processing the finding. Macie does this for all sensitive data findings that aren't archived automatically by a [suppression rule \(p. 203\)](#).

If you're the Macie administrator for an organization, publication is limited to findings from sensitive data discovery jobs that you ran. Only the account that creates a job can publish sensitive data findings that the job produces.

When Macie publishes sensitive data findings to Security Hub, it uses the [AWS Security Finding Format \(ASFF\)](#), which is the standard format for all findings in Security Hub. In the ASFF, the `Types` field indicates a finding's type. This field uses a taxonomy that's slightly different from the finding type taxonomy in Macie.

The following table lists the ASFF finding type for each type of sensitive data finding that Macie can create.

Macie finding type	ASFF finding type
SensitiveData:S3Object/Credentials	Sensitive Data Identifications/Passwords/ SensitiveData:S3Object-Credentials
SensitiveData:S3Object/CustomIdentifier	Sensitive Data Identifications/PII/ SensitiveData:S3Object-CustomIdentifier

Macie finding type	ASFF finding type
SensitiveData:S3Object/Financial	Sensitive Data Identifications/Financial/ SensitiveData:S3Object-Financial
SensitiveData:S3Object/Multiple	Sensitive Data Identifications/PII/ SensitiveData:S3Object-Multiple
SensitiveData:S3Object/Personal	Sensitive Data Identifications/PII/ SensitiveData:S3Object-Personal

Policy findings

If you configure Macie to publish [policy findings \(p. 141\)](#) to Security Hub, Macie automatically publishes each new policy finding that it creates and it does so immediately after it finishes processing the finding. If Macie detects a subsequent occurrence of an existing policy finding, it automatically publishes an update to the existing finding in Security Hub, using a publication frequency that you specify for your account. Macie performs these tasks for all policy findings that aren't archived automatically by a [suppression rule \(p. 203\)](#).

If you're the Macie administrator for an organization, publication is limited to policy findings for S3 buckets that are owned directly by your account. Macie doesn't publish policy findings that it creates or updates for member accounts in your organization. This helps ensure that you don't have duplicate findings data in Security Hub.

As is the case for sensitive data findings, Macie uses the AWS Security Finding Format (ASFF) when it publishes new and updated policy findings to Security Hub. In the ASFF, the `Types` field uses a taxonomy that's slightly different from the finding type taxonomy in Macie.

The following table lists the ASFF finding type for each type of policy finding that Macie can create. If Macie created or updated a policy finding in Security Hub on or after January 28, 2021, the finding has one of the following values for the ASFF `Types` field in Security Hub.

Macie finding type	ASFF finding type
Policy:IAMUser/S3BlockPublicAccessDisabled	Software and Configuration Checks/AWS Security Best Practices/Policy:IAMUser-S3BlockPublicAccessDisabled
Policy:IAMUser/S3BucketEncryptionDisabled	Software and Configuration Checks/AWS Security Best Practices/Policy:IAMUser-S3BucketEncryptionDisabled
Policy:IAMUser/S3BucketPublic	Effects/Data Exposure/Policy:IAMUser-S3BucketPublic
Policy:IAMUser/S3BucketReplicatedExternally	Software and Configuration Checks/AWS Security Best Practices/Policy:IAMUser-S3BucketReplicatedExternally
Policy:IAMUser/S3BucketSharedExternally	Software and Configuration Checks/AWS Security Best Practices/Policy:IAMUser-S3BucketSharedExternally

If Macie created or last updated a policy finding before January 28, 2021, the finding has one of the following values for the ASFF `Types` field in Security Hub:

- Policy:IAMUser/S3BlockPublicAccessDisabled
- Policy:IAMUser/S3BucketEncryptionDisabled
- Policy:IAMUser/S3BucketPublic
- Policy:IAMUser/S3BucketReplicatedExternally
- Policy:IAMUser/S3BucketSharedExternally

The values in the preceding list map directly to values for the **Finding type** (type) field in Macie.

Note

As you review and process policy findings in Security Hub, note the following exceptions:

- In certain AWS Regions, Macie began using ASFF finding types for new and updated findings as early as January 25, 2021.
- If you acted upon a policy finding in Security Hub before Macie began using ASFF finding types in your AWS Region, the value for the ASFF Types field of the finding will be one of the Macie finding types in the preceding list. It will not be one of the ASFF finding types in the preceding table. This is true for policy findings that you acted upon using the AWS Security Hub console or the **BatchUpdateFindings** operation of the AWS Security Hub API.

Latency for publishing findings

When Macie creates a new policy or sensitive data finding, it publishes the finding to Security Hub immediately after it finishes processing the finding.

When Macie detects a subsequent occurrence of an existing policy finding, it publishes an update to the existing Security Hub finding. The timing of the update depends on the publication frequency that you choose for your Macie account. By default, Macie publishes updates every 15 minutes. For more information, including how to change the setting for your account, see [Configuring publication settings for findings](#) (p. 227).

Retrying publication when Security Hub is not available

If Security Hub isn't available, Macie creates a queue of findings that haven't been received by Security Hub. When the system is restored, Macie retries publication until the findings are received by Security Hub.

Updating existing findings in Security Hub

After Macie publishes a policy finding to Security Hub, Macie updates the finding to reflect any additional occurrences of the finding or finding activity. Macie does this only for policy findings. Sensitive data findings, unlike policy findings, are all treated as new (unique) because they derive from individual sensitive data discovery jobs.

When Macie publishes an update to a policy finding, Macie updates the value for the **Updated At** (UpdatedAt) field of the finding. You can use this value to determine when Macie most recently detected a subsequent occurrence of the potential policy violation or issue that produced the finding.

Macie might also update the value for the **Types** (Types) field of a finding if the existing value for the field isn't an [ASFF finding type](#) (p. 221). This depends on whether you've acted upon the finding in Security Hub. If you haven't acted upon the finding, Macie changes the field's value to the appropriate ASFF finding type. If you've acted upon the finding, using either the AWS Security Hub console or the **BatchUpdateFindings** operation of the AWS Security Hub API, Macie doesn't change the field's value.

Examples of Amazon Macie findings in AWS Security Hub

When Amazon Macie publishes findings to AWS Security Hub, it uses the [AWS Security Finding Format \(ASFF\)](#). This is the standard format for all findings in Security Hub. The following examples use sample data to demonstrate the structure and nature of the findings data that Macie publishes to Security Hub in this format:

- [Example of a sensitive data finding \(p. 223\)](#)
- [Example of a policy finding \(p. 225\)](#)

Example of a sensitive data finding in Security Hub

Here's an example of a sensitive data finding that Macie published to Security Hub using the ASFF.

```
{
  "SchemaVersion": "2018-10-08",
  "Id": "5be50fce24526e670df77bc00example",
  "ProductArn": "arn:aws:securityhub:us-east-1::product/aws/macie",
  "ProductName": "Macie",
  "CompanyName": "Amazon",
  "Region": "us-east-1",
  "GeneratorId": "aws/macie",
  "AwsAccountId": "111122223333",
  "Types": [
    "Sensitive Data Identifications/PII/SensitiveData:S3Object-Personal"
  ],
  "CreatedAt": "2022-05-11T10:23:49.667Z",
  "UpdatedAt": "2022-05-11T10:23:49.667Z",
  "Severity": {
    "Label": "HIGH",
    "Normalized": 70
  },
  "Title": "The S3 object contains personal information.",
  "Description": "The object contains personal information such as first or last names, addresses, or identification numbers.",
  "ProductFields": {
    "JobArn": "arn:aws:macie2:us-east-1:111122223333:classification-job/698e99c283a255bb2c992feceexample",
    "S3Object.Path": "DOC-EXAMPLE-BUCKET1/2022 Sourcing.tsv",
    "S3Object.Extension": "tsv",
    "S3Bucket.effectivePermission": "NOT_PUBLIC",
    "OriginType": "SENSITIVE_DATA_DISCOVERY_JOB",
    "S3Object.PublicAccess": "false",
    "S3Object.Size": "14",
    "S3Object.StorageClass": "STANDARD",
    "S3Bucket.allowsUnencryptedObjectUploads": "TRUE",
    "JobId": "698e99c283a255bb2c992feceexample",
    "aws/securityhub/FindingId": "arn:aws:securityhub:us-east-1::product/aws/macie/5be50fce24526e670df77bc00example",
    "aws/securityhub/ProductName": "Macie",
    "aws/securityhub/CompanyName": "Amazon"
  },
  "Resources": [
    {
      "Type": "AwsS3Bucket",
      "Id": "arn:aws:s3::DOC-EXAMPLE-BUCKET1",
      "Partition": "aws",
      "Region": "us-east-1",
      "Details": {
```

```
      "AwsS3Bucket": {
        "OwnerId":
"7009a8971cd538e11f6b6606438875e7c86c5b672f46db45460ddcd08example",
        "OwnerName": "johndoe",
        "OwnerAccountId": "444455556666",
        "CreatedAt": "2020-12-30T18:16:25.000Z",
        "ServerSideEncryptionConfiguration": {
          "Rules": [
            {
              "ApplyServerSideEncryptionByDefault": {
                "SSEAlgorithm": "NONE"
              }
            }
          ]
        },
        "PublicAccessBlockConfiguration": {
          "BlockPublicAcls": true,
          "BlockPublicPolicy": true,
          "IgnorePublicAcls": true,
          "RestrictPublicBuckets": true
        }
      }
    },
    {
      "Type": "AwsS3Object",
      "Id": "arn:aws:s3:::DOC-EXAMPLE-BUCKET1/2022 Sourcing.tsv",
      "Partition": "aws",
      "Region": "us-east-1",
      "DataClassification": {
        "DetailedResultsLocation": "s3://macie-data-discovery-results/
AWSLogs/111122223333/Macie/us-east-1/
698e99c283a255bb2c992feceexample/111122223333/32b8485d-4f3a-3aa1-be33-
aa3f0example.jsonl.gz",
        "Result": {
          "MimeType": "text/tsv",
          "SizeClassified": 14,
          "AdditionalOccurrences": false,
          "Status": {
            "Code": "COMPLETE"
          },
          "SensitiveData": [
            {
              "Category": "PERSONAL_INFORMATION",
              "Detections": [
                {
                  "Count": 1,
                  "Type": "USA_SOCIAL_SECURITY_NUMBER",
                  "Occurrences": {
                    "Cells": [
                      {
                        "Column": 10,
                        "Row": 1,
                        "ColumnName": "Other"
                      }
                    ]
                  }
                }
              ]
            }
          ],
          "TotalCount": 1
        }
      },
      "CustomDataIdentifiers": {
        "Detections": [
        ],
        "TotalCount": 0
      }
    }
  ]
}
```

```
    }
  },
  "Details": {
    "AwsS3Object": {
      "LastModified": "2022-04-22T18:16:46.000Z",
      "ETag": "e8e1ca03ee8d006d457444445example",
      "VersionId": "SlBC72z5hArgexOJifxw_IN57EXAMPLE",
      "ServerSideEncryption": "NONE"
    }
  }
},
"WorkflowState": "NEW",
"Workflow": {
  "Status": "NEW"
},
"RecordState": "ACTIVE",
"FindingProviderFields": {
  "Severity": {
    "Label": "HIGH"
  },
  "Types": [
    "Sensitive Data Identifications/PII/SensitiveData:S3Object-Personal"
  ]
},
"Sample": false
}
```

Example of a policy finding in Security Hub

Here's an example of a new policy finding that Macie published to Security Hub in the ASFF.

```
{
  "SchemaVersion": "2018-10-08",
  "Id": "36ca8ba0-caf1-4fee-875c-37760example",
  "ProductArn": "arn:aws:securityhub:us-east-1::product/aws/macie",
  "ProductName": "Macie",
  "CompanyName": "Amazon",
  "Region": "us-east-1",
  "GeneratorId": "aws/macie",
  "AwsAccountId": "111122223333",
  "Types": [
    "Software and Configuration Checks/AWS Security Best Practices/Policy:IAMUser-S3BlockPublicAccessDisabled"
  ],
  "CreatedAt": "2022-04-24T09:27:43.313Z",
  "UpdatedAt": "2022-04-24T09:27:43.313Z",
  "Severity": {
    "Label": "HIGH",
    "Normalized": 70
  },
  "Title": "Block Public Access settings are disabled for the S3 bucket",
  "Description": "All Amazon S3 block public access settings are disabled for the Amazon S3 bucket. Access to the bucket is controlled only by access control lists (ACLs) or bucket policies.",
  "ProductFields": {
    "S3Bucket.effectivePermission": "PUBLIC",
    "S3Bucket.allowsUnencryptedObjectUploads": "FALSE",
    "aws/securityhub/FindingId": "arn:aws:securityhub:us-east-1::product/aws/macie/36ca8ba0-caf1-4fee-875c-37760example",
    "aws/securityhub/ProductName": "Macie",
    "aws/securityhub/CompanyName": "Amazon"
  },
}
```

```
"Resources": [
  {
    "Type": "AwsS3Bucket",
    "Id": "arn:aws:s3::DOC-EXAMPLE-BUCKET2",
    "Partition": "aws",
    "Region": "us-east-1",
    "Tags": {
      "Team": "Recruiting",
      "Division": "HR"
    },
    "Details": {
      "AwsS3Bucket": {
        "OwnerId":
"7009a8971cd538e11f6b6606438875e7c86c5b672f46db45460ddcd08example",
        "OwnerName": "johndoe",
        "OwnerAccountId": "444455556666",
        "CreatedAt": "2020-11-25T18:24:38.000Z",
        "ServerSideEncryptionConfiguration": {
          "Rules": [
            {
              "ApplyServerSideEncryptionByDefault": {
                "SSEAlgorithm": "NONE"
              }
            }
          ]
        },
        "PublicAccessBlockConfiguration": {
          "BlockPublicAcls": false,
          "BlockPublicPolicy": false,
          "IgnorePublicAcls": false,
          "RestrictPublicBuckets": false
        }
      }
    }
  },
  {
    "WorkflowState": "NEW",
    "Workflow": {
      "Status": "NEW"
    },
    "RecordState": "ACTIVE",
    "FindingProviderFields": {
      "Severity": {
        "Label": "HIGH"
      },
      "Types": [
        "Software and Configuration Checks/AWS Security Best Practices/Policy:IAMUser-
S3BlockPublicAccessDisabled"
      ]
    },
    "Sample": false
  }
]
```

Enabling and configuring AWS Security Hub integration

To use Amazon Macie integration with AWS Security Hub, you must enable Security Hub for your AWS account. For information about how to enable Security Hub, see [Setting up AWS Security Hub](#) in the *AWS Security Hub User Guide*.

When you enable both Macie and Security Hub, the integration is enabled automatically. This means that Macie automatically begins to publish new and updated policy findings to Security Hub. You don't need to take any additional steps to configure the integration.

You can optionally customize your configuration by choosing the frequency with which Macie publishes updates to policy findings in Security Hub. You can also choose to publish sensitive data findings to Security Hub in addition to policy findings. To learn how, see [Configuring publication settings for findings](#) (p. 227).

Stopping the publication of findings to AWS Security Hub

To stop publishing findings to AWS Security Hub, you can change the publication settings for your Amazon Macie account. To learn how, see [Choosing publication destinations for findings](#) (p. 227). You can also do this by using the Security Hub console or the Security Hub API. To learn how, see [Disabling and enabling the flow of findings from an integration \(console\)](#) or [Disabling the flow of findings from an integration \(Security Hub API, AWS CLI\)](#) in the *AWS Security Hub User Guide*.

Configuring publication settings for Amazon Macie findings

To support integration with other applications, services, and systems, Amazon Macie automatically publishes both policy findings and sensitive data findings to Amazon EventBridge as events. For information about how you can use EventBridge to monitor and process findings, see [EventBridge integration](#) (p. 215).

You can configure Macie to automatically publish findings to AWS Security Hub too, using destination options that you specify in the publication settings for your account. With these options, you can configure Macie to publish only policy findings, only sensitive data findings, or both policy and sensitive data findings to Security Hub. You can also configure Macie to stop publishing any findings to Security Hub. For information about how you can use Security Hub to monitor and process findings, see [Security Hub integration](#) (p. 219).

For policy findings, the timing with which Macie publishes a finding to another AWS service depends on whether the finding is new and on the publication frequency that you specify for your account. For sensitive data findings, the timing is always immediate—Macie publishes a sensitive data finding immediately after it finishes processing the finding. Unlike policy findings, Macie treats all sensitive data findings as new (unique) because they derive from individual sensitive data discovery jobs.

Note that Macie doesn't publish policy or sensitive data findings that are archived automatically by a [suppression rule](#) (p. 203). In other words, Macie doesn't publish suppressed findings to other AWS services.

Topics

- [Choosing publication destinations for findings](#) (p. 227)
- [Determining the publication frequency for findings](#) (p. 228)
- [Changing the publication frequency for findings](#) (p. 229)

Choosing publication destinations for findings

You can configure Macie to automatically publish policy and sensitive data findings to Security Hub in addition to EventBridge. By default, Macie publishes only new and updated policy findings to Security

Hub. To change or extend the default configuration, adjust the publication destination settings for your account.

When you adjust your destination settings, you choose the categories of findings that you want Macie to publish to Security Hub—only sensitive data findings, only policy findings, or both sensitive data and policy findings. You can also choose to stop publishing any category of finding to Security Hub.

If you change your destination settings, your change applies only to the current AWS Region. If you're the Macie administrator for an organization, your change applies only to your account. It doesn't apply to any associated member accounts. For more information, see [Managing multiple accounts \(p. 238\)](#).

To choose publication destinations for findings

1. Open the Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, choose **Settings**.
3. In the **Publication of findings** section, under **Destinations**, choose from the following options:
 - **Publish policy findings to** – Select the **Security Hub** check box to automatically publish new and updated policy findings to Security Hub. To stop publishing new and updated policy findings to Security Hub, clear this check box.
 - **Publish sensitive data findings to** – Select the **Security Hub** check box to automatically publish sensitive data findings to Security Hub. To stop publishing sensitive data findings to Security Hub, clear this check box.
4. Choose **Save**.

If you chose to publish any category of finding to Security Hub, make sure that you also enable Security Hub in the current Region and configure it to accept the findings that Macie publishes. Otherwise, you won't be able to access the findings in Security Hub. To learn how to accept findings in Security Hub, see [Managing product integrations](#) in the *AWS Security Hub User Guide*.

Determining the publication frequency for findings

In Macie, each finding has a unique identifier. Macie uses this identifier to determine when to publish a finding to another AWS service:

- **New findings** – When Macie creates a new policy or sensitive data finding, it assigns a unique identifier to the finding as part of processing the finding. Immediately after Macie finishes processing the finding, it publishes the finding as a new EventBridge event. Depending on the publication settings for your account, Macie also publishes the finding as a new finding in Security Hub.
- **Updated findings** – When Macie detects a subsequent occurrence of an existing policy finding, it updates the existing finding by adding details about the subsequent occurrence and incrementing the count of occurrences. Macie also publishes these updates to the existing EventBridge event and, depending on the publication settings for your account, the existing Security Hub finding. Macie does this only for policy findings. Sensitive data findings, unlike policy findings, are all treated as new (unique) because they derive from individual sensitive data discovery jobs.

By default, Macie publishes updated findings every 15 minutes as part of a recurring publication cycle. This means that any policy findings that are updated after the most recent publication cycle will be held, updated again as necessary, and included in the next publication cycle (approximately 15 minutes later). You can change this schedule by choosing a different publication frequency. For example, if you configure Macie to publish updated findings every hour and a publication occurs at 12:00, then any updates that occur after 12:00 are published at 13:00.

Note that neither of these cases applies to findings that are archived automatically by a [suppression rule \(p. 203\)](#). Macie doesn't publish suppressed findings to other AWS services.

Changing the publication frequency for findings

You can change the schedule that Macie uses to publish updates to existing policy findings in other AWS services. By default, Macie publishes updated findings every 15 minutes. If you change this schedule, your change applies only to the current AWS Region. If you're the Macie administrator for an organization, your change also applies to all associated member accounts in the Region. For more information, see [Managing multiple accounts](#) (p. 238).

To change the publication frequency for updated findings

1. Open the Macie console at <https://console.aws.amazon.com/macie/>.
2. In the navigation pane, choose **Settings**.
3. In the **Publication of findings** section, under **Update frequency for policy findings**, choose how often you want Macie to publish updated policy findings to other AWS services.
4. Choose **Save**.

Amazon EventBridge event schema for Amazon Macie findings

To support integration with other applications, services, and systems, such as monitoring or event management systems, Amazon Macie automatically publishes findings to Amazon EventBridge as events. EventBridge, formerly Amazon CloudWatch Events, is a serverless event bus service that delivers a stream of real-time data from applications and other AWS services to targets such as AWS Lambda functions, Amazon Simple Notification Service topics, and Amazon Kinesis streams. To learn more about EventBridge and EventBridge events, see the [Amazon EventBridge User Guide](#).

Note

If you currently use CloudWatch Events, note that EventBridge and CloudWatch Events are the same underlying service and API. However, EventBridge includes additional features that enable you to receive events from software as a service (SaaS) applications and your own applications. Because the underlying service and API are the same, the event schema for Macie findings is also the same.

Macie publishes events for all new findings and subsequent occurrences of existing policy findings, except findings that you archive automatically using [suppression rules](#) (p. 203). Each event is a JSON object that conforms to the EventBridge schema for AWS events and contains a JSON representation of a finding. Because the findings data is structured as an EventBridge event, you can more easily monitor, process, and act upon findings by using other applications, services, and tools.

Topics

- [Event schema](#) (p. 229)
- [Event example for a policy finding](#) (p. 230)
- [Event example for a sensitive data finding](#) (p. 233)

Event schema

The following example shows the schema of an [EventBridge event](#) for a Macie finding. For detailed descriptions of the fields that can be included in a finding event, see the [Finding table](#) in the *Amazon Macie API Reference*. The structure and fields of a finding event map closely to the **Finding** object of the Amazon Macie API.

```
{
  "version": "0",
  "id": "event ID",
  "detail-type": "Macie Finding",
  "source": "aws.macie",
  "account": "Amazon Web Services account ID (string)",
  "time": "event timestamp (string)",
  "region": "AWS Region (string)",
  "resources": [
    <-- ARNs of the resources involved in the event -->
  ],
  "detail": {
    <-- Details for a policy or sensitive data finding -->
  },
  "policyDetails": null, <-- Additional details for a policy finding or "null" for a
sensitive data finding -->
  "sample": Boolean,
  "archived": Boolean
}
```

Event example for a policy finding

The following example uses sample data to demonstrate the structure and nature of objects and fields in an EventBridge event for a policy finding.

In this example, the event reports a subsequent occurrence of an existing policy finding: default encryption was disabled for an S3 bucket. The following fields and values can help you determine that this is the case:

- The `type` field is set to `Policy:IAMUser/S3BucketEncryptionDisabled`.
- The `createdAt` and `updatedAt` fields have different values. This is one indicator that the event reports a subsequent occurrence of an existing finding. The values for these fields would be the same if the event reported a new finding.
- The `count` field is set to 2, which indicates that this is the second occurrence of the finding.
- The `category` field is set to `POLICY`.
- The value for the `classificationDetails` field is `null`, which helps differentiate this event for a policy finding from an event for a sensitive data finding. For a sensitive data finding, this value would be a set of objects and fields that provide information about how and what sensitive data was found.

Also note that the value for the `sample` field is `true`. This value emphasizes that this is an example event for use in the documentation.

```
{
  "version": "0",
  "id": "0948ba87-d3b8-c6d4-f2da-732a1example",
  "detail-type": "Macie Finding",
  "source": "aws.macie",
  "account": "123456789012",
  "time": "2021-04-29T23:12:15Z",
  "region": "us-east-1",
  "resources": [],
  "detail": {
    "schemaVersion": "1.0",
    "id": "64b917aa-3843-014c-91d8-937ffexample",
    "accountId": "123456789012",
    "partition": "aws",
    "region": "us-east-1",
    "type": "Policy:IAMUser/S3BucketEncryptionDisabled",
  }
}
```

```
    "title": "Encryption is disabled for the S3 bucket",
    "description": "Encryption is disabled for the Amazon S3 bucket. The data in the
bucket isn't encrypted
    using server-side encryption.",
    "severity": {
      "score": 1,
      "description": "Low"
    },
    "createdAt": "2021-04-29T15:46:02Z",
    "updatedAt": "2021-04-29T23:12:15Z",
    "count": 2,
    "resourcesAffected": {
      "s3Bucket": {
        "arn": "arn:aws:s3:::DOC-EXAMPLE-BUCKET1",
        "name": "DOC-EXAMPLE-BUCKET1",
        "createdAt": "2020-04-03T20:46:56.000Z",
        "owner": {
          "displayName": "johndoe",
          "id":
"7009a8971cd538e11f6b6606438875e7c86c5b672f46db45460ddcd08example"
        },
        "tags": [
          {
            "key": "Division",
            "value": "HR"
          },
          {
            "key": "Team",
            "value": "Recruiting"
          }
        ],
        "defaultServerSideEncryption": {
          "encryptionType": "NONE",
          "kmsMasterKeyId": null
        },
        "publicAccess": {
          "permissionConfiguration": {
            "bucketLevelPermissions": {
              "accessControlList": {
                "allowsPublicReadAccess": false,
                "allowsPublicWriteAccess": false
              },
              "bucketPolicy": {
                "allowsPublicReadAccess": false,
                "allowsPublicWriteAccess": false
              },
              "blockPublicAccess": {
                "ignorePublicAcls": true,
                "restrictPublicBuckets": true,
                "blockPublicAcls": true,
                "blockPublicPolicy": true
              }
            },
            "accountLevelPermissions": {
              "blockPublicAccess": {
                "ignorePublicAcls": false,
                "restrictPublicBuckets": false,
                "blockPublicAcls": false,
                "blockPublicPolicy": false
              }
            }
          },
          "effectivePermission": "NOT_PUBLIC"
        },
        "allowsUnencryptedObjectUploads": "FALSE"
      },
    },
  },
```

```
    "s3Object": null
  },
  "category": "POLICY",
  "classificationDetails": null,
  "policyDetails": {
    "action": {
      "actionType": "AWS_API_CALL",
      "apiCallDetails": {
        "api": "DeleteBucketEncryption",
        "apiServiceName": "s3.amazonaws.com",
        "firstSeen": "2021-04-29T15:46:02.401Z",
        "lastSeen": "2021-04-29T23:12:15.401Z"
      }
    },
    "actor": {
      "userIdentity": {
        "type": "AssumedRole",
        "assumedRole": {
          "principalId": "AROA1234567890EXAMPLE:AssumedRoleSessionName",
          "arn": "arn:aws:sts::123456789012:assumed-role/RoleToBeAssumed/MySessionName",
          "accountId": "111122223333",
          "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
          "sessionContext": {
            "attributes": {
              "mfaAuthenticated": false,
              "creationDate": "2021-04-29T10:25:43.511Z"
            }
          },
          "sessionIssuer": {
            "type": "Role",
            "principalId": "AROA1234567890EXAMPLE",
            "arn": "arn:aws:iam::123456789012:role/RoleToBeAssumed",
            "accountId": "123456789012",
            "userName": "RoleToBeAssumed"
          }
        }
      }
    },
    "root": null,
    "iamUser": null,
    "federatedUser": null,
    "awsAccount": null,
    "awsService": null
  },
  "ipAddressDetails": {
    "ipAddressV4": "192.0.2.0",
    "ipOwner": {
      "asn": "-1",
      "asnOrg": "ExampleFindingASNorg",
      "isp": "ExampleFindingISP",
      "org": "ExampleFindingORG"
    },
    "ipCountry": {
      "code": "US",
      "name": "United States"
    },
    "ipCity": {
      "name": "Ashburn"
    },
    "ipGeoLocation": {
      "lat": 39.0481,
      "lon": -77.4728
    }
  },
  "domainDetails": null
}
},
```

```
    "sample": true,  
    "archived": false  
  }  
}
```

Event example for a sensitive data finding

The following example uses sample data to demonstrate the structure and nature of objects and fields in an EventBridge event for a sensitive data finding.

In this example, the event reports a new sensitive data finding: an S3 object contains more than one category of sensitive data. The following fields and values can help you determine that this is the case:

- The `type` field is set to `SensitiveData:S3Object/Multiple`.
- The `createdAt` and `updatedAt` fields have the same values. Unlike policy findings, this is always the case for sensitive data findings. All sensitive data findings are considered new because they derive from individual sensitive data discovery jobs.
- The `count` field is set to `1`, which indicates that this is a new finding. Unlike policy findings, this is always the case for sensitive data findings. All sensitive data findings are considered unique because they derive from individual sensitive data discovery jobs.
- The `category` field is set to `CLASSIFICATION`.
- The `jobArn` and `jobId` fields indicate which sensitive data discovery job produced the finding.
- The value for the `policyDetails` field is `null`, which helps differentiate this event for a sensitive data finding from an event for a policy finding. For a policy finding, this value would be a set of objects and fields that provide information about a potential policy violation or issue with the security or privacy of an S3 bucket.

