Amazon Macie

User Guide

Amazon Macie: User Guide

Table of Contents

What Is Amazon Macie?

Amazon Macie is a security service that uses machine learning to automatically discover, classify, and protect sensitive data in AWS. Macie recognizes sensitive data such as personally identifiable information (PII) or intellectual property, and provides you with dashboards and alerts that give visibility into how this data is being accessed or moved.

Important
Currently, Macie is supported in the following regions:

- US East (Northern Virginia)
- US West (Oregon)

Features of Amazon Macie

Data Discovery and Classification

Amazon Macie enables you to identify business-critical data and analyze access patterns and user behavior:

- Continuously monitor new data in your AWS environment
- Use artificial intelligence to understand access patterns of historical data
- Automatically access user activity, applications, and service accounts
- Use natural language processing (NLP) methods to understand data
- Intelligently and accurately assign business value to data and prioritize business-critical data based on your unique organization
- Create your own security alerts and custom policy definitions

Data Security

Amazon Macie enables you to be proactive with security compliance and achieve preventive security:

- Identify and protect various data types, including PII, PHI, regulatory documents, API keys, and secret keys
- Verify compliance with automated logs that allow for instant auditing
- Identify changes to policies and access control lists
- Observe changes in user behavior and receive actionable alerts
- Receive notifications when data and account credentials leave protected zones
- Detect when large quantities of business-critical documents are shared internally and externally

Pricing for Macie

Pricing in Macie is based on the content sources classified or processed. For detailed information about Macie pricing, see Amazon Macie Pricing.

Accessing Macie

You can work with Macie in any of the following ways:

Macie Console

Sign in to the AWS Management Console and open the Macie console at https://us-east-1.redirection.macie.aws.amazon.com/.

The console is a browser-based interface to access and use Macie.

Concepts and Terminology

As you get started with Amazon Macie, you can benefit from learning about its key concepts.

Account

A standard AWS account that contains your AWS resources. When you sign up for Amazon Web Services (AWS), your AWS account is automatically signed up for all services in AWS, including Macie. The account that you use to sign in to AWS at the time when you first enable Macie is designated as the **master** account.

You can also integrate other accounts with Macie. These other accounts are called **member** accounts.

> **Note**
> No users from the member accounts are granted access to the Macie console. Only the master account users have access to the Macie console where they can configure Macie and monitor and protect the resources in both master and member accounts.

Alert

A notification about a potential security issue that is discovered by Macie. Alerts are displayed on the Macie console and provide a comprehensive narrative about all activity that occurred over the last 24 hours.

Macie provides the following types of alerts:

- *Basic alerts* - alerts that are generated by the security checks that Macie performs. There are two types of basic alerts in Macie:
  - Managed (Macie-curated) basic alerts which you cannot modify. You can only enable or disable the existing managed basic alerts.
  - Custom basic alerts which you can create and modify to your exact specifications.
- *Predictive alerts* - automatic alerts based on activity within your AWS infrastructure that deviates from the established normal activity baseline. More specifically, Macie continuously monitors IAM user and role activity within your AWS infrastructure and builds a model of the normal behavior. It then looks for deviations from that normal baseline and when such activity is detected, it generates automatic predictive alerts. For example a user uploading or downloading a large number of S3 objects in a single day might trigger an alert if that user typically downloads one or two S3 objects over the course of a week.

For more information about alerts, including alert categories and details about the contents of Macie alerts, see Amazon Macie Alerts (p. 18).

Data source

The origin or location of a set of data. To classify and protect your data, Macie analyzes and processes information from the following data sources:

AWS CloudTrail event logs, Including Amazon S3 Object-Level API Activity

AWS CloudTrail provides you with a history of AWS API calls for your account, including API calls made using the AWS Management Console, the AWS SDKs, the command line tools, and higher-level AWS services. AWS CloudTrail also allows you to identify which users and accounts called AWS APIs for services that support CloudTrail, the source IP address that the calls were made from, and when the calls occurred. For more information, see What is AWS CloudTrail?

For data classification purposes, Macie utilizes CloudTrail's ability to capture object-level API activity on S3 objects (data events). For more information, see Logging Data and Management Events for Trails.

Amazon S3

In this release, Macie analyzes and processes data stored in the Amazon S3 buckets. You can select the S3 buckets that contain objects that you want Macie to classify and monitor.

Amazon Simple Storage Service (Amazon S3) is storage for the Internet. Amazon S3 stores data as objects within buckets. An object consists of a file and optionally any metadata that describes that file. To store an object in Amazon S3, you upload the file you want to store to a bucket. Buckets are the containers for objects. For more information, see Getting Started with Amazon Simple Storage Service.

User

In the context of Macie a user is the AWS Identity and Access Management (IAM) identity that makes the request. Macie uses the CloudTrail `userIdentity` element to distinguish the following user types. For more information, see CloudTrail userIdentity Element.

- Root – The request was made with your AWS account credentials.
- IAM user – The request was made with the credentials of an IAM user.
- Assumed role – The request was made with temporary security credentials that were obtained with a role via a call to the AWS Security Token Service (AWS STS) AssumeRole API.
- Federated user – The request was made with temporary security credentials that were obtained via a call to the AWS STS GetFederationToken API.
- AWS account – The request was made by another AWS account.
- AWS service – The request was made by an AWS account that belongs to an AWS service.

When specifying a user in the Macie console, for example searching for a user in the **Users** tab or constructing a query in the **Research** tab, or whitelisting a user in a basic alert with the index of **Cloudtrail data**, you must use a special Macie format called **macieUniqueId**. The macieUniqueId is a combination of the IAM UserIdentity element and the recipientAccountId. For more information, see the list of UserIdentity elements above and the definition of recipientAccountId in the CloudTrail Record Contents. The examples below list various structures of macieUniqueId, depending on the user identity type:

- 123456789012:root
- 123456789012:user/Bob
- 123456789012:assumed-role/Accounting-Role/Mary

For more detailed examples, see Users in Macie (p. 32).

Setting Up Amazon Macie

Topics

Step 1 - Enable Macie

When you sign up for AWS, your AWS account is automatically signed up for all services in AWS, including Amazon Macie. If you don't have an account, use the following procedure to create one.

To sign up for AWS

1. Open https://aws.amazon.com/, and then choose **Create an AWS Account**.

   Note
   This might be unavailable in your browser if you previously signed into the AWS Management Console. In that case, choose **Sign in to a different account**, and then choose **Create a new AWS account**.
2. Follow the online instructions.

   Part of the sign-up procedure involves receiving a phone call and entering a PIN using the phone keypad.

When you launch the Macie console for the first time, choose **Get Started** and complete the following procedure to enable Macie.

Important
The AWS account that you use to enable Macie is automatically designated as your master account. For more information, see Concepts and Terminology (p. 3).

To enable Amazon Macie

On the **Enable Amazon Macie** page, complete the following steps:

1. Verify region preferences by reviewing the value in the drop-down menu under the **Region** section.

   Note
   The region that you are currently signed in to is automatically selected.
2. **Required:** create the Identity Access Management (IAM) roles that provide Macie with access to your AWS account. You can create these roles and the required policies by launching the AWS CloudFormation stack templates found at the URLs listed below. These roles only need to be created once for use in all regions. For more information about AWS CloudFormation and CloudFormation stacks, see What is AWS CloudFormation? and Working with Stacks.

   - For US East (Virginia): Macie CloudFormation template for a master account
   - US West (Oregon): Macie CloudFormation template for a master account
3. **Required**: make sure that AWS CloudTrail is enabled in your account. If CloudTrail is not enabled, you must navigate to the AWS CloudTrail console and enable AWS CloudTrail. For more information, see Overview for Creating a Trail.
4. Grant Macie the required permissions to access your CloudTrail data by checking the checkbox in the **Permissions** section.
5. Choose **Enable Macie**.

Step 2 - Integrate Amazon S3 with Macie

To classify and protect your data, Macie analyzes and processes information from AWS CloudTrail and Amazon S3. Enabling CloudTrail in your account is required in order to enable Macie. Integrating S3 with Macie (in other words, initially specifying one or more S3 buckets for Macie to monitor) is not required in order to enable Macie. However we strongly recommend that you integrate with S3 as part of setting up Macie and specify at least one S3 bucket that contains objects that you want Macie to classify and monitor. For more information and details about how Macie classifies your data, see Classify Your Data with Macie (p. 11).

When you integrate with S3, Macie creates a trail and a bucket to store the logs about the S3 object-level API activity (data events) that it will now analyze along with other CloudTrail logs that it processes.

Important
Macie has a default limit on the amount of data that it can classify in an AWS account. Once this data limit is reached, Macie stops classifying the data in this AWS account. The default data classification limit is 3TB. You can contact customer support and request an increase to the default limit.

You can use the following procedure to integrate with S3 as part of setting up Macie:

1. Log in to AWS with the credentials of the AWS account that is serving as your Macie master account.
2. Navigate to the Macie console's **Integrations** tab and choose the **Services** tab.
3. In the **Services** tab, select the account id (master or member) in the **Select an account** drop-down. The **Amazon S3** tile is then displayed.
4. Choose the **Add** button in the **Amazon S3** tile.
5. On the **Selected S3 buckets and prefixes** page, choose the edit icon, and then select either the full buckets (recommended) or bucket/prefix combinations for Macie to monitor. You can select up to 250 S3 buckets and prefixes.

   Note
   You can only select S3 buckets in your current AWS region.

   Important
   When you specify an S3 bucket, by default, Macie only classifies objects that are added to the bucket after your bucket selection is complete. However you can instruct Macie to classify all existing objects in the specified S3 bucket by checking the **Classify all** checkbox. If you decide to **Classify all** S3 objects in a bucket, make sure to note the following values:

   - **Total size** - total size of the data within this bucket, which is the sum of the sizes of all the objects in the bucket. This value is displayed in the **Total size** column on the **Select S3 buckets for Macie to monitor** page.
   - **Total processed** - total size of the data from the bucket that Macie will actually classify. This value is displayed in the **Total processed** column on the **Select S3 buckets for Macie to monitor** page.
   - **Total cost estimate** - the total content classification cost estimate for each S3 bucket. This value is displayed in the **Total cost estimate** column on the **Select S3 buckets for Macie to monitor** page.

   These values are greyed out if **Classify all** checkbox is not checked. These values are calculated based on the snapshot of the contents of your S3 buckets taken within the last 48 hours.
   The **Total cost estimate** for each bucket is based on the **Total processed** value for that bucket and the general Macie pricing of $5 per GB processed by the content classification engine. The total cost estimates are provided only for S3 buckets, not for prefixes. For more information, see Amazon Macie Pricing.