Also note that the value for the `sample` field is `true`. This value emphasizes that this is an example event for use in the documentation.

```
{  
  "version": "0",  
  "id": "14ddd0b1-7c90-b9e3-8a68-6a408example",  
  "detail-type": "Macie Finding",  
  "source": "aws.macie",  
  "account": "123456789012",  
  "time": "2022-04-20T08:19:10Z",  
  "region": "us-east-1",  
  "resources": [],  
  "detail": {  
    "schemaVersion": "1.0",  
    "id": "4ed45d06-c9b9-4506-ab7f-18a57example",  
    "accountId": "123456789012",  
    "partition": "aws",  
    "region": "us-east-1",  
    "type": "SensitiveData:S3Object/Multiple",  
    "title": "The S3 object contains multiple types of sensitive information.",  
    "description": "The object contains more than one type of sensitive information.",  
    "severity": {  
      "score": 3,  
      "description": "High"  
    },  
    "createdAt": "2022-04-20T18:19:10Z",  
    "updatedAt": "2022-04-20T18:19:10Z",  
    "count": 1,  
    "resourcesAffected": {  
      "s3Bucket": {
```

Amazon Macie User Guide
Event example for a sensitive data finding

```
    "arn": "arn:aws:s3:::DOC-EXAMPLE-BUCKET2",
    "name": "DOC-EXAMPLE-BUCKET2",
    "createdAt": "2020-05-15T20:46:56.000Z",
    "owner": {
      "displayName": "johndoe",
      "id":
"7009a8971cd538e11f6b6606438875e7c86c5b672f46db45460ddcd08example"
    },
    "tags":[
      {
        "key":"Division",
        "value":"HR"
      },
      {
        "key":"Team",
        "value":"Recruiting"
      }
    ],
    "defaultServerSideEncryption": {
      "encryptionType": "aws:kms",
      "kmsMasterKeyId": "arn:aws:kms:us-
east-1:123456789012:key/1234abcd-12ab-34cd-56ef-1234567890ab"
    },
    "publicAccess": {
      "permissionConfiguration": {
        "bucketLevelPermissions": {
          "accessControlList": {
            "allowsPublicReadAccess": false,
            "allowsPublicWriteAccess": false
          },
          "bucketPolicy":{
            "allowsPublicReadAccess": false,
            "allowsPublicWriteAccess": false
          },
          "blockPublicAccess": {
            "ignorePublicAcls": true,
            "restrictPublicBuckets": true,
            "blockPublicAcls": true,
            "blockPublicPolicy": true
          }
        },
        "accountLevelPermissions": {
          "blockPublicAccess": {
            "ignorePublicAcls": false,
            "restrictPublicBuckets": false,
            "blockPublicAcls": false,
            "blockPublicPolicy": false
          }
        }
      },
      "effectivePermission": "NOT_PUBLIC"
    },
    "allowsUnencryptedObjectUploads": "TRUE",
  },
  "s3Object":{
    "bucketArn": "arn:aws:s3:::DOC-EXAMPLE-BUCKET2",
    "key": "2022 Sourcing.csv",
    "path": "DOC-EXAMPLE-BUCKET2/2022 Sourcing.csv",
    "extension": ".csv",
    "lastModified": "2022-04-19T22:08:25.000Z",
    "versionId": "",
    "serverSideEncryption": {
      "encryptionType": "aws:kms",
      "kmsMasterKeyId": "arn:aws:kms:us-
east-1:123456789012:key/1234abcd-12ab-34cd-56ef-1234567890ab"
    }
  },
```



```
    "size": 4750,
    "storageClass": "STANDARD",
    "tags": [
      {
        "key": "Division",
        "value": "HR"
      },
      {
        "key": "Team",
        "value": "Recruiting"
      }
    ],
    "publicAccess": false,
    "etag": "6bb7fd4fa9d36d6b8fb8882caexample"
  }
},
"category": "CLASSIFICATION",
"classificationDetails": {
  "jobArn": "arn:aws:macie2:us-east-1:123456789012:classification-
job/3ce05dbb7ec5505def334104bexample",
  "jobId": "3ce05dbb7ec5505def334104bexample",
  "result": {
    "status": {
      "code": "COMPLETE",
      "reason": null
    },
    "sizeClassified": 4750,
    "mimeType": "text/csv",
    "additionalOccurrences": true,
    "sensitiveData": [
      {
        "category": "PERSONAL_INFORMATION",
        "totalCount": 65,
        "detections": [
          {
            "type": "USA_SOCIAL_SECURITY_NUMBER",
            "count": 30,
            "occurrences": {
              "lineRanges": null,
              "offsetRanges": null,
              "pages": null,
              "records": null,
              "cells": [
                {
                  "row": 2,
                  "column": 1,
                  "columnName": "SSN",
                  "cellReference": null
                },
                {
                  "row": 3,
                  "column": 1,
                  "columnName": "SSN",
                  "cellReference": null
                },
                {
                  "row": 4,
                  "column": 1,
                  "columnName": "SSN",
                  "cellReference": null
                }
              ]
            }
          }
        ]
      }
    ],
    {
      "type": "NAME",
```

```
        "count": 35,
        "occurrences": {
          "lineRanges": null,
          "offsetRanges": null,
          "pages": null,
          "records": null,
          "cells": [
            {
              "row": 2,
              "column": 3,
              "columnName": "Name",
              "cellReference": null
            },
            {
              "row": 3,
              "column": 3,
              "columnName": "Name",
              "cellReference": null
            }
          ]
        }
      ],
    },
    {
      "category": "FINANCIAL_INFORMATION",
      "totalCount": 30,
      "detections": [
        {
          "type": "CREDIT_CARD_NUMBER",
          "count": 30,
          "occurrences": {
            "lineRanges": null,
            "offsetRanges": null,
            "pages": null,
            "records": null,
            "cells": [
              {
                "row": 2,
                "column": 14,
                "columnName": "CCN",
                "cellReference": null
              },
              {
                "row": 3,
                "column": 14,
                "columnName": "CCN",
                "cellReference": null
              }
            ]
          }
        }
      ]
    }
  ],
  "customDataIdentifiers": {
    "totalCount": 0,
    "detections": []
  },
  "detailedResultsLocation": "s3://macie-data-discovery-results/
AWSLogs/123456789012/Macie/us-east-1/
3ce05dbb7ec5505def334104bexample/d48bf16d-0deb-3e49-9d8c-
d407cexample.jsonl.gz",
  "originType": "SENSITIVE_DATA_DISCOVERY_JOB"
},
```

```
    "policyDetails": null,  
    "sample": true,  
    "archived": false  
  }  
}
```

Managing multiple Amazon Macie accounts

If your AWS environment has multiple accounts, you can associate the Amazon Macie accounts in your environment and centrally manage them as an organization in Macie. With this configuration, a designated Macie administrator can assess and monitor the overall security posture of your organization's Amazon Simple Storage Service (Amazon S3) data estate, and run sensitive data discovery jobs to detect sensitive data in your organization's S3 buckets. The administrator can also perform various account management and administration tasks at scale, such as monitoring estimated usage costs and assessing account quotas.

In Macie, an organization consists of a designated Macie administrator account and one or more associated member accounts. You can associate the accounts in two ways, by integrating Macie with AWS Organizations or by sending and accepting membership invitations in Macie. We recommend that you integrate Macie with AWS Organizations.

AWS Organizations is a global account management service that enables AWS administrators to consolidate and centrally manage multiple AWS accounts. It provides account management and consolidated billing features that are designed to support budgetary, security, and compliance needs. It's offered at no additional charge and it integrates with multiple AWS services, including Macie, AWS Security Hub, and Amazon GuardDuty. To learn more, see the [AWS Organizations User Guide](#).

If you prefer to centrally manage multiple Macie accounts without using AWS Organizations, you can use membership invitations instead. If you send an invitation and it's accepted by another account, your account becomes the Macie administrator account for the other account. If you receive and accept an invitation, your account becomes a Macie member account and the Macie administrator account can access and manage certain settings, data, and resources for your Macie account.

Topics

- [Understanding the relationship between Amazon Macie administrator and member accounts \(p. 238\)](#)
- [Managing Amazon Macie accounts with AWS Organizations \(p. 240\)](#)
- [Managing Amazon Macie accounts by invitation \(p. 260\)](#)

Understanding the relationship between Amazon Macie administrator and member accounts

If you centrally manage multiple Amazon Macie accounts as an organization, the Macie administrator has access to Amazon Simple Storage Service (Amazon S3) inventory data, policy findings, and certain Macie settings and resources for associated member accounts. The administrator can also run sensitive data discovery jobs to detect sensitive data in S3 buckets that member accounts own. Support for specific tasks varies based on whether a Macie administrator account is associated with a member account through AWS Organizations or by invitation.

The following table provides details about the relationship between Macie administrator and member accounts. It indicates the default permissions for each type of account. To further restrict access to Macie features and operations, you can use custom [AWS Identity and Access Management \(IAM\) policies \(p. 286\)](#).

In the table:

- **Self** indicates that the account can't perform the task for any associated accounts.
- **Any** indicates that the account can perform the task for an individual associated account.
- **All** indicates that the account can perform the task and the task applies to all associated accounts.

The dash (–) indicates that the account can't perform the task.

Task	Designation		
	Administrator	Administrator	Member
	Through AWS Organizations	By invitation	
Enable Macie	Any	Self	Self
Review the organization's account inventory ¹ (p. 240)	All	All	–
Add a member account	Any	Any	–
Remove (disassociate) a member account	Any	Any	–
Disassociate from an administrator account ² (p. 240)	–	–	Self
Create sample findings	Self	Self	Self
Review metadata and statistics for S3 buckets	All	All	Self
Review policy findings	All	All	Self
Suppress (archive) policy findings ³ (p. 240)	All	All	–
Publish policy findings ⁴ (p. 240)	Self	Self	Self
Create and use allow lists	Self	Self	Self
Create and use custom data identifiers	Self	Self	Self
Create and run sensitive data discovery jobs ⁵ (p. 240)	Self	Self	Self
Review the details of sensitive data discovery jobs ⁶ (p. 240)	Self	Self	Self
Review sensitive data findings ⁷ (p. 240)	Self	Self	Self
Suppress (archive) sensitive data findings ⁷ (p. 240)	Self	Self	Self
Publish sensitive data findings ⁷ (p. 240)	Self	Self	Self
Configure publication destinations for findings	Self	Self	Self
Set the publication frequency for findings	All	All	Self
Configure a repository for sensitive data discovery results	Self	Self	Self

Configure Macie to retrieve sensitive data samples	Self	Self	Self
Review account quotas and estimated usage costs	All	All	Self
Suspend Macie ^{8 (p. 240)}	Any	Any	Self
Disable Macie ^{9 (p. 240)}	Self	Self	Self

1. The Macie administrator for an organization in AWS Organizations can review all accounts in the organization, including accounts that haven't enabled Macie. The administrator for an invitation-based organization can review only those accounts that they add to their inventory.
2. A member of an AWS Organizations organization can't perform this task. A member of an invitation-based organization can perform this task.
3. Only an administrator can suppress policy findings. If an administrator creates a suppression rule, Macie applies the rule to policy findings for all accounts in the organization unless the rule is configured to exclude specific accounts. If a member creates a suppression rule, Macie doesn't apply the rule to policy findings for the member's account.
4. Only the account that owns an affected resource can publish policy findings for the resource to AWS Security Hub. Both administrator and member accounts automatically publish policy findings for an affected resource to Amazon EventBridge.
5. A member can configure a job to analyze objects only in S3 buckets that their account owns. An administrator can configure a job to analyze objects in buckets that their account owns or a member account owns. For information about how quotas are applied and costs are calculated for multiple-account jobs, see [Understanding how estimated usage costs are calculated \(p. 278\)](#).
6. Only the account that creates a job can access the job's details. This includes job-related details in the S3 bucket inventory.
7. Only the account that creates a job can access, suppress, or publish sensitive data findings that the job produces.
8. For an administrator to perform this task for their own account, the administrator must first disassociate their account from all member accounts.
9. For an administrator to perform this task for their own account, the administrator must first disassociate their account from all member accounts and delete the associations between their account and all of those accounts. For a member to perform this task for their account, the member must first disassociate their account from its administrator account.

Managing Amazon Macie accounts with AWS Organizations

If you use AWS Organizations to centrally manage multiple AWS accounts, you can integrate Amazon Macie with AWS Organizations, and then centrally manage Macie for accounts in your organization. With this configuration, a designated Macie administrator can enable and manage Macie for as many as 5,000 accounts. The administrator can also access Amazon Simple Storage Service (Amazon S3) inventory data

and run sensitive data discovery jobs for the accounts. For details about tasks that the administrator can perform, see [Understanding the relationship between Amazon Macie administrator and member accounts](#) (p. 238).

To integrate Macie with AWS Organizations, you start by designating an account as the delegated Macie administrator account for the organization. The Macie administrator then enables Macie for other accounts in the organization, adds those accounts as Macie member accounts, and configures Macie settings and resources for the accounts.

Tip

If you already associated a Macie administrator account with member accounts by using invitations, you can designate that account as the delegated Macie administrator account for your organization in AWS Organizations. If you do this, all currently associated member accounts remain members and you can take full advantage of the benefits of managing accounts by using AWS Organizations. For more information, see [Transitioning from an invitation-based organization](#) (p. 243).

The topics in this section explain how to integrate Macie with AWS Organizations and how to administer and manage Macie for accounts in an organization.

Topics

- [Considerations and recommendations for using Amazon Macie with AWS Organizations](#) (p. 241)
- [Integrating and configuring an organization in Amazon Macie](#) (p. 243)
- [Reviewing Amazon Macie accounts for an organization](#) (p. 249)
- [Managing Amazon Macie member accounts for an organization](#) (p. 252)
- [Designating a different Amazon Macie administrator account for an organization](#) (p. 257)
- [Disabling Amazon Macie integration with AWS Organizations](#) (p. 259)

Considerations and recommendations for using Amazon Macie with AWS Organizations

Before you integrate Amazon Macie with AWS Organizations and configure your organization in Macie, consider the following requirements and recommendations.

Topics

- [Designating a Macie administrator account](#) (p. 241)
- [Changing or removing the designation of a Macie administrator account](#) (p. 242)
- [Adding and removing Macie member accounts](#) (p. 242)
- [Transitioning from an invitation-based organization](#) (p. 243)

Designating a Macie administrator account

While you determine which account should be the delegated Macie administrator account for your organization, keep the following in mind:

- An organization can have only one delegated Macie administrator account.
- An account can't be a Macie administrator and member account at the same time.
- Only the AWS Organizations management account for an organization can designate the delegated Macie administrator account for the organization, and only the management account can subsequently change or remove that designation.

- The AWS Organizations management account for an organization can also be the delegated Macie administrator account for the organization. However, we don't recommend this configuration based on AWS security best practices and the principle of least privilege. Users who have access to the management account for billing purposes are likely to be different from users who need access to Macie for information security purposes.

If you prefer this configuration, you must enable Macie for the organization's management account in at least one AWS Region before you designate the account as the delegated Macie administrator account. Otherwise, the account won't be able to access and manage Macie settings and resources for member accounts.

- Unlike AWS Organizations, Macie is a Regional service. This means that the designation of a Macie administrator account is a Regional designation. It also means that associations between Macie administrator and member accounts are Regional. For example, if the management account designates a Macie administrator account in the US East (N. Virginia) Region, the Macie administrator can manage Macie for member accounts only in that Region.

To centrally manage Macie accounts in multiple AWS Regions, the management account must log in to each Region where the organization currently uses or will use Macie, and then designate the Macie administrator account in each of those Regions. The Macie administrator can then configure the organization in each of those Regions. For a list of Regions where Macie is currently available, see [Amazon Macie endpoints and quotas](#) in the *Amazon Web Services General Reference*.

- An account can be associated with only one Macie administrator account at a time. If your organization uses Macie in multiple Regions, the designated Macie administrator account must be the same in all of those Regions. However, your organization's management account must designate the administrator account separately in each Region.
- If the Macie administrator's AWS account is suspended, isolated, or closed, all associated Macie member accounts are automatically removed as Macie member accounts but Macie isn't disabled for the accounts.

Changing or removing the designation of a Macie administrator account

Only the AWS Organizations management account for an organization can change or remove the designation of a delegated Macie administrator account for the organization.

If the management account removes the designation, all associated member accounts are removed as Macie member accounts but Macie isn't disabled for the accounts. For an account to also pause or stop using Macie, a user of the account must suspend (pause) or disable Macie for the account.

Adding and removing Macie member accounts

As you add, remove, and otherwise manage member accounts for your organization, keep the following in mind:

- A Macie administrator account can be associated with no more than 5,000 active (enabled) Macie member accounts in each AWS Region. If your organization exceeds this quota, the Macie administrator won't be able to add member accounts until they remove the necessary number of existing member accounts in the Region.

When an organization meets this quota, we notify the Macie administrator by creating AWS Health and Amazon CloudWatch events for their account. We also send email to the address that's associated with their account.

If you're the Macie administrator for an organization, you can determine how many active member accounts are currently associated with your account by using the **Accounts** page on the Amazon Macie

console or the [DescribeOrganizationConfiguration](#) operation of the Amazon Macie API. For more information, see [Reviewing Amazon Macie accounts for an organization](#) (p. 249).

- An account can be associated with only one Macie administrator account at a time. This means that an account can't accept a Macie invitation from another account if it's already associated with the Macie administrator account for an organization in AWS Organizations.

Similarly, if an account already accepted an invitation, the Macie administrator for an organization in AWS Organizations can't add the account as a Macie member account. The account must first disassociate from its current, invitation-based administrator account.

- To add the AWS Organizations management account as a Macie member account, a user of the management account must first enable Macie for the account. The Macie administrator isn't allowed to enable Macie for the management account.
- A member account can't disassociate from its Macie administrator account. Only the Macie administrator can remove an account as a Macie member account.
- If the Macie administrator removes a Macie member account, Macie continues to be enabled for the account. To also pause or stop using Macie, a user of the account must suspend (pause) or disable Macie for the account.

Transitioning from an invitation-based organization

If you already associated a Macie administrator account with member accounts by using Macie membership invitations, we recommend that you designate that account as the delegated Macie administrator account for your organization in AWS Organizations. This simplifies the transition from an invitation-based organization.

If you do this, all currently associated member accounts continue to be members. If a member account is part of your organization in AWS Organizations, the account's association automatically changes from **By invitation** to **Via AWS Organizations** in Macie. If a member account isn't part of your organization in AWS Organizations, the account's association continues to be **By invitation**. In both cases, the accounts continue to be associated with the delegated Macie administrator account as member accounts.

We recommend this approach because an account can't be associated with more than one Macie administrator account at the same time. If you designate a different account as the Macie administrator account for your organization in AWS Organizations, the designated administrator won't be able to manage accounts that are already associated with another Macie administrator account by invitation. Each member account must first disassociate from its current, invitation-based administrator account. The Macie administrator for your organization in AWS Organizations can then add the account as a Macie member account and begin managing the account.

After you integrate Macie with AWS Organizations and you configure your organization in Macie, you can optionally designate a different Macie administrator account for the organization. You can also continue to use invitations to associate and manage member accounts that aren't part of your organization in AWS Organizations.

Integrating and configuring an organization in Amazon Macie

To start using Amazon Macie with AWS Organizations, the AWS Organizations management account for the organization designates an account as the delegated Macie administrator account for the organization. This enables Macie as a trusted service in AWS Organizations. It also enables Macie in the current AWS Region for the designated administrator account, and it allows the designated administrator account to enable and manage Macie for other accounts in the organization in that same Region. For information about how these permissions are granted, see [Using AWS Organizations with other AWS services](#) in the *AWS Organizations User Guide*.

The delegated Macie administrator then configures the organization in Macie, primarily by adding the organization's accounts as Macie member accounts in the Region. The administrator can then access certain Macie settings, data, and resources for those accounts in that Region.

This topic explains how to designate a delegated Macie administrator for an organization and how to add the organization's accounts as Macie member accounts. Before you perform these tasks, ensure that you understand the [relationship between administrator and member accounts](#) (p. 238). It's also a good idea to review the [considerations and recommendations](#) (p. 241) for using Macie with AWS Organizations.

Tasks

- [Step 1: Verify your permissions](#) (p. 244)
- [Step 2: Designate the delegated Macie administrator account for the organization](#) (p. 245)
- [Step 3: Automatically enable and add new organization accounts as Macie member accounts](#) (p. 246)
- [Step 4: Enable and add existing organization accounts as Macie member accounts](#) (p. 248)

To integrate and configure the organization in multiple Regions, the AWS Organizations management account and the delegated Macie administrator repeat these steps in each additional Region.

Step 1: Verify your permissions

Before you designate the delegated Macie administrator account for your organization, verify that you (as a user of the AWS Organizations management account) are allowed to perform the following AWS Organizations actions:

- `organizations:DescribeOrganization`
- `organizations:EnableAWSServiceAccess`
- `organizations:ListAWSServiceAccessForOrganization`
- `organizations:RegisterDelegatedAdministrator`

These actions allow you to: retrieve information about your organization; integrate Macie with AWS Organizations; retrieve information about which AWS services you've integrated with AWS Organizations; and, designate a delegated Macie administrator account for your organization.

To grant these permissions, append the following statement to an existing Macie policy for your account:

```
{
  "Sid": "Grant permissions to designate a delegated Macie administrator",
  "Effect": "Allow",
  "Action": [
    "organizations:EnableAWSServiceAccess",
    "organizations:DescribeOrganization",
    "organizations:ListAWSServiceAccessForOrganization",
    "organizations:RegisterDelegatedAdministrator"
  ],
  "Resource": "*"
}
```

If you want to designate your AWS Organizations management account as the delegated Macie administrator account for the organization, your account also needs permission to perform the following AWS Identity and Access Management (IAM) action: `CreateServiceLinkedRole`. This action allows you to enable Macie for the management account. However, based on AWS security best practices and the principle of least privilege, we don't recommend that you do this.

If you decide to grant this permission, add the following statement to the IAM policy for your AWS Organizations management account:

```
{
  "Sid": "Grant permissions to enable Macie",
  "Effect": "Allow",
  "Action": [
    "iam:CreateServiceLinkedRole"
  ],
  "Resource": "arn:aws:iam::<111122223333>:role/aws-service-role/macie.amazonaws.com/AWSServiceRoleForAmazonMacie",
  "Condition": {
    "StringLike": {
      "iam:AWSServiceName": "macie.amazonaws.com"
    }
  }
}
```

In the statement, replace `111122223333` with the account ID for the management account.

If you want to administer Macie in a manually enabled AWS Region, also update the value for the Macie service principal in the `Resource` element and the `iam:AWSServiceName` condition key. The value must specify the Region code for the Region. For example, to administer Macie in the Middle East (Bahrain) Region, which has the Region code `me-south-1`, do the following:

- For the `Resource` element, replace

```
arn:aws:iam::<111122223333>:role/aws-service-role/macie.amazonaws.com/AWSServiceRoleForAmazonMacie
```

with

```
arn:aws:iam::<111122223333>:role/aws-service-role/macie.me-south-1.amazonaws.com/AWSServiceRoleForAmazonMacie
```

Where `111122223333` is the account ID for the management account and `me-south-1` specifies the Region code for the Region.

- For the `iam:AWSServiceName` condition key, replace `macie.amazonaws.com` with `macie.me-south-1.amazonaws.com`

Where `me-south-1` specifies the Region code for the Region.

For a list of Regions where Macie is currently available and the Region code for each one, see [Amazon Macie endpoints and quotas](#) in the *Amazon Web Services General Reference*. For information about manually enabled Regions, see [Managing AWS Regions](#) in the *Amazon Web Services General Reference*.

Step 2: Designate the delegated Macie administrator account for the organization

After you verify your permissions, you (as a user of the AWS Organizations management account) can designate the delegated Macie administrator account for your organization.

To designate the delegated Macie administrator account for an organization

To designate the delegated Macie administrator account for your organization, you can use the Amazon Macie console or the Amazon Macie API. Only a user of the AWS Organizations management account can perform this task.

Console

Follow these steps to designate the delegated Macie administrator account by using the Amazon Macie console.

To designate the delegated Macie administrator account

1. Log in to the AWS Management Console using your AWS Organizations management account.
2. By using the AWS Region selector in the upper-right corner of the page, select the Region in which you want to designate the delegated Macie administrator account for your organization.
3. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
4. Do one of the following, depending on whether Macie is enabled for your management account in the current Region:
 - If Macie isn't enabled, choose **Get started** on the welcome page.
 - If Macie is enabled, choose **Settings** in the navigation pane.
5. Under **Delegated administrator**, enter the 12-digit account ID for the AWS account that you want to designate as the Macie administrator account.
6. Choose **Delegate**.

Repeat the preceding steps in each additional Region in which you want to integrate your organization with Macie. You must designate the same Macie administrator account in each of those Regions.

API

To designate the delegated Macie administrator account programmatically, use the [EnableOrganizationAdminAccount](#) operation of the Amazon Macie API. To designate the account in multiple Regions, submit the designation for each Region in which you want to integrate your organization with Macie. You must designate the same Macie administrator account in each of those Regions.

When you submit the designation, use the required `adminAccountId` parameter to specify the 12-digit account ID for the AWS account to designate as the Macie administrator account for the organization. Also ensure that you specify the Region that the designation applies to.

To designate the Macie administrator account by using the [AWS Command Line Interface \(AWS CLI\)](#), run the `enable-organization-admin-account` command. For the `admin-account-id` parameter, specify the 12-digit account ID for the AWS account to designate. Use the `region` parameter to specify the Region that the designation applies to. For example:

```
C:\> aws macie2 enable-organization-admin-account --region us-east-1 --admin-account-id 111122223333
```

Where `us-east-1` is the Region that the designation applies to (the US East (N. Virginia) Region) and `111122223333` is the account ID for the account to designate.

After you designate the Macie administrator account for your organization, the Macie administrator can begin configuring the organization in Macie.

Step 3: Automatically enable and add new organization accounts as Macie member accounts

By default, Macie isn't automatically enabled for new accounts when the accounts are added to your organization in AWS Organizations. In addition, the accounts aren't automatically added as Macie member accounts. The accounts appear in the Macie administrator's account inventory. However, Macie isn't necessarily enabled for the accounts and the Macie administrator can't necessarily access Macie settings, data, and resources for the accounts.

If you're the delegated Macie administrator for the organization, you can change this configuration setting for your organization. If you turn on the **Auto-enable** setting, Macie is automatically enabled for

new accounts when the accounts are added to your organization in AWS Organizations, and the accounts are automatically associated with your Macie administrator account as member accounts. Turning on this setting doesn't affect existing accounts in your organization. To enable and manage Macie for existing accounts, you must manually add the accounts as Macie member accounts. The [next step \(p. 248\)](#) explains how to do this.

Note

If you turn on the **Auto-enable** setting, note the following exceptions:

- If a new account is already associated with a different Macie administrator account, Macie doesn't automatically add the account as a member account in your organization.

The account must disassociate from its current Macie administrator account before it can be part of your organization in Macie. You can then manually add the account. To identify accounts where this is the case, you can [review the account inventory \(p. 249\)](#) for your organization.

- If your organization reaches the quota of 5,000 Macie member accounts in an AWS Region, Macie automatically turns off this setting in the Region.

If this happens, we notify you by creating AWS Health and Amazon CloudWatch events for your Macie administrator account. We also send email to the address that's associated with that account. If the total number of accounts subsequently decreases to fewer than 5,000 accounts, Macie automatically turns on the setting again.

To automatically enable and add new organization accounts as Macie member accounts

To automatically enable and add new accounts as Macie member accounts, you can use the Amazon Macie console or the Amazon Macie API. Only the delegated Macie administrator for the organization can perform this task.

Console

To perform this task by using the console, you must be allowed to perform the following AWS Organizations action: `organizations:listAccounts`. This action allows you to retrieve and display information about the accounts in your organization. If you have these permissions, follow these steps to automatically enable and add new organization accounts as Macie member accounts.

To automatically enable and add new organization accounts

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. By using the AWS Region selector in the upper-right corner of the page, select the Region in which you want to automatically enable and add new accounts as Macie member accounts.
3. In the navigation pane, under **Settings**, choose **Accounts**.
4. On the **Accounts** page, next to **Add accounts**, turn on the **Auto-enable** setting.

Repeat the preceding steps in each additional Region in which you want to configure your organization in Macie.

To subsequently change this setting and stop enabling and adding new accounts automatically, repeat the preceding steps and turn off the **Auto-enable** setting.

API

To automatically enable and add new Macie member accounts programmatically, use the [UpdateOrganizationConfiguration](#) operation of the Amazon Macie API. When you submit your request, set the value for the `autoEnable` parameter to `true`. (The default value is `false`.) Also ensure that you specify the Region that your request applies to. To automatically enable and add new accounts in additional Regions, submit the request for each additional Region.

If you use the AWS CLI to submit the request, run the `update-organization-configuration` command and specify the `auto-enable` parameter to enable and add new accounts automatically. For example:

```
$ aws macie2 update-organization-configuration --region us-east-1 --auto-enable
```

Where `us-east-1` is the Region in which to automatically enable and add new accounts, the US East (N. Virginia) Region.

To subsequently change this setting and stop enabling and adding new accounts automatically, run the same command again and use the `no-auto-enable` parameter, instead of the `auto-enable` parameter, in each applicable Region.

Step 4: Enable and add existing organization accounts as Macie member accounts

When you integrate Macie with AWS Organizations, Macie isn't automatically enabled for all the existing accounts in your organization. In addition, the accounts aren't automatically associated with the delegated Macie administrator account as Macie member accounts.

Therefore, the final step of integrating and configuring your organization in Macie is to add existing organization accounts as Macie member accounts. When you add an existing account as a Macie member account, Macie is automatically enabled for the account and you (as the delegated Macie administrator) gain access to certain Macie settings, data, and resources for the account.

Note that you can't add an account that's currently associated with another Macie administrator account. To add the account, work with the account owner to first disassociate the account from its current administrator account. In addition, you can't add an existing account if Macie is currently suspended for the account. The account owner must first re-enable Macie for the account. Finally, if you want to add the AWS Organizations management account as a member account, a user of that account must first enable Macie for the account.

To enable and add existing organization accounts as Macie member accounts

To enable and add existing organization accounts as Macie member accounts, you can use the Amazon Macie console or the Amazon Macie API. Only the delegated Macie administrator for the organization can perform this task.

Console

To perform this task by using the console, you must be allowed to perform the following AWS Organizations action: `organizations:listAccounts`. This action allows you to retrieve and display information about the accounts in your organization. If you have these permissions, follow these steps to enable and add existing accounts as Macie member accounts.

To enable and add existing organization accounts

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. By using the AWS Region selector in the upper-right corner of the page, select the Region in which you want to enable and add existing accounts as Macie member accounts.
3. In the navigation pane, under **Settings**, choose **Accounts**.

The **Accounts** page opens and displays a table of the accounts that are associated with your Macie account. If an account is part of your organization in AWS Organizations, its **Type** is **Via AWS Organizations**. If an account isn't a Macie member account, its **Status** is **Not a member**.

4. In the **Accounts** table, select the check box for each account that you want to add as a Macie member account.

Tip

To more easily identify accounts to add, you can filter the table. To do this, place your cursor in the filter bar above the table, and then choose **Status**. Then choose **Status = Not a Member**.

5. On the **Actions** menu, choose **Add member**.
6. Confirm that you want to add the selected accounts as member accounts.

After you confirm the addition of the selected accounts, the status of the accounts changes to **Creating/Enabling**, and then **Enabled**.

Repeat the preceding steps in each additional Region in which you want to configure your organization in Macie.

API

To programmatically enable and add one or more existing accounts as Macie member accounts, use the [CreateMember](#) operation of the Amazon Macie API. When you submit your request, use the supported parameters to specify the 12-digit account ID and email address of each AWS account to enable and add. Also specify the Region that the request applies to. To enable and add existing accounts in additional Regions, submit the request for each additional Region.

To retrieve the account ID and email address of an AWS account to enable and add, you can optionally use the [ListMembers](#) operation of the Amazon Macie API. This operation provides details about the accounts that are associated with your Macie account, including accounts that aren't Macie member accounts. If the value for the `relationshipStatus` property of an account isn't `Enabled`, the account isn't a Macie member account.