The **Total processed** value for each bucket is calculated as follows:

- If an object's size is less than 1KB, 1KB is added to the **Total processed** value. Otherwise, the object's actual size is added to the **Total processed** value.
- If the object's size is greater than 20MB, 20MB is added to the **Total processed** value. Otherwise, the object's actual size is added to the **Total processed** value.
- For object in Amazon Glacier vaults, 0 is added to the **Total processed** value.

Note that it is possible for the **Total processed** value of an S3 bucket to be higher than the **Total size** value.

6. When you've finished your selections, choose **Review and Save**. And when you've finished reviewing your selections, choose **Save**.

Integrate Member Accounts and Amazon S3 with Amazon Macie

You can use the Macie console's **Integrations** tab to integrate member accounts with Macie and to integrate Amazon S3 with Macie for both your master account and member accounts. For more information about the master and member accounts, see Concepts and Terminology (p. 3).

Topics

Integrate Member Accounts with Macie

When you integrate member accounts with Macie you are enabling Macie to monitor resources and activity in these member accounts.

To integrate member accounts with Macie

1. Log in to AWS with the credentials of the AWS account that you want to integrate with Macie as a member account.
2. **Important**
   This is a required step if you're adding an AWS account as a Macie member account for the first time.
   You can skip this step if you're re-adding an AWS account as a Macie member account after you've disabled and re-enabled Macie.

   Create the IAM policies and roles that will grant this account the required permissions to be successfully integrated with Macie. The most efficient method of creating these roles and policies is by launching the AWS CloudFormation stack templates found at the URLs listed below. These roles only need to be created once for use in all regions. For more information about AWS CloudFormation and CloudFormation stacks, see What is AWS CloudFormation? and Working with Stacks.

   Important
   These CloudFormation stack templates for the member accounts differ from those used to enable the master accounts in Setting Up Amazon Macie (p. 5). Make sure to specify the master AWS account ID when running the stack templates below.

   - US East (Virginia): Macie CloudFormation template for a member account
   - US West (Oregon): Macie CloudFormation template for a member account
3. Log in to AWS with the Macie master account, navigate to the Macie console, and then choose the **Integrations** tab.
4. To integrate a member account, choose the **+** icon next to **Member accounts**.
5. In the **Add member AWS account(s)** pop up window, enter one or more AWS account IDs. Separate multiple account numbers with commas. Choose **Add accounts**.

   Important
   If you disable Macie, the service will no longer have access to the resources in your master or member accounts. If you decide to re-enable Macie, you will have to add member accounts to

it again using the steps above. However, the IAM policies and roles that you created for these member accounts in Step 2 above will remain intact after you disable Macie.

Specify Data for Macie to Monitor

In order for Macie to start monitoring and classifying your data, you must specify what data you want Macie to monitor and classify. You can use the **Integrations/Services** tab to specify the S3 buckets that contain the data that you want Macie to monitor.

Important
Currently, Macie can only monitor objects stored in Amazon S3 buckets.

Important
Macie has a default limit on the amount of data that it can classify in an AWS account. Once this data limit is reached, Macie stops classifying the data in this AWS account. The default data classification limit is 3TB. You can contact customer support and request an increase to the default limit.

You can integrate S3 with Macie (in other words, specifying one or more S3 buckets for Macie to monitor) during the initial Macie setup. For more information and instructions, see Setting Up Amazon Macie (p. 5). You can also use the the following procedure to integrate S3 with Macie at any time after you've enabled Macie.

1. Log in to AWS with the credentials of the AWS account that is serving as your Macie master account.
2. In the Macie console's **Integrations** tab, choose the **Services** tab.
3. In the **Services** tab, select the account id (master or member) in the **Select an account** drop-down. The **Amazon S3** tile is then displayed.
4. Choose the **Details** button in the **Amazon S3** tile.
5. On the **Selected S3 buckets and prefixes** page, choose the edit icon, and then select either the full buckets (recommended) or bucket/prefix combinations for Macie to monitor. You can select up to 250 S3 buckets and prefixes.

   Note
   You can only select S3 buckets in your current AWS region.

   Important
   When you specify an S3 bucket, by default, Macie only classifies objects that are added to the bucket after your bucket selection is complete. However you can instruct Macie to classify all existing objects in the specified S3 bucket by checking the **Classify all** checkbox. If you decide to **Classify all** S3 objects in a bucket, make sure to note the following values:

   - **Total size** - total size of the data within this bucket, which is the sum of the sizes of all the objects in the bucket. This value is displayed in the **Total size** column on the **Select S3 buckets for Macie to monitor** page.
   - **Total processed** - total size of the data from the bucket that Macie will actually classify. This value is displayed in the **Total processed** column on the **Select S3 buckets for Macie to monitor** page.
   - **Total cost estimate** - the total content classification cost estimate for each S3 bucket. This value is displayed in the **Total cost estimate** column on the **Select S3 buckets for Macie to monitor** page.

   These values are greyed out if **Classify all** checkbox is not checked. These values are calculated based on the snapshot of the contents of your S3 buckets taken within the last 48 hours.
   The **Total cost estimate** for each bucket is based on the **Total processed** value for that bucket and the general Macie pricing of $5 per GB processed by the content classification

engine. The total cost estimates are provided only for S3 buckets, not for prefixes. For more information, see Amazon Macie Pricing.
The **Total processed** value for each bucket is calculated as follows:

- If an object's size is less than 1KB, 1KB is added to the **Total processed** value. Otherwise, the object's actual size is added to the **Total processed** value.
- If the object's size is greater than 20MB, 20MB is added to the **Total processed** value. Otherwise, the object's actual size is added to the **Total processed** value.
- For object in Amazon Glacier vaults, 0 is added to the **Total processed** value.

Note that it is possible for the **Total processed** value of an S3 bucket to be higher than the **Total size** value.

6. When you've finished your selections, choose **Review and Save**. And when you've finished reviewing your selections, choose **Save**.

Encrypted Objects

If objects stored in your Amazon S3 buckets are encrypted, Macie might not be able to read and classify those objects:

- If your Amazon S3 objects are encrypted using Amazon S3-managed encryption keys (SSE-S3), Macie can read and classify the objects using the roles created during the setup process.
- If your Amazon S3 objects are encrypted using AWS KMS-managed keys (SSE-KMS), Macie can read and classify the objects only if you add the AWSMacieServiceCustomerServiceRole IAM role as a key user for the KMS customer master key (CMK). If you don't add the role as a key user for the KMS CMK, Macie cannot read and classify the objects. However, Macie still stores metadata on the object, including which KMS CMK was used to protect the object.
- If your Amazon S3 objects are encrypted using client-side encryption, Macie cannot read and classify the objects, but still stores metadata on the object.

Classify Your Data with Macie

Macie can help you classify your sensitive and business-critical data stored in the cloud. Currently, Macie analyzes and processes data stored in AWS S3 buckets. To classify your data, Macie also uses CloudTrail's ability to capture object-level API activity on S3 objects (data events). However, Macie only monitors CloudTrail data events if you specify at least one S3 bucket for Macie to monitor.

Once you specify the S3 bucket(s) for Macie to monitor, you enable Macie to continuously monitor and discover new data as it enters your AWS infrastructure. For more information on how to specify S3 buckets for Macie to monitor, see Specify Data for Macie to Monitor (p. 9).

> **Note**
> Macie's content classification engine processes up to the first 20 MB of an S3 object. For more information, see Specify Data for Macie to Monitor (p. 9).

Topics

Classify data with Macie

> **Important**
> Currently, Macie supports the following compression and archive file formats:

- BZIP
- GZIP
- LZO
- RAR
- SNAPPY
- AR
- CPIO
- Unix dump
- TAR
- zip
- XZ
- Pack200
- BZIP2
- 7z
- ARJ
- LZMA
- DEFLATE
- Brotli

Once Macie begins monitoring your data, it uses the following automatic content classification methods to identify and prioritize your sensitive and critical data and to accurately assign business value to your data.

Content Type

To classify your data objects by a content type, Macie uses an identifier that is embedded in the file header. Macie offers a set of managed (Macie-curated) content types, each with a designated risk level between 1 and 10.

Macie can assign only one content type to an object.

You cannot modify existing or add new content types. You can enable or disable the existing content types, thus instructing Macie to either include or exclude them in its data classification process.

To view, enable, or disable content types

1. In the Macie console, navigate to the **Settings** page.
2. In the **Classify data** section, choose **Content types**.
3. Choose any of the listed managed content types to view its details.

   To enable or disable a content type, on its details page, use the **Enabled/Disabled** drop-down menu, and then choose **Save**.

File Extension

Macie can also classify your objects by their file extensions. Macie offers a set of managed file extensions, each with a designated risk level between 1 and 10.

Macie can assign only one file extension to an object.

You cannot modify existing or add new file extensions. You can enable or disable the existing file extensions, thus instructing Macie to either include or exclude them in its data classification process.

To view, enable, or disable file extensions

1. In the Macie console, navigate to the **Settings** page.
2. In the **Classify data** section, choose **File extensions**.
3. Choose any of the listed managed file extensions to view its details.

   To enable or disable a file extension, on its details page, use the **Enabled/Disabled** drop-down menu, and then choose **Save**.

Theme

Object classification by theme is based on keywords that Macie searches for as it examines the contents of data objects. Macie offers a set of managed themes, each with a designated risk level between 1 and 10.

Macie can assign one or more themes to an object.

You cannot modify existing or add new themes. You can enable or disable the existing themes, thus instructing Macie to either include or exclude them in its data classification process.

To view, enable, or disable themes

1. In the Macie console, navigate to the **Settings** page.
2. In the **Classify data** section, choose **Themes**.

3. Choose any of the listed managed themes to view its details.

   To enable or disable a theme, on its details page, use the **Enabled/Disabled** drop-down menu, and then choose **Save**.

Regex

Object classification by regex is based on specific data or data patterns that Macie searches for as it examines the contents of data objects. Macie offers a set of managed regex, each with a designated risk level between 1 and 10.

Macie can assign one or more regex to an object.

You cannot modify existing or add new regex. You can enable or disable the existing regex, thus instructing Macie to either include or exclude them in its data classification process.

To view, enable, or disable regex

1. In the Macie console, navigate to the **Settings** page.
2. In the **Classify data** section, choose **Regex**.
3. Choose any of the listed managed regex to view its details.

   To enable or disable a regex, on its details page, use the **Enabled/Disabled** drop-down menu, and then choose **Save**.

Personally Identifiable Information (PII)

Object classification by PII is based on recognizing any personally identifiable artifacts based on industry standards such as NIST-80-122 and FIPS 199. Macie is able to recognize the following PII artifacts:

- Full names
- Mailing addresses
- Email addresses
- Credit card numbers
- IP addresses (IPv4 and IPv6)
- Drivers license IDs (USA)
- National identification numbers (USA)
- Birth dates

As part of PII object classification, Macie also assigns each matching object a PII impact of high, moderate, and low using the following criteria:

- High
  - >= 1 full name and credit card
  - >= 50 names or emails and any combination of other PII
- Moderate
  - >= 5 names or emails and any combination of other PII
- Low
  - 1-5 names or emails and any combination of PII
  - Any quantity of PII attributes above (without names or emails)

Support Vector Machine-Based Classifier

Another method that Macie uses to classify your S3 objects is a Support Vector Machine (SVM) classifier that classifies content inside your Macie-monitored S3 objects (text, token n-grams, and character n-grams) as well as their metadata features (document length, extension, encoding, headers) in order to achieve accurate classification of documents based on content. This Macie-managed classifier was trained against a large corpus of training data of various types and has been optimized to support accurate detection of various content types, including source code, application logs, regulatory documents, and database backups. Also, the classifier has the ability to generalize its detections. For example, if it detected a new kind of source code that doesn't match any of the types of source code that it is trained to recognize, it can generalize the detection as being just "source code".

> **Note**
> This data classification method is not surfaced in the Macie's Settings. The following list of artifacts is Macie-managed and cannot be edited, enabled, or disabled.

The SVM classifier in Macie is trained to detect the following content types:

- E-books
- Email
- Generic encryption keys
- Financial
  - SEC regulatory forms
- JSON
  - AWS CloudTrail logs
  - Jupyter notebooks
- Application logs
  - Apache format
  - AWS S3 server logs
  - Linux syslog
- Database
  - MongoDB backup
  - MySQLbackup
  - MySQL script
- Source code
  - F#
  - VimL
  - ActionScript
  - Assembly
  - Bash
  - Batchfile
  - C
  - Clojure
  - Cobol
  - CoffeeScript
  - CUDA
  - Erlang
  - Fortran
  - Go

  - Haskell
  - Java
  - JavaScript
  - LISP
  - Lua
  - Matlab
  - ObjectiveC
  - Perl
  - PHP
  - PowerShell
  - Processing
  - Python
  - R
  - Ruby
  - Scala
  - Swift
  - VHDL
- Web languages
  - CSS
  - HTML
  - XML

Object Risk Level

Through the automatic classification methods described above, a Macie-monitored object is assigned various risk levels based on each content type, file extension, theme, regex, PII, and SVM artifact that is assigned to it. The object's compound (final) risk level is then set to the highest value of its assigned risk levels.

Retention Duration for S3 Metadata

Macie stores metadata about your S3 objects for the default duration of 1 month. You can extend this duration up to 12 months.

Protect Your Data with Macie

Topics

Macie can help you monitor how your sensitive and business-critical data stored in the cloud is being used. Macie applies artificial intelligence to understand historical data access patterns and automatically assesses activity of users, applications and service accounts. This can help you detect unauthorized access and avoid data leaks.

Once you enable Macie it uses the following automated methods to protect your data:

CloudTrail events

Macie analyzes and processes a subset of CloudTrail-logged data and management events (API calls) that can occur within your infrastructure. Macie designates a risk level between 1 and 10 for each of the supported CloudTrail events.

You cannot modify existing or add new CloudTrail events to the Macie-managed list. You can enable or disable the supported CloudTrail events, thus instructing Macie to either include or exclude them in its data security process.

To view, enable, or disable supported CloudTrail events

1. In the Macie console, navigate to the **Settings** page.
2. In the **Protect data** section, choose **AWS CloudTrail events**.
3. Choose any of the listed events to view its details.

   To enable or disable an event, on its details page, use the **Enabled/Disabled** drop-down menu, and then choose **Save**.

CloudTrail errors

Macie analyzes and processes errors that can occur when Macie-supported subset of CloudTrail-logged data and management events (API calls) take place within your infrastructure. Macie designates a risk level between 1 and 10 for each of the supported CloudTrail errors.

You cannot modify existing or add new CloudTrail errors to the Macie-managed list. You can enable or disable the supported CloudTrail errors, thus instructing Macie to either include or exclude them in its data security process.

To view, enable, or disable supported CloudTrail errors

1. In the Macie console, navigate to the **Settings** page.
2. In the **Protect data** section, choose **AWS CloudTrail errors**.
3. Choose any of the listed errors to view its details.

To enable or disable an error, on its details page, use the **Enabled/Disabled** drop-down menu, and then choose **Save**.

Amazon Macie Alerts

An alert is a notification about a potential security issue discovered by Amazon Macie. This section describes the following information:

Topics

Basic and Predictive Macie Alerts

Macie generates two types of alerts:

- *Basic alerts* - alerts generated by the security checks that Macie performs. There are two types of basic alerts in Macie:
  - Managed (Macie-curated) basic alerts which you cannot modify. You can enable or disable the existing managed basic alerts.

    Note
    You can identify managed basic alerts by the value of `MacieDefault` in the **Created by** field in the **Basic alerts** list in the **Settings** tab.
  - Custom basic alerts which you can create and modify to your exact specifications. For more information, see Adding New and Editing Existing Custom Basic Alerts (p. 21).
- *Predictive alerts* - automatic alerts based on activity within your AWS infrastructure that deviates from the established 'normal' activity baseline. More specifically, Macie continuously monitors activity within your AWS infrastructure and builds a model of the 'normal' behavior. It then looks for deviations from that normal baseline and when such activity is detected, it generates automatic predictive alerts. For example a user uploading or downloading a large number of S3 objects in a single day might trigger an alert if that user typically downloads one or two S3 objects over the course of a week.

Alert Categories in Macie

Macie's basic alerts (managed and custom) can be of the following categories:

- **Configuration compliance** - related to compliance-controlled content, policy, configuration settings, control and data plane logging, and patch level.
- **Data compliance** - related to the discovery of compliance or security-controlled content, such as the existence of Personally Identifiable Information (PII), or access credentials.
- **File hosting** - related to you hosting possible malware, unsafe software, or attackers' command and control infrastructure through compromised hosts or storage services.

- **Service disruption** - configuration changes that can lead to you not being able to access resources in your own environment.
- **Ransomware** - potentially malicious software or activity designed to block your access to your own computer system until a sum of money is paid.
- **Suspicious access** - access to your resources from a risky anomalous IP address, user, or system, such as an attacker masquerading their connection through a compromised host.
- **Identity enumeration** - a series of API calls or accesses enumerating access levels to your systems that can possibly indicate the early stages of an attack or compromised credentials.
- **Privilege escalation** - successful or unsuccessful attempts to gain elevated access to resources that are normally protected from an application or user, or attempts to gain access to your system or network for an extended period of time.
- **Anonymous access** - attempted access to your resources from an IP address, user, or service with the intent to hide a user's true identity. Examples include the use of proxy servers, virtual private networks and other anonymity services such as Tor.
- **Open permissions** - identification of sensitive resources protected by potentially overly permissive (and thus risky) access control mechanisms.
- **Location anomaly** - an anomalous and risky location of the access attempt to your sensitive data.
- **Information loss** - an anomalous and risky access to your sensitive data.
- **Credentials loss** - possible compromise of your credentials.

To quickly view a list of your existing alerts of a particular category, choose that category from the **Categories** list on the Macie console's **Alerts** tab.

Severity Levels for Alerts in Macie

Each Macie alert has an assigned severity level. This reduces the need to prioritize one alert over another in your analyses. It can also help you determine your response when an alert highlights a potential problem. **Critical**, **High**, **Medium**, and **Low** levels all indicate a security issue that can result in compromised information confidentiality, integrity, and availability within your infrastructure. The **Informational** level simply highlights a security configuration detail of your Macie-monitored infrastructure. Following are recommended ways to respond to each:

- **Critical** – Describes a security issue that can result in a compromise of the information confidentiality, integrity, and availability within your infrastructure. We recommend that you treat this security issue as an emergency and implement an immediate remediation. The main difference between a Critical and High severity is that a Critical severity alert might be informing you of a security compromise of a large number of your resources or systems. A High severity alert is informing you of a security compromise of one or several of your resources or systems.
- **High** – Describes a security issue that can result in a compromise of the information confidentiality, integrity, and availability within your infrastructure. We recommend that you treat this security issue as an emergency and implement an immediate remediation.
- **Medium** – Describes a security issue that can result in a compromise of the information confidentiality, integrity, and availability within your infrastructure. We recommend that you fix this issue at the next possible opportunity, for example, during your next service update.
- **Low** - Describes a security issue that can result in a compromise of the information confidentiality, integrity, and availability within your infrastructure. We recommend that you fix this issue as part of one of your future service updates.
- **Informational** – Describes a particular security configuration detail of your infrastructure. Based on your business and organization goals, you can either simply make note of this information or use it to improve the security of your systems and resources.

Locating and Analyzing Macie Alerts

You can use the following procedure to locate and analyze existing alerts:

1. To view your generated alerts (including **Active** and **Archived** basic or predictive alerts), in the Macie console, navigate to the **Alerts** page.

   Each alert has a summary section that contains the following information:

   - Alert severity, which can be **Critical**, **High**, **Medium**, **Low**, or **Informational**. For more information, see Severity Levels for Alerts in Macie (p. 19).
   - A timestamp that indicates when the alert was generated or last updated.
   - The alert category - for more information, see Alert Categories in Macie (p. 18).
   - One of the following:
     - If the alert's index is **CloudTrail data**, a user that engaged in the activity that prompted Macie to generate the alert. For more information, see the definition of a 'user' concept in the context of Macie in Concepts and Terminology (p. 3)
     - If the alert's index is **S3 bucket properties** or **S3 objects**, a bucket name that was involved in or that contains the objects that were involved in the activity that prompted Macie to generate the alert.