To enable and add one or more existing accounts by using the AWS CLI, run the [create-member](#) command. Use the `region` parameter to specify the Region in which to enable and add the accounts. Use the `account` parameters to specify the account ID and email address for each AWS account to add. For example:

```
C:\> aws macie2 create-member --region us-east-1 --account={"accountId":  
"123456789012", "email": "janedoe@example.com"}
```

Where `us-east-1` is the Region in which to enable and add the account as a Macie member account (the US East (N. Virginia) Region), and the `account` parameters specify the account ID (`123456789012`) and email address (`janedoe@example.com`) for the account.

If your request succeeds, the status (`relationshipStatus`) of the specified account changes to `Enabled` in your account inventory.

Reviewing Amazon Macie accounts for an organization

After an AWS Organizations organization is [integrated and configured](#) (p. 243) in Amazon Macie, the organization's delegated Macie administrator can access an inventory of the organization's accounts in Macie. As the Macie administrator for an organization, you can use this inventory to review statistics and details for your organization's Macie accounts in an AWS Region. You can also use this inventory to [manage Macie member accounts](#) (p. 252) in a Region.

To review the Macie accounts for an organization

To review the accounts for your organization, you can use the Amazon Macie console or the Amazon Macie API.

Console

Follow these steps to review your organization's Macie accounts by using the Amazon Macie console.

To review your organization's accounts

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. By using the AWS Region selector in the upper-right corner of the page, select the Region in which you want to review your organization's accounts.
3. In the navigation pane, under **Settings**, choose **Accounts**.

The **Accounts** page opens and displays aggregated statistics and a table of the accounts that are associated with your Macie account in the current AWS Region.

At the top of the **Accounts** page, you'll find the following aggregated statistics.

Via AWS Organizations

Active reports the total number of accounts that are associated with your account through AWS Organizations and are currently Macie member accounts in your organization. Macie is enabled for these accounts and you're the Macie administrator of the accounts.

All reports the total number of accounts that are associated with your account through AWS Organizations, including accounts that aren't currently Macie member accounts.

By invitation

Active reports the total number of accounts that are associated with your account by Macie invitation and are currently Macie member accounts. (These accounts aren't associated with your account through AWS Organizations.) Macie is enabled for the accounts and you're the Macie administrator of the accounts because they accepted a Macie membership invitation from you.

All reports the total number of accounts that are associated with your account by Macie invitation, including accounts that haven't responded to an invitation from you.

Active/All

Active reports the total number of accounts that are currently Macie member accounts for your account, either through AWS Organizations or by Macie invitation. Macie is enabled for these accounts and you're the Macie administrator of the accounts.

All reports the total number of accounts that are associated with your account, either through AWS Organizations or by Macie invitation. This includes accounts that are part of your organization in AWS Organizations and aren't currently Macie member accounts, and any accounts that haven't responded to a Macie membership invitation from you.

In the table, you'll find details about each account in the current Region. The table includes all the accounts that are associated with your Macie account, either through AWS Organizations or by Macie invitation.

Account ID

The account ID and email address for the AWS account.

Name

The account name for the AWS account. This value is typically **N/A** for accounts that are associated with your account by Macie invitation.

Type

How the account is associated with your account, through AWS Organizations or by Macie invitation.

Status

The status of the relationship between your account and the account. For an account in an AWS Organizations organization (**Type** is **Via AWS Organizations**), possible values are:

- **Account suspended** – The AWS account is suspended.
- **Created/Enabling** – Macie is processing a request to enable and add the account as a Macie member account.
- **Enabled** – The account is a Macie member account. Macie is enabled for the account and you're the Macie administrator for the account.
- **Not a member** – The account is part of your organization in AWS Organizations but it isn't a Macie member account.
- **Paused (suspended)** – The account is a Macie member account but Macie is currently suspended for the account.
- **Region disabled** – The account is part of your organization in AWS Organizations but the current Region is disabled for the AWS account.
- **Removed (disassociated)** – The account was previously a Macie member account but was subsequently removed as a member account. You disassociated the account from your Macie administrator account. Macie continues to be enabled for the account.

Last action

When you or the associated account most recently performed an action that affected the relationship between your accounts.

To sort the table by a specific field, click the column heading for the field. To change the sort order, click the column heading again. To filter the table, place your cursor in the filter bar, and then add a filter condition for a field. To further refine the results, add filter conditions for additional fields.

API

To review your organization's accounts programmatically, use the [ListMembers](#) operation of the Amazon Macie API and be sure to specify the Region that your request applies to. To review the accounts in additional Regions, submit your request in each additional Region.

When you submit your request, use the `onlyAssociated` parameter to specify which accounts to include in the response. By default, Macie returns details about only those accounts that are Macie member accounts in the specified Region, either through AWS Organizations or by Macie invitation. To retrieve these details for all the accounts that are associated with your Macie account, including accounts that aren't member accounts, include the `onlyAssociated` parameter in your request and set the parameter's value to `false`.

To review your organization's accounts by using the [AWS Command Line Interface \(AWS CLI\)](#), run the `list-members` command. For the `only-associated` parameter, specify whether to include all associated accounts or only Macie member accounts. To include only member accounts, omit this parameter or set the parameter's value to `true`. To include all accounts, set this value to `false`. For example:

```
C:\> aws macie2 list-members --region us-east-1 --only-associated false
```

Where `us-east-1` is the Region that the request applies to, the US East (N. Virginia) Region.

If your request succeeds, Macie returns a `members` array. The array contains a `member` object for each account that meets the criteria specified in the request. In that object, the `relationshipStatus` field indicates the current status of the relationship between your account

and the other account in the specified Region. For an account in an AWS Organizations organization, possible values are:

- `AccountSuspended` – The AWS account is suspended.
- `Created` – Macie is processing a request to enable and add the account as a Macie member account.
- `Enabled` – The account is a Macie member account. Macie is enabled for the account and you're the Macie administrator for the account.
- `Paused` – The account is a Macie member account but Macie is currently suspended (paused) for the account.
- `RegionDisabled` – The account is part of your organization in AWS Organizations but the current Region is disabled for the AWS account.
- `Removed` – The account was previously a Macie member account but was subsequently removed as a member account. You disassociated the account from your Macie administrator account. Macie continues to be enabled for the account.

For information about other fields in the `member` object, see [Members](#) in the *Amazon Macie API Reference*.

Managing Amazon Macie member accounts for an organization

After an AWS Organizations organization is [integrated and configured \(p. 243\)](#) in Amazon Macie, the organization's delegated Macie administrator can access certain Macie settings, data, and resources for member accounts.

As the Macie administrator for an organization, you can also perform certain account management and administration tasks in Macie:

- Manage the status of Macie for individual accounts, including enabling or suspending Macie for an account.
- Add and remove Macie member accounts.
- Monitor Macie quotas and estimated usage costs for individual accounts and the organization overall.

You can also review Amazon Simple Storage Service (Amazon S3) inventory data and policy findings for Macie member accounts. And you can create and run sensitive data discovery jobs to detect sensitive data in S3 buckets that those accounts own. For a detailed list of tasks that you can perform, see [Understanding the relationship between Amazon Macie administrator and member accounts \(p. 238\)](#).

By default, Macie gives you visibility into relevant data and resources for all the Macie member accounts in your organization. You can also drill down to review data and resources for individual accounts. For example, if you [use the Summary dashboard \(p. 20\)](#) to assess your organization's Amazon S3 security posture, you can filter the data by account. Similarly, if you [monitor estimated usage costs \(p. 278\)](#), you can access breakdowns of estimated costs for individual member accounts.

In addition to tasks that are common to administrator and member accounts, you can centrally perform various administrative tasks for your organization.

Tasks

- [Adding Amazon Macie member accounts to an organization \(p. 253\)](#)
- [Suspending Amazon Macie for member accounts in an organization \(p. 254\)](#)
- [Removing Amazon Macie member accounts from an organization \(p. 256\)](#)

As the Macie administrator for an organization, you can perform these tasks by using the Amazon Macie console or the Amazon Macie API. If you prefer to use the console, note that you must be allowed to perform the following AWS Organizations action: `organizations:listAccounts`. This action allows you to retrieve and display information about accounts that are part of your organization in AWS Organizations.

Adding Amazon Macie member accounts to an organization

In some cases, you might need to manually add an account as a Macie member account. This is the case for accounts that you previously removed (disassociated) as member accounts. This is also the case if you didn't configure Macie to [automatically enable and add new accounts as member accounts \(p. 246\)](#) when accounts are added to your organization in AWS Organizations.

When you add an account as a Macie member account, Macie is enabled for the account in the current AWS Region, if it isn't already enabled in that Region, and the account is associated with your Macie administrator account as a member account in the Region. The member account doesn't receive an invitation or other notification that you established this relationship between your accounts.

Note that you can't add an account that's already associated with another Macie administrator account. The account must first disassociate from its current administrator account. In addition, you can't add the AWS Organizations management account as a member account unless the management account has already enabled Macie for the account. To learn about additional requirements, see [Considerations and recommendations for using Amazon Macie with AWS Organizations \(p. 241\)](#).

To add a Macie member account to an organization

To add one or more Macie member accounts to your organization, you can use the Amazon Macie console or the Amazon Macie API.

Console

Follow these steps to add one or more Macie member accounts by using the Amazon Macie console.

To add a Macie member account

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. By using the AWS Region selector in the upper-right corner of the page, select the Region in which you want to add a member account.
3. In the navigation pane, under **Settings**, choose **Accounts**. The **Accounts** page opens and displays a table of the accounts that are associated with your account.
4. (Optional) To more easily identify accounts that are part of your organization in AWS Organizations and aren't Macie member accounts, use the filter bar above the table to add the following filter conditions:

- **Type = Organization**
- **Status = Not a Member**

To also display accounts that you previously removed and might want to add as member accounts, also add a **Status = Removed** filter condition.

5. In the **Accounts** table, select the check box for each account that you want to add as a member account.
6. On the **Actions** menu, choose **Add member**.
7. Confirm that you want to add the selected number of accounts as member accounts.

After you confirm your selections, the status of the selected accounts changes to **Created/Enabling**, and then **Enabled** in your account inventory.

Repeat the preceding steps in each additional Region in which you want to add a member account.

API

To add one or more Macie member accounts programmatically, use the [CreateMember](#) operation of the Amazon Macie API.

When you submit your request, use the supported parameters to specify the 12-digit account ID and email address for each AWS account that you want to add. Also specify the Region that the request applies to. To add an account in additional Regions, submit your request in each additional Region.

To retrieve the account ID and email address of an account to add, you can correlate the output of the [ListAccounts](#) operation of the AWS Organizations API and the [ListMembers](#) operation of the Amazon Macie API. For the **ListMembers** operation of the Macie API, include the `onlyAssociated` parameter in your request and set the parameter's value to `false`. If the operation succeeds, Macie returns a `members` array that provides details about all the accounts that are associated with your Macie administrator account in the specified Region, including accounts that aren't currently member accounts. Note the following in the array:

- If the value for the `relationshipStatus` property of an account isn't `Enabled`, the account is associated with your account but it isn't a Macie member account.
- If an account isn't included in the array but is included in the output of the **ListAccounts** operation of the AWS Organizations API, the account is part of your organization in AWS Organizations but it isn't associated with your account and, therefore, isn't a Macie member account.

To add a member account by using the AWS CLI, run the [create-member](#) command. Use the `region` parameter to specify the Region in which to add the account. Use the `account` parameters to specify the account ID and email address for each account to add. For example:

```
C:\> aws macie2 create-member --region us-east-1 --account={"accountId":  
\"123456789012\", \"email\": \"janedoe@example.com\"}
```

Where `us-east-1` is the Region in which to add the account as a member account (the US East (N. Virginia) Region), and the `account` parameters specify the account ID (`123456789012`) and email address (`janedoe@example.com`) for the account.

If your request succeeds, the status (`relationshipStatus`) of the specified account changes to `Enabled` in your account inventory.

Suspending Amazon Macie for member accounts in an organization

As the Macie administrator for an organization in AWS Organizations, you can suspend Macie for a member account in your organization. If you do this, you can also re-enable Macie for the account at a later time.

When you suspend Macie for a member account:

- Macie loses access to and stops providing metadata about the account's Amazon S3 data in the current AWS Region.
- Macie stops performing all activities for the account in the Region. This includes monitoring S3 buckets and running sensitive data discovery jobs that are currently in progress.
- Macie cancels all sensitive data discovery jobs that were created by the account in the Region. A job can't be resumed or restarted after it's cancelled.

If you created jobs to analyze data that the member account owns, Macie doesn't cancel your jobs. Instead, the jobs skip resources that are owned by the account.

While an account is suspended, Macie retains the Macie session identifier, settings, and resources for the account in the applicable Region. For example, the account's findings remain intact and aren't affected for up to 90 days. Your organization doesn't incur Macie charges for the account in the applicable Region while Macie is suspended for the account in that Region.

To suspend Macie for a member account in an organization

To suspend Macie for a member account in an organization, you can use the Amazon Macie console or the Amazon Macie API.

Console

Follow these steps to suspend Macie for a member account by using the Amazon Macie console.

To suspend Macie for a member account

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. By using the AWS Region selector in the upper-right corner of the page, select the Region in which you want to suspend Macie for the member account.
3. In the navigation pane, under **Settings**, choose **Accounts**.
4. In the **Accounts** table, select the check box for the account to suspend.
5. On the **Actions** menu, choose **Suspend Macie**.
6. Confirm that you want to suspend Macie for the account.

After you confirm the suspension, the status of the account changes to **Paused (suspended)** in your account inventory.

Repeat the preceding steps in each additional Region in which you want to suspend Macie for the account.

API

To suspend Macie for a member account programmatically, use the [UpdateMemberSession](#) operation of the Amazon Macie API.

When you submit your request, use the `id` parameter to specify the 12-digit account ID for the AWS account that you want to suspend Macie for. For the `status` parameter, specify `PAUSED` as the new status for the Macie account. Also specify the Region that the request applies to. To suspend the account in additional Regions, submit your request in each additional Region.

To retrieve the account ID for the account to suspend, you can use the [ListMembers](#) operation of the Amazon Macie API. If you do this, consider filtering the results by including the `onlyAssociated` parameter in your request. If you set this parameter's value to `true`, Macie returns a `members` array that provides details about only those accounts that are currently member accounts.

To suspend Macie for a member account by using the AWS CLI, run the `update-member-session` command. Use the `region` parameter to specify the Region in which to suspend Macie and use the `id` parameter to specify the account ID for the AWS account to suspend Macie for. For the `status` parameter, specify `PAUSED`. For example:

```
C:\> aws macie2 update-member-session --region us-east-1 --id 123456789012 --status PAUSED
```

Where `us-east-1` is the Region in which to suspend Macie (the US East (N. Virginia) Region), `123456789012` is the account ID for the account to suspend Macie for, and `PAUSED` is the new status of Macie for the account.

If your request succeeds, Macie returns an empty response and the status of the specified account changes to `Paused` in your account inventory.

Removing Amazon Macie member accounts from an organization

If you want to stop accessing Macie settings, data, and resources for a member account, you can remove the account as a Macie member account. Note that only you can do this for the account. An AWS Organizations member account can't disassociate from its Macie administrator account.

When you remove a Macie member account, Macie remains enabled for the account in the current AWS Region. However, the account is disassociated from your Macie administrator account and it becomes a standalone Macie account. This means that you lose access to all Macie settings, data, and resources for the account, including metadata and policy findings for the account's Amazon S3 data. This also means that you can no longer use sensitive data discovery jobs to analyze objects in S3 buckets that the account owns. If you already created jobs to do this, the jobs skip buckets that the account owns.

After you remove a Macie member account, the account continues to appear in your account inventory. Macie doesn't notify the account's owner that you removed the account.

To remove a Macie member account from an organization

To remove a Macie member account from your organization, you can use the Amazon Macie console or the Amazon Macie API.

Console

Follow these steps to remove a Macie member account by using the Amazon Macie console.

To remove a Macie member account

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. By using the AWS Region selector in the upper-right corner of the page, select the Region in which you want to remove the member account.
3. In the navigation pane, under **Settings**, choose **Accounts**.
4. In the **Accounts** table, select the check box for the account that you want to remove as a member account.
5. On the **Actions** menu, choose **Disassociate account**.
6. Confirm that you want to remove the selected account as a member account.

After you confirm your selection, the status of the account changes to **Removed (disassociated)** in your account inventory.

Repeat the preceding steps in each additional Region in which you want to remove the member account.

API

To remove a Macie member account programmatically, use the [DisassociateMember](#) operation of the Amazon Macie API.

When you submit your request, use the `id` parameter to specify the 12-digit AWS account ID for the member account to remove. Also specify the Region that the request applies to. To remove the account in additional Regions, submit your request in each additional Region.

To retrieve the account ID for the member account to remove, you can use the [ListMembers](#) operation of the Amazon Macie API. If you do this, consider filtering the results by including the `onlyAssociated` parameter in your request. If you set this parameter's value to `true`, Macie returns a `members` array that provides details about only those accounts that are currently Macie member accounts.

To remove a Macie member account by using the AWS CLI, run the `disassociate-member` command. Use the `region` parameter to specify the Region in which to remove the account. Use the `id` parameter to specify the account ID for the member account to remove. For example:

```
C:\> aws macie2 disassociate-member --region us-east-1 --id 123456789012
```

Where `us-east-1` is the Region in which to remove the account (the US East (N. Virginia) Region) and `123456789012` is the account ID for the account to remove.

If your request succeeds, Macie returns an empty response and the status of the specified account changes to `Removed` in your account inventory.

Designating a different Amazon Macie administrator account for an organization

After an AWS Organizations organization is [integrated and configured](#) (p. 243) in Amazon Macie, the AWS Organizations management account can designate a different account as the delegated Macie administrator account for the organization.

To change the designation, you (as a user of the AWS Organizations management account) must have the [same permissions](#) (p. 244) that were required to initially designate a Macie administrator account for the organization. You must also be allowed to perform the following AWS Organizations action: `organizations:deregisterDelegatedAdministrator`. This action allows you to remove the current designation.

If your organization uses Macie in multiple AWS Regions, you must change the designation in each Region in which your organization uses Macie—the delegated Macie administrator account must be the same in all of those Regions. To learn about additional requirements, see [Considerations and recommendations for using Amazon Macie with AWS Organizations](#) (p. 241).

To designate a different Macie administrator account for your organization

To designate a different Macie administrator account for your organization, you can use the Amazon Macie console or a combination of the Amazon Macie and AWS Organizations APIs. Only a user of the AWS Organizations management account can change the designation.

Console

To change the designation by using the Amazon Macie console, follow these steps.

To designate a different Macie administrator account

1. Log in to the AWS Management Console using your AWS Organizations management account.
2. By using the AWS Region selector in the upper-right corner of the page, select the Region in which you want to change the designation.
3. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
4. Do one of the following, depending on whether Macie is enabled for your management account in the current Region:
 - If Macie isn't enabled, choose **Get started** on the welcome page.
 - If Macie is enabled, choose **Settings** in the navigation pane.
5. Under **Delegated administrator**, choose **Remove**. To change the designation, you must first remove the current designation.
6. Confirm that you want to remove the current designation.

7. Under **Delegated administrator**, enter the 12-digit account ID for the AWS account to designate as the new Macie administrator account for the organization.
8. Choose **Delegate**.

Repeat the preceding steps in each additional Region in which you integrated Macie with AWS Organizations.

API

To change the designation programmatically, you use two operations of the Amazon Macie API and one operation of the AWS Organizations API. This is because you have to remove the current designation in both Macie and AWS Organizations before you submit the new designation.

To remove the current designation:

1. Use the [DisableOrganizationAdminAccount](#) operation of the Macie API. For the required `adminAccountId` parameter, specify the 12-digit account ID for the AWS account that's currently designated as the Macie administrator account for the organization.
2. Use the [DeregisterDelegatedAdministrator](#) operation of the AWS Organizations API. For the `AccountId` parameter, specify the 12-digit account ID for the account that's currently designated as the Macie administrator account for the organization. This value should match the account ID that you specified in the preceding Macie request. For the `ServicePrincipal` parameter, specify the Macie service principal (`macie.amazonaws.com`).

After you remove the current designation, submit the new designation by using the [EnableOrganizationAdminAccount](#) operation of the Macie API. For the required `adminAccountId` parameter, specify the 12-digit account ID for the AWS account to designate as the new Macie administrator account for the organization.

To change the designation by using the [AWS CLI](#), run the `disable-organization-admin-account` command of the Macie API and the `deregister-delegated-administrator` command of the AWS Organizations API. These commands remove the current designation in Macie and AWS Organizations, respectively. For the `admin-account-id` and `account-id` parameters, specify the 12-digit account ID for the AWS account to remove as the current Macie administrator account. Use the `region` parameter to specify the Region that the removal applies to. For example:

```
C:\> aws macie2 disable-organization-admin-account --region us-east-1 --admin-account-id 111122223333 && aws organizations deregister-delegated-administrator --region us-east-1 --account-id 111122223333 --service-principal macie.amazonaws.com
```

Where:

- `us-east-1` is the Region that the removal applies to, the US East (N. Virginia) Region.
- `111122223333` is the account ID for the account to remove as the Macie administrator account.
- `macie.amazonaws.com` is the Macie service principal.

After you remove the current designation, submit the new designation by running the `enable-organization-admin-account` command of the Macie API. For the `admin-account-id` parameter, specify the 12-digit account ID for the AWS account to designate as the new Macie administrator account for the organization. Use the `region` parameter to specify the Region that the designation applies to. For example:

```
C:\> aws macie2 enable-organization-admin-account --region us-east-1 --admin-account-id 444455556666
```


Where `us-east-1` is the Region that the designation applies to (the US East (N. Virginia) Region) and `444455556666` is the account ID for the account to designate as the new Macie administrator account.

Disabling Amazon Macie integration with AWS Organizations

After an AWS Organizations organization is integrated with Amazon Macie, the AWS Organizations management account can subsequently disable the integration. As a user of the AWS Organizations management account, you can do this by disabling trusted service access for Macie in AWS Organizations.

When you disable trusted service access for Macie, the following occurs:

- Macie loses its status as a trusted service in AWS Organizations.
- The organization's Macie administrator account loses access to all Macie settings, data, and resources for all Macie member accounts in all AWS Regions.
- All Macie member accounts become standalone Macie accounts. If Macie was enabled for a member account in one or more Regions, Macie continues to be enabled for the account in those Regions. However, the account is no longer associated with a Macie administrator account in any Region.

For additional information about the results of disabling trusted service access, see [Using AWS Organizations with other AWS services](#) in the *AWS Organizations User Guide*.

To disable trusted service access for Macie

To disable trusted service access, you can use the AWS Organizations console or the AWS Organizations API. Only a user of the AWS Organizations management account can disable trusted service access for Macie. For details about the permissions that you need, see [Permissions required to disable trusted access](#) in the *AWS Organizations User Guide*.

Before you disable trusted service access, optionally work with the delegated Macie administrator for your organization to suspend or disable Macie for member accounts and to clean up Macie resources for those accounts.

Console

To disable trusted service access by using the AWS Organizations console, follow these steps.

To disable trusted service access

1. Log in to the AWS Management Console using your AWS Organizations management account.
2. Open the AWS Organizations console at <https://console.aws.amazon.com/organizations/>.
3. In the navigation pane, choose **Services**.
4. Under **Integrated services**, choose **Amazon Macie**.
5. Choose **Disable trusted access**.
6. Confirm that you want to disable trusted access.

API

To disable trusted service access programmatically, use the [DisableAWSServiceAccess](#) operation of the AWS Organizations API. For the `ServicePrincipal` parameter, specify the Macie service principal (`macie.amazonaws.com`).

To disable trusted service access by using the [AWS Command Line Interface \(AWS CLI\)](#), run the `disable-aws-service-access` command of the AWS Organizations API. For the `service-principal` parameter, specify the Macie service principal (`macie.amazonaws.com`). For example:

```
C:\> aws organizations disable-aws-service-access --service-principal
macie.amazonaws.com
```

Managing Amazon Macie accounts by invitation

You can centrally manage multiple Amazon Macie accounts in two ways, by [integrating Macie with AWS Organizations \(p. 240\)](#) or by using membership invitations. If you use membership invitations, a designated Macie administrator can manage Macie for as many as 1,000 accounts. The administrator can also access Amazon Simple Storage Service (Amazon S3) inventory data and run sensitive data discovery jobs for the accounts. For details about the tasks that administrators can perform, see [Understanding the relationship between Amazon Macie administrator and member accounts \(p. 238\)](#).

In an invitation-based organization, you associate Macie accounts with each other by sending and accepting membership invitations in Macie. If you send an invitation and it's accepted by another account, you become the Macie administrator for the other account and the other account becomes a member account in your organization. If you receive and accept an invitation, your account becomes a member account and the Macie administrator can access certain Macie settings, data, and resources for your account.

Tip

If you create an invitation-based organization in Macie, you can subsequently [transition to using AWS Organizations \(p. 262\)](#) instead. You can also use both methods at the same time to manage multiple Macie accounts. For example, if your AWS environment includes test accounts, you might exclude the accounts from your organization in AWS Organizations and manage them separately by invitation.

The topics in this section explain how to create and participate in an invitation-based organization, and how to perform various administrative tasks for the organization.

Topics

- [Considerations and recommendations for invitation-based organizations in Amazon Macie \(p. 260\)](#)
- [Creating and managing an invitation-based organization in Amazon Macie \(p. 263\)](#)
- [Reviewing Amazon Macie accounts for an invitation-based organization \(p. 271\)](#)
- [Designating a different Amazon Macie administrator account for an invitation-based organization \(p. 273\)](#)
- [Managing your membership in an invitation-based organization in Amazon Macie \(p. 274\)](#)

Considerations and recommendations for invitation-based organizations in Amazon Macie

Before you create or begin managing an invitation-based organization in Amazon Macie, consider the following requirements and recommendations.

Topics

- [Choosing a Macie administrator account \(p. 261\)](#)
- [Sending invitations and managing Macie member accounts \(p. 261\)](#)
- [Responding to and managing membership invitations \(p. 262\)](#)
- [Transitioning to AWS Organizations \(p. 262\)](#)

Choosing a Macie administrator account

While you determine which account should be the Macie administrator account for the organization, keep the following in mind:

- An organization can have only one Macie administrator account.
- An account can't be a Macie administrator and member account at the same time.
- Macie is a Regional service. This means that the association between a Macie administrator account and a member account is Regional—the association exists only in the AWS Region that an invitation is sent from and accepted in. For example, if the Macie administrator sends invitations in the US East (N. Virginia) Region and those invitations are accepted, the Macie administrator can manage the member accounts only in that Region.

To centrally manage Macie accounts in multiple AWS Regions, the Macie administrator can log in to each Region where the organization currently uses or will use Macie and send invitations to the appropriate accounts in each of those Regions. For a list of Regions where Macie is currently available, see [Amazon Macie endpoints and quotas](#) in the *Amazon Web Services General Reference*.

- A member account can be associated with only one Macie administrator account at a time. If your organization uses Macie in multiple Regions, this means that the Macie administrator account must be the same in all of those Regions. However, administrator and member accounts must send and accept invitations separately in each Region.
- If the Macie administrator's AWS account is suspended, isolated, or closed, all associated member accounts are automatically removed as member accounts but Macie continues to be enabled for those accounts.

Sending invitations and managing Macie member accounts

As the Macie administrator for an invitation-based organization, keep the following in mind when you send invitations and manage accounts in the organization:

- Before you send an invitation, ensure that you [understand the relationship between Macie administrator and member accounts \(p. 238\)](#).
- If you send an invitation, related data might be transferred across AWS Regions. This is the case because Macie verifies the receiving account's email address by using an email verification service that operates only in the US East (N. Virginia) Region.
- You can send an invitation to any active AWS account, including accounts that haven't enabled Macie. However, to accept or decline an invitation, the receiving account must enable Macie in the Region that the invitation was sent from.
- A Macie administrator account can be associated with no more than 1,000 accounts in each AWS Region. This includes accounts that haven't responded to invitations yet. If your account meets this quota, you can't add or invite additional accounts until you remove the necessary number of associated accounts, receive the necessary number of declined invitations, or a combination of the two.

To determine how many accounts are currently associated with your account, you can use the **Accounts** page on the Amazon Macie console or the [ListMembers](#) operation of the Amazon Macie API. For more information, see [Reviewing Amazon Macie accounts for an invitation-based organization \(p. 271\)](#).

- An account can be associated with only one Macie administrator account at a time. This means that an account can't accept your invitation if it's already associated with another Macie administrator account. The account must first disassociate from its current Macie administrator account.
- In an invitation-based organization, a member account can disassociate from its Macie administrator account at any time. If this happens, Macie continues to be enabled for the account and the account becomes a standalone Macie account. Macie doesn't notify you if a member account disassociates from

your administrator account. However, the account continues to appear in your account inventory and it has a status of **Member resigned**.

- If you remove a member account from your organization, Macie continues to be enabled for the account and the account becomes a standalone Macie account.

Responding to and managing membership invitations

As a recipient of an invitation or a member of an invitation-based organization, keep the following in mind when you respond to and manage invitations that you receive:

- Before you accept an invitation, ensure that you [understand the relationship between Macie administrator and member accounts \(p. 238\)](#).
- Your account can be associated with only one Macie administrator account at a time. If you accept an invitation and subsequently want to join another organization (by invitation or through AWS Organizations), you have to first disassociate your account from its current Macie administrator account. You can then join the other organization.
- To accept or decline an invitation, you have to enable Macie in the AWS Region that the invitation was sent from. The account that sent the invitation can't enable Macie in that Region for you. Declining an invitation is optional. If you decline an invitation, you can optionally disable Macie in the applicable Region after you decline the invitation.
- If you're a Macie administrator, you can't accept an invitation to become a member account—an account can't be a Macie administrator and member account at the same time. To become a member account, you must first disassociate your account from all of its member accounts by removing all member accounts from your current organization.
- Macie is a Regional service. If you accept an invitation, the association between your account and the Macie administrator account is Regional—the association exists only in the AWS Region that the invitation was sent from and accepted in.
- If you use Macie in multiple Regions, the Macie administrator account for your account has to be the same in all of those Regions. However, the Macie administrator has to send invitations to you separately in each Region, and you have to accept the invitations separately in each Region.
- You can disassociate your account from a Macie administrator account at any time. If you do this, Macie continues to be enabled for your account and your account becomes a standalone Macie account.
- If your Macie administrator removes your account from their organization, Macie continues to be enabled for your account and your account becomes a standalone Macie account.

Transitioning to AWS Organizations

After you create an invitation-based organization in Macie, you can transition to using AWS Organizations instead. To simplify the transition, we recommend that you designate the existing, invitation-based administrator account as the Macie administrator account for the organization in AWS Organizations.