       Important
       In Macie, each alert is based on one of the following:

       - For the alerts with the index of **CloudTrail data**, only one user - the IAM identity whose activity prompted Macie to generate the alert.
       - For the alerts with the index of **S3 bucket properties** or **S3 objects**, only one S3 bucket that was involved in or that contains objects that were involved in the activity that prompted Macie to generate the alert.
   - A number of comments that were left on the alert.
   - A total number of results, which can consist of a list of user sessions, or a list of 3 buckets, or a list of S3 objects that match the query that is included in the definition of the alert. For more information, see Adding New and Editing Existing Custom Basic Alerts (p. 21).
   - A number of views on the alert.
   - The AWS region in which the activity captured in this alert took place.
2. To analyze any alert further, choose the alert to expand its details pane. The following information is included in the alert details:

   - The alert summary that includes the description and the total number of results - a number of user sessions, or a number of S3 buckets, or a number of S3 objects that match the query that is included in the definition of the alert.
   - A list of the alert results. This is a list of user sessions, or a list of S3 buckets, or a list of S3 objects, depending the index that is specified in the definition for this alert. For more information, see Adding New and Editing Existing Custom Basic Alerts (p. 21).
     - If you specified **CloudTrail data** as the index, the alert details contain a list of user sessions that match the query specified in the alert definition for a particular user.
     - If you specified **S3 buckets** as the index, the alert details contain a list of S3 buckets that match the query specified in the alert definition for a particular user.
     - If you specified **S3 objects** as the index, the alert details contain a list of S3 objects that match the query specified in the alert definition for a particular user.

   You can choose each result to examine it further and view all its fields. For more information, see Researching AWS Data, Researching S3 Bucket Properties Data, or Researching S3 Objects Data sections in Using the Macie Research Tab (p. 36)

You can also use the **Research** looking glass icon to navigate to the **Research** tab and view the results of a particular alert there. The **Query Parser** in the **Research** tab is then pre-populated with the query that can be used to generate these results.

Adding New and Editing Existing Custom Basic Alerts

You can use the following procedure to add new and edit existing custom basic alerts:

1. In the Macie console, navigate to the **Settings** page and choose the icon for **Basic alerts**.
2. On the **Basic alerts** page, either choose the edit icon for the alert that you want to modify. Or, to add a new basic alert, choose **Add new**.
3. Do one of the following:

   - If you're editing the existing alert, make your changes, including enabling or disabling the alert, and then choose **Save**.
   - If you're adding a new alert, on the **Basic alert definition** page, specify the following:
     - Alert title - for example, "An S3 bucket has an IAM policy that grants 'read' rights to everyone."
     - Description for the alert - for example, "An IAM policy on an S3 bucket contains a clause that effectively grants 'read' access to any user. It is recommended that you audit this S3 bucket and its data and confirm that this is intentional.
     - Alert category - for more information, see Alert Categories in Macie (p. 18).
     - Alert query - a query that describes the activity that you want Macie to generate an alert about. For example, `s3_world_readability:"true"`. This query looks for an IAM policy on an S3 bucket that grants 'read' access to any user. For more information about constructing queries, see Constructing Queries in Macie (p. 36).

       Note
       You can use the looking glass icon next to an existing alert to navigate to the **Research** tab. This alert's query is then automatically displayed in the **Query Parser** and the results of this query are displayed in the **Research** tab.
   - Query index - this is the repository of data against which Macie will run the query specified in this alert. You can select either CloudTrail data, S3 buckets, or S3 objects. Depending on your selection, the alert will contain either a list of cloud trail user sessions (5-minute aggregates of raw CloudTrail data), a list of S3 buckets, or a list of S3 objects that match the activity that your alert defines.
   - A minimum number of activity matches that must occur before an alert is generated.
   - Alert severity - for more information, see Severity Levels for Alerts in Macie (p. 19)
   - Whitelisted users or whitelisted buckets, depending on the selected alert index. If you whitelist a user or a bucket, Macie will not generate an alert for this user or bucket when they are involved in the activity that the alert defines.

     Important
     In Macie, each alert is based on one of the following:
     - For the alerts with the index of **CloudTrail data**, only one user - the IAM identity whose activity prompted Macie to generate the alert.
     - For the alerts with the index of **S3 bucket properties** or **S3 objects**, only one S3 bucket that was involved in or that contains objects that were involved in the activity that prompted Macie to generate the alert.

     When whitelisting a user in a basic alert with the index of CloudTrail data, you must use a special Macie format called **macieUniqueId**. For example, 123456789012:root or

123456789012:user/Bob, or 123456789012:assumed-role/Accounting-Role/Mary, depending on the identity type of the user you want to whitelist. For more information, see the definition of the 'user' concept in Users in Macie (p. 32).

- Specify whether this alert is enabled or disabled.

Working with Existing Alerts

You can use the following procedure to archive or unarchive alerts or to choose edit the existing basic alerts.

1. In the Macie console, navigate to the **Alerts** page and locate the alert that you want to either archive, unarchive (if it's an archived alert), or edit.
2. Choose the down arrow in the alert summary pane and then choose either of the following:

   - **Archive**

     Note
     Or **Unarchive**, if this is an archived alert.
   - **Edit basic alert**

     Important
     This option is not available for predictive alerts. You cannot edit predictive alerts, which are automatically generated by Macie based on activity within your AWS infrastructure that deviates from the established 'normal' activity baseline. For more information, see Basic and Predictive Macie Alerts (p. 18).

Group Archiving Alerts

You can use the following procedure to group archive alerts:

1. In the Macie console's **Alerts** page, choose **Group Archive**.
2. In the **Group archive** window, use the available settings to archive or unarchive multiple alerts at the same time.

Whitelisting Users or Buckets for Basic Alerts

Macie allows you to whitelist users (if the alert's index is **CloudTrail data**) and buckets (if the alert's index is **S3 bucket properties** or **S3 objects**) for both Macie-managed and custom basic alerts.

Note
Macie does not allow you to whitelist users or buckets for predictive alerts.

You can use the following procedure to whitelist a specific user or a specific bucket that engaged in or was involved in the activity that prompted Macie to generate a specific alert.

Important
In Macie, each alert is based on one the following:

- For the alerts with the index of **CloudTrail data**, only one user - the IAM identity whose activity prompted Macie to generate the alert.
- For the alerts with the index of **S3 bucket properties** or **S3 objects**, only one S3 bucket that was involved in or that contains objects that were involved in the activity that prompted Macie to generate the alert.

Whitelist users or S3 buckets for custom basic alerts using the Alerts tab

1. In the Macie console's **Alerts** tab, locate the custom basic alert for which you want to whitelist a user or (S3 bucket) listed in the alert's summary.
2. Choose the down arrow in the alert summary pane and then choose **Whitelist user** (if this alert's index is **CloudTrail data**) or **Whitelist bucket** (if the alert's index is **S3 bucket properties** or **S3 objects**).
3. In the **Whitelist user** (or **Whitelist bucket**) window, verify the user or bucket that you want to whitelist (automatically pre-selected and matching the user or bucket listed in the alert's summary), and then choose **Submit**.

You can use the following procedure to whitelist multiple users or buckets at the same time for custom basic alerts.

Whitelist users or S3 buckets for custom basic alerts using the Settings tab

1. In the Macie console's **Settings** tab, choose **Basic alerts**, and then locate the custom basic alert for which you want to whitelist users or S3 buckets.
2. Choose the edit icon next to the alert.
3. Specify users or S3 buckets that you want to whitelist in either **Whitelisted users** (if this alert's index is **CloudTrail data**) or **Whitelisted buckets** (if the alert's index is **S3 bucket properties** or **S3 objects**) fields and then choose **Save**.

   Note
   When whitelisting a user in a basic alert with the index of CloudTrail data, you must use a special Macie format called **macieUniqueId**. For example, 123456789012:root or 123456789012:user/Bob, or 123456789012:assumed-role/Accounting-Role/Mary, depending on the identity type of the user you want to whitelist. For more information, see the definition of the 'user' concept in Users in Macie (p. 32).

Whitelist users or S3 buckets for Macie-managed basic alerts

1. In the Macie console's **Alerts** tab, locate the Macie-managed basic alert for which you want to whitelist users or S3 buckets.
2. Choose the down arrow in the alert summary pane and then choose **Whitelist user** (if this alert's index is **CloudTrail data**) or **Whitelist bucket** (if the alert's index is **S3 bucket properties** or **S3 objects**).
3. In the **Whitelist user** or **Whitelist bucket** window, check the **Clone and disable the default managed alert** checkbox and then choose **Submit**.
4. Navigate to the Macie console's **Settings** tab.

   Note that the original managed alert that you worked with in the previous step is now disabled. Note also that this alert has been cloned into a new custom basic alert. For example, if your original managed basic alert was called "An S3 bucket has an IAM policy that grants 'read' rights to everyone", this alert is now disabled and a new (cloned) custom basic alert called "An S3 bucket has an IAM policy that grants 'read' rights to everyone (modified)" is created.
5. Choose the edit icon next to the cloned custom basic alert.
6. Specify users or S3 buckets that you want to whitelist in either **Whitelisted users** (if this alert's index is **CloudTrail data**) or **Whitelisted buckets** (if the alert's index is **S3 bucket properties** or **S3 objects**) fields and then choose **Save**.

   Note
   When whitelisting a user in a basic alert with the index of CloudTrail data, you must use a special Macie format called **macieUniqueId**. For example, 123456789012:root or 123456789012:user/Bob, or 123456789012:assumed-role/Accounting-Role/Mary,

depending on the identity type of the user you want to whitelist. For more information, see the definition of the 'user' concept in Users in Macie (p. 32).

Using the Macie Dashboard

The Macie **Dashboard** draws a comprehensive picture of all of your Macie-monitored data and activity. This topic describes the metrics and views that you can use in the **Dashboard** to view your monitored data grouped by various interest points. Each metric and view provides you with one or more ways of navigating to the Macie console's **Research** tab, where you can construct and run queries in the query parser and conduct in-depth investigative research of your Macie-monitored data and activity.

Dashboard Metrics

The following **Dashboard** metrics allow you to view your monitored data grouped by several key interest points:

- **High-risk S3 objects** - While classifying data (p. 11), Macie assigns a risk value to each monitored data object. This is Macie's way of helping you identify and prioritize your sensitive data over other, less business-critical data. This metric allows you to see all of your Macie-monitored data objects with a risk levels of 8 through 10.
- **Total event occurrences** - As part of securing data (p. 16), Macie analyzes and processes CloudTrail-logged events (API calls) that occur within your infrastructure. This metric provides the total count of all Macie-monitored event occurrences that took place within your infrastructure since you enabled Macie.
- **Total user sessions** - A user session is a 5-minute aggregate of CloudTrail data. This metric provides the total count of all user sessions of CloudTrail data that Macie analyzed and processed since it was enabled.

Dashboard Views

Follow this procedure to use the predefined Macie **Dashboard** views and generate distinct subsets of your Macie-monitored data and activity:

To use Macie Dashboard views

1. Choose the corresponding icon to select any of the following views to display various subsets of your Macie-monitored data and activity:
   - S3 objects for a selected time range (p. 26)
   - S3 objects (p. 26)
   - S3 objects by PII (p. 27)
   - S3 objects by ACL (p. 27)
   - CloudTrail events and associated users (p. 28)
   - CloudTrail errors and associated users (p. 29)
   - Activity location (p. 30)
   - AWS CloudTrail events (p. 30)
   - Activity ISPs (p. 30)

- AWS CloudTrail user identity types (p. 31)

2. If present in the selected view, locate and move the **Minimum risk** slider to the desired value. The **Minimum risk** slider enables you to only view items with the assigned risk equal to and greater than the selected value.

S3 objects for selected time range

This view provides a visual representation of your monitored S3 objects that match the following search criteria:

- At least one of the object's assigned themes is of the top 20 most frequently assigned themes
- The object's assigned risk is either equal to or greater than the value selected on the **Minimum risk** slider
- S3 object was last modified during one of the following time ranges:
  - The past six months
  - Between the date when Macie was enabled and a date six months prior to today

To navigate from this view to the **Research** tab, you can select (double-click) any of the squares that represent the displayed time ranges or themes. Your selection is automatically translated into a query that is displayed in the query parser in the **Research** tab.

You can follow this sample procedure:

1. In the Macie **Dashboard**, select the **S3 objects over selected time range** view.
2. Set the **Minimum risk** slider to 5.
3. In the generated graph, double-click the square next to **Range: 0 - 6 months ago**.

   As a result, you are redirected to the **Research** tab with the following query automatically displayed in the query parser:

   ```
   themes:* AND dlp_risk:[5 TO *] AND @timestamp:[now-6M/M TO now]
   ```

   This query matches your selection to view the Macie-monitored S3 objects that are assigned one or more of the top 20 most frequently assigned themes, that have an assigned risk of 5 or higher, and that were last modified at some point in the past 6 months. The results of this query are also displayed. You can then modify the query result controls available on the **Research** tab and run the query again, and conduct in-depth investigative research of the generated results. For more information, see Using the Macie Research Tab (p. 36).

S3 objects

This view provides the complete list of your Macie-monitored S3 objects, grouped by the assigned themes. For each theme, a percentage that this theme represents of the total number of your Macie-monitored S3 objects is displayed, as well as the total count of the S3 objects that were assigned this theme.

To navigate from this view to the **Research** tab, you can choose the looking glass icon next to any of the displayed themes. Your selection is automatically translated into a query that is displayed in the query parser in the **Research** tab.

You can follow this sample procedure:

1. In the Macie **Dashboard**, select the **S3 objects** view.

2. From the generated list of S3 objects, choose the looking glass icon next to, for example, **json/aws_cloudtrail_logs**.

   As a result, you are redirected to the **Research** tab with the following query automatically displayed in the query parser:

   ```
   themes:"json/aws_cloudtrail_logs"
   ```

   This query matches your selection to view the Macie-monitored S3 objects with the assigned theme of **json/aws_cloudtrail_logs**. The results of this query are also displayed. You can then modify the query result controls available on the **Research** tab and run the query again, and conduct in-depth investigative research of the generated results. For more information, see Using the Macie Research Tab (p. 36).

S3 objects by PII

This view provides the following lists:

- **S3 objects by PII priority**

  This is a complete list of your S3 objects that contain PII artifacts, grouped by the Macie-assigned PII priority. For each PII priority level, a percentage that the number of objects with this level represents of the total number of the S3 objects with PII artifacts is displayed, as well as the total count of the S3 objects with this PII priority level.
- **S3 objects by PII types**

  This is a complete list of your S3 objects that contain PII artifacts, grouped by the PII artifact types. For each PII artifact type, a percentage that the number of objects with PII artifacts of this type represents of the total number of the S3 objects with PII artifacts is displayed, as well as the total count of the S3 objects with PII artifacts of this type.

For more information about PII-based object classification, see Classify Your Data with Macie (p. 11).

To navigate from this view to the **Research** tab, you can choose the looking glass icon next to any of the displayed PII impacts or PII types. Your selection is automatically translated into a query that is displayed in the query parser in the **Research** tab.

You can follow this sample procedure:

1. In the Macie **Dashboard**, select the **S3 objects by PII** view.
2. For example, let's generate a list of S3 objects with low PII priority. In the **S3 objects by PII priority** list, choose the looking glass icon next to the low PII priority.

   As a result, you are redirected to the **Research** tab with the following query automatically displayed in the query parser:

   ```
   pii_impact:"low"
   ```

   The results of this query are also displayed. You can then modify the query result controls available on the **Research** tab and run the query again, and conduct in-depth investigative research of the generated results. For more information, see Using the Macie Research Tab (p. 36).

S3 objects by ACL

This view provides the following lists:

- **S3 objects by ACL URIs**

  This is a complete list of URIs that appear in access control lists (ACL)s that are attached to your S3 objects. For each URI, a percentage that the number of objects with ACLs attached that contain this URI represents of the total number of the S3 objects with ACLs attached is displayed, as well as the total count of the S3 objects with ACLs attached that contain this URI.
- **S3 objects by ACL display names**

  This is a complete list of user display names that appear in ACLs that are attached to your S3 objects. For each display name, a percentage that the number of objects with ACLs attached that contain this display name represents of the total number of the S3 objects with ACLs attached is displayed, as well as the total count of the S3 objects with ACLs attached that contain this display name.
- **S3 objects by ACL permissions**

  This is a complete list of access permissions that appear in ACLs that are attached to your S3 objects. For each permissions level, a percentage that the number of objects with ACLs attached that contain this permission level represents of the total number of the S3 objects with ACLs attached is displayed, as well as the total count of the S3 objects with ACLs attached that contain this permission level.

To navigate from this view to the **Research** tab, you can choose the looking glass icon next to any of the displayed URIs, ACL display names, and ACL permissions. Your selection is automatically translated into a query that is displayed in the query parser in the **Research** tab.

You can follow this sample procedure:

1. In the Macie **Dashboard**, select the **S3 objects by ACL** view.
2. For example, let's generate a list of S3 objects with attached ACLs that contain full control permissions. In the **S3 objects by ACL permissions** list, choose the looking glass icon next to the **FULL_CONTROL** permission.

   As a result, you are redirected to the **Research** tab with the following query automatically displayed in the query parser:

   ```
   object_acl.Grants.Permission:"FULL_CONTROL"
   ```

   The results of this query are also displayed. You can then modify the query result controls available on the **Research** tab and run the query again, and conduct in-depth investigative research of the generated results. For more information, see Using the Macie Research Tab (p. 36).

CloudTrail events and associated users

This view provides the following lists:

- **AWS CloudTrail events**

  This is a visual representation of the top 20 (by assigned risk and based on the value selected on the **Minimum risk** slider) CloudTrail data and management events that occurred during the last 60 days. You can use the available **Daily** or **Weekly** radio buttons to modify the graph to view daily or weekly results.

  To navigate from this view to the **Research** tab, you can select (double-click) any square that represents a particular event that you would like to investigate further. The number in parenthesis next to the event name represents the number of user sessions (5-minute aggregates of CloudTrail data) in which this event is present. In the **Research** tab, your selection is automatically translated into a query that is displayed in the query parser.
- AWS Cloud trail associated users

This is a visual representation of the users associated with the top 20 (by assigned risk and based on the value selected on the **Minimum risk** slider) CloudTrail data and management events that occurred during the last 60 days. You can use the available **Daily** or **Weekly** radio buttons to modify the graph to view daily or weekly results.

To navigate from this view to the **Research** tab, you can select (double-click) any square that represents a particular error that you would like to investigate further.The number in parenthesis next to the user name represents the number of user sessions (5-minute aggregates of CloudTrail data) with which this user is associated. In the **Research** tab, your selection is automatically translated into a query that is displayed in the query parser.

You can follow this sample procedure:

1. In the Macie **Dashboard**, select the **CloudTrail events and associated users** view.
2. Set the **Minimum risk** slider to 1.
3. For example, let's generate a list of user sessions in which **PutRestApi** event is present. Double-click on the square next to **PutRestApi**.

   As a result, you are redirected to the **Research** tab with the following query automatically displayed in the query parser:

   ```
   eventNameIsp.key.keyword:"PutRestApi" AND @timestamp:[now-60d TO now]
   ```

   The results of this query are also displayed. You can then modify the query result controls available on the **Research** tab and run the query again, and conduct in-depth investigative research of the generated results. For more information, see Using the Macie Research Tab (p. 36).

CloudTrail errors and associated users

This view provides the following lists:

- **AWS CloudTrail errors**

  This is a visual representation of the top 20 (by assigned risk and based on the value selected on the **Minimum risk** slider) CloudTrail errors that occurred during the last 60 days. You can use the available **Daily** or **Weekly** radio buttons to modify the graph to view daily or weekly results.

  To navigate from this view to the **Research** tab, you can select (double-click) any square that represents a particular error that you would like to investigate further. The number in parenthesis next to the error name represents the number of user sessions (5-minute aggregates of CloudTrail data) in which this error is present. In the **Research** tab, your selection is automatically translated into a query that is displayed in the query parser.
- AWS Cloud trail associated users

  This is a visual representation of the users associated with the top 20 (by assigned risk and based on the value selected on the **Minimum risk** slider) CloudTrail errors that occurred during the last 60 days. You can use the available **Daily** or **Weekly** radio buttons to modify the graph to view daily or weekly results.

  To navigate from this view to the **Research** tab, you can select (double-click) any square that represents a particular error that you would like to investigate further. The number in parenthesis next to the user name represents the number of user sessions (5-minute aggregates of CloudTrail data) with which this user is associated. In the **Research** tab, your selection is automatically translated into a query that is displayed in the query parser.

You can follow this sample procedure:

1. In the Macie **Dashboard**, select the **CloudTrail errors and associated users** view.
2. Set the **Minimum risk** slider to 1.
3. For example, let's generate a list of user sessions in which **Client.InvalidPermission.NotFound** error is present. Double-click on the square next to **Client.InvalidPermission.NotFound**.

   As a result, you are redirected to the **Research** tab with the following query automatically displayed in the query parser:

   ```
   eventNameErrorCode.secondary:"Client.InvalidPermission.NotFound" AND
   @timestamp:[now-60d TO now]
   ```

   The results of this query are also displayed. You can then modify the query result controls available on the **Research** tab and run the query again, and conduct in-depth investigative research of the generated results. For more information, see Using the Macie Research Tab (p. 36).