If you do this, all currently associated member accounts continue to be members. If a member account is part of the organization in AWS Organizations, the account's association automatically changes from **By invitation** to **Via AWS Organizations** in Macie. If a member account isn't part of the organization in AWS Organizations, the account's association continues to be **By invitation**. In both cases, the accounts continue to be associated with the Macie administrator account as member accounts.

We recommend this approach because a member account can be associated with only one Macie administrator account at a time. If you designate a different account as the Macie administrator account for an organization in AWS Organizations, the designated administrator won't be able to manage accounts that are already associated with another Macie administrator account by invitation. Each member account must first disassociate from its current, invitation-based administrator account. Only

then can the Macie administrator for the AWS Organizations organization add the member account to their organization and begin managing Macie for the account.

After you integrate Macie with AWS Organizations and you configure your organization in Macie, you can optionally designate a different Macie administrator account for the organization. You can also continue to use invitations to associate and manage member accounts that aren't part of your organization in AWS Organizations.

Creating and managing an invitation-based organization in Amazon Macie

To create an invitation-based organization in Amazon Macie, you start by determining which account you want to be the Macie administrator account for the organization. You then use that account to add member accounts—you send membership invitations to other AWS accounts, inviting the accounts to join the organization as Macie member accounts in the current AWS Region. To create the organization in multiple Regions, send membership invitations from each Region in which the other accounts currently or will use Macie.

When an account accepts an invitation, it becomes a Macie member account that's associated with the Macie administrator account in the applicable Region. The Macie administrator account can then access certain Macie settings, data, and resources for the member account in that Region.

As the Macie administrator for an invitation-based organization, you can review Amazon Simple Storage Service (Amazon S3) inventory data and policy findings for member accounts. You can also create and run sensitive data discovery jobs to detect sensitive data in S3 buckets that member accounts own. For a detailed list of the tasks that you can perform, see [Understanding the relationship between Amazon Macie administrator and member accounts \(p. 238\)](#).

By default, Macie gives you visibility into relevant data and resources for your organization overall. You can also drill down to review data and resources for individual accounts in your organization. For example, if you [use the Summary dashboard \(p. 20\)](#) to assess your organization's Amazon S3 security posture, you can filter the data by account. Similarly, if you [monitor estimated usage costs \(p. 278\)](#), you can access breakdowns of estimated costs for individual member accounts.

In addition to tasks that are common to administrator and member accounts, you can centrally perform various administrative tasks for your organization. Before you perform these tasks, it's a good idea to review the [considerations and recommendations \(p. 260\)](#) for managing invitation-based organizations in Macie.

Tasks

- [Adding Amazon Macie member accounts to an invitation-based organization \(p. 263\)](#)
- [Suspending Amazon Macie for member accounts in an invitation-based organization \(p. 267\)](#)
- [Removing Amazon Macie member accounts from an invitation-based organization \(p. 268\)](#)
- [Deleting associations with other accounts \(p. 269\)](#)

Adding Amazon Macie member accounts to an invitation-based organization

As the Macie administrator for an invitation-based organization, you add member accounts to your organization by performing two primary steps:

1. Add the accounts to your account inventory in Macie. This associates the accounts with your account.
2. Send membership invitations to the accounts.

When an account accepts your invitation, it becomes a member account in your organization.

Step 1: Add the accounts

To add one or more accounts to your account inventory, you can use the Amazon Macie console or the Amazon Macie API.

Console

With the Amazon Macie console, you can add one account at a time, or add multiple accounts at the same time by uploading a comma-separated values (CSV) file. Follow these steps to add one or more accounts by using the console.

To add one account

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. By using the AWS Region selector in the upper-right corner of the page, select the Region in which you want to add an account.
3. In the navigation pane, under **Settings**, choose **Accounts**.
4. Choose **Add accounts**.
5. In the **Enter account details** section, choose the **Add account** tab. Then do the following:
 - For **Account ID**, enter the 12-digit account ID for the AWS account to add.
 - For **Email address**, enter the email address for the AWS account to add.
6. Choose **Add**, and then choose **Next**.

Macie adds the account to your account inventory. The account's type is **By invitation** and its status is **Created**. Repeat the preceding steps in each additional Region in which you want to add the account.

To add multiple accounts

1. By using a text editor, create a CSV file as follows:
 - a. Add the following header as the first line of the file: `Account ID,Email`
 - b. For each account, create a new line that has the 12-digit account ID for the AWS account to add and the email address for the account. Separate the entries with a comma, for example:
`111111111111,janedoe@example.com`

The email address must match the email address that's associated with the AWS account.

- c. Verify that the file's contents are formatted as shown in the following example, which contains the required header and information for three accounts:

```
Account ID,Email
111111111111,janedoe@example.com
222222222222,jorgesouza@example.com
333333333333,lijuan@example.com
```

- d. Save the file on your computer.
2. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
 3. By using the AWS Region selector in the upper-right corner of the page, select the Region in which you want to add the accounts.
 4. In the navigation pane, under **Settings**, choose **Accounts**.
 5. Choose **Add accounts**.
 6. In the **Enter account details** section, choose the **Upload list (CSV)** tab.

7. Choose **Browse**, and then select the CSV file that you created in step 1.
8. Choose **Add accounts**, and then choose **Next**.

Macie adds the accounts to your account inventory. Their type is **By invitation** and their status is **Created**. Repeat steps 3 through 8 in each additional Region in which you want to add the accounts.

API

To add one or more accounts programmatically, use the [CreateMember](#) operation of the Amazon Macie API. When you submit your request, use the supported parameters to specify the 12-digit account ID and email address for each AWS account to add. Also specify the Region that the request applies to. To add accounts in additional Regions, submit the request in each additional Region.

To add accounts by using the [AWS Command Line Interface \(AWS CLI\)](#), run the `create-member` command. Use the `region` parameter to specify the Region in which to add the accounts. Use the account parameters to specify the account ID and email address for each AWS account to add. For example:

```
C:\> aws macie2 create-member --region us-east-1 --account="{\"accountId\":  
\"111111111111\", \"email\": \"janedoe@example.com\"}"
```

Where `us-east-1` is the Region in which to add the account (the US East (N. Virginia) Region) and the account parameters specify the account ID (`111111111111`) and email address (`janedoe@example.com`) for the account to add.

If your request succeeds, Macie adds each account to your account inventory with a status of `Created` and you receive output similar to the following:

```
{  
  "arn": "arn:aws:macie2:us-east-1:123456789012:member/111111111111"  
}
```

Where `arn` is the Amazon Resource Name (ARN) of the resource that was created for the association between your account and the account that you added. In this example, `123456789012` is the account ID for the account that created the association and `111111111111` is the account ID for the account that was added.

Step 2: Send membership invitations to the accounts

After you add an account to your account inventory, you can invite the account to join your organization as a Macie member account. To do this, send a membership invitation to the account. When you send an invitation, an **Accounts** badge and notification appear on the Amazon Macie console for the recipient's account, if Macie is enabled for the account. Macie also creates an AWS Health event for the account.

Depending on whether you use the Amazon Macie console or API to send the invitation, Macie also sends the invitation to the email address that you specified for the recipient's account when you added the account. The email message indicates that you would like to become the Macie administrator for their account, and it includes the account ID for your AWS account and the recipient's AWS account. The message also explains how to access the invitation. You can optionally add custom text to the message.

To send a membership invitation to one or more accounts, you can use the Amazon Macie console or the Amazon Macie API.

Console

Follow these steps to send a membership invitation by using the Amazon Macie console.

To send a membership invitation

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. By using the AWS Region selector in the upper-right corner of the page, select the Region in which you want to send the invitation.
3. In the navigation pane, under **Settings**, choose **Accounts**.
4. In the **Accounts** table, select the check box for each account that you want to send the invitation to.

Tip

To more easily identify accounts that you added and haven't sent invitations to yet, you can filter the table. To do this, place your cursor in the filter bar above the table, and then choose **Status**. Then choose **Status = Created**.

5. On the **Actions** menu, choose **Invite**.
6. (Optional) In the **Message** box, enter any custom text that you want to include in the email message that contains the invitation. The text can contain as many as 80 alphanumeric characters.
7. Choose **Invite**.

To send the invitation in additional AWS Regions, repeat the preceding steps in each additional Region.

After you send the invitation, the status of a recipient account changes to **Email verification in progress** in your account inventory. If Macie can verify an account's email address, the account's status subsequently changes to **Invited**. If Macie can't verify the address, the account's status changes to **Email verification failed**. If this happens, work with the account owner to get the correct email address. Then [delete the association between your accounts \(p. 269\)](#), [add the account \(p. 264\)](#) again, and send the invitation again.

When a recipient accepts an invitation, the status of the recipient's account changes to **Enabled** in your account inventory. If a recipient declines an invitation, the recipient's account is disassociated from your account and removed from your account inventory.

API

To send an invitation programmatically, use the [CreateInvitations](#) operation of the Amazon Macie API. When you submit your request, use the supported parameters to specify the 12-digit account ID for each AWS account to send the invitation to. An account ID must match the account ID for an account in your account inventory. Otherwise, an error occurs. Also specify the Region to send the invitation from. To send the invitation from additional Regions, submit the request in each additional Region.

In your request, you can also specify whether to send the invitation as an email message, and whether to include custom text in that message. If you choose to send an email message, Macie sends the invitation to the email address that you specified for an account when you added the account to your account inventory. To send the invitation as an email message, omit the `disableEmailNotification` parameter or set the value for the parameter to `false`. (The default value is `false`.) To add custom text to the message, use the `message` parameter to specify the text to add. The text can contain as many as 80 alphanumeric characters.

To send invitations by using the AWS CLI, run the [create-invitations](#) command. Use the `region` parameter to specify the Region to send the invitation from. Use the `account-ids` parameter to specify the account ID for each AWS account to send the invitation to. For example:

```
C:\> aws macie2 create-invitations --region us-east-1 --account-ids=["111111111111",  
"222222222222","333333333333"]
```


Where `us-east-1` is the Region to send the invitation from (the US East (N. Virginia) Region) and the `account-ids` parameter specifies account IDs for three accounts to send the invitation to. To send an invitation as an email message too, also include the `no-disable-email-notification` parameter and optionally include the `message` parameter to specify custom text to add to the message.

After you send the invitation, the status of each recipient account changes to `EmailVerificationInProgress`. If Macie can verify an account's email address, the account's status subsequently changes to `Invited`. If Macie can't verify the address, the account's status changes to `EmailVerificationFailed`. If this happens, work with the account owner to get the correct address. Then [delete the association between your accounts \(p. 269\)](#), [add the account \(p. 264\)](#) again, and send the invitation again.

When a recipient accepts an invitation, the status of the recipient's account changes to `Enabled` in your account inventory. If a recipient declines an invitation, the recipient's account is disassociated from your account and removed from your account inventory.

Suspending Amazon Macie for member accounts in an invitation-based organization

As the Macie administrator for an organization, you can suspend Macie in a specific AWS Region for individual member accounts in your organization. Note, however, that you can't re-enable Macie for a member account after you suspend it. Only a user of the account can subsequently re-enable Macie for the account.

When you suspend Macie for a member account:

- Macie loses access to and stops providing metadata about the account's Amazon S3 data in the Region.
- Macie stops performing all activities for the account in the Region. This includes monitoring S3 buckets and running sensitive data discovery jobs that are currently in progress.
- Macie cancels all sensitive data discovery jobs that were created by the account in the Region. A job can't be resumed or restarted after it's cancelled.

If you created jobs to analyze data that the member account owns, Macie doesn't cancel those jobs. Instead, the jobs skip resources that are owned by the account.

While an account is suspended, Macie retains the Macie session identifier, settings, and resources for the account in the applicable Region. For example, the account's findings remain intact and aren't affected for up to 90 days. The account isn't charged for using Macie in the applicable Region while Macie is suspended for the account in that Region.

To suspend Macie for a member account in an invitation-based organization

To suspend Macie for a member account in an invitation-based organization, you can use the Amazon Macie console or the Amazon Macie API.

Console

Follow these steps to suspend Macie for a member account by using the Amazon Macie console.

To suspend Macie for a member account

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. By using the AWS Region selector in the upper-right corner of the page, select the Region in which you want to suspend Macie for a member account.
3. In the navigation pane, under **Settings**, choose **Accounts**.

4. In the **Accounts** table, select the check box for the account to suspend.
5. On the **Actions** menu, choose **Suspend Macie**.
6. Confirm that you want to suspend Macie for the selected account.

After you confirm the suspension, the status of the account changes to **Paused (suspended)** in your account inventory.

Repeat the preceding steps in each additional Region in which you want to suspend Macie for the account.

API

To suspend Macie for a member account programmatically, use the [UpdateMemberSession](#) operation of the Amazon Macie API. When you submit your request, use the `id` parameter to specify the 12-digit account ID of the AWS account that you want to suspend Macie for. For the `status` parameter, specify `PAUSED` as the new status for the Macie account. Also specify the Region that the request applies to. To suspend Macie in additional Regions, submit your request in each additional Region.

To retrieve the account ID for the member account, you can use the [ListMembers](#) operation of the Amazon Macie API. If you do this, consider filtering the results by including the `onlyAssociated` parameter in your request. If you set this parameter's value to `true`, Macie returns a `members` array that provides details about only those accounts that are currently member accounts for your administrator account.

To suspend Macie for a member account by using the AWS CLI, run the `update-member-session` command. Use the `region` parameter to specify the Region in which to suspend Macie and use the `id` parameter to specify the account ID for the account to suspend Macie for. For the `status` parameter, specify `PAUSED`. For example:

```
C:\> aws macie2 update-member-session --region us-east-1 --id 123456789012 --status PAUSED
```

Where `us-east-1` is the Region in which to suspend Macie (the US East (N. Virginia) Region), `123456789012` is the account ID for the account to suspend Macie for, and `PAUSED` is the new status of Macie for the account.

If your request succeeds, Macie returns an empty response and the status of the specified account changes to `Paused` in your account inventory.

Removing Amazon Macie member accounts from an invitation-based organization

As a Macie administrator, you can remove a member account from your organization. You do this by disassociating the account from your Macie administrator account.

If you remove a member account, Macie continues to be enabled for the account and the account continues to appear in your account inventory. However, the account becomes a standalone Macie account. Macie doesn't notify the account's owner when you remove the account. Therefore, consider contacting the account owner to ensure that they begin managing settings and resources for their account.

When you remove a member account, you lose access to all Macie settings, resources, and data for the account. This includes findings and metadata for S3 buckets that the account owns. In addition, you can no longer use sensitive data discovery jobs to analyze objects in S3 buckets that the account owns. If you already created jobs to do this, the jobs skip buckets that the account owns.

After you remove a member account, you can subsequently add it to your organization again by sending a new invitation to the account. You can also remove it from your account inventory completely by deleting the association between your accounts.

To remove a member account from an invitation-based organization

To remove a member account from your organization, you can use the Amazon Macie console or the Amazon Macie API.

Console

Follow these steps to remove a member account by using the Amazon Macie console.

To remove a member account

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. By using the AWS Region selector in the upper-right corner of the page, select the Region in which you want to remove the member account.
3. In the navigation pane, under **Settings**, choose **Accounts**.
4. In the **Accounts** table, select the check box for the account to remove.
5. On the **Actions** menu, choose **Disassociate account**.
6. Confirm that you want to remove the selected account as a member account.

After you confirm your selection, the status of the account changes to **Removed (disassociated)** in your account inventory.

Repeat the preceding steps in each additional Region in which you want to remove the member account.

API

To remove a member account programmatically, use the [DisassociateMember](#) operation of the Amazon Macie API. When you submit your request, use the `id` parameter to specify the 12-digit AWS account ID for the member account to remove. Also specify the Region that the request applies to. To remove the account in additional Regions, submit your request in each additional Region.

To retrieve the account ID for the account to remove, you can use the [ListMembers](#) operation of the Amazon Macie API. If you do this, consider filtering the results by including the `onlyAssociated` parameter in your request. If you set this parameter's value to `true`, Macie returns a `members` array that provides details about only those accounts that are currently member accounts for your account.

To remove a member account by using the AWS CLI, run the `disassociate-member` command. Use the `region` parameter to specify the Region in which to remove the account. Use the `id` parameter to specify the account ID for the account to remove. For example:

```
C:\> aws macie2 disassociate-member --region us-east-1 --id 123456789012
```

Where `us-east-1` is the Region in which to remove the account (the US East (N. Virginia) Region) and `123456789012` is the account ID for the account to remove.

If your request succeeds, Macie returns an empty response and the status of the specified account changes to `Removed` in your account inventory.

Deleting associations with other accounts

After you add an account to your account inventory, you can delete the association between your account and the other account. You can do this for any account in your inventory except:

- An account that's part of your organization in AWS Organizations. This type of association is controlled through AWS Organizations not Macie.
- A member account that accepted a Macie membership invitation to join your organization. If this is the case, you must [remove the member account \(p. 268\)](#) before you can delete the association.

When you delete an association, Macie removes the account from your account inventory. If you subsequently want to restore the association, you have to add the account again as if it were a completely new account.

To delete an association with another account

To delete an association between your account and another account, you can use the Amazon Macie console or the Amazon Macie API.

Console

To use the Amazon Macie console to delete an association with another account, follow these steps.

To delete an association

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. By using the AWS Region selector in the upper-right corner of the page, select the Region in which you want to delete the association.
3. In the navigation pane, under **Settings**, choose **Accounts**.
4. In the **Accounts** table, select the check box for the account whose association you want to delete.
5. On the **Actions** menu, choose **Delete account**.
6. Confirm that you want to delete the selected association.

Repeat the preceding steps in each additional Region in which you want to delete the association.

API

To delete an association with another account programmatically, use the [DeleteMember](#) operation of the Amazon Macie API. When you submit your request, use the `id` parameter to specify the 12-digit account ID for the AWS account to delete the association with. Also specify the Region that the request applies to. To delete the association in additional Regions, submit your request in each additional Region.

To retrieve the account ID for the account, you can use the [ListMembers](#) operation of the Amazon Macie API. If you do this, include the `onlyAssociated` parameter in your request and set the parameter's value to `false`. If the operation is successful, Macie returns a `members` array that provides details about all the accounts that are associated with your account, including accounts that aren't currently member accounts.

To delete an association with another account by using the AWS CLI, run the [delete-member](#) command. Use the `region` parameter to specify the Region in which to delete the association and the `id` parameter to specify the account ID for the account. For example:

```
C:\> aws macie2 delete-member --region us-east-1 --id 123456789012
```

Where `us-east-1` is the Region in which to delete the association with the other account (the US East (N. Virginia) Region) and `123456789012` is the account ID for the account.

If your request succeeds, Macie returns an empty response and the association between your account and the other account is deleted. The previously associated account is removed from your account inventory.

Reviewing Amazon Macie accounts for an invitation-based organization

To help you manage the accounts in your organization, Amazon Macie provides an inventory of the accounts that are associated with your Macie account in each AWS Region where you use Macie. By using this inventory, you can check the status of individual accounts and review account statistics and details for your organization. You can also manage the status of the relationship between your account and individual accounts.

To review accounts for an invitation-based organization

To review the accounts in your organization, you can use the Amazon Macie console or the Amazon Macie API.

Console

Follow these steps to review your organization's accounts by using the Amazon Macie console.

To review your organization's accounts

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. By using the AWS Region selector in the upper-right corner of the page, select the Region in which you want to review your organization's accounts.
3. In the navigation pane, under **Settings**, choose **Accounts**.

The **Accounts** page opens and displays aggregated statistics and a table of the accounts that are associated with your Macie account in the current AWS Region.

At the top of the **Accounts** page, you'll find the following aggregated statistics.

Via AWS Organizations

If you're the Macie administrator for an organization in AWS Organizations, **Active** reports the total number of accounts that are associated with your account through AWS Organizations and are currently Macie member accounts in your organization. Macie is enabled for these accounts and you're the Macie administrator of the accounts.

All reports the total number of accounts that are associated with your account through AWS Organizations, including accounts that aren't currently Macie member accounts.

By invitation

Active reports the total number of accounts that are currently Macie member accounts in your invitation-based organization. Macie is enabled for these accounts and you're the Macie administrator of the accounts because they accepted a membership invitation from you.

All reports the total number of accounts that are associated with your account by Macie invitation, including accounts that haven't responded to an invitation from you.

Active/All

Active reports the total number of accounts that are currently Macie member accounts for your account, either through AWS Organizations or by invitation. Macie is enabled for these accounts and you're the Macie administrator of the accounts.

All reports the total number of accounts that are associated with your account, either through AWS Organizations or by invitation. This includes accounts that haven't accepted a Macie

membership invitation from you. This also includes accounts that are associated with your account through AWS Organizations and aren't currently Macie member accounts.

In the table, you'll find details about each account in the current Region. The table includes all the accounts that are associated with your Macie account, either by Macie invitation or through AWS Organizations.

Account ID

The account ID and email address for the AWS account.

Name

The account name for the AWS account. This value is typically **N/A** for accounts that are associated with your account by invitation.

Type

How the account is associated with your account, by invitation or through AWS Organizations.

Status

The status of the relationship between your account and the account. For an account in an invitation-based organization (**Type** is **By invitation**), possible values are:

- **Account suspended** – The AWS account is suspended.
- **Created (Invite)** – You added the account but haven't sent a membership invitation to it.
- **Email verification failed** – You tried to send a membership invitation to the account but the specified email address isn't valid for the account.
- **Email verification in progress** – You sent a membership invitation to the account and Macie is processing the request.
- **Enabled** – The account is a member account. Macie is enabled for the account and you're the Macie administrator of the account.
- **Invited** – You sent a membership invitation to the account and the account hasn't responded to your invitation.
- **Member resigned** – The account was previously a member account. However, the account resigned from your organization by disassociating from your account.
- **Paused (suspended)** – The account is a member account but Macie is currently suspended for the account.
- **Region disabled** – The current Region is disabled for the AWS account.
- **Removed (disassociated)** – The account was previously a member account. However, you removed it as a member account by disassociating it from your account.

Last action

When you or the associated account most recently performed an action that affected the relationship between your accounts.

To sort the table by a specific field, click the column heading for the field. To change the sort order, click the column heading again. To filter the table, place your cursor in the filter bar, and then add a filter condition for a field. To further refine the results, add filter conditions for additional fields.

API

To review your organization's accounts programmatically, use the [ListMembers](#) operation of the Amazon Macie API and be sure to specify the Region that your request applies to. To review the details in additional Regions, submit your request in each additional Region.

When you submit your request, use the `onlyAssociated` parameter to specify which accounts to include in the response. By default, Macie returns details about only those accounts that are member

accounts in the specified Region, either by invitation or through AWS Organizations. To retrieve the details of all associated accounts, including accounts that aren't member accounts, include the `onlyAssociated` parameter in your request and set the parameter's value to `false`.

To review your organization's accounts by using the [AWS Command Line Interface \(AWS CLI\)](#), run the `list-members` command. For the `only-associated` parameter, specify whether to include all associated accounts or only member accounts. To include only member accounts, omit this parameter or set the parameter's value to `true`. To include all accounts, set this value to `false`. For example:

```
C:\> aws macie2 list-members --region us-east-1 --only-associated false
```

Where `us-east-1` is the Region that the request applies to, the US East (N. Virginia) Region.

If your request succeeds, Macie returns a `members` array. The array contains a `member` object for each account that meets the criteria specified in the request. In that object, the `relationshipStatus` field indicates the current status of the association between your account and the other account in the specified Region. For an account in an invitation-based organization, possible values are:

- `AccountSuspended` – The AWS account is suspended.
- `Created` – You added the account but haven't sent a membership invitation to it.
- `EmailVerificationFailed` – You tried to send a membership invitation to the account but the specified email address isn't valid for the account.
- `EmailVerificationInProgress` – You sent a membership invitation to the account and Macie is processing the request.
- `Enabled` – The account is a member account. Macie is enabled for the account and you're the Macie administrator of the account.
- `Invited` – You sent a membership invitation to the account and the account hasn't responded to your invitation.
- `Paused` – The account is a member account but Macie is currently suspended (paused) for the account.
- `RegionDisabled` – The current Region is disabled for the AWS account.
- `Removed` – The account was previously a member account. However, you removed it as a member account by disassociating it from your account.
- `Resigned` – The account was previously a member account. However, the account resigned from your organization by disassociating from your account.

For information about other fields in the `member` object, see [Members](#) in the *Amazon Macie API Reference*.

Designating a different Amazon Macie administrator account for an invitation-based organization

After you create and establish an invitation-based organization, you can change the Amazon Macie administrator account for the organization. To do this, administrators and members of the organization should take the following steps:

1. The current Macie administrator optionally exports the current inventory of active member accounts for the organization. This simplifies the transition by helping you identify member accounts that should continue to be part of the organization.

2. The current Macie administrator [removes all member accounts \(p. 268\)](#) from the current organization. This disassociates the accounts from the current administrator account but Macie continues to be enabled for the accounts.
3. The new Macie administrator [adds the previous member accounts \(p. 263\)](#) to the new organization. This associates the accounts with the new administrator account.
4. Each member account accepts the invitation to join the new organization. When an account accepts the invitation, the account becomes an active member account in the new organization. The new Macie administrator can then access Macie settings, data, and resources for the account.

If your organization uses Macie in multiple AWS Regions, perform the preceding steps in each of those Regions.

To export the current inventory of active member accounts, the current Macie administrator can use the Amazon Macie console or the Amazon Macie API. With the console, the current administrator can export the data as a comma-separated values (CSV) file. The new administrator can then use the console to upload the CSV file and add all the accounts (in bulk) to the new organization.

To export member account data by using the console

1. Log in to the AWS Management Console using the credentials for the current Macie administrator account.
2. By using the AWS Region selector in the upper-right corner of the page, select the Region in which you want to export the data.
3. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
4. In the navigation pane, under **Settings**, choose **Accounts**.
5. (Optional) To filter the **Accounts** table and show only those accounts that are currently active Macie member accounts in the organization, use the filter bar above the table to add the following filter conditions:
 - **Type = Invitation**
 - **Status = Enabled**
6. In the **Accounts** table, select the check box for each active member account to include in the exported data.
7. Choose **Export CSV**.
8. Specify a name and location for the file.

With the Amazon Macie API, the current Macie administrator can retrieve the data in JSON format. The new Macie administrator can then use that data to generate the list of account IDs and email addresses for the accounts to add and invite to the new organization. To retrieve the data in JSON format, use the [ListMembers](#) operation of the Amazon Macie API. If the operation succeeds, Macie returns a `members` array that provides details about all the accounts that are associated with the administrator's account. If an account is an active Macie member account in the current, invitation-based organization, the value for the `relationshipStatus` property of the account is `Enabled` and the `invitedAt` property specifies a date and time.

Managing your membership in an invitation-based organization in Amazon Macie

If you're invited to join an organization in Amazon Macie, you can optionally accept or decline the invitation. In Macie, an organization is a set of accounts that are centrally managed as a group of related accounts. An organization consists of one designated Macie administrator account and one or more associated member accounts.

If you accept an invitation, your account becomes a member account in the organization. When you accept, the account that sent the invitation becomes the Macie administrator account for your account—you associate your account with the other account and you enable an administrator-member relationship between the accounts. The Macie administrator account can then access certain Macie settings, data, and resources for your account in the applicable AWS Region. For more information, see [Understanding the relationship between Amazon Macie administrator and member accounts \(p. 238\)](#).

If you decline an invitation, the current status and settings for your Macie account aren't changed.

Topics

- [Responding to membership invitations for organizations \(p. 275\)](#)
- [Disassociating from an Amazon Macie administrator account \(p. 276\)](#)

Responding to membership invitations for organizations

When you receive an invitation to join an organization, Amazon Macie notifies you in several ways. By default, Macie sends the invitation to you as an email message. Macie also creates an AWS Health event for your AWS account. If you already use Macie in the AWS Region from which the invitation was sent, Macie also displays an **Accounts** badge and notification on the Macie console.

After you receive an invitation, you can optionally accept or decline the invitation. Before you respond, note the following:

- You can be a member of only one organization at a time. If you receive multiple invitations, you can accept only one. Or, if you're already a member of an organization, you have to disassociate your account from its current Macie administrator account before you can join a different organization.
- If you use Macie in multiple Regions, your account has to have the same Macie administrator account in all of those Regions. The Macie administrator has to send invitations to you separately from each Region, and you have to accept the invitations separately in each Region.
- To accept or decline an invitation, you have to enable Macie in the Region that the invitation was sent from. Declining an invitation is optional. If you enable Macie to decline an invitation, you can [disable Macie \(p. 327\)](#) in the Region after you decline the invitation. This helps ensure that you don't incur unnecessary charges for using Macie in the Region.

For additional considerations, see [Responding to and managing membership invitations \(p. 262\)](#).

To respond to a membership invitation for an organization

To respond to a membership invitation, you can use the Amazon Macie console or the Amazon Macie API.

Console

Follow these steps to respond to a membership invitation by using the Amazon Macie console.

To respond to a membership invitation

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. By using the AWS Region selector in the upper-right corner of the page, select the Region in which you received the invitation.
3. If you haven't enabled Macie in the Region, choose **Get started**, and then choose **Enable Macie**. You have to enable Macie before you can accept or decline an invitation.
4. In the navigation pane, under **Settings**, choose **Accounts**.
5. Under **Administrator account**, do one of the following:
 - To accept the invitation, turn on **Accept** (☑) next to the invitation. Then choose **Accept invitation** or **Update**, depending on whether you previously accepted another invitation.

- To decline the invitation, choose **Decline invitation** next to the invitation, and then confirm that you want to decline the invitation.

If you received and want to respond to the invitation in additional Regions, repeat the preceding steps in each additional Region.

API

To respond to an invitation programmatically, use the [AcceptInvitation](#) or [DeclineInvitations](#) operation of the Amazon Macie API, depending on whether you want to accept or decline the invitation. When you submit your request, be sure to specify the Region that the invitation was sent from. To respond to the invitation in additional Regions, submit your request in each additional Region.

In an `AcceptInvitation` request, use the `administratorAccountId` parameter to specify the 12-digit account ID for the AWS account that sent the invitation. Use the `invitationId` parameter to specify the unique ID for the invitation to accept.

In a `DeclineInvitations` request, use the `accountIds` parameter to specify the 12-digit account ID for the AWS account that sent the invitation to decline.

To retrieve the IDs, you can use the [ListInvitations](#) operation of the Amazon Macie API. If the operation succeeds, Macie returns an `invitations` array that provides details about invitations that you've received, including the account ID for the account that sent each invitation and the unique ID for each invitation. If the value for the `relationshipStatus` property of an invitation is `Invited`, you haven't responded to the invitation yet.