Activity location

This view includes a map that shows the locations of activity that Macie is monitoring for a selected time period. To view details, you can use the available time period pull-down menu (past 15 days, past 30 days, past 90 days, or past year), and then choose any location pin.

To navigate from this view to the **Research** tab, choose the number of events that appears in a tooltip window for a location pin. In the **Research** tab, your selection is automatically translated into a query that is displayed in the query parser. For example, you can auto-generate the following query to display a list of user sessions that occurred in the past 15 days in Seattle:

```
geoLocation.key:"Seattle:UnitedStates:47.6145:-122.348" AND @timestamp:[now-15d
TO now]
```

You can then modify the query result controls available on the **Research** tab and run the query again, and conduct in-depth investigative research of the generated results. For more information, see Using the Macie Research Tab (p. 36).

AWS CloudTrail events

AWS CloudTrail events

This view provides the complete list of your Macie-monitored CloudTrail data and management events. For each event, the total count of the user sessions (5-minute integrations of CloudTrail data) in which this event is present and the percentage that this total represents of the total number of user sessions is displayed.

To navigate from this view to the **Research** tab, you can choose the looking glass icon next to any of the displayed events. Your selection is automatically translated into a query that is displayed in the query parser in the **Research** tab. For example, you can auto-generate the following query to view all user sessions in which the AssumeRole event is present:

```
eventNameIsp.key.keyword:"AssumeRole"
```

You can then modify the query result controls available on the **Research** tab and run the query again, and conduct in-depth investigative research of the generated results. For more information, see Using the Macie Research Tab (p. 36).

Activity ISPs

Activity ISPs

This view provides the complete list of your Macie-monitored CloudTrail data and management events, grouped by the associated internet service providers (ISPs). For each ISP, the total count of the user sessions (5-minute integrations of CloudTrail data) in which this ISP is present and the percentage that this total represents of the total number of user sessions is displayed.

To navigate from this view to the **Research** tab, you can choose the looking glass icon next to any of the displayed themes. Your selection is automatically translated into a query that is displayed in the query parser in the **Research** tab. For example, you can auto-generate the following query to view all user sessions that are associated with Amazon:

```
eventNameIsp.secondary.keyword:"Amazon"
```

You can then modify the query result controls available on the **Research** tab and run the query again, and conduct in-depth investigative research of the generated results. For more information, see Using the Macie Research Tab (p. 36).

AWS CloudTrail user identity types

This view provides the complete list of your Macie-monitored CloudTrail data and management events, grouped by the user identity type (for more information, see the definition for 'user' in Concepts and Terminology (p. 3). For each user identity type, the total count of the user sessions (5-minute integrations of CloudTrail data) in which this user identity type is present and the percentage that this total represents of the total number of user sessions is displayed.

To navigate from this view to the **Research** tab, you can choose the looking glass icon next to any of the displayed themes. Your selection is automatically translated into a query that is displayed in the query parser in the **Research** tab. For example, you can auto-generate the following query to view all user sessions that contain requests that were originated by the **AssumedRole** user identity type:

```
userIdentityType.key:"AssumedRole"
```

You can then modify the query result controls available on the **Research** tab and run the query again, and conduct in-depth investigative research of the generated results. For more information, see Using the Macie Research Tab (p. 36).

Users in Macie

The **Users** tab can help you draw a comprehensive picture of all of the Macie-monitored data and activity for a particular selected user. This topic describes how to search for the users whose activity you want to investigate further in the **Users** tab and the views that you can use in this tab to see the selected users' monitored data grouped by various interest points. Each view provides you with one or more ways of navigating to the Macie console's **Research** tab, where you can construct and run queries in the Query Parser and conduct in-depth investigative research of the Macie-monitored data and activity for the selected users.

Topics

MacieUniqueID

In the context of Macie a user is the AWS Identity and Access Management (IAM) identity that makes a particular request. Macie uses the CloudTrail `userIdentity` element to distinguish the following user types. For more information, see CloudTrail userIdentity Element.

- Root – The request was made with your AWS account credentials.
- IAM user – The request was made with the credentials of an IAM user.
- Assumed role – The request was made with temporary security credentials that were obtained with a role via a call to the AWS Security Token Service (AWS STS) AssumeRole API.
- Federated user – The request was made with temporary security credentials that were obtained via a call to the AWS STS GetFederationToken API.
- AWS account – The request was made by another AWS account.
- AWS service – The request was made by an AWS account that belongs to an AWS service.

When specifying a user in the Macie console, for example searching for a user in the **Users** tab or constructing a query in the **Research** tab, or whitelisting a user in a basic alert with the index of **Cloudtrail data**, you must use a special Macie format called **macieUniqueId**. The macieUniqueId is a combination of the IAM UserIdentity element and the recipientAccountId. For more information, see CloudTrail userIdentity Element and the definition of recipientAccountId in CloudTrail Record Contents.

The examples below list various structures of macieUniqueId, depending on the user identity type:

userIdentity	MacieUniqueId
`"userIdentity": { "type": "AssumedRole" "arn": "arn:aws:sts::123456789012:assumed-role/Accounting-Role/Mary" }`	123456789012:assumed-role/accounting-role
`"userIdentity": {`	123456789012:user:bob

userIdentity	MacieUniqueId

<pre>"type": "IAMUser",
"arn":
"arn:aws:iam::123456789012:user/
Bob",
"userName": "Bob"
}</pre>	

<pre>"userIdentity": {
"type": "FederatedUser"
"arn":
"arn:aws:sts::123456789012:federated-
user/Alice",
"principalId":
"123456789012:Alice",
}</pre>
	123456789012:federated-user:alice

<pre>"recipientAccountId":
"123456789012",
"userIdentity": {
"type": "AWSAccount"
"accountId":
"ANONYMOUS_PRINCIPAL",
}</pre>
	123456789012:ANONYMOUS_PRINCIPAL

<pre>"macieUniqueId":
"123456789012:root:root",
"userIdentity": {
"type": "Root"
"sourceARN":
"arn:aws:iam::123456789012:root",
}</pre>
	123456789012:root:root

<pre>"userIdentity": {
"invokedBy":
"codepipeline.amazonaws.com",
"type": "AWSService"
}
"recipientAccountId":
"123456789012",</pre>
	123456789012:codepipeline.amazonaws.com

<pre>"recipientAccoundId":
"123456789012",
"userIdentity": {
"type": "AWSAccount"
"accountId":
"987654321098",
"principalId":
"AIDABCDEFGHI123456XYZ",
}</pre>
	123456789012:AIDABCDEFGHI123456XYZ

User Categories in Macie

Based on their activity (API calls), users in Macie are grouped into the following categories:

- **Platinum**: these IAM users or roles have a history of making high risk API calls indicative of an administrator or root user, such as creating users, authorizing security group ingress, or updating policies. These accounts should be monitored closely for signs of account compromise.
- **Gold**: these IAM users or roles have a history of making infrastructure-related API calls indicative of a power user, such as running instances or writing data to Amazon S3. These accounts should be monitored closely for signs of account compromise.
- **Silver**: these IAM users or roles have a history of issuing high quantities of medium-risk API calls, such as Describe* and List* operations, or read-only access requests to Amazon S3.
- **Bronze**: these IAM users or roles typically execute lower quantities of Describe* and List* API calls in the AWS environment.

Investigating Users

Follow this procedure to generate a comprehensive picture of all of the Macie-monitored data and activity for the specified user:

1. In the Macie console's Users tab, specify a user name in the Search field and press Enter.

   Note
   When specifying a user, you must use a special Macie format called **macieUniqueId**. For example, 123456789012:root or 123456789012:user/Bob, or 123456789012:assumed-role/Accounting-Role/Mary, depending on the identity type of the user you want to whitelist. For more information, see the definition of the 'user' concept in Concepts and Terminology (p. 3).
2. When the user data is generated, choose the corresponding icon to select any of the following views to display various subsets of this user's Macie-monitored data and activity:

   - High-risk CloudTrail events (p. 34)
   - High-risk CloudTrail errors (p. 35)
   - Activity location (p. 35)
   - CloudTrail events (p. 35)
   - Activity ISPs (p. 35)
   - CloudTrail user identity types (p. 35)
3. If present in the selected view, locate and move the **Minimum risk** slider to the desired value. The **Minimum risk** slider enables you to only view items with the assigned risk equal to and greater than the selected value.

High-risk CloudTrail Events

This view provides a visual representation of the top 20 (by assigned risk and based on the value selected on the **Minimum risk** slider) CloudTrail data and management events that occurred during the last 60 days for the selected user. You can use the available **Daily** or **Weekly** radio buttons to modify the graph to view daily or weekly results.

To navigate from this view to the **Research** tab, you can select (double-click) any square that represents a particular event that you would like to investigate further. The number in parenthesis next to the event name represents the number of user sessions (5-minute aggregates of CloudTrail data) in which this event is present. In the **Research** tab, your selection is automatically translated into a query that is displayed in the Query Parser. For more information, see Using the Macie Research Tab (p. 36)

High-risk CloudTrail Errors

This view provides a visual representation of the top 20 (by assigned risk and based on the value selected on the **Minimum risk** slider) CloudTrail errors that occurred during the last 60 days for the selected user. You can use the available **Daily** or **Weekly** radio buttons to modify the graph to view daily or weekly results.

To navigate from this view to the **Research** tab, you can select (double-click) any square that represents a particular error that you would like to investigate further. The number in parenthesis next to the error name represents the number of user sessions (5-minute aggregates of CloudTrail data) in which this error is present. In the **Research** tab, your selection is automatically translated into a query that is displayed in the query parser. For more information, see Using the Macie Research Tab (p. 36).

Activity Location

This view includes a map that shows the locations of activity that Macie is monitoring for a selected time period for the specified user. To view details, you can use the available time period pull-down menu (past 15 days, past 30 days, past 90 days, or past year), and then choose any location pin.

To navigate from this view to the **Research** tab, choose the number of events that appears in a tooltip window for a location pin. In the **Research** tab, your selection is automatically translated into a query that is displayed in the query parser. For more information, see Using the Macie Research Tab (p. 36).

CloudTrail Events

This view provides the complete list of Macie-monitored CloudTrail data and management events for the specified user. For each event, the total count of the user sessions (5-minute integrations of CloudTrail data) in which this event is present and the percentage that this total represents of the total number of user sessions is displayed.

To navigate from this view to the **Research** tab, you can choose the looking glass icon next to any of the displayed events. Your selection is automatically translated into a query that is displayed in the query parser in the **Research** tab. For more information, see Using the Macie Research Tab (p. 36).