To respond to an invitation by using the [AWS Command Line Interface \(AWS CLI\)](#), run the `accept-invitation` or `decline-invitations` command, depending on whether you want to accept or decline the invitation. Use the `region` parameter to specify the Region that the invitation was sent from. For example:

```
C:\> aws macie2 accept-invitation --region us-east-1 --administrator-account-id 123456789012 --invitation-id d8bdad0e203fd1242e0a4721bexample
```

Where `us-east-1` is the Region that the invitation was sent from (the US East (N. Virginia) Region), `123456789012` is the account ID for the account that sent the invitation, and `d8bdad0e203fd1242e0a4721bexample` is the unique ID for the invitation to accept.

If a request to accept an invitation succeeds, Macie returns an empty response. If a request to decline an invitation succeeds, Macie returns an empty `unprocessedAccounts` array.

After you decline an invitation, the invitation persists as a resource for your Macie account. You can optionally delete it by using the [DeleteInvitations](#) operation or, for the AWS CLI, the `delete-invitations` command.

Disassociating from an Amazon Macie administrator account

If you accepted an invitation to join an organization in Amazon Macie, you can subsequently resign from the organization by disassociating your account from its current Macie administrator account. Note that you can't do this if your account is a member account in an AWS Organizations organization. To resign from an AWS Organizations organization, work with your Macie administrator to remove your account from the organization.

If you disassociate your account from its Macie administrator account, the Macie administrator loses access to all settings, data, and resources for your Macie account. This includes metadata and policy

findings for Amazon S3 data that you own. This also means that the administrator can no longer analyze your Amazon S3 data with sensitive data discovery jobs.

When you disassociate your account, Macie continues to be enabled for your account in the applicable Region. However, your account becomes a standalone Macie account in the Region. The status of your account changes to **Member resigned** in the administrator's account inventory.

To disassociate from a Macie administrator account

To disassociate your account from its current Macie administrator account, you can use the Amazon Macie console or the Amazon Macie API.

Console

Follow these steps to disassociate your account from its Macie administrator account by using the Amazon Macie console.

To disassociate from an administrator account

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. By using the AWS Region selector in the upper-right corner of the page, select the Region in which you want to disassociate your account from its administrator account.
3. In the navigation pane, under **Settings**, choose **Accounts**.
4. Under **Administrator account**, turn off **Accept** () next to the invitation, and then choose **Update**.

The account and original invitation continue to appear on the **Accounts** page. If you decide to re-join the organization, you can use this page to accept the original invitation again.

If you want to disassociate your account from its Macie administrator account in additional Regions, repeat the preceding steps in each additional Region.

API

To disassociate your account from its Macie administrator account programmatically, use the [DisassociateFromAdministratorAccount](#) operation of the Amazon Macie API. When you submit your request, be sure to specify the Region that the request applies to. To disassociate from the account in additional Regions, submit your request in each additional Region.

To disassociate your account from its Macie administrator account by using the AWS CLI, run the [disassociate-from-administrator-account](#) command. Use the `region` parameter to specify the Region in which to disassociate from the account.

If your request succeeds, Macie returns an empty response.

After you disassociate from the account, the original invitation persists as a resource for your Macie account unless you delete it. If you decide to re-join the organization, you can use this resource to accept the original invitation again. Alternatively, you can delete the invitation by using the [DeleteInvitations](#) operation or, for the AWS CLI, the [delete-invitations](#) command.

Forecasting and monitoring Amazon Macie costs

To help you forecast and monitor your costs for using Amazon Macie, Macie calculates and provides estimated usage costs for your account. With this data, you can determine whether to adjust your use of the service or your account quotas. If you're currently participating in the 30-day free trial of Macie, you can use this data to estimate the cost of monitoring your Amazon Simple Storage Service (Amazon S3) data for security and access control after your free trial ends. You can also check the status of your trial.

You can review your estimated usage costs on the Amazon Macie console and access them programmatically with the Amazon Macie API. If you're the Macie administrator for an organization, you can review and access both aggregated data for your organization and breakdowns of the data for accounts in your organization.

In addition to the estimated usage costs that Macie provides, you can review and monitor your actual costs by using AWS Billing and Cost Management. AWS Billing and Cost Management provides features that are designed to help you track and analyze your costs for AWS services, and manage budgets for your account or organization. It also provides features that can help you forecast usage costs based on historical data. To learn more, see the [AWS Billing and Cost Management User Guide](#).

Topics

- [Understanding how estimated usage costs are calculated for Amazon Macie \(p. 278\)](#)
- [Reviewing estimated usage costs for Amazon Macie \(p. 280\)](#)
- [Participating in the Amazon Macie free trial \(p. 283\)](#)

Understanding how estimated usage costs are calculated for Amazon Macie

Amazon Macie pricing is based on two dimensions, preventative control monitoring and sensitive data discovery jobs.

Preventative control monitoring

These costs derive from maintaining your S3 bucket inventory and evaluating and monitoring your buckets for security and access control. You're charged based on the total number of buckets that Macie can access for your account. The charges are prorated per day.

Sensitive data discovery jobs

These costs derive from running sensitive data discovery jobs to analyze S3 objects and report sensitive data in those objects. You're charged based on the amount of uncompressed data that Macie analyzes in objects when a job runs. There's no charge for objects that Macie can't analyze for reasons such as use of an unsupported Amazon S3 storage class, use of an unsupported file or storage format, or permissions settings. For more information, see [Discovering sensitive data \(p. 44\)](#). In addition, these costs don't vary based on the number of findings that your jobs produce.

Note that these costs are restricted by the monthly [sensitive data discovery quota \(p. 328\)](#) for your account. (The default quota is 5 TB of data.) If a job is running and the analysis of eligible objects

reaches this quota, Macie automatically pauses the job until the next calendar month starts (and the monthly quota is reset for your account) or you increase the quota for your account.

If you're the Macie administrator for an organization, these costs are restricted by the monthly sensitive data discovery quota for each account that you analyze data for. The quota for a member account defines the maximum amount of data that your jobs and the member account's jobs can analyze for the account during a calendar month.

If a job is running and the analysis of eligible objects reaches this quota for a member account, Macie stops analyzing objects that are owned by the account. When Macie finishes analyzing objects for all other accounts that haven't met the quota, Macie automatically pauses the job. If it's a one-time job, Macie automatically resumes the job when the next calendar month starts or the quota is increased for all the affected accounts, whichever occurs first. If it's a periodic job, Macie automatically resumes the job when the next run is scheduled to start or the next calendar month starts, whichever occurs first. If a scheduled run starts before the next calendar month starts or the quota is increased for an affected account, Macie doesn't analyze objects that are owned by the account.

For detailed information and examples of usage costs, see [Amazon Macie pricing](#).

When you use Macie to review your estimated usage costs, it's important to understand how the cost estimates are calculated. Consider the following:

- The estimates are reported in US Dollars and are for the current AWS Region only. If you use Macie in multiple Regions, the data isn't aggregated for all the Regions in which you use Macie.
- On the console, the estimates are inclusive for the current calendar month to date. If you query the data programmatically with the Amazon Macie API, you can choose an inclusive time range for the estimates. This can be a rolling time range of the preceding 30 days or the current calendar month to date.
- The estimates don't reflect all the discounts that might apply to your account. The exception is discounts that derive from Regional volume pricing tiers, as described in [Amazon Macie pricing](#). If your account qualifies for this type of discount, the estimates reflect that discount.

If you're the Macie administrator for an organization, the estimates don't reflect combined usage volume discounts for your organization. For information about these discounts, see [Volume discounts](#) in the *AWS Billing and Cost Management User Guide*.

- For preventative control monitoring, the estimate is based on the average daily cost for the applicable time range. The cost is prorated per day.
- For sensitive data discovery jobs, the estimate is based on the amount of uncompressed data that your jobs have analyzed thus far during the applicable time range.
- If you're the Macie administrator for an organization and you run jobs that analyze data for a member account, the estimated cost of those jobs is included in the estimate for the applicable member account. The estimated cost isn't included in the estimate for your administrator account.
- If your account is a member account in an organization and your Macie administrator runs jobs that analyze your data, the estimated cost of those jobs is included in the estimate for your account.
- The estimates don't include costs that you incur for using other AWS services with certain Macie features. For example, using customer managed AWS KMS keys to decrypt S3 objects that you want to inspect for sensitive data.

Also note that Macie provides a monthly free tier for sensitive data discovery jobs. Each month, there's no charge for you to analyze up to 1 GB of data to discover and report sensitive data in S3 objects. If you analyze more than 1 GB of data during a given month, sensitive data discovery charges begin to accrue for your account after the first 1 GB of data. If you analyze less than 1 GB of data during a given month, the remaining allocation doesn't roll over to the next month. If your account is part of an organization, the free tier applies to each individual account in your organization. In other words, there's no charge for each account in your organization to analyze up to 1 GB of data each month.

Reviewing estimated usage costs for Amazon Macie

To review your current estimated usage costs for Amazon Macie, you can use the Amazon Macie console or the Amazon Macie API. Both the console and the API provide estimated costs for Macie pricing dimensions:

- **Preventative control monitoring** – This is the estimated cost of maintaining your S3 bucket inventory and evaluating and monitoring the buckets for security and access control.
- **Data discovery jobs** – This is the estimated cost of the sensitive data discovery jobs that you ran.

For more information about these dimensions, see [Understanding how estimated usage costs are calculated \(p. 278\)](#). For detailed information and examples of usage costs, see [Amazon Macie pricing](#).

In Macie, estimated usage costs are reported in US Dollars and apply only to the current AWS Region. If you use the console to review the data, the cost estimates are for the current calendar month to date (inclusively). If you query the data programmatically with the Amazon Macie API, you can specify an inclusive time range for the estimates, either a rolling time range of the preceding 30 days or the current calendar month to date.

Topics

- [Reviewing estimated usage costs on the Amazon Macie console \(p. 280\)](#)
- [Querying estimated usage costs with the Amazon Macie API \(p. 281\)](#)

Reviewing estimated usage costs on the Amazon Macie console

Follow these steps to review your estimated usage costs by using the Amazon Macie console.

To review your estimated usage costs on the console

1. Open the Macie console at <https://console.aws.amazon.com/macie/>.
2. By using the AWS Region selector in the upper-right corner of the page, select the Region in which you want to review your estimated costs.
3. In the navigation pane, choose **Usage**.

If you have a standalone Macie account or your account is a member account in an organization, the **Usage** page displays a breakdown of the estimated usage costs for your account.

If you're the Macie administrator for an organization, the **Usage** page lists accounts in your organization:

- In the table, the **Total** field indicates the total estimated cost for each account.
- The **Estimated costs** section shows the total estimated cost for your organization and a breakdown of those costs by pricing dimension.

To review the breakdown of estimated costs for a specific account in your organization, choose the account in the table. The **Estimated costs** section then shows this breakdown. To show this data for another account, choose the account in the table. To clear your account selection, choose **X** next to the account ID.

Querying estimated usage costs with the Amazon Macie API

To query your estimated usage costs programmatically, you can use the following operations of the Amazon Macie API:

- **GetUsageTotals** – This operation returns total estimated usage costs for your account, grouped by usage metric. If you're the Macie administrator for an organization, this operation returns aggregated cost estimates for all the accounts in your organization. To learn more about this operation, see [Usage Totals](#) in the *Amazon Macie API Reference*.
- **GetUsageStatistics** – This operation returns usage statistics and related data for your account, grouped by account and then by usage metric. The data includes total estimated usage costs, current account quotas, and, if applicable, the date and time when the 30-day free trial started. If you're the Macie administrator for an organization, this operation returns a breakdown of the data for all the accounts in your organization. You can customize your query by sorting and filtering the query results. To learn more about this operation, see [Usage Statistics](#) in the *Amazon Macie API Reference*.

When you use either operation, you can optionally specify an inclusive time range for the data. This time range can be a rolling time range of the preceding 30 days (`PAST_30_DAYS`) or the current calendar month to date (`MONTH_TO_DATE`). If you don't specify a time range, Macie returns the data for the preceding 30 days.

The following examples show how to query estimated usage costs and statistics by using the [AWS Command Line Interface \(AWS CLI\)](#). You can also query the data by sending HTTPS requests directly to Macie, or by using a current version of another AWS command line tool or an AWS SDK. For information about AWS tools and SDKs, see [Tools to Build on AWS](#).

Examples

- [Example 1: Querying total estimated usage costs \(p. 281\)](#)
- [Example 2: Querying usage statistics \(p. 282\)](#)

Example 1: Querying total estimated usage costs

To query total estimated usage costs by using the AWS CLI, run the `get-usage-totals` command and optionally specify a time range for the data. For example:

```
C:\> aws macie2 get-usage-totals --time-range MONTH_TO_DATE
```

Where `MONTH_TO_DATE` specifies the current calendar month to date as the time range for the data.

If the command runs successfully, you receive output similar to the following.

```
{
  "timeRange": "MONTH_TO_DATE",
  "usageTotals": [
    {
      "currency": "USD",
      "estimatedCost": "153.45",
      "type": "SENSITIVE_DATA_DISCOVERY"
    },
    {
      "currency": "USD",
      "estimatedCost": "10.50",
      "type": "DATA_INVENTORY_EVALUATION"
    }
  ]
}
```

```
]
}
```

Where `estimatedCost` is the total estimated usage cost for the associated usage metric (type): `SENSITIVE_DATA_DISCOVERY`, for analyzing S3 objects to detect sensitive data; and, `DATA_INVENTORY_EVALUATION`, for monitoring and evaluating S3 buckets for security and access control.

Example 2: Querying usage statistics

To query usage statistics by using the AWS CLI, run the `get-usage-statistics` command. You can optionally sort, filter, and specify a time range for the query results. The following example retrieves usage statistics for a Macie administrator account for the preceding 30 days. The results are sorted in ascending order by account ID.

For Linux, macOS, or Unix, using the backslash (\) line-continuation character to improve readability:

```
$ aws macie2 get-usage-statistics \
--sort-by '{"key":"accountId","orderBy":"ASC"}' \
--time-range PAST_30_DAYS
```

For Microsoft Windows, using the caret (^) line-continuation character to improve readability:

```
C:\> aws macie2 get-usage-statistics ^
--sort-by={"key":"accountId","orderBy":"ASC"} ^
--time-range PAST_30_DAYS
```

Where `PAST_30_DAYS` specifies the preceding 30 days as the time range for the data.

If the command runs successfully, Macie returns a `records` array. The array contains an object for each account that's included in the query results. For example:

```
{
  "records": [
    {
      "accountId": "111122223333",
      "freeTrialStartDate": "2020-05-20T12:26:36.917000+00:00",
      "usage": [
        {
          "currency": "USD",
          "estimatedCost": "94.53",
          "serviceLimit": {
            "isServiceLimited": false,
            "unit": "TERABYTES",
            "value": 50
          },
          "type": "SENSITIVE_DATA_DISCOVERY"
        },
        {
          "currency": "USD",
          "estimatedCost": "6.35",
          "type": "DATA_INVENTORY_EVALUATION"
        }
      ]
    },
    {
      "accountId": "444455556666",
      "freeTrialStartDate": "2020-05-18T16:26:36.917000+00:00",
      "usage": [
        {
```



```
    "currency": "USD",
    "estimatedCost": "153.45",
    "serviceLimit": {
      "isServiceLimited": false,
      "unit": "TERABYTES",
      "value": 50
    },
    "type": "SENSITIVE_DATA_DISCOVERY"
  },
  {
    "currency": "USD",
    "estimatedCost": "10.50",
    "type": "DATA_INVENTORY_EVALUATION"
  }
]
},
"timeRange": "PAST_30_DAYS"
}
```

Where `estimatedCost` is the total estimated usage cost for the associated usage metric (`type`) for an account: `SENSITIVE_DATA_DISCOVERY`, for analyzing S3 objects to detect sensitive data; and, `DATA_INVENTORY_EVALUATION`, for monitoring and evaluating S3 buckets for security and access control.

Participating in the Amazon Macie free trial

When you enable Amazon Macie for the first time, your AWS account is automatically enrolled in the 30-day free trial of Macie. This includes individual member accounts in an AWS Organizations organization.

During the free trial, there's no charge for using Macie in a specific AWS Region to generate and maintain an inventory of your Amazon Simple Storage Service (Amazon S3) buckets and to evaluate and monitor the buckets for security and access control. The applicable Region is the Region that's active when you enable Macie for your account. Although you can use Macie in most Regions, your account is eligible for the free trial in only one Region.

Note

The free trial doesn't include discovery of sensitive data. This means that you'll incur charges if you create and run sensitive data discovery jobs that analyze more than 1 GB of data during the free trial. (Macie provides a monthly free tier for jobs. Each month, there's no charge for you to analyze up to 1 GB of data in S3 objects. After the first 1 GB of data, costs accrue.) You might also incur charges for other AWS services that you use with certain Macie features—for example, using customer managed AWS KMS keys to decrypt S3 objects that you want to inspect for sensitive data.

After the 30-day free trial ends, charges begin to accrue for maintaining your S3 bucket inventory and evaluating and monitoring your buckets for security and access control.

To check your status and estimated costs during the free trial

During the free trial, you can check the status of your trial and review estimated usage costs for your account. The cost estimates are based on your use of Macie thus far during the free trial. They can help you understand what some of your usage costs might be after the free trial ends. For details about how Macie calculates these values, see [Understanding how estimated usage costs are calculated \(p. 278\)](#).

Follow these steps to review this data on the Amazon Macie console. You can also access this data programmatically by using the [GetUsageStatistics](#) operation of the Amazon Macie API.

1. Open the Macie console at <https://console.aws.amazon.com/macie/>.

2. By using the AWS Region selector in the upper-right corner of the page, select the Region in which you enrolled in the free trial.
3. In the navigation pane, choose **Usage**.

The **Usage** page indicates the number of remaining days in your free trial. It also shows a breakdown of your estimated usage costs in US Dollars:

- **Preventative control monitoring** – This is the total projected cost of maintaining your S3 bucket inventory and evaluating and monitoring your buckets for security and access control after the free trial ends.
- **Data discovery jobs** – This is the total estimated cost of any sensitive data discovery jobs that you ran. Sensitive data discovery isn't included in the free trial.

If you're the Macie administrator for an organization, the **Usage** page provides details about the Macie accounts in your organization:

- In the table, the **Free trial** field indicates whether an account is currently participating in the free trial. (This field is empty if the free trial has ended for an account.) The **Total** field indicates the total estimated cost for each account.
- The **Estimated costs** section shows estimated costs for your organization overall.

To review the breakdown of estimated costs for a specific account in your organization, choose the account in the table. The **Estimated costs** section then shows this breakdown. To show this data for another account, choose the account in the table. To clear your account selection, choose **X** next to the account ID.

Security in Amazon Macie

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from data centers and network architectures that are built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between AWS and you. The [shared responsibility model](#) describes this as security *of* the cloud and security *in* the cloud:

- **Security of the cloud** – AWS is responsible for protecting the infrastructure that runs AWS services in the AWS Cloud. AWS also provides you with services that you can use securely. Third-party auditors regularly test and verify the effectiveness of our security as part of the [AWS Compliance Programs](#). To learn about the compliance programs that apply to Amazon Macie, see [AWS Services in Scope by Compliance Program](#).
- **Security in the cloud** – Your responsibility is determined by the AWS services that you use. You are also responsible for other factors including the sensitivity of your data, your company's requirements, and applicable laws and regulations.

This documentation helps you understand how to apply the shared responsibility model when using Macie. The following topics show you how to configure Macie to meet your security and compliance objectives. You also learn how to use other AWS services that help you to monitor and secure your Macie resources.

Topics

- [Data protection in Amazon Macie \(p. 285\)](#)
- [Identity and access management for Amazon Macie \(p. 286\)](#)
- [Logging and monitoring in Amazon Macie \(p. 309\)](#)
- [Compliance validation for Amazon Macie \(p. 309\)](#)
- [Resilience in Amazon Macie \(p. 310\)](#)
- [Infrastructure security in Amazon Macie \(p. 310\)](#)
- [Amazon Macie and interface VPC endpoints \(AWS PrivateLink\) \(p. 310\)](#)

Data protection in Amazon Macie

The AWS [shared responsibility model](#) applies to data protection in Amazon Macie. As described in this model, AWS is responsible for protecting the global infrastructure that runs all of the AWS Cloud. You are responsible for maintaining control over your content that is hosted on this infrastructure. This content includes the security configuration and management tasks for the AWS services that you use. For more information about data privacy, see the [Data Privacy FAQ](#). For information about data protection in Europe, see the [AWS Shared Responsibility Model and GDPR](#) blog post on the [AWS Security Blog](#).

For data protection purposes, we recommend that you protect AWS account credentials and set up individual user accounts with AWS Identity and Access Management (IAM). That way each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

- Use multi-factor authentication (MFA) with each account.

- Use SSL/TLS to communicate with AWS resources. We recommend TLS 1.2 or later.
- Set up API and user activity logging with AWS CloudTrail.
- Use AWS encryption solutions, along with all default security controls within AWS services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing personal data that is stored in Amazon S3.
- If you require FIPS 140-2 validated cryptographic modules when accessing AWS through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see [Federal Information Processing Standard \(FIPS\) 140-2](#).

We strongly recommend that you never put confidential or sensitive information, such as your customers' email addresses, into tags or free-form fields such as a **Name** field. This includes when you work with Macie or other AWS services using the console, API, AWS CLI, or AWS SDKs. Any data that you enter into tags or free-form fields used for names may be used for billing or diagnostic logs. If you provide a URL to an external server, we strongly recommend that you do not include credentials information in the URL to validate your request to that server.

Encryption at rest

Amazon Macie securely stores your data at rest using AWS encryption solutions. Macie encrypts data, such as findings, using an AWS managed key from AWS Key Management Service (AWS KMS).

If you disable Macie, it permanently deletes all resources that it stores or maintains for you, such as sensitive data discovery jobs, custom data identifiers, and findings.

Encryption in transit

Macie encrypts all data in transit between AWS services.

Amazon Macie analyzes data from Amazon S3 and exports sensitive data discovery results to an S3 bucket. After Macie gets the information that it needs from the S3 objects, they are discarded.

Macie accesses Amazon S3 using a VPC endpoint powered by AWS PrivateLink. Therefore, traffic between Macie and Amazon S3 stays on the Amazon network and does not go over the public internet. For more information, see [AWS PrivateLink](#).

Identity and access management for Amazon Macie

AWS Identity and Access Management (IAM) is an AWS service that helps an administrator securely control access to AWS resources. IAM administrators control who can be *authenticated* (signed in) and *authorized* (have permissions) to use Macie resources. IAM is an AWS service that you can use with no additional charge.

Topics

- [Audience \(p. 287\)](#)
- [Authenticating with identities \(p. 287\)](#)
- [Managing access using policies \(p. 289\)](#)
- [How Amazon Macie works with AWS Identity and Access Management \(p. 291\)](#)
- [Identity-based policy examples for Amazon Macie \(p. 296\)](#)

- [Service-linked roles for Amazon Macie \(p. 302\)](#)
- [AWS managed policies for Amazon Macie \(p. 305\)](#)
- [Troubleshooting Amazon Macie identity and access \(p. 308\)](#)

Audience

How you use AWS Identity and Access Management (IAM) differs, depending on the work that you do in Macie.

Service user – If you use the Macie service to do your job, then your administrator provides you with the credentials and permissions that you need. As you use more Macie features to do your work, you might need additional permissions. Understanding how access is managed can help you request the right permissions from your administrator. If you cannot access a feature in Macie, see [Troubleshooting Amazon Macie identity and access \(p. 308\)](#).

Service administrator – If you're in charge of Macie resources at your company, you probably have full access to Macie. It's your job to determine which Macie features and resources your service users should access. You must then submit requests to your IAM administrator to change the permissions of your service users. Review the information on this page to understand the basic concepts of IAM. To learn more about how your company can use IAM with Macie, see [How Amazon Macie works with AWS Identity and Access Management \(p. 291\)](#).

IAM administrator – If you're an IAM administrator, you might want to learn details about how you can write policies to manage access to Macie. To view example Macie identity-based policies that you can use in IAM, see [Identity-based policy examples for Amazon Macie \(p. 296\)](#).

Authenticating with identities

Authentication is how you sign in to AWS using your identity credentials. For more information about signing in using the AWS Management Console, see [Signing in to the AWS Management Console as an IAM user or root user](#) in the *IAM User Guide*.

You must be *authenticated* (signed in to AWS) as the AWS account root user, an IAM user, or by assuming an IAM role. You can also use your company's single sign-on authentication or even sign in using Google or Facebook. In these cases, your administrator previously set up identity federation using IAM roles. When you access AWS using credentials from another company, you are assuming a role indirectly.

To sign in directly to the [AWS Management Console](#), use your password with your root user email address or your IAM user name. You can access AWS programmatically using your root user or IAM users access keys. AWS provides SDK and command line tools to cryptographically sign your request using your credentials. If you don't use AWS tools, you must sign the request yourself. Do this using *Signature Version 4*, a protocol for authenticating inbound API requests. For more information about authenticating requests, see [Signature Version 4 signing process](#) in the *AWS General Reference*.

Regardless of the authentication method that you use, you might also be required to provide additional security information. For example, AWS recommends that you use multi-factor authentication (MFA) to increase the security of your account. To learn more, see [Using multi-factor authentication \(MFA\) in AWS](#) in the *IAM User Guide*.

AWS account root user

When you create an AWS account, you begin with one sign-in identity that has complete access to all AWS services and resources in the account. This identity is called the AWS account *root user* and is accessed by signing in with the email address and password that you used to create the account. We

strongly recommend that you do not use the root user for your everyday tasks. Safeguard your root user credentials and use them to perform the tasks that only the root user can perform. For the complete list of tasks that require you to sign in as the root user, see [Tasks that require root user credentials](#) in the *AWS General Reference*.

Federated identity

As a best practice, require human users, including users that require administrator access, to use federation with an identity provider to access AWS by using temporary credentials.

A *federated identity* is a user from your enterprise user directory, a web identity provider, the AWS Directory Service, the Identity Center directory, or any user that accesses AWS by using credentials provided through an identity source. When federated identities access AWS accounts, they assume roles, and the roles provide temporary credentials.

For centralized access management, we recommend that you use AWS IAM Identity Center (successor to AWS Single Sign-On). You can create users and groups in IAM Identity Center, or you can connect and synchronize to a set of users and groups in your own identity source for use across all your AWS accounts and applications. For information about IAM Identity Center, see [What is IAM Identity Center?](#) in the *AWS IAM Identity Center (successor to AWS Single Sign-On) User Guide*.

IAM users and groups

An *IAM user* is an identity within your AWS account that has specific permissions for a single person or application. Where possible, we recommend relying on temporary credentials instead of creating IAM users who have long-term credentials such as passwords and access keys. However, if you have specific use cases that require long-term credentials with IAM users, we recommend that you rotate access keys. For more information, see [Rotate access keys regularly for use cases that require long-term credentials](#) in the *IAM User Guide*.

An *IAM group* is an identity that specifies a collection of IAM users. You can't sign in as a group. You can use groups to specify permissions for multiple users at a time. Groups make permissions easier to manage for large sets of users. For example, you could have a group named *IAMAdmins* and give that group permissions to administer IAM resources.

Users are different from roles. A user is uniquely associated with one person or application, but a role is intended to be assumable by anyone who needs it. Users have permanent long-term credentials, but roles provide temporary credentials. To learn more, see [When to create an IAM user \(instead of a role\)](#) in the *IAM User Guide*.

IAM roles

An *IAM role* is an identity within your AWS account that has specific permissions. It is similar to an IAM user, but is not associated with a specific person. You can temporarily assume an IAM role in the AWS Management Console by [switching roles](#). You can assume a role by calling an AWS CLI or AWS API operation or by using a custom URL. For more information about methods for using roles, see [Using IAM roles](#) in the *IAM User Guide*.

IAM roles with temporary credentials are useful in the following situations:

- **Federated user access** – To assign permissions to a federated identity, you create a role and define permissions for the role. When a federated identity authenticates, the identity is associated with the role and is granted the permissions that are defined by the role. For information about roles for federation, see [Creating a role for a third-party Identity Provider](#) in the *IAM User Guide*.

If you use IAM Identity Center, you configure a permission set. To control what your identities can access after they authenticate, IAM Identity Center correlates the permission set to a role in IAM. For

information about permissions sets, see [Permission sets](#) in the *AWS IAM Identity Center (successor to AWS Single Sign-On) User Guide*.

- **Temporary IAM user permissions** – An IAM user or role can assume an IAM role to temporarily take on different permissions for a specific task.
- **Cross-account access** – You can use an IAM role to allow someone (a trusted principal) in a different account to access resources in your account. Roles are the primary way to grant cross-account access. However, with some AWS services, you can attach a policy directly to a resource (instead of using a role as a proxy). To learn the difference between roles and resource-based policies for cross-account access, see [How IAM roles differ from resource-based policies](#) in the *IAM User Guide*.
- **Cross-service access** – Some AWS services use features in other AWS services. For example, when you make a call in a service, it's common for that service to run applications in Amazon EC2 or store objects in Amazon S3. A service might do this using the calling principal's permissions, using a service role, or using a service-linked role.
 - **Principal permissions** – When you use an IAM user or role to perform actions in AWS, you are considered a principal. Policies grant permissions to a principal. When you use some services, you might perform an action that then triggers another action in a different service. In this case, you must have permissions to perform both actions. To see whether an action requires additional dependent actions in a policy, see [Actions, resources, and condition keys for Amazon Macie](#) in the *Service Authorization Reference*.
 - **Service role** – A service role is an [IAM role](#) that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For more information, see [Creating a role to delegate permissions to an AWS service](#) in the *IAM User Guide*.
 - **Service-linked role** – A service-linked role is a type of service role that is linked to an AWS service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your IAM account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.
- **Applications running on Amazon EC2** – You can use an IAM role to manage temporary credentials for applications that are running on an EC2 instance and making AWS CLI or AWS API requests. This is preferable to storing access keys within the EC2 instance. To assign an AWS role to an EC2 instance and make it available to all of its applications, you create an instance profile that is attached to the instance. An instance profile contains the role and enables programs that are running on the EC2 instance to get temporary credentials. For more information, see [Using an IAM role to grant permissions to applications running on Amazon EC2 instances](#) in the *IAM User Guide*.

To learn whether to use IAM roles or IAM users, see [When to create an IAM role \(instead of a user\)](#) in the *IAM User Guide*.

Managing access using policies

You control access in AWS by creating policies and attaching them to AWS identities or resources. A policy is an object in AWS that, when associated with an identity or resource, defines their permissions. AWS evaluates these policies when a principal (user, root user, or role session) makes a request. Permissions in the policies determine whether the request is allowed or denied. Most policies are stored in AWS as JSON documents. For more information about the structure and contents of JSON policy documents, see [Overview of JSON policies](#) in the *IAM User Guide*.