Activity ISPs

This view provides the complete list of Macie-monitored CloudTrail data and management events, grouped by the associated internet service providers (ISPs) for the specified user. For each ISP, the total count of the user sessions (5-minute integrations of CloudTrail data) in which this ISP is present and the percentage that this total represents of the total number of user sessions is displayed.

To navigate from this view to the **Research** tab, you can choose the looking glass icon next to any of the displayed themes. Your selection is automatically translated into a query that is displayed in the query parser in the **Research** tab. For more information, see Using the Macie Research Tab (p. 36).

CloudTrail User Identity Types

This view provides the complete list of Macie-monitored CloudTrail data and management events, grouped by the user identity type (for more information, see the definition for 'user' in Concepts and Terminology (p. 3)) for the specified user. For each user identity type, the total count of the user sessions (5-minute integrations of CloudTrail data) in which this user identity type is present and the percentage that this total represents of the total number of user sessions is displayed.

To navigate from this view to the **Research** tab, you can choose the looking glass icon next to any of the displayed themes. Your selection is automatically translated into a query that is displayed in the query parser in the **Research** tab. For more information, see Using the Macie Research Tab (p. 36).

Using the Macie Research Tab

You can use the **Research** tab in the Macie console to construct and run queries in the Query Parser and conduct in-depth investigative research of your Macie-monitored data and activity. You can navigate to the **Research** tab at any time and construct queries from scratch in the empty parser. For more information, see Constructing Queries in Macie (p. 36). Or you can be redirected to the **Research** tab from various places throughout the Macie console, for example, any of the **Dashboard** views (see Using the Macie Dashboard (p. 25)) or the **Basic alerts** list (see Amazon Macie Alerts (p. 18)). When redirected to the **Research** tab from other places in the console, your data selection is translated into an automatically generated query that is displayed in the query parser.

Topics

Constructing Queries in Macie

Macie allows you to construct queries in the Query Parser in the **Research** tab. This Query Parser is a lexer which interprets a string into a Lucene Query using JavaCC. For more information about query syntax, see Apache Lucene - Query Parser Syntax.

The following are example queries for common searches:

- You can use the following query to search for any console login not that did not originate from IP addresses owned by Amazon: `eventNameIsp.compound:/ConsoleLogin:~(Amazon.*)/`
- You can use the following query to search for PII artifacts inside a public S3 bucket: `filesystem_metadata.bucket:"my-public-bucket" AND (pii_impact:"moderate" OR pii_impact:"high")`

Research Filters

In the Macie **Research** tab, you can apply the following filters to your searches:

Data index

The first **Research** tab filter (pull-down list), with the pre-selected default value of **Cloudtrail data**, allows you to specifying the index (or the data repository) that you want Macie to search through. This filter includes the following options:

- **CloudTrail data** - a collection of 5-minute aggregates of raw Cloudtrail data
- **S3 bucket properties** - a collection of metadata about the S3 buckets that Macie is monitoring
- **S3 objects** - a collection of metadata about the S3 objects that are stored in the buckets that Macie is monitoring

Number of Results to Display

The next **Research** tab filter with the pre-selected default value of **Top 10**, allows you to control the number of results to display when you do your initial search and the number of additional results to display if more results are available. This filter includes the following options:

- Top 10
- Top 50
- Top 100
- Top 500

Time Range

The third **Research** tab filter with the pre-selected default value of **Past 30 days**, allows you to define a time range for which you want to display your search results. This filter includes the following options:

- Past 7 days
- Past 30 days
- Past 90 days
- Past 365 days
- All
- Custom time range

Researching AWS CloudTrail Data

The following section describes the elements of the search results that get displayed when you use the **Research** tab to investigate your Macie-monitored CloudTrail data.

Complete the following steps in the **Research** tab:

1. Select **CloudTrail data** in the first filter pull-down list.
2. For this sample procedure, select **Top 10** in the second filter pull-down list.
3. For this sample procedure, select **Past 90** days in the third filter pull-down list.
4. Choose the button with the looking glass icon to start the search.

Your search produces the following elements:

- The **total number of results** that matched your CloudTrail data search for the selected time range.
- The **graphical representation** of CloudTrail data search results for the selected time range.

  Note
  If your data set is very large and you specify a very wide time range, you data might not render properly and this graph might not be displayed as one of the resulting elements of your search.

  Important
  You can use the graph to further narrow your search and generate and run a query that will produce a subset of the results generated by your original selections in the steps above. Double-click any of the graph's results and your selection is translated into a new query that's automatically displayed in the **Query Parser**, and the **Research** tab is refreshed with the results of this new query.

- **Search results summary** - this is a list of the most significant fields from your search. The first line includes the top (or bottom) 3 values for each field. The second line includes the top (or bottom) 10 values for each field.

  Important
  You can use the fields in the search results summary to further narrow your search and generate and run a query that will produce a subset of the results generated by your original selections in the steps above. Choose the first or the second line of results for any field, and in the expanded results breakdown, choose the looking glass icon next to any of the results. Your choices are then translated into a new query that's automatically displayed in the **Query Parser**, and the **Research** tab is refreshed with the results of this new query.
- A list of **user sessions** (5-minute aggregates of CloudTrail data) that match your search criteria. You can choose any user session to expand it and view its details.

The following table includes the complete list of fields that can appear in the results of your CloudTrail data searches. You can use this table to investigate the results of your CloudTrail data searches:

Field name	Definition
IP address intelligence	Security-related intelligence on the source IP address of the request.
Objects read	ARNs of S3 objects that are being read by the event.
User agent	The agent through which the request was made.
Error code	The AWS service error if the request returns an error.
Resources	A list of resources accessed in the event.
Macie unique ID	A format that is unique to Macie for specifying users. Macie unique ID is a combination of the IAM UserIdentity element and the recipientAccountId.
Recipient account ID	The ID of the AWS account that received the event.
Account ID	The ID of the AWS account that owns the entity that granted permissions for the request.
IP location	Key pair of city and country that hosts the IP address of the request.
Principal ID	A unique identifier for the entity that made the request.
Source IP address	The IP address that the request was made from.
Objects deleted	ARNs of S3 objects that are being deleted by the event.
ISP	The ISP where the event originated.
Source IP address	IP address that the AWS request came from
Objects written	ARNs of S3 objects that are being written by the event.
Event type	The type of the event that generated the event record (for example, AwsApiCall, AwsServiceEvent, or ConsoleSignin).
Resource owner account ID	Account ID of the S3 bucket owner.
User identity type	The type of the identity that made the request.

Field name	Definition
Event source	The service that the request was made to.
Event name	The name of the event for which the request was made.
Session name	The name of the session that was created when the request was made through an assumed role.
Source ARN	The ARN used to make the request.
AWS region	The AWS region in which the request is made.
Event name (ISP)	The name of the event for which the request was made and where Macie was able to successfully resolve an ISP.
Timestamp	The timestamp of the request.
Event name by error code	The event name and the corresponding error code.
Count of unique objects deleted	The count of unique S3 objects deleted by the event.
Count of unique objects written	The count of unique S3 objects written by the event.
Count of unique objects read	The count of unique S3 objects read by the event.
Count of unique event names	The count of unique event names.
Earliest event timestamp	The timestamp of the earliest event in the user session.
Latest event timestamp	The timestamp of the latest event in the user session.

Researching S3 Bucket Properties Data

The following section describes the elements of the search results that get displayed when you use the **Research** tab to investigate your Macie-monitored S3 bucket properties data.

Complete the following steps in the **Research** tab:

1. Select **S3 bucket properties** in the first filter pull-down list.
2. For this sample procedure, select **Top 10** in the second filter pull-down list.
3. For this sample procedure, select **Past 90** days in the third filter pull-down list.
4. Choose the button with the looking glass icon to start the search.

Your search results contain the following elements:

- The **total number of results** that matched your S3 bucket properties data search for the selected time range.
- The **graphical representation** of the S3 bucket properties data search results for the selected time range.

  Note
  If your data set is very large and you specify a very wide time range, you data might not render properly and this graph might not be displayed as one of the resulting elements of your search.

Important
You can use the graph to further narrow your search and generate and run a query that will produce a subset of the results generated by your original selections in the steps above. Double-click any of the graph's results and your selection is translated into a new query that's automatically displayed in the **Query Parser**, and the **Research** tab is refreshed with the results of this new query.

- **Search results summary** - this is a list of the most significant fields from your search. The first line includes the top (or bottom) 3 values for each field. The second line includes the top (or bottom) 10 values for each field.

  Important
  You can use the fields in the search results summary to further narrow your search and generate and run a query that will produce a subset of the results generated by your original selections in the steps above. Choose the first or the second line of results for any field, and in the expanded results breakdown, choose the looking glass icon next to any of the results. Your choices are then translated into a new query that's automatically displayed in the **Query Parser**, and the **Research** tab is refreshed with the results of this new query.

- A list of S3 buckets that match your search criteria. You can choose any bucket to expand it and view its details.

The following table includes the complete list of fields that can appear in the results of your S3 buckets searches. You can use this table to investigate the results of your S3 buckets searches:

Field name	Definition
S3 fully writable	True or false value (evaluated through an automated reasoning engine) indicative of unrestricted public 'write' access policy to the S3 bucket.
S3 fully readable	True or false value (evaluated through an automated reasoning engine) indicative of unrestricted public 'read' access policy to the S3 bucket.
Bucket tag value	The values in the tag set associated with the S3 bucket.
Bucket region	Specifies the AWS region where the S3 bucket resides.
Bucket permission	Permission given to the grantee for the S3 bucket.
Bucket tag key	The keys in the tag set associated with the S3 bucket.
Bucket grantee type	Describes the S3 bucket's grantee type, such as CanonicalUser or a Group.
Bucket grantee URI	Describes the S3 bucket's grantee URI, such as AllUsers or AuthenticatedAWSusers.
Bucket policy principal ARNs	A list of ARNs for the principals listed in the S3 bucket's IAM policy.
Bucket name	The name of the S3 bucket.
Bucket owner account ID	The ID of the AWS account that owns the S3 bucket.
Bucket owner name	The display name of the S3 bucket owner.
Bucket timestamp	The timestamp when Macie last updated bucket information.