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

Every IAM entity (user or role) starts with no permissions. By default, users can do nothing, not even change their own password. To give a user permission to do something, an administrator must attach a permissions policy to a user. Or the administrator can add the user to a group that has the intended permissions. When an administrator gives permissions to a group, all users in that group are granted those permissions.

IAM policies define permissions for an action regardless of the method that you use to perform the operation. For example, suppose that you have a policy that allows the `iam:GetRole` action. A user with that policy can get role information from the AWS Management Console, the AWS CLI, or the AWS API.

Identity-based policies

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see [Creating IAM policies](#) in the *IAM User Guide*.

Identity-based policies can be further categorized as *inline policies* or *managed policies*. Inline policies are embedded directly into a single user, group, or role. Managed policies are standalone policies that you can attach to multiple users, groups, and roles in your AWS account. Managed policies include AWS managed policies and customer managed policies. To learn how to choose between a managed policy or an inline policy, see [Choosing between managed policies and inline policies](#) in the *IAM User Guide*.

Resource-based policies

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are IAM *role trust policies* and Amazon S3 *bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must [specify a principal](#) in a resource-based policy. Principals can include accounts, users, roles, federated users, or AWS services.

Resource-based policies are inline policies that are located in that service. You can't use AWS managed policies from IAM in a resource-based policy.

Access control lists (ACLs)

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

Amazon S3, AWS WAF, and Amazon VPC are examples of services that support ACLs. To learn more about ACLs, see [Access control list \(ACL\) overview](#) in the *Amazon Simple Storage Service Developer Guide*.

Other policy types

AWS supports additional, less-common policy types. These policy types can set the maximum permissions granted to you by the more common policy types.

- **Permissions boundaries** – A permissions boundary is an advanced feature in which you set the maximum permissions that an identity-based policy can grant to an IAM entity (IAM user or role). You can set a permissions boundary for an entity. The resulting permissions are the intersection of entity's identity-based policies and its permissions boundaries. Resource-based policies that specify the user or role in the `Principal` field are not limited by the permissions boundary. An explicit deny in any of these policies overrides the allow. For more information about permissions boundaries, see [Permissions boundaries for IAM entities](#) in the *IAM User Guide*.
- **Service control policies (SCPs)** – SCPs are JSON policies that specify the maximum permissions for an organization or organizational unit (OU) in AWS Organizations. AWS Organizations is a service for grouping and centrally managing multiple AWS accounts that your business owns. If you enable all features in an organization, then you can apply service control policies (SCPs) to any or all of your accounts. The SCP limits permissions for entities in member accounts, including each AWS account root user. For more information about Organizations and SCPs, see [How SCPs work](#) in the *AWS Organizations User Guide*.

- **Session policies** – Session policies are advanced policies that you pass as a parameter when you programmatically create a temporary session for a role or federated user. The resulting session's permissions are the intersection of the user or role's identity-based policies and the session policies. Permissions can also come from a resource-based policy. An explicit deny in any of these policies overrides the allow. For more information, see [Session policies](#) in the *IAM User Guide*.

Multiple policy types

When multiple types of policies apply to a request, the resulting permissions are more complicated to understand. To learn how AWS determines whether to allow a request when multiple policy types are involved, see [Policy evaluation logic](#) in the *IAM User Guide*.

How Amazon Macie works with AWS Identity and Access Management

Before you use AWS Identity and Access Management (IAM) to manage access to Amazon Macie, learn which IAM features are available to use with Macie.

IAM features you can use with Amazon Macie

IAM feature	Macie support
Identity-based policies (p. 291)	Yes
Resource-based policies (p. 292)	No
Policy actions (p. 292)	Yes
Policy resources (p. 293)	Yes
Policy condition keys (p. 294)	Yes
Access control lists (ACLs) (p. 295)	No
Attribute-based access control (ABAC) – tags in policies (p. 295)	Yes
Temporary credentials (p. 295)	Yes
Cross-service principal permissions (p. 296)	No
Service roles (p. 296)	No
Service-linked roles (p. 296)	Yes

For a high-level view of how Macie and other AWS services work with most IAM features, see [AWS services that work with IAM](#) in the *IAM User Guide*.

Identity-based policies for Amazon Macie

Supports identity-based policies	Yes
----------------------------------	-----

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform,

on which resources, and under what conditions. To learn how to create an identity-based policy, see [Creating IAM policies](#) in the *IAM User Guide*.

With IAM identity-based policies, you can specify allowed or denied actions and resources as well as the conditions under which actions are allowed or denied. You can't specify the principal in an identity-based policy because it applies to the user or role to which it is attached. To learn about all of the elements that you can use in a JSON policy, see [IAM JSON policy elements reference](#) in the *IAM User Guide*.

Macie supports identity-based policies. For examples, see [Identity-based policy examples for Amazon Macie](#) (p. 296).

Resource-based policies within Amazon Macie

Supports resource-based policies	No
----------------------------------	----

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are IAM *role trust policies* and Amazon S3 *bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must [specify a principal](#) in a resource-based policy. Principals can include accounts, users, roles, federated users, or AWS services.

To enable cross-account access, you can specify an entire account or IAM entities in another account as the principal in a resource-based policy. Adding a cross-account principal to a resource-based policy is only half of establishing the trust relationship. When the principal and the resource are in different AWS accounts, an IAM administrator in the trusted account must also grant the principal entity (user or role) permission to access the resource. They grant permission by attaching an identity-based policy to the entity. However, if a resource-based policy grants access to a principal in the same account, no additional identity-based policy is required. For more information, see [How IAM roles differ from resource-based policies](#) in the *IAM User Guide*.

Macie doesn't support resource-based policies. That is to say, you can't attach a policy directly to a Macie resource.

Policy actions for Amazon Macie

Supports policy actions	Yes
-------------------------	-----

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Action` element of a JSON policy describes the actions that you can use to allow or deny access in a policy. Policy actions usually have the same name as the associated AWS API operation. There are some exceptions, such as *permission-only actions* that don't have a matching API operation. There are also some operations that require multiple actions in a policy. These additional actions are called *dependent actions*.

Include actions in a policy to grant permissions to perform the associated operation.

Policy actions for Macie use the following prefix before the action:

```
macie2
```

For example, to grant someone permission to access information about all the managed data identifiers that Macie provides, which is an action that corresponds to the `ListManagedDataIdentifiers` operation of the Amazon Macie API, include the `macie2:ListManagedDataIdentifiers` action in their policy:

```
"Action": "macie2:ListManagedDataIdentifiers"
```

To specify multiple actions in a single statement, separate them with commas. For example:

```
"Action": [
  "macie2:ListManagedDataIdentifiers",
  "macie2:ListCustomDataIdentifiers"
]
```

You can also specify multiple actions by using wildcards (*). For example, to specify all actions that begin with the word `List`, include the following action:

```
"Action": "macie2:List*"
```

However, as a best practice, you should create policies that follow the principle of least privilege. In other words, you should create policies that include only the permissions that are required to perform a specific task.

For a list of Macie actions, see [Actions defined by Amazon Macie](#) in the *Service Authorization Reference*. For examples of policies that specify Macie actions, see [Identity-based policy examples for Amazon Macie](#) (p. 296).

Policy resources for Amazon Macie

Supports policy resources	Yes
---------------------------	-----

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Resource` JSON policy element specifies the object or objects to which the action applies. Statements must include either a `Resource` or a `NotResource` element. As a best practice, specify a resource using its [Amazon Resource Name \(ARN\)](#). You can do this for actions that support a specific resource type, known as *resource-level permissions*.

For actions that don't support resource-level permissions, such as listing operations, use a wildcard (*) to indicate that the statement applies to all resources.

```
"Resource": "*" 
```

Macie defines the following resource types:

- Allow list
- Custom data identifier
- Filter or suppression rule, also referred to as a *findings filter*
- Member account
- Sensitive data discovery job, also referred to as a *classification job*

You can specify these types of resources in policies by using ARNs.

For example, to create a policy for the sensitive data discovery job that has the job ID `3ce05dbb7ec5505def334104bexample`, you can use the following ARN:

```
"Resource": "arn:aws:macie2:*:*:classification-job/3ce05dbb7ec5505def334104bexample"
```

Or, to specify all the sensitive data discovery jobs for a certain account, use a wildcard (*):

```
"Resource": "arn:aws:macie2:*:*:123456789012:classification-job/*"
```

Where `123456789012` is the account ID for the AWS account that created the jobs. As a best practice, however, you should create policies that follow the principle of least privilege. In other words, you should create policies that include only the permissions that are required to perform a specific task on a specific resource.

Some Macie actions can apply to multiple resources. For example, the `macie2:BatchGetCustomDataIdentifiers` action can retrieve the details of multiple custom data identifiers. In these cases, a principal must have permissions to access all the resources that the action applies to. To specify multiple resources in a single statement, separate the ARNs with commas:

```
"Resource": [
  "arn:aws:macie2:*:*:custom-data-identifier/12g4aff9-8e22-4f2b-b3fd-3063eexample",
  "arn:aws:macie2:*:*:custom-data-identifier/2d12c96a-8e78-4ca6-b1dc-8fd65example",
  "arn:aws:macie2:*:*:custom-data-identifier/4383a69d-4a1e-4a07-8715-208ddexample"
]
```

For a list of Macie resource types and the ARN syntax for each one, see [Resource types defined by Amazon Macie](#) in the *Service Authorization Reference*. To learn which actions you can specify with each resource type, see [Actions defined by Amazon Macie](#) in the *Service Authorization Reference*. For examples of policies that specify resources, see [Identity-based policy examples for Amazon Macie \(p. 296\)](#).

Policy condition keys for Amazon Macie

Supports service-specific policy condition keys	Yes
---	-----

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Condition` element (or *Condition block*) lets you specify conditions in which a statement is in effect. The `Condition` element is optional. You can create conditional expressions that use [condition operators](#), such as equals or less than, to match the condition in the policy with values in the request.

If you specify multiple `Condition` elements in a statement, or multiple keys in a single `Condition` element, AWS evaluates them using a logical AND operation. If you specify multiple values for a single condition key, AWS evaluates the condition using a logical OR operation. All of the conditions must be met before the statement's permissions are granted.

You can also use placeholder variables when you specify conditions. For example, you can grant an IAM user permission to access a resource only if it is tagged with their IAM user name. For more information, see [IAM policy elements: variables and tags](#) in the *IAM User Guide*.

AWS supports global condition keys and service-specific condition keys. To see all AWS global condition keys, see [AWS global condition context keys](#) in the *IAM User Guide*.

For a list of Macie condition keys, see [Condition keys for Amazon Macie](#) in the *Service Authorization Reference*. To learn which actions and resources you can use a condition key with, see [Actions defined by Amazon Macie](#). For examples of policies that use condition keys, see [Identity-based policy examples for Amazon Macie](#) (p. 296).

Access control lists (ACLs) in Amazon Macie

Supports ACLs	No
---------------	----

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

Amazon Simple Storage Service (Amazon S3) is an example of an AWS service that supports ACLs. To learn more, see [Access control list \(ACL\) overview](#) in the *Amazon Simple Storage Service User Guide*.

Macie doesn't support ACLs. That is to say, you can't attach an ACL to a Macie resource.

Attribute-based access control (ABAC) with Amazon Macie

Supports ABAC (tags in policies)	Yes
----------------------------------	-----

Attribute-based access control (ABAC) is an authorization strategy that defines permissions based on attributes. In AWS, these attributes are called *tags*. You can attach tags to IAM entities (users or roles) and to many AWS resources. Tagging entities and resources is the first step of ABAC. Then you design ABAC policies to allow operations when the principal's tag matches the tag on the resource that they are trying to access.

ABAC is helpful in environments that are growing rapidly and helps with situations where policy management becomes cumbersome.

To control access based on tags, you provide tag information in the [condition element](#) of a policy using the `aws:ResourceTag/key-name`, `aws:RequestTag/key-name`, or `aws:TagKeys` condition keys.

If a service supports all three condition keys for every resource type, then the value is **Yes** for the service. If a service supports all three condition keys for only some resource types, then the value is **Partial**.

For more information about ABAC, see [What is ABAC?](#) in the *IAM User Guide*. To view a tutorial with steps for setting up ABAC, see [Use attribute-based access control \(ABAC\)](#) in the *IAM User Guide*.

You can attach tags to Macie resources—allow lists, custom data identifiers, filter rules and suppression rules, member accounts, and sensitive data discovery jobs. You can also control access to these types of resources by providing tag information in the `Condition` element of a policy. For information about tagging Macie resources, see [Tagging Amazon Macie resources](#) (p. 315). For an example of an identity-based policy that controls access to a resource based on tags, see [Identity-based policy examples for Amazon Macie](#) (p. 296).

Using temporary credentials with Amazon Macie

Supports temporary credentials	Yes
--------------------------------	-----

Some AWS services don't work when you sign in using temporary credentials. For additional information, including which AWS services work with temporary credentials, see [AWS services that work with IAM](#) in the *IAM User Guide*.

You are using temporary credentials if you sign in to the AWS Management Console using any method except a user name and password. For example, when you access AWS using your company's single sign-on (SSO) link, that process automatically creates temporary credentials. You also automatically create temporary credentials when you sign in to the console as a user and then switch roles. For more information about switching roles, see [Switching to a role \(console\)](#) in the *IAM User Guide*.

You can manually create temporary credentials using the AWS CLI or AWS API. You can then use those temporary credentials to access AWS. AWS recommends that you dynamically generate temporary credentials instead of using long-term access keys. For more information, see [Temporary security credentials in IAM](#).

Macie supports the use of temporary credentials.

Cross-service principal permissions for Amazon Macie

Supports principal permissions	No
--------------------------------	----

When you use an IAM user or role to perform actions in AWS, you are considered a principal. Policies grant permissions to a principal. When you use some services, you might perform an action that then triggers another action in a different service. In this case, you must have permissions to perform both actions. To see whether an action requires additional dependent actions in a policy, see [Actions, resources, and condition keys for Amazon Macie](#) in the *Service Authorization Reference*.

Macie doesn't provide any actions that require permissions for additional, dependent actions in other AWS services.

Service roles for Amazon Macie

Supports service roles	No
------------------------	----

A service role is an [IAM role](#) that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For more information, see [Creating a role to delegate permissions to an AWS service](#) in the *IAM User Guide*.

Macie doesn't assume or use service roles. To perform actions on your behalf, Macie uses a service-linked role.

Service-linked roles for Amazon Macie

Supports service-linked roles	Yes
-------------------------------	-----

A service-linked role is a type of service role that is linked to an AWS service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your IAM account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.

Macie uses a service-linked role to perform actions on your behalf. For details about this role, see [Service-linked roles for Amazon Macie \(p. 302\)](#).

Identity-based policy examples for Amazon Macie

By default, users and roles don't have permission to create or modify Macie resources. They also can't perform tasks by using the AWS Management Console, AWS Command Line Interface (AWS CLI), or AWS

API. An IAM administrator must create IAM policies that grant users and roles permission to perform actions on the resources that they need. The administrator must then attach those policies for users that require them.

To learn how to create an IAM identity-based policy by using these example JSON policy documents, see [Creating IAM policies](#) in the *IAM User Guide*.

For details about actions and resource types defined by Macie, including the format of the ARNs for each of the resource types, see [Actions, resources, and condition keys for Amazon Macie](#) in the *Service Authorization Reference*.

When you create a policy, be sure to resolve security warnings, errors, general warnings, and suggestions from AWS Identity and Access Management Access Analyzer (IAM Access Analyzer) before you save the policy. IAM Access Analyzer runs policy checks to validate a policy against IAM [policy grammar](#) and [best practices](#). These checks generate findings and provide actionable recommendations to help you author policies that are functional and conform to security best practices. To learn about validating policies by using IAM Access Analyzer, see [IAM Access Analyzer policy validation](#) in the *IAM User Guide*. To review a list of the warnings, errors, and suggestions that IAM Access Analyzer can return, see [IAM Access Analyzer policy check reference](#) in the *IAM User Guide*.

Topics

- [Policy best practices \(p. 297\)](#)
- [Using the Amazon Macie console \(p. 298\)](#)
- [Example: Allow users to review their own permissions \(p. 298\)](#)
- [Example: Allow users to create sensitive data discovery jobs \(p. 299\)](#)
- [Example: Allow users to manage a sensitive data discovery job \(p. 300\)](#)
- [Example: Allow users to review findings \(p. 301\)](#)
- [Example: Allow users to review custom data identifiers based on tags \(p. 302\)](#)

Policy best practices

Identity-based policies determine whether someone can create, access, or delete Macie resources in your account. These actions can incur costs for your AWS account. When you create or edit identity-based policies, follow these guidelines and recommendations:

- **Get started with AWS managed policies and move toward least-privilege permissions** – To get started granting permissions to your users and workloads, use the *AWS managed policies* that grant permissions for many common use cases. They are available in your AWS account. We recommend that you reduce permissions further by defining AWS customer managed policies that are specific to your use cases. For more information, see [AWS managed policies](#) or [AWS managed policies for job functions](#) in the *IAM User Guide*.
- **Apply least-privilege permissions** – When you set permissions with IAM policies, grant only the permissions required to perform a task. You do this by defining the actions that can be taken on specific resources under specific conditions, also known as *least-privilege permissions*. For more information about using IAM to apply permissions, see [Policies and permissions in IAM](#) in the *IAM User Guide*.
- **Use conditions in IAM policies to further restrict access** – You can add a condition to your policies to limit access to actions and resources. For example, you can write a policy condition to specify that all requests must be sent using SSL. You can also use conditions to grant access to service actions if they are used through a specific AWS service, such as AWS CloudFormation. For more information, see [IAM JSON policy elements: Condition](#) in the *IAM User Guide*.
- **Use IAM Access Analyzer to validate your IAM policies to ensure secure and functional permissions** – IAM Access Analyzer validates new and existing policies so that the policies adhere to the IAM policy language (JSON) and IAM best practices. IAM Access Analyzer provides more than 100 policy

checks and actionable recommendations to help you author secure and functional policies. For more information, see [IAM Access Analyzer policy validation](#) in the *IAM User Guide*.

- **Require multi-factor authentication (MFA)** – If you have a scenario that requires IAM users or root users in your account, turn on MFA for additional security. To require MFA when API operations are called, add MFA conditions to your policies. For more information, see [Configuring MFA-protected API access](#) in the *IAM User Guide*.

For more information about best practices in IAM, see [Security best practices in IAM](#) in the *IAM User Guide*.

Using the Amazon Macie console

To access the Amazon Macie console, you must have a minimum set of permissions. These permissions must allow you to list and view details about the Macie resources in your AWS account. If you create an identity-based policy that is more restrictive than the minimum required permissions, the console won't function as intended for entities (IAM users or roles) with that policy.

You don't need to allow minimum console permissions for users that are making calls only to the AWS CLI or the AWS API. Instead, allow access to only the actions that match the API operation that you're trying to perform.

To ensure that IAM users and roles can use the Amazon Macie console, provide them with console access. For information about creating a user with console access, see [Creating an IAM user in your AWS account](#) in the *IAM User Guide*.

If you create a policy to allow IAM users or roles to use the Amazon Macie console, ensure that the policy includes the appropriate `macie2:List` actions for the resources that those users or roles need to access on the console. Otherwise, they won't be able to navigate to or display details about those resources on the console.

For example, to review the details of a sensitive data discovery job by using the console, a user must be allowed to perform the `macie2:DescribeClassificationJob` action for the job *and* the `macie2:ListClassificationJobs` action. If a user isn't allowed to perform the `macie2:ListClassificationJobs` action, the user won't be able to display a list of jobs on the **Jobs** page of the console, and therefore won't be able to choose the job to display its details. For the details to include information about a custom data identifier that the job uses, for example, the user must also be allowed to perform the `macie2:BatchGetCustomDataIdentifiers` action for the custom data identifier.

Example: Allow users to review their own permissions

This example shows how you might create a policy that allows IAM users to view the inline and managed policies that are attached to their user identity. This policy includes permissions to complete this action on the console or programmatically using the AWS CLI or AWS API.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ViewOwnUserInfo",
      "Effect": "Allow",
      "Action": [
        "iam:GetUserPolicy",
        "iam:ListGroupsForUser",
        "iam:ListAttachedUserPolicies",
        "iam:ListUserPolicies",
        "iam:GetUser"
      ],
      "Resource": ["arn:aws:iam::*:user/${aws:username}"]
    }
  ]
}
```



```
    },  
    {  
      "Sid": "NavigateInConsole",  
      "Effect": "Allow",  
      "Action": [  
        "iam:GetGroupPolicy",  
        "iam:GetPolicyVersion",  
        "iam:GetPolicy",  
        "iam:ListAttachedGroupPolicies",  
        "iam:ListGroupPolicies",  
        "iam:ListPolicyVersions",  
        "iam:ListPolicies",  
        "iam:ListUsers"  
      ],  
      "Resource": "*"   
    }  
  ]  
}
```

Example: Allow users to create sensitive data discovery jobs

This example shows how you might create a policy that allows an IAM user to create sensitive data discovery jobs.

In the example, the first statement grants `macie2:CreateClassificationJob` permissions to the user. These permissions allow the user to create jobs. The statement also grants `macie2:DescribeClassificationJob` permissions. These permissions allow the user to access the details of existing jobs. Although these permissions aren't required to create jobs, access to these details can help the user create jobs that have unique configuration settings.

The second statement in the example allows the user to create, configure, and review jobs by using the Amazon Macie console. The `macie2:ListClassificationJobs` permissions allow the user to display existing jobs on the **Jobs** page of the console. All other permissions in the statement allow the user to configure and create a job by using the **Create job** pages on the console.

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Sid": "CreateAndReviewJobs",  
      "Effect": "Allow",  
      "Action": [  
        "macie2:CreateClassificationJob",  
        "macie2:DescribeClassificationJob"  
      ],  
      "Resource": "arn:aws:macie2:*:*:classification-job/*"  
    },  
    {  
      "Sid": "CreateAndReviewJobsOnConsole",  
      "Effect": "Allow",  
      "Action": [  
        "macie2:ListClassificationJobs",  
        "macie2:ListAllowLists",  
        "macie2:ListCustomDataIdentifiers",  
        "macie2:ListManagedDataIdentifiers",  
        "macie2:SearchResources",  
        "macie2:DescribeBuckets"  
      ],  
      "Resource": "*"   
    }  
  ]  
}
```

Example: Allow users to manage a sensitive data discovery job

This example shows how you might create a policy that allows an IAM user to access the details of a particular sensitive data discovery job, the job whose ID is `3ce05dbb7ec5505def334104bexample`. The example also allows the user to change the status of the job as necessary.

The first statement in the example grants `macie2:DescribeClassificationJob` and `macie2:UpdateClassificationJob` permissions to the user. These permissions allow the user to retrieve the job's details and change the job's status, respectively. The second statement grants `macie2:ListClassificationJobs` permissions to the user, which allows the user to access the job by using the **Jobs** page on the Amazon Macie console.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ManageOneJob",
      "Effect": "Allow",
      "Action": [
        "macie2:DescribeClassificationJob",
        "macie2:UpdateClassificationJob"
      ],
      "Resource": "arn:aws:macie2:*:*:classification-
job/3ce05dbb7ec5505def334104bexample"
    },
    {
      "Sid": "ListJobsOnConsole",
      "Effect": "Allow",
      "Action": "macie2:ListClassificationJobs",
      "Resource": "*"
    }
  ]
}
```

You might also allow the user to access logging data (*log events*) that Macie publishes to Amazon CloudWatch Logs for the job. To do this, you can add statements that grant permissions to perform CloudWatch Logs (`logs`) actions on the log group and stream for the job. For example:

```
"Statement": [
  {
    "Sid": "AccessLogGroupForMacieJobs",
    "Effect": "Allow",
    "Action": [
      "logs:DescribeLogGroups",
      "logs:DescribeLogStreams"
    ],
    "Resource": "arn:aws:logs:*:*:log-group:aws/macie/classificationjobs"
  },
  {
    "Sid": "AccessLogEventsForOneMacieJob",
    "Effect": "Allow",
    "Action": "logs:GetLogEvents",
    "Resource": [
      "arn:aws:logs:*:*:log-group:aws/macie/classificationjobs/*",
      "arn:aws:logs:*:*:log-group:aws/macie/classificationjobs:log-
stream:3ce05dbb7ec5505def334104bexample"
    ]
  }
]
```

For information about managing access to CloudWatch Logs, see [Overview of managing access permissions to your CloudWatch Logs resources](#) in the *Amazon CloudWatch Logs User Guide*.

Example: Allow users to review findings

This example shows how you might create a policy that allows an IAM user to access findings data by using the Amazon Macie API or the Amazon Macie console.

In this example, the `macie2:GetFindings` and `macie2:GetFindingStatistics` permissions allow the user to retrieve the data. The `macie2:ListFindings` permissions allow the user to retrieve and review the data by using the **Summary** dashboard and the **Findings** pages on the Amazon Macie console.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ReviewFindings",
      "Effect": "Allow",
      "Action": [
        "macie2:GetFindings",
        "macie2:GetFindingStatistics",
        "macie2:ListFindings"
      ],
      "Resource": "*"
    }
  ]
}
```

You might also allow the user to create and manage filter rules and suppression rules for findings. To do this, you might include a statement that grants the following permissions: `macie2:CreateFindingsFilter`, `macie2:GetFindingsFilter`, `macie2:UpdateFindingsFilter`, and `macie2>DeleteFindingsFilter`. To allow the user to manage the rules by using the Amazon Macie console, also include `macie2:ListFindingsFilters` permissions in the policy. For example:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ReviewFindings",
      "Effect": "Allow",
      "Action": [
        "macie2:GetFindings",
        "macie2:GetFindingStatistics",
        "macie2:ListFindings"
      ],
      "Resource": "*"
    },
    {
      "Sid": "ManageRules",
      "Effect": "Allow",
      "Action": [
        "macie2:GetFindingsFilter",
        "macie2:UpdateFindingsFilter",
        "macie2:CreateFindingsFilter",
        "macie2>DeleteFindingsFilter"
      ],
      "Resource": "arn:aws:macie2:*:*:findings-filter/*"
    },
    {
      "Sid": "ListRulesOnConsole",
      "Effect": "Allow",
      "Action": "macie2:ListFindingsFilters",
      "Resource": "*"
    }
  ]
}
```

```
}
]
```

Example: Allow users to review custom data identifiers based on tags

In identity-based policies, you can use conditions to control access to Amazon Macie resources based on tags. This example shows how you might create a policy that allows a user to review custom data identifiers by using the Amazon Macie console or the Amazon Macie API. However, permission is granted only if the value for the `Owner` tag is the user's name.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ReviewCustomDataIdentifiersIfOwner",
      "Effect": "Allow",
      "Action": "macie2:GetCustomDataIdentifier",
      "Resource": "arn:aws:macie2:*:*:custom-data-identifier/*",
      "Condition": {
        "StringEquals": {"aws:ResourceTag/Owner": "${aws:username}"}
      }
    },
    {
      "Sid": "ListCustomDataIdentifiersOnConsoleIfOwner",
      "Effect": "Allow",
      "Action": "macie2:ListCustomDataIdentifiers",
      "Resource": "*",
      "Condition": {
        "StringEquals": {"aws:ResourceTag/Owner": "${aws:username}"}
      }
    }
  ]
}
```

You can attach this policy to IAM users in your account. If a user named `richard-roe` attempts to review the details of a custom data identifier, the custom data identifier must be tagged `Owner=richard-roe` or `owner=richard-roe`. Otherwise, the user is denied access. The condition tag key `Owner` matches both `Owner` and `owner` because condition key names aren't case sensitive. For more information, see [IAM JSON policy elements: Condition](#) in the *IAM User Guide*.

Service-linked roles for Amazon Macie

Amazon Macie uses an AWS Identity and Access Management (IAM) [service-linked role](#) named `AWSServiceRoleForAmazonMacie`. This service-linked role is an IAM role that's linked directly to Macie. It's predefined by Macie and it includes all the permissions that Macie requires to call other AWS services on your behalf. Macie uses this service-linked role in all the AWS Regions where Macie is available.

A service-linked role makes setting up Macie easier because you don't have to manually add the necessary permissions. Macie defines the permissions of this service-linked role, and unless defined otherwise, only Macie can assume the role. The defined permissions include the trust policy and the permissions policy, and that permissions policy can't be attached to any other IAM entity.

You must configure permissions to allow an IAM entity (such as a user, group, or role) to create, edit, or delete a service-linked role. For more information, see [Service-linked role permissions](#) in the *IAM User Guide*. You can delete a service-linked role only after you delete its related resources. This protects your resources because you can't inadvertently remove permission to access the resources.

For information about other services that support service-linked roles, see [AWS services that work with IAM](#) and look for the services that have **Yes** in the **Service-linked roles** column. Choose a **Yes** with a link to review the service-linked role documentation for that service.

Topics

- [Service-linked role permissions for Amazon Macie \(p. 303\)](#)
- [Creating the service-linked role for Amazon Macie \(p. 304\)](#)
- [Editing the service-linked role for Amazon Macie \(p. 304\)](#)
- [Deleting the service-linked role for Amazon Macie \(p. 304\)](#)
- [Supported AWS Regions for the Amazon Macie service-linked role \(p. 305\)](#)

Service-linked role permissions for Amazon Macie

Amazon Macie uses the service-linked role named `AWSServiceRoleForAmazonMacie`. This service-linked role trusts the `macie.amazonaws.com` service to assume the role.

The permissions policy for the role, which is named `AmazonMacieServiceRolePolicy`, allows Macie to perform tasks such as the following on the specified resources:

- Use Amazon S3 actions to retrieve information about S3 buckets and objects.
- Use Amazon S3 actions to retrieve S3 objects.
- Use AWS Organizations actions to retrieve information about associated accounts.
- Use Amazon CloudWatch Logs actions to log events for sensitive data discovery jobs.

The role is configured with the following permissions policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "iam:ListAccountAliases",
        "organizations:DescribeAccount",
        "organizations:ListAccounts",
        "s3:GetAccountPublicAccessBlock",
        "s3:ListAllMyBuckets",
        "s3:GetBucketAcl",
        "s3:GetBucketLocation",
        "s3:GetBucketLogging",
        "s3:GetBucketPolicy",
        "s3:GetBucketPolicyStatus",
        "s3:GetBucketPublicAccessBlock",
        "s3:GetBucketTagging",
        "s3:GetBucketVersioning",
        "s3:GetBucketWebsite",
        "s3:GetEncryptionConfiguration",
        "s3:GetLifecycleConfiguration",
        "s3:GetReplicationConfiguration",
        "s3:ListBucket",
        "s3:GetObject",
        "s3:GetObjectAcl",
        "s3:GetObjectTagging"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
```

```
    "Action": [
      "logs:CreateLogGroup"
    ],
    "Resource": [
      "arn:aws:logs:*:*:log-group:/aws/macie/*"
    ]
  },
  {
    "Effect": "Allow",
    "Action": [
      "logs:CreateLogStream",
      "logs:PutLogEvents",
      "logs:DescribeLogStreams"
    ],
    "Resource": [
      "arn:aws:logs:*:*:log-group:/aws/macie/*:log-stream:*"
    ]
  }
]
```

For details about updates to the `AmazonMacieServiceRolePolicy` policy, see [Amazon Macie updates to AWS managed policies \(p. 306\)](#).

You must configure permissions to allow an IAM entity (such as a user, group, or role) to create, edit, or delete a service-linked role. For more information, see [Service-linked role permissions](#) in the *IAM User Guide*.

Creating the service-linked role for Amazon Macie

You don't need to manually create the `AWSServiceRoleForAmazonMacie` service-linked role for Amazon Macie. When you enable Macie for your AWS account, Macie automatically creates the service-linked role for you.

If you delete the Macie service-linked role and then need to create it again, you can use the same process to re-create the role in your account. When you enable Macie again, Macie creates the service-linked role again for you.

Editing the service-linked role for Amazon Macie

Amazon Macie doesn't allow you to edit the `AWSServiceRoleForAmazonMacie` service-linked role. After a service-linked role is created, you can't change the name of the role because various entities might reference the role. However, you can edit the description of the role using IAM. For more information, see [Editing a service-linked role](#) in the *IAM User Guide*.

Deleting the service-linked role for Amazon Macie

If you no longer need to use Amazon Macie, we recommend that you manually delete the `AWSServiceRoleForAmazonMacie` service-linked role. When you disable Macie, Macie doesn't delete the role for you.

Before you delete the role, you must disable Macie in each AWS Region where you enabled it. You must also manually clean up the resources for the role. To delete the role, you can use the IAM console, the AWS CLI, or the AWS API. For more information, see [Deleting a service-linked role](#) in the *IAM User Guide*.

Note

If Macie is using the `AWSServiceRoleForAmazonMacie` role when you try to delete the resources, the deletion might fail. If that happens, wait a few minutes and then try the operation again.

If you delete the `AWSServiceRoleForAmazonMacie` service-linked role and need to create it again, you can create it again by enabling Macie for your account. When you enable Macie again, Macie creates the service-linked role again for you.

Supported AWS Regions for the Amazon Macie service-linked role

Amazon Macie supports using the `AWSServiceRoleForAmazonMacie` service-linked role in all the AWS Regions where Macie is available. For a list of Regions where Macie is currently available, see [Amazon Macie endpoints and quotas](#) in the *Amazon Web Services General Reference*.

AWS managed policies for Amazon Macie

To add permissions to users, groups, and roles, it is easier to use AWS managed policies than to write policies yourself. It takes time and expertise to [create IAM customer managed policies](#) that provide your team with only the permissions they need. To get started quickly, you can use our AWS managed policies. These policies cover common use cases and are available in your AWS account. For more information about AWS managed policies, see [AWS managed policies](#) in the *IAM User Guide*.

AWS services maintain and update AWS managed policies. You can't change the permissions in AWS managed policies. Services occasionally add additional permissions to an AWS managed policy to support new features. This type of update affects all identities (users, groups, and roles) where the policy is attached. Services are most likely to update an AWS managed policy when a new feature is launched or when new operations become available. Services do not remove permissions from an AWS managed policy, so policy updates won't break your existing permissions.

Additionally, AWS supports managed policies for job functions that span multiple services. For example, the `ViewOnlyAccess` AWS managed policy provides read-only access to many AWS services and resources. When a service launches a new feature, AWS adds read-only permissions for new operations and resources. For a list and descriptions of job function policies, see [AWS managed policies for job functions](#) in the *IAM User Guide*.

Amazon Macie provides two AWS managed policies, the `AmazonMacieFullAccess` policy and the `AmazonMacieServiceRolePolicy` policy.

Topics

- [AWS managed policy: AmazonMacieFullAccess \(p. 305\)](#)
- [AWS managed policy: AmazonMacieServiceRolePolicy \(p. 306\)](#)
- [Amazon Macie updates to AWS managed policies \(p. 306\)](#)

AWS managed policy: AmazonMacieFullAccess

You can attach the `AmazonMacieFullAccess` policy to your IAM identities.

This policy grants administrative permissions that allow an IAM identity (principal) to create the [Macie service-linked role \(p. 302\)](#) and perform all read and write actions for Amazon Macie. The policy must be attached to a principal before the principal can enable Macie for their account—a principal must be allowed to create the Macie service-linked role in order to enable Macie for their account.

Permissions details

This policy includes the following permissions:

- `macie2` – Allows principals to perform all read and write actions for Amazon Macie.
- `iam` – Allows principals to create service-linked roles. The `Resource` element specifies the service-linked role for Macie. The `Condition` element uses the `iam:AWSServiceName` [IAM condition key](#) and the `StringLike` [condition operator](#) to restrict permissions to the service-linked role for Macie.
- `pricing` – Allows principals to retrieve pricing data for their AWS account from AWS Billing and Cost Management. Macie uses this data to calculate and display estimated costs when principals create and configure sensitive data discovery jobs.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "macie2:*"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": "iam:CreateServiceLinkedRole",
      "Resource": "arn:aws:iam::*:role/aws-service-role/macie.amazonaws.com/
AWSServiceRoleForAmazonMacie",
      "Condition": {
        "StringLike": {
          "iam:AWSServiceName": "macie.amazonaws.com"
        }
      }
    },
    {
      "Effect": "Allow",
      "Action": "pricing:GetProducts",
      "Resource": "*"
    }
  ]
}
```

AWS managed policy: AmazonMacieServiceRolePolicy

You can't attach the `AmazonMacieServiceRolePolicy` policy to your IAM entities. This policy is attached to a service-linked role that allows Macie to perform actions on your behalf. For more information, see [Service-linked roles for Amazon Macie \(p. 302\)](#).

Amazon Macie updates to AWS managed policies

View details about updates to AWS managed policies for Amazon Macie since this service began tracking these changes. For automatic alerts about changes to this page, subscribe to the RSS feed on the [Macie document history \(p. 331\)](#) page.

Change	Description	Date
AmazonMacieFullAccess (p. 305) – Updates to an existing policy	In the <code>AmazonMacieFullAccess</code> policy, Macie updated the Amazon Resource Name (ARN) of the Macie service-linked role (<code>aws-service-role/macie.amazonaws.com/AWSServiceRoleForAmazonMacie</code>).	June 30, 2022
AmazonMacieServiceRolePolicy (p. 305) – Updates to an existing policy	Macie removed actions and resources for Amazon Macie Classic from the <code>AmazonMacieServiceRolePolicy</code> policy. Amazon Macie Classic has been discontinued and is no longer available. More specifically, Macie removed all AWS CloudTrail actions. Macie also removed all Amazon S3 actions for the following resources: <code>arn:aws:s3:::awsmacie-*</code> , <code>arn:aws:s3:::awsmacietrail-*</code> , and <code>arn:aws:s3:::*-awsmacietrail-*</code> .	May 20, 2022
AmazonMacieFullAccess (p. 305) – Updates to an existing policy	Macie added an AWS Billing and Cost Management (<code>pricing</code>) action to the <code>AmazonMacieFullAccess</code> policy. This action allows principals to retrieve pricing data for their account. Macie uses this data to calculate and display estimated costs when principals create and configure sensitive data discovery jobs. Macie also removed Amazon Macie Classic (<code>macie</code>) actions from the <code>AmazonMacieFullAccess</code> policy.	March 7, 2022
AmazonMacieServiceRolePolicy (p. 305) – Updates to an existing policy	Macie added Amazon CloudWatch Logs actions to the <code>AmazonMacieServiceRolePolicy</code> policy. These actions allow Macie to publish log events to CloudWatch Logs for sensitive data discovery jobs.	April 13, 2021
Macie started tracking changes	Macie started tracking changes for its AWS managed policies.	April 13, 2021

Troubleshooting Amazon Macie identity and access

The following information can help you diagnose and fix common issues that you might encounter when working with Amazon Macie and AWS Identity and Access Management (IAM).

Topics

- [I'm not authorized to perform an action in Amazon Macie \(p. 308\)](#)
- [I want to view my access keys \(p. 308\)](#)
- [I'm an AWS administrator and want to allow others to access Amazon Macie \(p. 308\)](#)
- [I want to allow people outside my AWS account to access my Amazon Macie resources \(p. 309\)](#)

I'm not authorized to perform an action in Amazon Macie

If the AWS Management Console tells you that you're not authorized to perform an action, then you must contact your administrator for assistance. Your administrator is the person that provided you with your user name and password.

The following example error occurs when the `mateojackson` IAM user tries to use the Amazon Macie console to review the details of a fictional `my-example-widget` resource but doesn't have the fictional `macie2:GetWidget` permissions.

```
User: arn:aws:iam::123456789012:user/mateojackson is not authorized to perform:
macie2:GetWidget on resource: my-example-widget
```

In this case, Mateo asks his administrator to update his policies to allow him to access the `my-example-widget` resource using the `macie2:GetWidget` action.

I want to view my access keys

After you create your IAM user access keys, you can view your access key ID at any time. However, you can't view your secret access key again. If you lose your secret key, you must create a new access key pair.

Access keys consist of two parts: an access key ID (for example, `AKIAIOSFODNN7EXAMPLE`) and a secret access key (for example, `wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY`). Like a user name and password, you must use both the access key ID and secret access key together to authenticate your requests. Manage your access keys as securely as you do your user name and password.

Important

Do not provide your access keys to a third party, even to help [find your canonical user ID](#). By doing this, you might give someone permanent access to your account.

When you create an access key pair, you are prompted to save the access key ID and secret access key in a secure location. The secret access key is available only at the time you create it. If you lose your secret access key, you must add new access keys to your IAM user. You can have a maximum of two access keys. If you already have two, you must delete one key pair before creating a new one. To view instructions, see [Managing access keys](#) in the *IAM User Guide*.

I'm an AWS administrator and want to allow others to access Amazon Macie

To allow others to access Macie, you must create an IAM entity (user or role) for the person or application that needs access. They will use the credentials for that entity to access AWS. You must then attach a policy to the entity that grants them the correct permissions in Macie.

To get started right away, see [Creating your first IAM delegated user and group](#) in the *IAM User Guide*.

I want to allow people outside my AWS account to access my Amazon Macie resources

You can create a role that users in other accounts or people outside of your organization can use to access your resources. You can specify who is trusted to assume the role. For services that support resource-based policies or access control lists (ACLs), you can use those policies to grant people access to your resources.

To learn more, consult the following:

- To learn whether Macie supports these features, see [How Amazon Macie works with AWS Identity and Access Management \(p. 291\)](#).
- To learn how to provide access to your resources across AWS accounts that you own, see [Providing access to an IAM user in another AWS account that you own](#) in the *IAM User Guide*.
- To learn how to provide access to your resources to third-party AWS accounts, see [Providing access to AWS accounts owned by third parties](#) in the *IAM User Guide*.
- To learn how to provide access through identity federation, see [Providing access to externally authenticated users \(identity federation\)](#) in the *IAM User Guide*.
- To learn the difference between using roles and resource-based policies for cross-account access, see [How IAM roles differ from resource-based policies](#) in the *IAM User Guide*.

Logging and monitoring in Amazon Macie

Amazon Macie integrates with AWS CloudTrail, which is a service that provides a record of actions that were taken in Macie by a user, a role, or another AWS service. This includes actions from the Amazon Macie console and programmatic calls to Amazon Macie API operations. By using the information collected by CloudTrail, you can determine which requests were made to Macie. For each request, you can identify when it was made, the IP address from which it was made, who made it, and additional details. For more information, see [Logging Amazon Macie API calls using AWS CloudTrail \(p. 312\)](#).

Compliance validation for Amazon Macie

Third-party auditors assess the security and compliance of Amazon Macie as part of multiple AWS compliance programs. These include SOC, PCI, FedRAMP, HIPAA, and others.

For a list of AWS services in scope of specific compliance programs, see [AWS Services in Scope by Compliance Program](#). For general information, see [AWS Compliance Programs](#).

You can download third-party audit reports using AWS Artifact. For more information, see [Downloading Reports in AWS Artifact](#).

Your compliance responsibility when using Macie is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. AWS provides the following resources to help with compliance:

- [Security and Compliance Quick Start Guides](#) – These deployment guides discuss architectural considerations and provide steps for deploying security- and compliance-focused baseline environments on AWS.
- [Architecting for HIPAA Security and Compliance Whitepaper](#) – This whitepaper describes how companies can use AWS to create HIPAA-compliant applications.

- [AWS Compliance Resources](#) – This collection of workbooks and guides might apply to your industry and location.
- [Evaluating Resources with Rules](#) in the *AWS Config Developer Guide* – AWS Config assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.
- [AWS Security Hub](#) – This AWS service provides a comprehensive view of your security state within AWS that helps you check your compliance with security industry standards and best practices.

Resilience in Amazon Macie

The AWS global infrastructure is built around AWS Regions and Availability Zones. Regions provide multiple physically separated and isolated Availability Zones, which are connected through low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

For more information about AWS Regions and Availability Zones, see [AWS Global Infrastructure](#).

Infrastructure security in Amazon Macie

As a managed service, Macie is protected by the AWS global network security procedures that are described in the [Amazon Web Services: Overview of Security Processes](#) whitepaper.

You use AWS published API calls to access Macie through the network. Clients must support Transport Layer Security (TLS) 1.0 or later. We recommend TLS 1.2 or later. Clients must also support cipher suites with perfect forward secrecy (PFS) such as Ephemeral Diffie-Hellman (DHE) or Elliptic Curve Ephemeral Diffie-Hellman (ECDHE). Most modern systems such as Java 7 and later support these modes.

Additionally, requests must be signed using an AWS access key ID and secret access key that is associated with an IAM principal. Or you can use the [AWS Security Token Service](#) (AWS STS) to generate temporary security credentials to sign requests.

Amazon Macie and interface VPC endpoints (AWS PrivateLink)

If you use Amazon Virtual Private Cloud (Amazon VPC) to host your AWS resources, you can establish a private connection between your VPC and Amazon Macie. Amazon VPC is an AWS service that you can use to launch AWS resources in a virtual network that you define. With a VPC, you have control over your network settings, such as the IP address range, subnets, route tables, and network gateways.

To connect your VPC to Macie, you create an *interface VPC endpoint* for Macie. Interface endpoints are powered by [AWS PrivateLink](#), a technology that enables you to privately access Amazon Macie APIs without an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection. Instances in your VPC don't need public IP addresses to communicate with Amazon Macie APIs. Traffic between your VPC and Macie doesn't leave the Amazon network.

Each interface endpoint is represented by one or more [elastic network interfaces](#) in your subnets. For more information, see [Access an AWS service using an interface VPC endpoint](#) in the *Amazon VPC User Guide*.

Topics

- [Considerations for Amazon Macie VPC endpoints \(p. 311\)](#)
- [Creating an interface VPC endpoint for Amazon Macie \(p. 311\)](#)

Considerations for Amazon Macie VPC endpoints

Amazon Macie supports VPC endpoints in all the AWS Regions where it's currently available except the Asia Pacific (Osaka) Region. For a list of Regions where Macie is currently available, see [Amazon Macie endpoints and quotas](#) in the *Amazon Web Services General Reference*. In addition, Macie supports making calls to all of its API actions from a VPC.

If you create an interface VPC endpoint for Macie, consider doing the same for other AWS services that provide VPC support and integrate with Macie, such as Amazon EventBridge and AWS Security Hub. Macie and those services can then use VPC endpoints for the integration. For example, if you create a VPC endpoint for Macie and a VPC endpoint for Security Hub, Macie can use its VPC endpoint when it publishes findings to Security Hub and Security Hub can use its VPC endpoint when it receives the findings. For information about services that support VPC endpoints, see [AWS services that integrate with AWS PrivateLink](#) in the *Amazon VPC User Guide*.

For additional considerations, see [Access an AWS service using an interface VPC endpoint](#) in the *Amazon VPC User Guide*.

Note that VPC endpoint policies are not supported for Macie. By default, full access to Macie is allowed through the endpoint. For more information, see [Identity and access management for VPC endpoints and VPC endpoint services](#) in the *Amazon VPC User Guide*.

Creating an interface VPC endpoint for Amazon Macie

You can create an interface VPC endpoint for the Amazon Macie service by using either the Amazon VPC console or the AWS Command Line Interface (AWS CLI). For more information, see [Create a VPC endpoint](#) in the *Amazon VPC User Guide*.

When you create a VPC endpoint for Macie, use the following service name:

- `com.amazonaws.region.macie2`

Where *region* is the Region code for the applicable AWS Region.

If you enable private DNS for the endpoint, you can make API requests to Macie using its default DNS name for the Region, for example, `macie2.us-east-1.amazonaws.com` for the US East (N. Virginia) Region.

For more information, see [Access an AWS service using an interface VPC endpoint](#) in the *Amazon VPC User Guide*.

Logging Amazon Macie API calls using AWS CloudTrail

Amazon Macie integrates with AWS CloudTrail, which is a service that provides a record of actions that were taken in Macie by a user, a role, or another AWS service. CloudTrail captures all API calls for Macie as events. The calls captured include calls from the Amazon Macie console and code calls to Amazon Macie API operations.

If you create a trail, you can enable continuous delivery of CloudTrail events to an Amazon Simple Storage Service (Amazon S3) bucket, including events for Macie. If you don't configure a trail, you can still review the most recent events by using **Event history** on the CloudTrail console. Using the information collected by CloudTrail, you can determine the request that was made to Macie, the IP address from which the request was made, who made the request, when it was made, and additional details.

To learn more about CloudTrail, see the [AWS CloudTrail User Guide](#).

Topics

- [Amazon Macie information in AWS CloudTrail \(p. 312\)](#)
- [Understanding Amazon Macie log file entries \(p. 313\)](#)

Amazon Macie information in AWS CloudTrail

AWS CloudTrail is enabled for your AWS account when you create the account. When activity occurs in Amazon Macie, that activity is recorded in a CloudTrail event along with other AWS events in **Event history**. You can review, search, and download recent events in your AWS account. For more information, see [Viewing events with CloudTrail Event History](#) in the *AWS CloudTrail User Guide*.

For an ongoing record of events in your AWS account, including events for Macie, create a trail. A *trail* enables CloudTrail to deliver log files to an S3 bucket. By default, when you create a trail by using the AWS CloudTrail console, the trail applies to all AWS Regions. The trail logs events from all Regions in the AWS partition and delivers the log files to the S3 bucket that you specify. Additionally, you can configure other AWS services to further analyze and act upon the event data collected in CloudTrail logs. For more information, see the following topics in the *AWS CloudTrail User Guide*:

- [Creating a trail for your AWS account](#)
- [CloudTrail supported services and integrations](#)
- [Configuring Amazon SNS notifications for CloudTrail](#)
- [Receiving CloudTrail log files from multiple Regions](#)
- [Receiving CloudTrail log files from multiple accounts](#)

All Macie actions are logged by CloudTrail and are documented in the [Amazon Macie API Reference](#). For example, calls to the `ListFindings` and `CreateFindingsFilter` actions generate entries in CloudTrail log files.

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root or AWS Identity and Access Management (IAM) user credentials.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another AWS service.

For more information, see [CloudTrail userIdentity element](#).

Understanding Amazon Macie log file entries

A trail is a configuration that enables delivery of events as log files to an S3 bucket that you specify. AWS CloudTrail log files contain one or more log entries. An event represents a single request from any source and includes information about the requested action, the date and time of the action, request parameters, and so on. CloudTrail log files aren't an ordered stack trace of the public API calls, so they don't appear in any specific order.

The following example shows a CloudTrail log entry that demonstrates the Amazon Macie `ListFindings` action.

```
{
  "eventVersion": "1.05",
  "userIdentity": {
    "type": "IAMUser",
    "principalId": "123456789012",
    "arn": "arn:aws:iam::123456789012:user/Mary",
    "accountId": "123456789012",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "userName": "Mary",
  },
  "eventTime": "2020-05-22T16:09:56Z",
  "eventSource": "macie2.amazonaws.com",
  "eventName": "ListFindings",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "198.51.100.1",
  "userAgent": "console.amazonaws.com",
  "requestParameters": {
    "sortCriteria": {
      "attributeName": "updatedAt",
      "orderBy": "DESC"
    },
    "findingCriteria": {
      "criterion": {
        "archived": {
          "eq": [
            "false"
          ]
        },
        "category": {
          "eq": [
            "POLICY"
          ]
        }
      }
    }
  },
  "maxResults": 10
},
"responseElements": null,
"requestID": "d58af6be-1115-4a41-91f8-ace03example",
"eventID": "ad97fac5-f7cf-4ff9-9cf2-d0676example",
"eventType": "AwsApiCall",
```

```
} "recipientAccountId": "123456789012"
```


Tagging Amazon Macie resources

A *tag* is a label that you can define and assign to AWS resources, including certain types of Amazon Macie resources. Tags can help you identify, categorize, and manage resources in different ways, such as by purpose, owner, environment, or other criteria. For example, you can use tags to apply policies, allocate costs, distinguish between versions of resources, or identify resources that support certain compliance requirements.

You can assign tags to the following types of Macie resources: allow lists, custom data identifiers; filter rules and suppression rules for findings; and, sensitive data discovery jobs. If you're the Macie administrator for an organization, you can also assign tags to member accounts in your organization. A resource can have as many as 50 tags.

Topics

- [Tagging fundamentals \(p. 315\)](#)
- [Using tags in IAM policies \(p. 316\)](#)
- [Adding tags to Amazon Macie resources \(p. 316\)](#)
- [Reviewing tags for Amazon Macie resources \(p. 318\)](#)
- [Editing tags for Amazon Macie resources \(p. 320\)](#)
- [Removing tags from Amazon Macie resources \(p. 322\)](#)

Tagging fundamentals

Each tag consists of a required *tag key* and an optional *tag value*, both of which you define. A *tag key* is a general label that acts as a category for more specific tag values. A *tag value* acts as a descriptor for a tag key.

For example, if you create custom data identifiers and sensitive data discovery jobs to analyze data at different points in a workflow (one set for staged data and another for production data), you might assign a `Stack` tag key to those resources. The value of the `Stack` tag key might be `Staging` for custom data identifiers and jobs that are designed to analyze staged data, and `Production` for the others.

As you define and assign tags to resources, keep the following in mind:

- Each resource can have a maximum of 50 tags.
- For each resource, each tag key must be unique and it can have only one tag value.
- Tag keys and values are case sensitive. As a best practice, we recommend that you define a strategy for capitalizing tags and implement that strategy consistently across your resources.
- A tag key can have a maximum of 128 characters. A tag value can have a maximum of 256 characters. The characters can be letters, numbers, spaces (representable in UTF-8), or the following symbols:
`_ . : / = + - @`
- The `aws:` prefix is reserved for use by AWS. You can't use it in any tag keys or values that you define. In addition, you can't change or remove tag keys or values that use this prefix. Tags that use this prefix don't count against the quota of 50 tags per resource.
- Any tags that you assign are available only for your AWS account and only in the AWS Region in which you assign them.

For additional restrictions, tips, and best practices, see [Tagging AWS resources](#) in the *Amazon Web Services General Reference*.

Important

Do not store confidential or other types of sensitive data in tags. Tags are accessible from many AWS services, including AWS Billing and Cost Management. They aren't intended to be used for sensitive data.

To add and manage tags for Macie resources, you can use the Amazon Macie console, the Amazon Macie API, the Tag Editor on the AWS Resource Groups console, or the AWS Resource Groups Tagging API. With Macie, you can add tags to a resource when you create the resource. You can also add and manage tags for individual existing resources. With Resource Groups, you can add and manage tags in bulk for multiple existing resources spanning multiple AWS services, including Macie. For more information, see the [AWS Resource Groups and Tags User Guide](#).

Using tags in IAM policies

After you start tagging resources, you can define tag-based, resource-level permissions in AWS Identity and Access Management (IAM) policies. By using tags in this way, you can implement granular control of which users and roles in your AWS account have permission to create and tag resources, and which users and roles have permission to add, edit, and remove tags more generally. To control access based on tags, you can use [tag-related condition keys](#) in the [Condition element](#) of IAM policies.

For example, you can create a policy that allows a user to have full access to all Amazon Macie resources, if the Owner tag for the resource specifies their user name:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ModifyResourceIfOwner",
      "Effect": "Allow",
      "Action": "macie2:*",
      "Resource": "*",
      "Condition": {
        "StringEqualsIgnoreCase": {"aws:ResourceTag/Owner": "${aws:username}"}
      }
    }
  ]
}
```

If you define tag-based, resource-level permissions, the permissions take effect immediately. This means that your resources are more secure as soon as they're created, and you can quickly start enforcing the use of tags for new resources. You can also use resource-level permissions to control which tag keys and values can be associated with new and existing resources. For more information, see [Controlling access to AWS resources using tags](#) in the *IAM User Guide*.

Adding tags to Amazon Macie resources

To add tags to an individual Amazon Macie resource, you can use the Amazon Macie console or the Amazon Macie API. To add tags to multiple Macie resources at the same time, use the [Tag Editor](#) on the AWS Resource Groups console or the tagging operations of the [AWS Resource Groups Tagging API](#).

Important


Adding tags to a resource can affect access to the resource. Before you add a tag to a resource, review any AWS Identity and Access Management (IAM) policies that might use tags to control access to resources.

Console

When you create an allow list, custom data identifier, or sensitive data discovery job, the Amazon Macie console provides options for adding tags to the resource. To add tags to a filter or suppression rule or a member account in an organization, you have to create the resource before you can add tags to it.

Follow these steps to add one or more tags to an existing resource by using the Amazon Macie console.

To add a tag to a resource

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. Depending on the type of resource that you want to add a tag to, do one of the following:
 - For an allow list, choose **Allow lists** in the navigation pane.
Then, in the table, select the check box for the list. Then choose **Manage tags** on the **Actions** menu.
 - For a custom data identifier, choose **Custom data identifiers** in the navigation pane.
Then, in the table, select the check box for the custom data identifier. Then choose **Manage tags** on the **Actions** menu.
 - For a filter or suppression rule, choose **Findings** in the navigation pane.
Then, in the **Saved rules** list, choose the edit icon () next to the rule. Then choose **Manage tags**.
 - For a member account in your organization, choose **Accounts** in the navigation pane.
Then, in the table, select the check box for the account. Then choose **Manage tags** on the **Actions** menu.
 - For a sensitive data discovery job, choose **Jobs** in the navigation pane.
Then, in the table, select the check box for the job. Then choose **Manage tags** on the **Actions** menu.
3. In the **Manage tags** window, choose **Edit tags**.
4. Choose **Add tag**.
5. In the **Key** box, enter the tag key for the tag to add to the resource. Then, in the **Value** box, optionally enter a tag value for the key.

A tag key can contain as many as 128 characters. A tag value can contain as many as 256 characters. The characters can be letters, numbers, spaces, or the following symbols: `_ . : / = + - @`
6. (Optional) To add another tag to the resource, choose **Add tag**, and then repeat the preceding step. You can assign as many as 50 tags to a resource.
7. When you finish adding tags, choose **Save**.

API

To create a resource and add one or more tags to it programmatically, use the appropriate `Create` operation for the type of resource that you want to create:

- **Allow list** – Use the [CreateAllowList](#) operation or, if you're using the AWS Command Line Interface (AWS CLI), run the `create-allow-list` command.
- **Custom data identifier** – Use the [CreateCustomDataIdentifier](#) operation or, if you're using the AWS CLI, run the `create-custom-data-identifier` command.

- **Filter or suppression rule** – Use the [CreateFindingsFilter](#) operation or, if you're using the AWS CLI, run the [create-findings-filter](#) command.
- **Member account** – Use the [CreateMember](#) operation or, if you're using the AWS CLI, run the [create-member](#) command.
- **Sensitive data discovery job** – Use the [CreateClassificationJob](#) operation or, if you're using the AWS CLI, run the [create-classification-job](#) command.

To add one or more tags to an existing resource, use the [TagResource](#) operation of the Amazon Macie API or, if you're using the AWS CLI, run the [tag-resource](#) command.

In your request, use the `tags` parameter to specify the tag key (`key`) and optional tag value (`value`) for each tag to add to the resource. The `tags` parameter specifies a string-to-string map of tag keys and their associated tag values.

For example, the following AWS CLI command adds a `Stack` tag key with a `Production` tag value to the specified job. This example is formatted for Microsoft Windows and it uses the caret (^) line-continuation character to improve readability.