Field name	Definition
Bucket IAM policy	The IAM policy document attached to the S3 bucket.
Bucket lifecycle configuration	Lifecycle rules attached to the S3 bucket.
Bucket ACL	ACL information attached to the S3 bucket.
Bucket logging configuration	Indicates whether logging is enabled on the S3 bucket.
Bucket versioning	The versioning state of the S3 bucket.
Bucket tagging	Contains the S3 bucket's TagSet and Tag elements.
Bucket activity	If you choose this field, your selection is automatically translated into a query displayed in the Query Parser. When you run this query, it generates all CloudTrail activity on this S3 bucket for the same time range as specified in your original search.
S3 bucket URL	A URL for the S3 bucket in the Amazon S3 console.

Researching S3 Objects Data

The following section describes the elements of the search results that get displayed when you use the **Research** tab to investigate your Macie-monitored S3 objects.

Complete the following steps in the **Research** tab:

1. Select **S3 objects** in the first filter pull-down list.
2. For this sample procedure, select **Top 10** in the second filter pull-down list.
3. For this sample procedure, select **Past 90** days in the third filter pull-down list.
4. Choose the button with the looking glass icon to start the search.

Your search results include the following elements:

- The **total number of results** that matched your S3 objects search for the selected time range.
- The **graphical representation** of the S3 objects search results for the selected time range.

  Note
  If your data set is very large and you specify a very wide time range, you data might not render properly and this graph might not be displayed as one of the resulting elements of your search.

  Important
  You can use the graph to further narrow your search and generate and run a query that will produce a subset of the results generated by your original selections in the steps above. Double-click any of the graph's results and your selection is translated into a new query that's automatically displayed in the **Query Parser**, and the **Research** tab is refreshed with the results of this new query.
- **Search results summary** - this is a list of the most significant fields from your search. The first line includes the top (or bottom) 3 values for each field. The second line includes the top (or bottom) 10 values for each field.

  Important
  You can use the fields in the search results summary to further narrow your search and generate and run a query that will produce a subset of the results generated by your original

selections in the steps above. Choose the first or the second line of results for any field, and in the expanded results breakdown, choose the looking glass icon next to any of the results. Your choices are then translated into a new query that's automatically displayed in the **Query Parser**, and the **Research** tab is refreshed with the results of this new query.

- A list of S3 objects that match your search criteria. You can choose any S3 object to expand it and view its details.

The following table includes the complete list of fields that can appear in the results of your S3 objects searches. You can use this table to investigate the results of your S3 objects searches:

Field name	Definition
Object PII artifacts	The PII artifacts discovered in the contents of the S3 object.
Grantee URI	The URI of the S3 object's grantee.
Bucket name	The name of the S3 bucket where the S3 object is stored.
Object file type	S3 object's file type based on its Macie-assigned file extension.
Account ID	The AWS account ID that owns the S3 object.
Last modified	The time when the S3 object was last modified.
Object theme risk	S3 object's risk level determined by its Macie-assigned theme.
Object PII details	The total counts of unique PII artifacts discovered in the contents of the S3 object.
Object key	The key name that uniquely identifies the object in the S3 bucket.
AWS region	The AWS region that hosts the S3 bucket where the object is stored.
Object MD5 digest	The base64-encoded 128-bit MD5 digest of the object.
Object ACL	ACL information attached to the S3 object.
Object size	S3 object's size in bytes.
Object storage class	Storage class used for storing the S3 object.
Object matching keywords	The matching keywords discovered in the contents of the S3 object for the Macie-assigned theme.
Object KMS key	If the x-amz-server-side-encryption is present and has the value of aws:kms, this indicates the ID of the Key Management Service (KMS) master encryption key that was used for the object.
Matching REGEX	The matching keywords discovered in the contents of the S3 object for the Macie-assigned regex.
Object activity	If you choose this field, your selection is automatically translated into a query displayed in the Query Parser. When you run this query, it generates all CloudTrail activity on

Field name	Definition
	this S3 object for the same time range as specified in your original search.
S3 object URL	A URL for the S3 object in the Amazon S3 console.
Object cp CLI command	The cp CLI command to copy the S3 object to another location locally or in S3.
Object PII priority	The PII priority of the S3 object.
Object mimetypes	A mimetype match that describes the S3 object's content file type based on its header.
Bucket owner	The display name of the owner of the S3 bucket where the S3 object is stored.
Object language code	Then language code of the contents of the S3 object.
Object risk level	S3 object's Macie-assigned risk level.
Object encryption	The encryption standard used for encrypting the S3 object.
Object theme	S3 object's Macie-assigned theme.
Path prefix or folder	The first level prefix or folder of the S3 bucket where the S3 object is stored.
Object content type	S3 object's Macie-assigned content type.

Save a Query as an Alert

You can use the following procedure to save a query that is displayed in the query parser as a basic alert. For more information about basic alerts, see Amazon Macie Alerts (p. 18).

1. In the Macie console's **Research** tab, either autogenerate or construct a query in the query parser.
2. Choose the **Save query as alert** icon.
3. Fill out the Basic alert definition form and then choose **Save**. For more information, see Adding New and Editing Existing Custom Basic Alerts (p. 21).

Favorite Queries

You can mark queries that you frequently run as favorite and quickly view a list of your favorite queries.

1. In the Macie console's **Research** tab, either autogenerate or construct a query in the query parser.
2. Choose the **Mark query as favorite** icon.
3. Fill out the **Favorite query definition** form by specifying the name and the description for the favorite query, and then choose **Save**.
4. To view the list of your favorite queries, in the Macie console's **Research** tab, choose the **Favorite queries** icon.

Access Control in Amazon Macie

The master account users have access to the Macie console where they can configure Macie and use it to monitor and protect the resources in both master and member accounts. (For more information about master and member accounts, see Concepts and Terminology (p. 3) and Integrate Member Accounts and Amazon S3 with Amazon Macie (p. 8)).

In order for the master account users to be able to use the Macie console, they must be granted the required permissions. To ensure this, you can use the following policy document to create and attach an IAM policy to any user identity type that belongs to your master Macie account. This policy grants master account users permissions to use the Macie console in its full capacity:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Action": [
                "macie:*"
            ],
            "Resource": [
                "*"
            ],
            "Effect": "Allow"
        }
    ]
}
```

AWS Managed (Predefined) Policies for Macie

AWS addresses many common use cases by providing standalone IAM policies that are created and administered by AWS. Managed policies grant necessary permissions for common use cases so you can avoid having to investigate what permissions are needed.

The following AWS managed policies, which you can attach to users in your account, are specific to Macie:

- **AmazonMacieFullAccess** - Provides full access to Macie.
- **AmazonMacieSetupRole** - Provides Macie with access to your AWS account.
- **AmazonMacieServiceRole** - Grants Macie read-only access to resource dependencies in your account in order to enable data analysis.

Disable Amazon Macie and Delete Collected Metadata

You can use the Macie general settings page in the Macie console to disable Macie.

If you disable Macie, it will no longer have access to the resources in the master account and all member accounts. You must add member accounts again if you decide to re-enable Macie.

If you disable Macie, it stops processing the resources in the master account and all member accounts. Once Macie is disabled, the metadata that Macie collected while monitoring the data in your master and member accounts is deleted. Within 90 days from disabling Macie, all of this metadata is expired from the Macie system backups.

> **Important**
> Disabling Macie does not prompt the deletion of your other data within your AWS accounts. Only the metadata that was collected by Macie while it monitored your accounts is deleted once Macie is disabled.

1. Navigate to the **Macie general settings** page by choosing the down arrow next to your signed in name.
2. In the **Macie general settings** page, check the following checkboxes:
   - **I understand that if I disable Macie, the service will no longer have access to the resources in the master account and all member accounts. You must add member accounts again if you decide to re-enable Macie.**
   - **I understand that if I disable Macie, the service will stop processing the resources in the master account and all member accounts. All metadata that Macie collected while monitoring the data in these accounts will be deleted.**
3. Choose **Disable Amazon Macie**.

Monitoring Amazon Macie Alerts with Amazon CloudWatch Events

Amazon Macie sends notifications based on Amazon CloudWatch Events when any change in the Macie alerts takes place. This includes newly generated alerts and updates to existing alerts. Notifications are sent for all Macie alert types, including predictive alerts and basic alerts, both managed and custom. For more information about alert types, see Amazon Macie Alerts (p. 18).

Macie sends notifications based on Amazon CloudWatch Events for the alerts generated in both master and member Macie accounts. However, only the master Macie account has access to the generated CloudWatch events. For more information about master and member accounts, see Concepts and Terminology (p. 3).

The CloudWatch event for Macie has the following format:

Note
In the sample format below, the fictional account ID of "111122223333" represents the ID of the master Macie account.

```
{
  "version": "0",
  "id": "CWE-event-id",
  "detail-type": "Macie Alert",
  "source": "aws.macie",
  "account": "111122223333",
  "time": "2017-04-24T22:28:49Z",
  "region": "us-east-1",
  "resources": [
    "arn:aws:macie:us-east-1:111122223333:trigger/trigger_id/alert/alert_id",
    "arn:aws:macie:us-east-1:111122223333:trigger/trigger_id"
  ],
  "detail": {
    "notification-type": "ALERT_CREATED",
    "name": "Scanning bucket policies",
    "tags": [
      "Custom_Alert",
      "Insider"
    ],
    "url": "https://lb00.us-east-1.macie.aws.amazon.com/111122223333/posts/alert_id",
    "alert-arn": "arn:aws:macie:us-east-1:111122223333:trigger/trigger_id/alert/alert_id",
    "risk-score": 8,
    "trigger": {
      "rule-arn": "arn:aws:macie:us-east-1:111122223333:trigger/trigger_id",
      "alert-type": "basic",
      "created-at": "2017-01-02 19:54:00.644000",
      "description": "Alerting on failed enumeration of large number of bucket policies",
      "risk": 8
    },
    "created-at": "2017-04-18T00:21:12.059000",
    "actor": "555566667777:assumed-role:superawesome:aroaidpldc7nsesfnheji",
    "summary": {ALERT_DETAILS_JSON}
  }
}
```

You can complete the following procedure to configure your master Macie account to receive CloudWatch events from Macie and to pipe those events into an Amazon Simple Queue Service (SQS) queue. Before completing this procedure, make sure to create the SQS queue for the events from Macie. For more information, see Tutorial: Creating an Amazon SQS Queue.

1. Open the CloudWatch console at https://console.aws.amazon.com/cloudwatch/.
2. In the navigation pane, choose **Events**, and then choose **Create rule**.
3. Choose **Edit** and type in the following event pattern for the Macie events:

```
{
   "source": [
   "aws.macie"
  ]
}
```

4. In the **Targets** pane, choose **Add target**, select **SQS queue** in the target drop-down menu, and then specify your queue for the events from Macie.

You should now be able to see Macie alerts in your specified queue in the SQS console.