```
C:\> aws macie2 tag-resource ^
--resource-arn arn:aws:macie2:us-east-1:123456789012:classification-
job/3ce05dbb7ec5505def334104bexample ^
--tags={"Stack\":"Production\"}
```

Where:

- `resource-arn` is the Amazon Resource Name (ARN) of the job to add a tag to.
- `Stack` is the tag key of the tag to add to the job.
- `Production` is the tag value for the specified tag key (`Stack`).

In the following example, the command adds several tags to the job:

```
C:\> aws macie2 tag-resource ^
--resource-arn arn:aws:macie2:us-east-1:123456789012:classification-
job/3ce05dbb7ec5505def334104bexample ^
--tags={"Owner\":"jane-doe\","CostCenter\":"12345\","Stack\":"Production\"}
```

For each tag in a `tags` map, both the `key` and `value` arguments are required. However, the value for the `value` argument can be an empty string. If you don't want to associate a tag value with a tag key, don't specify a value for the `value` argument. For example, the following AWS CLI command adds an `Owner` tag key with no associated tag value:

```
C:\> aws macie2 tag-resource ^
--resource-arn arn:aws:macie2:us-east-1:123456789012:classification-
job/3ce05dbb7ec5505def334104bexample ^
--tags={"Owner\":"\""} 
```

If a tagging operation succeeds, Macie returns an empty HTTP 204 response. Otherwise, Macie returns an HTTP 4xx or 500 response that indicates why the operation failed.

Reviewing tags for Amazon Macie resources

You can review the tags (both tag keys and tag values) for an Amazon Macie resource by using the Amazon Macie console or the Amazon Macie API. If you prefer to do this for multiple Macie resources at

the same time, you can use the [Tag Editor](#) on the AWS Resource Groups console or the [AWS Resource Groups Tagging API](#).

Console

Follow these steps to review a resource's tags by using the Amazon Macie console.

To review the tags for a resource

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. Depending on the type of resource whose tags you want to review, do one of the following:

- For an allow list, choose **Allow lists** in the navigation pane.

Then, in the table, select the check box for the list. Then choose **Manage tags** on the **Actions** menu.

- For a custom data identifier, choose **Custom data identifiers** in the navigation pane.

Then, in the table, select the check box for the custom data identifier. Then choose **Manage tags** on the **Actions** menu.

- For a filter or suppression rule, choose **Findings** in the navigation pane.

Then, in the **Saved rules** list, choose the edit icon (✎) next to the rule. Then choose **Manage tags**.

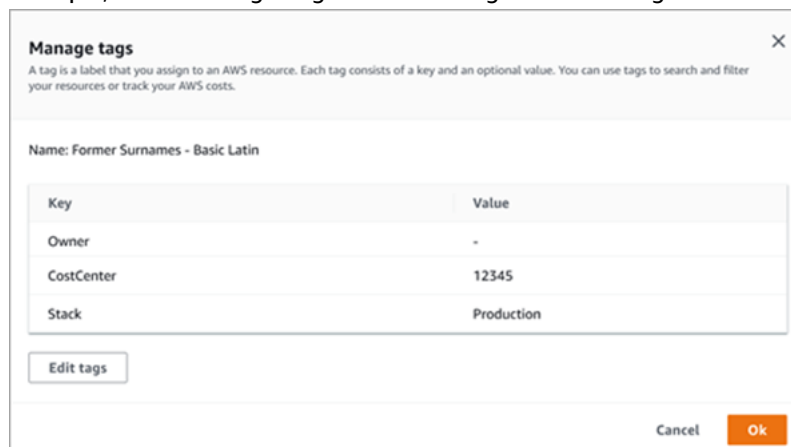
- For a member account in your organization, choose **Accounts** in the navigation pane.

Then, in the table, select the check box for the account. Then choose **Manage tags** on the **Actions** menu.

- For a sensitive data discovery job, choose **Jobs** in the navigation pane.

Then, in the table, select the check box for the job. Then choose **Manage tags** on the **Actions** menu.

The **Manage tags** window lists all the tags that are currently assigned to the resource. For example, the following image shows the tags that are assigned to a custom data identifier.



In this example, three tags are assigned to the custom data identifier: the **Owner** tag key with no associated tag value; the **CostCenter** tag key with **12345** as an associated tag value; and, the **Stack** tag key with **Production** as an associated tag value.

3. When you finish reviewing the tags, choose **Cancel** to close the window.

API

To retrieve and review the tags for an existing resource programmatically, use the [ListTagsForResource](#) operation of the Amazon Macie API. In your request, use the `resourceArn` parameter to specify the Amazon Resource Name (ARN) of the resource.

If you're using the AWS CLI, run the `list-tags-for-resource` command and use the `resource-arn` parameter to specify the ARN of the resource. For example:

```
C:\> aws macie2 list-tags-for-resource --resource-arn arn:aws:macie2:us-east-1:123456789012:classification-job/3ce05dbb7ec5505def334104bexample
```

If the operation succeeds, Macie returns a `tags` object that lists all the tags (both tag keys and tag values) for the resource. For example:

```
{
  "tags": {
    "Stack": "Production",
    "CostCenter": "12345",
    "Owner": ""
  }
}
```

Where `Stack`, `CostCenter`, and `Owner` are the tag keys that are assigned to the resource. `Production` is the tag value that's associated with the `Stack` tag key. `12345` is the tag value that's associated with the `CostCenter` tag key. The `Owner` tag key doesn't have an associated tag value.

To retrieve a list of all the Macie resources that have tags and all the tags that are assigned to each of those resources, use the [GetResources](#) operation of the AWS Resource Groups Tagging API. In your request, set the value for the `ResourceTypeFilters` parameter to `macie2`. To do this using the AWS CLI, run the `get-resources` command and set the value for the `resource-type-filters` parameter to `macie2`. For example:

```
C:\> aws resourcegroupstaggingapi get-resources --resource-type-filters "macie2"
```

If the operation succeeds, Resource Groups returns a `ResourceTagMappingList` array that contains the ARNs of all the Macie resources that have tags, and the tag keys and values that are assigned to each of those resources.

Editing tags for Amazon Macie resources

To edit the tags (tag keys or tag values) for an Amazon Macie resource, you can use the Amazon Macie console or the Amazon Macie API. To do this for multiple Macie resources at the same time, use the [Tag Editor](#) on the AWS Resource Groups console or the [AWS Resource Groups Tagging API](#).

Important


Editing the tags for a resource can affect access to the resource. Before you edit a tag key or value for a resource, review any AWS Identity and Access Management (IAM) policies that might use the tag to control access to resources.

Console

Follow these steps to edit a resource's tags by using the Amazon Macie console.

To edit the tags for a resource

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.

2. Depending on the type of resource whose tags you want to edit, do one of the following:
 - For an allow list, choose **Allow lists** in the navigation pane.
Then, in the table, select the check box for the list. Then choose **Manage tags** on the **Actions** menu.
 - For a custom data identifier, choose **Custom data identifiers** in the navigation pane.
Then, in the table, select the check box for the custom data identifier. Then choose **Manage tags** on the **Actions** menu.
 - For a filter or suppression rule, choose **Findings** in the navigation pane.
Then, in the **Saved rules** list, choose the edit icon () next to the rule. Then choose **Manage tags**.
 - For a member account in your organization, choose **Accounts** in the navigation pane.
Then, in the table, select the check box for the account. Then choose **Manage tags** on the **Actions** menu.
 - For a sensitive data discovery job, choose **Jobs** in the navigation pane.
Then, in the table, select the check box for the job. Then choose **Manage tags** on the **Actions** menu.
3. In the **Manage tags** window, choose **Edit tags**.
4. Do any of the following:
 - To add a tag value to a tag key, enter the value in the **Value** box next to the tag key.
 - To change an existing tag key, choose **Remove** next to the tag. Then choose **Add tag**. In the **Key** box that appears, enter the new tag key. Optionally enter an associated tag value in the **Value** box.
 - To change an existing tag value, choose **X** in the **Value** box that contains the value. Then enter the new tag value in the **Value** box.
 - To remove an existing tag value, choose **X** in the **Value** box that contains the value.
 - To remove an existing tag (both the tag key and tag value), choose **Remove** next to the tag.

A resource can have as many as 50 tags. A tag key can contain as many as 128 characters. A tag value can contain as many as 256 characters. The characters can be letters, numbers, spaces, or the following symbols: `_ . : / = + - @`
5. When you finish editing the tags, choose **Save**.

API

When you edit a tag for a resource programmatically, you overwrite the existing tag with new values. Therefore, the best way to edit a tag depends on whether you want to edit a tag key, a tag value, or both. To edit a tag key, [remove the current tag \(p. 322\)](#) and [add a new tag \(p. 316\)](#).

To edit or remove only the tag value that's associated with a tag key, overwrite the existing value by using the `TagResource` operation of the Amazon Macie API or, if you're using the AWS CLI, running the `tag-resource` command. In your request, use the `tags` parameter to specify the tag key whose tag value you want to change, and specify the new tag value for the key.

For example, the following command changes the tag value from `Production` to `Staging` for the `Stack` tag key that's assigned to the specified sensitive data discovery job. This example is formatted for Microsoft Windows and it uses the caret (^) line-continuation character to improve readability.

```
C:\> aws macie2 tag-resource ^
--resource-arn arn:aws:macie2:us-east-1:123456789012:classification-
job/3ce05dbb7ec5505def334104bexample ^
--tags={"Stack\":"Staging\"}

```

Where:

- `resource-arn` is the Amazon Resource Name (ARN) of the job.
- `Stack` is the tag key that's associated with the tag value to change.
- `Staging` is the new tag value for the specified tag key (`Stack`).

To remove the tag value for a tag key, don't specify a value for the `value` argument in the `tags` parameter. For example:

```
C:\> aws macie2 tag-resource ^
--resource-arn arn:aws:macie2:us-east-1:123456789012:classification-
job/3ce05dbb7ec5505def334104bexample ^
--tags={"Stack\":"\""}

```

If the operation succeeds, Macie returns an empty HTTP 204 response. Otherwise, Macie returns an HTTP 4xx or 500 response that indicates why the operation failed.

Removing tags from Amazon Macie resources

To remove tags from an Amazon Macie resource, you can use the Amazon Macie console or the Amazon Macie API. To do this for multiple Macie resources at the same time, use the [Tag Editor](#) on the AWS Resource Groups console or the [AWS Resource Groups Tagging API](#).

Important


Removing tags from a resource can affect access to the resource. Before you remove a tag, review any AWS Identity and Access Management (IAM) policies that might use the tag to control access to resources.

Console

Follow these steps to remove one or more tags from a resource by using the Amazon Macie console.

To remove a tag from a resource

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. Depending on the type of resource that you want to remove a tag from, do one of the following:
 - For an allow list, choose **Allow lists** in the navigation pane.
Then, in the table, select the check box for the list. Then choose **Manage tags** on the **Actions** menu.
 - For a custom data identifier, choose **Custom data identifiers** in the navigation pane.
Then, in the table, select the check box for the custom data identifier. Then choose **Manage tags** on the **Actions** menu.
 - For a filter or suppression rule, choose **Findings** in the navigation pane.

Then, in the **Saved rules** list, choose the edit icon () next to the rule. Then choose **Manage tags**.

- For a member account in your organization, choose **Accounts** in the navigation pane.
Then, in the table, select the check box for the account. Then choose **Manage tags** on the **Actions** menu.
 - For a sensitive data discovery job, choose **Jobs** in the navigation pane.
Then, in the table, select the check box for the job. Then choose **Manage tags** on the **Actions** menu.
3. In the **Manage tags** window, choose **Edit tags**.
 4. Do any of the following:
 - To remove only the tag value for a tag, choose **X** in the **Value** box that contains the value to remove.
 - To remove both the tag key and tag value (as a pair) for a tag, choose **Remove** next to the tag to remove.
 5. (Optional) To remove more tags from the resource, repeat the preceding step for each additional tag to remove.
 6. When you finish removing tags, choose **Save**.

API

To remove one or more tags from a resource programmatically, use the [UntagResource](#) operation of the Amazon Macie API. In your request, use the `resourceArn` parameter to specify the Amazon Resource Name (ARN) of the resource to remove a tag from. Use the `tagKeys` parameter to specify the tag key of the tag to remove. To remove only a specific tag value (not a tag key) from a resource, [edit the tag \(p. 320\)](#) instead of removing the tag.

If you're using the AWS CLI, run the `untag-resource` command and use the `resource-arn` parameter to specify the ARN of the resource to remove a tag from. Use the `tag-keys` parameter to specify the tag key of the tag to remove. For example, the following command removes the `Stack` tag (both the tag key and tag value) from the specified sensitive data discovery job:

```
C:\> aws macie2 untag-resource ^
--resource-arn arn:aws:macie2:us-east-1:123456789012:classification-
job/3ce05dbb7ec5505def334104bexample ^
--tag-keys Stack
```

Where `resource-arn` is the ARN of the job to remove a tag from, and `Stack` is the tag key of the tag to remove.

To remove multiple tags from a resource, add each additional tag key as an argument for the `tag-keys` parameter. For example:

```
C:\> aws macie2 untag-resource ^
--resource-arn arn:aws:macie2:us-east-1:123456789012:classification-
job/3ce05dbb7ec5505def334104bexample ^
--tag-keys Stack Owner
```

If the operation succeeds, Macie returns an empty HTTP 204 response. Otherwise, Macie returns an HTTP 4xx or 500 response that indicates why the operation failed.

Creating Amazon Macie resources with AWS CloudFormation

Amazon Macie integrates with AWS CloudFormation, which is a service that helps you model and set up your AWS resources so that you can spend less time creating and managing your resources and infrastructure. You create a template that describes all the AWS resources that you want (such as custom data identifiers), and AWS CloudFormation provisions and configures those resources for you.

When you use AWS CloudFormation, you can reuse your template to set up your Macie resources consistently and repeatedly. Describe your resources once, and then provision the same resources over and over in multiple AWS accounts and AWS Regions.

Topics

- [Amazon Macie and AWS CloudFormation templates \(p. 324\)](#)
- [Learn more about AWS CloudFormation \(p. 324\)](#)

Amazon Macie and AWS CloudFormation templates

To provision and configure resources for Amazon Macie and related services, you must understand [AWS CloudFormation templates](#). Templates are text files in JSON or YAML format. These templates describe the resources that you want to provision in your AWS CloudFormation stacks.

If you're unfamiliar with JSON or YAML, you can use AWS CloudFormation Designer, which is a graphic tool for creating and modifying AWS CloudFormation templates. With Designer, you can diagram your template resources by using a drag-and-drop interface, and then edit their details using the integrated JSON and YAML editor. For more information, see [What is AWS CloudFormation Designer?](#) in the *AWS CloudFormation User Guide*.

You can create AWS CloudFormation templates for the following types of Macie resources:

- Allow lists
- Custom data identifiers
- Filter rules and suppression rules for findings, also referred to as *findings filters*

For more information, including examples of JSON and YAML templates for these types of resources, see the [Amazon Macie resource type reference](#) in the *AWS CloudFormation User Guide*.

Learn more about AWS CloudFormation

To learn more about AWS CloudFormation, refer to the following resources:

- [AWS CloudFormation](#)
- [AWS CloudFormation User Guide](#)
- [AWS CloudFormation API Reference](#)

- [AWS CloudFormation Command Line Interface User Guide](#)

Suspending or disabling Amazon Macie

You can suspend or disable Amazon Macie in a specific AWS Region by using the Amazon Macie console or the Amazon Macie API. Macie then stops performing all activities for your account in that Region. You aren't charged for using Macie in the Region while it's suspended or disabled.

If you suspend or disable Macie, you can re-enable it at a later time.

Topics

- [Suspending Amazon Macie \(p. 326\)](#)
- [Disabling Amazon Macie \(p. 327\)](#)

Suspending Amazon Macie

If you suspend Amazon Macie, Macie retains the session identifier, settings, and resources for your account in the applicable AWS Region. For example, your existing findings remain intact and are retained for up to 90 days. However, when you suspend Macie, it stops performing all activities for your account in the applicable Region. This includes monitoring your Amazon Simple Storage Service (Amazon S3) data and running any sensitive data discovery jobs that are currently in progress. Macie also cancels all of your sensitive data discovery jobs in the Region.

Note

If your account is the Macie administrator account for an organization, you must remove all member accounts that are associated with your account before you suspend Macie for your account.

After you suspend Macie, you can re-enable it. You then regain access to your settings and resources in the applicable Region, and Macie resumes its activities for your account in that Region. This includes updating the S3 bucket inventory for your account and monitoring the buckets for security and access control. This doesn't include resuming or restarting your sensitive data discovery jobs. Sensitive data discovery jobs can't be resumed or restarted after they're cancelled.

This topic explains how to suspend Macie by using the Amazon Macie console. If you prefer to do this programmatically, you can use the [Account Administration](#) resource of the Amazon Macie API.

To suspend Macie

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. By using the AWS Region selector in the upper-right corner of the page, select the Region in which you want to suspend Macie.
3. In the navigation pane, choose **Settings**.
4. Choose **Suspend Macie**.
5. When prompted for confirmation, enter **Suspend**, and then choose **Suspend**.

To suspend Macie in additional Regions, repeat the preceding steps in each additional Region.

Disabling Amazon Macie

When you disable Amazon Macie, Macie stops performing all activities for your account in the applicable AWS Region. This includes monitoring your Amazon Simple Storage Service (Amazon S3) data and running any sensitive data discovery jobs that are currently in progress. Macie also deletes all the existing settings and resources that it stores or maintains for your account in the applicable Region, including your findings and sensitive data discovery jobs. Resources that you stored or published to other AWS services remain intact and aren't affected—for example, sensitive data discovery results in Amazon S3 and finding events in Amazon EventBridge.

Warning

If you disable Macie, you also permanently delete all of your existing findings, sensitive data discovery jobs, custom data identifiers, and other resources that Macie stores or maintains for your account in the applicable Region. These resources can't be recovered after they're deleted. To keep these resources and only pause your use of Macie, suspend Macie instead of disabling it.

If your account is part of an organization that centrally manages multiple Macie accounts, you must do the following before you disable Macie:

- If your account is a Macie member account, disassociate your account from its Macie administrator account.
- If your account is a Macie administrator account, remove all member accounts that are associated with your account and delete the associations between your account and those accounts.

This topic explains how to disable Macie by using the Amazon Macie console. If you prefer to do this programmatically, you can use the [Account Administration](#) resource of the Amazon Macie API.

To disable Macie

1. Open the Amazon Macie console at <https://console.aws.amazon.com/macie/>.
2. By using the AWS Region selector in the upper-right corner of the page, select the Region in which you want to disable Macie.
3. In the navigation pane, choose **Settings**.
4. Choose **Disable Macie**.
5. When prompted for confirmation, enter **Disable**, and then choose **Disable**.

To disable Macie in additional Regions, repeat the preceding steps in each additional Region.

Amazon Macie quotas

Your AWS account has certain default quotas, formerly referred to as *limits*, for each AWS service. These quotas are the maximum number of service resources or operations for your account. This topic lists the quotas that apply to Amazon Macie resources and operations for your account. Unless otherwise noted, each quota applies to your account in each AWS Region.

Some quotas can be increased, while others cannot. To request an increase to a quota, use the [Service Quotas console](#). To learn how to request an increase, see [Requesting a quota increase](#) in the *Service Quotas User Guide*. If a quota isn't available on the Service Quotas console, use the [service limit increase form](#) in AWS Support Center to request an increase to the quota.

Accounts

- Member accounts by invitation: 1,000
- Member accounts through AWS Organizations: 5,000

Findings

- Findings per run of a sensitive data discovery job: 100,000 + 5% of any remaining findings after the 100,000 threshold is met

This quota applies only to the Amazon Macie console and the Amazon Macie API. There isn't a quota for the number of finding events that Macie publishes to Amazon EventBridge or the number of sensitive data discovery results that Macie creates for each run of a job.

- Detection locations per sensitive data finding: 15
- Requests to retrieve and reveal sensitive data samples: 100 per day

This quota resets every 24 hours at 00:00:01 UTC+0.

- Size of an Amazon Simple Storage Service (Amazon S3) object to retrieve and reveal sensitive data samples from: 10 MB

This quota also applies to archive files.

- Filter rules and suppression rules per account: 1,000

Sensitive data discovery

- Monthly sensitive data discovery per account: 5 TB

This quota is adjustable. To increase the quota to as much as 1,000 TB (1 PB), use the [Service Quotas console](#). To request an increase for more than 1 PB, use the [service limit increase form](#).

- Custom data identifiers per account: 10,000
- Allow lists per account: 10, 1–5 allow lists that specify predefined text and 1–5 allow lists that specify regular expressions

Additional quotas apply to an allow list that specifies predefined text. The list can't contain more than 100,000 entries and the storage size of the list can't exceed 35 MB.

- S3 buckets per sensitive data discovery job: 1,000. If your account is the Macie administrator account for an organization, the buckets can span as many as 1,000 accounts in your organization.

This quota applies to a job only if you configure the job to analyze specific buckets that you select. It doesn't apply to jobs that use runtime bucket criteria to determine which buckets to analyze.

- Custom data identifiers per sensitive data discovery job: 30
- Allow lists per sensitive data discovery job: 10, 1–5 allow lists that specify predefined text and 1–5 allow lists that specify regular expressions
- Size of an individual file to analyze:
 - Adobe Portable Document Format (.pdf) file: 1,024 MB
 - Apache Avro object container (.avro) file: 8 GB
 - Apache Parquet (.parquet) file: 8 GB
 - GNU Zip compressed archive (.gz or .gzip) file: 8 GB
 - Microsoft Excel workbook (.xls or .xlsx) file: 512 MB
 - Microsoft Word document (.doc or .docx) file: 512 MB
 - Non-binary text file: 20 GB
 - TAR archive (.tar) file: 20 GB
 - ZIP compressed archive (.zip) file: 8 GB

If a file is larger than the applicable quota, Macie doesn't analyze any data in the file.

- Extraction and analysis of data in a compressed or archive file:
 - Storage size (compressed): 8 GB for a GNU Zip compressed archive (.gz or .gzip) file or ZIP compressed archive (.zip) file; 20 GB for a TAR archive (.tar) file
 - Nested archive depth: 10 levels
 - Extracted files: 1,000,000
 - Extracted bytes: 10 GB of data that uses a [supported file type or storage format \(p. 124\)](#)

If the metadata for a compressed or archive file indicates that the file contains more than 10 nested levels or exceeds the applicable quota for storage size or extracted bytes, Macie doesn't extract or analyze any data in the file.

If Macie begins to extract and analyze data in a compressed or archive file and subsequently determines that the file contains more than 1,000,000 files or exceeds the quota for extracted bytes, Macie stops analyzing data in the file and creates sensitive data findings and discovery results only for the data that was processed.

- Analysis of nested elements in structured data: 256 levels per file

This quota applies only to JSON (.json) and JSON Lines (.jsonl) files. If the nested depth of either type of file exceeds this quota, Macie doesn't analyze any data in the file.

- Detection locations per sensitive data discovery result: 1,000 per sensitive data detection type
- Detection of full names: 1,000 per file, including archive files

After Macie detects the first 1,000 occurrences of full names in a file, Macie stops incrementing the count and reporting location data for full names.

- Detection of mailing addresses: 1,000 per file, including archive files

After Macie detects the first 1,000 occurrences of mailing addresses in a file, Macie stops incrementing the count and reporting location data for mailing addresses.

AWS glossary

For the latest AWS terminology, see the [AWS glossary](#) in the *AWS General Reference*.

Document history for Amazon Macie

The following table describes the important changes to the documentation since the last release of Amazon Macie. For notification about updates to this documentation, you can subscribe to an RSS feed.

- **Latest documentation update:** August 30, 2022

Change	Description	Date
New feature (p. 331)	You can now create and use allow lists to specify text and text patterns that you want Macie to ignore when it inspects Amazon S3 objects for sensitive data. By using allow lists, you can define sensitive data exceptions for your particular scenarios or environment—for example, the names of public representatives for your organization, specific phone numbers, or sample data that your organization uses for testing.	August 30, 2022
New feature (p. 331)	To verify the nature of sensitive data that Macie finds in S3 objects, you can now configure and use Macie to retrieve samples of sensitive data reported by findings.	July 26, 2022
Updated functionality (p. 331)	In the AmazonMacieFullAccess policy , we updated the Amazon Resource Name (ARN) of the Macie service-linked role (<code>aws-service-role/macie.amazonaws.com/AWSServiceRoleForAmazonMacie</code>).	June 30, 2022
Updated functionality (p. 331)	We updated the AmazonMacieServiceRolePolicy policy , which is the policy that's attached to the Macie service-linked role (<code>AWSServiceRoleForAmazonMacie</code>). The policy no longer specifies actions and resources for Amazon Macie Classic. Amazon Macie Classic has been discontinued and is no longer available.	May 20, 2022

New functionality (p. 331)	Macie now includes the <code>OriginType</code> field in sensitive data findings that it publishes to AWS Security Hub . The <code>OriginType</code> field specifies how Macie found the sensitive data that produced a finding: <code>SENSITIVE_DATA_DISCOVERY_JOB</code> .	May 11, 2022
Updated content (p. 331)	Clarified how keyword and maximum match distance settings work for custom data identifiers .	April 22, 2022
New functionality (p. 331)	Macie now provides managed data identifiers that are designed to detect HTTP Basic Authorization headers, HTTP cookies, and JSON Web Tokens.	April 21, 2022
New content (p. 331)	Added descriptions and definitions of key concepts and terms for Macie.	March 16, 2022
New functionality (p. 331)	To calculate and display estimated costs when you create and configure sensitive data discovery jobs, Macie now retrieves pricing data for your AWS account from AWS Billing and Cost Management. To support this functionality, we added a Billing and Cost Management action to the AmazonMacieFullAccess policy .	March 7, 2022
New functionality (p. 331)	Macie now includes the <code>Sample</code> field in findings that it publishes to AWS Security Hub . The <code>Sample</code> field specifies whether a finding is a sample finding .	February 24, 2022
New content (p. 331)	Added information about using Amazon Virtual Private Cloud to establish a private connection between your VPC and Macie.	January 19, 2022

New functionality (p. 331)	You can now use the Amazon Macie console to assign and manage tags for custom data identifiers, filter and suppression rules for findings, sensitive data discovery jobs, and, if you're the Macie administrator for an organization, member accounts in your organization. A <i>tag</i> is a label that you optionally define and assign to certain types of AWS resources.	January 12, 2022
New content (p. 331)	Added information about using AWS Identity and Access Management to manage access to Macie.	December 20, 2021
New feature (p. 331)	When you create a custom data identifier , you can now define severity settings for sensitive data findings that it produces. With these settings, you can specify which severity to assign to a finding based on the number of occurrences of text that matches the custom data identifier's detection criteria.	November 4, 2021
New functionality (p. 331)	To learn about the different types of findings that Macie provides, you can generate sample findings . Sample findings use example data and placeholder values to demonstrate the kinds of information that Macie might include in each type of finding.	October 28, 2021
New functionality (p. 331)	Macie now includes the <code>OwnerAccountId</code> field in findings that it publishes to AWS Security Hub . This field specifies the account ID for the AWS account that owns the affected S3 bucket.	October 27, 2021
New content (p. 331)	Added information about centrally managing multiple Macie accounts . You can do this in two ways, by integrating Macie with AWS Organizations or by sending membership invitations from Macie.	October 13, 2021

New functionality (p. 331)	Your S3 bucket inventory now indicates if a bucket's permissions settings prevent Macie from retrieving information about the bucket or the bucket's objects and evaluating and monitoring the security and privacy of the bucket's data. In addition, we updated references to AWS KMS keys and customer managed keys to reflect current terminology.	October 5, 2021
New functionality (p. 331)	Macie now stores policy and sensitive data findings for 90 days instead of 30 days. If Macie created or updated a finding on or after August 31, 2021, you can access the finding for up to 90 days by using the Macie console or the Macie API. In certain AWS Regions, Macie began retaining findings for 90 days as early as September 27, 2021.	October 1, 2021
New feature (p. 331)	When you create a sensitive data discovery job , you can now specify which managed data identifiers you want the job to use when it analyzes S3 objects. With this feature, you can tailor a job's analysis to focus on certain types of sensitive data.	September 17, 2021
New functionality (p. 331)	Sensitive data findings now provide additional information to help you locate sensitive data in JSON and JSON Lines files.	July 6, 2021
Updated functionality (p. 331)	Macie now uses the <code>AwsS3Bucket</code> resource type in findings that it publishes to AWS Security Hub . (Macie previously set this value to <code>AWS::S3::Bucket</code> .) <code>AwsS3Bucket</code> is the resource type value that's used for S3 buckets in the AWS Security Finding Format (ASFF).	June 28, 2021

New feature (p. 331)	When you create a sensitive data discovery job , you can now define runtime criteria that determine which S3 buckets the job analyzes. With this feature, the scope of a job's analysis can dynamically adapt to changes to your bucket inventory.	May 15, 2021
New functionality (p. 331)	Your S3 bucket inventory and the Summary dashboard now provide encryption metadata and statistics indicating whether buckets require server-side encryption of new objects. In addition, you can now perform on-demand refreshes of object metadata for individual buckets in your bucket inventory.	April 30, 2021
New feature (p. 331)	You can now use Amazon CloudWatch Logs to monitor and analyze events that occur when you run sensitive data discovery jobs. To support this feature, we added CloudWatch Logs actions to the AWS managed policy for the Macie service-linked role .	April 14, 2021
Regional availability (p. 331)	Macie is now available in the AWS Asia Pacific (Osaka) Region.	April 5, 2021
New feature (p. 331)	You can now configure Macie to publish sensitive data findings to AWS Security Hub .	March 22, 2021
New content (p. 331)	Added information about monitoring and forecasting Macie costs and participating in the free trial.	February 26, 2021
Updated content (p. 331)	We replaced the term <i>master account</i> with the term <i>administrator account</i> . An administrator account is used to centrally manage multiple accounts .	February 12, 2021
New functionality (p. 331)	You can now refine the scope of sensitive data discovery jobs by using S3 object prefixes in custom include and exclude criteria.	February 2, 2021

Updated content (p. 331)	Macie now adheres to the finding type taxonomy of the AWS Security Finding Format (ASFF) when it publishes policy findings to AWS Security Hub.	January 28, 2021
New content (p. 331)	Added information about monitoring Amazon S3 data and assessing the security and privacy of that data.	January 8, 2021
Regional availability (p. 331)	Macie is now available in the AWS Africa (Cape Town) Region, the AWS Europe (Milan) Region, and the AWS Middle East (Bahrain) Region.	December 21, 2020
New functionality (p. 331)	If your account is a Macie administrator account, you can now create and run sensitive data discovery jobs that analyze data for as many as 1,000 buckets spanning as many as 1,000 accounts in your organization.	November 25, 2020
New functionality (p. 331)	Your S3 bucket inventory now indicates whether you've configured any one-time or periodic sensitive data discovery jobs to analyze data in a bucket. If you have, it also provides details about the job that ran most recently.	November 23, 2020
New content (p. 331)	Added information about filtering findings .	November 12, 2020
New functionality (p. 331)	Sensitive data findings now provide additional information to help you locate sensitive data in Apache Avro object containers, Apache Parquet files, and Microsoft Excel workbooks.	November 9, 2020
New feature (p. 331)	You can now use sensitive data findings to locate individual occurrences of sensitive data in S3 objects.	October 22, 2020
New feature (p. 331)	You can now pause and resume sensitive data discovery jobs .	October 16, 2020
New content (p. 331)	Added details about the severity scoring system for policy findings and sensitive data findings.	October 6, 2020

New features (p. 331)	You can now view statistics that indicate how much data Macie can analyze in individual S3 buckets when you run a sensitive data discovery job. In addition, you can now view the estimated cost of a job when you create a job.	September 3, 2020
New content (p. 331)	Added information about configuring, running, and managing sensitive data discovery jobs .	August 31, 2020
New functionality (p. 331)	Managed data identifiers can now detect certain types of personally identifiable information for Brazil.	July 31, 2020
Updated content (p. 331)	Added information about the supported syntax for regular expressions in custom data identifiers .	July 30, 2020
Updated content (p. 331)	Added keyword requirements for managed data identifiers , and increased the quota for the number of findings that each sensitive data discovery job can produce.	July 17, 2020
New content (p. 331)	Added information about using Amazon EventBridge and AWS Security Hub to monitor and process findings . This includes the EventBridge event schema for findings and event examples for policy and sensitive data findings.	June 22, 2020
New content (p. 331)	Added information about analyzing and suppressing findings .	June 17, 2020
New content (p. 331)	Added instructions for configuring Macie to store detailed discovery results in an S3 bucket .	June 2, 2020
New content (p. 331)	Added information about the types of sensitive data that Macie can detect, and encryption requirements for detecting sensitive data in Amazon S3 objects.	May 28, 2020
General availability (p. 331)	This is the initial public release of the <i>Amazon Macie User Guide</i> .	May 13, 2020